

Piecewise constant hazard estimation with the fused lasso

Manuel Rosenbaum, Jan Beyersmann & Michael Vogt¹

Ulm University

In applied time-to-event analysis, a flexible parametric approach is to model the hazard rate as a piecewise constant function of time. However, the change points and values of the piecewise constant hazard are usually unknown and need to be estimated. In this paper, we develop a fully data-driven procedure for piecewise constant hazard estimation. We work in a general counting process framework which nests a wide range of popular models in time-to-event analysis including Cox's proportional hazards model with potentially high-dimensional covariates, competing risks models as well as more general multi-state models. To construct our estimator, we set up a regression model for the increments of the Breslow estimator and then use fused lasso techniques to approximate the piecewise constant signal in this regression model. In the theoretical part of the paper, we derive the convergence rate of our estimator as well as some results on how well the change points of the piecewise constant hazard are approximated by our method. We complement the theory by both simulations and a real data example, illustrating that our results apply in rather general event histories such as multi-state models.

Key words: survival analysis, fused lasso, change points, piecewise constant hazard.

AMS 2020 subject classifications: 62J07, 62N02, 62P10.

1 Introduction

Hazard rates are a central quantity in time-to-event analysis. A flexible and useful parametric approach is to model the hazard as a piecewise constant function of time. A common example is the Cox model with a piecewise constant baseline hazard, which is often called the Poisson (or piecewise exponential) regression model for survival data (see e.g. p.192 in Andersen et al., 2021). In practice, the change points and function values of the piecewise constant hazard are unknown and need to be estimated. In this paper, we develop a fully data-driven method for estimating piecewise constant hazard functions. Our method is based on techniques from high-dimensional statistics, in particular on fused lasso techniques (Tibshirani et al., 2005), which – unlike many other machine learning techniques – are theoretically tractable and thus allow us to back up our estimation approach by quite comprehensive theory.

We work in a general counting process framework with a multiplicative intensity structure. Specifically, the intensity process is the product of a deterministic function of time

¹Address: Institute of Statistics, Department of Mathematics and Economics, Ulm University, 89081 Ulm, Germany. Email addresses: {manuel.rosenbaum,jan.beyersmann,m.vogt}@uni-ulm.de. Corresponding author is Michael Vogt.

– the (baseline) hazard rate we want to estimate – and a predictable process (which possibly incorporates covariate effects and may thus depend on unknown parameters). This framework nests a wide range of popular models in time-to-event analysis, including the Cox model with potentially high-dimensional covariates, competing risks models as well as more general multistate models. The counting process framework under consideration is introduced formally in Section 2 along with some leading examples of models that are nested in it.

Piecewise constant hazard models are regularly presented in the textbook literature on time-to-event analysis as a very flexible parametric modelling approach. However, guidance on choosing the number and location of the change points is either missing or rather informal. In the research literature, the majority of articles is restricted to simple survival models where the hazard only has a single jump or change point. Early examples are Matthews and Farewell (1982), Nguyen et al. (1984) and Loader (1991) who study maximum likelihood methods for estimating the change point. Further studies include Chang et al. (1994), Gijbels and Gürler (2003) and Zhao et al. (2009). Closely related to the problem of estimating the change point is the problem of testing the null hypothesis of a constant hazard against the alternative of a change point. This test problem is investigated in Matthews et al. (1985), Yao (1986), Worsley (1988), Henderson (1990) and Loader (1991) among others.

Only very few studies allow the piecewise constant hazard to have multiple change points. Goodman et al. (2011) and Han et al. (2014) propose sequential testing procedures to detect multiple change points in a simple survival model with right-censoring and no covariates. An extension of the approach in Goodman et al. (2011) to a two-sample setting is provided in He et al. (2013). Apart from these frequentist methods, Bayesian approaches to estimate a piecewise constant hazard with multiple change points can be found in Arjas and Gasbarra (1994) and Cooney and White (2021, 2023).

Another way to fit piecewise constant hazard models is via spline methods: Let the hazard be a function on the time interval $[0, \tau]$ and let $0 = u_0 < u_1 < u_2 < \dots < u_N < u_{N+1} = \tau$ be a partition of $[0, \tau]$. If the partition is chosen fine enough (i.e., if N is large enough), then the hazard can be approximated well by a piecewise constant function with jumps at the time points u_1, \dots, u_N . As such a piecewise constant function can be represented in terms of a spline basis with knots u_1, \dots, u_N , the hazard may be estimated by spline-based procedures. Specifically, one may estimate the spline coefficients together with the other model parameters by penalized maximum likelihood with a quadratic (i.e., ridge-type) penalty. This methodology is quite general in nature and can be applied to a wide range of models from time-to-event analysis. It is a specific variant of piecewise exponential additive modelling (PAM); see Bender et al. (2018) for an overview and Kopper et al. (2021) for a neural network assisted extension of PAM. The hazard estimator produced by this approach is a piecewise constant function with

potential jumps at u_1, \dots, u_N . Notably, one can show that the estimator will indeed have a jump at each point u_1, \dots, u_N in general (see e.g. Section 2 in Tibshirani, 2014). Hence, even if the underlying hazard is a simple piecewise constant function with a single change point, its estimator will have N change points in general. The drawback of this approach is thus that it does not produce parsimonious solutions.

Our approach is based on the fused lasso, which is closely related to spline methods (see Tibshirani, 2014, for details on the relation). However, in contrast to the spline approach from above, the fused lasso does not employ fixed knots but chooses the knots (and their number) adaptively. Put differently, it does not produce jumps at a large number of pre-specified knots. It rather selects the jump points in a data-driven way. For instance, if the underlying function is a simple step function with a single jump, the fused lasso attempts to find a parsimonious approximation with only a few change points (ideally one) in the vicinity of the true jump.

Roughly speaking, our estimation method works as follows: To start with, we estimate the cumulative hazard nonparametrically. To do so, we use an estimator off-the-shelf, in particular, the Breslow estimator (which reduces to the Nelson-Aalen estimator if there are no covariates). We then compute increments of the Breslow estimator over a fine grid of time points. These increments can be shown to follow a simple regression model with the following property: the regression function is a discretized version of the piecewise constant hazard we want to estimate. This allows us to construct an estimator of the hazard by applying techniques for estimating piecewise constant regression curves, in particular, fused lasso techniques. The main merits of our approach are as follows:

- (i) The approach is not restricted to a specific time-to-event model. It rather works in a general counting process framework which nests a wide range of popular models in time-to-event analysis.
- (ii) The approach does not presuppose any knowledge about the piecewise constant structure of the hazard. In particular, the location of change points, their number and the function values of the hazard are unknown.
- (iii) The approach produces parsimonious solutions, i.e., piecewise constant reconstructions of the hazard with only few parameters.
- (iv) The approach is fully data-driven: there are no free tuning parameters that need to be chosen in an adhoc fashion.

Our estimation methodology is developed in Section 3 and backed up by quite comprehensive theory in Section 4. Implementation details are dealt with in Section 5. The methods and theory are complemented by a simulation study in Section 6 and a real data example in Section 7. The data example considers a piecewise constant parametrization of the three-variate hazard measure in an illness-death multistate model, which

is used to jointly model so-called progression-free and overall survival, major time-to-event outcomes in oncology. The parametric fit is judged by comparison against standard nonparametric Kaplan-Meier estimators and may, e.g., be used for planning of randomised controlled trials and evaluation of operational trial characteristics.

2 Model setting

In its most general form, our model can be formulated as follows: We observe an n -dimensional counting process $N_{1:n} = \{N_{1:n}(t) : t \in [0, \tau]\}$ on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The process $N_{1:n}$ is adapted to the filtration $\{\mathcal{F}_t\}$, where \mathcal{F}_t is the σ -algebra generated by the observed data up to time t . Hence, \mathcal{F}_t represents the information available up to time t . Writing $N_{1:n} = (N_1, \dots, N_n)$, we regard $N_i = \{N_i(t) : t \in [0, \tau]\}$ as the counting process of the i -th observed subject in a sample of size n . The Doob-Meyer decomposition yields that for each i , $N_i(t) = \Lambda_i(t) + M_i(t)$, where Λ_i is a unique $\{\mathcal{F}_t\}$ -predictable process and M_i is a zero-mean local martingale with respect to $\{\mathcal{F}_t\}$. Assuming that Λ_i is absolutely continuous, we can write $\Lambda_i(t) = \int_0^t \lambda_i(s) ds$ with some $\{\mathcal{F}_t\}$ -predictable process λ_i . Plugging this into the Doob-Meyer decomposition gives

$$N_i(t) = \int_0^t \lambda_i(s) ds + M_i(t), \quad (2.1)$$

where λ_i is usually referred to as the intensity process of N_i . As common in the literature (cp. Aalen et al., 2008, Chapter 3.1.2), we assume that λ_i has the multiplicative structure

$$\lambda_i(t) = Z_i(t, \boldsymbol{\beta}) \alpha^*(t), \quad (2.2)$$

where α^* is a non-negative deterministic function and Z_i is an $\{\mathcal{F}_t\}$ -predictable process that may depend on certain unknown parameters $\boldsymbol{\beta} \in \mathbb{R}^d$ which need to be estimated from the data. The multiplicative structure (2.2) arises very naturally in many popular models in time-to-event analysis. Some leading examples are discussed below. In summary, we consider the general model

$$N_i(t) = \int_0^t Z_i(s, \boldsymbol{\beta}) \alpha^*(s) ds + M_i(t) \quad (2.3)$$

for each subject i and the corresponding cumulative model

$$\bar{N}(t) = \int_0^t \bar{Z}(s, \boldsymbol{\beta}) \alpha^*(s) ds + \bar{M}(t), \quad (2.4)$$

where $\bar{N}(t) = \sum_{i=1}^n N_i(t)$, $\bar{Z}(t, \boldsymbol{\beta}) = \sum_{i=1}^n Z_i(t, \boldsymbol{\beta})$ and $\bar{M}(t) = \sum_{i=1}^n M_i(t)$. Notably, the processes \bar{N} , \bar{Z} and \bar{M} depend on the sample size n . However, for ease of notation, we suppress this dependence throughout the paper.

The central statistical object in model (2.4) besides the parameter vector $\boldsymbol{\beta}$ is the function $\alpha^* : [0, \tau] \rightarrow \mathbb{R}_{\geq 0}$. In this paper, we assume that α^* has a piecewise constant structure: there are K time points $0 < \tau_1 < \dots < \tau_K < \tau$ such that

$$\alpha^*(t) = \begin{cases} \mathbf{a}_1 & \text{for } t \in [0, \tau_1) \\ \mathbf{a}_2 & \text{for } t \in [\tau_1, \tau_2) \\ \vdots & \\ \mathbf{a}_{K+1} & \text{for } t \in [\tau_K, \tau], \end{cases} \quad (2.5)$$

where $\mathbf{a}_1, \dots, \mathbf{a}_{K+1}$ are non-negative constants. The jump locations τ_1, \dots, τ_K , their number K as well as the function values $\mathbf{a}_1, \dots, \mathbf{a}_{K+1}$ are unknown. Our main goal is to construct an estimator of the piecewise constant function α^* .

Remark. Most probably, the assumption that the function α^* is piecewise constant does not hold exactly but only approximately in practice. We could reflect this in our methods and theory by allowing α^* to be a general function which can be approximated sufficiently well by a sequence of piecewise constant functions α_n^* with K_n jumps (where K_n grows with the sample size n and the approximation quality gets better with increasing n). It is possible to adapt our theory to this more general setting. However, as this would make the already quite intricate formulation of our results and proofs even more involved, we have decided to stick with the simpler assumption that α^* is exactly piecewise constant.

We now discuss some special cases of the general framework introduced above. In all considered settings, we assume that the data are independent and identically distributed (i.i.d.) across units i .

Setting A. To start with, we consider a simple survival model. Let T_i^* denote the survival time of subject i and assume that T_i^* has a density f^* . Our goal is to estimate the density f^* or the associated distribution function $F^*(t) = \int_0^t f^*(s) ds$. The situation is complicated by the fact that we do not fully observe the survival times T_i^* but only right-censored versions of them. More specifically, we observe the variables $T_i = T_i^* \wedge C_i$, where \wedge denotes the minimum and C_i is the censoring time, together with the variables $\delta_i = \mathbf{1}(T_i^* \leq C_i)$, which indicate whether there is censoring or not. The data sample thus has the form $\{(T_i, \delta_i) : i = 1, \dots, n\}$.

For theoretical analysis, a counting process formulation of the model is commonly used. In particular, we consider the counting processes N_i defined by

$$N_i(t) = \mathbf{1}(T_i \leq t, \delta_i = 1)$$

for $i = 1, \dots, n$. Under standard conditions discussed below, the model equations (2.3)

and (2.4) arise with

$$\alpha^*(t) := \frac{f^*(t)}{1 - F^*(t)} = \lim_{\delta \searrow 0} \frac{1}{\delta} \mathbb{P}(t \leq T_i^* < t + \delta \mid T_i^* \geq t) \quad (2.6)$$

and $Z_i(t) := \mathbf{1}(T_i \geq t)$ (which does not depend on any unknown parameters $\boldsymbol{\beta}$, that is, $Z_i(t, \boldsymbol{\beta}) = Z_i(t)$). Hence, the function α^* is identical to the hazard rate in the model. The hazard rate contains all the information we need. In particular, f^* and F^* can be retrieved from α^* by the formulas $f^*(t) = \alpha^*(t) \exp(-\int_0^t \alpha^*(s) ds)$ and $F^*(t) = 1 - \exp(-\int_0^t \alpha^*(s) ds)$.

The key assumption to guarantee (2.6) – i.e., to guarantee that the function α^* is identical to the hazard parameter of interest – is called *independent censoring* in the counting process literature; see e.g. Andersen et al. (1993) or Martinussen and Scheike (2006). A special case is the stronger assumption of *random censoring* where T_i^* and C_i are supposed to be stochastically independent.

Setting B. We now extend Setting A by incorporating covariates. There are various ways to do so. Presumably the most prominent way proposed by Cox leads to the proportional hazards model (Cox, 1972). Let the variables (T_i, δ_i) be defined as above and assume that we additionally observe a vector of d covariates $\mathbf{W}_i = (W_{i1}, \dots, W_{id})^\top$, which are time-independent for simplicity. The data sample is thus given by $\{(T_i, \delta_i, \mathbf{W}_i) : i = 1, \dots, n\}$. Notably, we do not assume the number of covariates d to be small relative to the sample size n . We rather allow for high-dimensional (sparse) cases where d is potentially much larger than n . A precise description of the high-dimensional setting under consideration is provided in Section 4.3. As in Setting A, we consider the counting processes

$$N_i(t) = \mathbf{1}(T_i \leq t, \delta_i = 1)$$

for $i = 1, \dots, n$. In the Cox model, the intensity process λ_i of N_i is assumed to have the form

$$\lambda_i(t) = \mathbf{1}(T_i \geq t) \exp(\boldsymbol{\beta}^\top \mathbf{W}_i) \alpha^*(t).$$

Setting $Z_i(t, \boldsymbol{\beta}) = \mathbf{1}(T_i \geq t) \exp(\boldsymbol{\beta}^\top \mathbf{W}_i)$, the process λ_i thus has the multiplicative structure (2.2), implying that model equations (2.3) and (2.4) are satisfied. As in Setting A, we assume the right-censoring to be independent. A special case is once again random censoring, meaning that T_i^* and C_i are stochastically independent given the covariates. Under the assumption of independent censoring, the conditional hazard

$$\alpha^*(t \mid \mathbf{w}) := \lim_{\delta \searrow 0} \frac{1}{\delta} \mathbb{P}(t \leq T_i^* < t + \delta \mid T_i^* \geq t, \mathbf{W}_i = \mathbf{w}) \quad (2.7)$$

is identical to $\alpha^*(t \mid \mathbf{w}) = \exp(\boldsymbol{\beta}^\top \mathbf{w}) \alpha^*(t)$ and α^* plays the role of a baseline hazard in the model.

Setting C. In Settings A and B, the observed subjects i only face one risk, e.g., the risk of dying when T_i^* is time-to-death. The competing risks model, in contrast, allows to incorporate different (competing) risks. Technically speaking, the survival times T_i^* have associated marks $\varepsilon_i \in \{1, 2, \dots, R\}$ which correspond to the R possible risks in the model. For instance, if T_i^* is time-to-death, there are R causes of death such as cardiovascular death ($\varepsilon_i = 1$), cancer death ($\varepsilon_i = 2$), and so on.

As in Settings A and B, we only observe a right-censored version $T_i = T_i^* \wedge C_i$ of the survival time T_i^* , where C_i denotes the censoring time. In addition, we allow for left-truncation in the model, where the left-truncation time L_i is assumed to be strictly smaller than C_i with probability 1. The time L_i is usually interpreted as delayed study entry, a common phenomenon in observational data. This means that subject i only becomes observable after time L_i *provided that* $T_i^* > L_i$. Thus, in the presence of left-truncation, the data are sampled from a conditional probability distribution, in particular, from the distribution given that study entry occurs. Defining δ_i as before and allowing for covariates \mathbf{W}_i as in Setting B, the observed data sample has the form $\{(T_i, \delta_i \cdot \varepsilon_i, L_i, \mathbf{W}_i) : i = 1, \dots, n\}$.

We now treat each event $r \in \{1, 2, \dots, R\}$ separately. We thus fix r and consider the event-specific counting processes N_{ir} for $i = 1, \dots, n$ defined by

$$N_{ir}(t) = \mathbf{1}(L_i < T_i \leq t, \delta_i \cdot \varepsilon_i = r).$$

As in Setting B, we assume a Cox model such that the intensity λ_{ir} of N_{ir} has the form

$$\lambda_{ir}(t) = \mathbf{1}(L_i < t \leq T_i) \exp(\boldsymbol{\beta}_r^\top \mathbf{W}_i) \alpha_r^*(t),$$

where $\boldsymbol{\beta}_r$ is a vector of event-specific regression coefficients. For given r , the n -dimensional counting process with the components N_{ir} for $i = 1, \dots, n$ thus satisfies the general model equations (2.3) and (2.4). Under the assumption of independent right-censoring and independent left-truncation, the event-specific conditional hazard

$$\alpha_r^*(t | \mathbf{w}) := \lim_{\delta \searrow 0} \frac{1}{\delta} \mathbb{P}(t \leq T_i^* < t + \delta, \varepsilon_i = r \mid T_i^* \geq t, \mathbf{W}_i = \mathbf{w})$$

is identical to $\alpha_r^*(t | \mathbf{w}) = \exp(\boldsymbol{\beta}_r^\top \mathbf{w}) \alpha_r^*(t)$ and the function α_r^* has the interpretation of an event-specific baseline hazard.

Setting D. As a final setting, we consider a general multi-state model where we observe a finite-state Markov process for each subject $i = 1, \dots, n$. More specifically, we observe a nonhomogeneous, time-continuous Markov process X_i with state space $\mathcal{R} = \{0, 1, 2, \dots, R\}$, right-continuous sample paths and, for ease of presentation, $\mathbb{P}(X_i(0) = 0) = 1$ for each i . A generalization to a non-degenerate initial distribution is immediate by conditioning on the initial states (Andersen et al., 1993, Section IV.4).

We treat each direct $\ell \rightarrow m$ transition with $\ell, m \in \mathcal{R}, \ell \neq m$, separately. For a given $\ell \rightarrow m$ transition, we define the counting process $N_{i,\ell \rightarrow m}$ for $i = 1, \dots, n$ by

$$N_{i,\ell \rightarrow m}(t) = \#\{s \leq t : X_i(s-) = \ell \text{ and } X_i(s) = m\},$$

which counts the number of direct $\ell \rightarrow m$ transitions in the time interval $[0, t]$. By Theorem II.6.8 in Andersen et al. (1993), the process $N_{i,\ell \rightarrow m}$ has the decomposition

$$N_{i,\ell \rightarrow m}(t) = \int_0^t \lambda_{i,\ell \rightarrow m}(s) ds + M_{i,\ell \rightarrow m}(t),$$

where $M_{i,\ell \rightarrow m}$ is a zero-mean local martingale with respect to the filtration $\{\mathcal{F}_t\}$ generated by the observed data and the intensity process $\lambda_{i,\ell \rightarrow m}$ has the form $\lambda_{i,\ell \rightarrow m}(s) = Z_{i,\ell}(s)\alpha_{\ell \rightarrow m}^*(s)$ with $Z_{i,\ell}(s) = \mathbf{1}(X_i(s-) = \ell)$ and the transition hazard

$$\alpha_{\ell \rightarrow m}^*(t) = \lim_{\delta \searrow 0} \frac{1}{\delta} \mathbb{P}(X_i(t + \delta) = m \mid X_i(t-) = \ell),$$

assuming that the limit in the above display exists. Hence, the processes $N_{i,\ell \rightarrow m}$ satisfy model equations (2.3) and (2.4). Right-censoring and left-truncation can be incorporated as additional states in the model (see Example III.3.3 in Andersen et al., 1993, for further details). Moreover, covariates can be incorporated, for example, by imposing a Cox model on the intensity process $\lambda_{i,\ell \rightarrow m}$ as in Settings B and C. For simplicity, we have however ignored covariates in our exposition.

Notably, Settings A–C can be formulated as special cases of the multi-state framework just described (with covariates added). Another popular multi-state model in applied survival analysis is the so-called illness-death model without recovery. This model is of major interest in oncology and underlies our application example in Section 7, where the model is introduced in detail.

3 Estimation strategy

We now describe how to construct an estimator of the piecewise constant hazard $\alpha^* : [0, \tau] \rightarrow \infty$. In the presence of left-truncation and right-censoring, the data usually become very sparse close to the end points of the interval $[0, \tau]$. Consider for instance a hypothetical clinical study with 0 and τ being the starting and closing date of the study, respectively. In such a situation, there are usually very few patients with event times close to τ . In case of delayed study entry, this issue arises (most probably) close to 0 as well. Hence, in the presence of left-truncation and right-censoring, reliable estimation of α^* is only possible on a subinterval $[\tau_{\min}, \tau_{\max}]$ of $[0, \tau]$. Notably, this is not a shortcoming of our methodology. It is rather an issue any (nonparametric) approach has to deal with. For example, it is well-known that in the presence of right-censoring, the Kaplan-Meier estimator becomes very unreliable towards the right-end

point τ . Usually, it levels out and the associated confidence intervals become huge, reflecting the sparsity of the data close to τ .

In what follows, we thus restrict attention to estimation of α^* on a suitable subinterval $[\tau_{\min}, \tau_{\max}]$. If there is only right-censoring, we can set $\tau_{\min} = 0$ and choose τ_{\max} somewhat smaller than τ , leaving us with the interval $[0, \tau_{\max}]$. If there is left-truncation on top, we also need to choose τ_{\min} somewhat larger than 0. For convenience, we assume that the interval $[\tau_{\min}, \tau_{\max}]$ (i) has length 1 and (ii) is so large that all change points of α^* lie in its interior, that is, $\tau_k \in (\tau_{\min}, \tau_{\max})$ for all $k = 1, \dots, K$. (i) can always be achieved by re-normalizing the data and (ii) is merely imposed to avoid additional notation specifying which change points lie in $(\tau_{\min}, \tau_{\max})$ and which do not.

We now turn to the construction of our estimator of α^* . We proceed in several steps:

Step 1. To start with, we require an estimator of the cumulative hazard function $A^*(t) = \int_0^t \alpha^*(s) ds$. We work with the standard Breslow estimator which is constructed as follows: By (2.4), it holds that $d\bar{N}(t) = \alpha^*(t)\bar{Z}(t, \boldsymbol{\beta})dt + d\bar{M}(t)$. With $J(t, \boldsymbol{\beta}) = \mathbf{1}(\bar{Z}(t, \boldsymbol{\beta}) > 0)$ and the convention that $0/0 := 0$, we thus get that

$$\int_0^t \frac{J(s, \boldsymbol{\beta})}{\bar{Z}(s, \boldsymbol{\beta})} d\bar{N}(s) = \int_0^t J(s, \boldsymbol{\beta}) \alpha^*(s) ds + \int_0^t \frac{J(s, \boldsymbol{\beta})}{\bar{Z}(s, \boldsymbol{\beta})} d\bar{M}(s). \quad (3.1)$$

Now suppose we have an estimator $\hat{\boldsymbol{\beta}}$ of the parameter vector $\boldsymbol{\beta} \in \mathbb{R}^d$ at hand. How to construct such an estimator depends on the specific setting under consideration. In many cases, it is possible to use a standard estimator off the shelf. In the Cox model of Settings B and C, for instance, $\hat{\boldsymbol{\beta}}$ can be chosen to be the partial likelihood estimator of $\boldsymbol{\beta}$ or an ℓ_1 -penalized version of it. From (3.1), it follows that

$$\int_0^t \frac{J(s, \hat{\boldsymbol{\beta}})}{\bar{Z}(s, \hat{\boldsymbol{\beta}})} d\bar{N}(s) = \int_0^t \alpha^*(s) ds + \Delta^J(t) + \Delta^\beta(t) + \eta(t), \quad (3.2)$$

where $A^*(t) = \int_0^t \alpha^*(s) ds$ is the cumulative hazard and

$$\begin{aligned} \Delta^J(t) &= \int_0^t (J(s, \boldsymbol{\beta}) - 1) \alpha^*(s) ds \\ \Delta^\beta(t) &= \int_0^t \frac{J(s, \hat{\boldsymbol{\beta}})}{\bar{Z}(s, \hat{\boldsymbol{\beta}})} d\bar{N}(s) - \int_0^t \frac{J(s, \boldsymbol{\beta})}{\bar{Z}(s, \boldsymbol{\beta})} d\bar{N}(s) \\ \eta(t) &= \int_0^t \frac{J(s, \boldsymbol{\beta})}{\bar{Z}(s, \boldsymbol{\beta})} d\bar{M}(s). \end{aligned}$$

The left-hand side of (3.2),

$$\hat{A}(t) := \int_0^t \frac{J(s, \hat{\boldsymbol{\beta}})}{\bar{Z}(s, \hat{\boldsymbol{\beta}})} d\bar{N}(s),$$

is the Breslow estimator of $A^*(t)$. As integrating over a counting process is the same as summing the integrand over the jump times of the process, it can be written as $\hat{A}(t) = \sum_{\{k: S_k \leq t\}} J(S_k, \hat{\beta}) / \bar{Z}(S_k, \hat{\beta})$, where $S_1 < S_2 < S_3 < \dots$ are the jump times of \bar{N} . According to (3.2), the distance between the Breslow estimator $\hat{A}(t)$ and the cumulative hazard $A^*(t)$ is the sum of three error terms: the bias term $\Delta^J(t)$, the approximation error $\Delta^\beta(t)$ which results from estimating β by $\hat{\beta}$ and the term $\eta(t)$ which (by using standard theory for counting processes) can be shown to be a zero-mean local martingale.

Step 2. We next derive a simple regression model for the increments of the Breslow estimator \hat{A} . Cover the interval $[\tau_{\min}, \tau_{\max}]$ with an equidistant grid of step length $1/n$. The grid points are given by $t_j = \tau_{\min} + j/n$ for $j = 0, \dots, n$ (taking into account that the length of $[\tau_{\min}, \tau_{\max}]$ is 1). Now consider the increments

$$Y_j := \frac{\hat{A}(t_j) - \hat{A}(t_{j-1})}{t_j - t_{j-1}} = n\{\hat{A}(t_j) - \hat{A}(t_{j-1})\}.$$

Using (3.2), we immediately obtain that

$$Y_j = \alpha_j^* + u_j \quad \text{with} \quad u_j = \Delta_j^\alpha + \Delta_j^J + \Delta_j^\beta + \eta_j \quad (3.3)$$

and $\alpha_j^* = \alpha^*(t_j)$ for $j = 1, \dots, n$, where

$$\begin{aligned} \Delta_j^\alpha &= \frac{\int_{t_{j-1}}^{t_j} \alpha^*(s) ds}{t_j - t_{j-1}} - \alpha^*(t_j) \\ \Delta_j^\ell &= \frac{\Delta^\ell(t_j) - \Delta^\ell(t_{j-1})}{t_j - t_{j-1}} \quad \text{for } \ell \in \{J, \beta\} \\ \eta_j &= \frac{\eta(t_j) - \eta(t_{j-1})}{t_j - t_{j-1}}. \end{aligned}$$

Hence, the increments Y_j follow a simple signal-plus-noise model, where the signal α_j^* is the hazard rate $\alpha^*(t_j)$ at time point $t_j = \tau_{\min} + j/n$ and u_j is the noise term. Put differently, the signal vector $\boldsymbol{\alpha}^* = (\alpha_1^*, \dots, \alpha_n^*)^\top$ in model (3.3) is identical to the discretized hazard rate $\boldsymbol{\alpha}^* = (\alpha^*(t_1), \dots, \alpha^*(t_n))^\top$. Notably, the variables Y_j , α_j^* and u_j in model (3.3) depend on n . A more precise model formulation thus reads $Y_{j,n} = \alpha_{j,n}^* + u_{j,n}$ with $\alpha_{j,n}^* = \alpha^*(\tau_{\min} + j/n)$. For ease of notation, we however suppress the dependence of Y_j , α_j^* and u_j on n in what follows.

Step 3. By assumption, the discretized hazard rate $\boldsymbol{\alpha}^* = (\alpha^*(t_1), \dots, \alpha^*(t_n))^\top$ is a piecewise constant vector. In order to estimate $\boldsymbol{\alpha}^*$ in model (3.3), we may thus use techniques to recover a piecewise constant signal in a regression model. One such technique is the fused lasso. Applying the fused lasso to the constructed data sample

$\{Y_1, \dots, Y_n\}$ yields the estimator $\hat{\boldsymbol{\alpha}}_\lambda = (\hat{\alpha}_{\lambda,1}, \dots, \hat{\alpha}_{\lambda,n})^\top$ defined by

$$\hat{\boldsymbol{\alpha}}_\lambda \in \operatorname{argmin}_{\boldsymbol{\alpha} \in \mathbb{R}^n} \left\{ \frac{1}{n} \sum_{j=1}^n (Y_j - a_j)^2 + \lambda \sum_{j=2}^n |a_j - a_{j-1}| \right\}, \quad (3.4)$$

where $\lambda > 0$ is a tuning parameter. An estimator of the hazard rate α^* on the interval $[\tau_{\min}, \tau_{\max}]$ is immediately obtained from $\hat{\boldsymbol{\alpha}}_\lambda$ by constant interpolation. Specifically, we define the estimator $\hat{\alpha}_\lambda : [\tau_{\min}, \tau_{\max}] \rightarrow \mathbb{R}$ by setting

$$\hat{\alpha}_\lambda(t) = \begin{cases} \hat{\alpha}_{\lambda,1} & \text{for } t \in [t_0, t_1) \\ \hat{\alpha}_{\lambda,j} & \text{for } t \in [t_j, t_{j+1}) \text{ and } j \in \{1, \dots, n-1\} \\ \hat{\alpha}_{\lambda,n} & \text{for } t = t_n \end{cases} \quad (3.5)$$

with $t_j = \tau_{\min} + j/n$.

4 Theoretical results

In what follows, we investigate the estimator $\hat{\boldsymbol{\alpha}}_\lambda$ (and thus also implicitly the interpolated version $\hat{\alpha}_\lambda$) from a theoretical point of view. More specifically, we derive results on the convergence rate of $\hat{\boldsymbol{\alpha}}_\lambda$ as well as results on how well the change points of the hazard rate α^* are approximated by those of $\hat{\boldsymbol{\alpha}}_\lambda$. Sections 4.1 and 4.2 provide the theory for the general counting process model of Section 2. In Section 4.3, we verify that the conditions required for the general theory are satisfied in Settings A–D under standard assumptions. All proofs are relegated to the Appendix. Asymptotic statements are always to be understood in the sense that the sample size n tends to infinity. Notably, the processes \bar{N} , \bar{M} and \bar{Z} as well as the processes derived from them (in particular, J , Δ^J , Δ^β and η) and the elements of the regression model (3.3) (namely, Y_j , α_j^* and u_j) all depend on the sample size n . To keep the notation concise, we however suppress their dependence on n throughout.

4.1 Assumptions

We impose the following assumptions on the general model from Section 2:

- (C1) The filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$ satisfies the “usual conditions”.
- (C2) The counting processes N_i are i.i.d. across i .
- (C3) For each i , the compensator Λ_i of the process N_i has the form $\Lambda_i(t) = \int_0^t \lambda_i(s) ds$ with some $\{\mathcal{F}_t\}$ -predictable process λ_i . Moreover, the process λ_i has the multiplicative structure (2.2), that is, $\lambda_i(t) = Z_i(t, \boldsymbol{\beta}) \alpha^*(t)$ with a non-negative deterministic function α^* and an $\{\mathcal{F}_t\}$ -predictable process $Z_i(\cdot, \boldsymbol{\beta})$.

(C4) The processes $\bar{Z}(\cdot, \boldsymbol{\beta})$ and $J(\cdot, \boldsymbol{\beta})/\bar{Z}(\cdot, \boldsymbol{\beta})$ are locally bounded for each n .

(C5) It holds that $\mathbb{P}(\inf_{t \in [\tau_{\min}, \tau_{\max}]} J(t, \boldsymbol{\beta}) = 0) = o(1)$.

(C6) For some natural number $\nu > 4$,

$$\mathbb{E} \left[\left\| \frac{J(\cdot, \boldsymbol{\beta})}{\bar{Z}(\cdot, \boldsymbol{\beta})} \right\|_{\infty}^{\nu} \right] = \mathbb{E} \left[\sup_{t \in [\tau_{\min}, \tau_{\max}]} \left| \frac{J(t, \boldsymbol{\beta})}{\bar{Z}(t, \boldsymbol{\beta})} \right|^{\nu} \right] \leq \frac{C_{\infty, \nu}}{n^{\nu}},$$

where $C_{\infty, \nu} < \infty$ is a fixed constant independent of n .

(C7) The approximation error $\Delta^{\beta}(t)$ in (3.2) is such that

$$\max_{1 \leq j \leq n} \left| \frac{\Delta^{\beta}(t_j) - \Delta^{\beta}(t_{j-1})}{t_j - t_{j-1}} \right| = O_p(\rho_n),$$

where $\rho_n = o(n^{-\xi})$ for some arbitrarily small but fixed $\xi > 0$.

(C1)–(C4) are standard conditions in the context of statistical models based on counting processes. (C5)–(C7) are quite abstract high-order conditions. We show below that they are fulfilled in Settings A–D under common assumptions.

4.2 General results

Our first theoretical result bounds the distance between the estimator $\hat{\boldsymbol{\alpha}}_{\lambda}$ and the discretized hazard rate $\boldsymbol{\alpha}^*$. More specifically, it provides a bound on $|\hat{\alpha}_{\lambda, j} - \alpha_j^*|$ for any $j \in \{1, \dots, n\}$. To formulate the result, we use the following notation: For $k \in \{1, \dots, K\}$, we let $n_k = \lceil (\tau_k - \tau_{\min})n \rceil$ denote the index where the discretized hazard rate $\boldsymbol{\alpha}^*$ has the k -th jump. Put differently, $n_1 < \dots < n_K$ are the K indices with $\alpha_{n_{k-1}}^* \neq \alpha_{n_k}^*$. (Note that $1 < n_1$ and $n_K < n$ for sufficiently large n because the change points of α^* all lie in the interior of $[\tau_{\min}, \tau_{\max}]$ by assumption.) In what follows, we call n_1, \dots, n_K the jump indices of the discretized hazard $\boldsymbol{\alpha}^*$. For notational convenience, we additionally set $n_0 = 1$ along with $n_{K+1} = n + 1$. For any j , we let $n_{k(j)}$ and $n_{k(j)+1}$ be the two jump indices with $n_{k(j)} \leq j \leq n_{k(j)+1} - 1$ and define $d_j = \min\{j + 1 - n_{k(j)}, n_{k(j)+1} - j\}$, which essentially gives the distance of j to the nearest jump index. Finally, we let $r_{k(j)} = n_{k(j)+1} - n_{k(j)}$ be the length of the interval between the $k(j)$ -th and the $(k(j) + 1)$ -th jump index.

Theorem 4.1. *Let (C1)–(C7) be satisfied and let $\kappa_n = c_n \max\{n^{2/\nu}, \sqrt{n}\rho_n\}$ with $\nu > 4$ from (C6), where $\{c_n\}$ is a slowly diverging sequence (e.g. $c_n = c_0 \log \log n$ with some constant $c_0 > 0$). Then for any positive value of the regularization parameter λ , the following holds with probability tending to 1:*

$$|\hat{\alpha}_{\lambda, j} - \alpha_j^*| \leq \max \left\{ \frac{\kappa_n}{\sqrt{d_j}}, \frac{\kappa_n^2}{4n\lambda}, \frac{2n\lambda}{r_{k(j)}} + \frac{2\kappa_n}{\sqrt{r_{k(j)}}} \right\} \quad \text{for } j \in \{1, \dots, n\}.$$

The elementwise bound of Theorem 4.1 allows us to derive the following bound on the ℓ_2 -error $\|\hat{\boldsymbol{\alpha}}_\lambda - \boldsymbol{\alpha}^*\|_2^2/n$.

Theorem 4.2. *Let (C1)–(C7) be satisfied and choose $\lambda = \kappa_n/\sqrt{n}$, where κ_n is defined as in Theorem 4.1. Then*

$$\frac{1}{n} \|\hat{\boldsymbol{\alpha}}_\lambda - \boldsymbol{\alpha}^*\|_2^2 = O_p\left(c_n^2 \log(n) \max\left\{n^{-1+\frac{4}{\nu}}, \rho_n^2\right\}\right).$$

We briefly give some remarks on Theorem 4.2:

- (a) As we will see in our analysis of Settings A–D, a leading case is that ν can be chosen as large as desired and $\rho_n = O(\log(n)/\sqrt{n})$. In this case, Theorem 4.2 yields that

$$\frac{1}{n} \|\hat{\boldsymbol{\alpha}}_\lambda - \boldsymbol{\alpha}^*\|_2^2 = O_p\left(\frac{1}{n^{1-\xi}}\right)$$

with $\xi > 0$ arbitrarily small but fixed. This is a fast convergence rate not far from optimal (see the discussion in Lin et al., 2017, for details on sharp rates for the fused lasso), which suggests that the general rate established in Theorem 4.2 is quite sharp.

- (b) Theorem 4.2 can alternatively be formulated in terms of the hazard function $\alpha^* : [\tau_{\min}, \tau_{\max}] \rightarrow \mathbb{R}$ rather than the discretized hazard rate $\boldsymbol{\alpha}^*$. To do so, let $\hat{\alpha}_\lambda : [\tau_{\min}, \tau_{\max}] \rightarrow \mathbb{R}$ be the interpolated version of $\hat{\boldsymbol{\alpha}}_\lambda$ defined in (3.5). Then Theorem 4.2 and some straightforward calculations imply that

$$\int_{\tau_{\min}}^{\tau_{\max}} \{\hat{\alpha}_\lambda(t) - \alpha^*(t)\}^2 dt = O_p\left(c_n^2 \log(n) \max\left\{n^{-1+\frac{4}{\nu}}, \rho_n^2\right\}\right).$$

We next turn to change point estimation. The change points of the estimator $\hat{\boldsymbol{\alpha}}_\lambda$ are given by the set

$$\hat{\mathcal{S}}_\lambda = \left\{ \tau_{\min} + \frac{j}{n} \in [\tau_{\min}, \tau_{\max}] : \hat{\alpha}_{\lambda, j-1} \neq \hat{\alpha}_{\lambda, j} \text{ for } j \in \{2, \dots, n\} \right\}.$$

Our goal is to understand how well the change points $\mathcal{S}^* = \{\tau_1, \dots, \tau_K\}$ of the hazard rate α^* are approximated by those in $\hat{\mathcal{S}}_\lambda$. Informally speaking, we show that the following two statements hold with high probability:

- (CP1) For each change point in \mathcal{S}^* , there exists at least one change point in $\hat{\mathcal{S}}_\lambda$ “close” to it.
- (CP2) Let $\hat{\tau} \in \hat{\mathcal{S}}_\lambda$ be a change point which is “far away” from any change point in \mathcal{S}^* . Then the jump size of $\hat{\boldsymbol{\alpha}}_\lambda$ at the change point $\hat{\tau}$ is “small”.

According to (CP1), $\hat{\boldsymbol{\alpha}}_\lambda$ has a jump close to each change point of α^* . In general,




Figure 1: Example of a piecewise constant hazard rate α^* (in black) and the fused lasso estimator $\hat{\alpha}_\lambda$ (in red).

however, we cannot guarantee that it has exactly one jump close to each change point. This means that our fused lasso estimator may reconstruct a change point in α^* by multiple jumps. Figure 1 gives an illustration. The black line is the true hazard rate α^* , which has a single change point at 0.5, and the red line is (the interpolated version $\hat{\alpha}_\lambda$ of) the estimator $\hat{\alpha}_\lambda$. As can be seen, $\hat{\alpha}_\lambda$ has two jumps close to the time point 0.5 which together reconstruct the change point of α^* . According to (CP2), $\hat{\alpha}_\lambda$ may have additional jumps far away from any change point of α^* . However, these additional jumps are negligible in the sense of being small. Figure 1 once again gives an illustration. The estimator $\hat{\alpha}_\lambda$ (in red) has a jump between 0.8 and 0.9. Nevertheless, it yields a good overall reconstruction of the step function α^* as this additional jump is small. In principle, it is possible to get rid of additional jumps far away from the true change points by applying post-processing methods to the fused lasso (see e.g. the Haar wavelet based method in Lin et al., 2017). However, these methods are rather of theoretical than practical value as they depend on additional tuning parameters which are extremely hard to choose in practice. We thus do not consider any such methods. In summary, (CP1) and (CP2) imply that the estimator $\hat{\alpha}_\lambda$ produces an accurate reconstruction of the piecewise constant hazard α^* , even though it may have more change points than α^* in general. Hence, our method is a suitable tool for (parsimonious) fitting of piecewise constant hazard curves, which is the main goal in many applications in time-to-event analysis (rather than precise estimation and interpretation of location and number of change points).

We now formalize the statements (CP1) and (CP2). To do so, we measure the distance between two discrete sets A and B by

$$d(A|B) = \max_{b \in B} \min_{a \in A} |a - b|.$$

The quantity $d(A|B)$ gives the maximal distance of an element in B to its closest element in A . Hence, it is small if for any element in B , there is an element in A close

to it. In order to formalize (CP1), we analyze the quantity $d(\hat{\mathcal{S}}_\lambda | \mathcal{S}^*)$. The following theorem shows that $d(\hat{\mathcal{S}}_\lambda | \mathcal{S}^*)$ is small for large n in the following sense:

Theorem 4.3. *Under the conditions of Theorem 4.2, it holds that*

$$d(\hat{\mathcal{S}}_\lambda | \mathcal{S}^*) = O_p\left(c_n^2 \log(n) \max\left\{n^{-1+\frac{4}{\nu}}, \rho_n^2\right\}\right).$$

In the leading case where ν can be chosen as large as desired and $\rho_n = O(\log(n)/\sqrt{n})$, Theorem 4.3 in particular yields that

$$d(\hat{\mathcal{S}}_\lambda | \mathcal{S}^*) = O_p\left(\frac{1}{n^{1-\xi}}\right)$$

with $\xi > 0$ arbitrarily small but fixed. Hence, there exists a small neighbourhood of order $1/n^{1-\xi}$ around each change point in \mathcal{S}^* which contains an element of $\hat{\mathcal{S}}_\lambda$. The next theorem formalizes (CP2).

Theorem 4.4. *Let the conditions of Theorem 4.2 be satisfied, let $\{\Delta_n\}$ be any sequence of positive real numbers and let $\hat{\mathcal{S}}_\lambda^{\text{far}}$ be the set of all change points $\hat{\tau} \in \hat{\mathcal{S}}_\lambda$ with $\min_{1 \leq k \leq K} |\hat{\tau} - \tau_k| \geq \Delta_n/n$. Then with probability tending to 1,*

$$|\hat{\alpha}_{\lambda,(\hat{\tau}-\tau_{\min})n} - \hat{\alpha}_{\lambda,(\hat{\tau}-\tau_{\min})n-1}| \leq C \frac{\kappa_n}{\sqrt{\Delta_n}} \quad \text{for all } \hat{\tau} \in \hat{\mathcal{S}}_\lambda^{\text{far}},$$

where C is a sufficiently large constant.

Theorem 4.4 essentially says the following: if the distance of a change point $\hat{\tau} \in \hat{\mathcal{S}}_\lambda$ to any change point τ_k of the hazard rate α^* is larger than Δ_n/n , then the estimator $\hat{\alpha}_\lambda$ cannot have a jump larger than $C\kappa_n/\sqrt{\Delta_n}$ at the change point $\hat{\tau}$. In the leading case where ν can be chosen as large as desired and $\rho_n = O(\log(n)/\sqrt{n})$, Theorem 4.4 simplifies slightly to: with probability tending to 1,

$$|\hat{\alpha}_{\lambda,(\hat{\tau}-\tau_{\min})n} - \hat{\alpha}_{\lambda,(\hat{\tau}-\tau_{\min})n-1}| \leq C \frac{n^\xi}{\sqrt{\Delta_n}} \quad \text{for all } \hat{\tau} \in \hat{\mathcal{S}}_\lambda^{\text{far}}$$

with $\xi > 0$ arbitrarily small but fixed. This in particular implies the following: if $\hat{\tau}$ remains bounded away from any change point τ_k of α^* , that is, if $\min_{1 \leq k \leq K} |\hat{\tau} - \tau_k| > c > 0$ for some $c > 0$ (with probability tending to 1), then the corresponding jump size $|\hat{\alpha}_{\lambda,(\hat{\tau}-\tau_{\min})n} - \hat{\alpha}_{\lambda,(\hat{\tau}-\tau_{\min})n-1}|$ is smaller than Cn^ξ/\sqrt{n} (with probability tending to 1).

4.3 Results for Settings A–D

We now show that the general theory from the previous Section 4.2 applies to Settings A–D. To do so, we verify that the high-order conditions (C1)–(C7) are fulfilled in these four settings under standard assumptions. As Setting A is nested as a special case in the Cox model of Setting B, we directly start with Setting B and do not treat Setting

A separately. A formal result for Setting A follows as a simple corollary to Proposition 4.5 below.

The Cox model of Setting B incorporates a vector of d covariates. Rather than restricting attention to the standard low-dimensional case where the dimension d is small and fixed, we consider the following high-dimensional setup: (i) The dimension d may be much larger than the sample size n . In particular, d may grow as any polynomial of n , that is, $d = O(n^a)$ with a arbitrarily large but fixed. (ii) The parameter vector $\boldsymbol{\beta}$ is such that $\|\boldsymbol{\beta}\|_1$ remains bounded as $n \rightarrow \infty$. In principle, it is possible to allow $\|\boldsymbol{\beta}\|_1$ to grow (sufficiently slowly) with n . However, to avoid certain technical complications, we make the stronger assumption that it is bounded. Condition (ii) can be regarded as a sparsity constraint. It is in particular fulfilled if the number of non-zero coefficients of $\boldsymbol{\beta}$ does not grow with the sample size and the entries of $\boldsymbol{\beta}$ are all bounded by a fixed constant $C < \infty$ in absolute value.

Proposition 4.5. *Consider the Cox model from Setting B and assume the following:*

(C_{B1}) *The data points $(T_i, \delta_i, \mathbf{W}_i)$ are i.i.d. across i .*

(C_{B2}) $\max_{1 \leq k \leq n} \mathbb{P}(t_{k-1} < T_i^* \leq t_k) \leq C/n$ *with some sufficiently large constant C .*

(C_{B3}) *The components N_i of the multivariate counting process $N_{1:n} = (N_1, \dots, N_n)$ have $\{\mathcal{F}_t\}$ -intensities $\lambda_i(t) = \mathbf{1}(t \leq T_i) \exp(\boldsymbol{\beta}^\top \mathbf{W}_i) \alpha^*(t)$.*

(C_{B4}) $\mathbb{P}(T_i \geq \tau_{\max}) = p$ *with some constant $p \in (0, 1)$.*

(C_{B5}) *The dimension d of the covariate vector \mathbf{W}_i grows at most polynomially in the sample size n , that is, $d = O(n^a)$ with some arbitrarily large but fixed constant a . Moreover, the covariates \mathbf{W}_i have bounded support, that is, $\|\mathbf{W}_i\|_\infty \leq C_W$ for all i and some constant $C_W < \infty$.*

(C_{B6}) *It holds that $\|\boldsymbol{\beta}\|_1 \leq C_\beta < \infty$ for some constant $C_\beta < \infty$ and $\hat{\boldsymbol{\beta}}$ is any estimator of $\boldsymbol{\beta}$ with the property that $\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_1 = O_p(\sqrt{\log(n)/n})$.*

Then (C1)–(C7) are satisfied with ν as large as desired, $\tau_{\min} = 0$ and $\rho_n = \log(n)/\sqrt{n}$.

We briefly comment on the assumptions of Proposition 4.5. The i.i.d. assumption in (C_{B1}) is standard in the literature. (C_{B2}) is rather mild as well. It is in particular fulfilled if the survival times T_i^* have a bounded density f^* . (C_{B3}) is merely a definition of the Cox model. (C_{B4}) requires that the probability of observing a right-censored survival time T_i (weakly) larger than τ_{\max} is strictly positive. This ensures that the data do not become too sparse towards the right end-point of the interval $[0, \tau_{\max}]$. Hence, (C_{B4}) essentially encapsulates that the data points T_i lie sufficiently dense in the interval $[0, \tau_{\max}]$ to perform reasonable estimation. The first part of (C_{B5}) restricts the dimension d of the covariate vector as already discussed above. The second part imposes boundedness on the covariates, which is a very common assumption in the

literature and not very problematic in practice as most covariates can be regarded as having bounded support. As shown in Huang et al. (2013), (C_B6) is satisfied (under appropriate regularity conditions not spelt out here) by an ℓ_1 -penalized version (i.e., a lasso version) of the partial likelihood estimator of $\boldsymbol{\beta}$. In the low-dimensional case where d does not grow with n , one may of course use the standard (unpenalized) partial likelihood estimator, which has the property that $\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_1 = O_p(1/\sqrt{n})$ (once again under certain regularity conditions not spelt out here).

We next turn to the competing risks model of Setting C, which allows for both right-censoring and left-truncation. In the presence of left-truncation, or put differently, delayed study entry, the data points $(T_i, \delta_i \cdot \varepsilon_i, L_i, \mathbf{W}_i)$ are drawn from the conditional probability distribution given study entry. Technically speaking, this means that the underlying probability space is not $(\Omega, \mathcal{F}, \mathbb{P})$ as in Settings A and B but rather the space $(\Omega, \mathcal{F}, \mathbb{P}^B)$, where \mathbb{P}^B is the conditional measure given study entry defined by $\mathbb{P}^B(A) = \mathbb{P}(A \cap B)/\mathbb{P}(B)$ for all $A \in \mathcal{F}$ with $B = \bigcap_{i=1}^n \{T_i^* > L_i\}$. We thus replace the unconditional measure \mathbb{P} by the conditional one \mathbb{P}^B . In particular, all probabilistic statements in the following proposition are with respect to \mathbb{P}^B .

Proposition 4.6. *Consider a specific event $r \in \{1, \dots, R\}$ in the competing risks model from Setting C and assume the following:*

(C_C1) *The data points $(T_i, \delta_i \cdot \varepsilon_i, L_i, \mathbf{W}_i)$ are i.i.d. across i .*

(C_C2) *$\max_{1 \leq k \leq n} \mathbb{P}^B(t_{k-1} < T_i^* \leq t_k) \leq C/n$ with some sufficiently large constant C .*

(C_C3) *The components N_{ir} of the multivariate counting process $N_{1:n,r} = (N_{1r}, \dots, N_{nr})$ have $\{\mathcal{F}_t\}$ -intensities $\lambda_{ir}(t) = \mathbf{1}(L_i < t \leq T_i) \exp(\boldsymbol{\beta}_r^\top \mathbf{W}_i) \alpha_r^*(t)$.*

(C_C4) *$\mathbb{P}^B(L_i < \tau_{\min} \wedge T_i \geq \tau_{\max}) = p$ with some constant $p \in (0, 1)$.*

(C_C5) *The dimension d of the covariate vector \mathbf{W}_i grows at most polynomially in the sample size n , that is, $d = O(n^a)$ with some arbitrarily large but fixed constant a . Moreover, the covariates \mathbf{W}_i have bounded support, that is, $\|\mathbf{W}_i\|_\infty \leq C_W$ for all i and some constant $C_W < \infty$.*

(C_C6) *It holds that $\|\boldsymbol{\beta}_r\|_1 \leq C_\beta < \infty$ for some constant $C_\beta < \infty$ and $\hat{\boldsymbol{\beta}}_r$ is any estimator of $\boldsymbol{\beta}_r$ with the property that $\|\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_r\|_1 = O_p(\sqrt{\log(n)/n})$.*

Then (C1)–(C7) are satisfied for the event r with ν as large as desired and $\rho_n = \log(n)/\sqrt{n}$.

The assumptions of Proposition 4.6 closely parallel those required for the analysis of Setting B in Proposition 4.5. We thus only give a brief comment on (C_C4). Compared to (C_B4), the statement in (C_C4) is a bit more complicated. This stems from the fact that we now also need to deal with left-truncation. However, the underlying idea is the same: (C_C4) ensures that the data on the interval $[\tau_{\min}, \tau_{\max}]$ are sufficiently dense to

perform reasonable estimation.

We finally analyze the Markov model with state space $\mathcal{R} = \{0, \dots, R\}$ from Setting D.

Proposition 4.7. *Consider a specific $\ell \rightarrow m$ transition with $\ell, m \in \mathcal{R}$, $\ell \neq m$, in the multi-state model from Setting D and assume the following:*

(C_{D1}) *For each i , X_i is a time-continuous Markov process with absolutely continuous intensity measure, right-continuous sample paths and $\mathbb{P}(X_i(0) = 0) = 1$. The processes X_i are i.i.d. across i .*

(C_{D2}) *It holds that $\mathbb{P}(X_i(t-) = \ell \text{ for all } t \in [\tau_{\min}, \tau_{\max}]) = p_\ell$ with some constant $p_\ell \in (0, 1)$.*

Then (C1)–(C6) are satisfied for the considered $\ell \rightarrow m$ transition with ν as large as desired. Moreover, (C7) is trivially satisfied with $\rho_n = 0$.

As the assumptions in (C_{D1}) are standard in the literature, we only comment on (C_{D2}). According to (C_{D2}), with positive probability, subject i is in state ℓ at (or more precisely, just before) time point t for any $t \in [\tau_{\min}, \tau_{\max}]$. This ensures that there are sufficiently many subjects in the risk set at any $t \in [\tau_{\min}, \tau_{\max}]$ to perform estimation. (C_{D2}) thus plays the same role as (C_{B4}) and (C_{C4}) in the results for Settings B and C.

5 Implementation

The estimator $\hat{\alpha}_\lambda$ depends on the tuning parameter λ . In what follows, we derive a data-driven rule for selecting λ in practice.

Cross-validation is presumably the most prominent technique for selecting the tuning parameters of lasso-type estimators. In our context, however, it should be treated with caution. The main reason is that our data (i.e., the constructed variables Y_j) are not independent but have a quite complicated dependence structure. As is well-known (see e.g. Altman, 1990; Opsomer et al., 2001; Rabinowicz and Rosset, 2022), cross-validation may perform very poorly when the data are dependent, and adjustments are needed to deal with the specific dependence structure at hand. Hence, simply using cross-validation algorithms off-the-shelf will most likely lead to unsatisfactory results in our case. Indeed, we have run different versions of cross-validation in our simulation exercises and have persistently obtained extremely bad results.

We thus pursue a different strategy for tuning parameter calibration. Our strategy exploits the fact that the fused lasso can be expressed in standard lasso form as follows: We first reparametrize the model $\mathbf{Y} = \boldsymbol{\alpha}^* + \mathbf{u}$, where $\mathbf{Y} = (Y_1, \dots, Y_n)^\top$, $\boldsymbol{\alpha}^* = (\alpha_1^*, \dots, \alpha_n^*)^\top$ and $\mathbf{u} = (u_1, \dots, u_n)^\top$. In particular, we set $\theta_1 = \alpha_1^*$ and $\theta_j = \alpha_j^* - \alpha_{j-1}^*$ for $2 \leq j \leq n$. This implies that $\boldsymbol{\alpha}^* = X\boldsymbol{\theta}$, where $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)^\top$ and the design matrix $X = (X_{ij} : 1 \leq i, j \leq n)$ has the entries $X_{ij} = \mathbf{1}(i \geq j)$. With this reparametrization,

we can reformulate the model as $\mathbf{Y} = X\boldsymbol{\theta} + \mathbf{u}$. The reparametrized fused lasso is

$$\hat{\boldsymbol{\theta}}_\lambda = (\hat{\theta}_{\lambda,1}, \dots, \hat{\theta}_{\lambda,n})^\top \in \operatorname{argmin}_{\boldsymbol{\vartheta} \in \mathbb{R}^n} \left\{ \frac{1}{n} \|\mathbf{Y} - X\boldsymbol{\vartheta}\|_2^2 + \lambda \|\boldsymbol{\vartheta}_{-1}\|_1 \right\}, \quad (5.1)$$

where $\boldsymbol{\vartheta}_{-1}$ denotes the vector $\boldsymbol{\vartheta}$ without the first component. The estimator $\hat{\boldsymbol{\theta}}_\lambda$ differs from the standard lasso only in the fact that the first coefficient is not penalized (that is, the penalty is $\|\boldsymbol{\vartheta}_{-1}\|_1$ rather than $\|\boldsymbol{\vartheta}\|_1$). Next let $\mathbf{Y}^c = (Y_1^c, \dots, Y_n^c)^\top$ be defined by $Y_i^c = Y_i - \bar{Y}$ with $\bar{Y} = n^{-1} \sum_{i=1}^n Y_i$ and let $X^c = (\mathbf{X}_2^c, \dots, \mathbf{X}_n^c)$ be the $n \times (n-1)$ matrix whose columns $\mathbf{X}_j^c = (X_{1j}^c, \dots, X_{nj}^c)^\top$ have the entries $X_{ij}^c = X_{ij} - \bar{X}_j$ with $\bar{X}_j = n^{-1} \sum_{i=1}^n X_{ij}$. Straightforward arguments show that the minimization problem (5.1) is equivalent to

$$\hat{\boldsymbol{\theta}}_{\lambda,-1} = (\hat{\theta}_{\lambda,2}, \dots, \hat{\theta}_{\lambda,n})^\top \in \operatorname{argmin}_{\boldsymbol{\vartheta} \in \mathbb{R}^{n-1}} \left\{ \frac{1}{n} \|\mathbf{Y}^c - X^c\boldsymbol{\vartheta}\|_2^2 + \lambda \|\boldsymbol{\vartheta}\|_1 \right\} \quad (5.2)$$

and $\hat{\theta}_{\lambda,1} = \bar{Y} - \sum_{j=2}^n \bar{X}_j \hat{\theta}_{\lambda,j}$. Importantly, (5.2) is a standard lasso problem. Put differently, $\hat{\boldsymbol{\theta}}_{\lambda,-1}$ is the standard lasso in the model $\mathbf{Y}^c = X^c\boldsymbol{\theta}_{-1} + \mathbf{u}^c$, where \mathbf{u}^c is defined analogously as \mathbf{Y}^c .

We now apply techniques for tuning parameter calibration to the standard lasso $\hat{\boldsymbol{\theta}}_{\lambda,-1}$ defined in (5.2). In particular, we work with techniques based on the lasso's effective noise. In the linear model $\mathbf{Y}^c = X^c\boldsymbol{\theta}_{-1} + \mathbf{u}^c$, the effective noise is given by $2\|(X^c)^\top \mathbf{u}^c\|_\infty/n$. Theory for the lasso suggests to choose the tuning parameter λ as a high quantile of the effective noise; see e.g. Belloni and Chernozhukov (2013) and Lederer and Vogt (2021). These quantiles, however, are usually not known in practice (as the distribution of the error terms \mathbf{u}^c is unknown). Nevertheless, they can be approximated in many cases. Hence, a general strategy for selecting λ is this: Approximate a high quantile (say the 90%-quantile) of the effective noise and then set λ equal to the computed quantile.

This strategy can be realized as follows in our context: Straightforward calculations show that the effective noise $2\|(X^c)^\top \mathbf{u}^c\|_\infty/n$ is equal to the maximum statistic

$$U_n := 2 \max_{2 \leq j \leq n} \left| -\frac{1}{n} \sum_{i < j} u_i + \frac{j-1}{n^2} \sum_{i=1}^n u_i \right|.$$

In order to approximate the quantiles of U_n , we run the following multiplier (or wild) bootstrap: We first compute the lasso residuals $\hat{u}_{\lambda_0,i} = Y_i - \hat{\alpha}_{\lambda_0,i}$ for $i = 1, \dots, n$, where λ_0 is a pilot tuning parameter whose choice is explained below. We then draw standard normal random vectors $\boldsymbol{\epsilon}^{(\ell)} = (\epsilon_1^{(\ell)}, \dots, \epsilon_n^{(\ell)})$ for $\ell = 1, \dots, L$, where L is a large number (say $L = 1000$). The values $\{\hat{u}_{\lambda_0,i} \cdot \epsilon_i^{(\ell)} : i = 1, \dots, n\}$ serve as the ℓ -th bootstrap sample

of the error terms $\{u_i : i = 1, \dots, n\}$. We next compute

$$U_n^{(\ell)} := 2 \max_{2 \leq j \leq n} \left| -\frac{1}{n} \sum_{i < j} \hat{u}_{\lambda_0, i} \epsilon_i^{(\ell)} + \frac{j-1}{n^2} \sum_{i=1}^n \hat{u}_{\lambda_0, i} \epsilon_i^{(\ell)} \right| \quad (5.3)$$

for each ℓ and determine the empirical q -quantile $z_n(q)$ of the bootstrap sample $\{U_n^{(\ell)} : \ell = 1, \dots, L\}$. Regarding $z_n(q)$ as an approximation of the theoretical q -quantile of U_n , we finally set $\lambda = z_n(q)$.

It remains to choose the starting value λ_0 in the above bootstrap procedure. The chosen value needs to be such that the lasso residuals $\hat{u}_{\lambda_0, i}$ give a reasonable approximation to the error terms u_i . This requires that the estimator $\hat{\alpha}_{\lambda_0}$ is reasonably close to α^* . However, it is not necessary that the number of jumps in $\hat{\alpha}_{\lambda_0}$ is close to the number of jumps in α^* . Indeed, even if $\hat{\alpha}_{\lambda_0}$ has many spurious small jumps which do not correspond to any real jump in α^* , the discretized curve $\hat{\alpha}_{\lambda_0}$ may still be close in shape to α^* . Hence, somewhat overestimating the number of change points in α^* should not do any harm. In most applications, it is natural to assume that the discretized hazard rate α^* has a rather small number of change points K . We may thus choose λ_0 such that $\hat{\alpha}_{\lambda_0}$ has a moderately large number of jumps K_{\max} (say K_{\max} between 10 and 20) which most likely exceeds the true number of jumps K a bit. This should equip us with an acceptable initial estimator $\hat{\alpha}_{\lambda_0}$ of the discretized hazard α^* .

In summary, our algorithm for selecting λ is as follows:

- Step 1. Set λ_0 to be the largest value such that the resulting estimator $\hat{\alpha}_{\lambda_0}$ has K_{\max} change points and compute the lasso residuals $\hat{u}_{\lambda_0, i} = Y_i - \hat{\alpha}_{\lambda_0, i}$ for $i = 1, \dots, n$.
- Step 2. For $\ell = 1, \dots, L$, draw standard normal random vectors $\epsilon^{(\ell)} = (\epsilon_1^{(\ell)}, \dots, \epsilon_n^{(\ell)})$, compute $U_n^{(\ell)}$ as defined in (5.3) and determine the empirical q -quantile $z_n(q)$ of the computed values.
- Step 3. Set $\lambda = z_n(q)$ with q close to 1.

6 Simulation study

In this section, we investigate the performance of our methods by Monte Carlo experiments. We simulate data from different designs that fit into Settings A and B:

- We consider two different piecewise constant hazard functions whose form is loosely inspired by the estimated hazard functions in the application of Section 7:

$$\begin{aligned} \alpha_{[1]}^*(t) &= 4 \cdot \mathbf{1}(t < 0.25) + 1 \cdot \mathbf{1}(t \geq 0.25) \\ \alpha_{[2]}^*(t) &= 4 \cdot \mathbf{1}(t < 0.2) + 1.5 \cdot \mathbf{1}(0.2 \leq t < 0.6) + 0.5 \cdot \mathbf{1}(t \geq 0.6). \end{aligned}$$

| | hazard | sample size | covariates |
|-------------|------------------|-----------------------------|------------|
| Scenario A1 | $\alpha_{[1]}^*$ | $n \in \{500, 1000, 2000\}$ | no |
| Scenario B1 | $\alpha_{[1]}^*$ | $n \in \{500, 1000, 2000\}$ | yes |
| Scenario A2 | $\alpha_{[2]}^*$ | $n \in \{500, 1000, 2000\}$ | no |
| Scenario B2 | $\alpha_{[2]}^*$ | $n \in \{500, 1000, 2000\}$ | yes |

Table 1: List of simulation scenarios.

- We either consider no covariates at all (Setting A) or a Cox model with two covariates $\mathbf{W}_i = (W_{i1}, W_{i2})^\top$ as defined in Setting B. The first covariate W_{i1} is a binary variable which takes the values -1 and 1 with probability 0.5 each, whereas the second covariate W_{i2} is continuous, in particular, uniformly distributed on $[-1, 1]$. The parameter vector is chosen as $\boldsymbol{\beta} = (\beta_1, \beta_2)^\top = (0.25, 1)^\top$.
- The right-censoring variables C_i are assumed to be independent from (T_i^*, \mathbf{W}_i) and are drawn from an exponential distribution with parameter 0.5 , resulting in around 20-25% censored data points in all considered simulation scenarios.
- For estimation, we consider the unit interval $[\tau_{\min}, \tau_{\max}] = [0, 1]$ and three different sample sizes $n \in \{500, 1000, 2000\}$.

The resulting simulation scenarios are listed in Table 1. Estimating the piecewise constant hazard in these scenarios is a non-trivial problem. This is reflected by fairly low signal-to-noise ratios in the regression model (3.3) which underlies our fused lasso method. The signal-to-noise ratio in model (3.3) is defined as the empirical variance of the signal vector $\boldsymbol{\alpha}^* = (\alpha_1^*, \dots, \alpha_n^*)^\top$ divided by that of the noise vector $\mathbf{u} = (u_1, \dots, u_n)^\top$. Formally, $\text{SNR} = \hat{V}(\boldsymbol{\alpha}^*)/\hat{V}(\mathbf{u})$, where $\hat{V}(\mathbf{w})$ is the empirical variance of a generic vector $\mathbf{w} = (w_1, \dots, w_n)^\top$. In the considered simulation scenarios, the SNR (averaged over 1000 simulation runs) ranges between 0.23 and 0.33. This shows that there is much more variation in the noise term than in the signal, implying that it is quite difficult to estimate the signal in this model.

Throughout the simulation study, we implement our estimator as described in Section 5 with $q = 0.9$, $K_{\max} = 20$ and $L = 100$. The results of our simulation experiments are reported in Figures 3–6 and Tables 2–3. As they are qualitatively very similar across the considered scenarios, we proceed as follows: we discuss the results for one of the scenarios – in particular, for Scenario A2 – in detail while keeping the discussion of the other scenarios rather short.

To start with, we have a closer look at a representative simulation run in Scenario A2 (with $n = 1000$). Figure 2 summarizes the results of this run. The red curve in the left-hand panel is the Breslow (i.e., Nelson-Aalen) estimator $\hat{A}_{[2]}(t)$ of the cumulative hazard $A_{[2]}^*(t) = \int_0^t \alpha_{[2]}^*(s) ds$. The middle panel shows the estimator $\hat{A}_{[2]}$ once again




Figure 2: Results of a representative simulation run in Scenario A2 with $n = 1000$. Left: Breslow (i.e., Nelson-Aalen) estimator. Middle: cumulative hazard in black and Breslow estimator in red. Right: hazard rate in black and its fused lasso estimator in red.

(in red) together with the piecewise linear cumulative hazard $A_{[2]}^*$ (in black), where the kinks of $A_{[2]}^*$ at $t = 0.2$ and $t = 0.6$ are indicated by vertical dotted lines. The right-hand panel depicts the hazard $\alpha_{[2]}^*$ in black and its fused lasso estimator in red. The left-hand (and middle) panel of Figure 2 demonstrate that it is not so easy to spot the piecewise linear structure of the function $A_{[2]}^*$ in the estimator $\hat{A}_{[2]}$. This illustrates that estimating the kinks and slope parameters of $A_{[2]}^*$ – or equivalently, the change points and function values of $\alpha_{[2]}^*$ – is not trivial at all. Nevertheless, the fused lasso estimator (depicted in the right-hand panel in red) does a pretty good job in reconstructing the piecewise constant hazard $\alpha_{[2]}^*$ (depicted in black). Several typical features of the estimator become visible here:

- (i) The fused lasso has a bias. This has two reasons: First, as any other member of the lasso family, it is a shrinkage estimator and thus is biased by construction. Second, our fused lasso method is based on the Breslow estimator. Hence, any imprecisions of this estimator are “inherited” by our method. In the case at hand, the Breslow estimator is sloping upwards more strongly than the cumulative hazard after $t = 0.6$. This induces a substantial upward bias in our fused lasso estimator after $t = 0.6$.
- (ii) The fused lasso often reconstructs jumps in the underlying hazard by a sequence of consecutive jumps rather than one big jump. In the case at hand, the fused lasso does not mimic the true change point at $t = 0.2$ by one big jump but rather by two jumps in quick succession.
- (iii) The fused lasso may have additional jumps far away from the true change points. However, these jumps tend to be small and thus negligible. In the case at hand, the fused lasso has additional small jumps between the two true change points at $t = 0.2$ and $t = 0.6$.

We will re-encounter these features of the fused lasso when discussing the results of our simulation experiments below.

Figure 5 depicts the results of 100 simulation runs in Scenario A2 for different sample sizes. The upper three panels show the 100 estimated hazard rates $\hat{\alpha}_\lambda$ (in grey) together with the true underlying hazard (in black). The lower three panels give the locations and jump heights of the change points estimated over the 100 simulation runs. Each grey circle represents one estimated change point. The x -axis specifies the location of the change point, whereas the y -axis gives the jump height. The larger red circles correspond to the location and jump height of the two true change points of the underlying hazard. As can be seen from the plots, the fused lasso estimates capture the piecewise constant step structure of the hazard pretty well, the approximation getting better as the sample size increases. The three features discussed in the context of Figure 2 also become visible here:

- (i) The fused lasso estimates have a bias which is most pronounced for $n = 500$ and gets smaller as the sample size increases. As already discussed, this bias has two sources: the shrinkage done by the fused lasso estimator and bias inherited from the Breslow estimator.
- (ii) The lower three panels show that many estimated jumps are located fairly close to the two true change points at $t = 0.2$ and $t = 0.6$. However, most of these jumps are substantially smaller in size than the true jumps (indicated by the two red dots). The reason is that the lasso often uses several smaller jumps in quick succession to reconstruct the true jumps at $t = 0.2$ and $t = 0.6$.
- (iii) Some of the fused lasso estimates have additional spurious jumps quite distant from the true change points $t = 0.2$ and $t = 0.6$. These jumps, however, tend to be comparatively small. (See also the plots for Scenarios A1 and B1 which give an even better illustration of this point.)

Having summarized the main simulation results for Scenario A2, we now briefly comment on the findings for the other scenarios. Scenario B2 differs from A2 only by incorporating covariates. The results for Scenario B2 in Figure 6 are very similar to those for A2 even though somewhat less precise. This is not surprising because the estimation of the parameter vector β produces additional approximation error, which gets reflected in somewhat less precise estimation results. The results for Scenarios A1/B1 (depicted in Figures 3 and 4) are qualitatively very similar to those for Scenarios A2/B2. In particular, the main points of our discussion apply to Scenarios A1/B1 equally well as to Scenarios A2/B2.

Tables 2 and 3 complement the visual results from Figures 3–6 by summary statistics. All reported numbers are based on 1000 simulation runs. Table 2 gives information on the approximation quality of the estimator $\hat{\alpha}_\lambda$: it reports the average of the squared ℓ_2 -error $\|\hat{\alpha}_\lambda - \alpha^*\|_2^2/n$ over the 1000 simulation runs along with the standard deviation in brackets. As expected, the average error (as well as the standard deviation) gets smaller

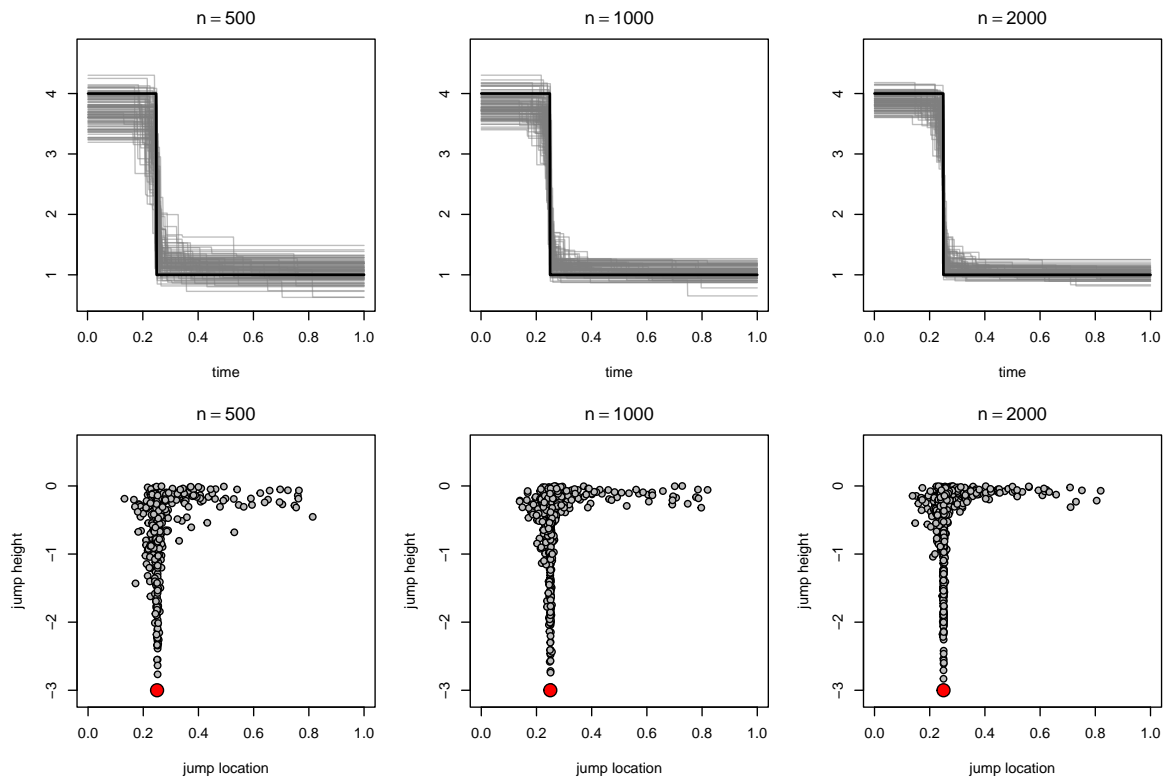


Figure 3: Simulation results for Scenario A1.

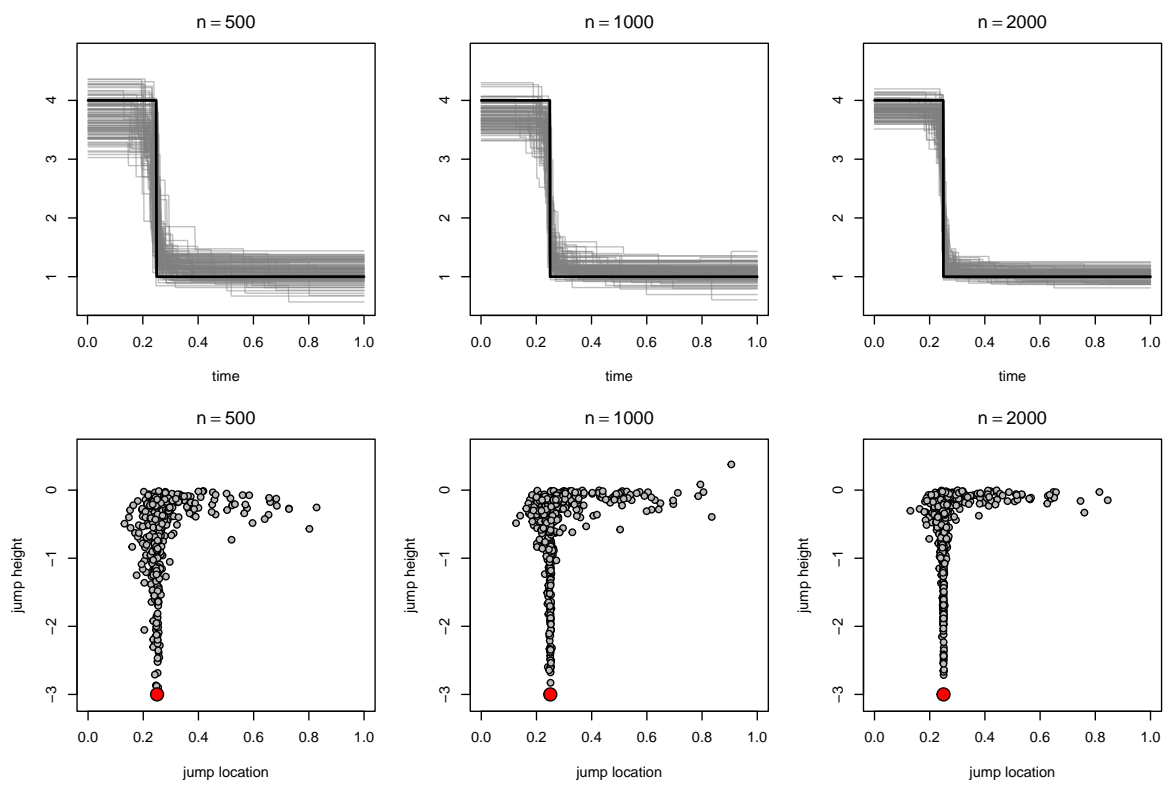


Figure 4: Simulation results for Scenario B1.

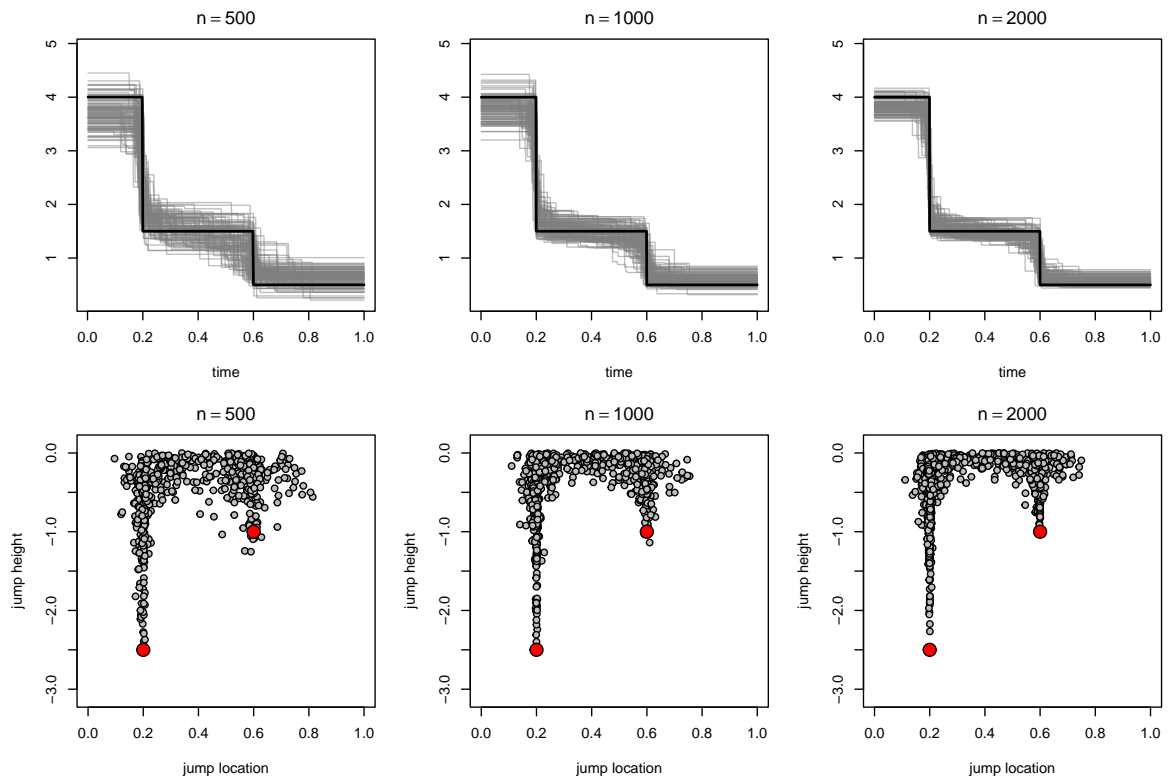


Figure 5: Simulation results for Scenario A2.

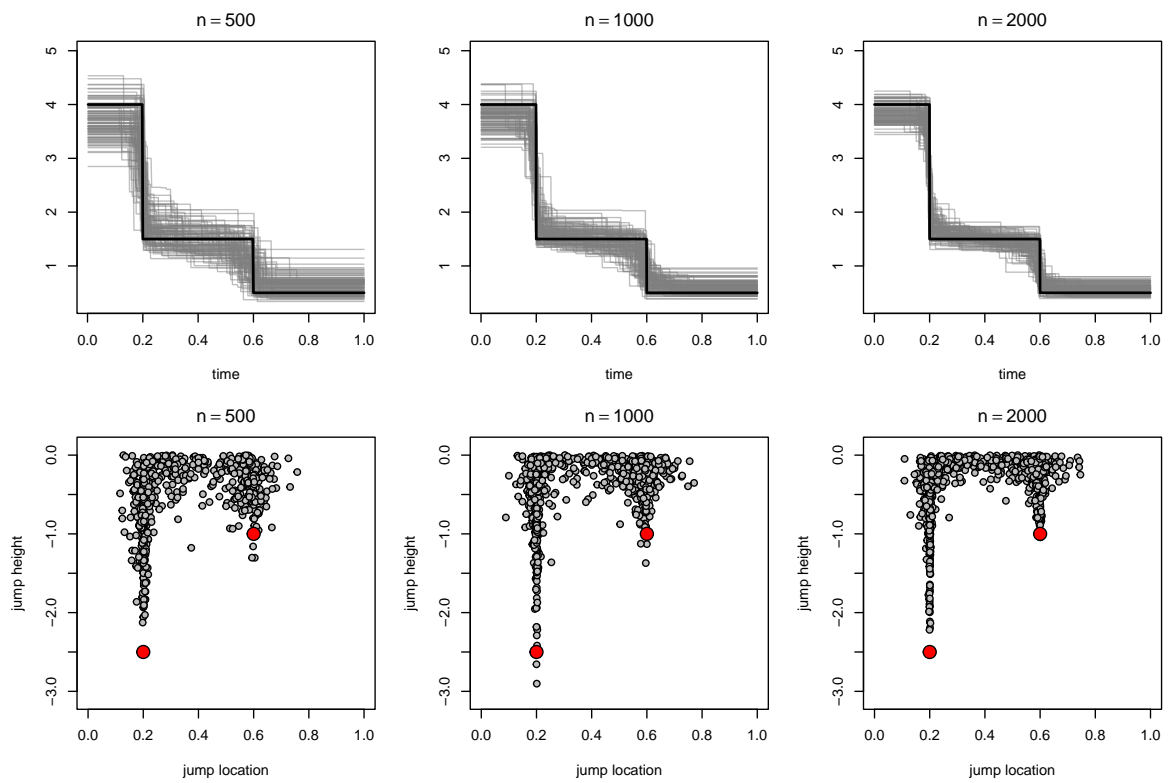


Figure 6: Simulation results for Scenario B2.

| | Scenario A1 | Scenario B1 | Scenario A2 | Scenario B2 |
|------------|---------------|---------------|---------------|---------------|
| $n = 500$ | 0.102 (0.062) | 0.124 (0.080) | 0.118 (0.056) | 0.134 (0.071) |
| $n = 1000$ | 0.058 (0.033) | 0.071 (0.042) | 0.069 (0.032) | 0.078 (0.039) |
| $n = 2000$ | 0.031 (0.017) | 0.037 (0.021) | 0.038 (0.017) | 0.044 (0.019) |

Table 2: Average values of the squared ℓ_2 -error $\|\hat{\alpha}_\lambda - \alpha^*\|_2^2/n$ with standard deviations in brackets.

| | Scenario A1 | Scenario B1 | Scenario A2 | Scenario B2 |
|------------|---------------|---------------|---------------|---------------|
| $n = 500$ | 0.002 (0.003) | 0.003 (0.003) | 0.016 (0.022) | 0.019 (0.026) |
| $n = 1000$ | 0.001 (0.001) | 0.001 (0.002) | 0.008 (0.009) | 0.009 (0.011) |
| $n = 2000$ | 0.001 (0.001) | 0.001 (0.001) | 0.004 (0.005) | 0.004 (0.005) |

Table 3: Average values of the asymmetric distance measure $d(\hat{\mathcal{S}}_\lambda, \mathcal{S}^*)$ with standard deviations in brackets.

as the sample size increases, implying that the approximation quality gets better with growing sample size. Table 3 is concerned with change point estimation. It reports the average value of the asymmetric distance measure $d(\hat{\mathcal{S}}_\lambda, \mathcal{S}^*)$ together with the standard deviation (in brackets). The distance measure $d(\hat{\mathcal{S}}_\lambda, \mathcal{S}^*)$ gives the following information: if $d(\hat{\mathcal{S}}_\lambda, \mathcal{S}^*)$ equals δ , then for each true change point, we can find an estimated change point which is not further away than δ . The average values and standard deviations of $d(\hat{\mathcal{S}}_\lambda, \mathcal{S}^*)$ in Table 3 are very small already for $n = 500$ and get even smaller with increasing sample size. Hence, our estimator tends to have change points very close to the true ones. This, however, does not exclude the existence of additional “spurious” change points far away from the true ones. According to Theorem 4.4, such “spurious” change points are negligible in the sense of having small jump size. As turns out, it is quite difficult to underpin this theoretical finding numerically as this would require to define precisely what is meant by a “spurious” change point. For instance, one could call an estimated change point “spurious” if it is further away from any true change point than a pre-specified threshold value, say 0.1. As such a definition is quite adhoc, we have not attempted anything of this kind. We rather think that the visual results in Figures 3–6 speak for themselves: they illustrate quite clearly that “spurious” jumps far away from the true change points tend to be comparatively small.

7 Empirical application

7.1 Background

The motivation behind the present empirical illustration is clinical trial planning in oncology. The gold standard to demonstrate clinical benefit of a novel therapy is evidence of delaying time-to-death, called ‘overall survival’ (OS) in the field. A second endpoint that is both observed earlier and used to demonstrate treatment benefit is

progression-free survival (PFS), i.e., the time until (diagnosed) progression of disease or death, whatever occurs first. In fact, it is not uncommon for contemporary trials to evaluate both endpoints, with PFS benefit leading to accelerated drug approval and authorities demanding subsequent follow-up of OS.

Trials are planned in terms of the OS and/or PFS number of events to be observed, typically assuming proportional hazards between treatment groups or even exponential distributions (Ohneberg and Schumacher, 2014; Oba and Kuchiba, 2019). If OS and PFS are co-primary endpoints, a common approach is to plan the number of events to be observed independently for OS and PFS, assuming proportional hazards. In what follows, we use the methodology developed earlier to come up with a more flexible and realistic trial design. We in particular propose a flexible parametric multi-state model with piecewise constant transition hazards to jointly model PFS and OS. We then fit the proposed model to clinical data from a randomised controlled trial in patients with non-small-cell lung cancer (Rittmeyer et al., 2017). Trial planning (e.g. calculating the number of events to be observed) could then be performed by simulating from the fitted model. In what follows, we focus on model fitting, in particular, on estimating the piecewise constant transition hazards, but we do not carry out any simulation exercises for trial planning. We refer to Beyer et al. (2020) and Danzer et al. (2022) as trial design examples using simulation in a multi-state context. Notably, we are not interested in a specific interpretation of the change points in the estimated transition hazards (as opposed to other studies in the change point literature such as Brazzale et al., 2019). The main motivation for fitting a piecewise constant hazard model is that this yields a flexible yet simple parametrization which is well suited for trial planning.

7.2 Multi-state modelling of PFS and OS

Following Meller et al. (2019), we model PFS and OS jointly via an illness-death model without recovery. Formally speaking, for each patient i in our sample, let X_i be a time-continuous Markov process with state space $\mathcal{R} = \{0, 1, 2\}$ and possible transitions $0 \rightarrow 1$, $0 \rightarrow 2$ and $1 \rightarrow 2$. We interpret state 0 as the initial state of being alive without progression, state 2 as the terminal state of death and state 1 as an intermediate state entered upon progression diagnosis. In this multi-state model, PFS is viewed as the waiting time in the initial state 0 and OS as the waiting time until reaching the absorbing state 2, i.e., $\text{PFS}_i = \inf\{t \geq 0 : X_i(t) \neq 0\}$ and $\text{OS}_i = \inf\{t \geq 0 : X_i(t) = 2\}$ for each patient i . In general, PFS and OS do not both follow proportional hazards at the same time in this model. Hence, our multi-state framework naturally accounts for the common concern in the field (Mukhopadhyay et al., 2022) that proportional hazards for both OS and PFS appear to be too simple for adequate modelling.

We do not observe the Markov processes X_i completely but right-censored versions of them, where C_i denotes the right-censoring time for patient i . In order to estimate the




Figure 7: Scatterplot of transition times in the atezolizumab arm of the data from Rittmeyer et al. (2017). Uncensored data points are depicted in black, right-censored points in grey. The x -axis displays arbitrary unit IDs, the y -axis displays time (in months).

transition hazards $\alpha_{0 \rightarrow 1}^*$, $\alpha_{0 \rightarrow 2}^*$ and $\alpha_{1 \rightarrow 2}^*$ by our methods, we can treat each transition ($0 \rightarrow 1$, $0 \rightarrow 2$ and $1 \rightarrow 2$) separately and reformulate the data in terms of the survival framework from Setting A as follows:

- Transition $0 \rightarrow 1$: $T_i^* = \text{PFS}_i$, $T_i = T_i^* \wedge C_i$ and $\delta_i = \mathbf{1}(T_i^* \leq C_i \text{ and } \text{PFS}_i < \text{OS}_i)$.
- Transition $0 \rightarrow 2$: $T_i^* = \text{PFS}_i$, $T_i = T_i^* \wedge C_i$ and $\delta_i = \mathbf{1}(T_i^* \leq C_i \text{ and } \text{PFS}_i = \text{OS}_i)$.
- Transition $1 \rightarrow 2$: $T_i^* = \text{OS}_i$, $T_i = T_i^* \wedge C_i$ and $\delta_i = \mathbf{1}(T_i^* \leq C_i)$. As patient i is at risk only after entering state 1, i.e., from time PFS_i onwards, we also have to deal with left-truncation. To do so, we only consider the subsample of patients i for whom $\text{PFS}_i < T_i$, thus treating PFS_i as a left-truncation time L_i .

Figure 7 is a scatterplot of the observed transition times T_i for the different transitions in the analyzed data set. The left-hand and middle plots indicate a possibly difficult estimation problem in that there are few observed $0 \rightarrow 2$ transitions and many rather early transitions into the intermediate state 1.

7.3 Re-analysis of clinical trial data

We consider data from Rittmeyer et al. (2017), available as supplement to Gandara et al. (2018), on 850 patients randomized in one-to-one fashion to either atezolizumab or docetaxel. In what follows, we restrict attention to the atezolizumab arm. (For the docetaxel arm, we found similar results not reported here.) As discussed by Rittmeyer et al. (2017), atezolizumab is a novel immunotherapy that aims at reestablishing anti-cancer immunity, while docetaxel is a more traditional therapy.

Our aim is to fit a parsimonious piecewise constant hazards parameterization of the illness-death model from above to the data at hand. The estimated piecewise constant transition hazards $\alpha_{0 \rightarrow 1}^*$, $\alpha_{0 \rightarrow 2}^*$ and $\alpha_{1 \rightarrow 2}^*$ are then transformed into survival functions of PFS and OS, respectively, using well-known solutions to Kolmogorov forward differen-

tial equations (see Andersen et al., 1993, Section II.6), and are visually compared with standard Kaplan-Meier plots of PFS and OS.

To implement our methods, we make the following choices: (i) To demonstrate that our methods are robust to the particular choice of the interval $[\tau_{\min}, \tau_{\max}]$, we carry out the data analysis for different intervals: letting $\tau(p)$ be the empirical p -quantile of the uncensored event times $\{T_i : \delta_i = 1\}$ defined above for the different transitions, we set $\tau_{\min} = \tau(0)$ for the transitions $0 \rightarrow 1$ and $0 \rightarrow 2$, $\tau_{\min} = \tau(0.025)$ for transition $1 \rightarrow 2$ (taking into account the implicit left truncation) and $\tau_{\max} = \tau(p)$ with $p \in \{0.8, 0.975, 0.99\}$. Note that for $p = 0.8$, τ_{\max} and thus the estimation window is quite small. We have added the value $p = 0.8$ for illustrative purposes but do not recommend to use such a small value in practice. (ii) To evaluate the influence of the tuning parameter λ on our estimation results, we run the fused lasso with different choices of λ . In particular, λ is chosen as explained in Section 5 with $q \in \{0.1, 0.5, 0.9\}$ (as well as $K_{\max} = 20$ and $L = 1000$). The larger q , the larger the tuning parameter λ and thus the smaller the number of change points in the estimated hazard. Whereas the choice $q = 0.9$ is in line with our theoretical considerations, $q = 0.1$ and $q = 0.5$ are quite low choices. In the application context at hand, it makes sense to consider q -values and thus λ -values considerably smaller than suggested by the theory for the following reason: while resulting in somewhat less parsimonious piecewise constant models with a larger number of change points, they may produce a better curve fit.

Figures 8–10 present the estimated piecewise constant hazards (left column) and the corresponding piecewise linear cumulative hazards (right column) for different choices of p and q . Each figure is structured as follows: The plots in each row correspond to a specific choice of $q \in \{0.1, 0.5, 0.9\}$. The choice of p is indicated by colour. In particular, the green, blue, red curve in each plot corresponds to the choice $p = 0.8, 0.975, 0.99$, respectively. The dark grey line in the right-hand plots is the Nelson-Aalen estimator of the cumulative hazard, the shaded area in light grey is the corresponding pointwise 95% confidence band. We briefly discuss the estimation results: (i) The estimated hazards in Figures 8–10 suggest that a model with constant transition hazards is too simple. (ii) Overall, our approach produces a good fit to the data. This becomes visible in the piecewise linear estimates of the cumulative hazards which are very similar to their nonparametric Nelson-Aalen counterparts for all choices of p and q . (iii) The estimated hazard curves are fairly similar for different choices of p . More notable are the differences in the estimates across q . In particular, the larger q , the less pronounced the estimated jumps. This is not surprising as larger q amounts to more shrinkage. (iv) For the $0 \rightarrow 1$ transition, the Nelson-Aalen estimator has a staircase-like structure at earlier times (i.e., close to the time point $t = 0$). This is most likely due to the fact that progression diagnosis comes in 'waves' during earlier times, a common feature in oncology trials. Notably, our piecewise linear estimates do not pick up this staircase

Hazard for $0 \rightarrow 1$ transition




Figure 8: Estimated hazard and cumulative hazard for the $0 \rightarrow 1$ transition with different choices of p and q .

Hazard for $0 \rightarrow 2$ transition




Figure 9: Estimated hazard and cumulative hazard for the $0 \rightarrow 2$ transition with different choices of p and q .

Hazard for $1 \rightarrow 2$ transition

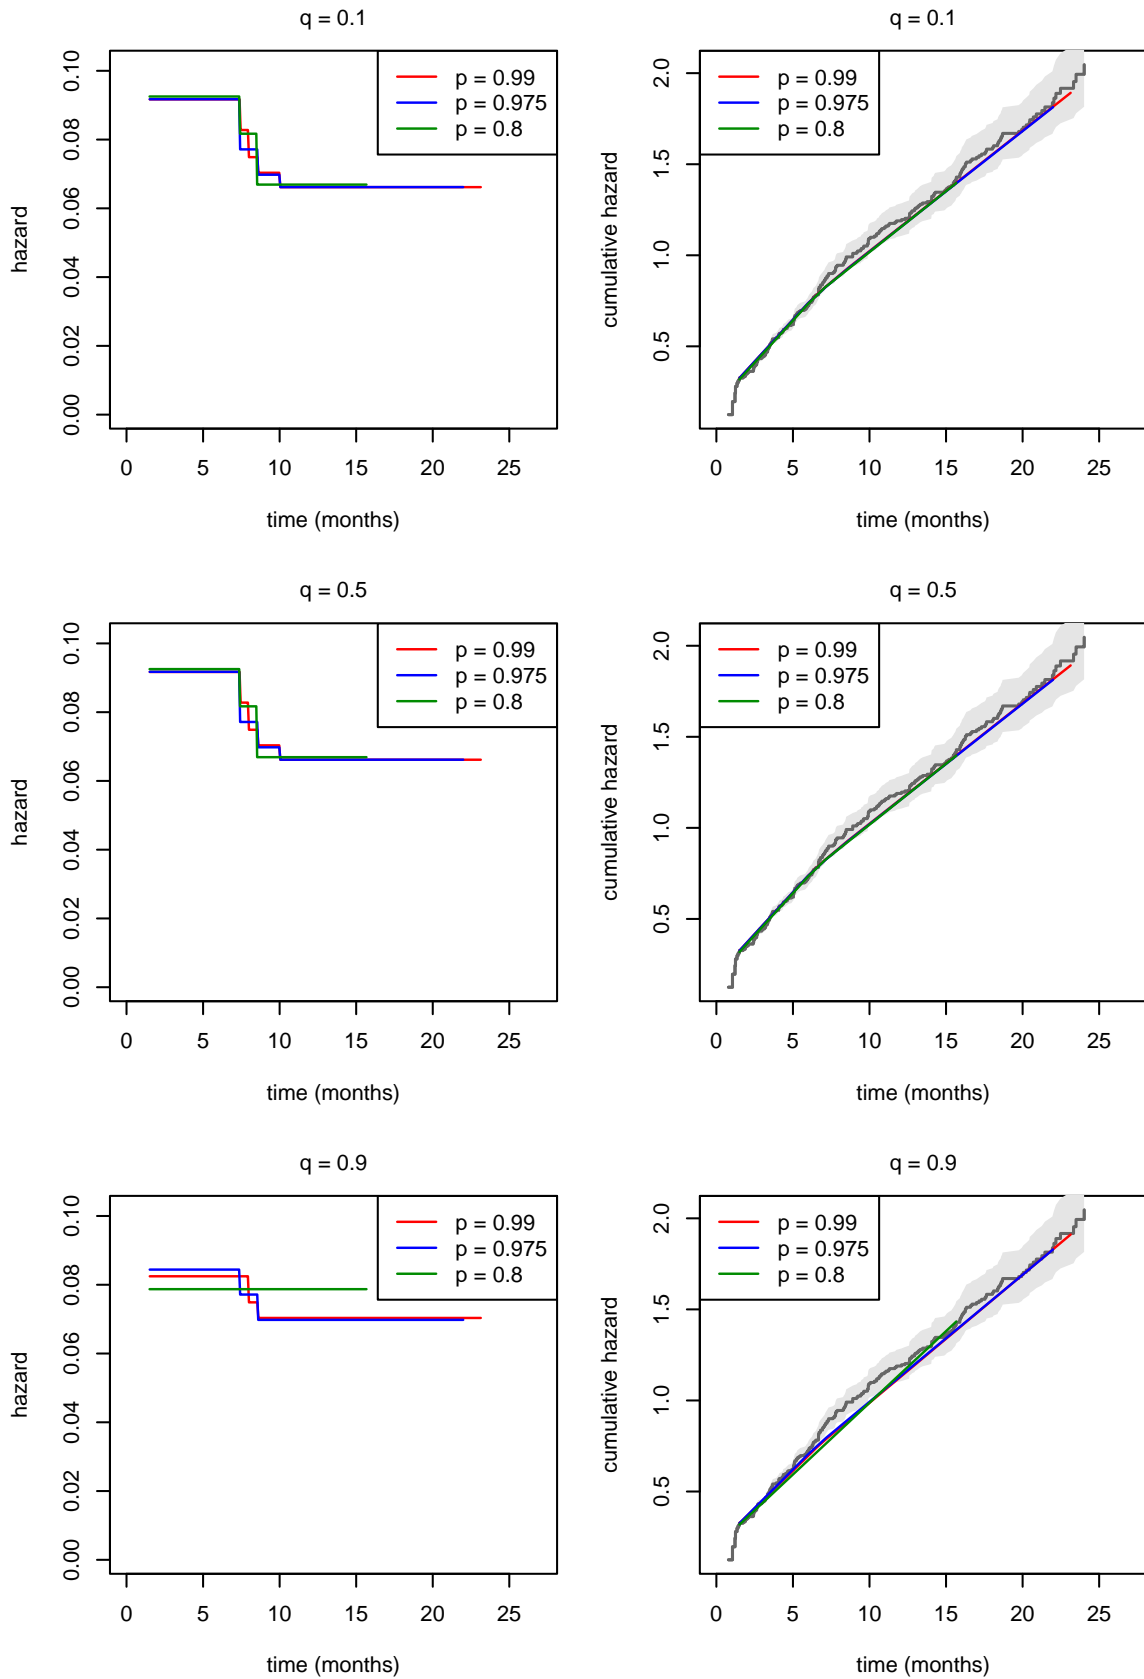


Figure 10: Estimated hazard and cumulative hazard for the $1 \rightarrow 2$ transition with different choices of p and q .

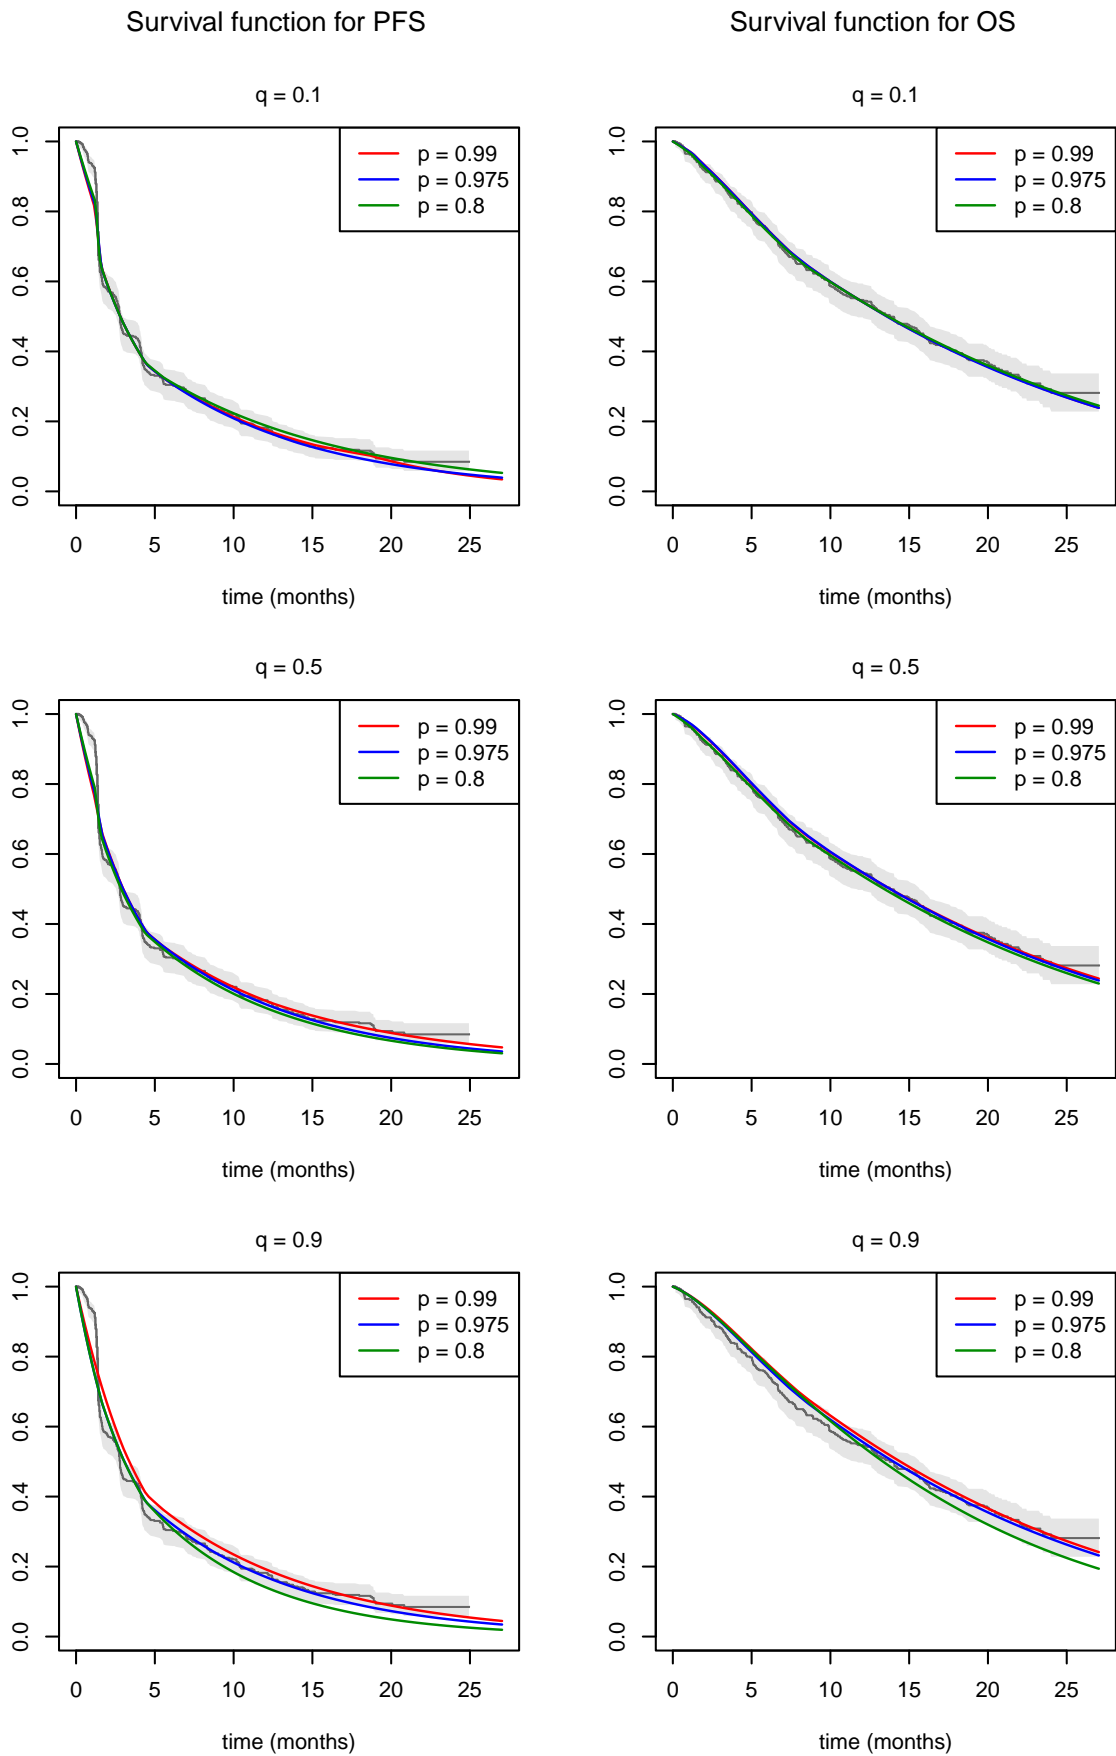


Figure 11: Estimated survival curves for PFS and OS.

structure but provide a smoother fit.

An aggregated assessment of the estimation results is offered by the survival function plots in Figure 11, which is structured analogously as the previous three figures. The plots depict estimators of the survival function for PFS (left column) and OS (right column) which are based on the estimated piecewise constant transition hazards. The dark grey line in the background is the Kaplan-Meier estimator of the survival function, the shaded area in light grey is the corresponding pointwise 95% confidence band. Titman and Sharples (2010) have suggested comparing the pointwise confidence intervals of the Kaplan-Meier estimator with the survival curves resulting from the parametric fit of the multi-state model as an informal model diagnostic. We here follow this suggestion: The estimated PFS survival curves yield an overall good fit to the non-parametric Kaplan-Meier estimator, the fit being somewhat less precise at earlier times (i.e., close to $t = 0$). Specifically, our estimates do not mimic the wave-like structure of the Kaplan-Meier estimator close to $t = 0$. This is not surprising as this wave-like structure is most probably due to the staircase structure of the Nelson-Aalen estimator for the $0 \rightarrow 1$ transition (discussed above) which is not picked up by our estimates. The estimated OS survival curves are quite close to the Kaplan-Meier estimator as well, the fit being somewhat worse for $q = 0.9$, in particular, for $(q, p) = (0.9, 0.8)$. To summarize, Figure 11 suggests that our method produces an overall adequate fit (except for the early 'wave' characteristic of the PFS Kaplan-Meier curve, which is a common phenomenon in oncology trials).

References

- AALEN, O., BORGAN, O. and GJESSING, H. (2008). *Survival and Event History Analysis: A Process Point of View*. Springer.
- ALTMAN, N. S. (1990). Kernel smoothing of data with correlated errors. *Journal of the American Statistical Association*, **85** 749–759.
- ANDERSEN, P., BORGAN, O., GILL, R. and KEIDING, N. (1993). *Statistical Models Based on Counting Processes*. Springer.
- ANDERSEN, P., POHAR PERME, M., VAN HOUWELINGEN, H., COOK, R., JOLY, P., MARTINUSSEN, T., TAYLOR, J., ABRAHAMOWICZ, M. and THERNEAU, T. (2021). Analysis of time-to-event for observational studies: guidance to the use of intensity models. *Statistics in Medicine*, **40** 185–211.
- ARJAS, E. and GASBARRA, D. (1994). Nonparametric Bayesian inference from right censored survival data, using the Gibbs sampler. *Statistica Sinica*, **4** 505–524.
- BELLONI, A. and CHERNOZHUKOV, V. (2013). Least squares after model selection in high-dimensional sparse models. *Bernoulli*, **19** 521–547.
- BENDER, A., GROLL, A. and SCHEIPL, F. (2018). A generalized additive model approach to time-to-event analysis. *Statistical Modelling*, **18** 299–321.

- BEYER, U., DEJARDIN, D., MELLER, M., RUFIBACH, K. and BURGER, H. U. (2020). A multistate model for early decision-making in oncology. *Biometrical Journal*, **62** 550–567.
- BRAZZALE, A. R., KÜCHENHOFF, H., KRÜGEL, S., SCHIERGENS, T. S., TRENTZSCH, H. and HARTL, W. (2019). Nonparametric change point estimation for survival distributions with a partially constant hazard rate. *Lifetime Data Analysis*, **25** 301–321.
- CHANG, I., CHEN, C. and HSIUNG, C. (1994). Estimation in change-point hazard rate models with random censorship. In *IMS Lecture Notes Monograph Series 23*. 78–92.
- COONEY, P. and WHITE, A. (2021). Change-point detection for piecewise exponential models. *arXiv:2112.03962*.
- COONEY, P. and WHITE, A. (2023). Extending beyond Bagust and Beale: Fully parametric piecewise exponential models for extrapolation of survival outcomes in health technology assessment. *Value in Health*, **26** 1510–1517.
- COX, D. R. (1972). Regression models and life tables (with discussion). *Journal of the Royal Statistical Society: Series B*, **34** 187–220.
- DANZER, M. F., TERZER, T., BERTHOLD, F., FALDUM, A. and SCHMIDT, R. (2022). Confirmatory adaptive group sequential designs for single-arm phase II studies with multiple time-to-event endpoints. *Biometrical Journal*, **64** 312–342.
- FLEMING, T. R. and HARRINGTON, D. P. (2005). *Counting Processes and Survival Analysis*. Wiley.
- GANDARA, D. R., PAUL, S. M., KOWANETZ, M., SCHLEIFMAN, E., ZOU, W., LI, Y., RITTMAYER, A., FEHRENBACHER, L., OTTO, G., MALBOEUF, C. ET AL. (2018). Blood-based tumor mutational burden as a predictor of clinical benefit in non-small-cell lung cancer patients treated with atezolizumab. *Nature Medicine*, **24** 1441–1448.
- GIJBELS, I. and GÜRLER, U. (2003). Estimation of a change point in a hazard function based on censored data. *Lifetime Data Analysis*, **9** 395–411.
- GOODMAN, M., LI, Y. and TIWARI, R. (2011). Detecting multiple change points in piecewise constant hazard functions. *Journal of Applied Statistics*, **38** 2523–2532.
- HAN, G., SCHELL, M. and KIM, J. (2014). Improved survival modeling in cancer research using a reduced piecewise exponential approach. *Statistics in Medicine*, **33** 59–73.
- HE, P., FANG, L. and SU, Z. (2013). A sequential testing approach to detecting multiple change points in the proportional hazards model. *Statistics in Medicine*, **32** 1239–1245.
- HENDERSON, R. (1990). A problem with the likelihood ratio test for a change-point hazard rate model. *Biometrika*, **77** 835–843.
- HUANG, J., SUN, T., YING, Z., YU, Y. and ZHANG, C.-H. (2013). Inequalities for the Lasso in the Cox model. *Annals of Statistics*, **41** 1142–1165.
- KOPPER, P., PÖLSTERL, S., WACHINGER, C., BISCHL, B., BENDER, A. and RÜGAMER, D. (2021). Semi-structured deep piecewise exponential models. *Proceedings of Machine Learning Research*.
- LEDERER, J. and VOGT, M. (2021). Estimating the lasso’s effective noise. *Journal of Machine*

- Learning Research*, **22** 1–32.
- LIN, K., SHARPNACK, J., RINALDO, A. and TIBSHIRANI, R. (2017). A sharp error analysis for the fused lasso, with application to approximate changepoint screening. In *Advances in Neural Information Processing Systems* (I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan and R. Garnett, eds.), vol. 30. Curran Associates, Inc.
- LOADER, C. (1991). Inference for a hazard rate change point. *Biometrika*, **78** 749–757.
- MARTINUSSEN, T. and SCHEIKE, T. (2006). *Dynamic Regression Models for Survival Data*. Springer.
- MATTHEWS, D. and FAREWELL, V. (1982). On testing for a constant hazard against a change-point alternative. *Biometrics*, **38** 463–468.
- MATTHEWS, D., FAREWELL, V. and PYKE, R. (1985). Asymptotic score-statistic processes and tests for constant hazard against a change-point alternative. *Annals of Statistics*, **13** 583–591.
- MELLER, M., BEYERSMANN, J. and RUFIBACH, K. (2019). Joint modeling of progression-free and overall survival and computation of correlation measures. *Statistics in Medicine*, **38** 4270–4289.
- MUKHOPADHYAY, P., YE, J., ANDERSON, K. M., ROYCHOUDHURY, S., RUBIN, E. H., HALABI, S. and CHAPPELL, R. J. (2022). Log-rank test vs MaxCombo and difference in restricted mean survival time tests for comparing survival under nonproportional hazards in immuno-oncology trials: a systematic review and meta-analysis. *JAMA Oncology*, **8** 1294–1300.
- NGUYEN, H., ROGERS, G. and WALKER, E. (1984). Estimation in change-point hazard rate models. *Biometrika*, **71** 299–304.
- OBA, K. and KUCHIBA, A. (2019). *Textbook of Clinical Trials in Oncology* (Ed. Halabi, S and Michiels, S), chap. Sample Size Calculations for Phase III Trials in Oncology. Boca Raton, FL: Chapman & Hall/ CRC.
- OHNEBERG, K. and SCHUMACHER, M. (2014). *Handbook of Survival Analysis* (Ed. Klein, J et al.), chap. Sample Size Calculations for Clinical Trials. Boca Raton, FL: Chapman & Hall/ CRC.
- OPSOMER, J., WANG, Y. and YANG, Y. (2001). Nonparametric regression with correlated errors. *Statistical Science*, **16** 134–153.
- RABINOWICZ, A. and ROSSET, S. (2022). Cross-validation for correlated data. *Journal of the American Statistical Association*, **117** 718–731.
- RITTMAYER, A., BARLESI, F., WATERKAMP, D., PARK, K., CIARDIELLO, F., VON PAWEL, J., GADGEEL, S. M., HIDA, T., KOWALSKI, D. M., DOLS, M. C. ET AL. (2017). Atezolizumab versus docetaxel in patients with previously treated non-small-cell lung cancer (OAK): a phase 3, open-label, multicentre randomised controlled trial. *The Lancet*, **389** 255–265.
- TIBSHIRANI, R. (2014). Adaptive piecewise polynomial estimation via trend filtering. *Annals of Statistics*, **42** 285–323.

- TIBSHIRANI, R., SAUNDERS, M., ROSSET, S., ZHU, J. and KNIGHT, K. (2005). Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B*, **67** 91–108.
- TITMAN, A. C. and SHARPLES, L. D. (2010). Model diagnostics for multi-state models. *Statistical Methods in Medical Research*, **19** 621–651.
- WOOD, A. T. A. (1999). Rosenthal’s inequality for point process martingales. *Stochastic Processes and their Applications*, **81** 231–246.
- WORSLEY, K. (1988). Exact percentage points of the likelihood-ratio test for a change-point hazard-rate model. *Biometrics*, **44** 259–263.
- YAO, Y.-C. (1986). Maximum likelihood estimation in hazard rate models with a change-point. *Communications in Statistics - Theory and Methods*, **15** 2455–2466.
- ZHANG, T. (2019). Element-wise estimation error of a total variation regularized estimator for change point detection. *arXiv:1901.00914*.
- ZHANG, T. and CHATTERJEE, S. (2023). Element-wise estimation error of generalized fused lasso. *Bernoulli*, **29** 2691–2718.
- ZHAO, X., WU, X. and ZHOU, X. (2009). A change-point model for survival data with long-term survivors. *Statistica Sinica*, **19** 377–390.

Technical Appendices

A Proof of main results

In what follows, we prove the theoretical results from Section 4. Throughout this and the following appendix, the symbols c and C denote generic constants that may take a different value on each occurrence. Moreover, the symbols c_j and C_j with subscript j (which may be either a natural number or a letter) are specific constants that are defined in the course of the appendices. Unless stated differently, the constants c , C , c_j and C_j do not depend on the sample size n .

Proof of Theorem 4.1

In order to prove Theorem 4.1, we make use of the following general result on the fused lasso due to Zhang and Chatterjee (2023); see also Theorem 3.1 in Zhang (2019).

Proposition A.1. *For any $\kappa_n > 0$ with the property that*

$$\max_{1 \leq k \leq \ell \leq n} \left| \frac{1}{\sqrt{\ell - k + 1}} \sum_{j=k}^{\ell} u_j \right| \leq \kappa_n,$$

it holds that

$$|\hat{\alpha}_{\lambda,j} - \alpha_j^*| \leq \max \left\{ \frac{\kappa_n}{\sqrt{d_j}}, \frac{\kappa_n^2}{4n\lambda}, \frac{2n\lambda}{r_{k(j)}} + \frac{2\kappa_n}{\sqrt{r_{k(j)}}} \right\}$$

for $j \in \{1, \dots, n\}$.

To make the paper as self-contained as possible, we provide a proof of Proposition A.1 in Appendix B. In order to apply Proposition A.1, we derive a (probabilistic) upper bound on the maximal partial sum $\max_{1 \leq k \leq \ell \leq n} |(\ell - k + 1)^{-1/2} \sum_{j=k}^{\ell} u_j|$ in the following proposition.

Proposition A.2. *Assume that (C1)–(C7) are satisfied and let*

$$\kappa_n = c_n \max \{n^{2/\nu}, \sqrt{n}\rho_n\},$$

where $\nu > 4$ is specified in (C6) and $\{c_n\}$ is a slowly diverging sequence (e.g. $c_n = c_0 \log \log n$ with some constant c_0). Then

$$\mathbb{P} \left(\max_{1 \leq k \leq \ell \leq n} \left| \frac{1}{\sqrt{\ell - k + 1}} \sum_{j=k}^{\ell} u_j \right| > \kappa_n \right) = o(1).$$

Proof. Since $u_j = \Delta_j^\alpha + \Delta_j^J + \Delta_j^\beta + \eta_j$, it holds that

$$\mathbb{P} \left(\max_{1 \leq k \leq \ell \leq n} \left| \frac{1}{\sqrt{\ell - k + 1}} \sum_{j=k}^{\ell} u_j \right| > \kappa_n \right) \leq P_n^\alpha + P_n^J + P_n^\beta + P_n^\eta,$$

where

$$\begin{aligned} P_n^\alpha &= \mathbb{P} \left(\max_{1 \leq k \leq \ell \leq n} \left| \frac{1}{\sqrt{\ell - k + 1}} \sum_{j=k}^{\ell} \Delta_j^\alpha \right| > \frac{\kappa_n}{4} \right) \\ P_n^J &= \mathbb{P} \left(\max_{1 \leq k \leq \ell \leq n} \left| \frac{1}{\sqrt{\ell - k + 1}} \sum_{j=k}^{\ell} \Delta_j^J \right| > \frac{\kappa_n}{4} \right) \\ P_n^\beta &= \mathbb{P} \left(\max_{1 \leq k \leq \ell \leq n} \left| \frac{1}{\sqrt{\ell - k + 1}} \sum_{j=k}^{\ell} \Delta_j^\beta \right| > \frac{\kappa_n}{4} \right) \\ P_n^\eta &= \mathbb{P} \left(\max_{1 \leq k \leq \ell \leq n} \left| \frac{1}{\sqrt{\ell - k + 1}} \sum_{j=k}^{\ell} \eta_j \right| > \frac{\kappa_n}{4} \right). \end{aligned}$$

We analyze these four probabilities one after the other:

- (i) As the change points $\tau_1 < \dots < \tau_K$ and their number K are fixed, we have the following for sufficiently large n , where we set $\tau_0 := \tau_{\min}$ and $\tau_{K+1} := \tau_{\max}$ for notational convenience: for each $j \in \{1, \dots, n\}$, either (a) $[t_{j-1}, t_j] \subseteq [\tau_{k-1}, \tau_k]$ for some k or (b) $t_{j-1} < \tau_k < t_j$ for some k . Case (a) occurs at least $n - K$ times, whereas case (b) occurs at most K times. It is straightforward to see that in case (a), $\Delta_j^\alpha = 0$, whereas in case (b),

$$\begin{aligned} |\Delta_j^\alpha| &= \left| \frac{\tau_k - t_{j-1}}{t_j - t_{j-1}} \mathbf{a}_k + \frac{t_j - \tau_k}{t_j - t_{j-1}} \mathbf{a}_{k+1} - \mathbf{a}_{k+1} \right| \\ &\leq \frac{\tau_k - t_{j-1}}{t_j - t_{j-1}} |\mathbf{a}_k - \mathbf{a}_{k+1}| \leq \max_{1 \leq k \leq K} |\mathbf{a}_k - \mathbf{a}_{k+1}| =: C_a < \infty. \end{aligned}$$

Since case (b) occurs at most K times, we obtain that

$$\max_{1 \leq k \leq \ell \leq n} \left| \frac{1}{\sqrt{\ell - k + 1}} \sum_{j=k}^{\ell} \Delta_j^\alpha \right| \leq KC_a.$$

With our choice of κ_n , this implies that $P_n^\alpha = 0$ for sufficiently large n .

- (ii) Since $\Delta_j^J = \{\Delta^J(t_j) - \Delta^J(t_{j-1})\} / \{t_j - t_{j-1}\}$ and $\Delta^J(t) = \int_0^t (J(s, \boldsymbol{\beta}) - 1) \alpha^*(s) ds$, we have

$$P_n^J \leq \mathbb{P} \left(\max_{1 \leq k \leq \ell \leq n} \sqrt{\ell - k + 1} \max_{k \leq j \leq \ell} |\Delta_j^J| > \frac{\kappa_n}{4} \right)$$

$$\begin{aligned}
&\leq \mathbb{P}\left(\max_{1 \leq j \leq n} |\Delta_j^J| > \frac{\kappa_n}{4\sqrt{n}}\right) \\
&\leq \mathbb{P}\left(\max_{1 \leq j \leq n} |\Delta_j^J| > 0\right) \\
&\leq \mathbb{P}\left(\inf_{t \in [\tau_{\min}, \tau_{\max}]} J(t, \boldsymbol{\beta}) = 0\right) = o(1)
\end{aligned}$$

by (C5).

(iii) By (C7), $\max_{1 \leq j \leq n} |\Delta_j^\beta| = O_p(\rho_n)$. Therefore,

$$\begin{aligned}
P_n^\beta &\leq \mathbb{P}\left(\max_{1 \leq k \leq \ell \leq n} \sqrt{\ell - k + 1} \max_{k \leq j \leq \ell} |\Delta_j^\beta| > \frac{\kappa_n}{4}\right) \\
&\leq \mathbb{P}\left(\max_{1 \leq j \leq n} |\Delta_j^\beta| > \frac{\kappa_n}{4\sqrt{n}}\right) = o(1),
\end{aligned}$$

as long as κ_n is such that $\kappa_n/\sqrt{n}\rho_n \rightarrow \infty$, which is ensured by our choice of κ_n .

(iv) A simple union bound yields that

$$P_n^\eta \leq \sum_{1 \leq k \leq \ell \leq n} P_{k,\ell}^\eta \quad \text{with} \quad P_{k,\ell}^\eta = \mathbb{P}\left(\left|\frac{1}{\sqrt{\ell - k + 1}} \sum_{j=k}^{\ell} \eta_j\right| > \frac{\kappa_n}{4}\right).$$

Since $\eta_j = \{\eta(t_j) - \eta(t_{j-1})\}/\{t_j - t_{j-1}\} = n\{\eta(t_j) - \eta(t_{j-1})\}$, it holds that $\sum_{j=k}^{\ell} \eta_j = n(\eta(t_\ell) - \eta(t_{k-1}))$. Hence, we can apply Markov's inequality to get that

$$P_{k,\ell}^\eta \leq \frac{n^\nu \mathbb{E}[|\eta(t_\ell) - \eta(t_{k-1})|^\nu]}{\{(\kappa_n/4)\sqrt{\ell - k + 1}\}^\nu}$$

for any $\nu > 0$. Our goal now is to derive a suitable upper bound on $\mathbb{E}|\eta(t_\ell) - \eta(t_{k-1})|^\nu$. To do so, we regard $\eta(t_\ell) - \eta(t_{k-1})$ as a stochastic integral of the form

$$\eta(t_\ell) - \eta(t_{k-1}) = \int_0^\tau H_{k,\ell}(s) d\bar{M}(s)$$

with

$$H_{k,\ell}(s) = \mathbf{1}(s \in (t_{k-1}, t_\ell]) \frac{J(s, \boldsymbol{\beta})}{\bar{Z}(s, \boldsymbol{\beta})}$$

and apply a version of Rosenthal's inequality to it. In particular, we use the following result due to Wood (1999). To formulate the result, we denote the intensity of the counting process \bar{N} by $\bar{\lambda}$, that is, $\bar{N}(t) = \int_0^t \bar{\lambda}(s) ds + \bar{M}(t)$ with $\bar{\lambda}(s) = \bar{Z}(s, \boldsymbol{\beta})\alpha^*(s)$.

Proposition A.3. *Consider the stochastic integral*

$$I_H(t) = \int_0^t H(s) d\bar{M}(s),$$

where $H = \{H(t) : t \in [0, \tau]\}$ is an $\{\mathcal{F}_t\}$ -predictable process. If (a) $\int_0^t \bar{\lambda}(s) ds < \infty$ almost surely for any time point $t \in [0, \tau]$, (b) $\mathbb{E}[\int_0^\tau |H(s)| \bar{\lambda}(s) ds] < \infty$ and (c) $\sup_{t \in [0, \tau]} \mathbb{E}[I_H^2(t)] < \infty$, then

$$\mathbb{E}[|I_H(\tau)|^\nu] \leq C_\nu \mathbb{E} \left[\{ \langle I_H \rangle(\tau) \}^{\nu/2} + \int_0^\tau |H(s)|^\nu \bar{\lambda}(s) ds \right]$$

for any $\nu > 2$, where C_ν is a finite constant depending only on ν .

In order to apply this lemma with $H = H_{k,\ell}$, we verify its conditions: As $\bar{Z}(\cdot, \beta)$ is predictable by (C3), so is the process $H_{k,\ell}$. Condition (a) is satisfied automatically in our framework (see e.g. Theorem 2.3.1 in Fleming and Harrington, 2005). Moreover, (b) holds since

$$\begin{aligned} \mathbb{E} \int_0^\tau |H_{k,\ell}(s)| \bar{\lambda}(s) ds &\leq \mathbb{E} \int_0^\tau \mathbf{1}(s \in (t_{k-1}, t_\ell]) J(s, \beta) \alpha^*(s) ds \\ &\leq \int_0^\tau \alpha^*(s) ds = A^*(\tau) \end{aligned}$$

and (c) follows from the bound

$$\begin{aligned} \mathbb{E} \left\{ \int_0^t H_{k,\ell}(s) d\bar{M}(s) \right\}^2 &= \mathbb{E} \int_0^t \left(\mathbf{1}(s \in (t_{k-1}, t_\ell]) \frac{J(s, \beta)}{\bar{Z}(s, \beta)} \right)^2 d\langle \bar{M} \rangle(s) \\ &= \mathbb{E} \int_0^t \left(\mathbf{1}(s \in (t_{k-1}, t_\ell]) \frac{J(s, \beta)}{\bar{Z}(s, \beta)} \right)^2 \bar{\lambda}(s) ds \\ &= \mathbb{E} \int_0^t |H_{k,\ell}(s)| \alpha^*(s) ds \\ &\leq \mathbb{E} \int_{\tau_{\min}}^{\tau_{\max}} \left| \frac{J(s, \beta)}{\bar{Z}(s, \beta)} \right| \alpha^*(s) ds \\ &\leq \left\{ \int_{\tau_{\min}}^{\tau_{\max}} \alpha^*(s) ds \right\} \mathbb{E} \sup_{s \in [\tau_{\min}, \tau_{\max}]} \left| \frac{J(s, \beta)}{\bar{Z}(s, \beta)} \right| < \infty, \end{aligned}$$

where $\mathbb{E} \sup_{s \in [\tau_{\min}, \tau_{\max}]} |J(s, \beta) / \bar{Z}(s, \beta)| < \infty$ by (C6). We now compute the upper bound of the Rosenthal inequality from Proposition A.3 in our case. Let ν be a natural number strictly larger than 4. Since

$$\mathbb{E} \left\| \frac{J(\cdot, \beta)}{\bar{Z}(\cdot, \beta)} \right\|_\infty^\nu \leq \frac{C_{\infty, \nu}}{n^\nu}$$

by (C6), it holds that

$$\begin{aligned}
\mathbb{E}\{\langle I_{H_{k,\ell}} \rangle(\tau)\}^{\nu/2} &= \mathbb{E}\left\{\int_0^\tau H_{k,\ell}^2(s)d\langle \bar{M} \rangle(s)\right\}^{\nu/2} \\
&= \mathbb{E}\left\{\int_0^\tau \left(\mathbf{1}_{(t_{k-1}, t_\ell]}(s) \frac{J(s, \boldsymbol{\beta})}{\bar{Z}(s, \boldsymbol{\beta})}\right)^2 \bar{\lambda}(s) ds\right\}^{\nu/2} \\
&= \mathbb{E}\left\{\int_0^\tau \mathbf{1}_{(t_{k-1}, t_\ell]}(s) \left|\frac{J(s, \boldsymbol{\beta})}{\bar{Z}(s, \boldsymbol{\beta})}\right| \alpha^*(s) ds\right\}^{\nu/2} \\
&\leq \left\{\int_0^\tau \mathbf{1}_{(t_{k-1}, t_\ell]}(s) \alpha^*(s) ds\right\}^{\nu/2} \mathbb{E}\left\|\frac{J(\cdot, \boldsymbol{\beta})}{\bar{Z}(\cdot, \boldsymbol{\beta})}\right\|_\infty^{\nu/2} \\
&\leq \left\{\frac{(\ell - k + 1)}{n} \|\alpha^*\|_\infty\right\}^{\nu/2} \mathbb{E}\left\|\frac{J(\cdot, \boldsymbol{\beta})}{\bar{Z}(\cdot, \boldsymbol{\beta})}\right\|_\infty^{\nu/2} \\
&\leq C_{\infty, \nu/2} \left\{\frac{(\ell - k + 1)}{n^2} \|\alpha^*\|_\infty\right\}^{\nu/2}
\end{aligned}$$

and

$$\begin{aligned}
\mathbb{E} \int_0^\tau |H_{k,\ell}(s)|^\nu \bar{\lambda}(s) ds &\leq \mathbb{E} \int_0^\tau \mathbf{1}_{(t_{k-1}, t_\ell]}(s) \left|\frac{J(s, \boldsymbol{\beta})}{\bar{Z}(s, \boldsymbol{\beta})}\right|^{\nu-1} \alpha^*(s) ds \\
&\leq \left\{\int_0^\tau \mathbf{1}_{(t_{k-1}, t_\ell]}(s) \alpha^*(s) ds\right\} \mathbb{E}\left\|\frac{J(\cdot, \boldsymbol{\beta})}{\bar{Z}(\cdot, \boldsymbol{\beta})}\right\|_\infty^{\nu-1} \\
&\leq \left\{\frac{(\ell - k + 1)}{n} \|\alpha^*\|_\infty\right\} \mathbb{E}\left\|\frac{J(\cdot, \boldsymbol{\beta})}{\bar{Z}(\cdot, \boldsymbol{\beta})}\right\|_\infty^{\nu-1} \\
&\leq C_{\infty, \nu-1} \left\{\frac{(\ell - k + 1)}{n^\nu} \|\alpha^*\|_\infty\right\}
\end{aligned}$$

with $\|\alpha^*\|_\infty = \sup_{t \in [0,1]} |\alpha^*(t)|$. We thus arrive at the bound

$$\begin{aligned}
\mathbb{E}|\eta(t_\ell) - \eta(t_{k-1})|^\nu &= \mathbb{E}\left|\int_0^\tau H_{k,\ell}(s)d\bar{M}(s)\right|^\nu \\
&\leq C_\nu \max\{C_{\infty, \nu/2}, C_{\infty, \nu-1}\} \\
&\quad \times \left[\left\{\frac{(\ell - k + 1)}{n^2} \|\alpha^*\|_\infty\right\}^{\nu/2} + \frac{(\ell - k + 1)}{n^\nu} \|\alpha^*\|_\infty\right].
\end{aligned}$$

Using this bound, we finally obtain that

$$P_n^\eta \leq \sum_{1 \leq k \leq \ell \leq n} \frac{n^\nu \mathbb{E}\left[|\eta(t_\ell) - \eta(t_{k-1})|^\nu\right]}{\{(\kappa_n/4)\sqrt{\ell - k + 1}\}^\nu} \leq \frac{Cn^2}{\kappa_n^\nu}$$

with the constant $C = 2^{2\nu+1} C_\nu \max\{C_{\infty, \nu/2}, C_{\infty, \nu-1}\} \max\{\|\alpha^*\|_\infty, \|\alpha^*\|_\infty^{\nu/2}\}$. With our choice of κ_n , this implies that $P_n^\eta = o(1)$. \square

Proof of Theorem 4.2

By Theorem 4.1, the following holds with probability tending to 1:

$$\begin{aligned} \sum_{j=n_k}^{n_{k+1}-1} (\hat{\alpha}_{\lambda,j} - \alpha_j^*)^2 &\leq \sum_{j=1}^{r_k} \left\{ \left[\frac{\kappa_n}{\sqrt{j}} \right]^2 + \left[\frac{\kappa_n}{\sqrt{r_k + 1 - j}} \right]^2 \right. \\ &\quad \left. + \frac{\kappa_n^4}{16(n\lambda)^2} + \frac{16(n\lambda)^2}{r_k^2} + \frac{16\kappa_n^2}{r_k} \right\} \\ &\leq C \left\{ \log(n)\kappa_n^2 + \frac{\kappa_n^4}{n\lambda^2} + n\lambda^2 + \kappa_n^2 \right\} \end{aligned}$$

for all $k \in \{0, \dots, K\}$ with some sufficiently large constant C , where we have used that $\sum_{j=1}^m j^{-1} \leq \log(m) + 1$ for any $m \in \mathbb{N}$ and $0 < c_r \leq r_k/n \leq C_r$ for all k and sufficiently large n with suitable constants $0 < c_r \leq C_r < \infty$. We thus obtain that

$$\begin{aligned} \frac{1}{n} \|\hat{\alpha}_\lambda - \alpha^*\|_2^2 &= \frac{1}{n} \sum_{j=1}^n (\hat{\alpha}_{\lambda,j} - \alpha_j^*)^2 \\ &= \frac{1}{n} \sum_{k=0}^K \sum_{j=n_k}^{n_{k+1}-1} (\hat{\alpha}_{\lambda,j} - \alpha_j^*)^2 \\ &\leq \frac{C}{n} \left\{ \log(n)\kappa_n^2 + \frac{\kappa_n^4}{n\lambda^2} + n\lambda^2 + \kappa_n^2 \right\} \end{aligned}$$

with probability tending to 1. Choosing λ such that $n\lambda^2 = \kappa_n^2$, that is, $\lambda = \kappa_n/\sqrt{n}$, we arrive at

$$\frac{1}{n} \|\hat{\alpha}_\lambda - \alpha^*\|_2^2 \leq \frac{C \log(n)\kappa_n^2}{n} = C \log(n)c_n^2 \max\{n^{-1+\frac{4}{\nu}}, \rho_n^2\}$$

with probability tending to 1. This completes the proof.

Proof of Theorem 4.3

Let $J^* = \{j \in \{2, \dots, n\} : \alpha_{j-1}^* \neq \alpha_j^*\}$ and note that $J^* = \{n_1, \dots, n_K\}$ with $n_k = \lceil (\tau_k - \tau_{\min})n \rceil$ for $k = 1, \dots, K$. Setting $n_0 = 1$ and $n_{K+1} = n + 1$ for convenience, we denote the smallest distance between two jump indices n_k and n_{k+1} of α^* by

$$D_n = \min_{0 \leq k \leq K} (n_{k+1} - n_k)$$

and the smallest jump height by

$$H_n = \min_{j \in J^*} |\alpha_j^* - \alpha_{j-1}^*|.$$

Theorem 4.3 is a direct consequence of Theorem 4 from Lin et al. (2017), which says the following:

Proposition A.4. *Let $\tilde{\boldsymbol{\alpha}} = (\tilde{\alpha}_1, \dots, \tilde{\alpha}_n)^\top$ be an estimator of $\boldsymbol{\alpha}^* = (\alpha_1^*, \dots, \alpha_n^*)^\top$ with the property that $n^{-1} \|\tilde{\boldsymbol{\alpha}} - \boldsymbol{\alpha}^*\|_2^2 = O_p(R_n)$ and let $\tilde{J} = \{j \in \{2, \dots, n\} : \tilde{\alpha}_{j-1} \neq \tilde{\alpha}_j\}$. Assume that $nR_n/H_n^2 = o(D_n)$. Then*

$$d(\tilde{J} | J^*) = O_p\left(\frac{nR_n}{H_n^2}\right).$$

A proof of Proposition A.4 is provided in Appendix B for completeness. We now verify that the assumptions of Proposition A.4 are fulfilled in our case. As the number of jumps K is fixed and finite, it holds that $D_n = \min_{0 \leq k \leq K} (\tau_{k+1} - \tau_k)n \geq c_D n$ with some constant $c_D > 0$. Similarly, $H_n \geq c_H$ with some $c_H > 0$. By Theorem 4.2,

$$\frac{1}{n} \|\hat{\boldsymbol{\alpha}}_\lambda - \boldsymbol{\alpha}^*\|_2^2 = O_p(R_n) \quad \text{with} \quad R_n = c_n^2 \log(n) \max\{n^{-1+\frac{4}{\nu}}, \rho_n^2\}.$$

Consequently,

$$\frac{nR_n}{H_n^2} \leq \frac{nR_n}{c_H^2} = \frac{1}{c_H^2} c_n^2 \log(n) \max\{n^{\frac{4}{\nu}}, n\rho_n^2\} = o(D_n),$$

which shows that the conditions of Proposition A.4 are fulfilled. Defining $\hat{J}_\lambda = \{j \in \{2, \dots, n\} : \hat{\alpha}_{\lambda, j-1} \neq \hat{\alpha}_{\lambda, j}\}$, we can now apply Proposition A.4 to obtain that

$$d(\hat{J}_\lambda | J^*) = O_p\left(c_n^2 \log(n) \max\{n^{\frac{4}{\nu}}, n\rho_n^2\}\right),$$

which is equivalent to

$$d(\hat{\mathcal{S}}_\lambda | \mathcal{S}^*) = O_p\left(c_n^2 \log(n) \max\{n^{-1+\frac{4}{\nu}}, \rho_n^2\}\right).$$

Proof of Theorem 4.4

Theorem 4.1 immediately yields that with probability tending to 1,

$$\begin{aligned} |\hat{\alpha}_{\lambda, \hat{\tau}_n} - \hat{\alpha}_{\lambda, \hat{\tau}_{n-1}}| &\leq |\hat{\alpha}_{\lambda, \hat{\tau}_n} - \alpha_{\hat{\tau}_n}^*| + |\hat{\alpha}_{\lambda, \hat{\tau}_{n-1}} - \alpha_{\hat{\tau}_{n-1}}^*| \\ &\leq 2 \max_{j \in \{\hat{\tau}_{n-1}, \hat{\tau}_n\}} \left\{ \frac{\kappa_n}{\sqrt{d_j}}, \frac{\kappa_n^2}{4n\lambda}, \frac{2n\lambda}{r_{k(j)}} + \frac{2\kappa_n}{\sqrt{r_{k(j)}}} \right\} \\ &\leq C \max\left\{ \frac{\kappa_n}{\sqrt{\Delta_n}}, \frac{\kappa_n}{\sqrt{n}} \right\} = C \frac{\kappa_n}{\sqrt{\Delta_n}} \end{aligned}$$

for all $\hat{\tau} \in \hat{\mathcal{S}}_\lambda^{\text{far}}$, where we have used that $\min_{1 \leq k \leq K} |\hat{\tau} - \tau_k| \geq \Delta_n/n$, $0 < c_r \leq r_k/n \leq C_r$ for all k and sufficiently large n with suitable constants $0 < c_r \leq C_r < \infty$, $\lambda = \kappa_n/\sqrt{n}$,

and Δ_n can always be chosen such that $\Delta_n \leq n$.

Proof of Proposition 4.5

Proposition 4.5 can be proven by a simplified version of the technical arguments for Proposition 4.6. Hence, we only prove the latter.

Proof of Proposition 4.6

In what follows, all probabilistic statements are to be understood with respect to the conditional probability measure \mathbb{P}^B given study entry. Formally, this measure is defined by $\mathbb{P}^B(A) = \mathbb{P}(A \cap B)/\mathbb{P}(B)$ for all $A \in \mathcal{F}$, where $B = \bigcap_{i=1}^n B_i$ with $B_i = \{T_i^* > L_i\}$. In addition to \mathbb{P}^B , we use the symbol \mathbb{E}^B to denote the expectation given study entry. As each competing risk $r \in \{1, \dots, R\}$ is analyzed separately, we let r be fixed throughout the proof and consider the corresponding multivariate counting process $N_{1:n,r} = (N_{1r}, \dots, N_{nr})$ on the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P}^B)$. Here, \mathcal{F}_t is the σ -algebra generated by the information available up to time t . Formally speaking, $\mathcal{F}_t = \sigma(N_{ir'}(s), \mathbf{1}(L_i < s \leq T_i) : 1 \leq i \leq n, 1 \leq r' \leq R, 0 \leq s \leq t) \vee \sigma(\mathbf{W}_1, \dots, \mathbf{W}_n)$.

Before we verify conditions (C1)–(C7), we prove an auxiliary lemma.

Lemma A.5. *Let $Z \sim \text{Bin}(n, q)$ be a binomially distributed random variable with $q \in (0, 1)$, that is,*

$$\mathbb{P}^B(Z = k) = \binom{n}{k} q^k (1 - q)^{n-k} \quad \text{for } k = 0, \dots, n.$$

Then for any $\nu \in \mathbb{N}$,

$$\mathbb{E}^B \left[\frac{\mathbf{1}(Z > 0)}{Z^\nu} \right] \leq \frac{(\nu + 1)!}{(nq)^\nu},$$

where we define

$$\frac{\mathbf{1}(Z > 0)}{Z^\nu} = \begin{cases} 1/Z^\nu & \text{for } Z > 0 \\ 0 & \text{for } Z = 0. \end{cases}$$

Proof. We have

$$\begin{aligned} \mathbb{E}^B \left[\frac{\mathbf{1}(Z > 0)}{Z^\nu} \right] &= \sum_{k=1}^n \frac{1}{k^\nu} \binom{n}{k} q^k (1 - q)^{n-k} \\ &= \sum_{k=1}^n \frac{1}{k^\nu} \frac{n!}{(n-k)!k!} q^k (1 - q)^{n-k} \\ &= \sum_{k=1}^n \frac{1}{k^\nu} \left(\prod_{i=1}^{\nu} \frac{(k+i)}{(k+i)} \right) \frac{n!}{(n-k)!k!} q^k (1 - q)^{n-k} \end{aligned}$$

$$\begin{aligned}
&= \sum_{k=1}^n \underbrace{\prod_{i=1}^{\nu} \frac{k+i}{k}}_{\leq \prod_{i=1}^{\nu} (1+i) = (\nu+1)!} \frac{n!}{(n-k)!(k+\nu)!} q^k (1-q)^{n-k} \\
&\leq (\nu+1)! \sum_{k=1}^n \frac{n!}{(n-k)!(k+\nu)!} q^k (1-q)^{n-k} \\
&= (\nu+1)! \sum_{k=1}^n \frac{n!}{[(n+\nu)-(k+\nu)]!(k+\nu)!} q^k (1-q)^{n-k} \\
&= (\nu+1)! \sum_{k=1}^n \left(\prod_{i=1}^{\nu} \frac{1}{n+i} \right) \frac{(n+\nu)!}{[(n+\nu)-(k+\nu)]!(k+\nu)!} \\
&\quad \times q^{k+\nu} q^{-\nu} (1-q)^{(n+\nu)-(k+\nu)} \\
&= (\nu+1)! q^{-\nu} \left(\prod_{i=1}^{\nu} \frac{1}{n+i} \right) \underbrace{\sum_{k=1}^n \binom{n+\nu}{k+\nu} q^{k+\nu} (1-q)^{(n+\nu)-(k+\nu)}}_{= \sum_{\ell=\nu+1}^{n+\nu} \binom{n+\nu}{\ell} q^{\ell} (1-q)^{(n+\nu)-\ell}} \\
&\quad \leq \sum_{\ell=0}^{n+\nu} \binom{n+\nu}{\ell} q^{\ell} (1-q)^{(n+\nu)-\ell} = 1 \\
&\leq (\nu+1)! q^{-\nu} \left(\prod_{i=1}^{\nu} \frac{1}{n+i} \right) \\
&\leq (\nu+1)! (nq)^{-\nu}. \quad \square
\end{aligned}$$

We now verify conditions (C1)–(C7).

Proof of (C1). We need to show that the family of σ -algebras $\{\mathcal{F}_t\}$ is increasing, right-continuous and complete (i.e. can be completed). This follows from standard arguments (see e.g. Section II.2 in Andersen et al., 1993). \square

Proof of (C2). By (C_{C1}), the data vectors $(T_i, \delta_i \cdot \varepsilon_i, L_i, \mathbf{W}_i)$ are i.i.d. across i (with respect to \mathbb{P}^B). This implies that the counting processes N_{ir} defined by $N_{ir}(t) = \mathbf{1}(L_i < T_i \leq t, \delta_i \cdot \varepsilon_i = r)$ are i.i.d. across i as well (with respect to \mathbb{P}^B). \square

Proof of (C3). Condition (C3) is a direct consequence of (C_{C3}). In particular, since the at-risk process $\mathbf{1}(L_i < \cdot \leq T_i)$ is adapted to $\{\mathcal{F}_t\}$ and left-continuous, it is $\{\mathcal{F}_t\}$ -predictable, which immediately implies that $Z_i(\cdot, \boldsymbol{\beta}_r) = \mathbf{1}(L_i < \cdot \leq T_i) \exp(\boldsymbol{\beta}_r^\top \mathbf{W}_i)$ and λ_{ir} are $\{\mathcal{F}_t\}$ -predictable as well. \square

Proof of (C₄). We have $\bar{Z}(t, \boldsymbol{\beta}_r) = \sum_{i=1}^n Z_i(t, \boldsymbol{\beta}_r)$ with $Z_i(t, \boldsymbol{\beta}_r) = \mathbf{1}(L_i < t \leq T_i) \exp(\boldsymbol{\beta}_r^\top \mathbf{W}_i)$ and $J(t, \boldsymbol{\beta}_r) = \mathbf{1}(\bar{Z}(t, \boldsymbol{\beta}_r) > 0)$. Since $\|\mathbf{W}_i\|_\infty \leq C_W < \infty$ by (C_{C5}) and $\|\boldsymbol{\beta}_r\|_1 \leq C_\beta < \infty$ by (C_{C6}), it holds that $|\boldsymbol{\beta}_r^\top \mathbf{W}_i| \leq \|\boldsymbol{\beta}_r\|_1 \|\mathbf{W}_i\|_\infty \leq C_\beta C_W < \infty$. From this, it immediately follows that the process $\bar{Z}(\cdot, \boldsymbol{\beta}_r)$ is bounded. With the convention

$0/0 := 0$, the process $J(\cdot, \boldsymbol{\beta}_r)/\bar{Z}(\cdot, \boldsymbol{\beta}_r)$ is bounded as well. This in particular implies that $\bar{Z}(\cdot, \boldsymbol{\beta}_r)$ and $J(\cdot, \boldsymbol{\beta}_r)/\bar{Z}(\cdot, \boldsymbol{\beta}_r)$ are locally bounded. \square

Proof of (C5). Let $J(t) = \mathbf{1}(\bar{Z}(t) > 0)$ with $\bar{Z}(t) = \sum_{i=1}^n \mathbf{1}(L_i < t \leq T_i)$ and write $\mathcal{T} = [\tau_{\min}, \tau_{\max}]$ for short. Since $0 < \min_{1 \leq i \leq n} \exp(\boldsymbol{\beta}_r^\top \mathbf{W}_i) \leq \max_{1 \leq i \leq n} \exp(\boldsymbol{\beta}_r^\top \mathbf{W}_i) < \infty$ with probability 1, it holds that

$$\mathbb{P}^B(J(t, \boldsymbol{\beta}_r) = J(t) \text{ for all } t \in \mathcal{T}) = 1,$$

which implies that

$$\mathbb{P}^B\left(\inf_{t \in \mathcal{T}} J(t, \boldsymbol{\beta}_r) = 0\right) = \mathbb{P}^B\left(\inf_{t \in \mathcal{T}} J(t) = 0\right).$$

Moreover, since $\bar{Z}(t) = \sum_{i=1}^n \mathbf{1}(L_i < t \wedge T_i \geq t) \geq \sum_{i=1}^n \mathbf{1}(L_i < \tau_{\min} \wedge T_i \geq \tau_{\max})$ for all $t \in \mathcal{T}$,

$$\begin{aligned} \mathbb{P}^B\left(\inf_{t \in \mathcal{T}} J(t) = 0\right) &\leq \mathbb{P}^B\left(\sum_{i=1}^n \mathbf{1}(L_i < \tau_{\min} \wedge T_i \geq \tau_{\max}) = 0\right) \\ &= \prod_{i=1}^n \left\{1 - \mathbb{P}^B(L_i < \tau_{\min} \wedge T_i \geq \tau_{\max})\right\} \\ &= (1 - p)^n = o(1), \end{aligned}$$

where we have used (C_{C4}) in the last line. \square

Proof of (C6). As before, let $\bar{Z}(t) = \sum_{i=1}^n \mathbf{1}(L_i < t \leq T_i)$, $J(t) = \mathbf{1}(\bar{Z}(t) > 0)$ and $\mathcal{T} = [\tau_{\min}, \tau_{\max}]$. Since $\mathbb{P}^B(J(t, \boldsymbol{\beta}_r) = J(t) \text{ for all } t \in \mathcal{T}) = 1$ as already shown in the proof of (C5) and $|\boldsymbol{\beta}_r^\top \mathbf{W}_i| \leq \|\boldsymbol{\beta}_r\|_1 \|\mathbf{W}_i\|_\infty \leq C_\beta C_W$ by (C_{C5}) and (C_{C6}), it holds that

$$\begin{aligned} &\mathbb{E}^B \left[\left\| \frac{J(\cdot, \boldsymbol{\beta}_r)}{\bar{Z}(\cdot, \boldsymbol{\beta}_r)} \right\|_\infty^\nu \middle| \mathbf{W}_1, \dots, \mathbf{W}_n \right] \\ &\leq \frac{1}{\min_{1 \leq i \leq n} \exp(\nu \boldsymbol{\beta}_r^\top \mathbf{W}_i)} \mathbb{E}^B \left[\sup_{t \in \mathcal{T}} \left| \frac{J(t)}{\bar{Z}(t)} \right|^\nu \middle| \mathbf{W}_1, \dots, \mathbf{W}_n \right] \\ &\leq \max_{1 \leq i \leq n} \exp(\nu |\boldsymbol{\beta}_r^\top \mathbf{W}_i|) \mathbb{E}^B \left[\sup_{t \in \mathcal{T}} \left| \frac{J(t)}{\bar{Z}(t)} \right|^\nu \middle| \mathbf{W}_1, \dots, \mathbf{W}_n \right] \\ &\leq \exp(\nu C_\beta C_W) \mathbb{E}^B \left[\sup_{t \in \mathcal{T}} \left| \frac{J(t)}{\bar{Z}(t)} \right|^\nu \middle| \mathbf{W}_1, \dots, \mathbf{W}_n \right] \end{aligned}$$

almost surely and therefore

$$\mathbb{E}^B \left[\left\| \frac{J(\cdot, \boldsymbol{\beta}_r)}{\bar{Z}(\cdot, \boldsymbol{\beta}_r)} \right\|_\infty^\nu \right] \leq \exp(\nu C_\beta C_W) \mathbb{E}^B \left[\sup_{t \in \mathcal{T}} \left| \frac{J(t)}{\bar{Z}(t)} \right|^\nu \right].$$

Next define $\bar{Z}_\mathcal{T} := \sum_{i=1}^n \mathbf{1}(L_i < \tau_{\min} \wedge T_i \geq \tau_{\max})$ and $J_\mathcal{T} := \mathbf{1}(\bar{Z}_\mathcal{T} > 0)$. Since

$\bar{Z}(t) \geq \bar{Z}_\mathcal{T}$ and $J(t) \geq J_\mathcal{T}$ for any $t \in \mathcal{T}$ and $J(t)/\bar{Z}(t)$ is upper bounded by 1 (with the convention $0/0 := 0$),

$$\begin{aligned}
\mathbb{E}^B \left[\sup_{t \in \mathcal{T}} \left| \frac{J(t)}{\bar{Z}(t)} \right|^\nu \right] &= \mathbb{E}^B \left[\sup_{t \in \mathcal{T}} \left| \frac{J(t)}{\bar{Z}(t)} \right|^\nu \mathbf{1} \left(\inf_{t \in \mathcal{T}} J(t) > 0 \right) \right] \\
&\quad + \mathbb{E}^B \left[\sup_{t \in \mathcal{T}} \left| \frac{J(t)}{\bar{Z}(t)} \right|^\nu \mathbf{1} \left(\inf_{t \in \mathcal{T}} J(t) = 0 \right) \right] \\
&\leq \mathbb{E}^B \left[\sup_{t \in \mathcal{T}} \left| \frac{J(t)}{\bar{Z}(t)} \right|^\nu \mathbf{1} \left(\inf_{t \in \mathcal{T}} J(t) > 0 \right) \right] \\
&\quad + \mathbb{P}^B \left(\inf_{t \in \mathcal{T}} J(t) = 0 \right) \\
&\leq \mathbb{E}^B \left[\sup_{t \in \mathcal{T}} \left| \frac{J(t)}{\bar{Z}(t)} \right|^\nu \mathbf{1} \left(\inf_{t \in \mathcal{T}} J(t) > 0 \wedge \inf_{t \in \mathcal{T}} J(t) = J_\mathcal{T} \right) \right] \\
&\quad + \mathbb{E}^B \left[\sup_{t \in \mathcal{T}} \left| \frac{J(t)}{\bar{Z}(t)} \right|^\nu \mathbf{1} \left(\inf_{t \in \mathcal{T}} J(t) > 0 \wedge \inf_{t \in \mathcal{T}} J(t) > J_\mathcal{T} \right) \right] \\
&\quad + \mathbb{P}^B \left(\inf_{t \in \mathcal{T}} J(t) = 0 \right) \\
&\leq \mathbb{E}^B \left[\left| \frac{J_\mathcal{T}}{\bar{Z}_\mathcal{T}} \right|^\nu \right] + \mathbb{P}^B(J_\mathcal{T} = 0) + \mathbb{P}^B \left(\inf_{t \in \mathcal{T}} J(t) = 0 \right).
\end{aligned}$$

From the proof of (C5), we already know that $\mathbb{P}^B(J_\mathcal{T} = 0) \leq (1-p)^n$ as well as $\mathbb{P}^B(\inf_{t \in \mathcal{T}} J(t) = 0) \leq (1-p)^n$. Moreover, since $\bar{Z}_\mathcal{T} = \sum_{i=1}^n \mathbf{1}(L_i < \tau_{\min} \wedge T_i \geq \tau_{\max})$ follows a binomial distribution $\text{Bin}(n, q)$ with $q = \mathbb{P}^B(L_i < \tau_{\min} \wedge T_i \geq \tau_{\max}) = p$ under (C4), we can apply Lemma A.5 to get that

$$\mathbb{E}^B \left[\left| \frac{J_\mathcal{T}}{\bar{Z}_\mathcal{T}} \right|^\nu \right] \leq \frac{(\nu+1)!}{\{np\}^\nu}.$$

Putting everything together, we arrive at

$$\mathbb{E}^B \left[\left\| \frac{J(\cdot, \boldsymbol{\beta}_r)}{\bar{Z}(\cdot, \boldsymbol{\beta}_r)} \right\|_\infty^\nu \right] \leq \exp(\nu C_\beta C_W) \left\{ \frac{(\nu+1)!}{\{np\}^\nu} + 2(1-p)^n \right\}$$

for any $\nu \in \mathbb{N}$. As $(1-p)^n$ converges to zero faster than $n^{-\nu}$, (C6) is an immediate consequence of this bound. \square

Proof of (C7). We show that

$$\max_{1 \leq k \leq n} |\Delta_k^\beta| = O_p \left(\frac{\log(n)}{\sqrt{n}} \right)$$

(where the notation $O_p(\cdot)$ and $o_p(\cdot)$ is to be understood with respect to \mathbb{P}^B). Our proof exploits the following facts about the random variables $\bar{Z}(t, \mathbf{b}) = \sum_{i=1}^n \mathbf{1}(L_i < t \leq T_i) \exp(\mathbf{b}^\top \mathbf{W}_i)$ and $J(t, \mathbf{b}) = \mathbf{1}(\bar{Z}(t, \mathbf{b}) > 0)$ with $\mathbf{b} \in \mathbb{R}^d$:

(a) It holds that

$$\mathbb{P}^B(J(T_i, \hat{\boldsymbol{\beta}}_r) = J(T_i, \boldsymbol{\beta}_r) \text{ for all } i = 1, \dots, n) = 1.$$

Proof. Since $J(t, \mathbf{b}) = \mathbf{1}(\sum_{i=1}^n \mathbf{1}(L_i < t \leq T_i) \exp(\mathbf{b}^\top \mathbf{W}_i) > 0)$, (a) is a direct consequence of the fact that $0 < \min_{1 \leq i \leq n} \exp(\boldsymbol{\beta}^\top \mathbf{W}_i) \leq \max_{1 \leq i \leq n} \exp(\boldsymbol{\beta}^\top \mathbf{W}_i) < \infty$ and $0 < \min_{1 \leq i \leq n} \exp(\hat{\boldsymbol{\beta}}_r^\top \mathbf{W}_i) \leq \max_{1 \leq i \leq n} \exp(\hat{\boldsymbol{\beta}}_r^\top \mathbf{W}_i) < \infty$ with probability 1. \square

(b) It holds that

$$\sup_{t \in \mathcal{T}} |\bar{Z}(t, \hat{\boldsymbol{\beta}}_r) - \bar{Z}(t, \boldsymbol{\beta}_r)| = O_p(\log(n)\sqrt{n}).$$

Proof. With the help of a Taylor expansion, we get that

$$\begin{aligned} & \sup_{t \in \mathcal{T}} |\bar{Z}(t, \hat{\boldsymbol{\beta}}_r) - \bar{Z}(t, \boldsymbol{\beta}_r)| \\ & \leq \sup_{t \in \mathcal{T}} \sum_{i=1}^n \mathbf{1}(L_i < t \leq T_i) |\exp(\hat{\boldsymbol{\beta}}_r^\top \mathbf{W}_i) - \exp(\boldsymbol{\beta}_r^\top \mathbf{W}_i)| \\ & \leq \sum_{i=1}^n |\exp(\hat{\boldsymbol{\beta}}_r^\top \mathbf{W}_i) - \exp(\boldsymbol{\beta}_r^\top \mathbf{W}_i)| \\ & = \sum_{i=1}^n |\exp(\tilde{\boldsymbol{\beta}}_r^\top \mathbf{W}_i) \mathbf{W}_i^\top (\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_r)| \\ & \leq \sum_{i=1}^n \exp(\tilde{\boldsymbol{\beta}}_r^\top \mathbf{W}_i) \|\mathbf{W}_i\|_\infty \|\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_r\|_1 \\ & = \sum_{i=1}^n \exp(\tilde{\boldsymbol{\beta}}_r^\top \mathbf{W}_i) \|\mathbf{W}_i\|_\infty \|\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_r\|_1 \mathbf{1} \left(\|\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_r\|_1 \leq \frac{\log(n)}{\sqrt{n}} \right) \\ & \quad + \sum_{i=1}^n \exp(\tilde{\boldsymbol{\beta}}_r^\top \mathbf{W}_i) \|\mathbf{W}_i\|_\infty \|\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_r\|_1 \mathbf{1} \left(\|\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_r\|_1 > \frac{\log(n)}{\sqrt{n}} \right), \end{aligned} \quad (\text{A.1})$$

where $\tilde{\boldsymbol{\beta}}_r$ is an intermediate value between $\hat{\boldsymbol{\beta}}_r$ and $\boldsymbol{\beta}_r$. In the case that $\|\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_r\|_1 \leq \log(n)n^{-1/2}$, $\exp(\tilde{\boldsymbol{\beta}}_r^\top \mathbf{W}_i) \leq \exp(\|\tilde{\boldsymbol{\beta}}_r\|_1 \|\mathbf{W}_i\|_\infty) \leq \exp(2C_\beta C_W)$ for sufficiently large n because $\tilde{\boldsymbol{\beta}}_r$ lies between $\hat{\boldsymbol{\beta}}_r$ and $\boldsymbol{\beta}_r$ and thus $\|\tilde{\boldsymbol{\beta}}_r\|_1 \leq \max\{\|\boldsymbol{\beta}_r\|_1, \|\hat{\boldsymbol{\beta}}_r\|_1\} \leq \|\boldsymbol{\beta}_r\|_1 + \log(n)n^{-1/2} \leq 2C_\beta$ for sufficiently large n . Consequently,

$$\begin{aligned} & \sum_{i=1}^n \exp(\tilde{\boldsymbol{\beta}}_r^\top \mathbf{W}_i) \|\mathbf{W}_i\|_\infty \|\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_r\|_1 \mathbf{1} \left(\|\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_r\|_1 \leq \frac{\log(n)}{\sqrt{n}} \right) \\ & \leq \frac{\exp(2C_\beta C_W) \log(n)}{\sqrt{n}} \sum_{i=1}^n \|\mathbf{W}_i\|_\infty = O_p(\log(n)\sqrt{n}). \end{aligned} \quad (\text{A.2})$$

Moreover, for any given $\varepsilon > 0$,

$$\begin{aligned} \mathbb{P} \left(\sum_{i=1}^n \exp(\tilde{\boldsymbol{\beta}}_r^\top \mathbf{W}_i) \|\mathbf{W}_i\|_\infty \|\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_r\|_1 \mathbf{1} \left(\|\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_r\|_1 > \frac{\log(n)}{\sqrt{n}} \right) > \varepsilon \right) \\ \leq \mathbb{P} \left(\|\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_r\|_1 > \frac{\log(n)}{\sqrt{n}} \right) = o(1), \end{aligned}$$

where the last equality follows from (C_C6), and thus

$$\sum_{i=1}^n \exp(\tilde{\boldsymbol{\beta}}_r^\top \mathbf{W}_i) \|\mathbf{W}_i\|_\infty \|\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_r\|_1 \mathbf{1} \left(\|\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_r\|_1 > \frac{\log(n)}{\sqrt{n}} \right) = o_p(1). \quad (\text{A.3})$$

Plugging (A.2) and (A.3) into (A.1) yields (b). \square

- (c) It holds that $\inf_{t \in \mathcal{T}} \bar{Z}(t, \boldsymbol{\beta}_r) \geq cn$ with probability tending to 1, where c is any positive constant with $c < p \exp(-C_\beta C_W)$.

Proof. Under our assumptions,

$$\begin{aligned} \left| \frac{\inf_{t \in \mathcal{T}} \bar{Z}(t, \boldsymbol{\beta}_r)}{n} \right| &= \inf_{t \in \mathcal{T}} \left\{ \frac{1}{n} \sum_{i=1}^n \exp(\boldsymbol{\beta}_r^\top \mathbf{W}_i) \mathbf{1}(L_i < t \leq T_i) \right\} \\ &\geq \frac{1}{n} \sum_{i=1}^n \exp(\boldsymbol{\beta}_r^\top \mathbf{W}_i) \mathbf{1}(L_i < \tau_{\min} \wedge T_i \geq \tau_{\max}) \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\exp(\boldsymbol{\beta}_r^\top \mathbf{W}_i) \mathbf{1}(L_i < \tau_{\min} \wedge T_i \geq \tau_{\max}) \right] + O_p \left(\frac{1}{\sqrt{n}} \right) \\ &\geq \exp(-C_\beta C_W) \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\mathbf{1}(L_i < \tau_{\min} \wedge T_i \geq \tau_{\max}) \right] + O_p \left(\frac{1}{\sqrt{n}} \right) \\ &= p \exp(-C_\beta C_W) + O_p \left(\frac{1}{\sqrt{n}} \right), \end{aligned}$$

where the second equality follows from a simple application of the weak law of large numbers. This immediately implies (c). \square

- (d) It holds that $\inf_{t \in \mathcal{T}} \bar{Z}(t, \hat{\boldsymbol{\beta}}_r) \geq cn$ with probability tending to 1, where c is any positive constant with $c < p \exp(-C_\beta C_W)$.

Proof. As by (b), $\inf_{t \in \mathcal{T}} \bar{Z}(t, \hat{\boldsymbol{\beta}}_r) \geq \inf_{t \in \mathcal{T}} \bar{Z}(t, \boldsymbol{\beta}_r) - \sup_{t \in \mathcal{T}} |\bar{Z}(t, \hat{\boldsymbol{\beta}}_r) - \bar{Z}(t, \boldsymbol{\beta}_r)| = \inf_{t \in \mathcal{T}} \bar{Z}(t, \boldsymbol{\beta}_r) + O_p(\log(n)\sqrt{n})$, we can use (c) to complete the proof. \square

We now bound $\max_{1 \leq k \leq n} |\Delta_k^\beta|$ with the help of (a)–(d). Using (a), we obtain that

$$\begin{aligned} \max_{1 \leq k \leq n} |\Delta_k^\beta| &= n \max_{1 \leq k \leq n} \left| \int \mathbf{1}(s \in (t_{k-1}, t_k]) \frac{J(s, \hat{\boldsymbol{\beta}}_r)}{\bar{Z}(s, \hat{\boldsymbol{\beta}}_r)} d\bar{N}_r(s) \right. \\ &\quad \left. - \int \mathbf{1}(s \in (t_{k-1}, t_k]) \frac{J(s, \boldsymbol{\beta}_r)}{\bar{Z}(s, \boldsymbol{\beta}_r)} d\bar{N}_r(s) \right| \end{aligned}$$

$$\begin{aligned}
&= n \max_{1 \leq k \leq n} \left| \sum_{\{i: t_{k-1} < T_i \leq t_k, \delta_i \cdot \varepsilon_i = r\}} \left\{ \frac{J(T_i, \hat{\beta}_r)}{\bar{Z}(T_i, \hat{\beta}_r)} - \frac{J(T_i, \beta_r)}{\bar{Z}(T_i, \beta_r)} \right\} \right| \\
&= n \max_{1 \leq k \leq n} \left| \sum_{\{i: t_{k-1} < T_i \leq t_k, \delta_i \cdot \varepsilon_i = r\}} J(T_i, \beta_r) \left\{ \frac{1}{\bar{Z}(T_i, \hat{\beta}_r)} - \frac{1}{\bar{Z}(T_i, \beta_r)} \right\} \right| \\
&\leq n \max_{1 \leq k \leq n} \left\{ \sum_{\{i: t_{k-1} < T_i \leq t_k, \delta_i \cdot \varepsilon_i = r\}} \left| \frac{1}{\bar{Z}(T_i, \hat{\beta}_r)} - \frac{1}{\bar{Z}(T_i, \beta_r)} \right| \right\} \\
&= n \max_{1 \leq k \leq n} \left\{ \sum_{\{i: t_{k-1} < T_i \leq t_k, \delta_i \cdot \varepsilon_i = r\}} \left| \frac{\bar{Z}(T_i, \beta_r) - \bar{Z}(T_i, \hat{\beta}_r)}{\bar{Z}(T_i, \hat{\beta}_r) \bar{Z}(T_i, \beta_r)} \right| \right\} \\
&\leq \frac{n}{\{\inf_{t \in \mathcal{T}} \bar{Z}(t, \hat{\beta}_r)\} \{\inf_{t \in \mathcal{T}} \bar{Z}(t, \beta_r)\}} \\
&\quad \times \max_{1 \leq k \leq n} \left\{ \sum_{\{i: t_{k-1} < T_i \leq t_k, \delta_i \cdot \varepsilon_i = r\}} |\bar{Z}(T_i, \beta_r) - \bar{Z}(T_i, \hat{\beta}_r)| \right\} \tag{A.4}
\end{aligned}$$

with probability 1. From (c) and (d), it directly follows that

$$\frac{n}{\{\inf_{t \in \mathcal{T}} \bar{Z}(t, \hat{\beta}_r)\} \{\inf_{t \in \mathcal{T}} \bar{Z}(t, \beta_r)\}} = O_p\left(\frac{1}{n}\right). \tag{A.5}$$

Moreover, condition (C_C2) and standard arguments for uniform convergence yield that $\max_{1 \leq k \leq n} \{\sum_{i=1}^n \mathbf{1}(t_{k-1} < T_i^* \leq t_k)\} = O_p(1)$. Combining this with (b), we get that

$$\begin{aligned}
&\max_{1 \leq k \leq n} \left\{ \sum_{\{i: t_{k-1} < T_i \leq t_k, \delta_i \cdot \varepsilon_i = r\}} |\bar{Z}(T_i, \beta_r) - \bar{Z}(T_i, \hat{\beta}_r)| \right\} \\
&\leq O_p(\log(n)\sqrt{n}) \max_{1 \leq k \leq n} \left\{ \sum_{i=1}^n \mathbf{1}(t_{k-1} < T_i^* \leq t_k) \right\} = O_p(\log(n)\sqrt{n}). \tag{A.6}
\end{aligned}$$

Plugging (A.6) and (A.5) into (A.4), we finally arrive at the desired result

$$\max_{1 \leq k \leq n} |\Delta_k^\beta| = O_p\left(\frac{\log(n)}{\sqrt{n}}\right). \quad \square$$

Proof of Proposition 4.7

As each direct $\ell \rightarrow m$ transition with $\ell, m \in \{0, \dots, R\}$, $\ell \neq m$, is analyzed separately, we let the pair (ℓ, m) be fixed throughout the proof and consider the corresponding multivariate counting process $N_{1:n, \ell \rightarrow m} = (N_{1, \ell \rightarrow m}, \dots, N_{n, \ell \rightarrow m})$ on the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$. Here, \mathcal{F}_t is the σ -algebra generated by the information available up to time t , which is given by $\mathcal{F}_t = \sigma(X_1(s), \dots, X_n(s) : 0 \leq s \leq t)$.

Proof of (C1). Standard arguments show that $\{\mathcal{F}_t\}$ is increasing, right-continuous and complete (i.e. can be completed); see Section II.2 in Andersen et al. (1993). \square

Proof of (C2). By (C_{D1}), the Markov processes X_i are i.i.d. across i , which implies that the counting processes $N_{i,\ell \rightarrow m}$ defined by $N_{i,\ell \rightarrow m}(t) = \#\{s \leq t : X_i(s-) = \ell \text{ and } X_i(s) = m\}$ are i.i.d. across i as well. \square

Proof of (C3). By Theorem II.6.8 in Andersen et al. (1993) and absolute continuity of the intensity measure as assumed in (C_{D1}), the process $N_{i,\ell \rightarrow m}$ has $\{\mathcal{F}_t\}$ -compensator $\Lambda_{i,\ell \rightarrow m}(t) = \int_0^t \lambda_{i,\ell \rightarrow m}(s) ds$, where $\lambda_{i,\ell \rightarrow m}(s) = Z_{i,\ell}(s) \alpha_{\ell \rightarrow m}^*(s)$ with $Z_{i,\ell}(s) = \mathbf{1}(X_i(s-) = \ell)$ and a deterministic, non-negative function $\alpha_{\ell \rightarrow m}^*$. Since $Z_{i,\ell}$ is adapted to $\{\mathcal{F}_t\}$ and left-continuous, it is $\{\mathcal{F}_t\}$ -predictable, which immediately implies that $\lambda_{i,\ell \rightarrow m}$ is $\{\mathcal{F}_t\}$ -predictable as well. \square

Proof of (C4). We have $\bar{Z}_\ell(t) = \sum_{i=1}^n Z_{i,\ell}(t)$ with $Z_{i,\ell}(t) = \mathbf{1}(X_i(t-) = \ell)$ and $J_\ell(t) = \mathbf{1}(\bar{Z}_\ell(t) > 0)$. Hence, with the convention $0/0 := 0$, the processes \bar{Z}_ℓ and J_ℓ/\bar{Z}_ℓ are obviously bounded, which in particular implies that they are locally bounded. \square

Proof of (C5). With $\mathcal{T} = [\tau_{\min}, \tau_{\max}]$, it holds that

$$\begin{aligned} \mathbb{P}\left(\inf_{t \in \mathcal{T}} J_\ell(t) = 0\right) &= \mathbb{P}\left(\inf_{t \in \mathcal{T}} \sum_{i=1}^n \mathbf{1}(X_i(t-) = \ell) = 0\right) \\ &\leq \mathbb{P}\left(\sum_{i=1}^n \inf_{t \in \mathcal{T}} \mathbf{1}(X_i(t-) = \ell) = 0\right) \\ &= \prod_{i=1}^n \left\{1 - \mathbb{P}(X_i(t-) = \ell \text{ for all } t \in \mathcal{T})\right\} \\ &= (1 - p_\ell)^n = o(1), \end{aligned}$$

where we have used (C_{D2}) in the last line. \square

Proof of (C6). As before, let $\bar{Z}_\ell(t) = \sum_{i=1}^n Z_{i,\ell}(t)$ with $Z_{i,\ell}(t) = \mathbf{1}(X_i(t-) = \ell)$ and $J_\ell(t) = \mathbf{1}(\bar{Z}_\ell(t) > 0)$. Moreover, define $\bar{Z}_{\ell,\mathcal{T}} = \sum_{i=1}^n \inf_{t \in \mathcal{T}} Z_{i,\ell}(t)$ and $J_{\ell,\mathcal{T}} = \mathbf{1}(\bar{Z}_{\ell,\mathcal{T}} > 0)$. With the convention $0/0 := 0$, it holds that

$$\begin{aligned} \mathbb{E}\left[\left\|\frac{J_\ell}{\bar{Z}_\ell}\right\|_\infty^\nu\right] &= \mathbb{E}\left[\sup_{t \in \mathcal{T}} \left|\frac{J_\ell(t)}{\bar{Z}_\ell(t)}\right|^\nu \mathbf{1}\left(\inf_{t \in \mathcal{T}} J_\ell(t) > 0\right)\right] + \mathbb{E}\left[\sup_{t \in \mathcal{T}} \left|\frac{J_\ell(t)}{\bar{Z}_\ell(t)}\right|^\nu \mathbf{1}\left(\inf_{t \in \mathcal{T}} J_\ell(t) = 0\right)\right] \\ &\leq \mathbb{E}\left[\sup_{t \in \mathcal{T}} \left|\frac{J_\ell(t)}{\bar{Z}_\ell(t)}\right|^\nu \mathbf{1}\left(\inf_{t \in \mathcal{T}} J_\ell(t) > 0\right)\right] + \mathbb{P}\left(\inf_{t \in \mathcal{T}} J_\ell(t) = 0\right) \\ &\leq \mathbb{E}\left[\sup_{t \in \mathcal{T}} \left|\frac{J_\ell(t)}{\bar{Z}_\ell(t)}\right|^\nu \mathbf{1}\left(\inf_{t \in \mathcal{T}} J_\ell(t) > 0 \wedge \inf_{t \in \mathcal{T}} J_\ell(t) = J_{\ell,\mathcal{T}}\right)\right] \\ &\quad + \mathbb{E}\left[\sup_{t \in \mathcal{T}} \left|\frac{J_\ell(t)}{\bar{Z}_\ell(t)}\right|^\nu \mathbf{1}\left(\inf_{t \in \mathcal{T}} J_\ell(t) > 0 \wedge \inf_{t \in \mathcal{T}} J_\ell(t) > J_{\ell,\mathcal{T}}\right)\right] \\ &\quad + \mathbb{P}\left(\inf_{t \in \mathcal{T}} J_\ell(t) = 0\right) \end{aligned}$$

$$\leq \mathbb{E} \left[\left| \frac{J_{\ell, \mathcal{T}}}{\bar{Z}_{\ell, \mathcal{T}}} \right|^\nu \right] + \mathbb{P}(J_{\ell, \mathcal{T}} = 0) + \mathbb{P} \left(\inf_{t \in \mathcal{T}} J_\ell(t) = 0 \right).$$

From the proof of (C5), we already know that $\mathbb{P}(J_{\ell, \mathcal{T}} = 0) \leq (1 - p_\ell)^n$ and $\mathbb{P}(\inf_{t \in \mathcal{T}} J_\ell(t) = 0) \leq (1 - p_\ell)^n$. Moreover, since $\bar{Z}_{\ell, \mathcal{T}} = \sum_{i=1}^n \inf_{t \in \mathcal{T}} \mathbf{1}(X_i(t-) = \ell)$ follows a binomial distribution $\text{Bin}(n, q)$ with $q = \mathbb{P}(X_i(t-) = \ell \text{ for all } t \in \mathcal{T}) = p_\ell$ under (C_D2), we can apply Lemma A.5 to get that

$$\mathbb{E} \left[\left| \frac{J_{\ell, \mathcal{T}}}{\bar{Z}_{\ell, \mathcal{T}}} \right|^\nu \right] \leq \frac{(\nu + 1)!}{\{np_\ell\}^\nu}.$$

Putting everything together, we arrive at

$$\mathbb{E} \left[\left\| \frac{J_\ell}{\bar{Z}_\ell} \right\|_\infty^\nu \right] \leq \frac{(\nu + 1)!}{\{np_\ell\}^\nu} + 2(1 - p_\ell)^n$$

for any $\nu \in \mathbb{N}$. As $(1 - p_\ell)^n$ converges to zero faster than $n^{-\nu}$, (C6) is an immediate consequence of this bound. \square

Proof of (C7). Since $\Delta_k^\beta = 0$ for all k in the present case (without covariates), (C7) is trivially fulfilled. \square

B Proof of auxiliary results

For completeness, we now provide proofs of Propositions A.1 and A.4.

Proof of Proposition A.1

The proof is essentially a reformulation of the arguments for Theorem 3.1 in Zhang (2019). As a preliminary step, we analyze the minimization problem

$$\underset{a_1, \dots, a_r \in \mathbb{R}}{\text{argmin}} \mathcal{L}(z_1, \dots, z_r, a_1, \dots, a_r \mid \underline{a}, \bar{a}) \quad (\text{B.1})$$

for given values $z_1, \dots, z_r, \underline{a}, \bar{a} \in \mathbb{R}$, where

$$\begin{aligned} & \mathcal{L}(z_1, \dots, z_r, a_1, \dots, a_r \mid \underline{a}, \bar{a}) \\ &= \sum_{j=1}^r (z_j - a_j)^2 + \gamma \left\{ \sum_{j=1}^{r-1} |a_{j+1} - a_j| + |a_1 - \underline{a}| + |a_r - \bar{a}| \right\} \end{aligned} \quad (\text{B.2})$$

with some penalization parameter $\gamma > 0$. We first derive certain statements about the structure of the solution vector

$$(\hat{a}_1, \dots, \hat{a}_r) \in \underset{a_1, \dots, a_r \in \mathbb{R}}{\text{argmin}} \mathcal{L}(z_1, \dots, z_r, a_1, \dots, a_r \mid \underline{a}, \bar{a})$$

in Lemmas B.1 and B.2 and then use these statements to obtain bounds on the components \hat{a}_k for $k \in \{1, \dots, r\}$ in Lemma B.3. For ease of notation, we write $\hat{a}_0 := \underline{a}$ and $\hat{a}_{r+1} := \bar{a}$ in what follows.

Lemma B.1. *Assume that*

$$\left| \frac{1}{\sqrt{r_2 - r_1 + 1}} \sum_{j=r_1}^{r_2} z_j \right| \leq C_z \quad (\text{B.3})$$

for some r_1, r_2 with $1 \leq r_1 \leq r_2 \leq r$. Then the solution vector $(\hat{a}_1, \dots, \hat{a}_r)$ has the following properties:

(i) If $\hat{a}_{r_1-1} < \hat{a}_{r_1} = \dots = \hat{a}_{r_2} > \hat{a}_{r_2+1}$, then

$$\hat{a}_{r_1} = \dots = \hat{a}_{r_2} \leq \frac{C_z^2}{4\gamma}.$$

(ii) If $\hat{a}_{r_1-1} > \hat{a}_{r_1} = \dots = \hat{a}_{r_2} < \hat{a}_{r_2+1}$, then

$$\hat{a}_{r_1} = \dots = \hat{a}_{r_2} \geq -\frac{C_z^2}{4\gamma}.$$

Lemma B.1 essentially says the following: if condition (B.3) is satisfied for all r_1, r_2 with $1 \leq r_1 \leq r_2 \leq r$, then the sequence $\hat{a}_0, \hat{a}_1, \dots, \hat{a}_r, \hat{a}_{r+1}$ cannot have a local minimum smaller than $-C_z^2/(4\gamma)$ or a local maximum larger than $C_z^2/(4\gamma)$.

Proof. We only prove (i) since (ii) can be verified by analogous arguments. For any $\varepsilon > 0$ smaller than $\min\{\hat{a}_{r_1} - \hat{a}_{r_1-1}, \hat{a}_{r_2} - \hat{a}_{r_2+1}\}$, define

$$\tilde{a}_j = \begin{cases} \hat{a}_j - \varepsilon & \text{for } j \in \{r_1, \dots, r_2\} \\ \hat{a}_j & \text{for } j \notin \{r_1, \dots, r_2\}. \end{cases} \quad (\text{B.4})$$

Direct calculations yield that

$$\sum_{j=1}^r (z_j - \hat{a}_j)^2 - \sum_{j=1}^r (z_j - \tilde{a}_j)^2 = 2\varepsilon \sum_{j=r_1}^{r_2} (\hat{a}_j - z_j) - (r_2 - r_1 + 1)\varepsilon^2$$

and

$$\begin{aligned} \sum_{j=0}^r |\hat{a}_{j+1} - \hat{a}_j| - \sum_{j=0}^r |\tilde{a}_{j+1} - \tilde{a}_j| &= |\hat{a}_{r_1} - \hat{a}_{r_1-1}| - |\hat{a}_{r_1} - \hat{a}_{r_1-1} - \varepsilon| \\ &\quad + |\hat{a}_{r_2+1} - \hat{a}_{r_2}| - |\hat{a}_{r_2+1} - \hat{a}_{r_2} + \varepsilon| = 2\varepsilon, \end{aligned}$$

where the last equality uses that $\hat{a}_{r_1} - \hat{a}_{r_1-1} > 0$, $\hat{a}_{r_2} - \hat{a}_{r_2+1} > 0$ and $0 < \varepsilon <$

$\min\{\hat{a}_{r_1} - \hat{a}_{r_1-1}, \hat{a}_{r_2} - \hat{a}_{r_2+1}\}$. By definition of $(\hat{a}_1, \dots, \hat{a}_r)$, it holds that

$$\mathcal{L}(z_1, \dots, z_r, \hat{a}_1, \dots, \hat{a}_r | \underline{a}, \bar{a}) \leq \mathcal{L}(z_1, \dots, z_r, \tilde{a}_1, \dots, \tilde{a}_r | \underline{a}, \bar{a}).$$

Hence, it follows that

$$\begin{aligned} 0 &\geq \mathcal{L}(z_1, \dots, z_r, \hat{a}_1, \dots, \hat{a}_r | \underline{a}, \bar{a}) - \mathcal{L}(z_1, \dots, z_r, \tilde{a}_1, \dots, \tilde{a}_r | \underline{a}, \bar{a}) \\ &= \left[\sum_{j=1}^r (z_j - \hat{a}_j)^2 + \gamma \left(|\hat{a}_1 - \underline{a}| + |\hat{a}_r - \bar{a}| + \sum_{j=1}^{r-1} |\hat{a}_{j+1} - \hat{a}_j| \right) \right] \\ &\quad - \left[\sum_{j=1}^r (z_j - \tilde{a}_j)^2 + \gamma \left(|\tilde{a}_1 - \underline{a}| + |\tilde{a}_r - \bar{a}| + \sum_{j=1}^{r-1} |\tilde{a}_{j+1} - \tilde{a}_j| \right) \right] \\ &= \left[\sum_{j=1}^r (z_j - \hat{a}_j)^2 - \sum_{j=1}^r (z_j - \tilde{a}_j)^2 \right] + \gamma \left[\sum_{j=0}^r |\hat{a}_{j+1} - \hat{a}_j| - \sum_{j=0}^r |\tilde{a}_{j+1} - \tilde{a}_j| \right] \\ &= 2\varepsilon \sum_{j=r_1}^{r_2} (\hat{a}_j - z_j) - (r_2 - r_1 + 1)\varepsilon^2 + 2\gamma\varepsilon, \end{aligned}$$

that is,

$$0 \geq \sum_{j=r_1}^{r_2} (\hat{a}_j - z_j) - \frac{(r_2 - r_1 + 1)\varepsilon}{2} + \gamma.$$

Letting $\varepsilon \rightarrow 0$, we obtain that

$$0 \geq \sum_{j=r_1}^{r_2} (\hat{a}_j - z_j) + \gamma = (r_2 - r_1 + 1)\hat{a}_{r_1} - \sum_{j=r_1}^{r_2} z_j + \gamma,$$

which in turn implies that

$$\hat{a}_{r_1} \leq \frac{\sum_{j=r_1}^{r_2} z_j - \gamma}{r_2 - r_1 + 1}.$$

Finally, using the assumption that $\sum_{j=r_1}^{r_2} z_j \leq C_z \sqrt{r_2 - r_1 + 1}$ and the inequality $\frac{a}{x} - \frac{b}{x^2} \leq \frac{a^2}{4b}$ for $x > 0$, we arrive at

$$\hat{a}_{r_1} \leq \frac{C_z \sqrt{r_2 - r_1 + 1} - \gamma}{r_2 - r_1 + 1} \leq \frac{C_z^2}{4\gamma}. \quad \square$$

Lemma B.2. *Assume that*

$$\max_{1 \leq r_1 \leq r_2 \leq r} \left| \frac{1}{\sqrt{r_2 - r_1 + 1}} \sum_{j=r_1}^{r_2} z_j \right| \leq C_z. \quad (\text{B.5})$$

Then the solution vector $(\hat{a}_1, \dots, \hat{a}_r)$ has one of the following patterns:

Pattern A: There exist $r_1, r_2 \in \{0, \dots, r+1\}$ with $r_1 \leq r_2$ such that

$$\hat{a}_0 \geq \dots \geq \hat{a}_{r_1-1} > \frac{C_z^2}{4\gamma} \text{ or } \hat{a}_0 \leq \dots \leq \hat{a}_{r_1-1} < -\frac{C_z^2}{4\gamma},$$

$$|\hat{a}_j| \leq \frac{C_z^2}{4\gamma} \text{ for all } j \in \{r_1, \dots, r_2\}, \text{ and}$$

$$\frac{C_z^2}{4\gamma} < \hat{a}_{r_2+1} \leq \dots \leq \hat{a}_{r+1} \text{ or } -\frac{C_z^2}{4\gamma} > \hat{a}_{r_2+1} \geq \dots \geq \hat{a}_{r+1}.$$

Pattern B: There exists $r_0 \in \{0, \dots, r\}$ such that either (a) or (b) holds:

$$(a) \hat{a}_0 \geq \dots \geq \hat{a}_{r_0} > \frac{C_z^2}{4\gamma} \text{ and } -\frac{C_z^2}{4\gamma} > \hat{a}_{r_0+1} \geq \dots \geq \hat{a}_{r+1}$$

$$(b) \hat{a}_0 \leq \dots \leq \hat{a}_{r_0} < -\frac{C_z^2}{4\gamma} \text{ and } \frac{C_z^2}{4\gamma} < \hat{a}_{r_0+1} \leq \dots \leq \hat{a}_{r+1}.$$

Pattern C: There exist $r_1, r_2 \in \{0, \dots, r+1\}$ with $r_1 \leq r_2$ such that either (a) or (b) holds:

$$(a) \hat{a}_j > \frac{C_z^2}{4\gamma} \text{ for all } j \text{ and } \hat{a}_0 \geq \dots > \hat{a}_{r_1} = \dots = \hat{a}_{r_2} < \dots \leq \hat{a}_{r+1}$$

$$(b) \hat{a}_j < -\frac{C_z^2}{4\gamma} \text{ for all } j \text{ and } \hat{a}_0 \leq \dots < \hat{a}_{r_1} = \dots = \hat{a}_{r_2} > \dots \geq \hat{a}_{r+1}.$$

Proof. Let $\mathcal{I} = \{0 \leq j \leq r+1 : |\hat{a}_j| \leq C_z^2/(4\gamma)\}$. From Lemma B.1, we already know that the sequence $\hat{a}_0, \hat{a}_1, \dots, \hat{a}_r, \hat{a}_{r+1}$ cannot have a local minimum smaller than $-C_z^2/(4\gamma)$ or a local maximum larger than $C_z^2/(4\gamma)$ under condition (B.5). This implies that \mathcal{I} must be an interval of the form $\mathcal{I} = \{r_1, \dots, r_2\}$ for some $0 \leq r_1 \leq r_2 \leq r+1$ or the empty set.

We first examine the case where \mathcal{I} is a non-empty interval of the form $\mathcal{I} = \{r_1, \dots, r_2\}$. Since the sequence $\hat{a}_0, \hat{a}_1, \dots, \hat{a}_r, \hat{a}_{r+1}$ cannot have a local maximum (minimum) larger than $C_z^2/(4\gamma)$ (smaller than $-C_z^2/(4\gamma)$), the values $\hat{a}_0, \dots, \hat{a}_{r_1-1}$ must either form a decreasing sequence with $\hat{a}_0 \geq \dots \geq \hat{a}_{r_1-1} > C_z^2/(4\gamma)$ or an increasing sequence with $\hat{a}_0 \leq \dots \leq \hat{a}_{r_1-1} < -C_z^2/(4\gamma)$. The values $\hat{a}_{r_2+1}, \dots, \hat{a}_{r+1}$ must fulfill an analogous restriction. We thus end up with pattern A.

Next consider the case where \mathcal{I} is the empty set and thus $|\hat{a}_j| > C_z^2/(4\gamma)$ for all $j \in \{0, \dots, r+1\}$. There are two subcases: either there are \hat{a}_j 's with different signs or all \hat{a}_j 's have the same sign. We start with the subcase where the \hat{a}_j 's have different signs. To avoid a violation of Lemma B.1 by producing a local maximum/minimum which is too large/small, the following must hold: The signs of the \hat{a}_j 's can only change once, that is, there exists r_0 such that \hat{a}_j is positive (negative) for all $j \leq r_0$ and negative (positive) for all $j > r_0$. Moreover, when the signs are first positive (negative) and then negative (positive), the sequence $\hat{a}_0, \hat{a}_1, \dots, \hat{a}_r, \hat{a}_{r+1}$ must be decreasing (increasing). We thus end up with pattern B. We now turn to the second subcase where all \hat{a}_j 's have the same sign. If $\hat{a}_j > C_z^2/(4\gamma)$ for all j , we must have that $\hat{a}_0 \geq \dots > \hat{a}_{r_1} = \dots = \hat{a}_{r_2} < \dots \leq \hat{a}_{r+1}$ for some indices r_1 and r_2 because otherwise the sequence $\hat{a}_0, \hat{a}_1, \dots, \hat{a}_r, \hat{a}_{r+1}$ would have a local maximum larger than $C_z^2/(4\gamma)$. Analogously, if $\hat{a}_j < -C_z^2/(4\gamma)$ for all j , we must

have that $\hat{a}_0 \leq \dots < \hat{a}_{r_1} = \dots = \hat{a}_{r_2} > \dots \geq \hat{a}_{r+1}$. As a result, we obtain pattern C. \square

With the help of Lemmas B.1 and B.2, we now prove that the entries of the solution vector $(\hat{a}_1, \dots, \hat{a}_r)$ can be bounded as follows.

Lemma B.3. *If*

$$\max_{1 \leq r_1 \leq r_2 \leq r} \left| \frac{1}{\sqrt{r_2 - r_1 + 1}} \sum_{j=r_1}^{r_2} z_j \right| \leq C_z, \quad (\text{B.6})$$

then

$$|\hat{a}_j| \leq \max \left\{ \frac{C_z}{\sqrt{j}}, \frac{C_z}{\sqrt{r+1-j}}, \frac{C_z^2}{4\gamma}, \frac{2\gamma}{r} + \frac{2C_z}{\sqrt{r}} \right\} \quad (\text{B.7})$$

for all $j \in \{1, \dots, r\}$.

Proof. From Lemma B.2, we already know that the vector $(\hat{a}_0, \hat{a}_1, \dots, \hat{a}_r, \hat{a}_{r+1})$ has pattern A, B or C. We now examine these three patterns one after the other.

First suppose that the vector has pattern A. Without loss of generality, assume that $\hat{a}_0 \geq \dots \geq \hat{a}_{r_1-1} > C_z^2/(4\gamma)$ and $C_z^2/(4\gamma) < \hat{a}_{r_2+1} \leq \dots \leq \hat{a}_{r+1}$ with $1 \leq r_1 \leq r_2 \leq r$. The other possibilities entailed by pattern A can be treated analogously. To start with, we have a closer look at the partial sequence $\hat{a}_0, \dots, \hat{a}_{r_1-1}$. For any $k \leq r_1 - 1$, define $k' := k'(k)$ as the biggest index such that $\hat{a}_k = \hat{a}_{k+1} = \dots = \hat{a}_{k'} \neq \hat{a}_{k'+1}$. Note that $k' \leq r_1 - 1$ since $\hat{a}_{r_1-1} \neq \hat{a}_{r_1}$. Defining

$$\tilde{a}_j = \begin{cases} \hat{a}_j - \varepsilon & \text{for } 1 \leq j \leq k' \\ \hat{a}_j & \text{else} \end{cases}$$

with $\varepsilon > 0$, we obtain that

$$\begin{aligned} \sum_{j=0}^r |\hat{a}_{j+1} - \hat{a}_j| - \sum_{j=0}^r |\tilde{a}_{j+1} - \tilde{a}_j| &= |\hat{a}_1 - \hat{a}_0| - |\hat{a}_1 - \hat{a}_0 - \varepsilon| \\ &\quad + |\hat{a}_{k'+1} - \hat{a}_{k'}| - |\hat{a}_{k'+1} - \hat{a}_{k'} + \varepsilon| = 0 \end{aligned}$$

for sufficiently small $\varepsilon > 0$ (taking into account that $\hat{a}_{k'+1} - \hat{a}_{k'} < 0$ and $\hat{a}_1 - \hat{a}_0 \leq 0$) and

$$\sum_{j=1}^r (z_j - \hat{a}_j)^2 - \sum_{j=1}^r (z_j - \tilde{a}_j)^2 = 2\varepsilon \sum_{j=1}^{k'} (\hat{a}_j - z_j) - k'\varepsilon^2.$$

This implies that

$$0 \geq \mathcal{L}(z_1, \dots, z_r, \hat{a}_1, \dots, \hat{a}_r \mid \underline{a}, \bar{a}) - \mathcal{L}(z_1, \dots, z_r, \tilde{a}_1, \dots, \tilde{a}_r \mid \underline{a}, \bar{a})$$

$$\begin{aligned}
&= \left[\sum_{j=1}^r (z_j - \hat{a}_j)^2 - \sum_{j=1}^r (z_j - \tilde{a}_j)^2 \right] + \gamma \left[\sum_{j=0}^r |\hat{a}_{j+1} - \hat{a}_j| - \sum_{j=0}^r |\tilde{a}_{j+1} - \tilde{a}_j| \right] \\
&= 2\varepsilon \sum_{j=1}^{k'} (\hat{a}_j - z_j) - k'\varepsilon^2
\end{aligned}$$

and thus

$$0 \geq \sum_{j=1}^{k'} (\hat{a}_j - z_j) - \frac{k'\varepsilon}{2}.$$

Letting $\varepsilon \rightarrow 0$, we further get that

$$0 \geq \sum_{j=1}^{k'} (\hat{a}_j - z_j) \geq k'\hat{a}_{k'} - \sum_{j=1}^{k'} z_j = k'\hat{a}_k - \sum_{j=1}^{k'} z_j.$$

Rearranging, inserting the bound (B.6) and taking into account that $k \leq k'$ yields

$$\hat{a}_k \leq \frac{C_z \sqrt{k'}}{k'} = \frac{C_z}{\sqrt{k'}} \leq \frac{C_z}{\sqrt{k}} \quad (\text{B.8})$$

for any $k \in \{1, \dots, r_1 - 1\}$. Analyzing the partial sequence $\hat{a}_{r_2+1}, \dots, \hat{a}_{r+1}$ in a similar way shows that

$$\hat{a}_k \leq \frac{C_z}{\sqrt{r-k+1}} \quad (\text{B.9})$$

for any $k \in \{r_2 + 1, \dots, r\}$. Combining the bounds (B.8) and (B.9) with the fact that $|\hat{a}_j| \leq C_z^2/(4\gamma)$ for $j \in \{r_1, \dots, r_2\}$, we finally arrive at the desired bound (B.7).

We next suppose that the vector $(\hat{a}_0, \hat{a}_1, \dots, \hat{a}_r, \hat{a}_{r+1})$ has pattern B. In this case, we can apply the same arguments as for pattern A to the partial sequences $\hat{a}_0, \dots, \hat{a}_{r_0}$ and $\hat{a}_{r_0+1}, \dots, \hat{a}_{r+1}$ to verify the bound (B.7).

We finally turn to the case that $(\hat{a}_0, \hat{a}_1, \dots, \hat{a}_r, \hat{a}_{r+1})$ has pattern C. Without loss of generality, we let $1 \leq r_1 \leq r_2 \leq r$ and consider case (a), that is, $\hat{a}_j > C_z^2/(4\gamma)$ for all j and $\hat{a}_0 \geq \dots > \hat{a}_{r_1} = \dots = \hat{a}_{r_2} < \dots \leq \hat{a}_{r+1}$. The other possibilities entailed by pattern C can be treated analogously. The same arguments as for pattern A yield that

$$\hat{a}_k \leq \frac{C_z}{\sqrt{k}} \quad \text{for } k \in \{1, \dots, r_1 - 1\}$$

and

$$\hat{a}_k \leq \frac{C_z}{\sqrt{r-k+1}} \quad \text{for } k \in \{r_2 + 1, \dots, r\}.$$

It thus remains to bound \hat{a}_k for $k \in \mathcal{I} = \{r_1, \dots, r_2\}$. Defining

$$\tilde{a}_j = \begin{cases} \hat{a}_j - \varepsilon & \text{for } j \in \mathcal{I} \\ \hat{a}_j & \text{for } j \notin \mathcal{I} \end{cases}$$

with $\varepsilon > 0$, we get that

$$\begin{aligned} & \sum_{j=0}^r |\hat{a}_{j+1} - \hat{a}_j| - \sum_{j=0}^r |\tilde{a}_{j+1} - \tilde{a}_j| \\ &= |\hat{a}_{r_1} - \hat{a}_{r_1-1}| - |\hat{a}_{r_1} - \hat{a}_{r_1-1} - \varepsilon| + |\hat{a}_{r_2+1} - \hat{a}_{r_2}| - |\hat{a}_{r_2+1} - \hat{a}_{r_2} + \varepsilon| = -2\varepsilon \end{aligned}$$

for sufficiently small $\varepsilon > 0$ (taking into consideration that $\hat{a}_{r_1} - \hat{a}_{r_1-1} < 0$ and $\hat{a}_{r_2+1} - \hat{a}_{r_2} > 0$) and

$$\sum_{j=1}^r (z_j - \hat{a}_j)^2 - \sum_{j=1}^r (z_j - \tilde{a}_j)^2 = 2\varepsilon \sum_{j=r_1}^{r_2} (\hat{a}_j - z_j) - (r_2 - r_1 + 1)\varepsilon^2.$$

We thus obtain that

$$\begin{aligned} 0 &\geq \mathcal{L}(z_1, \dots, z_r, \hat{a}_1, \dots, \hat{a}_r \mid \underline{a}, \bar{a}) - \mathcal{L}(z_1, \dots, z_r, \tilde{a}_1, \dots, \tilde{a}_r \mid \underline{a}, \bar{a}) \\ &= \left[\sum_{j=1}^r (z_j - \hat{a}_j)^2 - \sum_{j=1}^r (z_j - \tilde{a}_j)^2 \right] + \gamma \left[\sum_{j=0}^r |\hat{a}_{j+1} - \hat{a}_j| - \sum_{j=0}^r |\tilde{a}_{j+1} - \tilde{a}_j| \right] \\ &= 2\varepsilon \sum_{j=r_1}^{r_2} (\hat{a}_j - z_j) - (r_2 - r_1 + 1)\varepsilon^2 - 2\gamma\varepsilon, \end{aligned}$$

which in turn yields that

$$0 \geq \sum_{j=r_1}^{r_2} (\hat{a}_j - z_j) - \frac{(r_2 - r_1 + 1)\varepsilon}{2} - \gamma.$$

Letting $\varepsilon \rightarrow 0$, we can infer that

$$0 \geq \sum_{j=r_1}^{r_2} (\hat{a}_j - z_j) - \gamma = (r_2 - r_1 + 1)\hat{a}_{r_1} - \sum_{j=r_1}^{r_2} z_j - \gamma$$

and thus

$$\hat{a}_{r_1} \leq \frac{\sum_{j=r_1}^{r_2} z_j + \gamma}{r_2 - r_1 + 1}.$$

Rearranging this inequality, recalling that $\hat{a}_{r_1} = \hat{a}_k$ for all $k \in \mathcal{I}$ and inserting the bound (B.6) yields

$$\hat{a}_k \leq \frac{C_z \sqrt{r_2 - r_1 + 1} + \gamma}{r_2 - r_1 + 1} = \frac{C_z}{\sqrt{r_2 - r_1 + 1}} + \frac{\gamma}{r_2 - r_1 + 1} \quad (\text{B.10})$$

for any $k \in \mathcal{I}$. Since we deal with pattern 3(a), it additionally holds that

$$\hat{a}_k \leq \min \left\{ \frac{C_z}{\sqrt{r_1 - 1}}, \frac{C_z}{\sqrt{r - r_2}} \right\} \quad (\text{B.11})$$

for any $k \in \mathcal{I}$. Combining the bounds (B.10) and (B.11), we can conclude the following:

- If $r_1 - 1 \geq r/4$, then $\hat{a}_k \leq \frac{2C_z}{\sqrt{r}}$ for $k \in \mathcal{I}$.
- If $r - r_2 \geq r/4$, then $\hat{a}_k \leq \frac{2C_z}{\sqrt{r}}$ for $k \in \mathcal{I}$.
- If $r_1 - 1 \leq r/4$ and $r - r_2 \leq r/4$, then $r_2 - r_1 + 1 \geq r/2$ and thus $\hat{a}_k \leq \frac{C_z}{\sqrt{r/2}} + \frac{2\gamma}{r}$ for $k \in \mathcal{I}$.

Combining these bounds for $k \in \mathcal{I}$ with those derived for $k \notin \mathcal{I}$, we finally arrive at the desired bound (B.7). \square

To prove Proposition A.1, we now connect the minimization problem underlying the fused lasso to the minimization problem (B.1) analyzed above. For simplicity, we denote the fused lasso by $\hat{\boldsymbol{\alpha}}$, that is, we drop the subscript λ . It holds that

$$\hat{\boldsymbol{\alpha}} \in \underset{\mathbf{a} \in \mathbb{R}^n}{\text{argmin}} \mathcal{L}(\mathbf{Y}, \mathbf{a})$$

with

$$\mathcal{L}(\mathbf{Y}, \mathbf{a}) = \frac{1}{n} \sum_{j=1}^n (Y_j - a_j)^2 + \lambda \sum_{j=1}^{n-1} |a_{j+1} - a_j|.$$

Let $1 < n_1 < \dots < n_K \leq n$ be the indices where the signal vector $\boldsymbol{\alpha}^* = (\alpha_1^*, \dots, \alpha_n^*)$ has a jump (i.e., $\alpha_{n_{k-1}}^* \neq \alpha_{n_k}^*$ for $1 \leq k \leq K$), and additionally set $n_0 = 1$ along with $n_{K+1} = n + 1$. Given the indices n_1, \dots, n_K , any vector $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$ can be represented as $\mathbf{x} = (\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_K)$ with $\mathbf{x}_k = (x_{n_k}, \dots, x_{n_{k+1}-1})$ for $k \in \{0, \dots, K\}$. Obviously, for any $k \in \{1, \dots, K-1\}$, it holds that

$$\begin{aligned} \hat{\boldsymbol{\alpha}}_k &\in \underset{\mathbf{a}_k \in \mathbb{R}^{n_{k+1}-n_k+1}}{\text{argmin}} \mathcal{L}(\mathbf{Y}, \hat{\boldsymbol{\alpha}}_0, \dots, \hat{\boldsymbol{\alpha}}_{k-1}, \mathbf{a}_k, \hat{\boldsymbol{\alpha}}_{k+1}, \dots, \hat{\boldsymbol{\alpha}}_K) \\ &= \underset{\mathbf{a}_k \in \mathbb{R}^{n_{k+1}-n_k+1}}{\text{argmin}} \mathcal{L}(\mathbf{Y}_k, \mathbf{a}_k \mid \hat{\alpha}_{n_{k-1}}, \hat{\alpha}_{n_{k+1}}), \end{aligned}$$

where

$$\begin{aligned} &\mathcal{L}(\mathbf{Y}_k, \mathbf{a}_k \mid \hat{\alpha}_{n_{k-1}}, \hat{\alpha}_{n_{k+1}}) \\ &= \sum_{j=n_k}^{n_{k+1}-1} (Y_j - a_j)^2 + \gamma \sum_{j=n_k}^{n_{k+1}-2} |a_{j+1} - a_j| \\ &\quad + \gamma \{ |a_{n_k} - \hat{\alpha}_{n_{k-1}}| + |\hat{\alpha}_{n_{k+1}} - a_{n_{k+1}-1}| \} \end{aligned}$$

is the criterion function defined in (B.2) with $\gamma = n\lambda$. This shows that for any $k \in \{1, \dots, K-1\}$, the subvector $\hat{\boldsymbol{\alpha}}_k$ can be obtained by solving a subproblem with correctly specified end points, in particular, by minimizing the criterion function $\mathcal{L}(\mathbf{Y}_k, \mathbf{a}_k | \underline{a}, \bar{a})$ with the end points $\underline{a} = \hat{\alpha}_{n_{k-1}}$ and $\bar{a} = \hat{\alpha}_{n_{k+1}}$. It is easy to see that for any constant vector $\mathbf{c}_k = (c, \dots, c) \in \mathbb{R}^{n_{k+1}-n_k+1}$,

$$(\hat{\boldsymbol{\alpha}}_k - \mathbf{c}_k) \in \underset{\mathbf{a}_k \in \mathbb{R}^{n_{k+1}-n_k+1}}{\operatorname{argmin}} \mathcal{L}(\mathbf{Y}_k - \mathbf{c}_k, \mathbf{a}_k | \hat{\alpha}_{n_{k-1}} - c, \hat{\alpha}_{n_{k+1}} - c).$$

Hence, if we shift the input vector \mathbf{Y}_k as well as the end points by a constant c , the solution vector $\hat{\boldsymbol{\alpha}}_k$ gets shifted by the same amount. Since the k -th part of the signal vector $\boldsymbol{\alpha}_k^*$ is constant, this in particular implies that

$$(\hat{\boldsymbol{\alpha}}_k - \boldsymbol{\alpha}_k^*) \in \underset{\mathbf{a}_k \in \mathbb{R}^{n_{k+1}-n_k+1}}{\operatorname{argmin}} \mathcal{L}(\mathbf{u}_k, \mathbf{a}_k | \hat{\alpha}_{n_{k-1}} - \alpha_{n_k}^*, \hat{\alpha}_{n_{k+1}} - \alpha_{n_k}^*), \quad (\text{B.12})$$

where $\mathbf{u}_k = \mathbf{Y}_k - \boldsymbol{\alpha}_k^*$. Assuming that

$$\max_{1 \leq k \leq \ell \leq n} \left| \frac{1}{\sqrt{\ell - k + 1}} \sum_{j=k}^{\ell} u_j \right| \leq C_u,$$

we can now apply Lemma B.3 to obtain that

$$|\hat{\alpha}_j - \alpha_j^*| \leq \max \left\{ \frac{C_u}{\sqrt{d_j}}, \frac{C_u^2}{4n\lambda}, \frac{2n\lambda}{r_{k(j)}} + \frac{2C_u}{\sqrt{r_{k(j)}}} \right\} \quad (\text{B.13})$$

for $j \in \{n_1, \dots, n_K - 1\}$.

The subvectors $\hat{\boldsymbol{\alpha}}_0 - \boldsymbol{\alpha}_0^*$ and $\hat{\boldsymbol{\alpha}}_K - \boldsymbol{\alpha}_K^*$ solve a minimization problem similar to (B.12), where only one instead of two endpoints are given. Repeating the above arguments for these slightly different minimization problems yields the bound (B.13) also for $j \in \{1, \dots, n_1 - 1\} \cup \{n_K, \dots, n\}$. This completes the proof.

Proof of Proposition A.4

Fix some $\varepsilon > 0$. By assumption, we know that there is a constant $C = C(\varepsilon)$ and an integer $N_1 = N_1(\varepsilon)$ such that

$$\mathbb{P} \left(\frac{1}{n} \|\tilde{\boldsymbol{\alpha}} - \boldsymbol{\alpha}^*\|_2^2 \geq \frac{C}{2} R_n \right) \leq \varepsilon$$

for all $n \geq N_1$. We also know that there is an integer $N_2 = N_2(\varepsilon) > 0$ with $2\lceil CnR_n/H_n^2 \rceil < D_n$ for all $n \geq N_2$. Let $N = \max\{N_1, N_2\}$, take $n \geq N$ and let $r_n = \lceil CnR_n/H_n^2 \rceil$. Suppose that $d(\tilde{J} | J^*) > r_n$. Then by definition, there exists a jump index $n_k \in J^*$ such that no jump points of $\tilde{\boldsymbol{\alpha}}$ are within r_n of n_k , which implies that $\tilde{\alpha}_j$

is constant over $j \in \{n_k - r_n, \dots, n_k + r_n - 1\}$. With the notation

$$z = \tilde{\alpha}_{n_k - r_n} = \dots = \tilde{\alpha}_{n_k - 1} = \tilde{\alpha}_{n_k} = \dots = \tilde{\alpha}_{n_k + r_n - 1},$$

we obtain the lower bound

$$\frac{1}{n} \sum_{j=n_k - r_n}^{n_k + r_n - 1} (\tilde{\alpha}_j - \alpha_j^*)^2 = \frac{r_n}{n} (z - \alpha_{n_k - 1}^*)^2 + \frac{r_n}{n} (z - \alpha_{n_k}^*)^2 \geq \frac{r_n H_n^2}{2n} \geq \frac{C}{2} R_n,$$

where the first inequality holds because $(x - a)^2 + (x - b)^2 \geq (a - b)^2/2$ for all x (the quadratic in x here is minimized at $x = (a + b)/2$) and the second one holds because $r_n = \lceil CnR_n/H_n^2 \rceil$. As a consequence, we see that $d(\tilde{J} | J^*) > r_n$ implies

$$\frac{1}{n} \|\tilde{\boldsymbol{\alpha}} - \boldsymbol{\alpha}^*\|_2^2 \geq \frac{1}{n} \sum_{j=n_k - r_n}^{n_k + r_n - 1} (\tilde{\alpha}_j - \alpha_j^*)^2 \geq \frac{C}{2} R_n,$$

which yields that

$$\mathbb{P}\left(d(\tilde{J} | J^*) > r_n\right) \leq \mathbb{P}\left(\frac{1}{n} \|\tilde{\boldsymbol{\alpha}} - \boldsymbol{\alpha}^*\|_2^2 \geq \frac{C}{2} R_n\right) \leq \varepsilon$$

for all $n \geq N$. This completes the proof.