Stochastic modeling, sampling effects, and efficient estimation with applications to disability insurance

Oliver Lunding Sandqvist

Department of Mathematical Sciences
Faculty of Science
University of Copenhagen
—
PFA Pension

This thesis has been submitted to the PhD School of The Faculty of Science, University of Copenhagen.

Industrial PhD Thesis by:

Oliver Lunding Sandqvist
Halls Allé 12, 2. tv.
DK-1802 Frederiksberg C
oliver.sandqvist@outlook.dk
https://oliversandqvist.github.io

Assessment committee:

Professor Mogens Bladt (chairperson)
University of Copenhagen, Denmark

Professor Marcus Christiansen
Carl von Ossietzky Universität Oldenburg, Germany

Professor Katrien Antonio
KU Leuven, Belgium

Supervisors:

Associate Professor Christian Furrer
University of Copenhagen

PhD Lars Frederik Brandt
PFA Pension

Professor Mogens Steffensen
University of Copenhagen

Affiliated Professor Kristian Buchardt
AP Pension & University of Copenhagen

PhD Niklas Lindholm
PFA Pension

This thesis has been submitted to the PhD School of The Faculty of Science, University of Copenhagen on 29 January 2025. It has been written as part of an Industrial PhD project with File No. 1044-00144B funded jointly by PFA Pension and Innovation Fund Denmark.

Kristian Buchardt supervised until he changed affiliation from PFA Pension to AP Pension in January 2023, while Niklas Lindholm supervised from January 2023.

ISBN: 978-87-7125-240-8

Preface

This thesis has been prepared in partial fulfillment of the requirements for the PhD degree at the Department of Mathematical Sciences, Faculty of Science, University of Copenhagen. The work was conducted between February 2022 and January 2025 as an Industrial PhD project within Innovation Fund Denmark's program and with PFA Pension as the industrial partner. The research was supervised by Associate Professor Christian Furrer (University of Copenhagen), PhD Lars Frederik Brandt (PFA Pension), Professor Mogens Steffensen (University of Copenhagen), Affiliated Professor Kristian Buchardt (PFA Pension until January 2023, then AP Pension; University of Copenhagen since September 2023), and PhD Niklas Lindholm (PFA Pension).

The thesis consists of an introduction and five manuscripts produced during my studies. Each manuscript is a stand-alone scientific contribution and can be read independently. Discrepancies in notation consequently appear across the chapters.

Acknowledgments

The work contained in this thesis is the result of three wonderful years of PhD studies. I am grateful to PFA Pension and the Innovation Fund Denmark for making this possible by funding the project.

A sincere thanks goes to my supervisors. I thank Mogens, Lars, and Niklas for all their input and for continuously supporting the project. I thank Kristian for mentoring me before and during my first year of PhD studies and for providing feedback and ideas that shaped the project and ensured its industrial relevance. I thank Christian for all the attention and effort he has dedicated to the project and to me, for being ready to discuss any theoretical or practical problem I may have, and for providing guidance and advice when necessary while also giving me the independence to forge my own path.

From September to December of 2023, I had the pleasure of visiting the Department of Mathematical Statistics at Stockholm University. I kindly thank Mathias Lindholm and Filip Lindskog for their enthusiasm and hospitality during my stay. I hope the future holds opportunities for continued discussions. I am also grateful to Janne, Henning, Taariq, Niels, and the rest of the PhD group at the department for welcoming me as one of their own and for many nights of board games which made my time in Stockholm even more enjoyable.

I thank my colleagues at PFA Pension and the University of Copenhagen for creating such a positive and supportive environment. It has been a pleasure coming to the office each day. I also appreciate the interest and encouragement offered by my family and friends over the past three years.

Lastly, a heartfelt thanks to my wife Ditte for her love and support, for always cheering on me, and for getting me away from my studies from time to time. This journey has meant so much more because I have had you to share it with.

Oliver Lunding Sandqvist
Copenhagen, January 2025

List of papers

With the exception of Chapter 1, which has been prepared specifically for this thesis, each of the remaining five chapters corresponds to a self-contained manuscript.

Chapter 2: C. Furrer and O. L. Sandqvist (2025). Loss of earning capacity in Denmark – an actuarial perspective. DOI: 10.48550/arXiv.2501.11578.

Chapter 3: K. Buchardt, C. Furrer, and O. L. Sandqvist (2023). Transaction time models in multi-state life insurance. *Scandinavian Actuarial Journal* **2023**, 974–999. DOI: 10.1080/03461238.2023.2181708.

Chapter 4: O. L. Sandqvist (2025). A multistate approach to disability insurance reserving with information delays. DOI: 10.48550/arXiv.2312.14324.

Chapter 5: K. Buchardt, C. Furrer, and O. L. Sandqvist (2025). Estimation for multistate models subject to reporting delays and incomplete event adjudication with application to disability insurance. DOI: 10.48550/arXiv.2311.04318.

Chapter 6: O. L. Sandqvist (2024). Doubly robust inference with censoring unbiased transformations. DOI: 10.48550/arXiv.2411.04909.

There may be minor differences between the contents of a chapter and the corresponding manuscript.

Abstract

This thesis consists of a series of independent investigations related to multistate modeling, unified by their relevance to actuarial modeling of disability insurance policies. Chapter 1 sets the stage for the investigations in the subsequent chapters and provides an overview of the thesis's main contributions.

Chapter 2 complements the introduction, describing practical challenges and opportunities for modeling and risk management of disability insurance portfolios. It is highlighted that the presence of public benefits, claim settlement processes, and prevention initiatives increases the complexity of the insurance business substantially and that these aspects have received limited attention in the literature. Subsequently, potential approaches and avenues for future research are outlined.

Chapters 3 and 4 are concerned with adapting multistate reserving models to situations where the payments and information might not be fully up to date. In Chapter 3, we consider a general formulation where the payments in real-time arise as running payments based on the available information and backpay based on the arrival of new information. In this setting, it is possible to link the present value to the classic present value based on the contractual payments and to characterize the dynamics of the reserve. However, it is in general not possible to link the reserve to the classic multistate reserve, and this hence has to be investigated in concrete models. Chapter 4 proposes a model for disability insurance schemes and imposes relevant conditional independence assumptions in order to obtain explicit expressions for the reserves as natural modifications of the classic multistate reserves. The potential financial impact is illustrated by applying the methods to a novel dataset LEC-DK19 which is based on real data that has been anonymized and slightly altered. The estimation procedure developed in Chapter 5 is used to operationalize the model.

In Chapter 5, the focus is on estimation of multistate models affected by various forms of missingness including reporting delays and incomplete event adjudication. An estimation procedure is proposed, accommodating reporting delays by thinning and incomplete adjudication by imputation, making effective use of the available data. The large sample properties of the estimator are established based on M-estimation theory. We demonstrate the approach on simulated data and on LEC-DK19.

The final chapter, Chapter 6, considers efficient and robust statistical inference in the presence of right-censoring. A flexible nonparametric estimation procedure based on pseudo-values and cross-fitting is proposed, allowing one to leverage prediction methods from machine learning. Large sample properties of the procedure are established, showing that the approach is doubly robust with respect to those nuisance parameters that are needed to compute the pseudo-values. A simulation study investigates the performance of the approach. Finally, we apply the estimation procedure to conduct a regression discontinuity design using real data.

Resumé

Denne afhandling består af en række uafhængige studier relateret til flertilstandsmodellering med det til fælles, at de er relevante for aktuarmæssig modellering af tab-af-erhvervsevne policer. Kapitel 1 sætter scenen for studierne i de efterfølgende kapitler og giver et overblik over afhandlingens hovedbidrag.

Kapitel 2 supplerer indledningen ved at beskrive praktiske udfordringer og muligheder for modellering og risikostyring af tab-af-erhvervsevne forsikringsbestande. Det fremhæves, at tilstedeværelsen af offentlige ydelser, skadesbehandlingsprocesser og forebyggelsesinitiativer øger kompleksiteten af forsikringsforretningen betydeligt, og at disse aspekter har fået begrænset opmærksomhed i litteraturen. Herefter skitseres potentielle tilgange og muligheder for fremtidig forskning.

Kapitel 3 og 4 omhandler tilpasning af flertilstandsreserveringsmodeller til situationer, hvor betalingerne og informationen muligvis ikke er fuldstændig ajour. I Kapitel 3 betragter vi en generel formulering, hvor betalingerne i realtid fremkommer som løbende betalinger baseret på den tilgængelige information samt tilbagebetalinger baseret på nytilkommen information. I denne kontekst er det muligt at relatere nutidsværdien til den klassiske nutidsværdi baseret på de kontraktuelle betalinger samt at karakterisere dynamikken af reserven. Det er dog generelt ikke muligt at relatere reserven til den klassiske flertilstandsreserve, og det må derfor undersøges i konkrete modeller. Kapitel 4 foreslår en model for tab-af-erhvervsevne ordninger og indfører relevante betingede uafhængighedsantagelser for at opnå eksplicitte udtryk for reserverne som naturlige modifikationer af de klassiske flertilstandsreserver. Den potentielle økonomiske effekt illustreres ved at anvende metoderne på et nyt datasæt LEC-DK19, som er baseret på rigtige data, der er blevet anonymiseret og ændret en anelse. Estimationsproceduren udviklet i Kapitel 5 bruges til at operationalisere modellen.

I Kapitel 5 er fokus på estimation af flertilstandsmodeller, der er påvirket af forskellige former for ufuldstændighed, herunder anmeldelsesforsinkelser og ufuldstændige afgørelser om hændelser. En estimationsprocedure foreslås, som tager højde for anmeldelsesforsinkelser via udtynding og ufuldstændige afgørelser via imputering, hvilket gør effektiv brug af de tilgængelige data. De asymptotiske egenskaber for estimatoren etableres baseret på M-estimationsteori. Vi demonstrerer tilgangen på simulerede data og på LEC-DK19.

Det sidste kapitel, Kapitel 6, omhandler efficient og robust statistisk inferens, når der er højrecensurering. Vi foreslår en fleksibel ikke-parametrisk estimationsprocedure baseret på pseudo-værdier og cross-fitting, som tillader brugen af prædiktionsmetoder fra maskinlæring. De asymptotiske egenskaber for proceduren etableres, hvilket viser at tilgangen er dobbelt-robust med hensyn til hjælpeparametre, der er nødvendige for udregning af pseudo-værdierne. Et simulationsstudie udforsker, hvordan tilgangen præsterer. Endeligt anvendes estimationsproceduren til at udføre et regressionsdiskontinuitetsdesign ved brug af rigtige data.

Contents

Chapter 1

Introduction

The purpose of this thesis is to develop techniques for reserving and estimation motivated by challenges encountered in the actuarial modeling of disability insurance policies. The introduction provides some background for the research, highlighting concepts and results that are central to the investigations of the thesis in order to set the stage for the subsequent chapters. Additionally, an overview of the main contributions is presented.

1.1 Background

1.1.1 Classic multistate modeling

Life insurance policies are contractual agreements on the exchange of future premiums and benefits contingent on certain events. These events are typically related to the biometric state of the insured. Simple examples of policies could be *2 million DKK is paid to those left behind in case the insured dies before age 67* or *30,000 DKK is paid to the insured every month if and when their earning capacity is reduced by at least 50%.* The insurance company is mandated by regulation to determine and allocate the funds necessary to meet its obligations to the policyholders. This is referred to as *reserving.* Accurate reserves are important for charging a suitable premium, monitoring the portfolio gains and losses, and performing other risk management operations needed to maintain solvency.

Reserving

Since money can be invested, the insurance company generally needs to set aside less than 1 DKK today to cover a payment of 1 DKK in the future. More formally, one monetary unit invested in a savings account at time 0 evolves to $\kappa(t)$ by time t where κ is called an *accumulation function.* Basic economic arguments then imply that one monetary unit at time t evolves to $\kappa(s)/\kappa(t)$ at time s for $s \geq t$. Let $B(t)$ be the accumulated cash flow, defined as the benefits less premiums, at time t. The *present value* of the future cash flow when standing at time t is then

$$P(t) = \int_{(t,\infty)} \frac{\kappa(t)}{\kappa(s)} \,\mathrm{d}B(s).$$

This is the amount that the insurer would need at time t to meet all future obligations. The present value is generally not known from the information available at time t, so it cannot be used as a reserve. Motivated by diversification of insurance risk and arbitrage-free pricing of financial risk, and assuming that the insurance and financial risks are independent, a best estimate reserve can be defined as

$$V(t) = \mathbb{E}_{\mathbb{P}\otimes\mathbb{Q}}\left[P(t) \mid \mathcal{F}_t\right],$$

where $\mathcal{F}_t$ is a σ-algebra representing financial and insurance information available at time t, $\mathbb{P}$ is the physical measure for insurance risk, and $\mathbb{Q}$ is a risk-neutral measure for financial risk, see Brennan & Schwartz (1979a,b), Møller & Steffensen (2007), and Article 77 in EIOPA (2009). Risk margins such as those defined by the Solvency II regulatory framework can also be incorporated but are outside the scope of this presentation.

The reserve thus depends on three components: the time value of money, the cash flow, and the information. For this thesis, the financial component only plays a minor role. A common and reasonable assumption that is employed throughout is that the cash flow does not depend on financial risk and that the financial and insurance risks are independent. In this case, the financial component can be dealt with separately. One may even assume that κ is deterministic since the extension to stochastic κ that is independent of the insurance risk is straightforward. To obtain explicit expressions for the reserve, one further needs to impose assumptions on the structure of the information and the cash flow. Multistate models provide a natural and parsimonious way to do so. For an overview of the multistate modeling literature, see the introduction in Furrer (2020).

Multistate modeling

Let the biometric state of the insured be represented by a stochastic process $Y = \{Y(t)\}_{t\geq 0}$ which is piecewise constant and takes values in a finite state space $\mathcal{J}$. A prominent example used to model disabilities is illustrated in Figure 1.1.

In the classic Markov multistate modeling setup, one assumes that Y is Markov, that $\mathcal{F}_t = \sigma\{Y(s) : s \leq t\}$, and that the cash flow takes the form

$$\mathrm{d}B(t) = \sum_{j\in\mathcal{J}} 1\{Y(t) = j\} b_j(t)\,\mathrm{d}t + \sum_{j,k\in\mathcal{J}:j\neq k} b_{jk}(t) N_{jk}(\mathrm{d}t),$$

where $N_{jk}(t)$ is the number of transitions from j to k until and including time t. One may interpret b_j as sojourn payment rates and b_{jk} as transition payments.

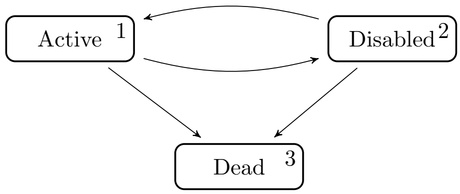

Figure 1.1: *Classic disability multistate model where $\mathcal{J} = \{1, 2, 3\}$. The arrows represent the possible transitions.*

When modeling disabilities, the Markov assumption is inadmissible since empirical evidence shows that the probability of reactivating or dying as disabled greatly depends on the duration spent in the disabled state. A common approach is then instead to use a semi-Markov setup. Here, it is assumed that (Y, U) is Markov, that the information is $\mathcal{F}_t = \sigma\{Y(s) : s \leq t\}$, and that the cash flow can be written as

$$\mathrm{d}B(t) = \sum_{j \in \mathcal{J}} 1\{Y(t) = j\} b_j\{t, U(t)\} \, \mathrm{d}t + \sum_{j,k \in \mathcal{J} : j \neq k} b_{jk}\{t, U(t-)\} N_{jk}(\mathrm{d}t),$$

where $U = \{U(t)\}_{t \geq 0}$ given by $U(t) = t - \sup\{s \in [0, t] : Y(s) \neq Y(t)\}$ is the duration since the last jump. Lump sum payments can also be incorporated but this requires slightly more involved notation, see Helwich (2008) and Adékambi & Christiansen (2017). Extensions of the model to include incidental policyholder behavior are also readily available, see Buchardt et al. (2015). In both the Markov and semi-Markov setup it is possible to choose a regular version of V that permits explicit integral expressions.

The introduction of a stochastic process Y generating the payments and information has several advantages. Firstly, one may incorporate structure about the payments that is a priori known from the insurance contract, avoiding the need to use data to rediscover this structure. Secondly, the strong intertemporal dependencies of running payments can be captured by modeling the dynamics of Y, where natural simplifying assumptions such as the Markov assumption are available. Lastly, one may exploit the powerful techniques that have been developed for actuarial and statistical modeling of multistate processes.

Information delays

Complexities arise in reserving when the process Y generating the contractual payments can no longer be observed in real-time. The motivation for studying this problem comes from disability insurance, where disabilities are affected by reporting delays and lengthy adjudications. In this case, it is not operational to choose $\mathcal{F}_t = \sigma\{Y(s) : s \leq t\}$ since this information is not always available at time t. Additionally, $B(t)$ might not have been paid out at time t so $P(t)$ is not always the relevant present value. The methods developed in Chapters 3 and 4 seek to retain the desirable properties of multistate modeling while accommodating delayed information about Y.

1.1.2 Marked point processes

In Section 1.1.1, the multistate model was represented with a piecewise constant stochastic process Y. An alternative representation that is used frequently throughout the thesis is that of a marked point process, which is often more convenient for mathematical derivations. If $(T_n)_{n\in\mathbb{N}}$ are the jump times of Y and $(Y_n)_{n\in\mathbb{N}}$ are the jump marks given by $Y_n = Y(T_n)$ then $(T_n, Y_n)_{n\in\mathbb{N}}$ is a marked point process. The marked point process and the initial state $Y(0)$ together generate the same information as the stochastic process Y, but the marked point process has the technical advantage that the index set is countable while it is uncountable for Y. Among other things, this makes it simpler to characterize its distribution, which can be done by iteratively specifying the conditional distribution of the jump times and jump marks. Marked point processes, despite their generality, admit explicit representations for key objects such as the compensator, likelihood, and infinitesimal generator.

Marked point processes also produce a comprehensive class of finite variation stochastic processes called *piecewise deterministic processes* which consists of all processes that at time t may be written as a function of t and the jump times and jump marks that have occurred before time t. Examples of piecewise deterministic processes are multistate models Y but also for example (Y, U). Properties of piecewise deterministic processes follow from properties of the underlying marked point process, so it is possible to take a marked point process as the starting point and construct all relevant objects as mappings of the marked point process.

Distribution

Let $H_n = (T_1, Y_1, \ldots, T_n, Y_n)$ be the first n jump times and jump marks, and assume the initial state is some deterministic value y_0 that is henceforth suppressed in the notation. The Ionescu-Tulcea theorem implies that the distribution of a marked point process can be specified by a sequence of Markov kernels which represent the regular conditional distributions of $T_{n+1} \mid H_n = h_n$ and $Y_{n+1} \mid H_n = h_n, T_{n+1} = t$ for $n \geq 0$; we denote these $P^{(n)}_{h_n}$ and $\pi^{(n)}_{h_n,t}$, respectively. Here we use the convention that H_0 generates the trivial σ-algebra and write $h_n = (t_1, y_1, \ldots, t_n, y_n)$ for a generic outcome of H_n. Regular conditional distributions are convenient to work with when one conditions on events of probability zero such as $(Y, U) = (y, u)$ for a fixed duration $u \geq 0$. In Chapter 3, they are used to obtain an expression for the dynamics of the so-called transaction time reserve resembling Thiele's differential equation. Furthermore, in Chapter 4 they are used to obtain explicit expressions for the proposed reserves as natural modifications of the classic multistate reserves described in Section 1.1.1.

In addition to information about the multistate model, one usually also has access to information about external covariates, and it is often of interest to quantify the effect of such covariates on the distribution of the multistate model. This can be incorporated into the setup without much difficulty. Baseline covariates, which are covariates that are registered at the start of an observation window and kept fixed hereafter, can be incorporated by specifying an extended initial state (y_0, w) for some outcome w of a stochastic covariate W. Each value of w then leads to a different collection of jump time and jump mark distributions. For time-dependent covariates, one may extend the marked point process such that the jump times and jump marks also include changes to the covariates.

Transition hazards

When the jump time distributions have density with respect to the Lesbesgue measure, except for possibly a point mass at ∞, they may be parameterized via their hazard functions. Denote by $\mu_{y_n\bullet}(s; h_n)$ the hazard function of $P^{(n)}_{h_n}$ which means

$$P^{(n)}_{h_n}((t,\infty]) = \exp\left\{ - \int_{(t_n,t]} \mu_{y_n\bullet}(s; h_n)\,\mathrm{d}s \right\}$$

for $t \geq t_n$. If the jump marks belong to a countable space $\mathcal{J}$ then one may further parameterize the mark distributions as

$$\pi^{(n)}_{h_n,t}(\{j\}) = \frac{\mu_{y_n j}(t; h_n)}{\mu_{y_n\bullet}(t; h_n)}$$

for $j \in \mathcal{J}$, where μ_{jk} are non-negative functions that satisfy $\sum_{k:k\neq j} \mu_{jk}(t;h) = \mu_{j\bullet}(t;h)$. Conversely, any set of non-negative functions μ_{jk} that are Lebesgue-integrable on compact subsets characterize the distribution of a marked point process via the above construction. The functions μ_{jk} are commonly referred to as *transition hazards*, *conditional hazards*, or simply *hazards*. Hazards are intimately connected with intensity processes; Theorem 4.5.2 of Jacobsen (2006) implies that for a marked point process constructed as above, the intensity of $N_{jk}(t) = \sum_{n=1}^{\infty} 1\{T_n \leq t\} 1\{Y_{n-1} = j, Y_n = k\}$ is $1\{Y_n = j\}\mu_{jk}(t; H_n)$ for $t \in (T_n, T_{n+1}]$. It is worth noting that one may generalize the above construction to allow for point masses in the jump time distribution by working with hazard measures instead of hazard functions. In this case, one recovers the distribution function via a product integral, see Theorem 4.1.1. in Jacobsen (2006).

Transition hazards provide a convenient parameterization for the jump time and jump mark distributions when they exist since the only restrictions on the hazards are local integrability and non-negativity. Realistic choices for the hazards generally satisfy local integrability, while non-negativity can be ensured by estimating $\log \mu_{ij}$

and subsequently applying the exponential function. This makes log hazards the natural estimands. In reserving, hazards also naturally arise via the compensator of the transition payments. Furthermore, reserves may be written as an integral over transition probabilities, and these probabilities may be computed efficiently in the semi-Markov case when hazards are available by solving differential equations, see Buchardt et al. (2015).

Likelihood and estimation

For real-life problems, the underlying distribution is not known and has to be inferred from data. The likelihood is a central object for many learning algorithms and it has a special structure for marked point processes. Let $\mathcal{F}_t = \sigma\{(T_m, Y_m) : T_m \leq t\}$ be the filtration generated by the marked point process. For elements $P, \tilde{P} \in \mathcal{P}$ of a statistical model $\mathcal{P}$, let P_t and $\tilde{P}_t$ be the restriction to $\mathcal{F}_t$ and assume local absolute continuity in the sense that $P_t \ll \tilde{P}_t$ for all $t \geq 0$. The *likelihood process* $\mathcal{L} = (\mathcal{L}_t)_{t\geq 0}$ is then defined as the process of Radon-Nikodym derivatives

$$\mathcal{L}_t = \frac{\mathrm{d}P_t}{\mathrm{d}\tilde{P}_t}.$$

If the statistical model is on the form $\mathcal{P} = \{P_\theta : \theta \in \Theta\}$ for $\Theta \subseteq \mathbb{R}^k$ then the model is said to be *parametric* otherwise it is said to be *nonparametric*. Classic statistical theory motivates estimators to be constructed as the argmax of the average observed log-likelihood where the argmax is over the set $\mathcal{P}$.

Let $P^{(n)}_{h_n}$ and $\pi^{(n)}_{h_n,t}$ be the conditional jump time and jump mark distributions for P and likewise use $\tilde{P}^{(n)}_{h_n}$ and $\tilde{\pi}^{(n)}_{h_n,t}$ for those generating $\tilde{P}$. Assume that the Radon-Nikodym derivatives $\mathrm{d}P^{(n)}_{h_n} / \mathrm{d}\tilde{P}^{(n)}_{h_n}$ and $\mathrm{d}\pi^{(n)}_{h_n,t} / \mathrm{d}\tilde{\pi}^{(n)}_{h_n,t}$ exist. Theorem 5.1.1 of Jacobsen (2006) then implies $P_t \ll \tilde{P}_t$ for all $t \geq 0$ and

$$\mathcal{L}_t = \left\{ \prod_{n=0}^{N(t)-1} \frac{\mathrm{d}P^{(n)}_{H_n}}{\mathrm{d}\tilde{P}^{(n)}_{H_n}}(T_{n+1}) \frac{\mathrm{d}\pi^{(n)}_{H_n,T_{n+1}}}{\mathrm{d}\tilde{\pi}^{(n)}_{H_n,T_{n+1}}}(Y_{n+1}) \right\} \frac{P^{(N(t))}_{H_{N(t)}}((t,\infty])}{\tilde{P}^{(N(t))}_{H_{N(t)}}((t,\infty])}$$

almost surely simultaneously for all $t \geq 0$. Here, $N(t) = \sum_{n=1}^{\infty} 1\{T_n \leq t\}$ is the number of jumps that have occurred by time t. The likelihood can alternatively be written in terms of compensating measures by using the product integral, see Section II.7.3 in Andersen et al. (1993).

When transition hazards exist, straightforward calculations give that the log-likelihood is

$$\log \mathcal{L}_t = \sum_{j,k \in \mathcal{J} : j \neq k} \int_0^t \log \mu_{jk}(s; H_{N(s-)}) \,\mathrm{d}N_{jk}(s) - \int_0^t 1\{Y(s) = j\} \mu_{jk}(s; H_{N(s)}) \,\mathrm{d}s$$

up to an additive constant. By discretizing the integrals, the optimization problem of computing the argmax becomes equivalent to the one for Poisson regression as noted in Friedman (1982) and Lindsey (1995). The form and properties of the above expression are largely unchanged by the presence of covariate information and various forms of missingness, including independent left-truncation and independent right-censoring. In such cases, the expression is referred to as a *partial likelihood* to distinguish it from the full likelihood which also contains the distribution of, for example, the missingness mechanism, see Section II.7.3 in Andersen et al. (1993). The partial likelihood is central to the estimation procedure developed in Chapter 5.

1.1.3 Causal inference

In addition to modeling the expected future losses, insurance companies may seek to mitigate some of their losses by improving the underlying risks. This has become increasingly topical in Denmark due to the sharp increase in the number of mental health-related disability claims. For disability insurance, mitigating actions may be taken to prevent disabilities or to reduce the duration or size of benefit payouts with return-to-work initiatives.

It would be highly advantageous for insurers to be able to quantify the effectiveness of their prevention and treatment initiatives. This would enable them to identify where their efforts will have the greatest impact and how many resources should be allocated to the initiatives, taking both costs and benefits into consideration. These initiatives are analogous to treatments administered in clinical trials and other types of interventional studies so the methodology developed there provides a natural starting point for assessing the impact. The gold standard for inferring a treatment effect is a randomized controlled trial. Insurers may however have financial or ethical objections to randomizing the use of their initiatives, in which case the treatment effect has to be inferred from observational data. This leads to hard statistical and identifiability problems. A fictional example is presented to help illustrate these challenges.

Example 1.1.1. (Max attempts impact evaluation.)
Max Risk is an actuary that has been tasked with assessing the impact of a recent prevention initiative. For every insured in the portfolio, Max has data on the form (Y, A, W) where $A \in \{0, 1\}$ is an indicator of whether they received the prevention initiative, $Y \in \{0, 1\}$ is an indicator of whether they reported a disability claim within the coverage period, and $W \in \mathbb{R}^p$ is a high-dimensional baseline covariate including age, gender, salary, marital status, address, etc.

Max is proficient with generalized linear models and hence posits a logistic regression model for the outcome as follows

$$Y \mid W, A \sim \mathrm{Bin}\{1, p(W, A)\}, \quad \operatorname{logit} p(W, A) = \alpha + \beta A + \gamma^T W.$$

They fit the model using their favorite statistical software and report the estimated coefficient $\hat{\beta}$ and a confidence interval as a summary of the treatment effect. The estimate is large and significant leading Max to conclude that the initiative was a success.

The Chief Actuary is skeptical of Max's analysis: What if the treatment effect depends on W? If so, the model is wrong; can the results be trusted? Since the dimension of W is large, it is difficult to manually check for interactions, so Max instead finds an old book on survey sampling and performs an alternative analysis based on the inverse probability of treatment weights. Max posits a logistic model for the treatment assignment as follows

$$A \mid W \sim \mathrm{Bin}\{1, p(W)\}, \quad \mathrm{logit}\, p(W) = \eta + \theta^T W$$

and estimates the average effect of the initiative as

$$\hat{\tau} = \frac{1}{n}\sum_{i=1}^{n} Y_i \frac{1\{A_i = 1\}}{\hat{p}(W_i)} - \frac{1}{n}\sum_{i=1}^{n} Y_i \frac{1\{A_i = 0\}}{1 - \hat{p}(W_i)}$$

where n is the number of insured in the sample. The estimator $\hat{\tau}$ converges to $\mathbb{E}[\mathbb{E}[Y \mid W, A = 1] - \mathbb{E}[Y \mid W, A = 0]]$ if $\hat{p}$ is uniformly consistent and the true function p is bounded away from 0 and 1. The estimated effect differs somewhat from the corresponding estimate derived from the logistic model for $Y \mid A, W$, namely

$$\tilde{\tau} = \frac{1}{n}\sum_{i=1}^{n} \hat{\mathbb{E}}[Y \mid W = W_i, A = 1] - \hat{\mathbb{E}}[Y \mid W = W_i, A = 0],$$

but it remains large and positive. Satisfied, the Chief Actuary recommends that the insurer should expand the use of the prevention initiative starting next year.

Sometime during next year, Max is standing at the water cooler talking to the person in charge of the prevention initiative. Max discovers that another factor that played a role in the treatment assignment A was how motivated the insured sounded when they were contacted. This information was however not registered in the insurer's database and was hence not included in the covariate W. This causes Max to worry about the validity of their analysis; if those who appeared motivated were also less likely to become disabled then they would perform better than the group not receiving treatment even if the initiative had no effect. Max hopes that this has not biased the analysis too severely. ◦

The scenario from Example 1.1.1 contains some common pitfalls associated with performing causal inference. The situation at the end of the example is analogous to the famous Simpson's paradox, highlighting that identifiability of the treatment effect requires additional knowledge about the data-generating mechanism beyond what can be learned from the available data. This is because unmeasured confounding variables may invalidate the causal interpretation of the results. This insight naturally leads to the question:

> *When is the treatment effect identifiable from observational data?*

Answering this question requires causal models. The requirements for identifiability of the treatment effect are strong, so in practice, if one deems it important to be able to quantify the effects, one should design the initiatives in such a way that the effect becomes identifiable.

Once identifiability holds, the causal effect is some function of the distribution of the observational data, and it can therefore be estimated using purely statistical methods. For example, under some conditions, the average treatment effect τ can be identified as $\tau = \mathbb{E}[\mathbb{E}[Y \mid W, A = 1] - \mathbb{E}[Y \mid W, A = 0]]$ which was the limit of $\hat{\tau}$ in Example 1.1.1. The way the estimand depends on the data distribution is however usually complicated and inference based on parametric maximum likelihood methods is regarded as problematic since they often perform poorly for high-dimensional covariates and since confidence bands and parameter values can be highly sensitive to model misspecification. This leads to the question:

> *How can one perform inference for the estimands coming from causal models in an efficient way without imposing strong distributional assumptions?*

The modern approach relies on semi-parametric efficiency theory with two popular frameworks being targeted maximum likelihood estimation and double machine learning as detailed in Van der Laan & Rose (2011) and Chernozhukov et al. (2018), respectively.

In summary, causal inference can be regarded as a two-step procedure. The initial step does not rely on data; rather, it involves leveraging subject-matter expertise to propose a causal model. Such a model depends on assumptions that are untestable from the observational data. The goal is to examine whether the causal effect of interest is identifiable from the distribution of the observational data under assumptions that are deemed reasonable. The second step consists of using statistical techniques to obtain an estimate of the resulting estimand. These two aspects are discussed in more detail below.

Identifiability

There exist several frameworks for causal modeling, with the most popular being graphical models and the Neyman–Rubin model, see the classic references Rubin (1974), Pearl (2009), and Hernán and Robins (2020). The causal framework used in this thesis is the Neyman–Rubin model as it is more natural for describing time-varying outcomes. Graphical models provide a natural way to identify treatment effects when one has more detailed knowledge about the causal influences between the treatment, outcome, and covariates. This is, however, rarely useful for insurance data which has an infamously low signal-to-noise ratio. In this thesis, causal models are used to discuss impact evaluation for prevention initiatives in Chapter 2 and in the data application of Chapter 6.

The Neyman–Rubin model is based on the notion of potential outcomes. When the observed data is on the form (Y, A, W) for a binary treatment variable A as in Example 1.1.1, the Neyman–Rubin model posits the existence of potential outcomes $Y^{(1)}$ and $Y^{(0)}$ corresponding to the value that Y would take if A had been 1 and 0, respectively, implying $Y = Y^{(A)}$. The average treatment effect is then defined as $\tau = \mathbb{E}[Y^{(1)} - Y^{(0)}]$.

One way to make τ identifiable is to assume that the data arises as a randomized experiment for each fixed value of W. More formally, one assumes $(Y^{(1)}, Y^{(0)}) \perp\!\!\!\perp A \mid W$ and that $\mathbb{P}(A = 1 \mid W)$ is bounded away from 0 and 1. In Example 1.1.1, the final paragraph indicates that the former assumption might be problematic; those who were assigned treatment likely had a different distribution of $(Y^{(1)}, Y^{(0)})$ compared to those that did not receive treatment, even for similar values of W. When the assumptions hold, it is straightforward to show that $\tau = \mathbb{E}[\mathbb{E}[Y \mid W, A = 1] - \mathbb{E}[Y \mid W, A = 0]]$. In the data application of Chapter 6, identifiability is instead achieved via a regression discontinuity design which assumes that the data was generated as a local randomized controlled trial along one of the coordinates of W. Both approaches boil down to having strata where people are exchangeable except that some receive treatment and others do not.

Targeted estimation

Estimation procedures for causal estimands typically rely on the ability to estimate related distributional quantities referred to as *nuisance parameters*. For example, as seen in Example 1.1.1, the estimators $\hat{\tau}$ and $\tilde{\tau}$ of τ depend on estimators of $\mathbb{E}[A \mid W]$ and $\mathbb{E}[Y \mid W, A]$, respectively. Estimating these conditional expectations can be done using various regression methods including machine learning algorithms which do not require strong distributional assumptions, can deal with high-dimensional covariates, and routinely outperform classical regression methods in terms of predictive accuracy. However, the regression estimators are tuned to have optimal performance for their respective prediction tasks which may lead to suboptimal performance for the estimand of interest. Influence functions may be used to construct estimators that remove this plug-in bias and achieve asymptotic efficiency in a local minimax sense as summarized in Kennedy (2022) and Hines et al. (2022).

An estimand can generally be represented as some functional $\psi : \mathcal{P} \mapsto \mathbb{R}$ for a statistical family $\mathcal{P}$. If this function is smooth, then it satisfies a distributional Taylor expansion

$$\psi(\hat{\mathbb{P}}) - \psi(\mathbb{P}) = -\int \varphi(x, \hat{\mathbb{P}}) \,\mathrm{d}\mathbb{P}(x) + R(\hat{\mathbb{P}}, \mathbb{P}), \tag{1.1.1}$$

where φ is the *efficient influence function* for the estimand ψ and R is a second-order remainder term. When the data consists of i.i.d. observations $X_1, \ldots, X_n$, this expansion suggests defining a debiased estimator $\hat{\psi}$ as $\hat{\psi} = \psi(\hat{\mathbb{P}}) + n^{-1} \sum_{i=1}^{n} \varphi(X_i, \hat{\mathbb{P}})$ referred to as the *one-step estimator*. There is a considerable body of theoretical and numerical work suggesting that debiased estimators are, in many scenarios, superior to plug-in estimators, see for example the references in Kennedy (2022). In fact, the debiased estimator is asymptotically efficient if the remainder R convergences sufficiently quickly to 0 and one uses cross-fitting such that $\hat{\mathbb{P}}$ and the average over φ uses different parts of the sample. Cross-fitting is not necessary if $\{\varphi(x, P) : P \in \mathcal{P}\}$ is a Donsker class.

Importantly, the remainder term R typically only involves second-order errors of the nuisance parameters. Since controlling R is the key to achieving asymptotic efficiency of the debiased estimator, this estimator inherits a robustness to misspecification in the nuisance parameters. For example, when the estimand ψ is the average treatment effect τ as in Example 1.1.1, one can show

$$
\begin{aligned}
R(\hat{\mathbb{P}}, \mathbb{P}) = \sum_{a=0}^{1} (-1)^{a+1} \int & \Big(\mathbb{P}(A = a \mid W = w) - \hat{\mathbb{P}}(A = a \mid W = w) \Big) \\
& \times \Big(\mathbb{E}[Y \mid W = w, A = a] - \hat{\mathbb{E}}[Y \mid W = w, A = a] \Big) \\
& \times \hat{\mathbb{P}}(A = a \mid W = w)^{-1} \, \mathbb{P}(W \in \mathrm{d}w)
\end{aligned}
$$

which by the triangle and Cauchy-Schwarz inequalities is bounded by a constant times the product of the error in the nuisance parameters. Furthermore, the one-step estimator is found to be

$$
\begin{aligned}
\hat{\psi} = \frac{1}{n} \sum_{i=1}^{n} & \hat{\mathbb{E}}[Y \mid W = W_i, A = 1] - \hat{\mathbb{E}}[Y \mid W = W_i, A = 0] \\
& + \frac{A_i - \hat{\mathbb{P}}(A = 1 \mid W = W_i)}{\hat{\mathbb{P}}(A = 1 \mid W = W_i)(1 - \hat{\mathbb{P}}(A = 1 \mid W = W_i))} (Y_i - \hat{\mathbb{E}}[Y \mid W = W_i, A = A_i]),
\end{aligned}
$$

which is the celebrated augmented inverse probability weighted estimator. As opposed to $\hat{\tau}$ and $\tilde{\tau}$ which are based on estimators of $\mathbb{E}[A \mid W]$ and $\mathbb{E}[Y \mid W, A]$, respectively, the one-step estimator combines these estimators to gain efficiency and robustness.

1.2 Contributions and overview

Here, we provide an overview of the thesis's scientific contributions. The mathematical style of this section is informal and the emphasis is on concepts and intuition. For a precise treatment, the reader is referred to the respective chapters. While each chapter constitutes a stand-alone scientific contribution, some similar themes permeate the thesis. In broad terms, the remaining chapters deal with ways of extending multistate methods to better meet the challenges faced in the actuarial modeling of disability insurance policies. Another shared theme is the effort to derive techniques that are rigorously grounded while remaining intuitive and of similar computational complexity to current methods. The papers can be read independently, but the recommended reading order follows the ordering of the chapters.

Chapter 2 contains a review article that complements the introduction in that it describes the complexities that disability insurance schemes give rise to and how the contributions of this thesis address some of these challenges. However, the discussion in Chapter 2 is broader and more oriented towards practice. The chapter is mainly concerned with illustrating and systematizing the contemporary Danish disability insurance landscape from an actuarial perspective to identify areas where the current literature is underdeveloped and where actuarial researchers and practitioners are well-positioned to contribute more and in turn meet the modern needs of life insurance companies.

Chapters 3 and 4 are concerned with stochastic modeling of reserves, while Chapters 5 and 6 focus on statistical aspects. The statistical papers differ significantly in both their approach and style, reflecting the differing demands of the pricing/reserving and impact evaluation tasks. By extension, this also reflects the different cultures of actuaries and statisticians. For pricing and reserving, in addition to being accurate at an individual and portfolio level, the estimates should be explainable to financial supervisory authorities, have limited computational complexity as reserves are updated frequently, be suitable for forecasting, and be simple and reliable to run in production. Altogether, this speaks for a parametric model. For impact evaluation, the primary goal is usually inference for a small selection of nonparametrically identifiable estimands, meaning that accurate quantification of the uncertainty in the estimates is key and forecasting is not necessary. Furthermore, since impact evaluations are typically updated less frequently than reserves and are usually not reported to supervisory authorities, the computational complexity of the methods can be higher. Altogether, this speaks for nonparametric methods.

1.2.1 Reserving with transaction time information

In Chapters 3 and 4, we propose ways to extend the multistate reserving approach to situations where information about the process driving the payments might not be up to date. Chapter 3 explores the general structure of the problem while Chapter 4 is concerned with obtaining explicit and operational expressions for disability insurance schemes.

Assume that the accumulated contractual payments $B = \{B(t)\}_{t\geq 0}$ are adapted to the information generated by some stochastic process $X = \{X(t)\}_{t\geq 0}$ and write

$B(t) = B(X,t)$ with a slight abuse of notation. If the insurer cannot observe X in real-time, for example, due to reporting delays, then $\mathcal{F}_t^X = \sigma\{X(s) : s \leq t\}$ is not available at time t and the insurer hence cannot ensure that $B(t)$ is paid out by time t. Instead, the insurer must through claim settlement processes retroactively ensure that the insured receives the benefits that they were eligible for. In practice, this is done via *backpay* which is a lump sum payment of overdue benefits that have been delayed by reporting delays and claims processing. As a consequence, the accumulated realized payments $\mathcal{B}$ may differ from B and the available information $\mathcal{F}^{\mathcal{Z}}$ may differ from $\mathcal{F}^X$.

These complications will lead to a mismatch in timing between releasing reserves and incoming losses if they are not accommodated in the actuarial reserving models. The approach in Chapters 3 and 4 is to introduce a reserve $\mathcal{V}$, which is based on $\mathcal{B}$ and $\mathcal{F}^{\mathcal{Z}}$, and to explore to which extent $\mathcal{V}$ can be linked to the classic reserve V which is based on B and $\mathcal{F}^X$. Thus, we define

$$P(t) = \int_{(t,\infty)} \frac{\kappa(t)}{\kappa(s)} \, \mathrm{d}B(s), \qquad \mathcal{P}(t) = \int_{(t,\infty)} \frac{\kappa(t)}{\kappa(s)} \, \mathrm{d}\mathcal{B}(s),$$
$$V(t) = \mathbb{E}[P(t) \mid \mathcal{F}_t^X], \qquad \mathcal{V}(t) = \mathbb{E}[\mathcal{P}(t) \mid \mathcal{F}_t^{\mathcal{Z}}],$$

for a deterministic accumulation function κ. We refer to the elements of $\mathcal{V}$ as the transaction time model and the elements of V as the valid time model.

To have any hope of linking $\mathcal{V}$ and V, one must impose some additional relations between the payments and/or the filtrations. The work in this thesis assumes that $\mathcal{B}$ satisfies

$$\int_{[0,t]} \frac{\kappa(t)}{\kappa(s)} \mathcal{B}(\mathrm{d}s) = \int_{[0,t]} \frac{\kappa(t)}{\kappa(s)} B(X^t, \mathrm{d}s) \tag{1.2.1}$$

where $X^t = \{X^t(s)\}_{s \geq 0}$ is a stochastic process corresponding to X, but based on the available transaction time information at time t, with the additional property that it is piecewise constant as a function of t. Intuitively, Equation (1.2.1) says that the transaction time cash flow is constructed such that the present value of the past payments equals the present value of the cash flow that would have arisen had payments been based on X^t which is the currently accepted version of X. A representation of $\mathcal{B}(\mathrm{d}t)$ is given in Section 4 of Chapter 3, showing that it consists of running payments based on X^t as well as backpay, properly accumulated with interest, whenever X^t and X^{t-} differ.

It is assumed that X^t converges to X as t goes to $+\infty$, establishing a link between the valid time and transaction time models. The main result of Chapter 3 is that

$$\mathcal{P}(t) = P(t) + \int_{[0,t]} \frac{\kappa(t)}{\kappa(s)} (B - \mathcal{B})(\mathrm{d}s). \tag{1.2.2}$$

Thus, the transaction time present value equals the valid time present value plus a term that corrects for erroneous payments before the current time. This is a simple relation, but it does not generally carry over to a simple relation between the reserves. Substituting Equation (1.2.2) in the definition of $\mathcal{V}$, it becomes clear that one should understand the distribution of X given $\mathcal{F}_t^{\mathcal{Z}}$ in order to obtain explicit expressions for $\mathcal{V}(t)$. There seems to be no canonical assumption so this has to be investigated in concrete models as in Chapter 4.

One could alternatively have started with Equation (1.2.2) as a definition of $\mathcal{P}$, but some properties come more naturally from the original definition. This is the case for the derivation of the dynamics of $\mathcal{V}$, which constitutes another important result of Chapter 3. Assuming that $\mathcal{F}_t^{\mathcal{Z}} = \sigma\{\mathcal{Z}(s) : s \leq t\}$ for a piecewise deterministic process $\mathcal{Z}$, one obtains the representation

$$\mathcal{V}(\mathrm{d}t) = \mathcal{V}(t-)\frac{\kappa(\mathrm{d}t)}{\kappa(t-)} - \mathcal{B}(\mathrm{d}t) + \int \mathcal{R}(t,\zeta)\, M(\mathrm{d}t, \mathrm{d}\zeta), \tag{1.2.3}$$

resembling a stochastic Thiele equation, compare with Norberg (1992) and Christiansen & Furrer (2021). This says that the transaction time reserve evolves due to interest rate gains, realized payments, and a term related to the sums at risk $\mathcal{R}$. The final term may be used for model validation by exploiting that it is a martingale if the model is correctly specified.

In Chapter 4, the approach is to assume a valid time semi-Markov setup and specify state spaces for X and $\mathcal{Z}$ in a manner which is sufficiently general to capture many disability insurance schemes encountered in practice. In order to make the distribution of X given $\mathcal{F}_t^{\mathcal{Z}}$ tractable, and to ultimately link $\mathcal{V}$ to V, various conditional independence assumptions are imposed, specifying what valid time information is sufficient for the transaction time information not to provide any information about the future trajectory of X. The primary technical contribution of the paper is a lemma which shows that a strong Markov property for a non-stopping time G can be obtained via relevant independence assumptions, roughly stating that if X is Markov and $\{X(s)\}_{s\geq t} \perp\!\!\!\perp \{X(s)\}_{s\leq t}, G \mid X(t)$ on the event $(G \leq t)$ for any $t \geq 0$, then the distribution of $\{X(s)\}_{s\geq G}$ given $\{X(s)\}_{s\leq G}$ equals the distribution of $\{X(s)\}_{s\geq g}$ given $\{X(s)\}_{s\leq g}$ evaluated in $g = G$.

Based on these assumptions, explicit expressions of $\mathcal{V}$ are derived, which is the main contribution of Chapter 4. For ease of exposition, we drop some terms which are of secondary importance, but the exact results may be found in the aforementioned chapter. The covered-but-not-reported reserve is approximately given by

$$\mathcal{V}(t) = V_a(t,t) + \int_{(0,t]} V_i(s,0)\mathbb{P}(\text{Reporting delay} > t-s)\mu_{ai}(s)\,\mathrm{d}s, \tag{1.2.4}$$

where $V_j(t,u)$ is the valid time reserve for state j, time t, and duration u, and μ_{ai} is the valid time disability hazard. Let $G(t)$ be the time from which the insured is eligible for disability benefits if the claim is (re)awarded when standing at time t and let $W(t)$ be the disability duration at time $G(t)$. The reported-but-not-paid reserve is then approximately

$$\mathcal{V}(t) = \mathbb{P}(\text{Claim is awarded} \mid \mathcal{F}_t^{\mathcal{Z}}) V_i(G(t), 0), \tag{1.2.5}$$

the paid-but-not-settled reserve for an insured currently receiving disability benefits is

$$\mathcal{V}(t) = V_i(t, W(t)), \tag{1.2.6}$$

and the paid-but-not-settled reserve for those not currently receiving benefits is approximately

$$\mathcal{V}(t) = \mathbb{P}(\text{Claim is reawarded} \mid \mathcal{G}_t) V_i(G(t), W(t)). \tag{1.2.7}$$

Defining $\mathcal{V}$ as a conditional expectation and deriving explicit expressions for the different stages of claim settlement, instead of specifying reserves for each stage separately, ensures that the reserves sum to a consistent portfolio reserve and that the reserves satisfy the dynamics from Equation (1.2.3). For estimation of the model constituents, it is argued that the statistical problem can be cast as a special case of the one investigated in Chapter 5, so the estimation procedure and asymptotic results carry over.

1.2.2 Estimation with transaction time information

In Chapter 5, we study the problem of parametric estimation and inference for multistate models where the observational scheme is affected by left-truncation, right-censoring, reporting delays, and incomplete event adjudication. We propose an estimation procedure that makes effective use of all the available data and derive large sample properties of the estimator.

Left-truncation and right-censoring together mean that the subject might only be in the sample during some sub-interval of the full observation window. Reporting delays have the effect that events that are not reported before the time of analysis are not part of the sample. We assume that events are reported in the order in which they occurred. This assumption is natural for individual modeling and is essential for the tractability of our approach. Finally, incomplete event adjudication has the effect that some reported events are removed from the sample after the time of analysis as they do not satisfy the criteria for true events. Differently from Chapter 4, no assumptions are imposed regarding how the hazards would be affected by information about reporting delays and adjudication processes.

Let θ be a parameter that characterizes the distribution of the multistate model. An estimation procedure for θ is derived based on the following line of reasoning.

If the multistate model was only affected by left-truncation and right-censoring and one had access to n i.i.d. observations, then an estimator could be defined as $\arg\max_\theta n^{-1}\sum_{i=1}^n \log \mathcal{L}_i(\theta)$ where $\mathcal{L}_i$ is the partial likelihood for subject i. This follows because $\mathbb{E}[\log \mathcal{L}(\theta)]$ is uniquely maximized in the true parameter θ_0 under weak assumptions and the average converges to the expectation by the law of large numbers. If the multistate model is also affected by reporting delays, then it is still possible to obtain an explicit formula for the partial log-likelihood $\mathcal{L}(\theta, f)$ which now also depends on f that parameterizes the distribution of the reporting delays. This is an important technical result for the approach in the paper. Since $\mathbb{E}[\mathcal{L}(\theta, f)]$ is uniquely maximized in the true parameter (θ_0, f_0) under weak assumptions, an estimator of θ could be obtained by $\arg\max_\theta n^{-1}\sum_{i=1}^n \mathcal{L}_i(\theta, \hat{f}_n)$ for an estimator $\hat{f}_n$ of f_0. To accommodate incomplete event adjudication, one may define $\mathbb{E}_g[\mathcal{L}(\theta, f) \mid Z]$, where Z generates the information available at the time of analysis and g parameterizes the distribution of the adjudication outcomes. Since

$$\mathbb{E}[\mathbb{E}_{g_0}[\mathcal{L}(\theta, f_0) \mid Z]] = \mathbb{E}[\mathcal{L}(\theta, f_0)]$$

and the right-hand-side is uniquely maximized in θ_0, one may define

$$\hat{\theta}_n = \arg\max_\theta \mathbb{E}_{\hat{g}_n}[\mathcal{L}(\theta, \hat{f}_n) \mid Z]. \tag{1.2.8}$$

The estimator $\hat{g}_n$ of g_0 can be based on classic event history analysis methods. Subsequently, one can let $\hat{f}_n = \arg\max_f \mathbb{E}_{\hat{g}_n}[\log \mathcal{L}(f) \mid Z]$ for a right-truncated log-likelihood $\log \mathcal{L}(f)$.

In total, the estimation procedure produces an estimator $(\hat{g}_n, \hat{f}_n, \hat{\theta}_n)$ of (g_0, f_0, θ_0). The formulation of this estimation procedure is the main conceptual contribution of the paper. Another contribution is the introduction of an approximation which reduces the computational complexity greatly and makes it so the estimation procedure can be cast in terms of simple adjustments of the observed exposures and occurrences.

One may recognize the estimation procedure as a so-called two-step M-estimation procedure and properties of the estimator may hence be derived using well-known methods. Under classic conditions, we show for $n \to \infty$ that

$$(\hat{g}_n, \hat{f}_n, \hat{\theta}_n) \to (g_0, f_0, \theta_0)$$

in probability and

$$n^{1/2}(\hat{\theta}_n - \theta_0) \to N(0, V)$$

in distribution for a variance matrix V, and that the estimators may be bootstrapped. This is the main result of Chapter 5.

1.2.3 Efficient inference under censoring and causal effects

In Chapter 6, we study the problem of obtaining efficient and robust statistical estimates in the presence of right-censoring. We propose a nonparametric estimation procedure that is model-agnostic such that machine learning estimators may be used to estimate relevant nuisance parameters. Subsequently, we establish large sample properties of the procedure for the purposes of inference. Based on these results, we extend the applicability of regression discontinuity designs to general right-censored data structures.

Let $X = \{X(t)\}_{t\geq 0}$ be a stochastic process, let $X^u = \{X(t \wedge u)\}_{t\geq 0}$ be the process stopped at u, and let (C, X^C) be the observed data for the right-censoring time C. The relevant outcome is $Y(X)$ and the estimand of interest is $\mathbb{E}[Y(X) \mid W]$ for baseline covariates W. There is a rich literature on efficient regression estimators when i.i.d. data on the form $(Y_1, W_1), \dots, (Y_n, W_n)$ is available. Instead of tailoring these methods to censored data, our approach is to transform the observed data into some pseudo-outcomes Y^* to remove the effect of censoring and then use methods designed for uncensored data.

For identifiability, let the time of analysis be η and assume $Y(X)$ only depends on X^η. Further, assume the censoring satisfies coarsening at random so that the density of $\mathbb{P}(C \in \mathrm{d}u \mid X)$ with respect to a fixed reference measure μ only depends on X^u. Finally, assume $\mathbb{P}(C \geq \eta \mid X)$ is uniformly bounded away from 0. With two candidate probability measures $\mathbb{P}_1$ and $\mathbb{P}_2$, define the pseudo-outcomes as

$$\begin{aligned} Y^*_{\mathbb{P}_1,\mathbb{P}_2}(C, X^C) = {} & \frac{Y(X)1_{(C\geq\eta)}}{\mathbb{P}_1(C \geq \eta \mid X)} \\ & + \int_{[0,\eta)} \frac{\mathbb{E}_2[Y(X) \mid X^u]}{\mathbb{P}_1(C > u \mid X)} \left\{ \mathrm{d}1_{(C\leq u)} - 1_{(C\geq u)} \frac{\mathbb{P}_1(C \in \mathrm{d}u \mid X)}{\mathbb{P}_1(C \geq u \mid X)} \right\}, \end{aligned}$$

where $\mathbb{E}_2$ is the expectation under $\mathbb{P}_2$. The motivation for this definition is that $n^{-1}\sum_{i=1}^n Y^*_{\hat{\mathbb{P}},\hat{\mathbb{P}}}(C_i, X_i^{C_i})$ is the one-step estimator for $\mathbb{E}[Y(X)]$. The estimand $\mathbb{E}[Y(X) \mid W]$ is not a smooth function of $\mathbb{P}$ and hence does not satisfy Equation (1.1.1), so it is not immediately clear from the general theory how to estimate it efficiently. Since the efficient estimator of $\mathbb{E}[Y(X)]$ is obtained by averaging these pseudo-values, the hope is that an efficient estimator of $\mathbb{E}[Y(X) \mid W]$ is obtained by regressing these pseudo-values on covariates.

A central result of the paper is the following product representation of the conditional bias,

$$\begin{aligned} & \mathbb{E}[Y^*_{\mathbb{P}_1,\mathbb{P}_2}(C, X^C) - Y(X) \mid W] \\ & = \mathbb{E}\Big[\int_{[0,\eta)} \{\mathbb{E}[Y(X) \mid X^u] - \mathbb{E}_2[Y(X) \mid X^u]\} \\ & \qquad \times \{\gamma_1(u \mid X) - \gamma(u \mid X)\} \frac{\mathbb{P}(C \geq u \mid X)}{\mathbb{P}_1(C > u \mid X)} \,\mathrm{d}\mu(u) \mid W\Big], \end{aligned} \tag{1.2.9}$$

where γ is the hazard of $C \mid X$ under $\mathbb{P}$ and γ_1 is the same but under $\mathbb{P}_1$. The conditional bias is important for deriving asymptotic results, and this product representation is used to show that the estimator is robust to deviations in the nuisance parameters $\mathbb{E}_2[Y(X) \mid X^u]$ and $\gamma_1(u \mid X)$.

Building on the approach from Kennedy (2023), we propose a cross-fitting estimation procedure and derive its large sample properties. We show under some regularity conditions that if the regression method $\hat{\mathbb{E}}_n$ converges at rate α to a Gaussian distribution, meaning that

$$n^{\alpha}\{\hat{\mathbb{E}}_n[Y^*_{\mathbb{P},\mathbb{P}}(C, X^C) \mid W = w] - \mathbb{E}[Y^*_{\mathbb{P},\mathbb{P}}(C, X^C) \mid W = w]\} \to N(\mu, \sigma^2)$$

in distribution for $n \to \infty$, then the cross-fitted estimator based on the estimated pseudo-outcomes $Y^*_{\hat{\mathbb{P}}_1,\hat{\mathbb{P}}_2}$ also converges to a Gaussian at rate α whenever

$$\|\gamma_1(u \mid X) - \gamma(u \mid X)\| \times \|\mathbb{E}_2[Y(X) \mid X^u] - \mathbb{E}[Y(X) \mid X^u]\| = O_{\mathbb{P}}(n^{-\beta})$$

for $\beta > \alpha$. Here $\|\cdot\|$ is a specific weighted $L_2(\mathbb{P})$-norm and stochastic boundedness $A_n = O_{\mathbb{P}}(a_n)$ roughly means that the probability of $|A_n/a_n|$ being large becomes small when n increases.

Finally, we extend regression discontinuity designs to our censored data setup. Assume that the covariate W is one-dimensional and define the conditional average treatment effect at w_0 as

$$\tau = \mathbb{E}[Y^{(1)} - Y^{(0)} \mid W = w_0]$$

where $Y^{(a)} = Y(X^{(a)})$ for the potential outcomes $X^{(a)}$. Inspired by the classic assumptions for the validity of regression discontinuity designs, assume

(i) $a^+ := \lim_{w \downarrow w_0} \mathbb{P}(A = 1 \mid W = w)$ and $a^- := \lim_{w \uparrow w_0} \mathbb{P}(A = 1 \mid W = w)$ exist and $a^+ \neq a^-$,

(ii) $\mathbb{E}[Y^{(1)} \mid W = w]$ and $\mathbb{E}[Y^{(0)} \mid W = w]$ are continuous in w at w_0,

(iii) $A \perp\!\!\!\perp (Y^{(1)} - Y^{(0)}) \mid W = w$ in the limit for $w \to w_0$.

One then obtains the identification

$$\tau = \frac{y^+ - y^-}{a^+ - a^-}$$

and the estimation procedure of this chapter may be used to estimate each of y^+, y^-, a^+, and a^-. Furthermore, the asymptotic distribution of the estimand follows as a special case of our general results.

Chapter 2

Loss of earning capacity in Denmark – an actuarial perspective

This chapter contains the manuscript *Furrer & Sandqvist (2025)*.

Abstract

We describe challenges and opportunities related to risk assessment and mitigation for loss of earning capacity insurance with a special focus on Denmark. The presence of public benefits, claim settlement processes, and prevention initiatives introduces significant intricacy to the risk landscape. Accommodating this requires the development of innovative approaches from researchers and practitioners alike. Actuaries are uniquely positioned to lead the way, leveraging their domain knowledge and mathematical-statistical expertise to develop equitable, data-driven solutions that mitigate risk and enhance societal well-being.

Keywords: Disability insurance; Mitigation; Multistate; Prevention; Public benefits

2.1 Introduction

The purpose of insurance is to protect against uncertain events that lead to financial losses, which is first and foremost achieved by the pooling of diversifiable risks. Disability benefits play a crucial role in ensuring income stability for individuals by reducing financial vulnerability during periods with reduced earning capacity as well as supporting part-time employment. The latter relates to the fundamental right to work. This right may be found in the Declaration of Philadelphia (1948)[1], in Article 23 of the Universal Declaration of Human Rights (1948), in Article 10 of the Valencia Declaration (1998)[2], and in Article 15 of the EU Charter of Fundamental

[1]Declaration concerning the aims and purposes of the International Labour Organisation.
[2]Declaration of Human Duties and Responsibilities.

Rights (2000). In a welfare state, the right to work is matched by a duty to work for those that are able – and support for the disabled and disadvantaged. Such is the situation in Denmark, where a well-rounded three-pillar system consisting of the public sector, labour market pensions, and private insurers offers a rich interplay between public benefits and insurance coverage (Kommissionen om tilbagetrækning og nedslidning, 2022, Chapter 6). The insurance coverages play a vital role in today's society, with approximately one in five Danes expected to rely on disability insurance at some point during their careers[3].

For classic risks such as mortality and longevity, a wide range of proven modeling approaches are available to the actuary. In comparison, disability insurance presents unique challenges. Disabilities can very in form, degree, and duration, and in the interplay with the public system, leading to complications in risk assessment and management. There is untapped potential in this field – and this potential can be harnessed by actuaries due to their unique combination of domain knowledge and mathematical as well as statistical competence.

We attribute the relatively unrealized potential to historical reasons first and foremost. Technological progress has just recently allowed the storage and processing of much larger quantities of data than in the past, which combined with advances in data analytics, machine learning, and task automation has greatly expanded the actuaries' and statisticians' sphere of influence. This also leads to increased competition in the market, both in relation to product design, pricing, and value creation, which puts additional pressure on actuaries to fulfill the potential.

To keep the presentation concise and closely aligned with our practical experience, we focus on the Nordic countries, especially Denmark. However, we expect the discussions to also be relevant for actuaries in other regions. For a broad introduction to the German health insurance landscape and actuarial practices, confer with Milbrodt and Röhrs (2016).

The paper is structured as follows. In Section 2.2, we describe and characterize the risk environment for loss of earning capacity. Subsequently, Section 2.3 contains a formalization and comparison of existing insurance coverage designs with a focus on the Danish disability insurance market. Based hereon, the next two sections are devoted to the topics of risk assessment through prediction and risk mitigation through prevention, respectively. In Section 2.4, the adequacy of current approaches to stochastic modeling, pricing, and reserving of disability risk is assessed. Next, in Section 2.5, methods for impact evaluation of prevention initiatives are discussed. Throughout, we identify avenues for further research and improved actuarial practice. Closing remarks are provided in Section 2.6.

---

[3] `https://danicapension.dk/en/personal/your-insurances/all-insurance-options/loss-of-earning-capacity` (accessed 14/1/2025).

2.2 Risks and agents

This section describes the different risks and agents involved in disability insurance, highlighting the complex interplays that arise. The main purpose is to identify and provide an overview of operational areas for which actuarial expertise is crucial to ensure successful risk assessment and mitigation. For this reason, we mainly place ourselves in the position of the insurer and the actuary. However, some of the insights we provide could also be relevant in a broader context, for instance in the design and implementation of public welfare reforms.

2.2.1 Background

Denmark is, like its northern neighbors, a welfare state with an elaborate social safety net, universal healthcare, and collective bargaining (Pedersen & Kuhnle, 2017). The social safety net seeks to cultivate the economic and social well-being of citizens with inadequate or no income, not least due to temporary or permanent incapacity to work. The wide range of public benefit schemes reflects efforts to balance measures that facilitate strengthened workforce participation with citizen welfare. This is not to insinuate that these two objectives are necessarily opposed to each other; on the contrary, there is a rich interplay between them. According to Bratsberg et al. (2013), for example, in Norway job loss is a major cause of disability program entry. In Denmark, public benefit schemes include temporary sickness benefits, but also long-term benefits like so-called flexjob benefits as well as early retirement pension. The system is widely recognized to be complex, frequently attracts political attention, and is under continuous development (Ekspertgruppen for fremtidens beskæftigelsesindsats, 2024).

Despite the extensive social safety net that characterizes the Nordic model, some individuals who become partly or completely incapacitated to work experience a considerable loss of income and inability to sustain their pension contributions. This is especially the case for high earners, but is expected to become increasingly relevant to the wider population. In Denmark, the increased privatization of welfare benefits can be seen as part of a larger movement towards a perhaps fundamentally different society characterized by its rich interplay between private and public actors and mixed solutions. We refer to Fischer & Kvist (2023) for a recent in-depth review and discussion of this trend.

In many ways, the key role of insurance is to ensure stability of an individual's income and consumption by risk-sharing through time and across individuals (Jensen et al., 2019). It is therefore no surprise that the market for so-called loss of earning capacity insurance (or for the sake of linguistic simplicity: *disability insurance*) has been growing and continues to grow. In Denmark, this includes private individual schemes managed by commercial insurance companies and supplements to company or occupational pension schemes offered by commercial or cooperative pension providers. Here is also the statutory workers' compensation insurance partly managed by the so-called Labour Market Insurance (Arbejdsmarkedets Erhvervssikring, AES) in accordance with the Workers' Compensation Act. The fact that the insurance schemes aim to ensure consumption smoothing entails an unfortunate spillover effect of complexity. Given that the public benefits are supposed to and constitute the main pillar, the design and appropriateness of insurance coverages must be considered residual hereto and under penalization by offset rules as to not under- or overcompensate the insured, see also Sections 3.1 and 5.2 in Jensen et al. (2019). This has given rise to complicated insurance coverages that interact closely with the public benefits, inheriting not least the complexities of the latter. For an overview of the Danish disability insurance market, see also Chapter 6 in Kommissionen om tilbagetrækning og nedslidning (2022).

Disability risk and insurance are also just from a purely biometric perspective more intricate than, say, mortality risk and life insurance. They are both long-term risks, but death is, after all, categorical, while disability comes in various forms and degrees. Proper risk assessment therefore necessitates an effective collaboration between case managers and medical professionals. Furthermore, while mortality risk has developed in a rather stable way, at least disregarding pandemics, disability risk has historically developed in a more volatile manner that directly affects disability insurers. An example hereof is the major increase in mental health claims, with the Danish insurer Velliv in 2022 reporting that 70 % of its payments to young disability claimants were stress related[4].

All in all, this paints a rather complicated picture of disability insurance in Denmark and the other Nordic countries. To exemplify and systematize the complexities, we introduce and study two realistic cases based on our actuarial experience in the Danish insurance and pensions industry. The cases are intended to reflect reality, while also bringing to light characteristics about the product and its risks that are of particular significance from an actuarial perspective.

2.2.2 Alex and Charlie

In the following, we present two fictitious, but quite illustrative, disability cases. The first case, about Alex, describes a possible trajectory in the public system and underscores the raison d'être for disability insurance in Denmark. The second case, about Charlie, focuses on the type of information that becomes available during the insurer's claims processing. The idea is to provide a better understanding of the data that the actuary can actually access for dynamic decision making and for actuarial and statistical modeling purposes.

[4] `https://finanswatch.dk/Finansnyt/Pension/Velliv/article13748200.ece` (accessed 14/1/2025).

Alex's accident

Alex, 33 years, falls on their bike on the way home from work and hits their head. In the first few weeks after the accident, they experience frequent headaches, but their work life is not immediately affected. In the next two months this changes as their condition deteriorates. They go on sick leave for a couple of months during which they receive full pay from their employer; the employer is partially compensated by the municipality. Alex's family physician assesses that their earning capacity has been reduced to one-third in their current job.[a]

Alex attempts to continue to work at their current employer, but only for 12 hours a week. However, this scheme is unsuccessful and they end up being terminated. For a while, Alex's only source of income is therefore sickness benefits combined with savings. The sickness benefits are eventually discontinued, at which time the municipality enrolls Alex in a vocational assessment and resource clarification program. During this program, they receive so-called resource benefits from the municipality.[b]

After about two years, Alex and the municipality identify a more appropriate career choice for Alex. The new job pays less than their original job, but they are able to work half of normal hours for full pay and they receive a so-called flex job salary supplement from the municipality.[c]

Alex worries about the rising inflation and the size of their pension since they had several years where they withheld pension contributions and are now paying a lower amount than before. Alex still suffers from severe and crippling headaches twice or thrice every month.[d]

[a]If Alex had had disability insurance coverage, they could have applied for disability benefits supported by their medical report. If awarded, their employer would be additionally compensated by insurance benefits, however in such a way that the total compensation does not exceed Alex's original salary. This means that it would be less expensive for the employer to retain Alex and, in the longer term, uncover alternative work arrangements.

[b]Resource benefits are generally lower than sickness benefits, so Alex's financial situation worsens as time passes, further negatively impacting their quality of life. The public benefits could have been complemented by insurance benefits, would Alex have been covered.

[c]Note that by this point, Alex has received three different types of public benefits, each subject to differing criteria: sickness benefits, resource benefits, and flex job salary supplement. Several other types of benefits exist, indicating the complexity of the public system. Even if Alex had been able to work normal hours, a disability

insurance could have continued to play some role, given the difference in salary between the new and the original job. This is because disability insurance benefits typically are based on the original salary just before disability.

[d]Premium exemption insurance, also known as premium waiver insurance, would ensure the continuation of pension contributions at the original level in case of a disability. This coverage aims to maintain the same level of welfare for the pensioner regardless of their disability history before retirement. Both the premium exemption insurance and the disability insurance may be subject to suitable indexation as to mitigate the influence of inflation.

Charlie's claim

The insurer receives a disability claim from Charlie, a 45 year old primary school teacher, on the ground of their reduced earning capacity. Charlie has recently returned from a three months sick leave due to work-related stress and burnout, and is currently only working part time, around 16 hours a week.[e]

The insurer's claims processing team reviews Charlie's medical records and rejects the claim, citing insufficient documentation of at least a 50 % reduction in earning capacity, which is an eligibility requirement stipulated in the insurance contract. The rejection is communicated to Charlie about four weeks after they submitted their claim.[f]

Several months later, Charlie reapplies with new medical documentation from their family physician as well as an independent psychiatrist, detailing severe stress and early-stage depression, which has led Charlie to now only working 12 hours a week. The insurer reassesses the claim and approves benefits retroactively, starting three months after the first day of Charlie' initial sick leave to take into account the waiting period stipulated in the insurance contract.[g]

The insurer regularly reviews Charlie's progress. After a doctor concludes that Charlie can return to work full hours, the insurer ceases benefits. However, a second medical opinion shows no improvement, and the insurer hence resumes payments.[h]

Over time, Charlie's condition worsens into a serious depression, and Charlie continues to receive insurance benefits. The insurer eventually refers Charlie to a vocational assessment. Charlie is reluctant to participate, citing a lack of energy. However, the insurer informs Charlie that the program is a prerequisite for receiving benefits, after which Charlie follows the recommendation.[i]

It is concluded that a job as a librarian at a nearby library would be optimal for Charlie's condition and Charlie fortunately manages to secure the job. The insurer regularly obtains updates from Charlie over the next 18 months as their mental health and work hours increase. Finally, Charlie reports that they feel much better and are now working full hours at the library, leading the insurer to cease all benefits.[j]

The next time the insurer hears from Charlie is at their retirement age. They are claiming their pension and look forward to spending more time with their grandchildren.

[e]There has thus been a reporting delay of more than three months since the onset of illness. Long reporting delays are often observed for disabilities and for multiple reasons, including the illness in itself, the need to collect medical documentation, an inattentiveness or unawareness of the insurance coverage, and the lack of a short-term monetary incentive due to the availability of sick leave.

[f]This constitutes a rejection of the claim, however it turns out to only be a temporary rejection given that Charlie later successfully reapplies. Temporary rejections are common as it is difficult to precisely determine the degree of loss of earning capacity. The involvement of medical professionals is required, and the counterfactual nature of the question of how many hours the insured *might be able to* work invokes an inevitable element of subjectivity.

[g]The reapplication is approved. In addition to continuing benefits from the time of adjudication, it also includes a backpay to take into account the claim history. The backpay may be accumulated with interest, the motivation being that the insurer should not have any monetary benefit from delaying the payout. The fact that payments frequently are awarded retroactively implies that the insurer might substantially under-reserve if backpay is ignored in the modeling.

[h]This constitutes a temporary reactivation followed by another successful reapplication resuming the benefits and likely also giving rise to another backpay related to the months for which no benefits where paid. This is once again an example of the complexity associated with adjudication of disabilities, in particular those related to mental health.

[i]The purpose with the vocational assessment is to identify other career choices that could allow Charlie to work more hours and hence receive less benefits. Such an assessment may be required both by the insurer and the municipality.

[j]This constitutes a permanent reactivation. Disability benefits typically continue until retirement, reactivation, or death, whichever occurs first.

2.2.3 Interplay between agents

The cases illustrate the main agents of disability insurance: the insured, the insurer, the employer, and the public system. In fact, the cases are described in a simplifying way. In the first case, about Alex's accident, an insurer was omitted from the narrative to better illustrate what would happen in the absence of insurance coverage. We may omit the insurer since the interplay between the insured and the public system is not directly affected by insurance coverage. However, counseling and financial support can have a significant impact on the trajectory of the insured's disability. In the second case, about Charlie's claim, the role of the public system was kept exogenous. This is because this system, from an insurer's point of view, is an external environment that they and the insured are adapting to. The case is basically a description of how information about a claim arrives at the insurer, which is via direct contact with the insured and the employer rather than the public system, as well as how the insurer may respond to this information.

Figure 2.1 provides a schematic representation of the interactions between the insured, the insurer, the employer, and the public system for disability insurance, where by an interaction we refer to a direct exchange between two or more agents. One central observation is that the graph is almost fully connected, which contributes to the many-faceted complexity of the situation. The only exception is that exchange between the public system and the insurer occurs indirectly, through the insured and the insured's employer.

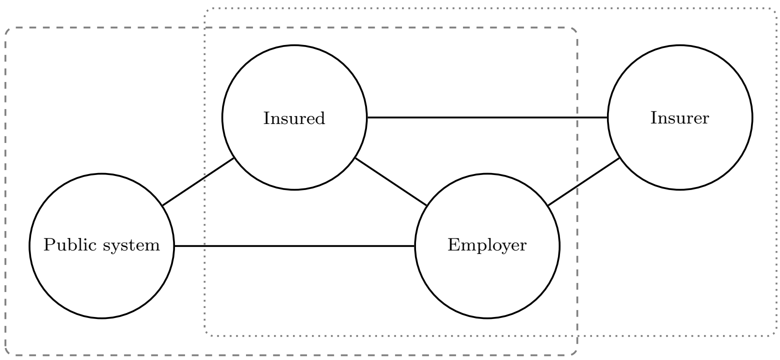

Figure 2.1: *The types of agents involved in disability insurance (nodes) and their interactions (edges). Gray outlines highlight the types of agents present in Alex's case (dashed) and Charlie's case (dotted), respectively.*

It should be mentioned in this context that the insurer may interact more or less directly with the public system through, for example, political channels and in order to influence the general conditions for disability insurance, but of course this does not relate to the individual coverage and claim and is therefore not represented in the figure. The insurer may also collect additional information, for instance from public registries, which could influence the way in which they interact with the insured. In general, the insurer plays a complementary role to the public system. This characteristic also presents itself in the areas of prevention and treatment, which we discuss next.

2.2.4 Prevention

Broadly speaking, insurance and prevention are two distinct and often conflicting responses to risks, and their successful integration is consequently deemed challenging (Dubois, 2011). Nonetheless, prevention has become a trendy topic in insurance, not least in regards to human health (Gauchon et al., 2020b). This is perhaps not too surprising; the dichotomy between insurance and prevention is less pronounced for disability risk than for, say, many non-life risks. This is not least due to the fact that injury and illness, besides the economic costs that loss of earning capacity insurance might cover, has additional human and social costs – and this reduces moral hazard; compare also with the discussion in Botzen et al. (2019).

In the context of disability insurance, prevention initiatives can take many different shapes and be initiated by different agents. In the following, we focus on the role of the insurer as a prevention (and not only insurance) provider. The diversity of potential disability prevention initiatives can be explored by considering various non-hierarchical taxonomies.

We can divide disability risk into two parts: frequency and severity risk. The former encompasses the probability of injury or illness causing substantial loss of earning capacity, while the latter describes the degree and duration of loss of earning capacity. There is clearly some overlap between the risks, so the division should not be considered strict. Prevention initiatives may target the frequency risk, severity risk, or both simultaneously. Sometimes, the nomenclature *primary prevention* is used for an initiative that predominantly targets frequency risk, while *secondary prevention* is used for an initiative that predominantly targets severity risk (Dubois, 2011; Kenkel, 2000). There is also the concepts of *self-protection* and *self-insurance*, see Ehrlich & Becker (1972), which we have seen used in place of primary and secondary prevention. However, we would like to warn against such interchangeable use. Self-protection and self-insurance also include initiatives that do not directly or even indirectly reduce risk, such as risk retention. To us, prevention implies affecting or intervening in the underlying risk environment. This is fundamentally different from loss protection for the insurer due to product design or reinsurance.

The primary role of the actuary in connection with prevention initiatives lies in using mathematical and statistical methods to perform impact evaluation, that is to measure the causal effects of the initiatives. It is, after all, within the actuary's role to assess which initiatives are appropriate from a cost perspective. This necessitates, however, that the actuary is involved in identifying potential prevention areas and in the design of appropriate initiatives. Paraphrasing Ronald Fisher: To consult the actuary after the initiative is completed is merely to ask them to conduct a post mortem examination.

As mentioned, it is possible to imagine many different prevention initiatives in connection to disability insurance. However, they all fall within one of five categories. There is the insurance coverage in itself, which provides peace of mind for the insured upon disability and consequently might have a reducing effect on the severity risk (Fischer & Kvist, 2023, p. 120). For occupational schemes, the insurer and the employer can cooperate on initiatives to improve the physical and psychological work environment and thereby reduce the frequency risk and, potentially, but likely to a lesser degree, also the severity risk.

The remaining three categories of initiatives are of increased actuarial interest. They encompass initiatives targeting the individual insured, but at different risk stages. First, there are health promotion initiatives, such as wearables (Spender et al., 2019), that aim to increase overall health and resilience. Second, prevention initiatives may be initiated just before a potential disability based on early warning indicators such as a health insurance claim or sick leave to either reduce the likelihood of disability or reduce the degree and shorten the duration. For example, the insurer may contact potential claimants and offer them additional short-term health services, such as immediate access to physical therapists and psychologists. Finally, prevention initiatives may be initialized during disability to improve the recovery rate of the insured, possibly based on pre-existing public programs. For example, the insurer might propose or even require that the insured completes a vocational assessment.

Across all categories, measuring the preventive impact of an initiative is challenging. It is necessary to separate the effect of the intervention from all other effects, including global health trends and local changes to underwriting and claims processing practices. The gold standard would be some sort of randomized trial, but this is often deemed ethically or practically undesirable. Instead, insurers might target the initiative towards high-risk individuals, adopting a cut-off selection rule. This, however, comes with additional challenges. The group subjected to the initiative would no longer be comparable to the group not subjected. Further, there is not necessarily a correlation between high-risk individuals and individuals for whom the intervention would have the largest positive effect. Therefore, causal insights are required to operationalize impact evaluations that are accurate and fit for purpose. We return to and expand on this key insight in Section 2.5.

2.2.5 Systematic and unsystematic risks

The discussion up until this point shows that disability insurance is complicated and subjects the insurer to many risks; this includes both systematic and idiosyncratic (unsystematic) types of risks. Idiosyncratic risks are specific to the insured or the insurance policy and thus manageable by diversification in accordance with the central limit theorem. Systematic risks affect the whole portfolio, or substantial segments of it, jointly and hence cannot be reduced by simply increasing the size of the portfolio. In Figure 2.2, we provide an overview of the main risks associated with disability insurance.

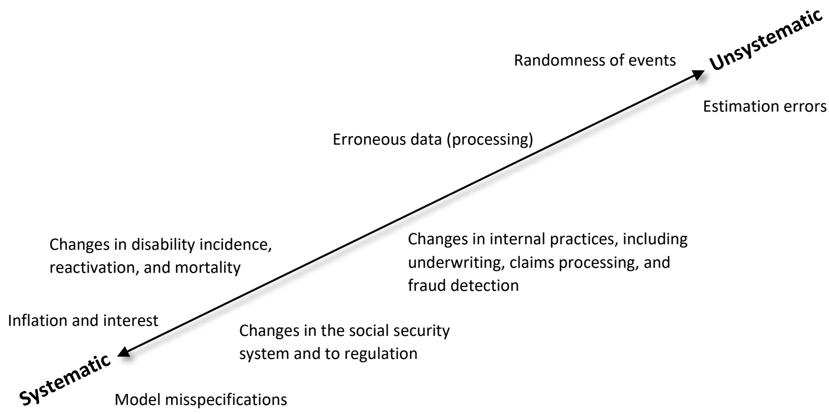

Figure 2.2: *Primary systematic and idiosyncratic (unsystematic) risks for disability insurance.*

The changes in the incidence of biometric events in particular, but also changes to internal practices and the social security system, can be abrupt, in which case one might speak of *shocks*. There is, however, also significant long-term risk, so-called *trend* risk. To mitigate said risk, it is essential to monitor biometric, economic, and societal trends and to adjust one's models accordingly. Reinsurance could also be an option, both in regards to biometric shocks and to counteract the interest rate risk associated with the potentially long cash flows. In general, the models should be well calibrated and supported by strong underwriting practices and fraud detection programs. Finally, prevention initiatives may reduce both unsystematic and systematic risk, confer with Subsection 2.2.4.

We should like to stress that actuaries are well-positioned to contribute to the key activities involved: product design, risk modeling, and optimal prevention. In the following three sections, we address each of these areas separately. The emphasize will be on relevant recent developments in actuarial science and associated opportunities for actuaries of the present and the future.

2.3 Product design

In this section, disability insurance coverages are formalized and compared with a focus on the Danish market. Our approach is initially descriptive, seeking to explore and understand the characteristics of existing products rather than discuss optimal design. Further, we are looking for common instead of distinguishing design trends, focusing on annuities. This simple-sounding task already reveals a number of challenges that hitherto have received limited to no attention in the actuarial literature.

Fundamental to the product is the specification of what constitutes lost earning capacity. The degree of lost earning capacity is usually taken to be the proportion of hours the insured is able to work compared to a standard number of working hours. The proportion may be based on the hours it is possible to work in the insured's own original occupation, a similar occupation, or any occupation. These specifications play an important role not least because they may influence the incidence rate, reactivation rate, and payment sizes considerably, which also has implications for subsequent actuarial modeling. Due to the rich interplay with the public system, it is often natural for the definition stipulated in the insurance contract to be at least partly aligned with the one used in the public system, which may result in the insurance coverages being subjected to the whims of politics. In Denmark, three-fifths of the insured receiving disability benefits in 2020 were awarded benefits based on eligibility for early retirement pension, while the remaining two-fifths were awarded based on the insurer's internal health assessment of the insured (Kommissionen om tilbagetrækning og nedslidning, 2022, Boks 6.4).

Having specified what constitutes loss of earning capacity, the product further depends on the specification of the size of the compensation in case of a disability. In Subsection 2.3.1, realistic contractual payments are formalized. Contractual payments are those that are stipulated in the insurance conditions. Further complexity arises in ensuring that the insured receives the benefits they were eligible for due to reporting delays and adjudications; this is the topic of Subsection 2.3.2.

2.3.1 Contractual payments

We here give a continuous time description of the contractual payments. This is a natural approach since the underlying biometric events occur in continuous time. In reality, payment streams feature lump-sum payments at discrete points in time, but payment rates are mathematically convenient and offer a useful approximation when the payment frequency is sufficiently high; this is the case for disability annuities with monthly payments.

Simple disability annuity

Multistate models provide a parsimonious way to formalize the contractual payments. Let $Y = \{Y(t)\}_{t \geq 0}$ be a stochastic process that is piecewise constant and which takes values in the state space $\mathcal{J}$ depicted in Figure 2.3. The definition of *Disabled* in the model is to be understood as being eligible for disability benefits according to the stipulations in the insurance conditions, and the contractual payments

$B = \{B(t)\}_{t\geq}$ of a simple disability annuity may then be cast as

$$\mathrm{d}B(t) = 1\{Y(t) = 2\}b\,\mathrm{d}t, \qquad B(0) = \pi,$$

where $\pi < 0$ is the initial premium and $b > 0$ is the disability benefit rate. This definition and corresponding mathematical expression hide the stipulations for eligibility. The latter usually require the disability to have occurred within the coverage period, which commonly spans one to three years, the proportion of lost earning capacity to be sufficiently high, for example 50 %, and a deferred (waiting) period of, say, three months to have passed.

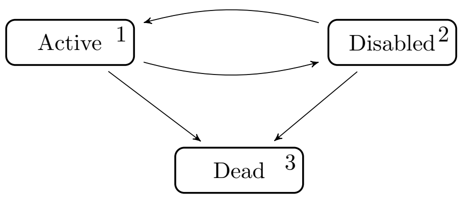

Figure 2.3: *State space $\mathcal{J}$ for a classic disability model. The arrows represent the possible transitions.*

Defining the disabled state as being eligible for disability benefits has the effect of placing the difficulty in modeling onto the probabilities of transitioning to and from the disabled state, rather than placing it in the payment rates. This implies, among other things, that the time of disablement and the transition probabilities become product dependent. However, using a product-independent definition of the time of disablement such as the first day of sickness leads to other complications in formalizing the contractual payments. Taking the deferred period as an example, one could add an indicator to the benefits that the duration in the disabled state has exceeded the deferred period, but this might not correctly describe what happens in practice since the deferred period is often annulled if the insured relapses after a short reactivation.

Thus, even in this simple setting where benefits are constant and claim settlement processes are not involved, the non-hierarchical nature of disability trajectories leads to disability insurance having a high level of complexity compared to other types of coverages. These complexities seem to be ignored in the classic literature on actuarial multistate modeling. In the formulations of Example 3.2 in Christiansen (2012) and Section 3.2.2 in Haberman & Pitacco (1998), it is for instance not possible to receive payments when entering the disabled state after the end of the coverage period and the deferred period resets after each entry to the disabled state. This means that temporary reactivations for the same underlying disability event are not accounted for. However, if temporary reactivations or multiple distinct disabilities are disregarded, the added complexities disappear, and one may thus give an approximate description of the contractual payments based on the classic literature by ignoring multiple disabilities and using the alternative state space depicted in Figure 2.4.

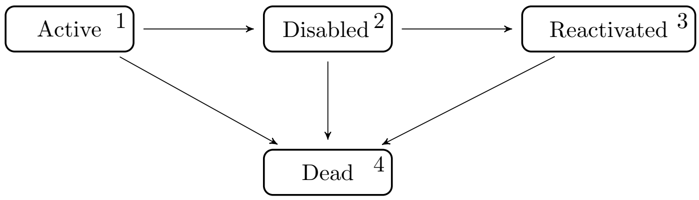

Figure 2.4: *State space for an alternative disability model with separate reactivated state. The arrows represent the possible transitions.*

When disability benefits are constant, underwriters should manually take people's financial situation into account and anticipate what benefits they will receive from other sources such as the public system. If the insured's total income would increase substantially in the event of a disability, as they may receive income from several sources, there is the potential for moral hazard as noted in Chapter 3 of Haberman & Pitacco (1998), but this effect is likely limited due to the reasons outlined in Subsection 2.2.4. Another downside, however, is that the insured has potentially over-insured themselves and therefore pays a premium that is too high in relation to the utility value of the coverage. On the other hand, if the insured's income substantially decreases in the event of a disability, then the insurance is not achieving its purpose of ensuring that the individual can maintain a similar or only slightly lower level of consumption. Accommodating these complications requires tailored insurance coverages since the public benefits change as the disability progresses. Such complicated coverages designed around the social security system, which are by now rather common, are the focal point of the next subsection.

Disability annuity

Using the state space from Figure 2.3 with the same definition of the *Disabled* state, a simple disability insurance product designed around the social security system has contractual payments on the form

$$\mathrm{d}B(t) = 1\{Y(t) = 2\} \min\{\max\{s - e_t - c_t, 0\}, d \times s\} \,\mathrm{d}t, \qquad B(0) = \pi,$$

where d is the percentage of salary covered, for example two-thirds, s is the salary just prior to the disability, e_t is the earning capacity at time t, and c_t is the compensation from social security and other insurance coverages at time t.

The Danish market also features more complicated coverages. Inflation regulation is, for instance, a common feature; here s is scaled with an appropriate inflation index as to maintain the purchasing power. Additionally, the deduction of benefits from other insurance coverages may first take place when the total income exceeds some higher proportion of the previous salary than d, say 80 %. The benefit rate might also be reduced whenever the insured is enrolled in public benefit programs that deduct compensation from insurance coverages; an example hereof is the resource clarification program. Furthermore, there is the entanglement with tax rules, pension contributions, and labor market contributions. The intention behind such modifications to the simple product is to stabilize the insured's total income level, while keeping the insurer's expenses – and therefore the premium rate – as low as possible.

These types of products are not well-studied in the actuarial literature. However, they can and should be studied using well-known methods from event history analysis and the mathematics of life insurance. It is worthwhile to note that it is not necessary to explicitly forecast e_t and c_t to calculate reserves and premiums; it suffices to model the average payments at future times (by the law of iterated expectations). This way of thinking might also be applied to incorporate some of the complexities described in Subsection 2.3.1, rather than just placing it on the transition probabilities. The actuarial community would benefit from more work in these areas.

2.3.2 Claim settlement payments

While the insurance contract determines the payments that the insured is eligible to, as described in Subsection 2.3.1, the actual process of awarding the insured these benefits is complicated by the claim application process.

There is both a reporting delay incurred by the insured (from the occurrence of the event until the insured notifies the insured) and a further delay incurred by the insurer due to its adjudication process (from when the insurer is notified until eligibility is determined). If the insurer determines that the insured is eligible for benefits, the insurer has to ensure that the insured is compensated according to the stipulations in the insurance contract. This is done by awarding backpay from the end of the deferred period to the time where the decision to award benefits is made – in addition to the annuity payments from that time until the insured reactivates, dies, or reaches their retirement age. In determining the size of backpay, the time value of money has to be taken into account. Therefore, backpay is usually accumulated with interest matching the periods that the benefits pertain to. A similar phenomenon occurs if benefits are stopped by the insurer, but the insured reapplies and successfully demonstrates that the payments were wrongfully stopped. The realized cash flow $\mathcal{B} = \{\mathcal{B}(t)\}_{t\geq 0}$ for the simple disability annuity of Subsection 2.3.1 consequently takes the form

$$\mathrm{d}\mathcal{B}(t) = 1\{Z(t) = 2\}b\,\mathrm{d}t + \text{backpay}(t)\,\mathrm{d}N(t)$$

where Z is a stochastic process taking values in the same space as Y, but only visiting the state *Disabled* when the insured is actually receiving annuity benefits, and where $N(t)$ counts the number of times backpay has been awarded before time t. The backpay takes the form

$$\text{backpay}(t) = \int_{\alpha(t)}^{\beta(t)} \exp\left(\int_s^t r(v)\,\mathrm{d}v\right) b\,\mathrm{d}s$$

with $\alpha(t)$ and $\beta(t)$ delimiting the relevant period and r denoting the interest rate.

These complications have only recently been addressed in the multistate modeling literature on disability insurance, namely in Buchardt et al. (2023, 2025) and Sandqvist (2025). One may therefore reasonably ask whether they can be ignored. For pricing or reserving at the inception of the contract, the answer is to a large degree 'yes', but for risk management and reserving during and after the coverage period, the answer turns out to be a resounding 'no' unless the reporting and adjudication delays are very short, see Subsection 2.4.4 for detailed explanations as to why.

In the context of estimation, the presence of reporting delays as well as incomplete adjudications biases the sample. Fitting a model to have good predictive performance on the observed data, for instance using out-of-sample-based methods, leads to a biased model – because the sample is biased. Reporting delays result in the sample containing too few observed disabilities, while later rejections upon adjudication result in the sample containing too many observed disabilities. It is therefore not even clear whether one over- or underestimates the disability incidence rate. The actuary plays an important role in recognizing and taking these effects into account when performing reserving and impact evaluations for disability insurance. In fact, actuarial reserving as a field has in many ways evolved around accommodating such sampling effects, which in addition to reporting delays and incomplete adjudication encompasses left-truncation and right-censoring stemming from the finite observation window and from insured entering and leaving the portfolio.

2.4 Actuarial modeling

As seen in Section 2.3, formalizing the contractual payments for a single insured is inherently difficult, and this also complicates the risk modeling. In this section, we discuss and compare relevant actuarial models, providing new insights into the adequacy of current stochastic modeling approaches when applied to disability insurance coverages. The overall aim is to obtain sound actuarial models for risk management, while maintaining as much simplicity as possible; they should have predictive prowess, and it should be possible to validate them in- and out-of-sample. This aim epitomizes the identity and role of the actuary since achieving it requires a synthesis of insight about the design of the products, the associated biometric and financial risks, as well as the available data and its consequences for the probabilistic and statistical aspects of the model.

2.4.1 Reserving and pricing

Premiums are said to be actuarially fair if they equal the expected expenses. From a narrow mathematical standpoint, obtaining actuarially fair premiums is therefore a special case of reserving. Reserves represent the value of the future cash flow, with market-consistent valuation leading to reserves V on the form

$$V(t) = \mathbb{E}\left[P(t) \mid \mathcal{F}_t\right], \tag{2.4.1}$$

$$P(t) = \int_t^\infty \exp\left(-\int_t^s r(v)\,\mathrm{d}v\right) \mathrm{d}B(s), \tag{2.4.2}$$

where r is the interest rate, $\mathrm{d}B$ is the cash flow, $\mathcal{F}_t$ is the information that is known at time t, and time is measured as the duration since the inception of the contract. This is the (conditional) expectation of the present value of future benefits with respect to some market-consistent (equivalent) probability measure. Actuarially fair premiums are those for which $V(0) + B(0) = 0$, which in our setting implies an initial premium $\pi = -V(0)$.

In practice, however, pricing and reserving models usually have different targets. For pricing, it is important to have precise premium sizes on an individual level to avoid adverse selection. For reserving, both accurate *size* and *timing* are important. The insurer has to charge enough money in the beginning to pay the later claims, which favors prices that are accurate on a portfolio level. Regulation compels the insurer to set aside additional capital to maintain solvency even in cases that are considerably worse than current best estimates, which may be thought of as accounting for the market price of risk. These are aspects related to the size of the reserve. To have good timing, the reserve has to be close to the present value of future benefits at each point over the duration of the contract period, not just at the beginning. Otherwise, there will be a mismatch between the release of reserves and incoming losses. This again leads to swings in the insurer's financial results and uncertainty about the insurer's financial situation that makes it difficult to run an efficient operation. Furthermore, the timing of the reserve is crucial for hedging of the financial risks associated with the future cash flow. If possible, the insurer would also want good timing on a more granular level to be able to monitor the business on a sub-portfolio level.

In addition to size and timing, actuaries also care about *model complexity* and *statistical complexity.* Model complexity comes in many forms. We here restrict our attention to the complexity of the mathematical concepts involved in specifying the model, and not for example the explainability of the model and its output. This notion of model complexity is relevant for the time and skill needed to understand the model, but also influences the risk of implementation errors and the ease with which the model may be communicated to non-experts. Statistical complexity is taken to include the hardness of both the statistical learning problem and the practical issues of modeling choices and implementation. In the following three subsections, different model paradigms are discussed with these considerations in mind; a comparison may be found in Subsection 2.4.5.

2.4.2 Non-life individual reserving models

Since formalizing the contractual payments is difficult, one may ask whether it is more favorable to disregard the known structure of the cash flow $\mathrm{d}B$ and also learn this from the data. This is similar to the individual reserving approach from non-life insurance, and one could hence use recent methods from this area such as Crevecoeur et al. (2022b) or Yang et al. (2024). Their approach is to model the full real-time development of the claim and all time-varying covariates; note, however, that Yang et al. (2024) assume the entire path of the time-varying covariates to be known. A brief summary goes as follows. A high-dimensional stochastic process $X = \{X(s) : s \leq t\}$, which represents all available policy and claim level information, is introduced, and the available information is specified as $\mathcal{F}_t = \sigma\{X(s) : s \leq t\}$. The distinction between contractual and realized payments is abandoned, meaning $\mathrm{d}B$ is taken to just be $\mathrm{d}\mathcal{B}$; compare to Subsections 2.3.1 and 2.3.2. This choice implies that $B(t)$ is recoverable from the information $\mathcal{F}_t$, so the present value in (2.4.2) is indeed the relevant present value, and the reserve of (2.4.1) is operational. However, the reserve does not admit a closed-form expression, and since calculating them numerically via differential equations is undesirable due to the high dimensionality of the model, Monte Carlo methods are used.

This approach makes the probabilistic description of the model relatively simple, but results in a learning problem with high statistical and computational complexity. In particular, the curse of dimensionality implies that the appropriate amount of data needed to estimate the model grows quickly with the dimensionality of the model, not least since the signal-to-noise ratio is low in insurance applications. Additionally, some of the assumptions that are made to reduce the dimensionality of the model in these specific papers would be problematic for disability insurance. More specifically, the Poisson assumption of Crevecoeur et al. (2022b) is unpleasant since one can in reality only be compensated for at most one disability at any time. Further, in Yang et al. (2024) it is assumed that the different payment sizes and waiting times are independent given the settlement time, which does not describe disability annuities well since payments occur monthly after the first payment and the payment sizes are similar for payments that are temporally close.

Another challenge with the non-life approach is that the individual model components are tuned to be accurate for their separate regression problems, and this may lead to substantial bias for the aggregate reserves due to plug-in bias. In Crevecoeur et al. (2022a), it is suggested to remove this aggregate bias by rescaling the predictions such that the aggregate predictions historically agree with the aggregate observations. An alternative that does not seem to have been explored in an actuarial context, but which might be useful, is to remove the plug-in bias for the portfolio reserve using a so-called one-step estimator, confer with Section 4 of Kennedy (2022).

2.4.3 Aggregate reserving models

General considerations

A different approach, which also does not utilize the known structure of the contractual payments, is to use aggregate models for which $\mathrm{d}B$ represents the realized cash flow of the portfolio instead of a single individual, and for which the information used for reserving, written $\mathcal{F}$, is kept at portfolio level (even if more granular information is available). Aggregate models have the advantage that they target the portfolio level reserve and that they are comparatively simple to implement.

However, aggregate models have four general disadvantages. First, they cannot use granular data, which may impede their predictive performance. Second, they are slow to capture shifts in the covariate composition of the portfolio or new trends such as an increase in mental health-related claims. Third, they lack explainability in the sense that the models offer little assistance in pinpointing underlying reasons for deviations between observations and predictions. Together, these three points make it difficult to monitor the business on a more granular level and to make ad-hoc adjustments based on additional external expert information. Fourth, they do not provide consistent estimates when used to impute right-censored claims in order to obtain a regression sample for fitting pricing models as is often done in non-life insurance; see the discussion in Crevecoeur et al. (2022b), especially Section 3.1.

Application to disability insurance

There are also some additional challenges associated with using aggregate models for disability insurance reserving. The aggregate reserving literature has generally evolved around the assumption that the available data consists of payments occurring in a run-off triangle consisting of accident years (rows) and development years (columns). The run-off triangle is commonly, but not always, assumed to be aggregated on some time grid, say monthly or yearly. Since disabilities frequently lead to several decades of benefits, while the eligibility criteria change rather frequently, one has in most situations too few years of data to be able to estimate the outstanding liabilities using chain-ladder, or variants thereof, since no complete run-off is observed. It is hence almost always necessary to impose some structural assumptions that allows one to extrapolate the future expenses for late development years.

Individual-level data shows that insured who have received benefits for a couple of years have very low reactivation and mortality rates, so claims that are open after a couple of years generally run until the retirement age. Furthermore, in Denmark, the social security benefits also tend to stabilize after some years since insured who remain disabled find a suitable flex job or retire early. Therefore, a pragmatic approach could be to first model the average age of insured who have open claims after some fixed number of years, and then in a second step to cast subsequent payments as constant until the retirement age subtracted the aforementioned average age. This would likely lead to a reasonable size of the reserves, but poor timing.

If data on individual ages is available, a perhaps preferable alternative is to include age as a covariate in the aggregate model. Some approaches for incorporating such individual information in chain-ladder models are given in Wüthrich (2018) and Delong et al. (2021). It also seems like it would be possible to extend the approach of Bischofberger et al. (2020) to include covariate information such as age, but the setup comes with other strong assumptions that are undesirable for modeling disability insurance, namely that payments are independent and all claims consist of single payments. If portfolio reserving is the only objective, we expect that the best size and timing can be obtained by leveraging granular data, but using a model that targets the portfolio reserve as discussed here and at the end of Subsection 2.4.2. This, however, comes at the cost of a significant increase in complexity compared to traditional aggregate models.

2.4.4 Multistate individual reserving models

Literature and challenges

The actuarial literature on individual disability modeling using multistate methodology is substantial, confer with Taylor (1971), Segerer (1993), Haberman & Pitacco (1998), Christiansen (2012), Djehiche & Löfdahl (2014), Aro et al. (2015), Sandqvist (2025), and the references therein. Nonetheless, the increasing complexity outlined in Subsection 2.3.1 arising from temporary reactivations and several distinct disabilities has only just received some awareness, see Sandqvist (2025), and the situation where payments are on the form described in Subsection 2.3.1 also warrants further exploration. The latter situation can to some extent be compared to the situation where the payment rate depends on the degree of disability as in Subsection 4.2 of Segerer (1993), so that the annuity formula is simply scaled with the mean degree of disability. Further, this approach can be extended to model mean benefits in the face of deduction from social security benefits, confer with Remark 2.1 of Sandqvist (2025). We believe that this is an important avenue for actuaries and actuarial scientists to explore.

There is also the effect of the claim settlement process to take into account, confer with Subsection 2.3.2. The implications of reporting delays on the estimation of disability incidents has been studied in Kaminsky (1987), König et al. (2011), and Aro et al. (2015). In the context of reserving, it is noted in Segerer (1993) that insurers can set up IBNR reserves if they have the data to support them, but additional details are not provided. It is only recently in Buchardt et al. (2023) and Sandqvist (2025) that the problem has been attempted to be tackled in a somewhat general manner. The topic is, however, frequently explored in non-life insurance contexts with approaches similar to those described in Subsection 2.4.2, see for example Norberg (1993, 1999), Haastrup & Arjas (1996), Antonio & Plat (2014), Lopez et al. (2019), Okine et al. (2022), Crevecoeur et al. (2022a,b), and Bücher & Rosenstock (2024).

The classic multistate approaches to disability insurance modeling have focused on reserves on the form (2.4.1) with the specifications that $\mathrm{d}B$ are the contractual payments and the information $\mathcal{F}_t$ equals $\sigma\{Y(s) : s \leq t\}$ for the biometric state process Y generating the contractual payments. This leads to two fundamental issues for reserving whenever there are non-negligible reporting delays and adjudication processes, and these challenges do not seem to have been discussed prior to Buchardt et al. (2023) and Sandqvist (2025). The issues are as follows. First, if there are reporting delays and adjudication processes then the choice $\mathcal{F}_t = \sigma\{Y(s) : s \leq t\}$ is not operational for reserving since this information might not be available at time t. For example, the insured can become disabled before time t, but report this after time t. Similarly, since $\{Y(s) : s \leq t\}$ might not have been observed at time t, the insurer might not have paid $B(t)$ at time t, so the present value in (2.4.2) is not the relevant present value. For example, the insurer might not have paid any disability benefits by time t, while the insured is disabled and eligible for benefits before time t, which would then lead to backpay after time t.

Recent contributions

Recently, Buchardt et al. (2023) and Sandqvist (2025) introduced a multistate framework that seeks to accommodate the challenges outlined in Subsection 2.4.4. The idea is to keep V defined as in the classic multistate modeling framework, but introduce another reserve $\mathcal{V}$ according to

$$\begin{aligned}\mathcal{V}(t) &= \mathbb{E}\left[\mathcal{P}(t) \mid \mathcal{G}_t\right], \\ \mathcal{P}(t) &= \int_t^\infty \exp\left(-\int_t^s r(v)\,\mathrm{d}v\right) \mathrm{d}\mathcal{B}(s),\end{aligned}$$

where $\mathrm{d}\mathcal{B}$ is the realized cash flow and $\mathcal{G}$ is the available policy- and claim-level information. In other words, this is essentially the same reserve as used in non-life insurance contexts, confer with Subsection 2.4.2, and hence uses the correct present value with an operational information. The modeling paradigm, however, deviates from the non-life approach by retaining the a priori known structure of $\mathrm{d}B$ and introducing relevant structural assumptions regarding the relation between $\mathrm{d}B$ and $\mathrm{d}\mathcal{B}$ as well as between $\mathcal{F}$ and $\mathcal{G}$. These assumptions allow one to express $\mathcal{V}$ in terms of the classic reserve V with certain natural modifications.

In particular, as briefly mentioned at the end of Section 2.3, one has $\mathcal{V}(0) = V(0)$ since $\mathcal{G}_0 = \mathcal{F}_0$ and $\mathcal{P}(0) = P(0)$. The former is because at policy initiation the state of the insured is known and no additional claim settlement information is available, while the latter follows from the fact that no backpay may relate to events before policy initiation. Therefore, the effect of the claim settlement process may be ignored at policy initiation. However, if there are significant reporting and adjudication delays, then there may be substantial discrepancies between the two filtrations and the two present values after time t. This cannot be ignored if the reserves are to have good timing.

A disadvantage of this modeling paradigm is that it usually necessitates rather strong independence assumptions which if violated may impair both the size and timing of the reserves. It is possible to accommodate certain violations of the assumptions, confer with Remark 3.9 and 3.10 of Sandqvist (2025), but only at the cost of additional model complexity. Another disadvantage is the need for quite granular data to be available to the actuary.

The statistical and computational complexity is comparable to current state-of-the-art models in the multistate life insurance literature and therefore moderate. Furthermore, the probabilistic description of the model is rather complicated, but not substantially more demanding than for regular multistate models. This means that if one is already using multistate models for reserving, it can be considerably easier to adopt the methodology of Buchardt et al. (2023) and Sandqvist (2025) compared to switching to a model inspired by non-life insurance methods. This may be of particular interest to insurers based in Denmark, not least due to the rich tradition for multistate modeling as reflected in the literature emanating from Copenhagen.

In comparing the multistate paradigm to the non-life individual paradigm described in Subsection 2.4.2, advantages are found to be that the structure of the contractual payments is exploited, the dimension of the model is likely considerably lower, and that the reserve may be calculated efficiently via differential equations rather than Monte Carlo methods. Furthermore, the resulting reserve is composed of modifications of classic reserves and hence, is easier to interpret and adjust on an ad-hoc basis. Based on Theorem 3.4 in Sandqvist (2025), the reserve for an insured with no reported disability claim is approximately

$$\mathcal{V}(t) = V_a(t) + \int_0^t V_i(s,0)\mathbb{P}(\text{Reporting delay} > t-s)\mu_{ai}(s)\,\mathrm{d}s,$$

where $V_a(t)$ is the classic reserve for an insured that is active at time t, $V_i(s,u)$ is the classic reserve for an insured that became disabled at time $s-u$, and μ_{ai} is the disability incidence rate. The first term corresponds to disabilities that are covered but not incurred, while the second term corresponds to disabilities that have occurred but have yet to be reported.

The situation of a reported disability claim is more involved, since one most distinguish between three cases: reported-but-not-paid, currently eligible for disability benefits, and previously eligible for disability benefits. Given the information available at time t, let $G(t)$ denote the time point from which the insured is eligible for disability benefits if the claim is (re)awarded and let $W(t)$ be the corresponding disability duration at $G(t)$. The reported-but-not-paid reserve then approximately reads

$$\mathcal{V}(t) = \mathbb{P}(\text{Claim is awarded} \mid \mathcal{G}_t)V_i(G(t),0).$$

This follows from Theorem 3.7 in Sandqvist (2025). Furthermore, the reserve for an insured currently eligible for disability benefits is

$$\mathcal{V}(t) = V_i(t, W(t)),$$

confer with Theorem 3.8 in Sandqvist (2025). Finally, the reserve for an insured previously eligible for disability benefits is approximately

$$\mathcal{V}(t) = \mathbb{P}(\text{Claim is reawarded} \mid \mathcal{G}_t)V_i(G(t), W(t)).$$

For estimation of the parameters required to calculate $\mathcal{V}$, in Sandqvist (2025) it is established that this is a special case of the statistical problem studied in Buchardt et al. (2025), so that their methods and results may be directly applied. Subject to certain simplifications, this leads to the following solution for the estimation of the disability and reactivation hazards. The exposures $E(t_i) = \int_{t_i}^{t_{i+1}} 1\{Y(s)=1\}\,\mathrm{d}s$ for the disability hazard have to be multiplied by $\mathbb{P}(\text{Reporting delay} \leq t-t_i)$, where t is now the time of statistical analysis. Furthermore, reported-but-not-paid occurrences qualify as disability occurrences, but only after scaling with $\mathbb{P}(\text{Claim is awarded} \mid \mathcal{G}_t)$, while (temporary) terminations of disability benefits qualify as (permanent) reactivations, but only after scaling with $1-\mathbb{P}(\text{Claim is reawarded} \mid \mathcal{G}_t)$. Finally, the scaling factors $\mathbb{P}(\text{Claim is awarded} \mid \mathcal{G}_t)$ and $\mathbb{P}(\text{Claim is reawarded} \mid \mathcal{G}_t)$ can be estimated by somewhat standard event history analysis methods. Altogether, this gives rise to a two-step estimation procedure in which estimation of the scaling factors precedes estimation of the hazards.

2.4.5 Comparison

Table 2.1 offers a summary of the points discussed in Subsections 2.4.2–2.4.4. It is not intended to capture all nuances, but rather to indicate the expected relative potential of the different modeling paradigms. For the prediction column, the entries represent the performance on a portfolio and individual level, respectively. Other aspects discussed but not represented in the table are explainability and computational complexity. It is up to the actuary to determine which modeling paradigms is best suited for risk assessment and management, including pricing and reserving, taking into account the resources, challenges, and strategies that are pertinent to their situation.

	Prediction		**Complexity**	
	Size	**Timing**	**Model**	**Statistical**
Non-life	poor/good	poor/good	moderate	high
Aggregate	decent/poor	decent/poor	low	low
Multistate	decent/decent	decent/good	high	moderate

Table 2.1: *Characteristics of the different actuarial modeling paradigms: non-life individual reserving models, aggregate reserving models, and multistate individual reserving models. The entries for size and timing are on the form portfolio/individual performance. For example, aggregate models are expected to yield reserves with decent size on a portfolio level and poor size on an individual level.*

2.5 Impact evaluation for prevention

The actuary might not have to play a prominent role whenever a prevention initiative merely constitutes an implementation of existing principles based on empirical evidence from, for instance, the health sciences. In any case, it is natural and advisable to base the design of initiatives not only on the insurer's own data, but also on the broader scientific literature. However, the situation of insurers' often differs from those in the health sciences literature: They are a different kind of health partner than, for example, an employer or care provider; their target demographic may differ from, for example, the wider population; and their basis and options for action may look different. It is therefore important to be aware that empirical insights from the literature cannot necessarily be applied one-to-one to the insurer's situation. Consequently, insurers would benefit greatly from being able to quantify the impact of prevention and treatment initiatives using their own data.

Both portfolio-level and more granular impact evaluations are valuable. A portfolio-level impact evaluation can reveal whether an initiative in force is cost effective, in the sense that the health promotion (measured as a reduction in disability benefit payout) exceeds the costs of administering the initiative. A more granular impact evaluation, on the other hand, could provide valuable feedback regarding the optimal design of the initiative and its primary target group.

While prevention is a topical subject in insurance, there does not exist much actuarial literature dedicated to it. In Gauchon et al. (2020b, 2021), a ruin theoretic setup is considered, and the optimal amount to invest into prevention is studied with prescience about how changes in investment affect the rate of claim arrivals. To operationalize this theoretic contribution, a natural question is: *How can one infer the effect of investments into prevention on the claim arrivals and the claim sizes?* Furthermore, since the same monetary amount could be invested in many different ways, further relevant question are: *Which initiatives are most effective in reducing the risks? Are these effects heterogeneous across different groups of insured – and in which sense?*

In Gauchon et al. (2020a) and to address the second and third question, it is proposed to cluster insured in terms of their consumption of different health services, for example psychiatry and radiology consultation, and then target prevention programs to high-risk individuals. This is a good starting point: if the risk is high, it may be easier to achieve a substantial risk reduction. However, this might not be a cost-effective solution. For example, insured at moderate risk could perhaps be relatively more receptive to the prevention initiative. There are clear parallels to the following classic case from sales:

> *Consider a department of telemarketers selling insurance. At first, the department makes calls at random. Over time, the department identifies that young women in particular are inclined to buy insurance during the call. The department therefore starts calling only younger women, but this causes sales figures to drop. By conducting interviews, it is found that almost all of the women who buy the insurance would have bought it regardless of the phone call. On the other hand, men, for example, are less likely to buy the insurance – but may be convinced through a phone call. The company therefore changes its strategy to never call young women, resulting in an increase in sales.*

The example shows that interventions should not necessarily target those with a high likelihood of buying (high risk) or even a high likelihood of buying given they receive the call (high risk given the intervention), but perhaps rather those for which the likelihood of buying increases the most by receiving a call (high impact given the intervention).

A systematic approach to impact evaluation, encompassing the three aforementioned questions, is causal inference – because the questions are inherently causal: They ask what the impact of an intervention is on the claim frequency and the claim sizes, and how this effect varies depending on the covariates, henceforth denoted W. In general, the causal effect of a prevention initiative (or treatment), henceforth denoted A, on some function of the future cash flow (or outcome), henceforth denoted Y, can only be inferred if there exist similar individuals and only some receive the treatment, while the others do not. Situations with detailed knowledge of the causal relationship between Y, A, and W constitute an exception; here graphical models, for example, may provide other ways to infer the causal effect. Insurance applications, on the other hand, typically fall outside this exception due to the complexity of the underlying risks (life trajectories), the low signal-to-noise ratio, and restrictions on the availability of relevant covariates to the insurer. For an introduction to causal inference, we refer to the modern classics Pearl (2009) and Hernán and Robins (2020).

One situation where causal inference, however, remains possible is when the assumptions of *no unmeasured confounding* and *positivity* are satisfied. Assume for simplicity that A is binary and that the outcome Y is fully observed and not for instance subject to right-censoring. Further, let $Y^{(1)}$ and $Y^{(0)}$ correspond to what the outcome would have been if A was set to 1 and 0, respectively, which are the so-called potential outcomes. Then $Y = Y^{(A)}$, and no unmeasured confounding states that

$$(Y^{(1)}, Y^{(0)}) \perp\!\!\!\perp A \mid W,$$

so that treatment assignment and how the subject would react to treatment are independent conditional on W. Positivity states that $\mathbb{P}(A = 1 \mid W = w)$ is uniformly bounded away from both one and zero as a function of w. In other words, for each possible realization of the covariates, there is a strictly positive probability both receiving and not receiving treatment. Under these assumptions, the conditions of a randomized controlled trial are satisfied for each possible realization of the covariates. Therefore, the conditional average treatment effect (CATE) becomes

$$\begin{aligned}\text{CATE} &:= \mathbb{E}[Y^{(1)} \mid W] - \mathbb{E}[Y^{(0)} \mid W] \\ &= \mathbb{E}[Y \mid W, A = 1] - \mathbb{E}[Y \mid W, A = 0].\end{aligned}$$

Another situation in which causal inference is possible is if the condition for a *regression discontinuity design* is satisfied, namely that there is a discontinuity in $w \mapsto \mathbb{P}(A = 1 \mid W = w)$, where W is now assumed to be one-dimensional. This creates a local randomized trial for subjects with covariates close to the discontinuity. If W is one-dimensional and the discontinuity is at w_0, then the conditional average treatment effect at w_0 is identified as

$$\begin{aligned}\text{CATE}(w_0) &:= \mathbb{E}[Y^{(1)} \mid W = w_0] - \mathbb{E}[Y^{(0)} \mid W = w_0] \\ &= \frac{\lim_{w \downarrow w_0} \mathbb{E}[Y \mid W = w] - \lim_{w \uparrow w_0} \mathbb{E}[Y \mid W = w]}{\lim_{w \downarrow w_0} \mathbb{P}(A = 1 \mid W = w) - \lim_{w \uparrow w_0} \mathbb{P}(A = 1 \mid W = w)}.\end{aligned}$$

An efficient estimation approach that applies to both situations with no unmeasured confounding and to regression discontinuity designs, and which allows the outcome to be affected by right-censoring, is proposed in Sandqvist (2024).

It becomes apparent from the above discussion that the data-generating mechanism has to admit a particular structure in order for causal inference to be possible. This may happen as a coincidental consequence of internal practices at the insurer. For example, if only subjects with an estimated risk over a certain threshold are targeted by a prevention initiative, one is perhaps in the setting of a regression discontinuity design. Similarly, if the assignment of treatment exclusively depends on a combination of covariates already available to the insurer and exogenous randomness, say how many claims were reported in the past month, no unmeasured confounding and positivity may hold, and causal inference is again possible.

However, rather than just relying on happy accidents and pure luck, we advise actuaries to be actively involved in the design of prevention initiatives and in such a way that their impact becomes quantifiable and hence optimizable. Actuaries have a unique opportunity in this regard due to their close connection to the management of their organization, their knowledge of the insurance coverages' properties and functionalities, and their statistical expertise not least for data subject to sampling effects.

2.6 Outlook

Our review confirms that disability insurance is an area of high practical complexity, not least in Denmark, but also with great potential for meaningful insurance mathematical innovation. To stimulate research and promote better practices, we have sought to illustrate and systematize these complexities; to review and assess the adequacy of current approaches; and to identify avenues for future research and development. We find that modeling the interplay with public benefits and the associated implications for product design constitute an underdeveloped and promising topic for future research. Additionally, there is a need for empirical studies related to the design and impact evaluation of prevention initiatives. Actuaries must continue to play a crucial role in this context, embracing another dimension to their role, while maintaining strong ethics and safeguarding fairness. We aptly recall the message of Barabas et al. (2018): The core ethical debate surrounding risk assessments is not simply one of bias or accuracy, but one of purpose: away from prediction and towards prevention.

Acknowledgments and declarations of interest

Oliver Lunding Sandqvist is an Industrial PhD student at PFA Pension with funding from the Innovation Fund Denmark under File No. 1044-00144B. Christian Furrer gratefully acknowledges support from Fynske Købstæders Fond. The authors would like to thank Mogens Steffensen for helpful comments on an earlier version of the manuscript. The expressed opinions are attributable solely to us and do not necessarily reflect the views of our past, current, and future employers.

Chapter 3

Transaction time models in multi-state life insurance

This chapter contains the manuscript *Buchardt, Furrer, and Sandqvist (2023)*.

Abstract

In life insurance contracts, benefits and premiums are typically paid contingent on the biometric state of the insured. Due to delays between the occurrence, reporting, and settlement of changes to the biometric state, the state process is not fully observable in real-time. This fact implies that the classic multi-state models for the biometric state of the insured are not able to describe the development of the policy in real-time, which encompasses handling of incurred-but-not-reported and reported-but-not-settled claims. We give a fundamental treatment of the problem in the setting of continuous-time multi-state life insurance by introducing a new class of models: transaction time models. The relation between the transaction time model and the classic model is studied and a result linking the present values in the two models is derived. The results and their practical implications are illustrated for disability coverages, where we obtain explicit expressions for the transaction time reserve in specific models.

Keywords: Prospective reserves; Disability insurance; Claims reserves; Valid and real-time; Piecewise deterministic processes

3.1 Introduction

The payments stipulated in life insurance contracts are usually an agreement on what payments are to be made for different possible outcomes of the biometric state of the insured (e.g. whether the insured is active, disabled, dead, etc.). For this reason, multi-state life insurance models take modeling of the biometric state of the insured as their starting point. The multi-state approach to life insurance dates back to at least Hoem (1969). Here, the prospective reserve is defined as the discounted probability-weighted future payments, which, as noted in Norberg (1991), corresponds to the expected present value of future payments given the information generated by the biometric state process. The introduction of an underlying stochastic process generating the payments introduces structure to the problem of predicting the cash flow at future points in time, due to the temporal dependencies of the process. This added structure of the payments is not in itself an assumption when the payments stipulated in the insurance contracts are formulated in terms of the biometric state of the insured. It is rather a way to introduce more a priori knowledge about the workings of the product into the mathematical model. All other things being equal, this makes the models more powerful.

Consequently, multi-state modeling seems a natural approach to modeling life insurance products. However, in the multi-state modeling literature, one also often assumes that the biometric state process generating the payments equals the process that generates the available information, see e.g. Norberg (1991), Buchardt et al. (2015), Djehiche & Löfdahl (2016), Bladt et al. (2020), and Christiansen (2021a) to name a few. This is rarely the case, since information about changes to the biometric state can be delayed or erroneous. A simple example of this phenomenon is the delay that occurs when an insured becomes disabled; it might take some time for the insured to report the event to the insurer. Between the occurrence of the disability and the time of reporting, the claim is an IBNR (Incurred-But-Not-Reported) claim. As long as the insured has not reported the disability, the insurer will continue believing that the insured is active. Hence, the information that the insurer has is different from the full information about the biometric state of the insured. To describe this phenomenon in more detail, and discuss how to approach reserving under the insurer's available information, we introduce the concepts of *valid time* and *transaction time* in the next section: Essentially, the valid time of an event is the time that it occurs, while the transaction time is the time that the event is registered in the insurers records.

It turns out that these concepts are also useful in clarifying the similarities and differences between life and non-life insurance products as well as between the models employed in the respective fields. The fact that payments in life insurance are deterministic functions of the biometric state process makes it so one does not have to estimate a separate distribution for the payment sizes; once the distribution of the state process is specified, the distribution of the payments follows. This is not the case in non-life insurance, and one therefore resorts to modeling the distribution of the observed payments directly. However, as will be explained, the biometric state process is a valid time object, while the observed payments are transaction time objects. This fact leads to key differences in the life and non-life insurance models. One such key difference is that it is more straightforward to formulate

IBNR and RBNS (Reported-But-Not-Settled) models in non-life insurance, as one can construct the RBNS model entirely in transaction time. The IBNR model may subsequently be constructed in two steps. First, one models the occurrence times of claims, which are valid time objects, as well as the corresponding reporting delays. Second, one leverages the RBNS model to find the expected cash flow of a claim conditional on the time of occurrence and the reporting delay. In life insurance models, one has to link the valid time payments to the transaction time concepts of IBNR and RBNS, and it is not obvious how to do this.

Our main contributions are: the introduction of the *basic bi-temporal structure assumptions* defined in Section 3.4, the derivation of Theorem 3.5.4 that links the present values in valid and transaction time, and an application of this theorem, namely Example 3.5.8. The first of these contributions establishes an explicit link between transaction and valid time processes. The second utilizes this link to obtain a tractable relation between the present values in valid and transaction time. To further obtain a tractable relation between the valid and transaction time reserves, more structure on how the transaction time information affects the distribution of the valid time process needs to be imposed. This is exactly what is explored in a simple example involving RBNS claims, culminating in Example 3.5.8, which constitutes the third main contribution. The example is kept simple for illustrative purposes, but our general framework also allows for the study of intricate examples that provide a more complete picture of IBNR and RBNS reserving. Such applications are the raison d'être of the framework.

The paper is structured as follows. In Section 3.2, the terms valid time and transaction time are given more precise definitions and discussed in the context of life insurance. An overview of the use of valid time and transaction time information in the insurance literature is provided, and similarities and differences between the situation in life insurance and non-life insurance are made explicit. The section ends by defining the class of piecewise deterministic processes, which constitute the basic building blocks for our model constructions. Section 3.3 constructs a model for the insurance contract in valid time similarly to how the classic life insurance multi-state models are constructed. Section 3.4 introduces the novel concept of a transaction time model corresponding to a valid time model, and this transaction time model is constructed. In Section 3.5, the valid and transaction time reserves are defined, and a result relating the transaction time present value to the valid time present value is derived. In a model for disability insurance where coverage depends on the origin of the disability, we show how this relation can be utilized to obtain a relation between the corresponding reserves. Finally, the dynamics of the transaction time reserve is derived and discussed.

3.2 Valid and transaction time

We now introduce the terms valid time and transaction time. These concepts are used to describe data that arises from a time-varying process. We outline how these types of data are currently being used in the life and non-life insurance literature. Subsequently, we introduce a class of stochastic processes which we use to model processes generating valid time and transaction time data.

Valid and transaction time data

The terms valid time and transaction time provide a natural terminology for describing information that is registered with delays and uncertainty. Valid time and transaction time are concepts stemming from the design of databases, specifically temporal databases, where time-varying information is recorded. The valid/transaction time taxonomy was developed in Snodgrass & Ahn (1985). There, valid time is defined as the time that an event occurs in reality, while transaction time is defined as the time when the data concerning the event was stored in the database. Hence, valid time is concerned with when events occurred (historical information), while transaction time is concerned with when events were observed (rollback information).

As noted by Snodgrass & Ahn (1985), an important difference between valid time and transaction time are the types of information updates that are permitted. A transaction time may be added to the database, but is never allowed to be changed after the fact due to the forward motion of time. In contrast, a valid time is always subject to change, since discrepancies between the history as it actually occurred and the representation of the history as stored in the database will often be detected after the fact. The authors argue that both valid time and transaction time are needed to fully capture time-varying behavior.

A database that contains both valid time and transaction time is called a *bi-temporal* database. Such a database supplies both historical and rollback information. Historical information e.g. *“Where was Taylor employed during 2010?”* is supplied by valid time, while rollback information e.g. *“In 2010, where did the database believe Taylor was employed?”* is supplied by transaction time. Since there may have been changes to the database after 2010, the answers to these questions may be different. The combination of valid time and transaction time supplies information on the form *“In 2015, where did the database believe Taylor was employed during 2010?”*.

To further clarify the concepts introduced above, a detailed example for total permanent disability insurance (or critical illness insurance) is given in Example 3.2.1, while Example 3.2.2 is devoted to disability insurance with coverage that depends on the origin of disability.

Example 3.2.1. (Bi-temporal insurance data: Total permanent disability insurance.)

Consider the following scenario: On 1/1/2020, Taylor buys a total permanent disability insurance effective immediately with a risk period of one year, which pays a sum b if they become disabled before the end of the risk period. For this, Taylor agrees to pay premiums at a rate π during the risk period while active. Taylor becomes disabled on 1/3/2020 and reports this to the insurer on 1/5/2020, two months later. On 1/6/2020, one month later, the insurer has finished processing the claim and awards Taylor disability benefits. Furthermore, Taylor is reimbursed for the premiums paid between 1/3/2020 and 1/6/2020.

If the insurer uses a bi-temporal database (valid time and transaction time), the database will at 1/6/2020 contain the following entries:

Name	State	Valid From	Valid Till	Entered	Superseded
Taylor	Active	1/1/2020	∞	1/1/2020	1/5/2020
Taylor	Active	1/1/2020	1/3/2020	1/5/2020	∞
Taylor	RBNS	1/3/2020	∞	1/5/2020	1/6/2020
Taylor	Disabled	1/3/2020	∞	1/6/2020	∞

Hence we see that the database records not only what happened in the 'real world', but also what was officially recorded at different times. Note that when it is not known when the information is valid till, the database by convention records the timestamp ∞. This is likewise the case when it is unknown when the entry will be superseded. Hence, to acquire the most recent belief about when events occurred, one would extract the rows where `Superseded` was ∞.

If the database was uni-temporal (valid time), the entries at 1/6/2020 would be:

Name	State	Valid From	Valid Till
Taylor	Active	1/1/2020	1/3/2020
Taylor	Disabled	1/3/2020	∞

Similarly, at 1/4/2020 the entries would be:

Name	State	Valid From	Valid Till
Taylor	Active	1/1/2020	∞

even though Taylor is already disabled at this time, due to the fact that this has not been reported to the insurer yet.

From these tables, we can see that, as described in Snodgrass & Ahn (1985), different information updates are permitted for a bi-temporal and a uni-temporal database. The uni-temporal database, in contrast to the bi-temporal database, only records what happened in the 'real world' based on the newest information. Previous records are modified or deleted. ∘

Example 3.2.2. (Bi-temporal insurance data: Disability insurance with different origins.)

Consider the following scenario: On 1/1/2020, Jessie buys a disability insurance effective immediately, that pays a rate b_{i_1} if they are affected by a work-related disability (WD) and a rate b_{i_2} if they are affected by a non-work-related disability (NWD). For this, Jessie agrees to pay the premium rate π while active. Jessie becomes disabled on 1/5/2020 and reports this to the insurer instantly. The insurer immediately evaluates the disability to have its origin outside of the workplace and therefore pays the rate b_{i_2} starting 1/5/2020. At 1/7/2020, the decision is reevaulated and it is concluded that the disability has its origin at the workplace. Consequently, there is a payment between the insurer and Jessie corresponding to the difference in rates between 1/5/2020 and 1/7/2020, and onward Jessie receives the rate b_{i_1}. Nothing else occurs before 1/1/2021.

If the insurer uses a bi-temporal database (valid time and transaction time), the database will at 1/1/2021 contain the following entries:

Name	State	Valid From	Valid Till	Entered	Superseded
Jessie	Active	1/1/2020	∞	1/1/2020	1/5/2020
Jessie	Active	1/1/2020	1/5/2020	1/5/2020	∞
Jessie	NWD	1/5/2020	∞	1/5/2020	1/7/2020
Jessie	WD	1/5/2020	∞	1/7/2020	∞

If the database was uni-temporal (valid time), the entries at 1/6/2020 would be:

Name	State	Valid From	Valid Till
Jessie	Active	1/1/2020	1/5/2020
Jessie	NWD	1/5/2020	∞

At 1/8/2020 the entries would be:

Name	State	Valid From	Valid Till
Jessie	Active	1/1/2020	1/5/2020
Jessie	WD	1/5/2020	∞

Similar to Example 3.2.1, we see a clear need for a bi-temporal database compared to a uni-temporal database. ∘

In Example 3.2.1, the claim is first an IBNR and later an RBNS claim, while in Example 3.2.2, the claim is an RBNS claim, but since reporting occurred with no delay, it is not an IBNR claim beforehand. In the following, we illustrate our methods and results on the latter example. Our methods are, however, not constrained to bi-temporal data on the above form, but may be applied to essentially any kind of bi-temporal insurance data.

Practitioners should be well-acquainted with bi-temporal insurance data. Bi-temporal data is important for internal use, as it is needed for reproducibility of statistical analyses when these are based on queries to a database. This is because reproducibility requires rollback information, since one has to recreate the information that the database had at a previous point in time. It also enables one to understand the difference between two otherwise identical analyses, performed at two different points in time. It is also important for external use, since auditors and regulatory authorities may inquire about financial reports from foregone years, making it important for insurers to be able to recreate the prerequisites that a given financial report was based on. As an example of this, Danish life insurance companies are required by law to publish all figures in the balance sheet and income statement of their financial reports for both the current and the previous year. Key figures have to be reported for the past five years. If prior financial reports have been affected by serious errors, the newest report has to publish figures for previous years as if the error had not been committed, so long it is practically feasible, cf. § 86 of Erhvervsministeriet (2015).

Valid and transaction time information in insurance

Inspired by the valid and transaction time taxonomy introduced above, and with a slight abuse of the terminology, we define a *valid time process* as a stochastic process that represents the true historical information. We use the notation X_s for the value of the valid time process at time s. With another slight abuse of the terminology, we define a *transaction time process* as a bi-temporal stochastic process that represents both historical and rollback information. We use the notation X_s^t for the value of the process at time s based on the observations available at time t. When referring to models based on valid time and transaction time processes, we use the terms *valid time model* and *transaction time model*, respectively. Current multi-state models for the biometric state of the insured are seen to be valid time models, as they model a process that describes when events occur without regard to when that information is observed by the insurer. However, in practice there is typically at least some delay in information concerning occurrence, reporting and settlement of claims. This necessitates additional model components, say in the form of models for IBNR and RBNS claims. Such model components are called claims reserving models or simply IBNR and RBNS models.

Claims reserving based on individual claims data have been subject to much study in non-life insurance, see e.g. Norberg (1993, 1999), Haastrup & Arjas (1996), Antonio & Plat (2014), Badescu et al. (2016, 2019), Lopez et al. (2019), Bischofberger et al. (2020), Delong et al. (2021), Crevecoeur et al. (2022a,b), and Okine et al. (2022). Early research, including Norberg (1993, 1999), focuses on joint modeling of all aspects of a claim and the subsequent computation of relevant conditional expectations of future cash flows as high-dimensional integrals with respect to the joint distribution. This may pose significant statistical and numerical challenges, and initially only limited attention was given to practical implementation. Recent research, including Crevecoeur et al. (2022a,b) and Okine et al. (2022), instead focuses on how the expected future cash flows depend on previous observations, typically payment times and payment sizes, and the corresponding factorization of the joint distribution into conditional and marginal parts. In particular, this improves interpretability and readily allows for dynamic reserving where current information is incorporated into the best estimate of future liabilities. This research has hitherto largely consisted of proposing statistical procedures and performing data-driven investigations. Regarding the latter, one attempts to identify the best predictive models by exploring which covariates are advantageous to include (thereby also determining the degree of individualization/collectivization) as well as exploring which statistical models to apply (e.g. parametric models such as GLMs, non-parametric models such as empirical distributions, or machine learning methods such as neural networks).

The primary reason that the life insurance literature has not been similarly occupied with finding suitable statistical models is, as explained in Section 3.1, that the payments stipulated in life insurance contracts are specified in terms of an underlying valid time state process. This implies that the conditional distribution of future jump times given previous observations can be obtained by estimating the distribution of the valid time state process, and a model for the conditional distribution of future payment sizes then follows automatically. Since the state process may almost always be formulated as a piecewise deterministic process, a point which we discuss in more detail later, the estimation may be performed using standard methods from survival and event history analysis. The state process formulation also allows for more explanatory models compared with the purely predictive models of non-life insurance that result from modeling the transaction time payment process directly.

The problem with the current approach in life insurance of ignoring the transaction time state process is that the alternative state process, namely the valid time state process, is not directly observable. In practice, one partly accounts for this via improvised IBNR corrections and RBNS corrections (on an aggregate level). Compared to the vast literature on IBNR and RBNS models in non-life insurance, the corresponding problem of claims reserving in life insurance has received limited attention. In the following, we seek to amend this.

Since the payments stipulated in a life insurance contract are generated by a valid time process, while the actually observed data is generated by a transaction time process, it becomes essential to establish an in-depth understanding of the relation between valid and transaction time processes in life insurance. The main purpose of this paper is to develop a conceptual and mathematical framework where this relation can be formulated and explored.

Piecewise deterministic processes

In order to achieve as much generality as possible, we take *piecewise deterministic processes* (PDPs) as the starting point for our models. The data generated by a stochastic process and recorded in a database only uniquely determines the path of the stochastic process if the stochastic process is a PDP. Informally, these are processes that have finitely many jumps on finite time intervals and which develop deterministically between the random jump times. This is because the value of the process is only recorded in the database at certain discrete times (e.g. daily, monthly, or at jump times), from which the entire path of the process must be inferred. If the data stems from e.g. a Brownian motion, which is not piecewise deterministic, then the data can only provide an approximation of the path taken, and the gaps between recorded values must be filled using some algorithm, e.g. modeling the process to be linearly evolving between recorded values. In this sense, PDPs provide the most general class of processes for which valid time and transaction time processes can be constructed.

We now define piecewise deterministic processes following Subsection 3.3 in Jacobsen (2006). To keep the exposition from becoming unnecessarily technical, we omit some mathematical details such as the exact form of measurability for certain functions. The interested reader may consult Section 3 of Jacobsen (2006) for an explicit construction of the basic measurable spaces. In general, further mappings from these spaces are taken to have the image of the mappings equipped with the pushforward σ-algebra as the codomain space. Let the background probability space be denoted by $(\Omega, \mathbb{F}, \mathbb{P})$. Let $(E, \mathcal{E})$ be a measurable space called the *mark space*. Introduce the *irrelevant* mark ∇ denoting the mark of a jump that does not occur in finite time, and set $\overline{E} = E \cup \{\nabla\}$.

Definition 3.2.3. *(Simple point processes.)*

A simple point process (SPP) is a sequence $\mathcal{T} = (T_n)_{n\in\mathbb{N}}$ *of* $[0,\infty]$*-valued random variables defined on* $(\Omega, \mathbb{F}, \mathbb{P})$ *such that*

(i) $\mathbb{P}(0 < T_1 \leq T_2 \leq \dots) = 1$

(ii) $\mathbb{P}(T_n < T_{n+1}, T_n < \infty) = \mathbb{P}(T_n < \infty)$

(iii) $\mathbb{P}(\lim_{n\to\infty} T_n = \infty) = 1.$ △

If condition (iii) is removed, one obtains the class of simple point processes allowing for *explosion*. In this paper, we limit the study to processes without explosion.

Definition 3.2.4. *(Marked point processes.)*

A marked point process (MPP) with mark space E is a double-sequence $(\mathcal{T}, \mathcal{Y}) = ((T_n)_{n\in\mathbb{N}}, (Y_n)_{n\in\mathbb{N}})$ of $(0,\infty]$-valued random variables T_n and $\overline{E}$-valued random variable Y_n defined on $(\Omega, \mathbb{F}, \mathbb{P})$ such that $\mathcal{T} = (T_n)_{n\in\mathbb{N}}$ is an SPP and such that

(i) $\mathbb{P}(Y_n \in E, T_n < \infty) = \mathbb{P}(T_n < \infty)$

(ii) $\mathbb{P}(Y_n = \nabla, T_n = \infty) = \mathbb{P}(T_n = \infty)$. △

For a given MPP $(\mathcal{T}, \mathcal{Y})$, define

$$\langle t \rangle = \sum_{n=1}^{\infty} 1_{(T_n \leq t)}$$

being the number of events in the time interval $[0, t]$, and define

$$H_t = (T_1, ..., T_{\langle t \rangle}; Y_1, ..., Y_{\langle t \rangle})$$

being the jump times and marks observed up until and including time t. We refer to H_t as the *MPP history* at time t.

Definition 3.2.5. *(Piecewise deterministic process.)*

A piecewise deterministic process with state space $(E, \mathcal{E})$ is an E-valued stochastic process X satisfying

$$X_t = f^{\langle t \rangle}_{H_t|x_0}(t),$$

where $X_0 = x_0$ is non-random, and for every $n \in \mathbb{N}_0$, $f^n_{h_n|x_0}(t)$ is a measurable E-valued function of $h_n = (t_1, ..., t_n; y_1, ..., y_n)$ with $t_n < \infty$, of $t \geq t_n$, and of x_0, satisfying the conditions

$$f^n_{h_n|x_0}(t_n) = y_n$$

for all h_n and $f^0_{|x_0}(0) = x_0$. △

To ensure the existence of relevant conditional distributions going forward, we henceforth assume that the mark space $(E, \mathcal{E})$ is a Borel space. We refer to the functions $f^n_{h_n|x_0}$ as *evolution functions* of X. The explicit dependence on n is standard and serves to highlight that the functional relation may depend on the cardinality of h_n. However, to ease notation, we henceforth suppress x_0 and n, writing for instance simply f_{h_n}. It is easily seen that the class of piecewise deterministic processes encompasses the usual choices for the biometric state process, cf. Example 3.2.6, 3.2.7, 3.2.8, and 3.2.9.

Example 3.2.6. (Pure jump process.)

Let $(E, \mathcal{E})$ be a Borel space, and let X be a pure jump process taking values in E which models the biometric state process of the insured. Assume that the initial state of X is fixed, and denote this initial state by x_0. A pure jump process is here taken to mean a càdlàg stochastic process with only finitely many jumps in any finite time interval satisfying

$$X_t = \sum_{0<s\leq t} \Delta X_s$$

for $\Delta X_s = X_s - X_{s-}$. We define an MPP $(\mathcal{T}, \mathcal{Y}) = ((T_n)_{n\in\mathbb{N}}, (Y_n)_{n\in\mathbb{N}})$ from X by letting T_n be the time of the n'th jump of X and setting $Y_n = X_{T_n}$. It is then easy to show that X can be reconstructed from the MPP. To do this, let $f_{h_n}(t) = y_n$, and let H be the MPP history of $(\mathcal{T}, \mathcal{Y})$. Then

$$f_{H_t}(t) = Y_{\langle t\rangle} = X_{T_{\langle t\rangle}} = X_t,$$

so X is a PDP. ∘

Example 3.2.7. (Pure jump Markov process.)

Following the previous example, let $\mathcal{F}_t^X = \sigma(X_s,\ 0 \leq s \leq t)$ be the filtration generated by X. For computational tractablility, one often assumes that X is a *Markov process*, meaning for every $s \leq t$ and $C \in \mathcal{E}$:

$$\mathbb{P}(X_t \in C \mid \mathcal{F}_s^X) = \mathbb{P}(X_t \in C \mid X_s)$$

$\mathbb{P}$-a.s., so that the behavior of the process at a future time point is only dependent on the past behavior of the process through the current state of the process. This is called the *Markov property*. Pure jump Markov processes X are a generalization of the usual choices of the state process Z found in the multi-state life insurance literature, as is shown in Example 3.2.8 and Example 3.2.9. ∘

Example 3.2.8. (Continuous-time Markov chain.)

A continuous-time Markov chain Z with fixed initial state is defined as a piecewise constant Markov process, see e.g. Chapter 2 of Norris (1998) or Section 7.2 of Jacobsen (2006). In the life insurance literature, the Markov chain is usually assumed to take values in a finite state space, see e.g. Norberg (1991). A continuous-time Markov chain Z on a finite state space $\{1, 2, ..., J\}$ for $J \in \mathbb{N}$ can thus be constructed from a pure jump Markov process X with $E = \{1, 2, ..., J\}$ by simply setting $Z = X$. ∘

Example 3.2.9. (Continuous-time semi-Markov process.)

In many applications, the most recent jump time contains valuable information and hence it may be necessary to include it as a coordinate of X in order to ensure that the Markov property of X holds. For a pure jump process Z, define W as the time of the last jump

$$W_t = \sup\{0 \leq s \leq t : Z_s \neq Z_t\}.$$

A continuous-time semi-Markov process Z is defined as a pure jump process on a finite state space with the property that (Z, W), equivalently (Z, U) with

$$U_t = t - W_t$$

the duration of sojourn in the current state, is a Markov process, see e.g. Section 2D in Helwich (2008). Semi-Markov models were introduced to life insurance independently by Janssen and Hoem, see Janssen (1966) and Hoem (1972). In the life insurance literature, the parameterization (Z, U) is more common than (Z, W).

For $K \in \mathbb{N}$ define the projection functions $\pi_k : \mathbb{R}^K \mapsto \mathbb{R}$ via $\pi_k(x_1, x_2, ..., x_K) = x_k$. If X is a pure jump Markov process with $E = \{1, 2, ..., J\} \times [0, \infty)$ that satisfies $\pi_2 X_t = T_{\langle t \rangle}$, then a continuous-time semi-Markov process Z may be constructed from X by setting $Z = \pi_1 X$. ◦

3.3 Valid time model

We now introduce the classic multi-state life insurance models. These are valid time models, cf. the discussion in Section 3.2. The valid time setup described below is essentially standard in the multi-state life insurance literature, although it is usually only formulated for Markov or semi-Markov processes, see the classics Hoem (1969), Hoem (1972), and Norberg (1991).

State process

Let the valid time state process $(X_t)_{t \geq 0}$ be an $\mathbb{R}^d$-valued stochastic process for $d \in \mathbb{N}$. We assume that X is a PDP with fixed initial state x_0. The jump times are denoted by τ_n. Let

$$\langle\!\langle t \rangle\!\rangle = \sum_{n=1}^{\infty} 1_{(\tau_n \leq t)}$$

denote the number of jumps by time t and denote by

$$H_t = (\tau_1, ..., \tau_{\langle\!\langle t \rangle\!\rangle}; X_{\tau_1}, ..., X_{\tau_{\langle\!\langle t \rangle\!\rangle}})$$

the MPP history of the process X at time t. The symbols τ_n and $\langle\!\langle t \rangle\!\rangle$ are used here instead of T_n and $\langle t \rangle$ introduced in the PDP section as the latter are reserved for the transaction time process introduced in Section 3.4. Further, let f_{h_n} be the evolution functions of X, so that

$$X_t = f_{H_t}(t).$$

The information generated by complete observation of the valid time process is given by the filtration $\mathcal{F}_t^X = \sigma(X_s, 0 \leq s \leq t)$.

Cash flow

We now define the valid time cash flow that represents the contractual payments. For later uses, especially to define the transaction time cash flow and to formulate links between present values in valid and transaction time, we need the valid time cash flow to be decoupled from the valid time state process. Denote by $\mathcal{A} \subseteq 2^{[0,\infty)}$ some sufficiently regular collection of sets. For the purposes of this paper, it is sufficient that $\mathcal{A}$ contains the intervals on $[0, \infty)$. Write x_A for $(x_s)_{s\in A}$ with $A \in \mathcal{A}$. We restrict our attention to x_A for which there exist $(x_s)_{s\in A^c}$ such that $(x_s)_{s\geq 0}$ lies in the image of X; here A^c denotes the complement of A. Similarly, write X_A for $(X_s)_{s\in A}$. Assume the existence of measurable functions

$$(x_A, t) \mapsto B(x_A, t) \in \mathbb{R}$$

for $t \geq 0$. We interpret $B(x_A, t)$ as the payments generated by the path x on $A \cap [0, t]$. Assume that $t \mapsto B(x_A, t)$ is a càdlàg finite variation function for any x_A, so that the measures $B(x_A, \mathrm{d}t)$ are well-defined. We further assume that the composition $B(X_A, t)$ is incrementally adapted to $\mathcal{F}^X$, meaning that $B(X_A, t) - B(X_A, s)$ is $\sigma(X_v, v \in (s, t])$-measurable for any interval $(s, t] \subseteq [0, \infty)$, cf. Definition 2.1 in Christiansen (2021b). For shorthand, we write $B(\mathrm{d}t) = B(X_{[0,\infty)}, \mathrm{d}t)$. The assumption of incremental adaptedness states that the aggregated payments in $(s, t]$ only depends on the path of X on $(s, t]$. The property incremental adaptedness does not have a critical technical function in this paper, but it clarifies the role of the state process X as the process that determines the contractual payments at a given instance in time. Therefore, we may discuss changes in information and payments through changes to X, which allows for practical interpretations that are more intuitive. From a purely mathematical point of view, one could also have based the analysis on the underlying MPP. We name $(B(t))_{t\geq 0}$ the accumulated cash flow in valid time. Note the use of the symbol B for two different objects, namely the stochastic process $t \mapsto B(t)$ and the deterministic function $t \mapsto B(x_A, t)$. It should always be clear from the context which object we are referring to.

To allow for the valuation of cash flows, we need the time value of money. A detailed treatment of this financial constituent of the model may be found in Norberg (1990). Let $t \mapsto \kappa(t)$ be some deterministic strictly positive càdlàg *accumulation function* with initial value $\kappa(0) = 1$. The corresponding *discount function* is $t \mapsto \frac{1}{\kappa(t)}$. We let $x_A \mapsto B^\circ(x_A)$ be the time 0 value of the accumulated payments for the path x_A, i.e. we set

$$B^\circ(x_A) = \int_{(0,\infty)} \frac{1}{\kappa(v)} B(x_A, \mathrm{d}v),$$

presupposing that this object exists (is finite).

Remark 3.3.1. (Cash flow terminology.)

In the life insurance literature, the stochastic process $(B(t))_{t\geq 0}$ defined above is sometimes also referred to as the payment function, the stream of net payments, the payment process, or the stochastic cash flow, see e.g. Norberg (1991) and Buchardt et al. (2015). We use the terminology that $B(t)$ is the accumulated cash flow at time t while the stochastic measure $B(\mathrm{d}t)$ is the cash flow. The latter should not be confused with the expected cash flow $A_x(t, \mathrm{d}s)$ defined by $A_x(t,s) = \mathbb{E}[B(s) - B(t) \mid X(t) = x]$, $s \geq t$. In the literature, the expected cash flow is sometimes also ambiguously referred to as the cash flow. ▽

Example 3.3.2. (Cash flow for the usual choices of state processes.)

For a set A on the form $[0,v)$, $[0,v]$, or $[0,\infty)$ for $v \geq 0$ and a pure jump Markov process X on a finite state space $E = \{1, 2, ..., J\}$, see also Example 3.2.8, one usually specifies the payments as

$$B(x_A, \mathrm{d}t) = \sum_{j=1}^{J} 1_A(t) 1_{(x_{t-}=j)} B_j(\mathrm{d}t) + \sum_{\substack{j,k=1 \\ j\neq k}}^{J} 1_A(t) b_{jk}(t) n_{jk}(x_A, \mathrm{d}t),$$

where $n_{jk}(x_A, t) = \#\{s \in [0,t] \cap A : x_{s-} = j, x_s = k\}$, while $t \mapsto B_j(t)$ are càdlàg finite variation functions modeling sojourn payments and $t \mapsto b_{jk}(t)$ are finite-valued Borel-measurable functions modeling transition payments. In this case,

$$B(\mathrm{d}t) = \sum_{j=1}^{J} 1_{(X_{t-}=j)} B_j(\mathrm{d}t) + \sum_{\substack{j,k=1 \\ j\neq k}}^{J} b_{jk}(t) N_{jk}(\mathrm{d}t)$$

for $N_{jk}(t) = n_{jk}(X_{[0,\infty)}, t)$. In the semi-Markov case of Example 3.2.9, where $X = (Z, W)$ is a pure Markov jump process (with W the time of the last jump), one usually specifies the payments as

$$B(x_A, \mathrm{d}t) = \sum_{j=1}^{J} 1_A(t) 1_{(z_{t-}=j)} B_{j,w_{t-}}(\mathrm{d}t) + \sum_{\substack{j,k=1 \\ j\neq k}}^{J} 1_A(t) b_{(j,w_{t-})(k,t)}(t) n_{jk}(z_A, \mathrm{d}t),$$

where $x_s = (z_s, w_s)$, while $t \mapsto B_{j,w}(t)$ and $t \mapsto b_{(j,w)(k,t)}(t)$ satisfy the same regularity conditions as in the Markov case. Then

$$B(\mathrm{d}t) = \sum_{j=1}^{J} 1_{(Z_{t-}=j)} B_{j,W_{t-}}(\mathrm{d}t) + \sum_{\substack{j,k=1 \\ j\neq k}}^{J} b_{(j,W_{t-})(k,t)}(t) N_{jk}(\mathrm{d}t)$$

for $N_{jk}(t) = n_{jk}(Z_{[0,\infty)}, t)$. ○

Example 3.3.3. (Continuous compound interest.)

Under continuous compound interest with force of interest $t \mapsto r(t)$, a deposit of one unit currency in a savings account at time 0 has at time t accumulated to

$$\kappa(t) = \exp\left(\int_{(0,t]} r(s)\,\mathrm{d}s\right),$$

see e.g. Norberg (1990). ○

Example 3.3.4. (Disability insurance with different origins: Valid time model.)

We construct a valid time model for the product described in Example 3.2.2. The state of the insured is modeled as a semi-Markov process $X = (Y, U)$, cf. Example 3.2.9, where Y takes values in the state space depicted in Figure 3.1.

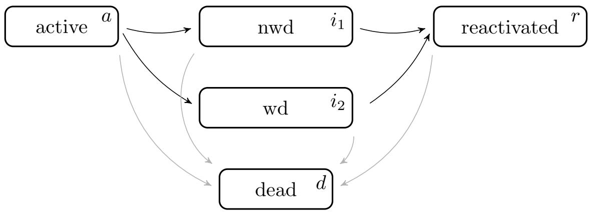

Figure 3.1: *The Y-component of the biometric state process X takes values in $\{a, i_1, i_2, r, d\}$. The absence of an arrow between states indicates that a direct jump between these states is impossible. To reduce clutter, the arrows into the dead state are made semi-transparent.*

The cash flow is assumed to consist of a premium rate $\pi < 0$ while active and a disability rate $b_{i_k} > 0$ while disabled in state i_k:

$$B(x_A, \mathrm{d}t) = 1_A(t)\pi 1_{(y_{t-}=a)}\,\mathrm{d}t + 1_A(t)\sum_{k=1}^{2} b_{i_k} 1_{(y_{t-}=i_k)}\,\mathrm{d}t,$$

so that

$$B(\mathrm{d}t) = \pi 1_{(Y_{t-}=a)}\,\mathrm{d}t + \sum_{k=1}^{2} b_{i_k} 1_{(Y_{t-}=i_k)}\,\mathrm{d}t.$$

The deterministic valid time cash flow $B(x_A, \mathrm{d}t)$ is used to specify the transaction time cash flow in Example 3.4.3. Since there are no payments in the reactivated or dead states, the semi-Markov assumption is, for valuation purposes, actually not a restriction. ○

3.4 Transaction time model

We now introduce the transaction time models corresponding to the valid time models introduced in Section 3.3. The transaction time setup described below is novel, and it allows for cash flows that are tailored to describe the payments that occur in real-time, like the ones that would result from the cases described in Examples 3.2.1 and 3.2.2.

State process

Let $(Z_t)_{t\geq 0}$ be an $\mathbb{R}^q$-valued pure jump process for $q \in \mathbb{N}$. We think of Z as describing the claim settlement process of the policy, which is observed in real-time by the insurer. The process Z jumps to another state when new information is made available to the insurer, for example through communication with the insured or due to internal decisions from the insurer. The jump times are denoted by T_n. We write

$$\langle t \rangle = \sum_{n=1}^{\infty} 1_{(T_n \leq t)}$$

for the number of jumps of Z that have happened at time t. We further introduce a doubly indexed stochastic process H_s^t for $0 \leq s \leq t$, which we name the *transaction time MPP history corresponding to* H_s. The idea is to interpret H_s^t as the value of H_s based on the transaction time information available at time t. We extend the definition to $s > t$ by letting $H_s^t = H_t^t$, akin to how in Example 3.2.1, the most recent history is set to be valid until ∞. The role of Z is to contain transaction time information that may not yet have resulted in a change to the transaction time MPP. An example could be a reported claim occurrence that has yet to be awarded or rejected. We briefly note that the main example contained of this paper is sufficiently simple that one could have omitted Z in the formulation and simply used H_s^t. This is, however, not possible in the general case. To imbue H_s^t and Z with the desired interpretations outlined above, we specify the dependencies to the corresponding valid time process:

(i) We only allow changes to the transaction time MPP history corresponding to H_s to occur at jumps of Z, so the process Z is the driver of new transaction time information arriving. This is formalized as

$$H_s^t = H_s^{T_{\langle t \rangle}}$$

for all $t, s \geq 0$.

(ii) We assume that there is a finite time after which no new information arrives (e.g. the time of death of the insured). This is formalized as

$$Z_t = Z_\eta$$

for all $t \geq \eta$ with $\mathbb{P}(\eta < \infty) = 1$. This condition is satisfied in any practical application. We name η the *absorption time.*

(iii) When there are no future changes to the transaction time MPP history corresponding to H_s, the observations are taken to be the true historical information. These observations constitute the finalized timeline that the insurer will observe. This is formalized as

$$H_s^t = H_s$$

for all $t \geq \eta$ and $s \geq 0$. This is the fundamental link between the valid time and transaction time models.

We name these assumptions the *basic bi-temporal structure assumptions.*

The transaction time state process is then defined as the doubly indexed stochastic process X_s^t satisfying

$$X_s^t = f_{H_s^t}(s),$$

with the interpretation that X_s^t is the value of X_s based on the available transaction time information at time t. At time t, the insurer has observed H_s^s and Z_s for all $s \leq t$. The insurer's available information is therefore generated by a process $(\mathcal{Z}_t)_{t \geq 0}$ given by

$$t \mapsto \mathcal{Z}_t = \big(Z_t, H_t^t\big),$$

and we define the *transaction time information* to be the filtration $\mathcal{F}_t^{\mathcal{Z}} = \sigma(\mathcal{Z}_s, 0 \leq s \leq t)$. Note that $\mathcal{Z}$ is a PDP by Example 3.2.6, since it is a pure jump process which furthermore can be embedded into the Borel space $(\mathbb{R}^\infty, \mathbb{B}(\mathbb{R}^\infty))$.

Remark 3.4.1. (Transaction and valid time filtrations.)

Note that in general, it holds that $\mathcal{F}_t^{\mathcal{Z}} \not\subseteq \mathcal{F}_t^X$ and $\mathcal{F}_t^X \not\subseteq \mathcal{F}_t^{\mathcal{Z}}$. This corresponds to $(\mathcal{Z}_s)_{0 \leq s \leq t}$ not being known from $(X_s)_{0 \leq s \leq t}$ and vice versa. In other words, the same valid time realizations can stem from different transaction time realizations, and transaction time realizations in a period do not generally determine the valid time realizations in that same period, due to the possibility of new information arriving later. ▽

Cash flow

We now define the transaction time cash flow that represents the payments occurring in real-time. We denote the accumulated cash flow in transaction time by $(\mathcal{B}(t))_{t \geq 0}$ and define it as

$$\mathcal{B}(\mathrm{d}t) = B(X_{[0,\infty)}^t, \mathrm{d}t) + \mathrm{d}\bigg(\sum_{0<s\leq t} \kappa(s) \Big(B^\circ(X_{[0,s)}^s) - B^\circ(X_{[0,s)}^{s-}) \Big) \bigg), \quad \mathcal{B}(0) = B(0).$$

The first term is well-defined as it may be written as $\sum_{n=0}^\infty 1_{(\langle t \rangle = n)} B(X_{[0,\infty)}^{T_n}, \mathrm{d}t)$ with the convention $T_0 = 0$. Hence the cash flow in transaction time consists of running payments similar to the ones in the valid time model, but here determined by X^t instead of X, as well as lump sum payments when some past X values are changed based on the newest transaction time information. The lump sum payments are commonly known as *backpay*. Backpay makes the accumulated historic payments congruent with the latest MPP history. Furthermore, and closely connected to the principle of no arbitrage, the backpay is accumulated to the time of payout according to κ, so that the insured is no better or worse off than if the payment had not been delayed and had been deposited in a savings account immediately after payout. Note also that the transaction time cash flow is incrementally adapted to $\mathcal{F}^{\mathcal{Z}}$, meaning that $\mathcal{Z}$ is the process that determines the transaction time payments at a given instance in time.

Note that since the first term of $\mathcal{B}$ is evaluated in X^t and not X^{t-}, the payment at time t based on the most recent information is included in the first term and should not be included in the backpay, which is why the right endpoint is excluded in the interval $[0, s)$ that appears in the second term. Note also that $X^s_{[0,s)} = X^{s-}_{[0,s)}$ unless Z jumps at time s. This implies that backpay can only be paid at jumps of Z.

Remark 3.4.2. (Cash flow modeling in non-life insurance.)

As described in Section 3.2, using non-life insurance methods, one would model the expected future payments arising from the real-time cash flow $\mathcal{B}$ by disregarding the state process X^t_s and the structure it imposes on the cash flow, and instead find suitable statistical models that predict $\mathcal{B}(\mathrm{d}t)$ directly. $\triangledown$

In Example 3.4.3, we consider the situation where it could be difficult to determine the origin or cause of a disability. This results in retroactive changes, which we can describe using a transaction time model.

Example 3.4.3. (Disability insurance with different origins: Transaction time model.)

We assume a simple transaction time model, namely that the Z-process takes values in the state space illustrated in Figure 3.2.

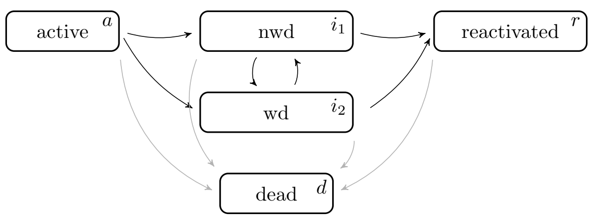

Figure 3.2: *The Z-component of the transaction time process $\mathcal{Z}$ takes values in $\{a, i_1, i_2, r, d\}$. To reduce clutter, the arrows into the dead state are made semi-transparent.*

We further assume that the transitions in the transaction time model equal those in the valid time model, but that the origin of disability i_k is not necessarily correctly identified when the claim is reported. The presumed origin may change while the insured is still disabled, but not after the insured has reactivated or died.

Let $\theta_j = \inf\{s \geq 0 : Z_s = j\}$ be first hitting times, let $\theta_i = \min\{\theta_{i_1}, \theta_{i_2}\}$ be the first hitting time of a disabled state, and write $i(t)$ for the most recently visited disabled state before time t. Define $x \wedge y = \min\{x, y\}$. By the specification of the model, H_s^t contains (θ_d, d) on $(\theta_d \leq t \wedge s)$, it contains (θ_r, r) on $(\theta_r \leq t \wedge s)$, and it contains (θ_i, i_k) on $(\theta_i \leq t \wedge s, i(t) = i_k)$. Finally, let $N_{jk}^Z(t)$ be the number of transitions from state j to state k for the process Z on $[0, t]$. The transaction time cash flow then reads

$$\begin{aligned}\mathcal{B}(\mathrm{d}t) &= \pi 1_{(Y_{t-}^t = a)}\,\mathrm{d}t + \sum_{k=1}^{2} b_{i_k} 1_{(Y_{t-}^t = i_k)}\,\mathrm{d}t \\ &\quad + \sum_{\substack{j,k\in\{1,2\}\\ j\neq k}} \left(\int_{[\theta_i, t)} \frac{\kappa(t)}{\kappa(s)} (b_{i_k} - b_{i_j})\,\mathrm{d}s \right) N_{i_j i_k}^Z(\mathrm{d}t).\end{aligned}$$

$\circ$

3.5 Reserving

Having defined the state process and cash flow in the valid and transaction time models, we are now in a position to define the prospective present value and prospective reserve in the respective models. The main result of the paper, Theorem 3.5.4, linking the present values in the two models, is deduced and discussed. Finally, the dynamics of the transaction time reserve is derived and its role in model validation is briefly considered.

Valid time reserve

Let the classic valid time prospective present value $(P(t))_{t\geq 0}$ be defined by

$$P(t) = \int_{(t,\infty)} \frac{\kappa(t)}{\kappa(s)} B(\mathrm{d}s).$$

We assume that $P(t)$ is well-defined and belongs to $L^1(\Omega, \mathbb{F}, \mathbb{P})$ for every $t \geq 0$. Then we can define the corresponding expected present value by

$$V(t) = \mathbb{E}\big[P(t) \mid \mathcal{F}_t^X\big] \tag{3.5.1}$$

for any $t \geq 0$. As a function of t, this is also known as the *prospective reserve* in valid time. We consider a version of V presumed to be $\mathcal{F}^X$-adapted and $\mathbb{P}$-a.s. right-continuous.

Remark 3.5.1. (Reserve for Markov state process.)

Note that $P(t)$ is $\sigma(X_s, t \leq s < \infty)$-measurable, so if X is a Markov-process, it follows that

$$V(t) = \mathbb{E}[P(t) \mid X_t]. \qquad \nabla$$

For random variables Y_1 and Y_2, we say that $Y_1 = Y_2$ on $A \in \mathbb{F}$ if $Y_1(\omega) = Y_2(\omega)$ for $\mathbb{P}$-almost all $\omega \in A$. We immediately have the following representation of the present value:

Proposition 3.5.2. *(Valid time present value.)*

For $t \geq 0$, it holds that

$$P(t) = \kappa(t)\Big(B^\circ(X^\eta_{[0,\infty)}) - B^\circ(X^\eta_{[0,t]})\Big)$$

on the event $(\eta < \infty)$.

Proof. Using the definitions introduced above, we find that

$$\begin{aligned}
P(t) &= \kappa(t) \int_{(t,\infty)} \frac{1}{\kappa(s)} B(\mathrm{d}s) \\
&= \kappa(t)\bigg(\int_{[0,\infty)} \frac{1}{\kappa(s)} B(\mathrm{d}s) - \int_{[0,t]} \frac{1}{\kappa(s)} B(\mathrm{d}s)\bigg) \\
&= \kappa(t)\bigg(\int_{[0,\infty)} \frac{1}{\kappa(s)} B(X_{[0,\infty)}, \mathrm{d}s) - \int_{[0,t]} \frac{1}{\kappa(s)} B(X_{[0,\infty)}, \mathrm{d}s)\bigg) \\
&= \kappa(t)\Big(B^\circ(X_{[0,\infty)}) - B^\circ(X_{[0,t]})\Big).
\end{aligned}$$

Now on the event $(\eta < \infty)$ we have that $X_s = X^\eta_s$ for all $s \geq 0$, from which the result follows. □

Transaction time reserve

The transaction time prospective present value $(\mathcal{P}(t))_{t\geq 0}$ is defined as

$$\mathcal{P}(t) = \int_{(t,\infty)} \frac{\kappa(t)}{\kappa(s)} \mathcal{B}(\mathrm{d}s).$$

We assume $\mathcal{P}(t)$ is well-defined and belongs to $L^1(\Omega, \mathbb{F}, \mathbb{P})$ for any $t \geq 0$. Let $\mathcal{V}$ be the corresponding expected present value in the transaction time model

$$\mathcal{V}(t) = \mathbb{E}\big[\mathcal{P}(t) \mid \mathcal{F}^{\mathcal{Z}}_t\big]. \tag{3.5.2}$$

Just as for V, we consider a version of $\mathcal{V}$ presumed to be $\mathcal{F}^{\mathcal{Z}}$-adapted and $\mathbb{P}$-a.s. right-continuous. Since the backpay $B^\circ(X^s_{[0,s)}) - B^\circ(X^{s-}_{[0,s)})$ telescopes, we also get the following representation of the transaction time present value $\mathcal{P}$:

Proposition 3.5.3. *(Transaction time present value.)*
For $t \geq 0$, it holds that

$$\mathcal{P}(t) = \kappa(t)\Big(B^\circ(X^\eta_{[0,\infty)}) - B^\circ(X^t_{[0,t]})\Big)$$

on the event $(\eta < \infty)$.

Proof. Note that the equality is trivially satisfied on $(t \geq \eta)$ by Proposition 3.5.2 and the observation that $P(t) = \mathcal{P}(t)$ on $(t \geq \eta)$ since $X^t = X^\eta$ implies that the running payments agree and there is no backpay after time t. Hence, what remains to be shown is that the equality also holds on $(t < \eta)$. We therefore let all the remaining calculations be on $(t < \eta)$. We re-index $T_k = T_{k+\langle t\rangle}$ for $k \in \mathbb{N}$, so T_k now refers to the k'th jump of Z after time t. Let n_η be a random variable $\mathbb{P}$-a.s. taking values in $\mathbb{N}$, with the defining feature that $T_{n_\eta} = \eta$, so that n_η is the number of jumps after time t until Z is absorbed. For notational convenience, introduce

$$\mathcal{P}^\circ(t) = \frac{1}{\kappa(t)}\mathcal{P}(t)$$

and $\beta_s = B^\circ(X^s_{[0,s)}) - B^\circ(X^{s-}_{[0,s)})$. We then have on the event $(\eta < \infty)$:

$$\begin{aligned}
\mathcal{P}^\circ(t) &= \int_{(t,\infty)} \frac{1}{\kappa(s)} \sum_{n=0}^{\infty} 1_{(\langle s\rangle = n)} B(X^{T_n}_{[0,\infty)}, \mathrm{d}s) + \sum_{t<s<\infty} \beta_s \\
&= \int_{(t,T_1)} \frac{1}{\kappa(s)} B(X^t_{[0,\infty)}, \mathrm{d}s) + \beta_{T_1} \\
&\quad + \sum_{n=1}^{n_\eta - 1}\Big(\int_{[T_n,T_{n+1})} \frac{1}{\kappa(s)} B(X^{T_n}_{[0,\infty)}, \mathrm{d}s) + \beta_{T_{n+1}}\Big) \\
&\quad + \int_{[\eta,\infty)} \frac{1}{\kappa(s)} B(X^\eta_{[0,\infty)}, \mathrm{d}s)
\end{aligned}$$

by decomposing the integrals between jumps of Z and using that β is only non-zero at jumps of Z. Hence we can write

$$\begin{aligned}
\mathcal{P}^\circ(t) &= B^\circ(X^t_{[0,T_1)}) - B^\circ(X^t_{[0,t]}) + B^\circ(X^{T_1}_{[0,T_1)}) - B^\circ(X^t_{[0,T_1)}) \\
&\quad + \sum_{n=1}^{n_\eta - 1}\Big(B^\circ(X^{T_n}_{[0,T_{n+1})}) - B^\circ(X^{T_n}_{[0,T_n)}) + B^\circ(X^{T_{n+1}}_{[0,T_{n+1})}) - B^\circ(X^{T_n}_{[0,T_{n+1})})\Big) \\
&\quad + B^\circ(X^\eta_{[0,\infty)}) - B^\circ(X^\eta_{[0,\eta)}) \\
&= B^\circ(X^{T_1}_{[0,T_1)}) - B^\circ(X^t_{[0,t]}) + \sum_{n=1}^{n_\eta - 1}\Big(B^\circ(X^{T_{n+1}}_{[0,T_{n+1})}) - B^\circ(X^{T_n}_{[0,T_n)})\Big) \\
&\quad + B^\circ(X^\eta_{[0,\infty)}) - B^\circ(X^\eta_{[0,\eta)}).
\end{aligned}$$

Observe that the sum telescopes, so we have

$$\begin{aligned}\mathcal{P}^\circ(t) &= B^\circ(X^{T_1}_{[0,T_1)}) - B^\circ(X^{t}_{[0,t]}) + B^\circ(X^{\eta}_{[0,\eta)}) - B^\circ(X^{T_1}_{[0,T_1)}) \\ &\quad + B^\circ(X^{\eta}_{[0,\infty)}) - B^\circ(X^{\eta}_{[0,\eta)}) \\ &= B^\circ(X^{\eta}_{[0,\infty)}) - B^\circ(X^{t}_{[0,t]}).\end{aligned}$$

Consequently,

$$\mathcal{P}(t) = \kappa(t)\mathcal{P}^\circ(t) = \kappa(t)\Big(B^\circ(X^{\eta}_{[0,\infty)}) - B^\circ(X^{t}_{[0,t]})\Big)$$

as desired. □

Relation between reserves

Using Proposition 3.5.2 and Proposition 3.5.3, the following theorem is now immediate:

Theorem 3.5.4. *(Representations of transaction time present value.)*
For $t \geq 0$, it holds that

$$\begin{aligned}\mathcal{P}(t) &= P(t) + \kappa(t)\Big(B^\circ(X^{\eta}_{[0,t]}) - B^\circ(X^{t}_{[0,t]})\Big) \\ &= P(t) + \sum_{t<s<\infty} \kappa(t)\Big(B^\circ(X^{s}_{[0,t]}) - B^\circ(X^{s-}_{[0,t]})\Big)\end{aligned}$$

on the event $(\eta < \infty)$.

Proof. From Propositions 3.5.2 and 3.5.3, it follows that on $(\eta < \infty)$:

$$\mathcal{P}(t) - P(t) = \kappa(t)\Big(B^\circ(X^{\eta}_{[0,t]}) - B^\circ(X^{t}_{[0,t]})\Big),$$

which implies

$$\mathcal{P}(t) = P(t) + \kappa(t)\Big(B^\circ(X^{\eta}_{[0,t]}) - B^\circ(X^{t}_{[0,t]})\Big).$$

This proves the first equality. The second equality corresponds to showing that

$$B^\circ(X^{\eta}_{[0,t]}) - B^\circ(X^{t}_{[0,t]}) = \sum_{t<s<\infty} \Big(B^\circ(X^{s}_{[0,t]}) - B^\circ(X^{s-}_{[0,t]})\Big).$$

This is trivially satisfied on $(t \geq \eta)$ since both the left- and right-hand side are zero. Using the same notation as the proof of Proposition 3.5.3 and the convention $T_0 = t$, we see on $(t < \eta)$:

$$\begin{aligned}\sum_{t<s<\infty} \Big(B^\circ(X^{s}_{[0,t]}) - B^\circ(X^{s-}_{[0,t]})\Big) &= \sum_{n=1}^{n_\eta} \Big(B^\circ(X^{T_n}_{[0,t]}) - B^\circ(X^{T_n-}_{[0,t]})\Big) \\ &= \sum_{n=1}^{n_\eta} \Big(B^\circ(X^{T_n}_{[0,t]}) - B^\circ(X^{T_{n-1}}_{[0,t]})\Big) \\ &= B^\circ(X^{\eta}_{[0,t]}) - B^\circ(X^{t}_{[0,t]})\end{aligned}$$

since the sum telescopes. This establishes the second equality and thus completes the proof. □

The assertion of Theorem 3.5.4 is quite intuitive, and one could alternatively have formulated the setup by taking $\mathcal{P}(t) = P(t) + \kappa(t)(B^\circ(X^\eta_{[0,t]}) - B^\circ(X^t_{[0,t]}))$ as a definition and proceeded from there. Defining the transaction time present value through the cash flow as $\mathcal{P}(t) = \int_{(t,\infty)} \frac{\kappa(t)}{\kappa(s)} \mathcal{B}(\mathrm{d}s)$ seems, however, the more principled approach. Furthermore, the cash flow representation is easier to work with in some situations; confer also with the proof of Theorem 3.5.10. Regardless, the representation in Theorem 3.5.4 tends to be more convenient when linking the transaction and valid time reserves as can be seen, for instance, in Example 3.5.8.

Remark 3.5.5. (Stylized illustration of Theorem 3.5.4.)

To better explain the contents of Theorem 3.5.4, we give a stylized example. Suppose that one is situated at time t, that benefits have been paid in $[0, t-1]$, and that the finalized payments consist of benefits during all of $[0, t+1]$. Suppose further that the benefits concerning $(t-1, t+1]$ occur as backpay at time $t+1$. Then the payments concerning $(t, t+1]$ appear in $P(t)$, while the payments concerning $(t-1, t]$ appear in $\kappa(t)\big(B^\circ(X^\eta_{[0,t]}) - B^\circ(X^t_{[0,t]})\big) = \kappa(t)\big(B^\circ(X^{t+1}_{[0,t]}) - B^\circ(X^{t+1-}_{[0,t]})\big)$. ▽

Remark 3.5.6. (Relation between reserves in valid and transaction time.)

By taking the conditional expectation given $\mathcal{F}^{\mathcal{Z}}_t$ of the expression from Theorem 3.5.4, we conclude that the expected present value in transaction time $\mathcal{V}$ is different from the classic expected present value in valid time V in two fundamental ways:

1. It reserves additionally to previously wrongly settled payments, so it is no longer strictly prospective in valid time, in the sense that payments may relate to valid time events that lie before the current point in time.

2. It conditions on the filtration $\mathcal{F}^{\mathcal{Z}}$, which is observable, compared to $\mathcal{F}^X$, which is only partially observable.

If there had been no previously wrongly settled payments, and if the conditional expectations given the two filtrations had been equal, then the reserves would also have been identical. Even though the relation between the present values is relatively simple, this does not translate into a simple relation between $\mathcal{V}$ and V in the general case. This is because we have so far imposed very little structure on the model for $\mathcal{Z}$, so how the conditional distribution of $\mathcal{Z}_s$ given $\mathcal{F}^{\mathcal{Z}}_t$ for $s \geq t$ depends on $\mathcal{F}^{\mathcal{Z}}_t$ can be almost arbitrarily complicated.

Note that another important consequence of Theorem 3.5.4 (or rather Proposition 3.5.3) is that one does not need the distribution of $\mathcal{Z}$ to calculate $\mathcal{V}(t)$. The conditional distribution of $X_{[0,\infty)}$ given $\mathcal{F}_t^{\mathcal{Z}}$ is sufficient, since

$$\mathcal{V}(t) = \mathbb{E}\big[\kappa(t)B^\circ(X_{[0,\infty)}) \mid \mathcal{F}_t^{\mathcal{Z}}\big] - \kappa(t)B^\circ(X_{[0,t]}^t).$$

Further, since $X_{[0,\infty)}$ equals $X_{[0,\infty)}^\eta$ almost surely with respect to $\mathbb{P}$ and $X_{[0,\infty)}^\eta$ is $\mathcal{F}_\infty^{\mathcal{Z}}$-measurable, the distribution of $\mathcal{Z}$ determines the conditional distribution of $X_{[0,\infty)}$ given $\mathcal{F}_t^{\mathcal{Z}}$ for any outcome of $\mathcal{Z}_{[0,t]}$. Consequently, the conditional distribution might be the natural modeling object. ▽

Remark 3.5.7. (Non-monotone information.)

Write $V(t) = g(X_{[0,t]})$ for a measurable function g, which exists by the Doob-Dynkin lemma. Continuing the discussion from Remark 3.5.6, standard practice seems to be to use the individual reserve $V^t(t) = g(X_{[0,t]}^t)$ at time t and use IBNR and RBNS factors on an aggregate level to correct for the fact that typically $X \neq X^t$. Note that the information that one uses for reserving is then non-monotone, since for $0 \leq s \leq t$, it holds that $X_{[0,s]}^s$ is generally unknown from $X_{[0,t]}^t$ and vice versa. Reserves in the presence of non-monotone information have been studied in Christiansen & Furrer (2021). In this, stochastic Thiele differential equations for prospective reserves are derived subject to information deletions, i.e. non-monotone information. These might be useful for studying the properties of the reserves V^t currently used in practice. ▽

Example 3.5.8. (Disability insurance with different origins: Reserving.)

We here describe reserving in the transaction time model from Example 3.4.3 and find explicit expressions for $\mathcal{V}$. To obtain intuitive formulas, we impose additional structure on the model for $\mathcal{Z}$ in the form of a conditional independence assumption.

Write $V_j(t,u) = \mathbb{E}[P(t) \mid X_t = (j,u)]$ for the state-wise valid time reserves in the semi-Markov setup. These may be calculated using known methods, see e.g. Christiansen (2012) and Buchardt et al. (2015). Note that on $(Z_t = a)$, we have $\mathcal{F}_t^X = \mathcal{F}_t^{\mathcal{Z}}$ and $\mathcal{P}(t) = P(t)$ and thus $\mathcal{V}(t) = V_a(t,t)$. On $(Z_t \in \{r,d\})$, we have $\mathcal{P}(t) = P(t) = 0$ and thus $\mathcal{V}(t) = 0$. Hence only the case $(Z_t \in \{i_1, i_2\})$ corresponding to an RBNS claim requires consideration. Using Theorem 3.5.4, we get

$$\begin{aligned}\mathcal{V}(t) &= \mathbb{E}\left[P(t) + \kappa(t)\int_{[\theta_i,t]} \frac{1}{\kappa(s)}(b_{Y_s} - b_{Z_t})\,\mathrm{d}s \,\middle|\, \mathcal{F}_t^{\mathcal{Z}}\right] \\ &= \frac{\kappa(t)}{\kappa(\theta_i)}\mathbb{E}\big[P(\theta_i) \mid \mathcal{F}_t^{\mathcal{Z}}\big] - \int_{(\theta_i,t]} \frac{\kappa(t)}{\kappa(s)} b_{Z_t}\,\mathrm{d}s.\end{aligned}$$

Assume that

$$X_{[0,\infty)} \perp\!\!\!\perp \mathcal{F}_t^{\mathcal{Z}} \mid \mathcal{F}_t^X.$$

In other words, if one had known the true value of the biometric state process, additional transaction time information is superfluous. Then by the law of iterated expectations and the semi-Markov property,

$$\mathbb{E}\big[P(\theta_i) \mid \mathcal{F}_t^{\mathcal{Z}}\big] = \sum_{k=1}^{2} \mathbb{P}(Y_{\theta_i} = i_k \mid \mathcal{F}_t^{\mathcal{Z}}) \bigg(\frac{\kappa(\theta_i)}{\kappa(t)} V_{i_k}(t, t-\theta_i) + \int_{(\theta_i, t]} \frac{\kappa(\theta_i)}{\kappa(s)} b_{i_k} \,\mathrm{d}s \bigg).$$

To conclude, on $(Z_t \in \{i_1, i_2\})$ it holds that

$$\mathcal{V}(t) = \sum_{k=1}^{2} \mathbb{P}(Y_{\theta_i} = i_k \mid \mathcal{F}_t^{\mathcal{Z}}) \bigg(V_{i_k}(t, t-\theta_i) + \int_{(\theta_i, t]} \frac{\kappa(t)}{\kappa(s)} (b_{i_k} - b_{Z_t}) \,\mathrm{d}s \bigg), \quad (3.5.3)$$

which is an explicit expression for the RBNS reserve. The probabilities $\mathbb{P}(Y_{\theta_i} = i_k \mid \mathcal{F}_t^{\mathcal{Z}})$ for $k \in \{1, 2\}$ may be calculated as absorption probabilities by extending the state space of the transaction time process to include separate reactivated and dead states for each of the disability origins.

As noted in Remark 3.5.6, the transaction time reserves differ from the valid time reserves through both the present values and the conditioning information. If the payment rates in the two disabled states, that is b_{i_1} and b_{i_2}, are equal, then the present values are equal, i.e. $\mathcal{P}(t) = P(t)$, and we obtain on $(Z_t \in \{i_1, i_2\})$ that

$$\mathcal{V}(t) = \sum_{k=1}^{2} \mathbb{P}(Y_{\theta_i} = i_k \mid \mathcal{F}_t^{\mathcal{Z}}) V_{i_k}(t, t-\theta_i).$$

If we additionally assume that the transition rates from states i_1 and i_2 are equal, then the difference due to differing conditioning information also disappears, and we get

$$\mathcal{V}(t) = V_{i_1}(t, t-\theta_i) = V(t),$$

meaning that the valid and transaction time reserves agree.

Note that it is easy to extend the model to include n disabled states $i_1, \ldots, i_n$. We may also extend the example to allow for transition between the disabled states in the valid time model. The disabled states could then represent more diverse and complex phenomena such as the degree of lost earning capacity or diagnoses. Such a model is depicted in Figure 3.3.

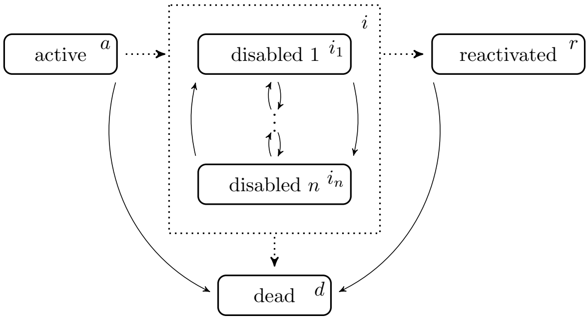

Figure 3.3: *Valid time model from Figure 3.1 extended to n disabled states and allowing for transition between the disabled states. To reduce clutter, transitions to and from the disabled states $\{i_1, \dots, i_n\}$ are represented with a single dotted arrow.*

Allowing for transition between the disabled states comes at a cost in the form of increased complexity of the transaction time model, since the latter needs to be able to generate the valid time process. One option is to let the Z-component take values in the state space depicted in Figure 3.4.

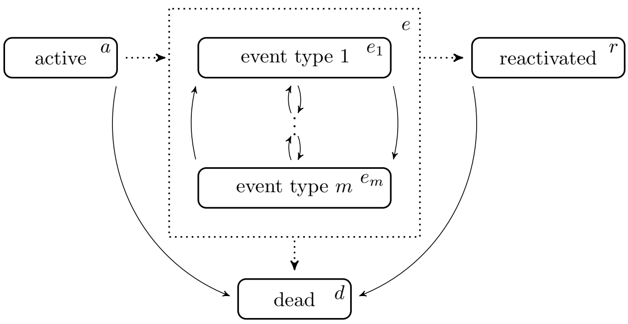

Figure 3.4: *The Z-component of a transaction time process $\mathcal{Z}$ that can generate the extended valid time model of Figure 3.3. To reduce clutter, transitions to and from the RBNS states $\{e_1, \dots, e_m\}$ are represented with a single dotted arrow.*

As noted in Remark 3.5.6, we do not need to specify the full distribution of changes in H_s^t at jumps of Z between states $e_1, \dots, e_m$. This distribution only affects the reserves through the distribution of the valid time process conditional on the observed information, and it therefore suffices to model this. On $(Z_t \in \{e_1, \dots, e_m\})$, corresponding to an RBNS claim, one may show that

$$\begin{aligned}\mathcal{V}(t) = \sum_{k=1}^{n} \Bigg(&\int_{[0,t]} V_{i_k}(t,u) \mathbb{P}(Y_t = i_k, U_t \in \mathrm{d}u \mid \mathcal{F}_t^{\mathcal{Z}}) \\ &+ \int_{(\theta_i, t]} \frac{\kappa(t)}{\kappa(s)} \mathbb{P}(Y_s = i_k \mid \mathcal{F}_t^{\mathcal{Z}})(b_{i_k} - b_{Y_s^t}) \, \mathrm{d}s \Bigg).\end{aligned}$$

This identity is comparable to (3.5.3), but it is more complex and requires one to model the path the insured takes through the disabled states between time θ_i and time t. This may be seen from the dependence on U_t in the first term and the dependence on Y_s for $s \in (\theta_i, t]$ in the second term.

Finally, it is easy to extend the example to allow for general semi-Markov payments of the form described in Example 3.3.2. This includes risk periods, waiting periods, and transition payments. ○

The above example serves as a simple theoretical demonstration of the potential of our general framework. As there is no reporting delay, there is no IBNR reserve in Example 3.5.8. The RBNS reserve of the example also only differs from the valid time disability reserve due to the imperfect observation of the disability type. To capture the full picture of IBNR and RBNS reserving, one would need to explore more intricate transaction time models with both reporting delays and claim adjudications. While this extension is outside the scope of this paper, our general framework readily allows for such continued studies. We stress that such applications are the main motivation for introducing the transaction time framework. In addition to developing model extensions, it could also be relevant to develop estimation procedures for the above example as well as for more complicated models.

Reserve dynamics

The study of reserve dynamics is of great importance, especially in relation to model validation, Cantelli's theorem and reserve-dependent payments, Hattendorff's theorem on non-correlation between losses, and the emergence and decomposition of surplus as well as sensitivity analyses, cf. Section 1 in Christiansen & Furrer (2021). We conclude this section by deriving the dynamics of the transaction time reserve $\mathcal{V}$ and valid time reserve V following the same procedure as for the classic reserve, see e.g. Christiansen & Djehiche (2020). Essentially, this amounts to applying an explicit martingale representation theorem to $\mathcal{V}$; the idea of applying martingale representation techniques dates back to Norberg (1992). The dynamics of the prospective reserves bears a resemblance to Thiele's differential equation; one might even say it constitutes a stochastic version of Thiele's differential equation. The literature, however, seems to reserve the term *stochastic Thiele equation* for the stochastic differential equation related to the so-called state-wise prospective reserves, see e.g. Christiansen & Furrer (2021).

Recall the definitions $\mathcal{V}(t) = \mathbb{E}\big[\mathcal{P}(t) \mid \mathcal{F}_t^{\mathcal{Z}}\big]$ and $\mathcal{Z}_t = (Z_t, H_t^t)$. Define a random counting measure $\mu_{\mathcal{Z}}$ corresponding to $\mathcal{Z}$ by

$$\mu_{\mathcal{Z}}(C) = \sum_{n=1}^{\infty} 1_C(T_n, \mathcal{Z}_{T_n}), \qquad C \in \mathbb{B}([0,\infty)) \otimes \mathbb{B}(\mathbb{R}^{\infty}),$$

and let $\Lambda_{\mathcal{Z}}$ be its *compensating measure*, given in Definition 4.3.2 (iii) of Jacobsen (2006). By Theorem 4.5.2 of Jacobsen (2006), if $\mathbb{E}[\mu_{\mathcal{Z}}([0,t] \times D)] < \infty$ for all $t \geq 0$ and $D \in \mathbb{B}(\mathbb{R}^\infty)$, we have that

$$t \mapsto \mu_{\mathcal{Z}}([0,t] \times D) - \Lambda_{\mathcal{Z}}([0,t] \times D)$$

is a martingale for any $D \in \mathbb{B}(\mathbb{R}^\infty)$. Let $\xi_n = (T_1, ..., T_n; \mathcal{Z}_{T_1}, ..., \mathcal{Z}_{T_n})$ be the MPP history of $\mathcal{Z}$ at time T_n.

Write $\zeta = (\zeta_z, \zeta_h)$ for a generic realization of $\mathcal{Z}_t$, where the coordinates ζ_z and ζ_h pertain to Z_t and H_t^t, respectively. Finally, define the sums at risk in the transaction time model for a jump of $\mathcal{Z}$ to ζ at time t:

$$\begin{aligned}\mathcal{R}(t,\zeta) = \sum_{n=1}^{\infty} 1_{(T_n < t \leq T_{n+1})} \Big(\kappa(t) \Big(& B^\circ((f_{\zeta_h}(s))_{0 \leq s \leq t}) - B^\circ(X_{[0,t]}^{t-}) \Big) \\ & + \mathbb{E}[\mathcal{P}(t) \mid \xi_n, (T_{n+1}, \mathcal{Z}_{T_{n+1}}) = (t,\zeta)] - \mathbb{E}[\mathcal{P}(t) \mid \xi_n, T_{n+1} > t] \Big).\end{aligned}$$

This is a difference in payments and reserves at time t between a jump-to-ζ and a remain-in-$\mathcal{Z}_{t-}$ scenario.

Remark 3.5.9. (Definition of non-standard conditional expectations.)

One should be careful about the definition of $\mathbb{E}[\mathcal{P}(t) \mid \xi_n, T_{n+1} > t]$ and similar quantities outside $(T_{n+1} > t)$, confer with e.g. Christiansen & Furrer (2021). In this paper, it corresponds to the version

$$\mathbb{E}[\mathcal{P}(t) \mid \xi_n, T_{n+1} > t] = \frac{\mathbb{E}[\mathcal{P}(t) 1_{(T_{n+1} > t)} \mid \xi_n]}{\mathbb{E}[1_{(T_{n+1} > t)} \mid \xi_n]}$$

under the convention $0/0 = 0$ and where the expectations are the regular conditional expectations constructed in Jacobsen (2006). That this version is the relevant one follows from the proof of Theorem 3.5.10. $\triangledown$

We then have the following theorem:

Theorem 3.5.10. *(Transaction time reserve dynamics.)*

For $t \geq 0$, it holds that

$$\mathcal{V}(\mathrm{d}t) = \mathcal{V}(t-) \frac{\kappa(\mathrm{d}t)}{\kappa(t-)} - \mathcal{B}(\mathrm{d}t) + \int_{\mathbb{R}^\infty} \mathcal{R}(t,\zeta) \, (\mu_{\mathcal{Z}} - \Lambda_{\mathcal{Z}})(\mathrm{d}t, \mathrm{d}\zeta). \tag{3.5.4}$$

Proof. Introduce

$$\mathcal{P}^\circ(t) = \frac{1}{\kappa(t)} \mathcal{P}(t).$$

We have that $\mathcal{P}^\circ(0) = \mathcal{P}(0)$, which is assumed integrable, so we can define

$$t \mapsto M_t = \mathbb{E}\big[\mathcal{P}^\circ(0) \mid \mathcal{F}_t^{\mathcal{Z}}\big],$$

which is a martingale. Since $\mathcal{V}$ is presumed $\mathcal{F}^{\mathcal{Z}}$-adapted and $\mathbb{P}$-a.s. right-continuous, the same holds for a version of M, since $M_t = \frac{1}{\kappa(t)}\mathcal{V}(t) + \mathcal{P}^\circ(0) - \mathcal{P}^\circ(t)$. Then a martingale representation theorem, namely Theorem 4.6.1 of Jacobsen (2006), gives the existence of predictable processes S_s^ζ such that

$$M_t = M_0 + \int_{(0,t]\times\mathbb{R}^\infty} S_s^\zeta \, (\mu_{\mathcal{Z}} - \Lambda_{\mathcal{Z}})(\mathrm{d}s, \mathrm{d}\zeta)$$

$\mathbb{P}$-a.s. simultaneously over t. Using the adaptedness of M, we can, as in the proof of the aforementioned Theorem 4.6.1, use Proposition 4.2.1(biii) of Jacobsen (2006) to write

$$M_t = \sum_{n=0}^{\infty} 1_{(T_n \leq t < T_{n+1})} g_{\xi_n}^n(t)$$

for measurable functions $(h_n, t) \mapsto g_{h_n}^n(t)$. Due to M being a conditional expectation, we can use Corollary 4.2.2 of Jacobsen (2006) to identify

$$g_{\xi_n}^n(t) = \mathbb{E}[\mathcal{P}^\circ(0) \mid \xi_n, T_{n+1} > t]$$

on $(T_n \leq t < T_{n+1})$. To identify the function for all (h_n, t), we first observe that according to Remark 4.2.3 of Jacobsen (2006),

$$g_{\xi_n}^n(t) = \frac{\mathbb{E}[\mathcal{P}^\circ(0) 1_{(T_{n+1}>t)} \mid \xi_n]}{\mathbb{E}[1_{(T_{n+1}>t)} \mid \xi_n]}$$

on $(T_n \leq t < T_{n+1})$. Define the functions $(h_n, t) \mapsto g_{h_n}^n(t)$ as

$$g_{h_n}^n(t) = \frac{\mathbb{E}[\mathcal{P}^\circ(0) 1_{(T_{n+1}>t)} \mid \xi_n = h_n]}{\mathbb{E}[1_{(T_{n+1}>t)} \mid \xi_n = h_n]}$$

on the set $D_n = \{(h_n, t) : \mathbb{E}[1_{(T_{n+1}>t)} \mid \xi_n = h_n] \neq 0\}$ and zero otherwise. These functions are well-defined since the conditional expectations are regular and fixed. By the above calculations, they satisfy the required identity of Proposition 4.2.1(biii) in Jacobsen (2006). For the measurability condition, note first that $\mathbb{E}[1_{(T_{n+1}>t)} \mid \xi_n = h_n]$ is measurable as a function of h_n since it is a regular conditional expectation, and that it is jointly measurable as a function of (h_n, t) since it is right-continuous as a function of t for any h_n by the dominated convergence theorem. This implies that D_n is measurable. By the same arguments, $\mathbb{E}[\mathcal{P}^\circ(0) 1_{(T_{n+1}>t)} \mid \xi_n = h_n]$ is seen to be jointly measurable as a function of (h_n, t). From this we may conclude that $(h_n, t) \mapsto g_{h_n}^n(t)$ is measurable, so it especially satisfies Proposition 4.2.1(biii) in Jacobsen (2006). In the following, we write $\mathbb{E}[\mathcal{P}^\circ(0) \mid \xi_n, T_{n+1} > t]$ for $g_{\xi_n}^n(t)$, but this is merely notation; calculations with $\mathbb{E}[\mathcal{P}^\circ(0) \mid \xi_n, T_{n+1} > t]$ actually use the properties of $g^n_{h_n}(t)$.

The proof of the aforementioned Theorem 4.6.1 furthermore gives that

$$S^\zeta_t = \sum_{n=0}^{\infty} 1_{(T_n < t \leq T_{n+1})} \Big(g^{n+1}_{(\xi_n,(t,\zeta))}(t) - g^n_{\xi_n}(t) \Big),$$

so that

$$\begin{aligned} S^\zeta_t &= \sum_{n=0}^{\infty} 1_{(T_n < t \leq T_{n+1})} \big(\mathbb{E}[\mathcal{P}^\circ(0) \mid \xi_n, (T_{n+1}, \mathcal{Z}_{T_{n+1}}) = (t,\zeta)] \\ &\quad - \mathbb{E}[\mathcal{P}^\circ(0) \mid \xi_n, T_{n+1} > t] \big) \\ &= \sum_{n=0}^{\infty} 1_{(T_n < t \leq T_{n+1})} \big(\mathbb{E}[\mathcal{P}^\circ(t-) \mid \xi_n, (T_{n+1}, \mathcal{Z}_{T_{n+1}}) = (t,\zeta)] \\ &\quad - \mathbb{E}[\mathcal{P}^\circ(t-) \mid \xi_n, T_{n+1} > t] \big) \end{aligned}$$

using that $\mathcal{P}^\circ(0) - \mathcal{P}^\circ(t-) = \int_{(0,t)} \frac{1}{\kappa(s)} \mathcal{B}(ds)$ are ξ_n-measurable on $(T_n < t \leq T_{n+1})$. Therefore, the dynamics of M is

$$\begin{aligned} \mathrm{d}M_t &= \int_{\mathbb{R}^\infty} S^\zeta_t \, (\mu_{\mathcal{Z}} - \Lambda_{\mathcal{Z}})(\mathrm{d}t, \mathrm{d}\zeta) \\ &= \sum_{n=0}^{\infty} \int_{\mathbb{R}^\infty} 1_{(T_n < t \leq T_{n+1})} \Big(\mathbb{E}[\mathcal{P}^\circ(t-) \mid \xi_n, (T_{n+1}, \mathcal{Z}_{T_{n+1}}) = (t,\zeta)] \\ &\quad - \mathbb{E}[\mathcal{P}^\circ(t-) \mid \xi_n, T_{n+1} > t] \Big) \, (\mu_{\mathcal{Z}} - \Lambda_{\mathcal{Z}})(\mathrm{d}t, \mathrm{d}\zeta). \end{aligned}$$

Using that

$$\mathcal{P}^\circ(t) - \mathcal{P}^\circ(0) = - \int_{(0,t]} \frac{1}{\kappa(s)} \mathcal{B}(\mathrm{d}s) \tag{3.5.5}$$

is $\mathcal{F}^{\mathcal{Z}}$-adapted, we get

$$\begin{aligned} \mathbb{E}[\mathcal{P}^\circ(0) \mid \mathcal{F}^{\mathcal{Z}}_t] - \mathbb{E}[\mathcal{P}^\circ(0) \mid \mathcal{F}^{\mathcal{Z}}_0] &= \mathbb{E}[\mathcal{P}^\circ(t) \mid \mathcal{F}^{\mathcal{Z}}_t] - \mathbb{E}[\mathcal{P}^\circ(0) \mid \mathcal{F}^{\mathcal{Z}}_0] \\ &\quad - (\mathcal{P}^\circ(t) - \mathcal{P}^\circ(0)), \end{aligned}$$

which upon rearrangement becomes

$$\mathbb{E}[\mathcal{P}^\circ(t) \mid \mathcal{F}^{\mathcal{Z}}_t] - \mathbb{E}[\mathcal{P}^\circ(0) \mid \mathcal{F}^{\mathcal{Z}}_0] = \mathcal{P}^\circ(t) - \mathcal{P}^\circ(0) + \mathbb{E}[\mathcal{P}^\circ(0) \mid \mathcal{F}^{\mathcal{Z}}_t] - \mathbb{E}[\mathcal{P}^\circ(0) \mid \mathcal{F}^{\mathcal{Z}}_0].$$

Introducing

$$\mathcal{V}^\circ(t) = \frac{1}{\kappa(t)} \mathcal{V}(t) = \mathbb{E}[\mathcal{P}^\circ(t) \mid \mathcal{F}^{\mathcal{Z}}_t],$$

we can write this as

$$\mathcal{V}^\circ(t) - \mathcal{V}^\circ(0) = \mathcal{P}^\circ(t) - \mathcal{P}^\circ(0) + M_t - M_0.$$

The identity (3.5.5) furthermore gives

$$\mathcal{P}^\circ(\mathrm{d}t) = -\frac{1}{\kappa(t)}\mathcal{B}(\mathrm{d}t).$$

The above calculations imply

$$\begin{aligned}
\mathcal{V}^\circ(\mathrm{d}t) &= \mathrm{d}M_t + \mathcal{P}^\circ(\mathrm{d}t) \\
&= \sum_{n=0}^{\infty} \int_{\mathbb{R}^\infty} 1_{(T_n < t \leq T_{n+1})} \Big(\mathbb{E}[\mathcal{P}^\circ(t-) \mid \xi_n, (T_{n+1}, \mathcal{Z}_{T_{n+1}}) = (t,\zeta)] \\
&\quad - \mathbb{E}[\mathcal{P}^\circ(t-) \mid \xi_n, T_{n+1} > t] \Big) \, (\mu_{\mathcal{Z}} - \Lambda_{\mathcal{Z}})(\mathrm{d}t, \mathrm{d}\zeta) - \frac{1}{\kappa(t)}\mathcal{B}(\mathrm{d}t).
\end{aligned}$$

The time t payment $\frac{1}{\kappa(t)}\mathcal{B}(\{t\})$ can be taken out of both intergrands, and this amounts to $B^\circ((f_{\zeta_h}(s))_{0\leq s\leq t}) - B^\circ(X^{t-}_{[0,t]})$. It is the difference in the payment at time t between a jump and a remain scenario when $\mathcal{Z}$ jumps to ζ. Taking out the time t payment, we get $\mathcal{P}^\circ(t-) = \frac{1}{\kappa(t)}\mathcal{B}(\{t\}) + \mathcal{P}^\circ(t)$, and using integration by parts, we finally have

$$\begin{aligned}
\mathcal{V}(\mathrm{d}t) &= \mathrm{d}(\kappa(t)\mathcal{V}^\circ(t)) \\
&= \mathcal{V}^\circ(t-)\kappa(\mathrm{d}t) + \kappa(t)\mathcal{V}^\circ(\mathrm{d}t) \\
&= \mathcal{V}(t-)\frac{\kappa(\mathrm{d}t)}{\kappa(t-)} - \mathcal{B}(\mathrm{d}t) \\
&\quad + \sum_{n=1}^{\infty} \int_{\mathbb{R}^\infty} 1_{(T_n < t \leq T_{n+1})} \Big(B^\circ((f_{\zeta_h}(s))_{0\leq s\leq t}) - B^\circ(X^{t-}_{[0,t]}) \\
&\quad + \mathbb{E}[\mathcal{P}(t) \mid \xi_n, (T_{n+1}, \mathcal{Z}_{T_{n+1}}) = (t,\zeta)] \\
&\quad - \mathbb{E}[\mathcal{P}(t) \mid \xi_n, T_{n+1} > t] \Big) \, (\mu_{\mathcal{Z}} - \Lambda_{\mathcal{Z}})(\mathrm{d}t, \mathrm{d}\zeta),
\end{aligned}$$

which yields the desired result by definition of the sums at risk. □

Theorem 3.5.10 shows that the transaction time reserve $\mathcal{V}$ changes with interest accrual $\mathcal{V}(t-)\frac{\kappa(\mathrm{d}t)}{\kappa(t-)}$, actual benefits less premiums $\mathcal{B}(\mathrm{d}t)$ and a martingale term $\int_{\mathbb{R}^\infty} \mathcal{R}(t,\zeta)\,(\mu_{\mathcal{Z}} - \Lambda_{\mathcal{Z}})(\mathrm{d}t, \mathrm{d}\zeta)$, which is the sums at risk integrated with respect to the underlying compensated random counting measure. The martingale term may be interpreted as stochastic noise since it is a mean-zero process, and may thus be used for model validation and back-testing purposes. Actual applications are outside the scope of this paper.

One could alternatively have derived Theorem 3.5.10 from Theorem 7.1 in Christiansen (2021b), which is an explicit martingale representation theorem that holds even when the information being conditioned on is non-monotone. The proof presented here is however more concise, as our information $\mathcal{F}^{\mathcal{Z}}$ is monotone, so more standard results apply. Theorem 3.5.10 is similar to Proposition 3.2 in Christiansen

& Djehiche (2020), but differs among other things by not being restricted to state processes taking values in a finite space.

Remark 3.5.11. (Dynamics of valid time reserve.)

Define the random counting measure μ_X corresponding to X by

$$\mu_X(C) = \sum_{n=1}^{\infty} 1_C(\tau_n, X_{\tau_n}), \qquad C \in \mathbb{B}([0,\infty)) \otimes \mathbb{B}(\mathbb{R}^d).$$

Let Λ_X be the compensating measure for μ_X, and let $\gamma_n = (\tau_1, ..., \tau_n; X_{\tau_1}, ..., X_{\tau_n})$ be the MPP history of X at time τ_n. By the same calculations as for Theorem 3.5.10, we find the dynamics of the valid time reserve V:

$$V(\mathrm{d}t) = V(t-)\frac{\kappa(\mathrm{d}t)}{\kappa(t-)} - B(\mathrm{d}t) + \int_{\mathbb{R}^d} R(t,y)\ (\mu_X - \Lambda_X)(\mathrm{d}t, \mathrm{d}y) \tag{3.5.6}$$

for the sums at risk

$$\begin{aligned} R(t,y) = \sum_{n=1}^{\infty} 1_{(\tau_n < t \leq \tau_{n+1})} \Big(& B\big((f_{(H_{t-},(t,y))}(s))_{0\leq s\leq t}, \{t\}\big) - B\big((f_{H_{t-}}(s))_{0\leq s\leq t}, \{t\}\big) \\ & + \mathbb{E}[P(t) \mid \gamma_n, (\tau_{n+1}, X_{\tau_{n+1}}) = (t,y)] - \mathbb{E}[P(t) \mid \gamma_n, \tau_{n+1} > t] \Big). \end{aligned}$$

This result is again similar to Proposition 3.2 in Christiansen & Djehiche (2020), but still differs among other things by not being restricted to state processes taking values in a finite space. The conditional expectations are to be interpreted as in Remark 3.5.9.

Suppose now that X is a pure Markov jump process on a finite state space $E = \{1, 2, ..., J\}$ with payments specified as in Example 3.3.2. In other words, the valid time payments consist of deterministic sojourn payments $t \mapsto B_j(t)$ and deterministic transition payments $t \mapsto b_{jk}(t)$. Then

$$B\big((f_{(H_{t-},(t,y))}(s))_{0\leq s\leq t}, \{t\}\big) - B\big((f_{H_{t-}}(s))_{0\leq s\leq t}, \{t\}\big) = b_{X_{t-}y}(t).$$

Furthermore,

$$\mu_X(\mathrm{d}t, \{k\}) = N_{X_{t-}k}(\mathrm{d}t)$$

and, since X is Markovian,

$$\Lambda_X(\mathrm{d}t, \{k\}) = \Lambda_{X_{t-}k}(\mathrm{d}t)$$

for suitably regular cumulative transition rates $t \mapsto \Lambda_{jk}(t)$. Consequently, by invoking the Markov property, the dynamics (3.5.6) read

$$\begin{aligned} V(\mathrm{d}t) = V(t-)\frac{\kappa(\mathrm{d}t)}{\kappa(t-)} - B(\mathrm{d}t) + \sum_{\substack{j,k=1 \\ j\neq k}}^{J} \mathbb{1}_{(X_{t-}=j)} \big(& b_{jk}(t) + \mathbb{E}[P(t) \,|\, X_t = k] \\ & - \mathbb{E}[P(t) \,|\, X_t = j]\big)\big(N_{jk}(\mathrm{d}t) - \Lambda_{jk}(\mathrm{d}t)\big). \end{aligned} \tag{3.5.7}$$

This constitutes a significant simplification. $\triangledown$

In comparing (3.5.4) with (3.5.6), it is apparent that the transaction and valid time reserves admit comparable dynamics. In both cases, there is a contribution due to interest accrual, a contribution from benefits less premiums, and finally a martingale term. In general, the dynamics of the transaction time reserve are more complicated than that of the valid time reserve – for two reasons. First, the martingale term is more involved, which stems from the fact that the model for $\mathcal{Z}$ is typically more elaborate than that for X. Second, the accumulated cash flow in transaction time $\mathcal{B}$ is a complicated function of, among other things, the accumulated cash flow in valid time B. The difference might be particularly striking under the quite common assumption that X is a pure Markov jump process on a finite state space $E = \{1, 2, \ldots, J\}$ and the valid time payments consist of deterministic sojourn and transition payments. In this case, the dynamics of the valid time reserve simplify, cf. (3.5.7), but there is in general no reason why this simplification should carry over to the transaction time reserve – unless further assumptions are imposed.

Acknowledgments and declarations of interest

We would like to thank an anonymous referee for very helpful comments and suggestions. Oliver Lunding Sandqvist's research has partly been funded by the Innovation Fund Denmark (IFD) under File No. 1044-00144B. The authors declare no conflicts of interest.

Chapter 4

A multistate approach to disability insurance reserving with information delays

This chapter contains the manuscript *Sandqvist (2025)*.

Abstract

Disability insurance claims are often affected by lengthy reporting delays and adjudication processes. The classic multistate life insurance modeling framework is ill-suited to handle such information delays since the cash flow and available information can no longer be based on the biometric multistate process determining the contractual payments. We propose a new individual reserving model for disability insurance schemes which describes the claim evolution in real-time. Under suitable independence assumptions between the available information and the underlying biometric multistate process, we show that these new reserves may be calculated as natural modifications of the classic reserves. For estimation of the model constituents, we employ the procedure proposed in Buchardt et al. (2025). A real data application shows the practical relevance of our concepts and results.

Keywords: Multistate life insurance; Claims reserving; Incurred-but-not-reported; Reported-but-not-settled

4.1 Introduction

Reserves are fundamental to the insurance industry, and recently, reserving for disability insurance schemes has become a topic of considerable interest for Danish insurers due to new regulation, worsening risks, and heightened price competition. Disability insurance and similar insurance schemes such as workers' compensation insurance generally work by covering disabilities of an insured that occur in a prespecified *coverage period* in exchange for a premium. Disabilities are covered in the sense that benefits are paid out if the insured becomes disabled with a disability that qualifies for a payout per the criteria specified in the insurance contract. The most prominent schemes pay benefits as long as the insured is disabled and is below the retirement age to compensate for lost wages. Usually, disability benefits will not be paid starting from disablement, but only once the disability has lasted a period of time called the *qualifying period* or *waiting period.* In fact, many disabilities will start payout even later due to reporting and adjudication delays. Reporting delays are defined as the time between the occurrence and reporting of an event. For Danish insurance companies, disabilities generally have long reporting delays compared to other insurance events such as deaths. The adjudication delay is defined as the time between when a claim is reported and when it is adjudicated. During the adjudication process, the insurance company evaluates whether the insured is eligible for disability benefits or not. This can be a lengthy process when there is a need to obtain further clinical assessments of the claimed disability.

These characteristics situate disability insurance somewhere between traditional life and non-life insurance schemes: the long cash flows associated with the possibility of paying benefits from disablement until retirement are similar to the characteristics of other life insurance schemes, while information delays are features that have so far primarily been explored in the non-life part of the insurance reserving literature. In this paper, we propose a model that is tailored to accommodate both of these features.

Our proposed model can in many ways be seen as an extension of the classic semi-Markov models that have dominated the actuarial literature on disability insurance, see for example Janssen (1966), Hoem (1972), Haberman & Pitacco (1998), Helwich (2008), Christiansen (2012), and Buchardt et al. (2015). Such models have also been used extensively in the biostatistical literature, see for example Lagakos et al. (1978), Andersen et al. (1993), Dabrowska (1995), Hougaard (2000), and Spitoni et al. (2012) as well as the references therein. The semi-Markov models have been popular in the disability insurance literature for several reasons. First and foremost, they allow the intensity of mortality and reactivation from a disability to depend on the duration since disablement, which is crucial in practice. In addition, the contractual payments in some cases depend on the duration since the last jump, for example due to a qualifying period, which can be handled in a semi-Markov setup. Finally, semi-Markov models, and multistate models more generally, provide a natural and parsimonious way to represent the information contained in an insurance contract and to capture the intertemporal dependencies of the cash flow. We seek to retain these attractive properties while accommodating the effects of reporting and adjudication delays.

As noted in Buchardt et al. (2023), the fundamental challenge in this endeavor is that contractual payments refer to when events occur (e.g. the time of death or the time of disablement) without any regard to when this information is observed by the insurer. On the other hand, the usual multistate life insurance modeling literature assumes that one can observe the process driving the contractual payments fully and in real-time. Therefore, when the information needed to determine the contractual payments at a given time is not available to the insurer at that time due to information delays, the problem falls outside the usual multistate life insurance modeling framework.

The paper Buchardt et al. (2023) has established a framework intended to deal with these complications, distinguishing between and linking the so-called valid time model, which models when events occur, and the so-called transaction time model, which models what is observed by the insurer. While some relations between the models stay simple in all cases, the relation between the reserves can be almost arbitrarily complicated and hence has to be investigated in specific models. In their Example 5.8, they derive an explicit relation in a simple example, but remark: "*To capture the full picture of IBNR and RBNS reserving, one would need to explore more intricate transaction time models with both reporting delays and claim adjudications*". Here, IBNR stands for incurred-but-not-reported while RBNS stands for reported-but-not-settled. In this paper, we do exactly this, applying the framework to derive explicit and tractable expressions for the reserves of fairly general disability insurance schemes under suitable assumptions. We also give detailed discussions on the reasonableness of the assumptions and the practical relevance of the results. The reserves are operationalized by employing the estimation procedure from Buchardt et al. (2025) who has studied parametric estimation of multistate models subject to reporting delays and adjudications. In addition to providing operational expressions for disability insurance reserves, a main contribution of the paper is to provide intuition for how transaction time information may affect the reserves in a realistic setting, allowing one to adjust the models when certain assumptions are not met, and serving as a basis for future work in this area.

The way reporting delays and adjudication processes are incorporated in our model shares some similarities with parts of the non-life insurance literature on individual reserving models, especially those formulated in the recent string of papers Crevecoeur et al. (2019), Verbelen et al. (2022), Crevecoeur et al. (2022a), and Crevecoeur et al. (2022b). The first two papers explore estimation of the claim frequency subject to IBNR claims. Both assume an underlying Poisson process driving the claim frequency and form a thinned version by deleting claims that are unreported by the time of analysis. In the first paper, the maximum likelihood estimator is obtained by assuming piecewise constant rates, while in the second paper, one treats the deleted claims as missing under an EM-algorithm. The third paper explores reserving and estimation of RBNS claims by modeling the conditional distribution of the full claim development, consisting of all payments and auxiliary characteristics of the claim, given the historical development. The model is calibrated using (weighted) maximum likelihood estimation. The last paper explores reserving and estimation of both IBNR and RBNS, using much of the framework that had been developed in the previous papers. Their reserves do not have closed-form solutions so Monte-Carlo simulation is used. All the models are formulated in discrete time.

Comparing with our approach, a formal difference is that we formulate the models in continuous time. The effect of IBNR on claim frequency is treated in a similar manner, but additional survival probabilities appear in our multistate approach compared to the Poisson model. The primary difference regarding IBNR however stems from how the payments are treated. In the non-life insurance models, the conditional expectation of the ultimate payment given the reporting delay is computed using Monte-Carlo simulation of the full real-time development of the claim while we instead are able to use the known form of the contractual payments. For RBNS modeling, Crevecoeur et al. (2022a) and Crevecoeur et al. (2022b) similarly propose to model the full real-time development of a claim conditional on historical developments. Having to model the full development of a claim results in a larger number of model elements, and thus a greater risk of misspecification. This risk of bias accumulation is acknowledged in Crevecoeur et al. (2022a) where it is suggested to re-scale each time layer of the model to ensure that the sum of the predictions equals the sum of the observed values in the training data.

Our approach requires additional conditional independence assumptions between the observed information and the underlying biometric state process driving the contractual payments, but in return, one only needs two extra model elements in addition to what is usually modeled in the multistate approach, namely the reporting delay distribution and the adjudication probabilities. Furthermore, one obtains relatively simple closed-form expressions for the reserves, eliminating the need for Monte-Carlo simulation. The derivation of the reserves is based on stochastic analysis that falls outside of existing results and techniques, because we are led to analyze the biometric state process stopped at a random time that is not a stopping time with respect to the filtration of interest, namely the filtration generated by the biometric state process. Such complications did not arise in the simple model from Example 5.8 of Buchardt et al. (2023), and the treatment of the resulting mathematical complexities is another main contribution of the paper.

It is also relevant to consider whether the reserves could be based on aggregate models (e.g., chain ladder Mack (1993, 1999)) rather than individual reserving models given their popularity with practitioners, see e.g. Lopez et al. (2018) and the references therein. For aggregate models to be applicable, steady-state assumptions have to hold on an aggregate level. Steady-state assumptions at a portfolio level are unsuitable for disability insurance since disability claims frequently lead to several decades of benefit payments causing the proportion of long-lasting disabilities in the portfolio to rise for many decades. Assuming that an aggregate reserving model was available, it would likely still suffer from certain robustness issues. For example, a model based on chain ladder would be slow to capture trends such as the sharp rise in mental health-related disabilities that has been observed in recent years, while it is straightforward to include a calendar time effect in the proposed reserving models. As another example, consider an IBNR reserve that arises as some transformation of the classic semi-Markov reserves for the policies that are currently in the portfolio. Then an influx of new policies would lead to an unwarranted increase in reserves; the aggregate IBNR reserve should initially remain unchanged since disabilities that occurred before entering the portfolio do not lead to disability benefits. Covariate shifts in the portfolio would also violate steady-state assumptions, while individual reserving models are robust to such shifts whenever the covariate is included in the model.

A general disadvantage of individual models is that they often lead to higher estimation risk since more elements have to be estimated. They may also lead to higher model risk since more assumptions are needed to construct the models. In this paper, we seek to accommodate the former by deriving models that do not require many new model elements. To accommodate the latter, we provide detailed discussions on how to adjust the models when central assumptions of the setup are violated. Methods for detecting deviations between the models and the realized outcomes are given in Remark 4.B.1 and Theorem 5.10 of Buchardt et al. (2023), making it possible to monitor the estimation and model risk. The proposed models thus possess many properties that could make them attractive for practitioners.

The paper is structured as follows. Section 4.2 describes our valid time and transaction time model for disability insurance schemes. Section 4.3 concerns reserving and contains the main results. Estimation is discussed in Section 4.4. Section 4.5 contains a real data application. Section 4.6 concludes. Lengthy proofs are deferred to Appendix 4.A and the straightforward extension to stochastic interest rates is given in Appendix 4.B.

4.2 Setup

4.2.1 Disability insurance in valid time

Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ be a filtered background probability space. The biometric state of the insured is governed by a non-explosive pure jump process $Y : \Omega \times \mathbb{R}_+ \mapsto \mathcal{J}$ on a finite state space $\mathcal{J} = \{1, 2, ..., J\}$ for $J \in \mathbb{N}$ with deterministic initial state y_0. Denote by N the corresponding multivariate counting process with components

$N_{jk} : \Omega \times \mathbb{R}_+ \mapsto \mathbb{N}_0$ $(j, k \in \mathcal{J}, k \neq j)$ given by

$$N_{jk}(t) = \#\{s \in (0, t] : Y_{s-} = j,\ Y_s = k\}.$$

For $\mathcal{A} \subseteq \mathcal{J}$, let $\tau_{\mathcal{A}} : \Omega \to \overline{\mathbb{R}}_+$ be the first hitting time of $\mathcal{A}$ such that $\tau_{\mathcal{A}} = \inf\{t \geq 0 : Y_t \in \mathcal{A}\}$. The information generated by Y is represented by the filtration $\mathcal{F}_t^Y = \sigma(Y_s, s \leq t)$. We shall also need the future information $\mathcal{F}^{t,Y} = \sigma(Y_s, s \geq t)$. Let $U : \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}_+$ be the duration in the current state,

$$U_t = t - \sup\{s \in (0, t] : Y_s \neq Y_t\}.$$

A life insurance contract between the insured and the insurer is stipulated by the specification of the accumulated cash flow $B : \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}$ representing the accumulated benefits less premiums. We refer to this as the valid time cash flow or the contractual payments, and assume it is on the usual semi-Markov form,

$$B(\mathrm{d}t) = \sum_{j=1}^{J} 1_{(Y_{t-}=j)} B_{j,t-U_{t-}}(\mathrm{d}t) + \sum_{\substack{j,k=1 \\ j \neq k}}^{J} b_{jk}(t, U_{t-}) N_{jk}(\mathrm{d}t), \quad B(0) \in \mathbb{R},$$

where $B_{j,w} : \mathbb{R}_+ \mapsto \mathbb{R}$ $(j \in \mathcal{J}, w \geq 0)$ are measurable, càdlàg and of finite variation, and $b_{jk}(t, u)$ $(j, k \in \mathcal{J}, j \neq k)$ are measurable and bounded. Since $t - U_{t-}$ is piecewise constant, the above expression is well-defined. We bundle all the processes that determine the payments into $X_t = (t, Y_t, U_t)$. Note that $\mathcal{F}_t^X = \mathcal{F}_t^Y$ since U is constructed from the history of Y. Like Buchardt et al. (2023), we name the model for X and B the *valid time model*. In this paper, we assume that the biometric state process Y takes values in the state space $\mathcal{J}$ depicted in Figure 4.1.

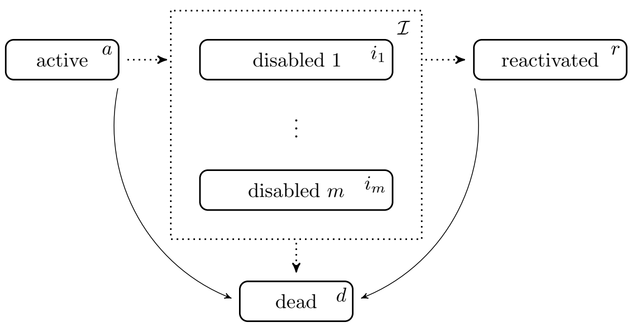

Figure 4.1: *The state process Y takes values in $\mathcal{J} = \{a, i_1, \ldots, i_m, r, d\}$, being an illness-death model with m disabled states $\mathcal{I} = \{i_1, \ldots, i_m\}$ and a separate reactivated state. To reduce clutter, all transitions to and from $\mathcal{I}$ are illustrated as single dotted arrows. Transition between the disabled states is not possible.*

Note that we have here labeled the states with letters instead of integers to stay consistent with the actuarial literature. This should cause no confusion in what follows. We assume $y_0 = a$ since only non-disabled are offered the insurance. We assume that Y is a semi-Markov process with measurable transition hazards $\mu_{jk} : \mathbb{R}_+^2 \mapsto \mathbb{R}_+$ ($j, k \in \mathcal{J}, j \neq k$) which are Lebesgue-integrable on compact subsets of $\mathbb{R}_+^2$ so that the intensity process for N_{jk} is given by $\lambda_{jk}(t) = 1_{(Y_{t-}=j)}\mu_{jk}(t, U_{t-})$. That Y is semi-Markov implies that X is Markov. The assumption regarding the existence of transition hazards (as opposed to cumulative transition hazards) could be removed using the techniques of Jacobsen (2006) or Helwich (2008) and it would similarly not be difficult to allow for an uncountable number of disabled states e.g. $\mathcal{I} = (0, 1]$ representing the degree of lost earning capacity. The choice and implications of the chosen state space are discussed in Remark 4.2.1.

Remark 4.2.1. (Valid time state space for disability insurance contract.)
In the multistate modeling literature, one usually allows for a general finite state space. We restrict our attention to the particular state space depicted in Figure 4.1 because, as was noted in the introduction, the relation between the valid time and transaction time reserves can be highly model-specific. The state space is intended to be general enough to capture most common disability insurance schemes. The hierarchical structure is imposed to simplify the transaction time model construction by making it so that there is only one disability and reactivation time to keep track of, as well as making the implementation in Section 4.5 easier since one can avoid implementing the semi-Markov Kolmogorov forward differential equations known from Buchardt et al. (2015) and instead use Thiele's differential equations successively.

Modeling disability insurance contracts using the model from Figure 4.1 implies that at most one disability can occur, that the disability type does not change after disablement, and that a reactivation of this disability is permanent. For contracts where it is important to model temporary reactivations, one might instead prefer to use a non-hierarchical illness-death model where reactivations are modeled as jumps back into state a instead of into the separate reactivated state r, see e.g. Figure 3 in Helwich (2008) or Example 2.1 in Christiansen (2012).

When coverage periods are short, as is usually the case for disability insurance schemes, and the disability hazard is small, ignoring the possibility of several disabilities can be reasonable. Even if one uses the non-hierachical model, it can be complicated to allow for several distinct disabilities if the insurance contract includes a coverage period. To see this, consider the situation where the insured becomes disabled within the coverage period of a disability annuity, reactivates, and becomes disabled again outside of the coverage period. Whether the insured is qualified for disability payments for the second disability depends on whether or not it was caused by the disability event in the coverage period.

The most natural way to capture this is to choose the disability event times to be those that lead to payout when estimating the disability hazard or to extend X

such that it contains information about which disability event is causing the current disablement. An alternative would be to let the payment rate in the disabled state be the average payment rate conditional on the historical development of X, see Remark 4.3.11 for more details. These approaches would all lead to models with an intricate dependence on the past coverage periods and the historical development of X. In Remark 4.3.11 we propose ways to obtain consistent reserves in situations where there may be several disabilities and/or transition between the disability types without having to use the non-hierarchical state space. ∇

In order to formulate the transaction time model in the next section, it is convenient to introduce some marked point process notation. In general, all our processes are assumed to be constructed according to the canonical approach of Jacobsen (2006), which among other things implies a specific regular conditional distribution used in the conditional distributions and conditional expectations. We note that X takes values in a Borel-space which we denote $(E, \mathcal{E})$. Write ∇ for the irrelevant mark and $\overline{E} = E \cup \{\nabla\}$. Let

$$K = \{((t_n)_{n\in\mathbb{N}}, (x_n)_{n\in\mathbb{N}}) \in \overline{\mathbb{R}}_+^{\mathbb{N}} \times \overline{E}^{\mathbb{N}} : t_1 \leq t_2 \leq \ldots \uparrow \infty, t_n < t_{n+1} \text{ if } t_n < \infty, \\ \text{and } x_n \in E \text{ iff } t_n < \infty\}$$

denote the space of sequences of jump times and jump marks and let this be equipped with the σ-algebra $\mathcal{K}$ generated by the coordinate projections

$$T_n^\circ((t_k)_{k\in\mathbb{N}}, (x_k)_{k\in\mathbb{N}}) = t_n, \quad X_n^\circ((t_k)_{k\in\mathbb{N}}, (x_k)_{k\in\mathbb{N}}) = x_n$$

for $n \in \mathbb{N}$. Let the stochastic process $H : \Omega \times \mathbb{R}_+ \mapsto K$ be the marked point process history of X. The value H_t consists of the ordered sequences of jump times $(\tau_{\{j\}} \times 1_{(\tau_{\{j\}} \leq t)})_{j\in\mathcal{J}}$ and corresponding jump marks $(X_{\tau_{\{j\}}} \times 1_{(\tau_{\{j\}} \leq t)})_{j\in\mathcal{J}}$ followed by a sequence of ∞ and ∇ respectively. Note that this representation of the jump times and jump marks only holds when the model is hierarchical. Since X is a piecewise deterministic process, there exists a measurable function $f : K \times \mathbb{R}_+ \mapsto E$ with the property that $X_t = f_{H_t}(t)$.

4.2.2 Disability insurance in transaction time

As was pointed out in Buchardt et al. (2023), it may sometimes be unreasonable to assume that the insurer has observed $\mathcal{F}_t^X$ at time t, since there can be reporting and processing delays for disability claims. In this case, we also cannot assume that $B(t)$ has been paid out at time t. Consequently, we introduce a stochastic process $\mathcal{Z} : \Omega\times\mathbb{R}_+ \mapsto S$, where $(S, \mathcal{S})$ is a Borel-space, which generates the insurer's available information $\mathcal{F}_t^{\mathcal{Z}}$. We furthermore introduce the stochastic process $\mathcal{B} : \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}$ modeling the accumulated observed payments which by construction will be $\mathcal{F}^{\mathcal{Z}}$-adapted, measurable, càdlàg and of finite variation. We refer to it as the transaction time cash flow. The model for $\mathcal{Z}$ and $\mathcal{B}$, which we now specify, is called the *transaction time model*.

Information

As a first coordinate of $\mathcal{Z}$ we define the right-continuous pure jump process $Z^{(1)}$: $\Omega \times \mathbb{R}_+ \mapsto \mathcal{J}^{(1)}$ taking values in the state space $\mathcal{J}^{(1)} = \{1, 2, 3, 4, 5\}$ depicted in Figure 4.2.

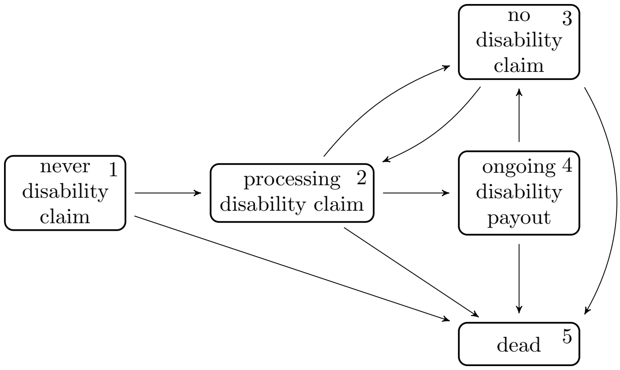

Figure 4.2: *State space $\mathcal{J}^{(1)}$ for the process $Z^{(1)}$.*

The process $Z^{(1)}$ represents the state of the claim settlement and it holds that $Z^{(1)}_0 = 1$. We introduce another coordinate of $\mathcal{Z}$ denoted $Z^{(2)} : \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}_+$ which represents the time of the disability event as reported by the insured in connection with a claim. We require $t \mapsto Z^{(2)}_t$ to be increasing and piecewise constant, and that its value can only increase upon a jump of $Z^{(1)}$ into state 2. Furthermore, we require $Z^{(2)}_t \leq t$ and that $Z^{(2)}_t$ stays constant after a nonzero amount of disability benefits have been awarded. How disability benefits are awarded is formalized later in this section. The interpretation is that when a disability claim is reported, the insurer also reports at which past time the disability occurred. Furthermore, different disability claims are allowed, but only until one of the claims is awarded, and the claims must furthermore always be reported in the same order as their chronological ordering. We let $Z^{(2)}_0 = 0$ as a convention. We also introduce a coordinate $Z^{(3)} : \Omega \times \mathbb{R}_+ \mapsto \mathcal{I}$ which represents the disability type that is reported in connection with a claim, assume that it only changes when $Z^{(2)}$ changes, and set $Z^{(3)}_0 = i_1$ as a convention. Finally, denote the counting processes related to $Z^{(1)}$ by $N^{(1)}_{jk} : \Omega \times \mathbb{R}_+ \mapsto \mathbb{N}_0$ $(j, k \in \mathcal{J}^{(1)}, j \neq k)$ and denote by $T_{\{j\}} = \inf\{s \in [0, \infty) : Z^{(1)}_s = j\}$ the first hitting time of state j by $Z^{(1)}$.

In addition to observing $Z = (Z^{(1)}, Z^{(2)}, Z^{(3)})$, the insurer observes what is being awarded to the insured; for example, whether a jump from state 2 to state 3 of $Z^{(1)}$ was accompanied by a payout of disability benefits in the form of *backpay* or not. The term backpay refers to a payout of overdue payments that have been delayed by reporting and processing delays, and such payments appear in the transaction time cash flow $\mathcal{B}$ constructed later in this section.

Knowing what is awarded to the insured however contains more information than simply knowing the realized payments since awarding disability benefits may not immediately lead to the commencement of payments if the adjudication is completed before the qualifying period ends. What is awarded to the insured is encoded in the bi-temporal stochastic process $H : \Omega \times \mathbb{R}_+^2 \mapsto K$, where H_s^t is interpreted as the value of H_s based on the information available at time t. We sometimes refer to t as the observational time and s as the historical time. We also introduce $X : \Omega \times \mathbb{R}_+^2 \mapsto E$ given by $X_s^t = (s, Y_s^t, U_s^t) = f_{H_s^t}(s)$, which is interpreted as the value of X_s based on the available information at time t. Similarly, introduce the bi-temporal counting processes $N_{jk} : \Omega \times \mathbb{R}_+^2 \mapsto \mathbb{N}_0$ which we denote $N_{jk}^t(s)$ with analogous interpretation. In total, we let $\mathcal{Z}_t = (Z_t, H_t^t)$.

To specify a model for H_s^t, we introduce an auxiliary stochastic process $G : \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}_+$, where G_t marks the beginning of the period where the insured would be eligible for additional disability when standing at time t, and which is given by

$$\mathrm{d}G_t = \mathrm{d}Z_t^{(2)} + 1_{(Z_t^{(1)}=4)}\,\mathrm{d}t + \sum_{k\in\{3,4,5\}} \delta_t^{2k}\,\mathrm{d}N_{2k}^{(1)}(t), \qquad G_0 = 0,$$

for stochastic processes $\delta^{2k} : \Omega\times\mathbb{R}_+ \mapsto \mathbb{R}_+$ with $k \in \{3,4,5\}$ satisfying $\delta_t^{24} = t - G_{t-}$ and $\delta_t^{2k} \in [0, t - G_{t-})$ when $k \in \{3,5\}$. The fact that $Z_t^{(2)}$ was required to be constant after a non-zero amount of disability claims have been awarded corresponds to saying that it stays constant after time t if $G_t > Z_t^{(2)}$. The interpretation of the specification of G is that the insured is disabled if they have an ongoing disability payout (in other words: the insurance company is not able to retract disability benefits that have been paid out), which is captured by G increasing with a rate of 1 in state 4 as well as G jumping to the value t if a jump from state 2 to state 4 occurs at time t. If the insured was eligible for additional disability benefits but is no longer disabled at the time of payout, this is captured by an increase in G upon a jump from state 2 to state 3 or to state 5.

From G, we can create other processes of interest such as $W : \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}_+$, being the number of time units the insured has been eligible for disability benefits using the current information, which is given by

$$W_t = G_t - Z_t^{(2)}.$$

We use W instead of writing expressions in terms of G whenever it eases interpretation. Some immediate properties are that G_t and W_t are increasing and $W_t \leq G_t \leq t$.

We now specify H_s^t. The bi-temporal process H_s^t contains $(T_{\{5\}}, d)$ on $(T_{\{5\}} \leq s, T_{\{5\}} \leq t)$, $(Z_t^{(2)}, Z_t^{(3)})$ on $(Z_t^{(2)} \leq s, 0 < W_t)$, and (G_t, r) on $(G_t \leq s, 0 < W_t, Z_t^{(1)} \neq 4, G_t \neq T_{\{5\}})$. Thus, events enter H_s^t when the corresponding jump time exceeds s and some t-related criterion is satisfied: Death has to have occurred before time t for the death event to be included, the insured has to have been deemed eligible for disability benefits for a nonzero amount of time by time t for the disability event to be included, and if it additionally holds that the payout of disability benefits has been stopped at time t and this wasn't due to death then the reactivation event is also included. As detailed later in this section, the transaction time cash flow $\mathcal{B}$ is constructed such that it is always in accordance with H_s^t. Therefore, H_s^t determines the payments and its specification is consequently essential to how the proposed transaction time model works.

To complete the modeling setup for the observations, we need to specify how $\mathcal{Z}$ is related to X. Note that $t \mapsto H_s^t$ is constant between jumps of $Z^{(1)}$, and assume $T_{\{5\}}$ is finite almost surely and that $H_s^t = H_s$ for $t \geq T_{\{5\}}$. In other words, $Z^{(1)}$ is the driver of new information arriving and the valid time process X_s is the limit of the transaction time process X_s^t when the observation time t tends to infinity. It follows that the *basic bi-temporal structure assumptions* introduced in Buchardt et al. (2023) are satisfied, and we may hence use their results. Note that with this specification of the transaction time model, once a disability claim has triggered some benefits, everything that happens afterward relates to this disability and no other disabilities can be reported. This is similar in spirit to how it was assumed that at most one disability could occur in the valid time model depicted in Figure 4.1.

Remark 4.2.2. (Granularity of the information.)
The above transaction time model presumes that the actuary has access to relatively granular information. It might be that the actuary does not receive information about reported claims, but is only notified about payouts. In that case, one may work with alternative transaction time models such as the $Z^{(1)}$ model depicted in Figure 4.3.

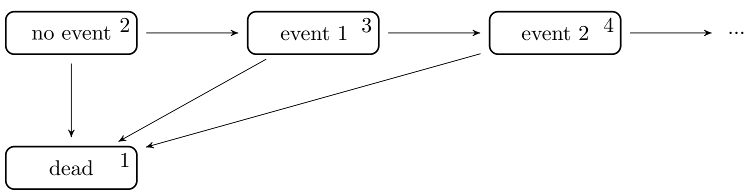

Figure 4.3: *Coarser state space for the process $Z^{(1)}$.*

Here the different events could represent starting running payments, stopping running payments, or awarding backpay. It would be natural to let $Z^{(2)} : \Omega \times \mathbb{R}_+ \mapsto \{0, 1\}$ be an indicator of whether there is running payments in the current state, and allow for a stochastic amount (including zero) of disability benefits to be awarded when jumping from one state to the next. The methods presented in this paper for the granular case can be adapted to handle this coarser case as well. Our methods do not apply if individual data is not available. $\triangledown$

Payments

We now specify how the transaction time cash flow $\mathcal{B}$ is related to the valid time cash flow B. We first introduce the time value of money. A detailed treatment may be found in Norberg (1990). Let $\kappa : \mathbb{R}_+ \mapsto \mathbb{R}_+$ be some deterministic strictly positive càdlàg *accumulation function* with initial value $\kappa(0) = 1$. A common choice is $\kappa(t) = \exp\left(\int_{(0,t]} r(v)\,\mathrm{d}v\right)$ for some deterministic integrable function $r : \mathbb{R}_+ \mapsto \mathbb{R}$ called the force of interest. The corresponding *discount function* is $t \mapsto 1/\kappa(t)$. Introduce the auxiliary stochastic process $\mathcal{B}' : \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}$ satisfying

$$\mathcal{B}'(t) = \sum_{j=1}^{J} \int_{[0,t]} \frac{\kappa(t)}{\kappa(s)} 1_{(Y^t_{s-}=j)} B_{j,s-U^t_{s-}}(\mathrm{d}s) + \sum_{\substack{j,k=1\\ j\neq k}}^{J} \int_{[0,t]} \frac{\kappa(t)}{\kappa(s)} b_{jk}(s, U^t_{s-}) N^t_{jk}(\mathrm{d}s)$$

which is the time t value of the payments generated by $(X^t_s)_{s\leq t}$. We then specify the transaction time cash flow $\mathcal{B} : \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}$ such that

$$\int_{[0,t]} \frac{\kappa(t)}{\kappa(s)} \mathcal{B}(\mathrm{d}s) = \mathcal{B}'(t). \tag{4.2.1}$$

The payments $\mathcal{B}$ are constructed such that the accumulated payments in real-time are always congruent with the most recent marked point process history. The way discounting is incorporated is closely connected to the notion of no arbitrage; the insurance company should not be able to keep the additional interest that they would earn by delaying the payout of benefits, so payments are accumulated with interest from the relevant historical time to the current observational time.

An explicit construction of a $\mathcal{B}$ satisfying Equation (4.2.1) and further discussions are given in Buchardt et al. (2023). That their definition of $\mathcal{B}$ satisfies Equation (4.2.1) follows by using their Proposition 5.3 upon dividing with $\kappa(t)$ on both sides and taking the difference between evaluating in 0 and t to see

$$\int_{(0,t]} \frac{1}{\kappa(s)} \mathcal{B}(\mathrm{d}s) = \frac{1}{\kappa(t)} \mathcal{B}'(t) - \mathcal{B}(0)$$

and then isolating $\mathcal{B}'(t)$.

In Appendix 4.B, we extend the results of Section 4.3 to stochastic interest rates that are independent of the valid and transaction time models and furthermore provide a way to validate the non-financial parts of the model. While these results are important for practical applications, they are relatively straightforward to derive and are somewhat orthogonal to the rest of the paper. Consequently, they have been deferred to the appendix.

4.3 Reserving

We now introduce the valid time and transaction time reserves, which are the focal point of the remainder of the paper. The present value in the valid time model $P : \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}$ and the present value in the transaction time model $\mathcal{P} : \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}$ are defined as

$$P(t) = \int_{(t,\infty)} \frac{\kappa(t)}{\kappa(s)} B(\mathrm{d}s), \quad \mathcal{P}(t) = \int_{(t,\infty)} \frac{\kappa(t)}{\kappa(s)} \mathcal{B}(\mathrm{d}s),$$

respectively. We assume $P(t)$ and $\mathcal{P}(t)$ are integrable for any t. The corresponding valid time reserve $V : \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}$ and transaction time reserve $\mathcal{V} : \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}$ are then defined as

$$V(t) = \mathbb{E}[P(t) \mid \mathcal{F}_t^X], \quad \mathcal{V}(t) = \mathbb{E}[\mathcal{P}(t) \mid \mathcal{F}_t^{\mathcal{Z}}].$$

Since $P(t)$ is $\mathcal{F}^{t,X}$-measurable and X is Markov, we get $V(t) = \mathbb{E}[P(t) \mid X_t]$ almost surely. We also introduce the state-wise reserves $V_j : \mathbb{R}_+^2 \mapsto \mathbb{R}$ given by $V_j(t,u) = \mathbb{E}[P(t) \mid X_t = (t,j,u)]$ $(t \geq 0, j \in \mathcal{J}, u \geq 0)$, which are measurable functions satisfying $V_{Y_t}(t, U_t) = V(t)$ almost surely. The choice of $V_j(t,u)$ is not unique, but we follow the convention from the literature, which is to choose the version where the transition probabilities satisfy the Chapmann-Kolmogorov equations surely, confer with Jacobsen (2006) p. 158-159 for the construction of such a version. In most applications, the specific choice of the state-wise reserves will not matter, since they will always be evaluated in X_t. This is not the case in our application, so we make this choice explicit. The choice is also important if one is interested in path properties of $V(t)$, see e.g. Christiansen & Furrer (2021). Our choice of state-wise reserves agrees with that of Christiansen & Furrer (2021), confer with the proof of Theorem 5.10 in Buchardt et al. (2023).

For $K \in \mathcal{A}$ we write

$$V_j(t,u;K) = \mathbb{E}[P(t) \mid X_t = (t,j,u), K]$$

for the reserve when we additionally condition on the event K. It is defined as $V_j(t,u;K) = \mathbb{E}[1_K P(t) \mid X_t = (t,j,u)]/\mathbb{P}(K \mid X_t = (t,j,u))$ with the convention $0/0 = 0$. This equals $\mathbb{E}[P(t) \mid X_t = (t,j,u), 1_K]$ on the event K.

Categorization

Using the setup introduced in Section 4.2.2, we can categorize the reserve at different points in time according to the usual claims reserving terminology known from e.g. Norberg (1993).

- On $(\text{Settled}_t) = (Z_t^{(1)} = 5)$, the claim and reserve are classified as *settled*, because even if some payments may remain in the dead state, the total claim size is known exactly.

- On $(\mathrm{RBNS}_t) = (Z_t^{(1)} \in \{2,3,4\})$, the claim and reserve are *reported-but-not-settled.*
- On $(\mathrm{CBNR}_t) = (Z_t^{(1)} = 1)$, the claim and reserve are *covered-but-not-reported.*

In the latter case, we will have an IBNR contribution for policies where the disability has already occurred, and a *covered-but-not-incurred* (CBNI) contribution for policies where no disability has occurred yet. We therefore also introduce the events $(\mathrm{IBNR}_t) = (\mathrm{CBNR}_t, \tau_{\mathcal{I}} \le t)$ and $(\mathrm{CBNI}_t) = (\mathrm{CBNR}_t, \tau_{\mathcal{I}} > t)$ which are however not known from the information available at time t. We also define $(\mathrm{RBNSr}_t) = (\mathrm{RBNS}_t, W_t > 0)$ and $(\mathrm{RBNSi}_t) = (\mathrm{RBNS}_t, W_t = 0)$, being RBNS where benefits have and have not been awarded respectively. In the former case, the time of disability is known and the time of reactivation is not fully determined, while in the latter case, both are not fully determined. The latter case is sometimes referred to as reported-but-not-paid in the literature, see e.g. Bettonville et al. (2021), and the former could consequently be called paid-but-not-settled, but these terms are not used in the current paper.

Independence

For the reserves to be tractable, we need some restriction on the conditional distribution of X given $\mathcal{F}_t^{\mathcal{Z}}$. Conditional independence criteria provide a natural way to impose such restrictions. We find it desirable to assume that the transaction time information only affects the distribution of future values X by affecting the probability that a certain valid time outcome was the true realization, and thus provides no additional information if the true valid time outcome was known. This leads to tractable reserves and is often a reasonable assumption, and even if it is not, the violation can often be remedied by extending the valid time model, see the discussion in Remark 4.3.10. When the insured is dead all is known and so no independence assumption is needed. Formally, we thus assume:

Assumption 4.3.1. *(Influence of transaction time information.)*
On (CBNR_t)

$$\sigma((X_s)_{s \ge G_t}) \perp\!\!\!\perp \mathcal{F}_t^{\mathcal{Z}} \mid 1_{(\tau_{\{d\}} \le t)}, X_{\tau_{\mathcal{I}}} 1_{(\tau_{\mathcal{I}} \le t)}.$$

On (RBNS_t)

$$\sigma((X_s)_{s \ge G_t}) \perp\!\!\!\perp \mathcal{F}_t^{\mathcal{Z}} \mid 1_{(\tau_{\{d\}} \le t)}, X_{G_t}. \qquad \diamond$$

Here X_{G_t} is understood as the composite stochastic variable $\omega \mapsto (X_{G_t(\omega)})(\omega)$. Note that on both events, it holds that $\tau_{\{d\}} > t$ and one could thus have replaced $1_{(\tau_{\{d\}} \le t)}$ by the event $(\tau_{\{d\}} > t)$ in the conditioning. This assumption states that the distribution of the valid time behavior of two subjects after time G_t is exchangeable whenever the values of the variables in the conditioning agree for these subjects no matter the rest of the transaction time information. The effect of this assumption is that the transaction time information $\mathcal{F}_t^{\mathcal{Z}}$ only affects the distribution of the variables entering in the conditioning and not the rest of the valid time process.

As Assumption 4.3.1 stands, there is still some transaction time information remaining via X_{G_t} in the second case; for example $(X_{G_t} = (G_t, i, 0))$ for $i \in \mathcal{I}$ implies that the disability starting at time G_t is not awarded at time t. What is needed to remove this final piece of transaction time information is a strong Markov-type property at the random time G_t. This situation is non-standard since the random time where the process is stopped is not a stopping time, see however Yackel (1968) where a random time change with a non-stopping time is used to obtain a Markov process from a semi-Markov process. The phenomenon is nevertheless similar to how left-truncated processes are usually studied conditional on some event having occurred prior, where the event is measurable with respect to an enlarged filtration stopped at the left-truncation time, but might not be measurable with respect to the self-exciting filtration, see for example Section III.3 of Andersen et al. (1993). To obtain the strong Markov property at G_t, we impose Assumption 4.3.2.

Assumption 4.3.2. *(Conditional independence of stopped valid time process.)*
$\forall v, t \geq 0$:

$$\mathcal{F}^{v,X} \perp\!\!\!\perp \mathcal{F}_v^X \vee \sigma(X_{G_t}) \mid 1_{(\tau_{\{d\}} \leq t)}, X_v$$

on $(G_t \leq v, \mathrm{RBNS}_t)$. $\diamond$

This is analogous to how the Markov property allows one to discard $(X_u)_{u \leq v}$ in the conditional distribution of $(X_s)_{s \geq v}$ when also conditioning on X_v. Here we however need to keep the knowledge that death has not occurred. Under this assumption, we get the following strong Markov property.

Lemma 4.3.3. *(Strong Markov property at G_t.)*
Under Assumption 4.3.2,

$$\mathbb{P}((X_s)_{s \geq G_t} \in \cdot \mid (X_s)_{s \leq G_t}, 1_{(\tau_{\{d\}} \leq t)}) = \mathbb{P}((X_s)_{s \geq x_1} \in \cdot \mid X_{x_1} = x, 1_{(\tau_{\{d\}} \leq t)})\Big|_{x = X_{G_t}}$$

on (RBNS_t) *where* $x = (x_1, x_2, x_3)$.

The proof of Lemma 4.3.3 is long and is hence deferred to the appendix. Lemma 4.3.3 states that there is no extra knowledge gained about the distribution of X knowing that a path $(X_s)_{s \leq x_1}$ came from a transaction time realization with $x = X_{G_t}$ compared with just conditioning on $(X_s)_{s \leq x_1}$ when knowledge about survival until time t is retained.

4.3.1 CBNR reserve

We first consider being on (CBNR_t) and calculate the transaction time reserve. This is only of interest if $\mathbb{P}(\mathrm{CBNR}_t) > 0$, so we assume that this is the case. Introduce the *IBNR-factor* $I : \mathbb{R}_+^2 \times \mathcal{I} \mapsto [0,1]$ defined as

$$I_i(s,t) = \mathbb{P}(\mathrm{CBNR}_t \mid \tau_{\mathcal{I}} = s, Y_{\tau_{\mathcal{I}}} = i).$$

Note that this probability can also be expressed as the probability that the delay between the disability event and the first disability claim is larger than $t-s$. Introduce also the transition probabilities

$$p_{jk}(s,t,u,z) = \mathbb{P}(Y_t = k, U_t \leq z \mid Y_s = j, U_s = u).$$

The reserve for the CBNR case is given in Theorem 4.3.4.

Theorem 4.3.4. *(CBNR reserve.)*
On (CBNR_t)*, we have*

$$\begin{aligned}\mathcal{V}(t) &= V_a(t,t) \times \mathbb{P}(\tau_{\mathcal{I}} > t \mid \mathrm{CBNR}_t) \\ &\quad + \sum_{i\in\mathcal{I}} \int_{(0,t]} \left(\frac{\kappa(t)}{\kappa(s)} V_i(s,0;(\tau_{\{d\}} > t)) + \frac{\kappa(t)}{\kappa(s)} b_{ai}(s,s) - \int_{(s,t]} \frac{\kappa(t)}{\kappa(v)} B_{a,0}(\mathrm{d}v) \right) \\ &\qquad \times \frac{p_{aa}(0,s,0,\infty)}{\mathbb{P}(\mathrm{CBNR}_t)} I_i(s,t) \mu_{ai}(s,s)\,\mathrm{d}s.\end{aligned}$$

Furthermore,

$$\mathbb{P}(\tau_{\mathcal{I}} > t \mid \mathrm{CBNR}_t) = 1 - \int_{(0,t]\times\mathcal{I}} \frac{I_i(s,t)}{\mathbb{P}(\mathrm{CBNR}_t)} \, (\tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}})(\mathbb{P})(\mathrm{d}s, \mathrm{d}i)$$

and

$$\begin{aligned}\mathbb{P}(\mathrm{CBNR}_t) = &\frac{\int_{(t,\infty)\times\mathcal{I}} I_i(s,t) \, (\tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}})(\mathbb{P})(\mathrm{d}s, \mathrm{d}i)}{1 - \mathbb{P}(\tau_{\mathcal{I}} = \infty \mid \tau_{\mathcal{I}} > t, \tau_{\{d\}} > t)} \\ &+ \int_{(0,t]\times\mathcal{I}} I_i(s,t) \, (\tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}})(\mathbb{P})(\mathrm{d}s, \mathrm{d}i).\end{aligned}$$

Inserting $(\tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}})(\mathbb{P})(\mathrm{d}s, \mathrm{d}i) = p_{aa}(0,s,0,\infty)\mu_{ai}(s,s)\,\mathrm{d}s$ for $s \in (0,\infty)$ and $i \in \mathcal{I}$ as well as substituting

$$1 - \mathbb{P}(\tau_{\mathcal{I}} = \infty \mid \tau_{\mathcal{I}} > t, \tau_{\{d\}} > t) = \sum_{i\in\mathcal{I}} \int_{(t,\infty)} p_{aa}(t,s,t,\infty)\mu_{ai}(s,s)\,\mathrm{d}s,$$

and using Remark 4.3.6 to calculate $V_i(s,0;(\tau_{\{d\}} > t))$, one sees that the reserve in Theorem 4.3.4 is computable using only the usual valid time hazards and the new model element $I_i(s,t)$ for $s < \infty$.

Proof. Write

$$\begin{aligned}\mathcal{V}(t) &= \mathbb{E}[1_{(\tau_{\mathcal{I}}>t)}\mathcal{P}(t) + 1_{(\tau_{\mathcal{I}}\leq t)}\mathcal{P}(t) \mid \mathcal{F}_t^{\mathcal{Z}}] \\ &= \mathbb{E}[\mathcal{P}(t) \mid \mathcal{F}_t^{\mathcal{Z}}, \tau_{\mathcal{I}} > t]\mathbb{P}(\tau_{\mathcal{I}} > t \mid \mathcal{F}_t^{\mathcal{Z}}) + \mathbb{E}[1_{(\tau_{\mathcal{I}}\leq t)}\mathcal{P}(t) \mid \mathcal{F}_t^{\mathcal{Z}}].\end{aligned}$$

The first term is the CBNI reserve and the second term is the IBNR reserve.

We start by treating the CBNI reserve. Note that on $(\mathrm{CBNR}_t, \tau_{\mathcal{I}} > t)$, it holds that $\mathcal{P}(t) = P(t)$ by Theorem 5.4 of Buchardt et al. (2023), since it then holds that $(X_s)_{0\leq s\leq t} = (X_s^t)_{0\leq s\leq t}$. This leads to

$$\begin{aligned}1_{(\mathrm{CBNR}_t)}\mathbb{E}[\mathcal{P}(t) \mid \mathcal{F}_t^{\mathcal{Z}}, \tau_{\mathcal{I}} > t] &= 1_{(\mathrm{CBNR}_t)}\mathbb{E}[P(t) \mid \tau_{\{d\}} > t, \tau_{\mathcal{I}} > t] \\ &= 1_{(\mathrm{CBNR}_t)}V_a(t,t),\end{aligned}$$

by the first part of Assumption 4.3.1 and using that $P(t)$ is $\sigma((X_s)_{s\geq G_t})$-measurable since $G_t \leq t$.

For the IBNR reserve, we find by Theorem 5.4 of Buchardt et al. (2023) that on $(\mathrm{CBNR}_t, \tau_{\mathcal{I}} \leq t)$ it holds

$$\begin{aligned}\mathcal{P}(t) &= P(t) + \int_{[0,t]} \frac{\kappa(t)}{\kappa(s)}(B-\mathcal{B})(\mathrm{d}s) \\ &= \frac{\kappa(t)}{\kappa(\tau_{\mathcal{I}}-)}P(\tau_{\mathcal{I}}-) - \int_{[\tau_{\mathcal{I}},t]} \frac{\kappa(t)}{\kappa(s)}B_{a,0}(\mathrm{d}s) \\ &= \frac{\kappa(t)}{\kappa(\tau_{\mathcal{I}})}P(\tau_{\mathcal{I}}) + \frac{\kappa(t)}{\kappa(\tau_{\mathcal{I}})}b_{aY_{\tau_{\mathcal{I}}}}(\tau_{\mathcal{I}},\tau_{\mathcal{I}}) - \int_{(\tau_{\mathcal{I}},t]} \frac{\kappa(t)}{\kappa(s)}B_{a,0}(\mathrm{d}s)\end{aligned}$$

since the payments are at least equal until $\tau_{\mathcal{I}}$ on this event. Hence

$$\begin{aligned}&1_{(\mathrm{CBNR}_t)}\mathbb{E}[\mathcal{P}(t)1_{(\tau_{\mathcal{I}}\leq t)} \mid \mathcal{F}_t^{\mathcal{Z}}] \\ &= 1_{(\mathrm{CBNR}_t)}\mathbb{E}\bigg[1_{(\tau_{\mathcal{I}}\leq t)}\bigg(\frac{\kappa(t)}{\kappa(\tau_{\mathcal{I}})}P(\tau_{\mathcal{I}}) + \frac{\kappa(t)}{\kappa(\tau_{\mathcal{I}})}b_{aY_{\tau_{\mathcal{I}}}}(\tau_{\mathcal{I}},\tau_{\mathcal{I}}) \\ &\qquad\qquad - \int_{(\tau_{\mathcal{I}},t]} \frac{\kappa(t)}{\kappa(s)}B_{a,0}(\mathrm{d}s)\bigg) \,\bigg|\, \mathcal{F}_t^{\mathcal{Z}}\bigg] \\ &= 1_{(\mathrm{CBNR}_t)}\mathbb{E}\bigg[1_{(\tau_{\mathcal{I}}\leq t)}\bigg(\frac{\kappa(t)}{\kappa(\tau_{\mathcal{I}})}\mathbb{E}\big[P(\tau_{\mathcal{I}}) \mid X_{\tau_{\mathcal{I}}} \vee \mathcal{F}_t^{\mathcal{Z}}\big] + \frac{\kappa(t)}{\kappa(\tau_{\mathcal{I}})}b_{aY_{\tau_{\mathcal{I}}}}(\tau_{\mathcal{I}},\tau_{\mathcal{I}}) \\ &\qquad\qquad - \int_{(\tau_{\mathcal{I}},t]} \frac{\kappa(t)}{\kappa(s)}B_{a,0}(\mathrm{d}s)\bigg) \,\bigg|\, \mathcal{F}_t^{\mathcal{Z}}\bigg]\end{aligned}$$

by the tower property. Now note that on $(\mathrm{CBNR}_t, \tau_{\mathcal{I}} \leq t)$, we have

$$\begin{aligned}\mathbb{E}\big[P(\tau_{\mathcal{I}}) \mid X_{\tau_{\mathcal{I}}} \vee \mathcal{F}_t^{\mathcal{Z}}\big] &= \mathbb{E}\big[P(\tau_{\mathcal{I}}) \mid X_{\tau_{\mathcal{I}}}, \tau_{\{d\}} > t\big] \\ &= \frac{\mathbb{E}\big[P(s)1_{(\tau_{\{d\}}>t)} \mid X_s = (s,i,0)\big]}{\mathbb{P}(\tau_{\{d\}} > t \mid X_s = (s,i,0))}\bigg|_{s=\tau_{\mathcal{I}}, i=Y_{\tau_{\mathcal{I}}}} \\ &= V_{Y_{\tau_{\mathcal{I}}}}(\tau_{\mathcal{I}}, 0; (\tau_{\{d\}} > t))\end{aligned}$$

by the first part of Assumption 4.3.1 and the usual strong Markov property, see Theorem 7.5.1 of Jacobsen (2006), using $(\tau_{\{d\}} > t) \in \sigma(X_t)$. Hence we obtain

$$\begin{aligned}
&\mathbb{E}[\mathcal{P}(t)1_{(\tau_{\mathcal{I}} \leq t)} \mid \mathcal{F}_t^{\mathcal{Z}}] \\
&= \int_{(0,t]\times\mathcal{I}} \bigg(\frac{\kappa(t)}{\kappa(s)} V_i(s,0;(\tau_{\{d\}} > t)) \\
&\quad + \frac{\kappa(t)}{\kappa(s)} b_{ai}(s,s) - \int_{(s,t]} \frac{\kappa(t)}{\kappa(v)} B_{a,0}(\mathrm{d}v) \bigg) (\tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}} \mid \mathcal{F}_t^{\mathcal{Z}})(\mathbb{P})(\mathrm{d}s, \mathrm{d}i)
\end{aligned}$$

on (CBNR_t). Note that on (CBNR_t), we have

$$(\tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}} \mid \mathcal{F}_t^{\mathcal{Z}})(\mathbb{P})(\mathrm{d}s, \mathrm{d}i) = (\tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}} \mid \mathrm{CBNR}_t)(\mathbb{P})(\mathrm{d}s, \mathrm{d}i).$$

Using Bayes' theorem, see for example Theorem 1.31 in Schervish (1995), we can write

$$\begin{aligned}
(\tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}} \mid \mathrm{CBNR}_t)(\mathbb{P})(\mathrm{d}s, \mathrm{d}i) &= \mathbb{P}(\mathrm{CBNR}_t \mid \tau_{\mathcal{I}} = s, Y_{\tau_{\mathcal{I}}} = i) \frac{(\tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}})(\mathbb{P})(\mathrm{d}s, \mathrm{d}i)}{\mathbb{P}(\mathrm{CBNR}_t)} \\
&= I_i(s,t)\mu_{ai}(s,s) \frac{p_{aa}(0,s,0,\infty)}{\mathbb{P}(\mathrm{CBNR}_t)} \, \mathrm{d}s
\end{aligned}$$

for $s \in [0,\infty)$ and $i \in \mathcal{I}$. This also implies

$$\mathbb{P}(\tau_{\mathcal{I}} > t \mid \mathrm{CBNR}_t) = 1 - \sum_{i\in\mathcal{I}} \int_{(0,t]} I_i(s,t)\mu_{ai}(s,s) \frac{p_{aa}(0,s,0,\infty)}{\mathbb{P}(\mathrm{CBNR}_t)} \, \mathrm{d}s.$$

For the final part, note

$$\begin{aligned}
\mathbb{P}(\mathrm{CBNR}_t) &= \mathbb{E}[\mathbb{P}(\mathrm{CBNR}_t \mid \tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}})] \\
&= \int_{(0,\infty)\times\mathcal{I}} I_i(s,t) \, (\tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}})(\mathbb{P})(\mathrm{d}s, \mathrm{d}i) + \mathbb{P}(\mathrm{CBNR}_t, \tau_{\mathcal{I}} = \infty).
\end{aligned}$$

We have

$$\begin{aligned}
\mathbb{P}(\mathrm{CBNR}_t, \tau_{\mathcal{I}} = \infty) &= \mathbb{P}(\tau_{\mathcal{I}} = \infty \mid \mathrm{CBNR}_t)\mathbb{P}(\mathrm{CBNR}_t) \\
&= \mathbb{P}(\tau_{\mathcal{I}} = \infty \mid \tau_{\mathcal{I}} > t, \tau_{\{d\}} > t)\mathbb{P}(\tau_{\mathcal{I}} > t \mid \mathrm{CBNR}_t)\mathbb{P}(\mathrm{CBNR}_t)
\end{aligned}$$

using the first part of Assumption 4.3.1 in the second equality. Inserting these expressions, isolating for $\mathbb{P}(\mathrm{CBNR}_t)$, and simplifying gives

$$\begin{aligned}
\mathbb{P}(\mathrm{CBNR}_t) = &\frac{\int_{(t,\infty)\times\mathcal{I}} I_i(s,t) \, (\tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}})(\mathbb{P})(\mathrm{d}s, \mathrm{d}i)}{1 - \mathbb{P}(\tau_{\mathcal{I}} = \infty \mid \tau_{\mathcal{I}} > t, \tau_{\{d\}} > t)} \\
&+ \int_{(0,t]\times\mathcal{I}} I_i(s,t) \, (\tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}})(\mathbb{P})(\mathrm{d}s, \mathrm{d}i).
\end{aligned}$$

Collecting the results, we obtain the statement of the theorem. $\square$

Remark 4.3.5. (Relation to non-life insurance Poisson models.)
For the IBNR term, the time s disability rate $\mu_{ai}(s,s)$ has to be multiplied by the IBNR-factor $\mathbb{P}(\mathrm{CBNR}_t \mid \tau_{\mathcal{I}} = s, Y_{\tau_{\mathcal{I}}} = i)$ similarly to the Poisson process model in Norberg (1999). Heuristically, one has to hold a disability reserve for the expected number of disabilities $\mu_{ai}(s,s)\,\mathrm{d}s$ at a prior time s times the proportion of insured that have yet to report their claim by time t, which is $\mathbb{P}(\mathrm{CBNR}_t \mid \tau_{\mathcal{I}} = s, Y_{\tau_{\mathcal{I}}} = i)$. The extra factor $p_{aa}(0,s,0,\infty)/\mathbb{P}(\mathrm{CBNR}_t)$ adjusts for the fact that there can be at most one disability occurrence in this model as opposed to a Poisson process model where there can be several occurrences. $\triangledown$

Remark 4.3.6. (Conditional semi-Markov model.)

In Section 6 of Hoem (1972), the author obtains an expression for the transition probabilities and hazards of a semi-Markov process conditional on not having entered a specific absorbing part of the state space before a given time. For our purposes, choose the transient states to be $\mathcal{J}\backslash\{d\}$. Calculating the transition probabilities that appear in $V_i(\,\cdot,\cdot\,;(\tau_{\{d\}} > t))$ can then be done as usual upon switching to the hazards $\mu_{\ell j}(s,u;t)$, where

$$\mu_{\ell j}(s,u;t) = \mu_{\ell j}(s,u)\frac{\sum_{k\in\mathcal{J}\backslash\{d\}} p_{jk}(s,t,0,\infty)}{\sum_{k\in\mathcal{J}\backslash\{d\}} p_{\ell k}(s,t,u,\infty)}$$

for $s < t$ and

$$\mu_{\ell j}(s,u;t) = \mu_{\ell j}(s,u)$$

for $s \geq t$. Consequently, the hazard of jumping to states where there is a higher probability of remaining in $\mathcal{J}\backslash\{d\}$ is increased and the hazard is decreased when there is a lower probability of remaining in $\mathcal{J}\backslash\{d\}$. After time t, the conditioning provides no additional information, and the hazards are equal to the hazards in the unconditional model. $\triangledown$

Remark 4.3.7. (Benefits of modeling reporting delays stochastically.)

One could also have considered a simpler model where reporting delays were deterministic. In the data application, the numerical value of such simple reserves are compared with those obtained with the proposed methods, see also Remark 4.5.1. A disadvantage of such an approach is that reporting delays are stochastic in reality, so the validity of the model would be less clear. Similarly, one might overlook the fact that the classic disability reserve is only the relevant "claim size" in the IBNR reserve if independence assumptions like Assumption 4.3.1 and 4.3.2 hold, see also the discussion in Remark 4.3.10.

There are also some disadvantages related to the size and timing of the reserve that might result from using a model with a deterministic reporting delay. A first-order error is, as discussed at the end of the introduction and in Remark 4.5.1, if the size of the portfolio increases by $x\%$ then the simple IBNR reserve becomes $x\%$ too high and vice versa. Changes to the size of the portfolio are relatively common for disability insurance due to the short coverage periods, leading the insured to have frequent opportunities to change their insurance provider. Second-order errors arise since the timing of the disability and the covariate dependence are handled slightly more imprecisely in the simple model. $\triangledown$

4.3.2 RBNS reserve

We now consider being on (RBNS_t). The reserve for the RBNSi case is given in Theorem 4.3.8.

Theorem 4.3.8. *(RBNSi reserve.)*
We have

$$\begin{aligned}\mathcal{V}(t) = &\frac{\kappa(t)}{\kappa(G_t)} V_a(G_t, G_t; (\tau_{\{d\}} > t)) \times (1 - \mathbb{P}(X_{G_t} = (G_t, i, 0) \mid \mathcal{F}_t^{\mathcal{Z}})) \\ &+ \left(\frac{\kappa(t)}{\kappa(G_t)} V_i(G_t, 0; (\tau_{\{d\}} > t)) + \frac{\kappa(t)}{\kappa(G_t)} b_{ai}(G_t, G_t) \right) \\ &\times \mathbb{P}(X_{G_t} = (G_t, i, 0) \mid \mathcal{F}_t^{\mathcal{Z}}) - \int_{(G_t, t]} \frac{\kappa(t)}{\kappa(s)} B_{a,0}(\mathrm{d}s)\end{aligned}$$

on (RBNSi_t) *with* $i = Z_t^{(3)}$.

Utilizing Remark 4.3.6 to calculate $V_i(G_t, 0; (\tau_{\{d\}} > t))$, and analogously calculate $V_a(G_t, G_t; (\tau_{\{d\}} > t))$, we see that the RBNSi reserve may be calculated using the usual valid time model and the new model element $\mathbb{P}(X_{G_t} = (G_t, Z_t^{(3)}, 0) \mid \mathcal{F}_t^{\mathcal{Z}})$ which we name the *adjudication probability*. This gives the probability that the reported disability claim will ultimately be awarded.

Proof. By similar calculations to the IBNR-case, we find

$$\mathcal{P}(t) = \frac{\kappa(t)}{\kappa(G_t-)} P(G_t-) - \int_{[G_t, t]} \frac{\kappa(t)}{\kappa(s)} \mathcal{B}(\mathrm{d}s).$$

The latter term is $\mathcal{F}_t^{\mathcal{Z}}$-measurable and on ($\mathrm{RBNSi}_t$) satisfies

$$\int_{[G_t, t]} \frac{\kappa(t)}{\kappa(s)} \mathcal{B}(\mathrm{d}s) = \int_{[G_t, t]} \frac{\kappa(t)}{\kappa(s)} B_{a,0}(\mathrm{d}s)$$

so the interesting part is $\kappa(t)/\kappa(G_t-) \times P(G_t-)$. On ($\mathrm{RBNSi}_t$) we have

$$\begin{aligned}\frac{\kappa(t)}{\kappa(G_t-)} P(G_t-) = &\frac{\kappa(t)}{\kappa(G_t)} P(G_t) + 1_{(X_{G_t} = (G_t, Z_t^{(3)}, 0))} \frac{\kappa(t)}{\kappa(G_t)} b_{aZ_t^{(3)}}(G_t, G_t) \\ &+ \frac{\kappa(t)}{\kappa(G_t)} B_{a,0}(\{G_t\}),\end{aligned}$$

which implies

$$\begin{aligned}
\mathcal{V}(t) &= \mathbb{E}\left[\frac{\kappa(t)}{\kappa(G_t)}P(G_t) + 1_{(X_{G_t}=(G_t,Z_t^{(3)},0))}\frac{\kappa(t)}{\kappa(G_t)}b_{aZ_t^{(3)}}(G_t,G_t) \mid \mathcal{F}_t^{\mathcal{Z}}\right] \\
&\quad - \int_{(G_t,t]}\frac{\kappa(t)}{\kappa(s)}B_{a,0}(\mathrm{d}s) \\
&= \frac{\kappa(t)}{\kappa(G_t)}\mathbb{E}[P(G_t) \mid \mathcal{F}_t^{\mathcal{Z}}, X_{G_t} = (G_t,a,G_t)]\mathbb{P}(X_{G_t} = (G_t,a,G_t) \mid \mathcal{F}_t^{\mathcal{Z}}) \\
&\quad + \frac{\kappa(t)}{\kappa(G_t)}\mathbb{E}[P(G_t) \mid \mathcal{F}_t^{\mathcal{Z}}, X_{G_t} = (G_t,Z_t^{(3)},0)]\mathbb{P}(X_{G_t} = (G_t,Z_t^{(3)},0) \mid \mathcal{F}_t^{\mathcal{Z}}) \\
&\quad + \frac{\kappa(t)}{\kappa(G_t)}b_{aZ_t^{(3)}}(G_t,G_t)\mathbb{P}(X_{G_t} = (G_t,Z_t^{(3)},0) \mid \mathcal{F}_t^{\mathcal{Z}}) - \int_{(G_t,t]}\frac{\kappa(t)}{\kappa(s)}B_{a,0}(\mathrm{d}s)
\end{aligned}$$

on (RBNSi$_t$). By the second part of Assumption 4.3.1 and Lemma 4.3.3, we get on (RBNSi$_t$) that

$$\begin{aligned}
\mathbb{E}[P(G_t) \mid \mathcal{F}_t^{\mathcal{Z}}, X_{G_t} = (G_t,a,G_t)] &= \mathbb{E}[P(G_t) \mid \tau_{\{d\}} > t, X_{G_t} = (G_t,a,G_t)] \\
&= V_a(G_t,G_t;(\tau_{\{d\}} > t))
\end{aligned}$$

and

$$\begin{aligned}
\mathbb{E}[P(G_t) \mid \mathcal{F}_t^{\mathcal{Z}}, X_{G_t} = (G_t,Z_t^{(3)},0)] &= \mathbb{E}[P(G_t) \mid \tau_{\{d\}} > t, X_{G_t} = (G_t,Z_t^{(3)},0)] \\
&= V_{Z_t^{(3)}}(G_t,0;(\tau_{\{d\}} > t))
\end{aligned}$$

using that all of $(X_s)_{s \leq G_t}$ is known in the conditioning for both cases. Collecting the results, we arrive at the desired expression. □

The reserve for the RBNSr case is given in Theorem 4.3.9.

Theorem 4.3.9. *(RBNSr reserve.)*
We have

$$\begin{aligned}
\mathcal{V}(t) = V_r(t,t-G_t) &\times (1 - \mathbb{P}(X_{G_t} = (G_t,i,W_t) \mid \mathcal{F}_t^{\mathcal{Z}})) \\
&+ \Bigg(\frac{\kappa(t)}{\kappa(G_t)}V_i(G_t,W_t;(\tau_{\{d\}} > t)) \\
&\quad - 1_{(Z_t^{(1)} \neq 4)}\Big(\frac{\kappa(t)}{\kappa(G_t)}b_{ir}(G_t,W_t) + \int_{(G_t,t]}\frac{\kappa(t)}{\kappa(s)}B_{r,G_t}(\mathrm{d}s)\Big)\Bigg) \\
&\times \mathbb{P}(X_{G_t} = (G_t,i,W_t) \mid \mathcal{F}_t^{\mathcal{Z}})
\end{aligned}$$

on (RBNSr$_t$) *with* $i = Z_t^{(3)}$.

Similarly to Theorem 4.3.8, one can use Remark 4.3.6 and the adjudication probability $\mathbb{P}(X_{G_t} = (G_t,Z_t^{(3)},W_t) \mid \mathcal{F}_t^{\mathcal{Z}})$ to calculate the RBNSr transaction time reserve. Note also that if $Z_t^{(1)} = 4$ then $\mathbb{P}(X_{G_t} = (G_t,Z_t^{(3)},W_t) \mid \mathcal{F}_t^{\mathcal{Z}}) = 1$ and furthermore $G_t = t$, $Z_t^{(3)} = Y_t$, and $W_t = U_t$. Thus, the expression collapses to $\mathcal{V}(t) = V_{Y_t}(t,U_t)$ which is the classic valid time disability reserve.

Proof. As in the RBNSi-case, we find

$$\mathcal{P}(t) = \frac{\kappa(t)}{\kappa(G_t-)} P(G_t-) - \int_{[G_t,t]} \frac{\kappa(t)}{\kappa(s)} \mathcal{B}(\mathrm{d}s).$$

The latter term is $\mathcal{F}_t^{\mathcal{Z}}$-measurable and on (RBNSr$_t$) satisfies

$$\begin{aligned}\int_{[G_t,t]} \frac{\kappa(t)}{\kappa(s)} \mathcal{B}(\mathrm{d}s) &= \frac{\kappa(t)}{\kappa(G_t)} B_{Z_t^{(3)}, G_t - W_t}(\{G_t\}) \\ &\quad + 1_{(Z_t^{(1)} \neq 4)} \Big(\frac{\kappa(t)}{\kappa(G_t)} b_{Z_t^{(3)} r}(G_t, W_t) + \int_{(G_t,t]} \frac{\kappa(t)}{\kappa(s)} B_{r,G_t}(\mathrm{d}s) \Big)\end{aligned}$$

so the interesting part is $\kappa(t)/\kappa(G_t-) \times P(G_t-)$. Proceeding in a similar manner as before, we see on (RBNSr$_t$) that

$$\begin{aligned}\frac{\kappa(t)}{\kappa(G_t-)} P(G_t-) &= \frac{\kappa(t)}{\kappa(G_t)} P(G_t) + 1_{(X_{G_t} = (G_t, r, 0))} \frac{\kappa(t)}{\kappa(G_t)} b_{Z_t^{(3)} r}(G_t, W_t) \\ &\quad + \frac{\kappa(t)}{\kappa(G_t)} B_{Z_t^{(3)}, G_t - W_t}(\{G_t\}).\end{aligned}$$

Now on (RBNSr$_t$) we have

$$\begin{aligned}\mathcal{V}(t) = &\frac{\kappa(t)}{\kappa(G_t)} \mathbb{E}\Big[P(G_t) \mid \mathcal{F}_t^{\mathcal{Z}}, X_{G_t} = (G_t, Z_t^{(3)}, W_t)\Big] \mathbb{P}(X_{G_t} = (G_t, Z_t^{(3)}, W_t) \mid \mathcal{F}_t^{\mathcal{Z}}) \\ &+ \frac{\kappa(t)}{\kappa(G_t)} \mathbb{E}\left[P(G_t) \mid \mathcal{F}_t^{\mathcal{Z}}, X_{G_t} = (G_t, r, 0)\right] \mathbb{P}(X_{G_t} = (G_t, r, 0) \mid \mathcal{F}_t^{\mathcal{Z}}) \\ &- 1_{(Z_t^{(1)} \neq 4)} \Bigg(\frac{\kappa(t)}{\kappa(G_t)} b_{Z_t^{(3)} r}(G_t, W_t) \mathbb{P}(X_{G_t} = (G_t, Z_t^{(3)}, W_t) \mid \mathcal{F}_t^{\mathcal{Z}}) \\ &+ \int_{(G_t,t]} \frac{\kappa(t)}{\kappa(s)} B_{r,G_t}(\mathrm{d}s) \Bigg).\end{aligned}$$

By the second part of Assumption 4.3.1 and Lemma 4.3.3, we see on (RBNSr$_t$) that

$$\begin{aligned}&\mathbb{E}\Big[P(G_t) \mid \mathcal{F}_t^{\mathcal{Z}}, X_{G_t} = (G_t, Z_t^{(3)}, W_t)\Big] \\ &= \mathbb{E}\Big[P(G_t) \mid \tau_{\{d\}} > t, X_{G_t} = (G_t, Z_t^{(3)}, W_t)\Big] \\ &= V_{Z_t^{(3)}}(G_t, W_t; (\tau_{\{d\}} > t))\end{aligned}$$

and

$$\begin{aligned}\mathbb{E}\left[P(G_t) \mid \mathcal{F}_t^{\mathcal{Z}}, X_{G_t} = (G_t, r, 0)\right] &= \mathbb{E}\left[P(G_t) \mid \tau_{\{d\}} > t, X_{G_t} = (G_t, r, 0)\right] \\ &= V_r(G_t, 0; (\tau_{\{d\}} > t)) \\ &= \frac{\kappa(G_t)}{\kappa(t)} V_r(t, t - G_t) + \int_{(G_t,t]} \frac{\kappa(G_t)}{\kappa(s)} B_{r,G_t}(\mathrm{d}s)\end{aligned}$$

using that $(X_s)_{s \leq G_t}$ is known in the former case, and towering on $(X_s)_{s<G_t}$ in the latter case, observing that the expectation given $(X_s)_{s \leq G_t}$ only depends on X_{G_t} and not $(X_s)_{s<G_t}$. Putting everything together, we arrive at the claimed result. □

Remark 4.3.10. (Weakening the independence assumptions.)
In the absence of Assumption 4.3.1 and 4.3.2, the valid time hazards and reserves could be influenced by additional transaction time information such as the disability reporting delay $U_{\mathcal{I}}$. This could be relevant if e.g. longer reporting delays were indicative of a more serious disability such that the intensity of N_{ir} with respect to the filtration $t \mapsto \mathcal{F}_t^X \vee \sigma(U_{\mathcal{I}})$ was $\lambda_{ir}(t) = 1_{(Y_{t-}=i)}\mu_{ir}(t, U_{t-}, U_{\mathcal{I}})$ with μ_{ir} being decreasing as a function of the last argument. Imposing an assumption like Assumption 4.3.1 when also conditioning on $U_{\mathcal{I}}$ would then for example lead to the award-term of the RBNSi reserve becoming

$$\mathbb{E}[P(G_t) \mid \mathcal{F}_t^{\mathcal{Z}}, X_{G_t} = (G_t, i, 0)] = V_i(G_t, 0, U_{\mathcal{I}}; (\tau^d > t))$$

with obvious notation. We briefly note that it is not a priori clear whether one would expect long reporting delays to be indicative of more or less severe disabilities. One could hypothesize that people with severe disabilities would find it more demanding to submit insurance claims. On the other hand, they might not need as much time to collect medical evidence if it is self-evident that they will be approved for disability benefits.

An alternative way to weaken Assumption 4.3.1 would be to incorporate more of the transaction time information in the valid time model e.g. by using different disabled states for different disability severities instead of a single disabled state. If the reporting delays only affect our estimate of the future trajectory of the valid time process X through the information they give us about the severity of the disability, Assumption 4.3.1 would be satisfied in the larger valid time model that incorporates information about severity of the disability. $\triangledown$

Remark 4.3.11. (On the simplifying assumptions.)
Some notable simplifications that have been made in order to arrive at tractable transaction time reserves are the independence assumptions (which were discussed in Remark 4.3.10), that there can be at most one disability event in the coverage period which reaches the payout stage, that the time and type of the disability event is completely known once benefits have started, and that the insurer cannot retract disability benefits. While it is true that none of these assumptions fully hold in practice, we believe that they are not seriously violated, and they only contribute with second-order effects compared to the main effects that have been included.

For example, the probability of experiencing q disabilities with independent causes is roughly equal to the disability hazard raised to the q'th power, so while more than one disability may occur in reality it is very uncommon. Similarly, there might be situations where the insurer and insured do not agree on the time of disablement leading to the time of disablement being changed to an earlier date after the insurer has started paying benefits. There may also be situations where retraction of disability benefits occurs if the insured willfully withheld information about their reactivation. However, the changes to the event times are probably sufficiently small and infrequent that the effect is negligible compared to other sources of error stemming from modeling, estimation, and forecasting error related to the biometric and financial model constituents. Accommodating transitions between the disabled states could be important depending on what the disability types represent. For example, they could represent different severities. If the severity of a disability changes often and different severities lead to substantially different payouts then this would be important to include in the model. We hence discuss possible remedies in the next paragraph.

If the aforementioned effects are sufficiently large to warrant explicit modeling, our models may still serve as a starting point. To accommodate the possibility of the time of disability changing after benefits have been awarded, one could for example add a term to the award-part of the RBNSi reserve equal to the the probability that the insured will be awarded benefits from an earlier time than G_t multiplied with the average additionally awarded amount for such cases. To accommodate multiple disabilities, a pragmatic approach could be to place the disability periods end-to-end in the estimation phase, such that an additional disability was treated as an annulment of a reactivation. This skews the time value of the cash flow but otherwise results in consistent reserves. It will however likely make the individual predictions less precise since the duration dependence will stem from a more heterogeneous population. For example, a long disability duration could stem from a single disability from which the insured has not reactivated, or it could be that they have reactivated from a long disability but then recently become disabled again.

If transitions between the disabled states were possible in the valid time state space one could formulate a transaction time model similar to the one specified towards the end of Example 5.8 in Buchardt et al. (2023). This would however result in a substantially more complicated model. We instead propose to let disability type i_k represent disabilities that start out as type i_k and estimate the reactivation, disabled mortality, and reactivated mortality hazards consistently with this. If the valid time disability payments also depend on the disability type, we propose to model these payments conditional on the initial disability type and the disability duration. By the tower property, this leads to the same reserves but in practice requires that one also estimates these conditional payments which brings the approach closer to the non-life insurance literature where the benefit sizes also have to be modeled. This approach may also be useful in other situations where the benefit size depends on more than the state and duration process; it could for example depend on whether the insured is receiving benefits from the government or other insurance companies. Taking the conditional expectation of the payments given the state and duration brings the problem back into something that can be represented in the usual semi-Markov framework.

As illustrated here, a benefit of having interpretable closed-form expressions for the reserves is that it is possible to reason about how to adjust the model when the underlying assumptions change. $\triangledown$

4.4 Estimation

To compute the transaction time reserves, one needs to estimate the valid time transition hazards, the IBNR-factor, and the adjudication probabilities. The IBNR-factor and adjudication probabilities are new model elements, and one hence needs to find a suitable way to estimate these. In addition, standard estimation procedures also do not apply for the valid time transition hazards since the data is contaminated by reporting delays and incomplete event adjudication.

For simplicity, we limit the discussion to the situation where there is at most one reported disability claim in the sense that $Z^{(2)}$ can increase only once. In this case, the delay between the disability time $\tau_{\mathcal{I}}$ and the time of the first reported disability $T_{\{2\}}$ equals the reporting delay of the disability event, the latter being the difference between $\tau_{\mathcal{I}}$ and the last time where a disability is reported. This is sufficient for the application in Section 4.5 and makes the statistical problem a special case of the one studied in Buchardt et al. (2025). The methods however easily generalize to the case where several distinct disabilities may be reported, in which case both the IBNR-factor and disability reporting delay distribution would need to be estimated in order to compute the IBNR reserve and estimate the disability hazard, respectively.

The data structure in Buchardt et al. (2025) consists of events that are reported with a delay and which may be confirmed or annulled upon adjudication. Their proposed estimation algorithm is a two-step procedure, where the first step is to estimate the adjudication probabilities and the reporting delay distribution, and the second step uses these to estimate the valid time hazards while correcting for contamination. Due to the assumed model for $\mathcal{Z}$, the adjudication probability $\mathbb{P}(X_{G_t} = (G_t, Z^{(3)}_t, W_t) \mid \mathcal{F}^{\mathcal{Z}}_t)$ can be calculated as an absorption probability for a suitable multistate model as in Buchardt et al. (2025). We henceforth refer to the transition hazards in the adjudication multistate model as adjudication hazards. Furthermore, estimating the IBNR-factor is equivalent to estimating the disability reporting delay distribution since

$$I_i(s,t) = \mathbb{P}(\mathrm{CBNR}_t \mid \tau_{\mathcal{I}} = s, Y_{\tau_{\mathcal{I}}} = i) = \mathbb{P}(T_{\{2\}} - \tau_{\mathcal{I}} > t - s \mid \tau_{\mathcal{I}} = s, Y_{\tau_{\mathcal{I}}} = i).$$

Imposing a parametric model for all three model elements hence results in an estimation problem that can be handled using Buchardt et al. (2025). More details are given in Appendix 4.C.

The estimator of the valid time hazards described in Appendix 4.C and employed in the data application outlined in Section 4.5 corresponds to the *Poisson approximation* from Buchardt et al. (2025) rather than the full estimator described in Section 3.3 of that paper. The approximation has the advantage of being simpler to implement and allowing one to estimate the different valid time hazards separately. Because the data employed in Section 4.5 only contains deaths recorded during the adjudication period, it is not possible to run the full estimation procedure of Buchardt et al. (2025). The mortality rates are nevertheless needed in order to calculate the reserves in the data application, so in Section 4.5 we employ the hazards given in Table 4.1 which are inspired by those published by the Danish Financial Supervisory Authority (FSA).

Remark 4.4.1. (Restricting the information used for adjudications.)
Note that Buchardt et al. (2025) allows one to reduce the information that the adjudication hazards depend on e.g. such that the reactivation adjudication hazards are not conditional on the disability reporting delay. In this paper, the adjudication probabilities are however defined conditional on the full transaction time filtration $\mathcal{F}^{\mathcal{Z}}$, so to compute these, the adjudication hazards must also depend on all of $\mathcal{F}^{\mathcal{Z}}$. One can of course still impose structural assumptions on the adjudication hazards such that they only depend on parts of $\mathcal{F}^{\mathcal{Z}}$. $\triangledown$

Remark 4.4.2. (Estimation and independence assumptions.)
The independence imposed via Assumption 4.3.1 and 4.3.2 is not used in the estimation procedure of Buchardt et al. (2025), and one hence still obtains consistent and asymptotically normal estimators if these assumptions do not hold. We conjecture that one could derive estimators that are more efficient than the ones suggested here by exploiting these independence assumptions. $\triangledown$

4.5 Data application

To illustrate our methods, we calculate transaction time reserves at time η for a subset of the LEC-DK19 (Loss of Earning Capacity – Denmark 2019) data set which was introduced in Buchardt et al. (2025). The data includes information on disability exposure and occurrences, reactivation exposure and occurrences, reporting delays for disabilities, and adjudication exposure and occurrences related to both disabilities and reactivations. The data window $[0, \eta]$ is $[31/01/2015, 01/09/2019]$. Available covariates are gender and age.

We note that the biometric data conforms with the valid time state space from Figure 4.1 with a single disabled state i_1 which we henceforth refer to as i for notational convenience. Furthermore, the adjudication data conforms with the adjudication multistate model described in Appendix 4.C. In the following we estimate the relevant model elements using the method proposed in Section 4.4 and subsequently calculate the reserves using the results from Section 4.3. The implementation is written in `R` (R Development Core Team, 2023) and is available on GitHub (`https://github.com/oliversandqvist/Web-appendix-disability-reserving`).

For estimation of the hazards, we let all the covariates enter in a linear predictor with log link and assume that all hazards are variationally independent such that there is no overlap in parameters between different hazards. The disability hazard is regressed on age, gender, and calendar time, while the reactivation hazard is regressed on the same covariates but also the duration as disabled. As noted in Section 4.C.5, the data does not permit a reasonable estimate of the death hazards, and we thus simply employ the death hazards from Table 4.1 when calculating reserves.

The adjudication hazards for an RBNSi claim are regressed on age, gender, duration since the disability event, duration since the disability event was reported, and whether or not the claim has been (temporarily) rejected previously. The adjudication hazards for an RBNSr claim are regressed on age, gender, duration since the disability event, and duration since the reactivation event. As noted in Remark 4.4.1, these hazards should depend on all of $\mathcal{F}^{\mathcal{Z}}$, so the fact that we for example do not regress the reactivation adjudication hazards on the disability reporting delay should be understood as the implicit assumption that the true value of that regressor is known to be zero.

For the disability reporting delay, we impose a Weibull proportional reverse time hazard distribution which has distribution function $t \mapsto (1 - \exp(-(\lambda t)^k))^{\exp(W^T \beta)}$ for covariates W and parameters (λ, k, β). For covariates, we use age at disability onset and gender. The data does not contain observations of reactivation reporting delays, but we take this to be an artifact of the data rather than a violation of Assumption 4.3.2 and hence proceed as if this had not been the case, recalling that the reactivation reporting delay distribution is only needed for estimation and not for reserving.

With these specifications, we note that the estimation procedure becomes identical to the one in Section 6 of Buchardt et al. (2025) and we may hence use their estimates; confer with Section 6 and Section G of the supplementary material in Buchardt et al. (2025) for the parameter values. We however keep the calendar time effect fixed at its value at η so as to not overextrapolate the observed calendar time trend when calculating the prospective reserves.

	Age	Male	Female	$\min\{\text{Duration}, 5\}$
μ_{ad}, μ_{rd}	0.09	-9.50	-9.80	-
μ_{id}	0.09	-6.40	-6.80	-0.25

Table 4.1: *Death hazards based on estimates published by the Danish FSA.*

For reserving, we sample 100 random insured at time η and compute reserves for each of the categories CBNR, RBNSi, and RBNSr. The data contains around 250,000 insured that are in the portfolio at time η, but it would take a long time to compute the reserve for everyone using our proof-of-concept `R` implementation. The terms that take the longest to compute are valid time active reserves and IBNR reserves which for a single insured take a couple of seconds to evaluate on a regular laptop. Insurance companies have access to optimized calculation kernels and greater computing power with some presently relying on semi-Markov models for their reserves, so this is not a limitation of our proposed approach but rather of this specific implementation. Furthermore, it is possible to speed up computations by finding suitable approximations of, for example, the IBNR reserve given in Theorem 4.3.4. For CBNR and RBNSr, we sample the 100 insured without replacement, but since there are only 59 insured in the RBNSi category at time η, we sample these with replacement.

The transaction time reserves are compared with a naive approach where X_η^η is plugged into the valid time reserves. The naive approach thus leads to reserves that are sufficient to cover disabilities that have an ongoing payout at time η and disabilities that occur after time η but ignores IBNR-claims, claims that are under adjudication, and possible reapplications.

The transaction time reserves are further compared with a simple approach where ad-hoc adjustments of the valid time reserves are made to adjust for IBNR and incomplete adjudication. The simple approach takes the CBNR reserve to be a valid time active reserve where the coverage period is extended with the average disability reporting delay. The average delay d is found to be around 0.53 years when the average is based on disabilities that occurred at least two years before time η to limit the effect of right-truncation. The heuristic is that the insurance company should cover disabilities that arrive up to d years after the end of the coverage period since these disabilities occurred within the coverage period if the reporting delay was deterministically equal to d. The RBNSi reserve is chosen to be a valid time disability reserve with duration 0 since most claims get awarded. For the RBNSr reserve, a valid time disability reserve is used if there are ongoing disability payments and a valid time active reserve is used otherwise to accommodate those who will apply for additional benefits in the future.

The same hazards are used in all of the approaches so the differences between the results only reflect the reserving methodologies and not the estimation. Since the data does not contain information about benefit type or size we set $B(\mathrm{d}t) = 1_{(Y_{t-}=i)}1_{(t-U_{t-}\leq\eta+3)}1_{(a+t\leq 67)}\,\mathrm{d}t$, where a is the age at time 0, corresponding to a unit disability annuity until retirement at age 67 with a coverage period of 3 years. Note that this specification of the cash flow also implies that the insured are covered for disabilities occurring before time η. We finally assume a constant force of interest $r \equiv 0.02$. The reserves are calculated by plugging in the estimated model elements into Theorem 4.3.4, 4.3.8, and 4.3.9. The state-wise valid time reserves entering into these expressions are calculated by solving Thiele's differential equation iteratively over the states by exploiting the hierarchical structure of $\mathcal{J}$.

Remark 4.5.1. (Alternative simple CBNR reserves.)
An alternative but similar simple approach for the CBNR category would be to reserve $V_a(\eta,\eta) + V_i(\eta - d, 0) \times \mu_{ai}(\eta - d, \eta - d) \times d$. Heuristically, with a constant reporting delay d, the probability of having an unreported disability is the probability of having a disability occur in $(\eta - d, \eta]$ which is roughly $\mu_{ai}(\eta - d, \eta - d) \times d$, and for each of these one reserves $V_i(\eta - d, 0)$. This can also be motivated by applying relevant approximations to the expression in Theorem 4.3.4. The resulting reserve is 9.22, which brings the difference between the proposed and simple method down by a factor of 2/3. Another alternative simple reserve could be $\exp(r \times d) \times V_a(\eta - d, \eta - d)$ which results in a reserve of 9.21.

The remaining difference can generally be attributed to the timing of the disability and the covariate dependence being incorporated slightly more imprecisely in the simple model. Since most disability claims are short, the backpay may constitute a considerable part of the payment, making it important to correctly assess the timing of the disability. Note that the performance of the simple methods would likely deteriorate in more complicated situations with inflow/outflow of insured, non-constant interest, and calendar time effects, confer with Remark 4.3.7. $\triangledown$

The results are given in Table 4.2. The settled category is not depicted since the reserve is zero in all cases. The largest relative difference arises for the RBNSi reserve. The naive method under-reserves since it ignores the reported disability and reserves as if the insured was active because no disability benefits have been awarded yet. Surprisingly, the proposed method also leads to larger reserves than the simple method which always reserves the valid time disability reserve. This happens because the adjudication probabilities are close to one (they have an average of around 0.9) and there are many older insured where the effect of conditioning on not having died and reserving from time G_t instead of t leads to moderately larger reserves.

The second largest relative difference is seen for the CBNR reserve. Here, the probability weighting of the active reserve is observed to be very close to one, so the main difference between the proposed and naive method is the IBNR contribution. Thus, the naive method under-reserves as it neglects this term. The simple method much closer to the proposed method as expected, but there is still a sizeable relative difference.

The smallest relative difference for the non-settled cases is seen for the RBNSr reserve. This is because, in the data set, there are considerably more insured receiving running benefits than are reactivated. In fact, only four reactivated insured were sampled and none of them had applied for additional benefits. As the population mix shifts toward a higher proportion of reactivated subjects, the difference between the proposed and naive method is likely to grow.

	(CBNR_η)	(RBNSi_η)	(RBNSr_η)
Proposed method	9.29	502.44	961.68
Simple method	8.98	481.20	945.66
Naive method	7.73	6.23	945.35
Simple method difference	0.32 (3.53%)	21.23 (4.41%)	16.02 (1.69%)
Naive method difference	1.56 (20.19%)	496.21 (7961.48%)	16.33 (1.73%)

Table 4.2: *Reserves for 100 randomly sampled insured from each of the claims settlement categories except the settled category where the reserve is identically zero.*

To explore the practical implications of our results for a full insurance portfolio, we approximate the average reserve in each of the categories for the wider insurance portfolio by the average reserve of the 100 sampled insured. Multiplying the average reserve in a given category with the total number of insured in that category, we obtain the results depicted in Figure 4.4. Comparing with Table 4.2, we see that despite the RBNSi category having the largest absolute and relative differences for the 100 sampled insured, the difference on the portfolio level is comparable with that of the RBNSr category, which had the smallest relative difference for the 100 sampled insured. This is because there are considerably more insured in the RBNSr category at time η. Similarly, the CBNR category, which showed the smallest absolute difference in Table 4.2, leads to the largest difference on the portfolio level due to this category being by far the largest. In total, the naive and simple method leads to portfolio reserves that are around 11.1% and 2.7% smaller than the proposed method, respectively.

To gain further insight into the financial implications, consider that in 2019, the annual reports of two large Danish insurers showed portfolio reserves for health and disability insurance obligations amounting to 9,351 and 17,606 million DKK respectively. Therefore, a 2.7% increase in these reserves would equate to approximately 250 million DKK for the former and 470 million DKK for the latter.

It would be highly relevant to compare the different reserves with observed claim developments to see which best describe the data. This is however not possible with the current data since the disability and reactivation occurrences and exposures are available for a single valuation date only, and their values at different valuation dates also cannot be inferred from the adjudication data since the same id does not refer to the same insured across the individual data tables. This is therefore left as a topic for future research.

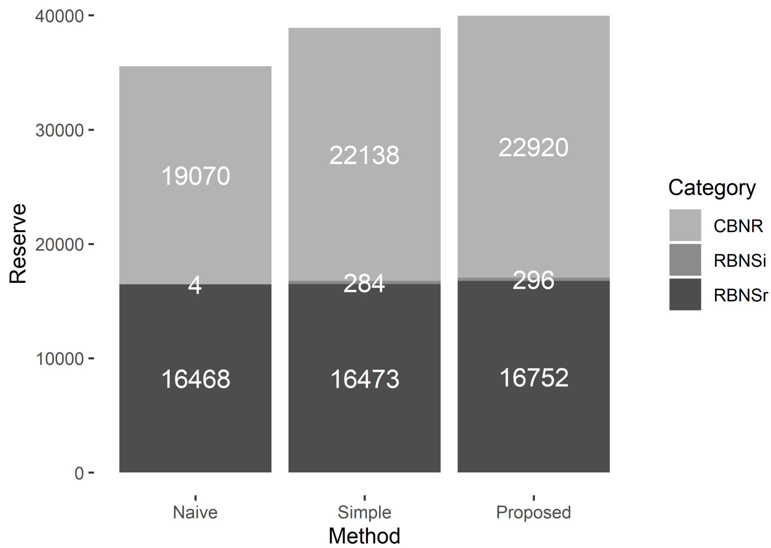

Figure 4.4: *The approximated portfolio level reserve decomposed by category.*

4.6 Conclusion

This paper develops an individual reserving model for disability insurance in the presence of information delays caused by reporting delays and adjudication processes. We have introduced suitable conditional independence assumptions that lead to tractable and interpretable reserves, which may be calculated using the usual valid time hazards, the IBNR-factor, and the adjudication probabilities. The reserves are tailored to the features of disability insurance schemes by accommodating reporting delays and adjudication processing while preserving the advantages offered by valid time multistate models, namely that contractual payments are an a priori known function of a multistate process whose intertemporal distribution is well-understood. It is argued that the estimation procedure from Buchardt et al. (2025) may be used to estimate the model constituents of the reserves. Finally, the practical potential of our models is illustrated through an application to a real insurance data set.

Acknowledgments and declarations of interest

This research has partly been funded by the Innovation Fund Denmark (IFD) under File No. 1044-00144B. The author declares no conflicts of interest. I would like to thank Kristian Buchardt and my supervisor Christian Furrer for many fruitful discussions. I would also like to thank them both for their helpful comments on an earlier version of the manuscript. I am also very grateful to an anonymous referee for comments and suggestions that improved the paper significantly.

4.A Proof of Lemma 4.3.3

To state and prove Proposition 4.A.1, which immediately implies Lemma 4.3.3, we first introduce some marked point process notation. Let $(\mathcal{M}, \mathcal{H}, (\mathcal{H}_t)_{t\geq 0})$ be the canonical space of counting measures, see Section 2 and 4.2 in Jacobsen (2006). Let $\mu \in \mathcal{M}$ be the underlying random counting measure for X. Write (T_n, X_n) for the jump times and marks of μ. For a given random counting measure m, we let $\tau_n(m)$ and $\eta_n(m)$ be the n'th jump time and jump mark respectively implying $T_n = \tau_n(\mu)$ and $X_n = \eta_n(\mu)$. For any measurable random time R, define the truncated measure $\theta_R\mu : (R < \infty) \to \mathcal{M}$ by

$$\theta_R\mu(\omega) = \sum_{n: R(\omega) < T_n(\omega) < \infty} \varepsilon_{(T_n(\omega), X_n(\omega))}$$

where $\varepsilon_{(t,x)}$ is the Dirac measure in (t, x). We write $T_{R,n} = \tau_n(\theta_R\mu)$ and $X_{R,n} = \eta_n(\theta_R\mu)$ for the jump times and jump marks determining $\theta_R\mu$. Thus, if $R(\omega) < \infty$,

$$\theta_R\mu(\omega) = \sum_{n: T_{R,n}(\omega) < \infty} \varepsilon_{(T_{R,n}(\omega), X_{R,n}(\omega))}$$

with $R(\omega) < T_{R,1}(\omega) \leq T_{R,2}(\omega) \leq \dots$. Also let $T_{R,0} = R$ and $X_{R,0} = X_R$. We denote by Q^x the distribution of $\theta_{x_1}\mu$ given $X_{x_1} = x$ constructed as the time-inhomogeneous case in Jacobsen (2006) p. 157-158. Note that this actually corresponds to $Q^{x_1,x}$ in the notation of Jacobsen (2006). Even though X is time-homogeneous, this is usually not the case when we condition on $1_{(\tau_{\{d\}} \leq t)}$ which is why we employ the time-inhomogeneous construction. The function governing the behavior of X between jumps is denoted ϕ_{vs}, meaning that if $v \leq s$ and no jumps occurred in $(v, s]$, one has $X_s = \phi_{vs}(X_v)$. The jump time and jump mark Markov-kernels for Q^x are denoted $F_{t_n,y_n}(v)$ and $r_t(\phi_{t_n t}(y_n), C)$ respectively. The interpretation is that $F_{t_n,y_n}(v)$ gives the probability that the next jump has occurred by time v given that the previous jump happened at time t_n with mark y_n, while $r_t(y, C)$ gives the probability that an event occurring at time t from state y ends up in the set C. Note $F_{\infty,\nabla}(v) = 0$ for any $v \in [0, \infty)$ and $r_\infty(y, C) = 1_{(\nabla \in C)}$.

We let t be fixed but arbitrary. Let $\tilde{Q}^x$ be constructed as Q^x but according to the modified Markov kernels where one additionally conditions on $1_{(\tau_{\{d\}} \leq t)}$, hence now being stochastic. One could, of course, have removed this additional stochasticity by replacing the indicator with the event $(\tau_{\{d\}} > t)$ in the conditioning on the relevant event (RBNS_t). We however stick with this construction since it more easily generalizes to other cases where one wants to keep additional information that is not deterministically known on the relevant event. It holds that $\tilde{Q}^x(H) = \mathbb{P}(\theta_{x_1}\mu \in H \mid X_{x_1} = x, 1_{(\tau_{\{d\}} \leq t)})$ almost surely for any $H \in \mathcal{H}$. Note that for every possible outcome of the conditioning information, the transition kernels of $\tilde{Q}^x$ stay on the Markov form: either the conditioning is superfluous because it relates to an event that occurred before the previous jump time, or it is future-measurable and one can hence use the formula for conditional probabilities in the jump time and jump mark kernels, use the Markov property, and then use the formula for conditional probabilities in reverse. This shows that the transition kernels indeed only depend on (t_n, y_n) and not on $t_1, ..., t_{n-1}$ and $y_1, ..., y_{n-1}$. They will be denoted $\tilde{r}_t(y, C)$ and $\tilde{F}_{t_n,y_n}(v)$. For notational convenience, we write $\mathbb{1}(t) = 1_{(\mathrm{RBNS}_t)}$. We now state and prove Proposition 4.A.1, which immediately implies the statement in Lemma 4.3.3.

Proposition 4.A.1. *(Strong Markov type property at G_t.)*
Under Assumption 4.3.2, one has

$$\mathbb{1}(t)\mathbb{P}(\theta_{G_t}\mu \in \cdot \mid (X_s)_{s\leq G_t}, 1_{(\tau_{\{d\}}\leq t)}) = \mathbb{1}(t)\tilde{Q}^{X_{G_t}}(\cdot) \tag{4.A.1}$$

$\mathbb{P}$-*a.s.*

This is an almost sure equality between probability measures on $(\mathcal{M}, \mathcal{H})$. We note that the proof does not use many properties of our model for X and G_t, and similar arguments may thus be used to show strong Markov properties for other Markov processes and random times provided that an independence assumption similar to Assumption 4.3.2 is imposed.

Proof. The proof is inspired by the proof of Theorem 7.5.1 in Jacobsen (2006) with the necessary changes to adjust for the fact that Assumption 4.3.2 is different than the usual Markov independence assumption.

As in Jacobsen (2006), we note that showing Equation (4.A.1) is equivalent to showing for $n \geq 1$ and all measurable and bounded function $f_i : (0, \infty] \times \overline{E} \to \mathbb{R}$ that the following holds:

$$\mathbb{1}(t)\mathbb{E}\left[\prod_{i=1}^{n} f_i(T_{G_t,i}, X_{G_t,i}) \mid (X_s)_{s\leq G_t}, 1_{(\tau_{\{d\}}\leq t)}\right] = \mathbb{1}(t)\tilde{E}^{X_{G_t}}\left[\prod_{i=1}^{n} f_i(\tau_i, \eta_i)\right] \tag{4.A.2}$$

which we will prove by induction on n. The proof consists of four steps:

(i) For a discretization of G_t to $G_t(M)$ with $M \in \mathbb{N}$, show the result (4.A.1) for X stopped at $T_{G_t(M),n-1}$ for any $n \geq 1$:

$$\begin{aligned}&1_{(T_{G_t(M),n-1}<\infty)}\mathbb{1}(t)\mathbb{P}(\theta_{T_{G_t(M),n-1}}\mu \in \cdot \mid (X_s)_{s\leq T_{G_t(M),n-1}}, X_{G_t}, 1_{(\tau_{\{d\}}\leq t)})\\ &= 1_{(T_{G_t(M),n-1}<\infty)}\mathbb{1}(t)\tilde{Q}^{X_{T_{G_t(M),n-1}}}(\cdot).\end{aligned}$$

(ii) Show convergence for $M \to \infty$:

$$\begin{aligned}&\lim_{M\to\infty} 1_{(T_{G_t(M),n-1}<\infty)}\mathbb{1}(t)\tilde{E}^{X_{T_{G_t(M),n-1}}}[f_n(\tau_1, \eta_1)]\\ &= 1_{(T_{G_t,n-1}<\infty)}\mathbb{1}(t)\tilde{E}^{X_{T_{G_t,n-1}}}[f_n(\tau_1, \eta_1)].\end{aligned}$$

(iii) Discretize G_t to $G_t(M)$ and use dominated convergence with (i) and (ii) to conclude:

$$\begin{aligned}&\mathbb{1}(t)\mathbb{E}[f_n(T_{G_t,n}, X_{G_t,n}) \mid (X_s)_{s\le T_{G_t,n-1}}, X_{G_t}, 1_{(\tau_{\{d\}}\le t)}]\\&= \mathbb{1}(t)\tilde{E}^{X_{T_{G_t,n-1}}}[f_n(\tau_1,\eta_1)].\end{aligned}$$

(iv) Use (iii) and induction over n to finish the proof.

(i) : Let $H \in \mathcal{H}$ be given. Define

$$G_t(M) = \sum_{m=1}^{\infty} t_{Mm} 1_{(t_{M(m-1)}\le G_t < t_{Mm})}$$

for $t_{Mm} = m2^{-M}$. Had we not known that $G_t < \infty$ almost surely, one would have to add $\infty 1_{(G_t=\infty)}$ here and show the result on $(G_t < \infty)$.

Take an F on the form

$$F = ((X_s)_{s\le T_{G_t(M),n-1}} \in B) \cap (X_{G_t} \in C) \cap (1_{(\tau_{\{d\}}\le t)} \in D)$$

and note that the collection of such sets constitutes an intersection-stable generator for

$$\sigma((X_s)_{s\le T_{G_t(M),n-1}}) \vee \sigma(X_{G_t}) \vee \sigma(1_{(\tau_{\{d\}}\le t)})$$

containing Ω. Write $F_{M,m} = F \cap (G_t(M) = t_{Mm})$ and note that

$$F_{M,m} \in \sigma((X_s)_{s\le T_{t_{Mm},n-1}}) \vee \sigma(X_{G_t}) \vee \sigma(1_{(\tau_{\{d\}}\le t)}).$$

Here we used that $G_t(M)$ is G_t-measurable and that X_{G_t} contains G_t as a coordinate. Now we get

$$\begin{aligned}&\int_{F_{M,m}} 1_{(T_{G_t(M),n-1}<\infty)}\mathbb{1}(t)\tilde{Q}^{X_{T_{G_t(M),n-1}}}(H)\, d\mathbb{P}\\
&= \int_{F_{M,m}} 1_{(T_{t_{Mm},n-1}<\infty)}\mathbb{1}(t)\tilde{Q}^{X_{T_{t_{Mm},n-1}}}(H)\, d\mathbb{P}\\
&= \int_{F_{M,m}} 1_{(T_{t_{Mm},n-1}<\infty)}\mathbb{1}(t)\mathbb{P}(\theta_{T_{t_{Mm},n-1}}\mu \in H \mid (X_s)_{s\le T_{t_{Mm},n-1}}, X_{G_t}, 1_{(\tau_{\{d\}}\le t)})\, d\mathbb{P}\\
&= \int_{F_{M,m}} 1_{(T_{t_{Mm},n-1}<\infty)}\mathbb{1}(t)1_{(\theta_{T_{t_{Mm},n-1}}\mu\in H)}\, d\mathbb{P}\\
&= \mathbb{P}\left(F_{M,m} \cap (T_{t_{Mm},n-1} < \infty) \cap (\mathrm{RBNS}_t) \cap (\theta_{T_{t_{Mm},n-1}}\mu \in H)\right)\\
&= \mathbb{P}\left(F_{M,m} \cap (T_{G_t(M),n-1} < \infty) \cap (\mathrm{RBNS}_t) \cap (\theta_{T_{G_t(M),n-1}}\mu \in H)\right),\end{aligned}$$

where the second equality is Lemma 4.A.2 with $T = T_{t_{Mm},n-1}$. The third equality uses $F_{M,m}\cap(T_{t_{Mm},n-1} < \infty)\cap(\mathrm{RBNS}_t) \in \sigma((X_s)_{s\le T_{t_{Mm},n-1}})\vee\sigma(X_{G_t})\vee\sigma(1_{(\tau_{\{d\}}\le t)})$;

recall in particular that $(\mathrm{RBNS}_t) = (T_{\{2\}} \leq t, T_{\{5\}} > t)$ and $(T_{\{2\}} \leq t)$ is known from $\sigma(X_{G_t})$ since $(T_{\{2\}} \leq t) = (G_t > 0)$. Summing over $m \geq 1$ and using $(G_t < \infty)$ almost surely gives

$$\begin{aligned}&\int_F 1_{(T_{G_t(M),n-1}<\infty)} \mathbb{1}(t) \tilde{Q}^{X_{T_{G_t(M),n-1}}}(H)\, d\mathbb{P} \\ &= \mathbb{P}\left(F \cap (T_{G_t(M),n-1} < \infty) \cap (\mathrm{RBNS}_t) \cap (\theta_{T_{G_t(M),n-1}} \mu \in H)\right).\end{aligned}$$

As the left- and right-hand side are finite measures, uniqueness of finite measures on intersection stable classes containing Ω gives that the equation holds for any $F \in \sigma((X_s)_{s \leq T_{G_t(M),n-1}}) \vee \sigma(X_{G_t}) \vee \sigma(1_{(\tau_{\{d\}} \leq t)})$. As it also holds that $1_{(T_{G_t(M),n-1}<\infty)} \mathbb{1}(t) \tilde{Q}^{X_{T_{G_t(M),n-1}}}(H)$ is $\sigma((X_s)_{s \leq T_{G_t(M),n-1}}) \vee \sigma(X_{G_t}) \vee \sigma(1_{(\tau_{\{d\}} \leq t)})$-measurable, it satisfies the defining properties of

$$\mathbb{E}[1_{(T_{G_t(M),n-1}<\infty)} 1_{(\mathrm{RBNS}_t)} 1_{(\theta_{T_{G_t(M),n-1}} \mu \in H)} \mid (X_s)_{s \leq T_{G_t(M),n-1}}, X_{G_t}, 1_{(\tau_{\{d\}} \leq t)}],$$

so we conclude

$$\begin{aligned}&1_{(T_{G_t(M),n-1}<\infty)} \mathbb{1}(t) \mathbb{P}(\theta_{T_{G_t(M),n-1}} \mu \in H \mid (X_s)_{s \leq T_{G_t(M),n-1}}, X_{G_t}, 1_{(\tau_{\{d\}} \leq t)}) \\ &= \mathbb{E}[1_{(T_{G_t(M),n-1}<\infty)} 1_{(\mathrm{RBNS}_t)} 1_{(\theta_{T_{G_t(M),n-1}} \mu \in H)} \mid (X_s)_{s \leq T_{G_t(M),n-1}}, X_{G_t}, 1_{(\tau_{\{d\}} \leq t)}] \\ &= 1_{(T_{(G_t)(M),n-1}<\infty)} \mathbb{1}(t) \tilde{Q}^{X_{T_{(G_t)(M),n-1}}}(H).\end{aligned}$$

(ii) : Note that we can write

$$\begin{aligned}&1_{(T_{G_t,n-1}<\infty)} \tilde{E}^{X_{T_{G_t,n-1}}}[f_n(\tau_1, \eta_1)] \\ &= 1_{(T_{G_t,n-1}<\infty)} \int_{(T_{G_t,n-1},\infty]} \int_E f_n(v,y)\, \tilde{r}_v(\phi_{T_{G_t,n-1}v}(X_{T_{G_t,n-1}}), dy) \tilde{F}_{X_{T_{G_t},n-1}}(dv)\end{aligned}$$

with a similar expression for $1_{(T_{G_t(M),n-1}<\infty)} \tilde{E}^{X_{T_{G_t(M),n-1}}}[f_n(\tau_1, \eta_1)]$.

On $(\overline{N}_{T_{G_t,n-1}} = \overline{N}_{T_{G_t(M),n-1}})$ and $(T_{G_t(M),n-1} < \infty)$, it holds for $v \geq T_{G_t(M),n-1}$, since $T_{G_t(M),n-1} \geq T_{G_t,n-1}$, that

$$\begin{aligned}\phi_{T_{G_t,n-1}v}(X_{T_{G_t,n-1}}) &= \phi_{T_{G_t(M),n-1}v}(\phi_{T_{G_t,n-1}T_{G_t(M),n-1}}(X_{T_{G_t,n-1}})) \\ &= \phi_{T_{G_t(M),n-1}v}(X_{T_{G_t(M),n-1}})\end{aligned}$$

and

$$\overline{\tilde{F}}_{X_{T_{G_t,n-1}}}(v) = \overline{\tilde{F}}_{X_{T_{G_t,n-1}}}(T_{G_t(M),n-1}) \overline{\tilde{F}}_{X_{T_{G_t(M),n-1}}}(v).$$

Therefore, it holds on $(\overline{N}_{T_{G_t,n-1}} = \overline{N}_{T_{G_t(M),n-1}})$ that

$$\begin{aligned}
&1_{(T_{G_t(M),n-1}<\infty)}\tilde{E}^{X_{T_{G_t(M),n-1}}}[f_n(\tau_1,\eta_1)] \\
&= 1_{(T_{G_t(M),n-1}<\infty)}\times \\
&\quad \int_{(T_{(G_t(M),n-1},\infty]}\int_{\overline{E}} f_n(v,y)\,\tilde{r}_v(\phi_{T_{G_t(M),n-1}v}(X_{T_{G_t(M),n-1}}),dy)\tilde{F}_{T_{G_t(M),n-1}}(dv) \\
&= 1_{(T_{G_t(M),n-1}<\infty)} \times \frac{1}{\overline{\tilde{F}}_{X_{T_{G_t,n-1}}}(T_{G_t(M),n-1})}\times \\
&\quad \int_{(T_{G_t(M),n-1},\infty]}\int_{\overline{E}} f_n(v,y)\,\tilde{r}_v(\phi_{T_{G_t,n-1}v}(X_{T_{G_t,n-1}}),dy)\tilde{F}_{T_{G_t,n-1}}(dv) \\
&\overset{M\to\infty}{\to} 1_{(T_{G_t,n-1}<\infty)}\tilde{E}^{X_{T_{G_t,n-1}}}[f_n(\tau_1,\eta_1)].
\end{aligned}$$

Note that

$$1_{\left(\overline{N}_{T_{G_t,n-1}}=\overline{N}_{T_{G_t(M),n-1}}\right)} \overset{M\to\infty}{\to} 1,$$

since $\overline{N}$ is right-continuous and $T_{G_t(M),n-1} \downarrow T_{G_t,n-1}$. We thus get

$$\begin{aligned}
&\lim_{M\to\infty} 1_{(T_{G_t(M),n-1}<\infty)}\tilde{E}^{X_{T_{G_t(M),n-1}}}[f_n(\tau_1,\eta_1)] \\
&= \lim_{M\to\infty} 1_{\left(\overline{N}_{T_{G_t,n-1}}=\overline{N}_{T_{G_t(M),n-1}}\right)}1_{(T_{G_t(M),n-1}<\infty)}\tilde{E}^{X_{T_{G_t(M),n-1}}}[f_n(\tau_1,\eta_1)] \\
&= \lim_{M\to\infty} 1_{\left(\overline{N}_{T_{G_t,n-1}}=\overline{N}_{T_{G_t(M),n-1}}\right)}1_{(T_{G_t,n-1}<\infty)}\tilde{E}^{X_{T_{G_t,n-1}}}[f_n(\tau_1,\eta_1)] \\
&= 1_{(T_{G_t,n-1}<\infty)}\tilde{E}^{X_{T_{G_t,n-1}}}[f_n(\tau_1,\eta_1)].
\end{aligned}$$

We are now in a position to show the result using dominated convergence.

(iii) : Take F on the form

$$F = ((X_s)_{s\le T_{G_t,n-1}} \in B) \cap (X_{G_t} \in C) \cap (1_{(\tau_{\{d\}}\le t)} \in D).$$

Such sets constitute an intersection stable generator of $\sigma((X_s)_{s\le T_{G_t,n-1}}) \vee \sigma(X_{G_t}) \vee \sigma(1_{(\tau_{\{d\}}\le t)})$ containing Ω. Now let

$$F_M = ((X_s)_{s\le T_{G_t(M),n-1}} \in B) \cap (X_{G_t} \in C) \cap (1_{(\tau_{\{d\}}\le t)} \in D)$$

as then

$$F_M \in \sigma((X_s)_{s\le T_{G_t(M),n-1}}) \vee \sigma(X_{G_t}) \vee \sigma(1_{(\tau_{\{d\}}\le t)})$$

and $1_{F_M} \overset{M\to\infty}{\to} 1_F$ by using that X is right-continuous, that $T_{G_t(M),n-1} \ge T_{G_t,n-1}$, and that $\lim_{M\to\infty} T_{G_t(M),n-1} = T_{G_t,n-1}$. We now see using dominated convergence and the results of (i) and (ii):

$$
\begin{aligned}
&\int_F 1_{(T_{G_t,n-1}<\infty)}\mathbb{1}(t)\tilde{E}^{X_{T_{G_t,n-1}}}[f_n(\tau_1,\eta_1)]d\mathbb{P} \\
&\overset{\text{(ii)}}{=} \lim_{M\to\infty}\int_{F_M} 1_{(T_{G_t(M),n-1}<\infty)}\mathbb{1}(t)\tilde{E}^{X_{T_{G_t(M),n-1}}}[f_n(\tau_1,\eta_1)]d\mathbb{P} \\
&\overset{\text{(i)}}{=} \lim_{M\to\infty}\int_{F_M} 1_{(T_{G_t(M),n-1}<\infty)}\mathbb{1}(t) \\
&\qquad\qquad \times \mathbb{E}[f_n(T_{G_t(M),n},X_{G_t(M),n}) \mid (X_s)_{s\le T_{G_t(M),n-1}},X_{G_t},1_{(\tau_{\{d\}}\le t)}]d\mathbb{P} \\
&= \lim_{M\to\infty}\int_{F_M} 1_{(T_{G_t(M),n-1}<\infty)}\mathbb{1}(t)f_n(T_{G_t(M),n},X_{G_t(M),n})d\mathbb{P} \\
&= \int_F 1_{(T_{G_t,n-1}<\infty)}\mathbb{1}(t)f_n(T_{G_t,n},X_{G_t,n})d\mathbb{P}.
\end{aligned}
$$

In the last equality, we used that $(T_{G_t(M),n},X_{G_t(M),n})(\omega) = (T_{G_t,n},X_{G_t,n})(\omega)$ for $M = M(\omega)$ sufficiently large. Now since the left-hand side and right-hand side are finite measures, when seen as a function of F, that are equal on an intersection stable generator including Ω, we can conclude that they are equal on all of $\sigma((X_s)_{s\le T_{G_t,n-1}}) \vee \sigma(X_{G_t}) \vee \sigma(1_{(\tau_{\{d\}}\le t)})$. It is also easily seen that $1_{(T_{G_t,n-1}<\infty)}\mathbb{1}(t)\tilde{E}^{X_{T_{G_t,n-1}}}[f_n(\tau_1,\eta_1)]$ is $\sigma((X_s)_{s\le T_{G_t,n-1}}) \vee \sigma(X_{G_t}) \vee \sigma(1_{(\tau_{\{d\}}\le t)})$-measurable, so we may conclude

$$
\begin{aligned}
&1_{(T_{G_t,n-1}<\infty)}\mathbb{1}(t)\mathbb{E}[f_n(T_{G_t,n},X_{G_t,n}) \mid (X_s)_{s\le T_{G_t,n-1}},X_{G_t},1_{(\tau_{\{d\}}\le t)}] \\
&= \mathbb{E}[1_{(T_{G_t,n-1}<\infty)}\mathbb{1}(t)f_n(T_{G_t,n},X_{G_t,n}) \mid (X_s)_{s\le T_{G_t,n-1}},X_{G_t},1_{(\tau_{\{d\}}\le t)}] \\
&= 1_{(T_{G_t,n-1}<\infty)}\mathbb{1}(t)\tilde{E}^{X_{T_{G_t,n-1}}}[f_n(\tau_1,\eta_1)].
\end{aligned}
$$

We can further strengthen this to

$$
\mathbb{1}(t)\mathbb{E}[f_n(T_{G_t,n},X_{G_t,n}) \mid (X_s)_{s\le T_{G_t,n-1}},X_{G_t},1_{(\tau_{\{d\}}\le t)}] = \mathbb{1}(t)\tilde{E}^{X_{T_{G_t,n-1}}}[f_n(\tau_1,\eta_1)] \tag{4.A.3}
$$

since also

$$
\begin{aligned}
&1_{(T_{G_t,n-1}=\infty)}\mathbb{E}[f_n(T_{G_t,n},X_{G_t,n}) \mid (X_s)_{s\le T_{G_t,n-1}},X_{G_t},1_{(\tau_{\{d\}}\le t)}] \\
&= 1_{(T_{G_t,n-1}=\infty)}f_n(\infty,\nabla) \\
&= 1_{(T_{G_t,n-1}=\infty)}\tilde{E}^{X_{T_{G_t,n-1}}}[f_n(\tau_1,\eta_1)].
\end{aligned}
$$

Now we are in a position to use induction over n in Equation (4.A.2).

(iv) : Using Equation (4.A.3) with $n = 1$ gives the result from Equation (4.A.2) for

$n = 1$. For $n \geq 2$, assume the result holds for $n-1$ and observe

$$
\begin{aligned}
&\mathbb{1}(t)\mathbb{E}\left[\prod_{i=1}^{n} f_i(T_{G_t,i}, X_{G_t,i}) \,\middle|\, (X_s)_{s\leq G_t}, 1_{(\tau_{\{d\}}\leq t)}\right] \\
&= \mathbb{1}(t)\mathbb{E}\Bigg[\mathbb{E}[f_n(T_{G_t,n}, X_{G_t,n}) \mid (X_s)_{s\leq T_{G_t,n-1}}, X_{G_t}, 1_{(\tau_{\{d\}}\leq t)}] \\
&\qquad\qquad \times \prod_{i=1}^{n-1} f_i(T_{G_t,i}, X_{G_t,i}) \,\Bigg|\, (X_s)_{s\leq G_t}, 1_{(\tau_{\{d\}}\leq t)}\Bigg] \\
&\overset{\text{(iii)}}{=} \mathbb{1}(t)\mathbb{E}\left[\tilde{E}^{X_{T_{G_t,n-1}}}[f_n(\tau_1,\eta_1)] \times \prod_{i=1}^{n-1} f_i(T_{G_t,i}, X_{G_t,i}) \,\middle|\, (X_s)_{s\leq G_t}, 1_{(\tau_{\{d\}}\leq t)}\right] \\
&= \mathbb{1}(t)\tilde{E}^{X_{G_t}}\left[\tilde{E}^{\eta_{n-1}}[f_n(\tau_1,\eta_1)] \times \prod_{i=1}^{n-1} f_i(\tau_i,\eta_i)\right]
\end{aligned}
$$

where the last equality follows by the induction hypothesis. Continuing the calculations, we see

$$
\mathbb{1}(t)\tilde{E}^{X_{G_t}}\left[\prod_{i=1}^{n-1} f_i(\tau_i,\eta_i) \times \tilde{E}^{\eta_{n-1}}[f_n(\tau_1,\eta_1)]\right] = \mathbb{1}(t)\tilde{E}^{X_{G_t}}\left[\prod_{i=1}^{n} f_i(\tau_i,\eta_i)\right],
$$

by writing out the expectations using the Markov kernels. Hence we have shown

$$
\mathbb{1}(t)\mathbb{E}\left[\prod_{i=1}^{n} f_i(T_{G_t,i}, X_{G_t,i}) \mid (X_s)_{s\leq G_t}, 1_{(\tau_{\{d\}}\leq t)}\right] = \mathbb{1}(t)\tilde{E}^{X_{G_t}}\left[\prod_{i=1}^{n} f_i(\tau_i,\eta_i)\right]
$$

so Equation (4.A.2) hold for all $n \in \mathbb{N}$ by induction, which proves the desired result. $\square$

For the proof of Proposition 4.A.1, we needed Lemma 4.A.2, which corresponds to Proposition 4.A.1 if the random time G_t was replaced by a $\mathcal{F}^X$-stopping time in some of the terms.

Lemma 4.A.2. *(Strong Markov type property at stopping time.) Under Assumption 4.3.2 it holds that*

$$
\begin{aligned}
&1_{(T<\infty)}1_{(G_t\leq T)}\mathbb{1}(t)\mathbb{P}(\theta_T\mu \in \cdot \mid (X_s)_{s\leq T}, X_{G_t}, 1_{(\tau_{\{d\}}\leq t)}) \\
&= 1_{(T<\infty)}1_{(G_t\leq T)}\mathbb{1}(t)\tilde{Q}^{X_T}(\cdot)
\end{aligned} \tag{4.A.4}
$$

for any $\mathcal{F}^X$-stopping time T.

The proof is very similar to the proof of Proposition 4.A.1.

Proof. This is equivalent to showing for all $n \geq 1$ and all measurable bounded functions f_i, $1 \leq i \leq n$ that

$$\begin{aligned}&1_{(T<\infty)}1_{(G_t\leq T)}\mathbb{1}(t)\mathbb{E}\left[\prod_{i=1}^{n} f_i(T_{T,i}, X_{T,i}) \,\Big|\, (X_s)_{s\leq T}, X_{G_t}, 1_{(\tau_{\{d\}}\leq t)}\right]\\ &= 1_{(T<\infty)}1_{(G_t\leq T)}\mathbb{1}(t)\tilde{E}^{X_T}\left[\prod_{i=1}^{n} f_i(\tau_i, \eta_i)\right].\end{aligned}$$

Define

$$T(M) = \sum_{m=1}^{\infty} t_{Mm}1_{(t_{M(m-1)}\leq T<t_{Mm})} + \infty 1_{(T=\infty)}$$

with $t_{Mm} = m2^{-M}$. Note that $(T(M) < \infty) = (T < \infty)$. We partition the proof into four parts corresponding to the four parts of the proof of Proposition 4.A.1.

(i) : First we show the result (4.A.4) with the discrete random time $T(M)$ in place of T, i.e.

$$\begin{aligned}&1_{(T<\infty)}1_{(G_t\leq T(M))}\mathbb{1}(t)\mathbb{P}(\theta_{T(M)}\mu \in \cdot \mid (X_s)_{s\leq T(M)}, X_{G_t}, 1_{(\tau_{\{d\}}\leq t)})\\ &= 1_{(T<\infty)}1_{(G_t\leq T(M))}\mathbb{1}(t)\tilde{Q}^{X_{T(M)}}(\cdot).\end{aligned} \tag{4.A.5}$$

Take $H \in \mathcal{H}$ and F on the form

$$F = ((X_s)_{s\leq T(M)} \in B) \cap (X_{G_t} \in C) \cap (1_{(\tau_{\{d\}}\leq t)} \in D).$$

Note that the collection of such sets constitutes an intersection-stable generator for $\sigma((X_s)_{s\leq T(M)}) \vee \sigma(X_{G_t}) \vee \sigma(1_{(\tau_{\{d\}}\leq t)})$ containing Ω. Write $F_{M,m} = F \cap (T(M) = t_{Mm})$, and note that $F_{M,m} \in \mathcal{F}^X_{t_{Mm}} \vee \sigma(X_{G_t}) \vee \sigma(1_{(\tau_{\{d\}}\leq t)})$ and also $(\mathrm{RBNS}_t) \cap (G_t \leq t_{Mm}) \in \mathcal{F}^X_{t_{Mm}} \vee \sigma(X_{G_t}) \vee \sigma(1_{(\tau_{\{d\}}\leq t)})$. Then

$$\begin{aligned}&\int_{F_{M,m}} 1_{(G_t\leq T(M))}\mathbb{1}(t)\tilde{Q}^{X_{T(M)}}(H)d\mathbb{P}\\ &= \int_{F_{M,m}} 1_{(G_t\leq t_{Mm})}\mathbb{1}(t)\tilde{Q}^{X_{t_{Mm}}}(H)d\mathbb{P}\\ &= \int_{F_{M,m}} 1_{(G_t\leq t_{Mm})}\mathbb{1}(t)\mathbb{P}(\theta_{t_{Mm}}\mu \in H \mid X_{t_{Mm}}, 1_{(\tau_{\{d\}}\leq t)})d\mathbb{P}\\ &= \int_{F_{M,m}} 1_{(G_t\leq t_{Mm})}\mathbb{1}(t)\mathbb{P}(\theta_{t_{Mm}}\mu \in H \mid \mathcal{F}^X_{t_{Mm}} \vee \sigma(X_{G_t}) \vee \sigma(1_{(\tau_{\{d\}}\leq t)}))d\mathbb{P}\\ &= \int_{F_{M,m}} 1_{(G_t\leq t_{Mm})}\mathbb{1}(t)1_{(\theta_{t_{Mm}}\mu\in H)}d\mathbb{P}\\ &= \mathbb{P}(F_{M,m} \cap (G_t \leq T(M)) \cap (\mathrm{RBNS}_t) \cap (\theta_{T(M)}\mu \in H)).\end{aligned}$$

The third equality is Assumption 4.3.2 with $v = t_{Mm}$. Summing over $m \geq 1$ and using uniqueness of finite measures on intersection-stable generators gives the result for $T(M)$ on $(T(M) < \infty) = (T < \infty)$. This now gives Equation (4.A.5).

(ii) : Note that for any bounded measurable function f, we have

$$\tilde{E}^{X_T}[f(\tau_1, \eta_1)] = \int_{(T,\infty]} \int_{\overline{E}} f(v, y)\, \tilde{r}_v(\phi_{Tv}(X_T), \mathrm{d}y)\, \tilde{F}_{X_T}(\mathrm{d}v)$$

with a similar expression for $\tilde{E}^{X_{T(M)}}[f(\tau_1, \eta_1)]$. On $(\overline{N}_T = \overline{N}_{T(M)})$, it holds for $v \geq T(M)$, since also $T(M) \geq T$, that

$$\begin{aligned} \phi_{Tv}(X_T) &= \phi_{T(M)v}(\phi_{TT(M)}(X_T)) \\ &= \phi_{T(M)v}(X_{T(M)}) \end{aligned}$$

and

$$\overline{\tilde{F}}_{X_T}(v) = \overline{\tilde{F}}_{X_T}(T(M))\overline{\tilde{F}}_{X_{T(M)}}(v).$$

Hence, we have on $(\overline{N}_T = \overline{N}_{T(M)}) \cap (T < \infty)$

$$\begin{aligned} \tilde{E}^{X_{T(M)}}[f(\tau_1, \eta_1)] &= \int_{(T(M),\infty]} \int_{\overline{E}} f(v, y)\, \tilde{r}_v(\phi_{T(M)v}(X_{T(M)}), \mathrm{d}y)\, \tilde{F}_{X_{T(M)}}(\mathrm{d}v) \\ &= \frac{1}{\overline{\tilde{F}}_{X_T}(T(M))} \int_{(T(M),\infty]} \int_{\overline{E}} f(v, y)\, \tilde{r}_v(\phi_{Tv}(X_T), \mathrm{d}y)\, \tilde{F}_{X_T}(\mathrm{d}v) \\ &\overset{M\to\infty}{\to} \int_{(T,\infty]} \int_{\overline{E}} f(v, y)\, \tilde{r}_v(\phi_{Tv}(X_T), \mathrm{d}y)\, \tilde{F}_{X_T}(\mathrm{d}v) \\ &= \tilde{E}^{X_T}[f(\tau_1, \eta_1)]. \end{aligned}$$

Since also $1_{(\overline{N}_{T(M)} = \overline{N}_T)} \to 1$ for $M \to \infty$, we have

$$\begin{aligned} \lim_{M\to\infty} 1_{(T<\infty)} \tilde{E}^{X_{T(M)}}[f(\tau_1, \eta_1)] &= \lim_{M\to\infty} 1_{(\overline{N}_{T(M)} = \overline{N}_T)} 1_{(T<\infty)} \tilde{E}^{X_{T(M)}}[f(\tau_1, \eta_1)] \\ &= \lim_{M\to\infty} 1_{(\overline{N}_{T(M)} = \overline{N}_T)} 1_{(T<\infty)} \tilde{E}^{X_T}[f(\tau_1, \eta_1)] \\ &= 1_{(T<\infty)} \tilde{E}^{X_T}[f(\tau_1, \eta_1)]. \end{aligned}$$

We are now in a position to use dominated convergence.

(iii) : Take F on the form $F = ((X_s)_{s\leq T} \in B) \cap (X_{G_t} \in C) \cap (1_{(\tau_{\{d\}} \leq t)} \in D)$ and write

$$F_M = ((X_s)_{s\leq T(M)} \in B) \cap (X_{G_t} \in C) \cap (1_{(\tau_{\{d\}} \leq t)} \in D).$$

Note that $F_M \in \sigma((X_s)_{s\leq T(M)}) \vee \sigma(X_{G_t}) \vee \sigma(1_{(\tau_{\{d\}} \leq t)})$ and $\lim_{M\to\infty} 1_{F_M} = 1_F$ since X is right-continuous and $T(M) \downarrow T$. Similarly, $\lim_{M\to\infty} 1_{(G_t \leq T(M))} = 1_{(G_t \leq T)}$

since the indicator $s \mapsto 1_{(G_t \le s)}$ is right-continuous and $T(M) \downarrow T$. By dominated convergence and the results from (i) and (ii), we obtain:

$$
\begin{aligned}
&\int_F 1_{(T<\infty)} 1_{(G_t \le T)} \mathbb{1}(t) \tilde{E}^{X_T}[f(\tau_1, \eta_1)] d\mathbb{P} \\
&\overset{\text{(ii)}}{=} \lim_{M\to\infty} \int_{F_M} 1_{(T<\infty)} 1_{(G_t \le T(M))} \mathbb{1}(t) \tilde{E}^{X_{T(M)}}[f(\tau_1, \eta_1)] d\mathbb{P} \\
&\overset{\text{(i)}}{=} \lim_{M\to\infty} \int_{F_M} 1_{(T<\infty)} 1_{(G_t \le T(M))} \mathbb{1}(t) \\
&\qquad\qquad \times \mathbb{E}[f(T_{T(M),1}, X_{T(M),1}) \mid (X_s)_{s \le T(M)}, X_{G_t}, 1_{(\tau_{\{d\}} \le t)}] d\mathbb{P} \\
&= \lim_{M\to\infty} \int_{F_M} 1_{(T<\infty)} 1_{(G_t \le T(M))} \mathbb{1}(t) f(T_{T(M),1}, X_{T(M),1}) d\mathbb{P} \\
&= \int_F 1_{(T<\infty)} 1_{(G_t \le T)} \mathbb{1}(t) f(T_{T,1}, X_{T,1}) d\mathbb{P},
\end{aligned}
$$

where the third equality follows by

$$
F_M \cap (T < \infty) \cap (G_t \le T(M)) \cap (\mathrm{RBNS}_t) \in \sigma((X_s)_{s \le T(M)} \vee \sigma(X_{G_t}) \vee \sigma(1_{(\tau_{\{d\}} \le t)})
$$

and the last equality follows since $\lim_{M\to\infty} 1_{F_M} = 1_F$. By uniqueness of finite measures on intersection stable generators, this shows that

$$
\begin{aligned}
&1_{(T<\infty)} 1_{(G_t \le T)} \mathbb{1}(t) \tilde{E}^{X_T}[f(\tau_1, \eta_1)] \\
&= \mathbb{E}[1_{(T<\infty)} 1_{(G_t \le T)} \mathbb{1}(t) f(T_{T,1}, X_{T,1}) \mid (X_s)_{s \le T}, X_{G_t}, 1_{(\tau_{\{d\}} \le t)}] \\
&= 1_{(T<\infty)} 1_{(G_t \le T)} \mathbb{1}(t) \mathbb{E}[f(T_{T,1}, X_{T,1}) \mid (X_s)_{s \le T}, X_{G_t}, 1_{(\tau_{\{d\}} \le t)}].
\end{aligned}
$$

The result also holds when removing the indicator $1_{(T<\infty)}$ since

$$
\begin{aligned}
&1_{(T=\infty)} 1_{(G_t \le T)} \mathbb{1}(t) \tilde{E}^{X_T}[f(\tau_1, \eta_1)] \\
&= 1_{(T=\infty)} 1_{(G_t \le T)} \mathbb{1}(t) f(\infty, \nabla) \\
&= 1_{(T=\infty)} 1_{(G_t \le T)} \mathbb{1}(t) \mathbb{E}[f(T_{T,1}, X_{T,1}) \mid (X_s)_{s \le T}, X_{G_t}, 1_{(\tau_{\{d\}} \le t)}].
\end{aligned}
$$

(iv) : Using the result of (iii) with $f = f_1$ gives the base case of the induction. For

$n \geq 2$, assume that the result hold for $n-1$. Observe then

$$
\begin{aligned}
&1_{(T<\infty)}1_{(G_t\leq T)}\mathbb{1}(t)\mathbb{E}\left[\prod_{i=1}^{n} f_i(T_{T,i}, X_{T,i}) \mid (X_s)_{s\leq T}, X_{G_t}, 1_{(\tau_{\{d\}}\leq t)}\right] \\
&= \mathbb{E}\Bigg[1_{(T<\infty)}1_{(G_t\leq T)}\mathbb{1}(t)\mathbb{E}[f_n(T_{T,n}, X_{T,n}) \mid (X_s)_{s\leq T_{T,n-1}}, X_{G_t}, 1_{(\tau_{\{d\}}\leq t)}] \\
&\qquad \times \prod_{i=1}^{n-1} f_i(T_{T,i}, X_{T,i}) \mid (X_s)_{s\leq T}, X_{G_t}, 1_{(\tau_{\{d\}}\leq t)}\Bigg] \\
&\overset{\text{(iii)}}{=} \mathbb{E}\Bigg[1_{(T<\infty)}1_{(G_t\leq T)}\mathbb{1}(t)\tilde{E}^{X_{T_{T,n-1}}}[f_n(\tau_1, \eta_1)] \\
&\qquad \times \prod_{i=1}^{n-1} f_i(T_{T,i}, X_{T,i}) \mid (X_s)_{s\leq T}, X_{G_t}, 1_{(\tau_{\{d\}}\leq t)}\Bigg] \\
&= 1_{(T<\infty)}1_{(G_t\leq T)}\mathbb{1}(t)\mathbb{E}\Bigg[\tilde{E}^{X_{T_{T,n-1}}}[f_n(\tau_1, \eta_1)] \\
&\qquad \times \prod_{i=1}^{n-1} f_i(T_{T,i}, X_{T,i}) \mid (X_s)_{s\leq T}, X_{G_t}, 1_{(\tau_{\{d\}}\leq t)}\Bigg] \\
&= 1_{(T<\infty)}1_{(G_t\leq T)}\mathbb{1}(t)\tilde{E}^{X_T}\left[\prod_{i=1}^{n-1} f_i(\tau_i, \eta_i)\tilde{E}^{\eta_{n-1}}[f_n(\tau_1, \eta_1)]\right]
\end{aligned}
$$

where the first equality follows by the tower property and the second follows by (iii) with $f = f_n$ and the $\mathcal{F}^X$-stopping time $T_{T,n-1}$ using also that $(G_t \leq T) \subseteq (G_t \leq T_{T,n-1})$. The last equality follows by the induction hypothesis. We finally see,

$$
\begin{aligned}
&1_{(T<\infty)}1_{(G_t\leq T)}\mathbb{1}(t)\tilde{E}^{X_T}\left[\prod_{i=1}^{n-1} f_i(\tau_i, \eta_i)\tilde{E}^{\eta_{n-1}}[f_n(\tau_1, \eta_1)]\right] \\
&= 1_{(T<\infty)}1_{(G_t\leq T)}\mathbb{1}(t)\tilde{E}^{X_T}\left[\prod_{i=1}^{n} f_i(\tau_i, \eta_i)\right]
\end{aligned}
$$

by writing out the expectations using the Markov kernels. This concludes the induction and the proof.

□

4.B Stochastic interest rate

We consider the extension of the results to models with stochastic interest rates. Assume that $\kappa(t) = \exp\left(\int_{(0,t]} r(v)\,\mathrm{d}v\right)$ with $r : \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}$ being a stochastic process. We assume that there exists an equivalent martingale measure for the financial market which we denote $\mathbb{Q}$. Define the time-t *forward interest rate* as

$$f(t, s) = -\frac{\partial \log P(t, s)}{\partial s}$$

where for $0 \le t \le s$ we have $P(t,s) = \mathbb{E}^{\mathbb{Q}}\left[\kappa(t)/\kappa(s) \mid \mathcal{F}_t^r\right]$ is the price at time t of a zero-coupon bond paying one unit at time s. Then $P(t,s) = \exp\left(\int_{(t,s]} f(t,v)\,\mathrm{d}v\right)$. We introduce

$$\kappa_u(t) = \exp\left(\int_{(0,t]} r(u,v)\,\mathrm{d}v\right)$$

for

$$r(u,v) := \begin{cases} r(v), & v \le u \\ f(u,v), & v > u \end{cases}$$

being the realized interest rate before time u and the forward interest rate after time u. Note for $t \le u \le s$ that

$$\begin{aligned}\mathbb{E}^{\mathbb{Q}}\left[\frac{\kappa(s)}{\kappa(t)} \,\middle|\, \mathcal{F}_u^r\right] &= \frac{\kappa(u)}{\kappa(t)}\mathbb{E}^{\mathbb{Q}}\left[\frac{\kappa(s)}{\kappa(u)} \,\middle|\, \mathcal{F}_u^r\right] \\ &= \exp\left(\int_{(t,u]} r(v)\,\mathrm{d}v\right)\exp\left(\int_{(u,s]} f(u,v)\,\mathrm{d}v\right) \\ &= \frac{\kappa_u(s)}{\kappa_u(t)}.\end{aligned}$$

The extension of our results to a stochastic interest rate is simple if there is no dependence or conditional dependence between the filtrations $\mathcal{F}^r$ and $\mathcal{F}^{\mathcal{Z}}$. Redefine the transaction time reserve as

$$\mathcal{V}(t) = \mathbb{E}[\mathcal{P}(t) \mid \mathcal{F}_t^{\mathcal{Z}} \vee \mathcal{F}_t^r]$$

where now $\mathbb{E} = \mathbb{E}^{\mathbb{P}\otimes\mathbb{Q}}$. By Theorem 5.4 of Buchardt et al. (2023), we see

$$\begin{aligned}\mathcal{V}(t) = \mathbb{E}\Big[\int_{[0,\infty)} \frac{\kappa(t)}{\kappa(s)} B(\mathrm{d}s) \,\Big|\, \mathcal{F}_t^{\mathcal{Z}} \vee \mathcal{F}_t^r\Big] \\ - \mathbb{E}\Big[\sum_{j=1}^J \int_{[0,t]} \frac{\kappa(t)}{\kappa(s)} 1_{(Y_{s-}^t=j)} B_{j,s-U_{s-}^t}(\mathrm{d}s) \\ + \sum_{\substack{j,k=1\\ j\neq k}}^J \int_{[0,t]} \frac{\kappa(t)}{\kappa(s)} b_{jk}(s, U_{s-}^t) N_{jk}^t(\mathrm{d}s) \,\Big|\, \mathcal{F}_t^{\mathcal{Z}} \vee \mathcal{F}_t^r\Big].\end{aligned}$$

Using the independence between $\mathcal{F}^r$ and $\mathcal{F}^{\mathcal{Z}}$, standard arguments now give

$$\begin{aligned}\mathcal{V}(t) = \mathbb{E}\Big[\int_{[0,\infty)} \frac{\kappa_t(t)}{\kappa_t(s)} B(\mathrm{d}s) \,\Big|\, \mathcal{F}_t^{\mathcal{Z}} \vee \mathcal{F}_t^r\Big] \\ - \mathbb{E}\Big[\sum_{j=1}^J \int_{[0,t]} \frac{\kappa_t(t)}{\kappa_t(s)} 1_{(Y_{s-}^t=j)} B_{j,s-U_{s-}^t}(\mathrm{d}s) \\ + \sum_{\substack{j,k=1\\ j\neq k}}^J \int_{[0,t]} \frac{\kappa_t(t)}{\kappa_t(s)} b_{jk}(s, U_{s-}^t) N_{jk}^t(\mathrm{d}s) \,\Big|\, \mathcal{F}_t^{\mathcal{Z}} \vee \mathcal{F}_t^r\Big].\end{aligned}$$

This expression is identical to the expression for $\mathcal{V}$ in the case of a deterministic interest rate, but where κ is replaced with κ_t. Now since $\mathcal{F}^r$ is completely exogeneous to $\mathcal{F}^{\mathcal{Z}}$, all the calculations in Section 4.3 remain valid with κ replaced by κ_t, and hence we conclude that the results presented in Theorem 4.3.4, 4.3.8, and 4.3.9 still hold true if one substitutes κ_t for κ in all terms.

Remark 4.B.1. (Run-off plots.)
Run-off plots are a common graphical tool used to validate the reserves. The intuition is that if the reserves are correctly specified, the reserve should be converted to payments such that the sum of the two stays constant over time, see Figure 4.5 for an illustration. We show that if interest is handled appropriately, this approach is justified.

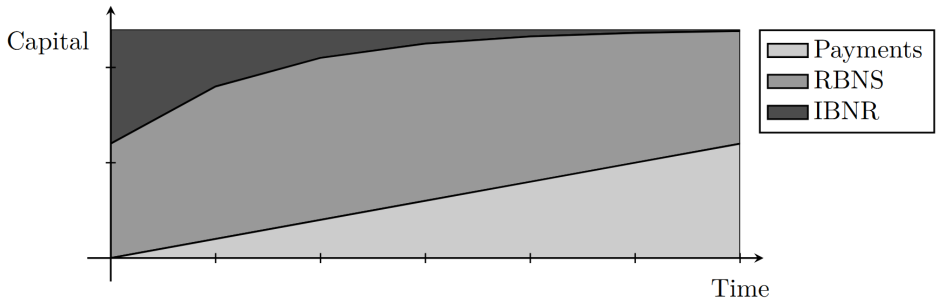

Figure 4.5: *Illustration of a run-off plot with origin at the end of the coverage period, where reserves consist of only IBNR and RBNS contributions.*

Assume that reserves are calculated with a fixed frequency (e.g. monthly) which we take to be the unit of the time scale. Let $\mathcal{V}^{\mathrm{pc}}(t) = \mathbb{E}[\mathcal{P}(t) \mid \mathcal{F}_t^{\mathcal{Z}} \vee \mathcal{F}_{t-1}^r]$ be the reserve calculated with the previous interest rate curve. Since increasing the number of policies does not diversify the market risk, we focus on validating the non-financial model elements. For one policy, straightforward calculations give

$$\begin{aligned} &\frac{\kappa_{t-1}(t-1)}{\kappa_{t-1}(t)}\mathcal{V}^{\mathrm{pc}}(t) - \mathcal{V}(t-1) \\ &= -\int_{(t-1,t]} \frac{\kappa_{t-1}(t-1)}{\kappa_{t-1}(s)}\mathcal{B}(\mathrm{d}s) + \mathbb{E}[\mathcal{P}(t-1) \mid \mathcal{F}_t^{\mathcal{Z}} \vee \mathcal{F}_{t-1}^r] - \mathcal{V}(t-1). \end{aligned} \tag{4.B.1}$$

Hence, the change in the reserves where the interest rate curve is kept fixed at the previous interest rate curve and discounted one time unit is equal to the realized payments discounted to the previous time unit plus a term which has mean zero conditional on $\mathcal{F}_{t-1}^{\mathcal{Z}} \vee \mathcal{F}_{t-1}^r$. In other words, the latter term has mean zero if the non-financial part of the model is correctly specified no matter the prior financial and non-financial developments. Since this non-financial risk can be diversified away by increasing the size of the portfolio, it becomes negligible when summing Equation (4.B.1) over many policies. Hence, the height of the stacked curves of $t \mapsto \mathcal{V}(0) + \sum_{m=1}^{t} \kappa_{m-1}(m-1)/\kappa_{m-1}(m) \times \mathcal{V}^{\mathrm{pc}}(m) - \mathcal{V}(m-1)$ and $t \mapsto \mathcal{B}(0) + \sum_{m=1}^{t} \int_{(m-1,m]} \kappa_{m-1}(m-1)/\kappa_{m-1}(s)\, \mathcal{B}(\mathrm{d}s)$ summed over the policies is excepted to stay constant throughout when the number of policies used in the sample is sufficiently large. $\triangledown$

4.C Estimation

This section details how to embed the statistical problem of this paper into the one from Buchardt et al. (2025) and provides a more accessible exposition of their estimation procedure applied to the current setting.

4.C.1 Statistical model

The disability and reactivation events may both be affected by reporting delays, but the death event is not. The disability reporting delay is $U_{\mathcal{I}} : \Omega \mapsto \mathbb{R}_+$ with $U_{\mathcal{I}} = 1_{(\tau_{\mathcal{I}} < \infty)}(T_{\{2\}} - \tau_{\mathcal{I}})$. Estimating the IBNR-factor $I_i(s,t)$ when $s < \infty$ is equivalent to estimating the disability reporting delay distribution since

$$I_i(s,t) = \mathbb{P}(\mathrm{CBNR}_t \mid \tau_{\mathcal{I}} = s, Y_{\tau_{\mathcal{I}}} = i) = \mathbb{P}(U_{\mathcal{I}} > t - s \mid \tau_{\mathcal{I}} = s, Y_{\tau_{\mathcal{I}}} = i).$$

Let $R : \Omega \mapsto \mathbb{R}_+$ with $R = \inf\{s \geq 0 : G_s = G_\infty\}$ be the final time where payments are stopped. The reactivation reporting delay $U_r : \Omega \mapsto \mathbb{R}_+$ is $U_r = 1_{(\tau_{\{r\}} < \infty)}(R - \tau_{\{r\}})$. Note that there is no reporting delay when the insurer terminates the running payments, but a jump of $Z^{(1)}$ from state 2 to state 3 or state 5 which triggers backpay may lead to a reactivation reporting delay. Estimation of the reactivation reporting delay distribution is not needed for reserve calculation but is needed for our proposed estimator of the valid time hazards.

Both the disability and reactivation events also have non-trivial adjudications while death events do not. One can calculate the adjudication probability $\mathbb{P}(X_{G_t} = (G_t, Z_t^{(3)}, W_t) \mid \mathcal{F}_t^{\mathcal{Z}})$ as an absorption probability for a multistate model on the state space $\mathcal{J}^\omega = \{1,2,3,4,5\}$ depicted in Figure 4.6 with a set of $\mathcal{F}^{\mathcal{Z}}$-predictable transition intensities $\omega_{jk} : \Omega \times \mathbb{R}_+ \mapsto \mathbb{R}_+$ ($j,k \in \mathcal{J}^\omega, j \neq k$). Disability benefits are awarded if and only if the process is absorbed in state 3 or 5. The multistate model of Figure 4.6 starts each time a disability or reactivation event is reported. A disability adjudication starts in state 1, while a reactivation adjudication starts in state 2. Note that since the adjudication hazards are $\mathcal{F}^{\mathcal{Z}}$-predictable, they can and will be different when the adjudication pertains to a disability or a reactivation event. The shared notation for the adjudication hazards and the state space is chosen for parsimony.

Let $\sigma(t) = \inf\{v \geq 0 : W_v > W_t\}$ be the next time where disability benefits are awarded after time t. The multistate model of Figure 4.6 corresponds to modeling $s \mapsto \mathcal{Z}_s$ on the interval $(t, \sigma(t)]$ using the self-exciting filtration $\mathcal{F}^{\mathcal{Z}}$ except that the mark at time $\sigma(t)$ is the reduced mark $(Z_{\sigma(t)}, 1_{(W_{\sigma(t)}>W_t)})$ as opposed to the full mark $\mathcal{Z}_{\sigma(t)}$. When the object of interest is the adjudication probability, there is no need to model the full mark, which would entail modeling how much backpay was awarded, and as a consequence also whether the insured receives running benefits after the award or not.

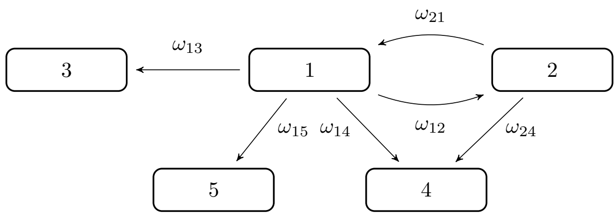

Figure 4.6: *Multistate model for adjudications. Active report is 1, inactive report is 2, awarded is 3, dead without award is 4, and dead with award is 5.*

Let the observation window be $[0, \eta]$ and let the valid time process X be subject to independent left-truncation and right-censoring, hence being observed on a random interval $(V, C] \subseteq [0, \eta]$ even in the absence of reporting delays and adjudication processes. Events that occurred in $(V, C]$ can be reported and adjudicated up until time η. Baseline covariates, meaning covariates that are known at time 0 in the valid time and transaction time filtrations, can easily be incorporated here and in the rest of the paper by conditioning on them throughout. Define the parameter spaces $\mathbb{G}$, $\mathbb{F}$, and Θ being subsets of Euclidean spaces. The statistical model is denoted $\mathcal{P} = \{P_{(g,f,\theta)} : (g, f, \theta) \in \mathbb{G} \times \mathbb{F} \times \Theta\}$. The adjudication intensities under $P_{(g,f,\theta)}$ only depend on (g, f, θ) through g and are denoted $t \mapsto \omega_{jk}(t; g)$. The reporting delay distributions under $P_{(g,f,\theta)}$ only depend on (g, f, θ) through f and are denoted $t \mapsto \mathbb{P}(U_{\mathcal{I}} \leq t \mid \tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}}, f)$ and $t \mapsto \mathbb{P}(U_r \leq t \mid \tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}}, \tau_{\{r\}}, f)$. The valid time hazards under $P_{(g,f,\theta)}$ only depend on (g, f, θ) through θ and are denoted $(t, u) \mapsto \mu_{jk}(t, u; \theta)$. We now describe the proposed estimators.

4.C.2 Adjudication probabilities

Let $N^{\omega}_{jk} : \Omega \times \mathbb{R}_+ \mapsto \mathbb{N}_0$ denote the counting processes on $\mathcal{J}^{\omega}$ where jumps due to starting a new adjudication are excluded. The objective function is the log-likelihood, which for one insured is

$$\ell_{\omega}(g) = \sum_{\substack{j,k \in \mathcal{J}^{\omega} \\ j \neq k}} \int_{(0,\eta]} \log\{\omega_{jk}(t; g)\}\, N^{\omega}_{jk}(\mathrm{d}t) - \int_{(0,\eta]} \omega_{jk}(t; g)\, \mathrm{d}t.$$

For n i.i.d. insured, the log-likelihood which we denote $\ell^{(n)}_{\omega}(g)$ is a sum over n such terms. The estimator is the maximum likelihood estimator $\hat{g}_n = \arg\max_{g \in \mathbb{G}} \ell^{(n)}_{\omega}(g)$. The argmax can be found using existing `glm` software packages, see Section D in the supplementary material of Buchardt et al. (2025) for details.

4.C.3 Reporting delay distribution

The reporting delays are right-truncated since only jumps reported before η are part of the sample. We first discuss disability reporting delays and subsequently reactivation reporting delays. Since we for disability reporting delays condition on $\tau_{\mathcal{I}}$, only disability claims that will be awarded should be included in the estimation of the reporting delay distribution for disability events. We accommodate this by weighting the relevant objective function with the adjudication probability. Assume that $U_{\mathcal{I}}$ given $(\tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}})$ has density with respect to a common reference measure across $\mathcal{P}$. We informally denote this density evaluated at a point u by $\frac{\mathrm{d}}{\mathrm{d}u}\mathbb{P}(U_{\mathcal{I}} \leq u \mid \tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}}, f)$. For one subject, the objective function for an observed disability reporting delay u is

$$\begin{aligned}\ell_{U_{\mathcal{I}}}(f;\hat{g}_n) &= \mathbb{P}(X_{Z_\eta^{(2)}} = (Z_\eta^{(2)}, Z_\eta^{(3)}, 0) \mid \mathcal{F}_\eta^{\mathcal{Z}}, \hat{g}_n) \\ &\quad \times \log\left\{\frac{\frac{\mathrm{d}}{\mathrm{d}u}\mathbb{P}(U_{\mathcal{I}} \leq u \mid \tau_{\mathcal{I}} = Z_\eta^{(2)}, Y_{\tau_{\mathcal{I}}} = Z_\eta^{(3)}, f)}{\mathbb{P}(U_{\mathcal{I}} \leq \eta - \tau_{\mathcal{I}} \mid \tau_{\mathcal{I}} = Z_\eta^{(2)}, Y_{\tau_{\mathcal{I}}} = Z_\eta^{(3)}, f)}\right\}.\end{aligned}$$

Here the first term is the adjudication probability $\mathbb{P}(X_{G_\eta} = (G_\eta, Z_\eta^{(3)}, 0) \mid \mathcal{F}_\eta^{\mathcal{Z}}, \hat{g}_n)$ on (RBNSi$_\eta$) and is 1 if the disability has been awarded. The objective function for n i.i.d. subjects is a sum of n such terms and we denote this by $\ell_{U_{\mathcal{I}}}^{(n)}(f;\hat{g}_n)$.

Similarly, assume U_r given $(\tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}}, \tau_{\{r\}})$ has density with respect to a common reference measure across $\mathcal{P}$ which we informally denote $\frac{\mathrm{d}}{\mathrm{d}w}\mathbb{P}(U_r \leq w \mid \tau_{\mathcal{I}}, Y_{\tau_{\mathcal{I}}}, \tau_{\{r\}}, f)$. For one subject, the objective function for an observed reactivation reporting delay w is

$$\begin{aligned}\ell_{U_r}(f;\hat{g}_n) &= (1 - \mathbb{P}(X_{G_\eta} = (G_\eta, Z_\eta^{(3)}, W_\eta) \mid \mathcal{F}_\eta^{\mathcal{Z}}, \hat{g}_n)) \\ &\quad \times \log\left\{\frac{\frac{\mathrm{d}}{\mathrm{d}w}\mathbb{P}(U_r \leq w \mid \tau_{\mathcal{I}} = Z_\eta^{(2)}, Y_{\tau_{\mathcal{I}}} = Z_\eta^{(3)}, \tau_{\{r\}} = G_\eta, f)}{\mathbb{P}(U_r \leq \eta - \tau_{\{r\}} \mid \tau_{\mathcal{I}} = Z_\eta^{(2)}, Y_{\tau_{\mathcal{I}}} = Z_\eta^{(3)}, \tau_{\{r\}} = G_\eta, f)}\right\}.\end{aligned}$$

Let the objective function for n i.i.d. subjects be denoted $\ell_{U_r}^{(n)}(f;\hat{g}_n)$. The estimator is then $\hat{f}_n = \arg\max_{f\in\mathbb{F}} \ell_{U_{\mathcal{I}}}^{(n)}(f;\hat{g}_n) + \ell_{U_r}^{(n)}(f;\hat{g}_n)$.

Remark 4.C.1. (Chain ladder estimator.)
The reporting delay distribution can alternatively be estimated using chain ladder, see for example Section 5 in Bücher & Rosenstock (2024), but the large-sample properties are then not a special case of Buchardt et al. (2025) since they consider parametric estimators. Note that chain ladder in this case does not estimate a reserve but rather an element of the individual reserve, namely the IBNR-factor. As seen in Section 4.3, an assumption like Assumption 4.3.1 is needed for the IBNR reserve to approximately decompose into an IBNR-factor-adjusted frequency multiplied with a classic valid time disability reserve as the associated claim size. $\triangledown$

4.C.4 Valid time hazards

Make a partition of the observation window $0 = t_0 < t_1 < \cdots < t_L = \eta$ and let $O_{jk}(\ell) = N^{\eta}_{jk}(t_{\ell+1}) - N^{\eta}_{jk}(t_\ell)$ and $E_j(\ell) = \int_{(t_\ell, t_{\ell+1}]} 1_{(Y^{\eta}_s = j)}\,\mathrm{d}s$ $(\ell = 0, \ldots, L-1)$ be the occurrences and exposures based on H^η for a single insured. To describe how the duration process is affected by accepting or rejecting a disability claim, we introduce the auxiliary durations $U_j(\ell) = U^{\eta}_{t_\ell}$ for $j \in \mathcal{J}$. Note that there is initially no dependence on j.

If there is an unadjudicated disability at time η, the occurrences, exposures, and durations are modified:

$$
\begin{aligned}
O_{ai}(\ell) &\leftarrow \mathbb{P}(X_{G_\eta} = (G_\eta, i, 0) \mid \mathcal{F}^{\mathcal{Z}}_\eta, \hat{g}_n), && G_\eta \in (t_\ell, t_{\ell+1}],\\
E_a(\ell) &\leftarrow (t_{\ell+1} - t_\ell) \times (1 - \mathbb{P}(X_{G_\eta} = (G_\eta, i, 0) \mid \mathcal{F}^{\mathcal{Z}}_\eta, \hat{g}_n)), && G_\eta \in (0, t_\ell],\\
E_i(\ell) &\leftarrow (t_{\ell+1} - t_\ell) \times \mathbb{P}(X_{G_\eta} = (G_\eta, i, 0) \mid \mathcal{F}^{\mathcal{Z}}_\eta, \hat{g}_n), && G_\eta \in (0, t_\ell],\\
U_i(\ell) &\leftarrow t_\ell - G_\eta, && G_\eta \in (0, t_\ell],
\end{aligned}
$$

for $i = Z^{(3)}_\eta$. If there is an unadjudicated reactivation at time η, the occurrences, exposures, and durations are modified:

$$
\begin{aligned}
O_{ir}(\ell) &\leftarrow (1 - \mathbb{P}(X_{G_\eta} = (G_\eta, i, W_\eta) \mid \mathcal{F}^{\mathcal{Z}}_\eta, \hat{g}_n)), && G_\eta \in (t_\ell, t_{\ell+1}],\\
E_i(\ell) &\leftarrow (t_{\ell+1} - t_\ell) \times \mathbb{P}(X_{G_\eta} = (G_\eta, i, W_\eta) \mid \mathcal{F}^{\mathcal{Z}}_\eta, \hat{g}_n), && G_\eta \in (0, t_\ell],\\
E_r(\ell) &\leftarrow (t_{\ell+1} - t_\ell) \times (1 - \mathbb{P}(X_{G_\eta} = (G_\eta, i, W_\eta) \mid \mathcal{F}^{\mathcal{Z}}_\eta, \hat{g}_n)), && G_\eta \in (0, t_\ell],\\
U_i(\ell) &\leftarrow t_\ell - Z^{(2)}_\eta, && G_\eta \in (0, t_\ell],
\end{aligned}
$$

for $i = Z^{(3)}_\eta$. Finally, in all cases, the exposures are modified with the reporting delay distribution:

$$
\begin{aligned}
E_{ai}(\ell) &\leftarrow E_a(\ell) \times \mathbb{P}(U_{\mathcal{I}} \leq \eta - \tau_{\mathcal{I}} \mid \tau_{\mathcal{I}} = t_\ell, Y_{\tau_{\mathcal{I}}} = i, \hat{f}_n),\\
E_{ir}(\ell) &\leftarrow E_i(\ell) \times \mathbb{P}(U_r \leq \eta - \tau_{\{r\}} \mid \tau_{\mathcal{I}} = Z^{(2)}_\eta, Y_{\tau_{\mathcal{I}}} = i, \tau_{\{r\}} = t_\ell, \hat{f}_n),
\end{aligned}
$$

for all $i \in \mathcal{I}$. For the remaining transitions set $E_{jk}(\ell) \leftarrow E_j(\ell)$.

The objective function $\ell(\theta; \hat{g}_n, \hat{f}_n)$ is the log-likelihood resulting from assuming that $(O_{jk}(\ell))_{j,k,\ell}$ are independent Poisson distributed random variables with mean values $(\mu_{jk}(t_\ell, U_j(\ell); \theta) E_{jk}(\ell))_{j,k,\ell}$. Let the objective function for n i.i.d. subjects be denoted $\ell^{(n)}(\theta; \hat{g}_n, \hat{f}_n)$. The estimator is $\hat{\theta}_n = \arg\max_{\theta \in \Theta} \ell^{(n)}(\theta; \hat{g}_n, \hat{f}_n)$, which corresponds to the Poisson approximation from Buchardt et al. (2025). The argmax can be found using existing `glm` software packages when the modified occurrences, exposures, and durations have been constructed. Note that if two occurrences have the same mean value, these occurrences and their corresponding exposures can be summed without changing the objective function. This aggregation may save memory and speed up computations.

4.C.5 Asymptotic properties

The Poisson approximation introduces bias that does not vanish asymptotically, but which is small when the hazards and reporting delays are small, see Section 3.3 of Buchardt et al. (2025) or Section B.2 of their supplementary material. The approximation error for the disability hazard is expected to be negligible since the hazard appears to be smaller than 10^{-2} by some margin. There is also a small approximation bias coming from implementing the Poisson approximation via occurrences and exposures, see Section D in the supplementary material of Buchardt et al. (2025) for details. Denote by $B_n = \hat{\theta}_n - \hat{\theta}_n^{\text{full}}$ the approximation bias, where $\hat{\theta}_n^{\text{full}}$ is the estimator based on the full procedure. We use $\overset{\text{a}}{\sim}$ to denote asymptotic distribution. The asymptotic properties of the estimators are given in Proposition 4.C.2.

Proposition 4.C.2. *Under Assumptions 1-3 and 5-8 from Buchardt et al. (2025), which are standard integrability and smoothness conditions,*

$$(\hat{g}_n, \hat{f}_n, \hat{\theta}_n) \overset{\text{a}}{\sim} \mathcal{N}((0, 0, B_n), \Sigma)$$

for a non-singular covariance matrix Σ.

Proof. This follows directly from Theorem 1 of Buchardt et al. (2025). □

Chapter 5

Estimation for multistate models subject to reporting delays and incomplete event adjudication with application to disability insurance

This chapter contains the manuscript *Buchardt, Furrer, and Sandqvist (2025).*

Abstract

Accurate forecasting of an insurer's outstanding liabilities is vital for the solvency of insurance companies and the financial stability of the insurance sector. For health and disability insurance, the liabilities are intimately linked with the biometric event history of the insured. Complete observation of event histories is often impossible due to sampling effects such as right-censoring and left-truncation, but also due to reporting delays and incomplete event adjudication. In this paper, we develop a parametric two-step M-estimation method that takes the aforementioned effects into account, treating the latter two as partially exogenous. The approach is valid under weak assumptions and allows for complicated dependencies between the event history, reporting delays, and adjudication while remaining relatively simple to implement. The estimation approach has desirable properties which are demonstrated by theoretical results and numerical experiments.

In the application, we introduce and consider a large portfolio of disability insurance policies. We find that properly accounting for the sampling effects has a large impact on the number of disabilities and reactivations that an insurer would forecast, allowing for a more accurate assessment of the insurer's liabilities and improved risk management.

Keywords: Event history analysis; Health insurance; Incomplete event adjudication; Reporting delay; Two-step M-estimation

5.1 Introduction

Health and disability insurance is a fundamental pillar of modern healthcare systems, providing individuals with financial protection against the costs of medical treatment and the loss of income due to injury or illness. In 2023, the health insurance industry of the United States collected more than \$1 trillion in premiums and paid out hundreds of billions of dollars in medical claims according to the National Association of Insurance Commissioners (2023). For insurers, managing future liabilities such as outstanding disability benefits is crucial not only for their own financial sustainability but also for the stability of the entire health insurance market. Accurate forecasting of these liabilities, so-called reserving, is therefore vital in ensuring that insurers can meet their obligations and maintain solvency.

Insurance benefits for health and disability insurance coverages are directly tied to the insured's biometric state, making it possible to focus on modeling this underlying process. That approach is typically preferable to modeling individual payments and their intertemporal dependencies directly; this is, on the other hand, usually required for individual reserving models in non-life insurance, see e.g. Yang et al. (2024). Multistate models provide a natural and parsimonious way to model the biometric state process of an individual and are therefore the focus of this paper. By modeling individual policies, the insurer may leverage granular data to predict and forecast outstanding liabilities.

We develop estimation methods for multistate models that allow actuaries and statisticians to employ all recently generated data while accommodating the contamination induced by reporting delays and incomplete event adjudications. This assists insurers in performing timely operational adjustment and risk mitigation when faced with emerging health trends such as the major increase in claims related to mental illnesses that has been observed in recent years, see for example Section 1 of Swiss Re (2022).

Longitudinal biometric data such as those arising from insurance policies consists of records on the occurrence and timing of certain events. The analysis of such data is typically referred to as event history analysis. Complete observation of the event history is often rendered impossible due to sampling effects, encompassing censoring and truncation as well as situations where information about individuals under study is not up to date or correct; this may for example be due to periodic sampling, reporting delays, or incomplete event adjudication. Incomplete event adjudication refers to the situation where it is undetermined, at the time of analysis, whether some reported events satisfy predetermined criteria for being true events. Ignoring these mechanisms leads to misleading and biased analyses. Censoring and truncation are most fundamental to the field of event history analysis and have so far received the most attention.

Our main focus is on handling reporting delays and incomplete event adjudication in the estimation of the conditional hazards of a multistate model, but we also allow for right-censoring and left-truncation. The goal is to use the available information to obtain as reliable predictions as possible. The available information consists of event dates and their corresponding event types, reporting delays, and adjudication information, but only for those events that are reported before the time of analysis.

5.1.1 Relevant literature

A naive way to handle reporting delays is back-censoring, see Casper & Cook (2012). Here, the right-censoring date is set back by some fixed amount, corresponding to only using data that is older than a given date. The approach does not introduce systematic bias if the right-censoring time is set back by an amount larger than the longest delay but is inefficient in that it may discard many observations. Since observations for later times are discarded, tail-estimates and trends may be especially impaired. This impedes forecasting and hence the ability to predict outstanding liabilities.

More efficient approaches have been proposed for survival models, see Hu & Tsiatis (1996) and Van der Laan & Hubbard (1998), and recurrent event models, see Kalbfleisch & Lawless (1989), Pagano et al. (1994), Becker & Cui (1997), Casper & Cook (2012), and Verbelen et al. (2022). Most of these focus on non-parametric estimation with inverse probability of censoring weighting being the most popular approach, but other approaches include an EM-algorithm for a multinomial model in Pagano et al. (1994), maximum likelihood via thinning of a Poisson process in Kalbfleisch & Lawless (1989), and a mix of these in Verbelen et al. (2022). In most of the aforementioned works, the marginal mean of the counting process is modeled, except in Kalbfleisch & Lawless (1989) and Verbelen et al. (2022) where hazards are modeled. Finally, some works take estimation of the reporting delay distribution as the main focus and either forgo estimation of the event hazards, see Lagakos et al. (1988) and Kalbfleisch & Lawless (1991), or use simple methods to adjust the event rates after estimating the reporting delay distribution, see Esbjerg et al. (1999).

The incorporation of reporting delays has, to the best of our knowledge, not hitherto been studied for non-competing-risks multistate models such as the illness-death models that are relevant for modeling disability insurance events. Additionally, we are the first to explore estimation under a monotone reporting assumption, which requires events to be reported in the same order in which they occurred. This assumption is natural when modeling individuals. Recurrent event models in the literature have so far been based on the assumption that reporting delays are independent of the event process; this assumption only seems plausible for aggregate models.

For incomplete event adjudication, a simple but inefficient solution is to use only confirmed events or all unrefuted events while back-censoring. For the method to be unbiased, one has to back-censor by more than the maximal sum of the reporting and processing delays, in which case the unrefuted and confirmed events are the same. Other approaches are explored in Cook (2000), Cook & Kosorok (2004), and Bladt & Furrer (2024), but only for single events without reporting delays and using methods that do not easily generalize to hazard estimation for multistate models, and where the adjudication outcome distribution is either exogeneously given or is assumed to only depend on the information available at the event time. In addition, limited attention was given to the estimation of the adjudication probabilities: Cook (2000) and Cook & Kosorok (2004) suggest logistic regression based on completed adjudications, while Bladt & Furrer (2024) relies on expert judgments.

Both reporting delays and incomplete event adjudication are prevalent in health and disability insurance. Disabilities only become known to the insurer when they are reported by the insured, at which point the event date also becomes known, as it determines the date from which the insured is eligible for disability benefits. Adjudication is performed primarily by the insurer, who considers whether the claim satisfies the criteria for disability benefits and determines at what stage the insured has recovered sufficiently to terminate benefits. This is often complicated by the need to obtain clinical assessments of the claimed disability.

5.1.2 Model considerations

In order to produce accurate predictions, actuarial forecasts should be dynamic in the sense that they are updated as new information arrives. Hazard rates provide a convenient way to parameterize all relevant conditional distributions, so that insurance reserves and other conditional estimands may be calculated using well-known integral or differential equations, see Hoem & Aalen (1978), Møller (1993), Norberg (2005), and Adékambi & Christiansen (2017). Consequently, this paper focuses on hazard estimation.

Baseline covariates and the event history may have a large effect on the distribution of future jumps, so we take a regression approach to incorporate these effects. Furthermore, we find parametric models advantageous for several reasons. Firstly, we need to extrapolate the estimates to times and durations that exceed the observation window. This holds both for the parameters of interest, being those related to the biometric events, and the nuisance parameters related to the reporting delay and adjudication processes. This makes the otherwise popular Kaplan–Meier, Aalen–Johansen, and Cox-type estimators less attractive.

Secondly, feature engineering can be used in combination with existing expert knowledge to guide the parametric specification since the focus is on prediction rather than inference. This may be helpful in counteracting some of the challenging characteristics of disability insurance data, such as the relative rareness of events, the low signal-to-noise rate, and the moderate dimension of covariates. Many of these characteristics are shared with nowcasting and forecasting of epidemics, see for example Noufaily et al. (2016), Psotka et al. (2020), and Stoner et al. (2023), which also constitutes a relevant area of application for our methods. As discussed in Cook (2000), Cook & Kosorok (2004), and Casper & Cook (2012), another possible area of application is interim stages of clinical trials, but here semi- and non-parametric models may be preferred, given the focus is often on inference for non-parametrically identifiable marginal estimands.

5.1.3 Our contribution

The contributions of this paper are twofold. Firstly, we develop a parametric approach that encompasses both reporting delays and incomplete event adjudication and which applies to hazard estimation for multistate models in general. The approach has many properties that are desirable for insurance applications and relaxes several assumptions from the literature. We also introduce an approximation that simplifies the estimation procedure considerably and which performs well in many situations encountered in practice. In addition, we show that our estimators are consistent and asymptotically normal under suitable regularity conditions and that they may be bootstrapped.

Secondly, the proposed methods are applied to a new data set based on a large Danish disability insurance portfolio. This application is noteworthy for two reasons. First, the data is unique in that it seems to be the first insurance data set to include information on event adjudication and the first disability or health insurance data set to contain information on reporting delays. Having access to such data and accounting for the time of analysis is essential to ensure unbiased analyses. Second, the analyses show that properly accommodating for reporting delays and incomplete event adjudication has a substantial effect on the estimated level and calendar-time dependence of the disability and reactivation hazards. This underscores the importance of our concepts and results. The impact of our contribution is further accentuated by taking the size and societal importance of the health and disability insurance sector into account.

The paper is organized as follows. Section 5.2 introduces the model. Section 5.3 discusses estimation of the parameters. Section 5.4 presents theoretical large-sample results. Section 5.5 contains numerical experiments, demonstrating desirable finite sample performance and stability under misspecification. The data application is presented in Section 5.6.

5.2 Model Specification

We consider a time-continuous multistate model with states $1, \dots, J$, which may be represented by a marked point process $(T_m^*, Y_m^*)_{m \geq 1}$, where $T_m^* \in (0, \infty]$ are jump times and $Y_m^* \in \{1, \dots, J\}$ are jump marks. Let $\mathbf{J}^*(t)$ denote the time t event history, containing the events that have occurred up until and including time t. The multistate model may equivalently be represented by a multivariate counting process $\mathbf{N}^*$ with components N_{jk}^*, so that $N_{jk}^*(t)$ counts the number of transitions from state j to state k in the interval $[0, t]$. Let $\mathbf{X}$ be baseline covariates, including for instance the initial state, taking values in a subset of $\mathbb{R}^p$. We suppose that N_{jk}^* satisfies a general intensity model

$$\mathbb{E}_{\boldsymbol{\theta}}\{N_{jk}^*(\mathrm{d}t) \mid \mathcal{F}_{t-}^*\} = \lambda_{jk}^*(t; \mathbf{X}, \boldsymbol{\theta}) \,\mathrm{d}t = I_j^*(t) \mu_{jk}^*\{t; \mathbf{J}^*(t-), \mathbf{X}, \boldsymbol{\theta}\} \,\mathrm{d}t, \tag{5.2.1}$$

where $\mathbb{E}_{\boldsymbol{\theta}}$ denotes expectation under the distribution with parameter $\boldsymbol{\theta} \in \Theta \subseteq \mathbb{R}^d$, $\mathcal{F}^*$ is the filtration generated by $\mathbf{N}^*$ and $\mathbf{X}$, and $I_j^*(t)$ is the indicator that the the last mark in $\mathbf{J}^*(t-)$ is j or the initial state is j if $\mathbf{J}^*(t-)$ is empty. Assume that the true parameter $\boldsymbol{\theta}_0$ belongs to the parameter set Θ. We seek to estimate $\boldsymbol{\theta}_0$, but subject to various forms of missingness and contamination. First, the observation of jumps is subject to delays. Second, not all jump times, jump marks, and delays are observed due to random left-truncation and right-censoring. Finally, some reported jumps are annulled due to event adjudication. In the following, we describe these mechanisms in more detail.

Let $\eta > 0$ be the deterministic time where the statistical analysis is conducted. Denote by $U_m^* \in [0, \infty)$ the random reporting delay associated with the m'th jump. We form a new thinned marked point process $(T_m, Y_m)_{m \geq 1}$ from $(T_m^*, Y_m^*)_{m \geq 1}$ by deleting all jumps with $T_m^* + U_m^* > \eta$. Let $(U_m)_{m \geq 1}$, $\mathbf{J}$, and N_{jk} be respectively the reporting delays, event history and counting processes associated with $(T_m, Y_m)_{m \geq 1}$. Throughout, we make the following assumption:

Assumption 5.2.1. *Reporting is monotone, meaning* $T_m^* + U_m^* \leq T_{m+1}^* + U_{m+1}^*$.

Assumption 5.2.1 is closely related to the concept of monotone missingness from the missing data literature, confer with Little (2021). Dropping this assumption would lead to a substantial increase in complexity, since even if $T_m < \infty$, there might still be a disagreement between $\mathbf{J}^*(T_m)$ and $\mathbf{J}(T_m)$. This assumption is natural when modeling individual subjects. Thinning has hitherto only been used for aggregate models when this assumption is replaced by an independence assumption. The true distribution function of U_m^* given $\mathbf{J}^*(T_m^*)$ and $\mathbf{X}$ is denoted $t \mapsto \mathrm{pr}_U\{t; \mathbf{J}^*(T_m^*), \mathbf{X}, \mathbf{f}_0\}$ for a parameter $\mathbf{f}_0$. That the components of $(U_m^*)_{m \geq 1}$ may have different distributions is captured by $\mathbf{J}^*(T_m^*)$ in the conditioning.

Observation of $\mathbf{N}$ is, as mentioned, also subject to left-truncation and right-censoring. To be precise, there is a random right-censoring time $C \leq \eta$, which could for example be the drop-out time of the subject, and a random left-truncation time $V < C$ where the subject enters the study if some event A has occurred prior to V. As an example, V could be when the subject enters the insurance portfolio, and the event A could be that the subject is still alive at that time. A comparable situation would be one where subjects not satisfying A or with $V > \eta$ were also considered part of the study but with no exposure. For example, if one had access to panel data, the current formulation would filter away subjects that are not in the portfolio during the observation window while the alternative formulation would keep the observations but let them have no exposure. In our applications, the event A is redundant as it is contained in the relevant filtration, but we allow for this extra generality which may be useful in other applications, confer with Section III.3 of Andersen et al. (1993).

Let $(T_m^c, Y_m^c)_{m\geq 1}$ be $(T_m, Y_m)_{m\geq 1}$ but where jumps outside of $(V, C]$ are deleted and let $(U_m^c)_{m\geq 1}$, N_{jk}^c, and $\mathbf{J}^c$ be the corresponding reporting delays, counting processes, and event history, respectively. We introduce the filtrations $\mathcal{F}_t = \sigma\{\mathbf{X}, (T_m, Y_m) : T_m \leq t\}$ and $\mathcal{F}_t^c = \sigma\{\mathbf{X}, \mathcal{G}, C \wedge t, (T_m^c, Y_m^c) : T_m^c \leq t\}$, where $\mathcal{G}$ is left-truncation information satisfying $A \in \mathcal{G}$ and V is $\mathcal{G}$-measurable; the delays are not included in these filtrations as the influence of prior delays on the intensities is not of interest when estimation of μ_{jk}^* is the final goal. Also, these filtrations are not directly observable (for instance, T_m is reported at time $T_m + U_m$, but already enters in the filtrations at time T_m), but they play a technical role in allowing us to formulate a criterion function from which to construct estimators.

Note that we allow for jumps to be reported between the right-censoring time C and the time of analysis η; this choice is appropriate for actuarial applications where subjects may leave the portfolio and afterwards report claims for events that occurred while they were still in the portfolio. In other applications, such as for medical trials, thinning according to $T_m^* + U_m^* > C$ might be more appropriate, as patients may be completely lost to follow-up at C. We stress that the results of this paper then still apply upon changing η to C in the reporting delay distribution, but under the additional assumption that C is independent of $(T_m^*, Y_m^*)_{m\geq 1}$. Furthermore, in some applications it might be relevant to also allow for periodic, but immediate, ascertainment of the current state. This has been explored for a survival model in Hu & Tsiatis (1996). We do not pursue this extension.

The following assumption, which we make throughout, is comparable with the assumption of independent filtering, confer with Section III.4 in Andersen et al. (1993).

Assumption 5.2.2. *Left-truncation and right-censoring are independent in relation to the marked point process $(T_m, Y_m)_{m\geq 1}$ in the sense that*

$$\mathbb{E}_{\boldsymbol{\theta},\mathbf{f}}\{N_{jk}^c(\mathrm{d}t) \mid \mathcal{F}_{t-}^c; A\} = 1_{(V,C]}(t)\mathbb{E}_{\boldsymbol{\theta},\mathbf{f}}\{N_{jk}(\mathrm{d}t) \mid \mathcal{F}_{t-}\}$$

for $(\boldsymbol{\theta}, \mathbf{f}) \in \Theta \times \mathbb{F}$*, where the notation on the left-hand side signifies conditioning on* $\mathcal{F}_{t-}^c$ *and* A*. Similarly, left-truncation and right-censoring are independent in relation to the reporting delays in the sense that* $\mathbb{P}_{\mathbf{f}}(U_m^c \leq t \mid \mathcal{F}_{T_m^c}^c; A) = \mathbb{P}_{\mathbf{f}}(U_\ell \leq t \mid \mathcal{F}_{T_\ell})$ *where* ℓ *is the index satisfying* $T_\ell = T_m^c$*. Here* $\mathbb{P}_{\mathbf{f}}$ *denotes the probability under the distribution with parameter* $\mathbf{f}$*.*

The third and final mechanism relates to incomplete event adjudication. The idea is to think of $(T_m^c, Y_m^c)_{m\geq 1}$ as obtained by thinning another marked point process $(\tilde{T}_m, \tilde{Y}_m)_{m\geq 1}$ accompanied by delays $(\tilde{U}_m)_{m\geq 1}$ satisfying $\tilde{T}_m + \tilde{U}_m \leq \eta$. To be precise, if $\xi_m \in \{0,1\}$ is the adjudication outcome of the m'th event, then $(T_m^c, Y_m^c)_{m\geq 1}$ is formed from $(\tilde{T}_m, \tilde{Y}_m)_{m\geq 1}$ by deleting all jumps with $\xi_m = 0$. Incomplete event adjudication thus becomes a missing data problem for the adjudication outcomes $(\xi_m)_{m\geq 1}$. The full available information is thus

$$\mathcal{H}_t^{\text{obs}} = \sigma\{\mathbf{X}, \mathcal{G}, C \wedge t, (\tilde{T}_m, \tilde{Y}_m, \tilde{U}_m, \mathcal{A}_{m,t}) : \tilde{T}_m + \tilde{U}_m \leq t\},$$

where $\mathcal{A}_{m,t}$ is a filtration representing adjudication information, such that when σ_m is the time where ξ_m becomes known, then σ_m is an $\mathcal{A}_{m,t}$-stopping time and ξ_m is $\mathcal{A}_{m,\sigma_m}$-measurable. Apart from σ_m and ξ_m, the filtration $\mathcal{A}_{m,t}$ may contain additional covariates that affect adjudication probabilities. It is possible to let the left-truncation information in $\mathcal{H}^{\text{obs}}$ be a superset of $\mathcal{G}$ without any additional effort.

The following assumption serves to reduce the modeling task:

Assumption 5.2.3. *Only the most recent jump of* $(\tilde{T}_m, \tilde{Y}_m)_{m\geq 1}$ *may be unadjudicated at any given time, meaning* $\sigma_m \leq \tilde{T}_{m+1} + \tilde{U}_{m+1}$*.*

No mathematical issues would arise from removing this assumption, but one would then need to find suitable estimators for the joint distribution of the adjudication outcomes instead of for single outcomes.

We also introduce a restricted (non-monotone) information $\mathcal{H}_t \subset \mathcal{H}_t^{\text{obs}}$, which we use as the conditioning information in the distribution of the adjudication outcomes. The point is that it may sometimes be convenient to have the option to use less information than all available information. It is important that the restricted information contains the confirmed jumps so that they are not treated as missing values, but which additional information one lets the adjudication probabilities depend on is optional. As will be seen in Section 5.3, $\mathcal{H}_t$ is needed for all $t \leq \eta$ to estimate the adjudication model, while $\mathcal{H}_\eta$ is sufficient to estimate the reporting delays and biometric hazards. Denote by $\langle t \rangle$ the number reported jumps of $(\tilde{T}_m, \tilde{Y}_m)_{m\geq 1}$ before time $C \wedge t$. We find a natural choice of $\mathcal{H}_t$ to be

$$\begin{aligned}\mathcal{H}_t = \sigma\{\mathbf{X}, \mathcal{G}, C \wedge t, (\tilde{T}_m, \tilde{Y}_m), (\tilde{T}_{\langle t \rangle}, \tilde{Y}_{\langle t \rangle}, \tilde{U}_{\langle t \rangle}, \mathcal{A}_{\langle t \rangle, t})\\ : \tilde{T}_m + \tilde{U}_m \leq t, \xi_m = 1, m < \langle t \rangle\}\end{aligned}$$

consisting of the confirmed jumps by time t as well as the most recent reported jump with its reporting delay and adjudication information; information about prior reporting delays and adjudication processes is removed. We suppose that the distribution of the adjudication outcome $\xi_{\langle t \rangle}$ given $\mathcal{H}_t$ is characterized by a parameter $\mathbf{g}_0$. The parameter set for $(\mathbf{f}, \mathbf{g})$ is assumed to be a product space $\mathbb{F} \times \mathbb{G}$ containing $(\mathbf{f}_0, \mathbf{g}_0)$.

It is worth noting that we place no distributional assumptions on $(\tilde{T}_m, \tilde{Y}_m)_{m \geq 1}$ apart from those inherited from $(T_m^c, Y_m^c)_{m \geq 1}$. Likewise, we do not want to impose criteria for how our process of interest $(T_m^*, Y_m^*)_{m \geq 1}$ would be affected by subsampling on adjudication information or reporting delays. We therefore choose to treat incomplete event adjudication and reporting delays as partially exogenous mechanisms. This naturally leads to a two-step method, where the first step concerns incomplete event adjudication and reporting delays, while the second step involves estimation of $\boldsymbol{\theta}_0$. The method is outlined in Section 5.3, and the main theoretical results, including weak consistency and asymptotic normality, are given in Section 5.4. An alternative approach described in Section C of the Supplementary material incorporates the delays in the second step; this improves efficiency, but comes at the cost of additional independence assumptions.

5.3 Estimation

5.3.1 General considerations

For our two-step method, we propose a two-step M-estimation procedure. Results on the asymptotics of such procedures may be found in Murphy & Topel (1985) and Newey & McFadden (1994) for the parametric case and Ichimura & Lee (2010), Kristensen & Salanié (2017), and Delsol & Van Keilegom (2020) for the semi-parametric case. We treat $\mathbf{f}$ and $\mathbf{g}$ as nuisance parameters and $\boldsymbol{\theta}$ as the parameter of interest. Discussion of concrete estimators is postponed to Section 5.3.2.

For notational simplicity, we henceforward suppress the baseline covariates and use the symbol $\bullet$ to signify summation over the corresponding index. Disregarding for a moment incomplete event adjudication, in the sense that we work under a filtration where the adjudication outcomes for all reported jumps are known, and fixing the parameter of the distribution of the reporting delays at $\mathbf{f}$, we obtain from Assumption 5.2.2 and Section II.7.3 in Andersen et al. (1993) the partial likelihood $L(\boldsymbol{\theta};\mathbf{f})$ for one subject under $\mathcal{F}^c_\eta$ as the following product integral

$$L(\boldsymbol{\theta};\mathbf{f}) = \prod_{t=V}^{C}\prod_{j=1}^{J}\left\{\prod_{k\neq j}\Lambda_{jk}(\mathrm{d}t;\boldsymbol{\theta},\mathbf{f})^{\Delta N_{jk}(t)}\right\}\left\{(1-\Lambda_{j\bullet}(\mathrm{d}t;\boldsymbol{\theta},\mathbf{f}))^{1-\Delta N_{j\bullet}(t)}\right\},$$

where Λ_{jk} is the compensator of N_{jk} i.e. $\Lambda_{jk}(\mathrm{d}t;\boldsymbol{\theta},\mathbf{f}) = \mathbb{E}_{\boldsymbol{\theta},\mathbf{f}}\{N_{jk}(\mathrm{d}t) \mid \mathcal{F}_{t-}\}$. Denote the corresponding log-likelihood by $\ell(\boldsymbol{\theta};\mathbf{f}) = \log L(\boldsymbol{\theta};\mathbf{f})$. We henceforth assume that $\mathbb{E}[\ell(\boldsymbol{\theta};\mathbf{f})]$ is uniquely maximized in $(\boldsymbol{\theta}_0,\mathbf{f}_0)$ which enables the use of M-estimation techniques. This holds under weak assumptions. For example, Section II.7.2 in Andersen et al. (1993) implies, under some integrability and smoothness conditions, that $(\boldsymbol{\theta}_0,\mathbf{f}_0)$ is a unique maximum of $\mathbb{E}[\ell(\boldsymbol{\theta};\mathbf{f})]$ as long as the compensators are non-constant in each coordinate of $(\boldsymbol{\theta},\mathbf{f})$.

Seeking a tractable expression of L in terms of (5.2.1), we introduce the weighted hazard

$$\nu_{jk}\{t;\mathbf{J}(t-),\boldsymbol{\theta},\mathbf{f}\} = \mu^*_{jk}\{t;\mathbf{J}(t-),\boldsymbol{\theta}\}\mathrm{pr}_U\left[\eta-t;\{\mathbf{J}(t-),t,k\},\mathbf{f}\right]$$

with $\{\mathbf{J}(t-),t,k\}$ being the event history containing the jumps $\mathbf{J}(t-)$ and a jump to state k at time t. Let V_t denote the duration since the most recent jump of $(T_m,Y_m)_{m\geq 1}$ at time t, let $I_j(t)$ be the indicator that the last mark in $\mathbf{J}(t-)$ is j or the initial state is j if $\mathbf{J}(t-)$ is empty. Introduce the survival probability in the current state $P^*\{t;\mathbf{J}(t-),\boldsymbol{\theta}\}$ satisfying

$$I_j(t)P^*\{t;\mathbf{J}(t-),\boldsymbol{\theta}\} = I_j(t)\exp\left[-\int_{t-V_{t-}}^{t}\mu^*_{j\bullet}\{s;\mathbf{J}(s-),\boldsymbol{\theta}\}\,\mathrm{d}s\right].$$

In Section B of the Supplementary material, we derive the following result.

Lemma 5.3.1. *It holds that* $\Lambda_{jk}(\mathrm{d}t;\boldsymbol{\theta},\mathbf{f}) = I_j(t)\mu_{jk}\{t;\mathbf{J}(t-),\boldsymbol{\theta},\mathbf{f}\}\,\mathrm{d}t$ *with*

$$\mu_{jk}\{t;\mathbf{J}(t-),\boldsymbol{\theta},\mathbf{f}\} = \gamma_j(t)\times\nu_{jk}\{t;\mathbf{J}(t-),\boldsymbol{\theta},\mathbf{f}\}$$

where $\gamma_j(t)$ *equals*

$$\frac{1-\int_{t-V_{t-}}^{t} P^*\{s;\mathbf{J}(s-),\boldsymbol{\theta}\}\mu_{j\boldsymbol{\cdot}}^*\{s;\mathbf{J}(s-),\boldsymbol{\theta}\}\,\mathrm{d}s}{\mathrm{pr}_U\{\eta-(t-V_{t-});\mathbf{J}(t-V_{t-}),\mathbf{f}\}-\int_{t-V_{t-}}^{t} P^*\{s;\mathbf{J}(s-),\boldsymbol{\theta}\}\nu_{j\boldsymbol{\cdot}}\{s;\mathbf{J}(s-),\boldsymbol{\theta},\mathbf{f}\}\,\mathrm{d}s}.$$

We now return to the problem that $\ell(\boldsymbol{\theta};\mathbf{f})$ is actually not computable using the available information $\mathcal{H}_\eta^{\text{obs}}$ due to incomplete event adjudication. Furthermore, estimation of $\mathbf{f}$ based on $\ell(\boldsymbol{\theta};\mathbf{f})$ is inefficient since it only utilizes $(T_m,Y_m)_{m\geq 1}$ and not $(U_m)_{m\geq 1}$. We therefore propose two-step M-estimation. Denote by $\mathbf{Z}_i^{\text{obs}}$ and $\mathbf{Z}_i$ $(i=1,\ldots,n)$ the subject i outcomes of $\mathcal{H}_\eta^{\text{obs}}$ and $\mathcal{H}_\eta$ respectively, and let $\mathbf{Z}^{\text{obs}}$ and $\mathbf{Z}$ be corresponding generic outcomes. Assume $\mathbf{Z}_i^{\text{obs}}$ are independent and identically distributed, which implies that the same holds for $\mathbf{Z}_i$. We specify the objective function to be the imputed likelihood $\ell(\mathbf{Z},\boldsymbol{\theta};\mathbf{f},\mathbf{g}) = \mathbb{E}_{\mathbf{g}}\{\ell(\boldsymbol{\theta};\mathbf{f}) \mid \mathbf{Z}\}$. For the relation to imputation, see Section A of the Supplementary material. Our two-step procedure is:

1. Estimate $(\mathbf{f}_0,\mathbf{g}_0)$ by $(\hat{\mathbf{f}}_n,\hat{\mathbf{g}}_n)$ using suitable estimators based on $(\mathbf{Z}_i^{\text{obs}})_{i=1}^n$.
2. Estimate $\boldsymbol{\theta}_0$ by $\hat{\boldsymbol{\theta}}_n = \arg\max_{\boldsymbol{\theta}}\sum_{i=1}^n \ell(\mathbf{Z}_i,\boldsymbol{\theta};\hat{\mathbf{f}}_n,\hat{\mathbf{g}}_n)$.

This is motivated by the observation $\mathbb{E}\{\ell(\mathbf{Z},\boldsymbol{\theta};\mathbf{f}_0,\mathbf{g}_0)\} = \mathbb{E}\{\ell(\boldsymbol{\theta};\mathbf{f}_0)\}$, which is uniquely maximized in $\boldsymbol{\theta}=\boldsymbol{\theta}_0$. Recall that $\mathcal{H}_\eta^{\text{obs}}$ and $\mathcal{H}_\eta$ are such that only $\xi_{\langle\eta\rangle}$ can be missing in $\ell(\boldsymbol{\theta};\mathbf{f})$, and that we wish to use adjudication probabilities based on the information $\mathcal{H}_\eta$ instead of $\mathcal{H}_\eta^{\text{obs}}$, leading us to condition on $\mathbf{Z}$ instead of $\mathbf{Z}^{\text{obs}}$ in the objective function. We therefore let $w(1,\mathbf{Z};\mathbf{g}) = \mathbb{E}_{\mathbf{g}}\{\xi_{\langle\eta\rangle} \mid \mathbf{Z}\}$ and $w(0,\mathbf{Z};\mathbf{g}) = 1-w(1,\mathbf{Z};\mathbf{g})$. Write $\ell(\boldsymbol{\theta};\mathbf{f}) = \ell(\boldsymbol{\theta};\mathbf{f},\xi_{\langle\eta\rangle})$, where the last argument signifies whether the jump $(\tilde{T}_{\langle\eta\rangle},\tilde{Y}_{\langle\eta\rangle})$ is included in the likelihood or not. We thus have

$$\ell(\mathbf{Z},\boldsymbol{\theta};\mathbf{f},\mathbf{g}) = w(0,\mathbf{Z};\mathbf{g})\ell(\boldsymbol{\theta};\mathbf{f},0) + w(1,\mathbf{Z};\mathbf{g})\ell(\boldsymbol{\theta};\mathbf{f},1). \tag{5.3.1}$$

In Section 5.3.2, we outline the first step, while Section 5.3.3 is dedicated to the second. The derivation of asymptotic properties is postponed to Section 5.4.

5.3.2 Estimation of f and g

Assumption 5.2.1 implies that $\mathbb{P}_{\mathbf{f}}\{U_m\leq t \mid \mathbf{J}(T_m)\} = \mathrm{pr}_U\{t;\mathbf{J}(T_m),\mathbf{f}\}/\mathrm{pr}_U\{\eta - T_m;\mathbf{J}(T_m),\mathbf{f}\}$ on the event $(T_m<\infty)$ for $t\leq\eta-T_m$. Note that one is not conditioning on $\mathbf{J}(T_m)$ on the right-hand side but rather inputting the values into the regular conditional probability pr_U. Thus, the parameter $\mathbf{f}$ could be estimated if $(T_m, Y_m)_{m\geq 1}$ rather than $(\tilde{T}_m, \tilde{Y}_m)_{m\geq 1}$ were available. We therefore also employ two-step M-estimation for the nuisance parameters, first constructing $\hat{\mathbf{g}}_n$ and subsequently using an imputed likelihood to construct $\hat{\mathbf{f}}_n$. Practical aspects of the implementation are discussed in Section D of the Supplementary material.

We first consider estimation of $\mathbf{g}$. Since there is time dependence in adjudication processing, and since we see no need for infinite-variation processes, we choose to model the adjudication information $\mathcal{A}_{m,t}$ of the m'th reported jump as being generated by a marked point process $(\tau_{m,\ell}, Y_{m,\ell})_{\ell\geq 1}$ with finite state space $1, \dots, K$. The time σ_m is specified as a first hitting time of a subset of the states $1, \dots, K$, and the value of ξ_m implied by hitting one of these states is non-random. The corresponding counting processes are denoted $N_{m,jk}$ and the indicator that the process is in state j at time $t-$ is denoted $I_{m,j}(t)$. This multistate model starts when the corresponding jump is reported and ends when the next jump is reported, meaning $\tilde{T}_m + \tilde{U}_m < \tau_{m,\ell} \leq \tilde{T}_{m+1} + \tilde{U}_{m+1}$ for all ℓ. We choose a log-likelihood type objective function for $\mathbf{g}$ and parameterize via conditional hazards $\omega_{jk}(t; \mathcal{H}_{t-}, \mathbf{g})$ given $\mathcal{H}_{t-}$, that is

$$\begin{aligned}\ell_\xi(\mathbf{Z}^{\text{obs}}, \mathbf{g}) = \sum_{m=1}^{\langle \eta \rangle} \sum_{j,k=1}^{K} \int_0^\eta &\log \omega_{jk}(t; \mathcal{H}_{t-}, \mathbf{g})\, N_{m,jk}(\mathrm{d}t) \\ &- \int_0^\eta I_{m,j}(t)\omega_{jk}(t; \mathcal{H}_{t-}, \mathbf{g})\, \mathrm{d}t.\end{aligned}$$

Note that the log-likelihood depends on $\mathcal{H}_t$ for all $t \leq \eta$; this information is contained in $\mathcal{H}_\eta^{\text{obs}}$ but not in $\mathcal{H}_\eta$. Estimation now proceeds by maximizing $\sum_{i=1}^n \ell_\xi(\mathbf{Z}_i^{\text{obs}}, \mathbf{g})$ with respect to $\mathbf{g}$ over a suitable parametric class. This maximization is standard. Other modeling choices for $\mathbf{g}$ are possible, and while our proofs of the asymptotic properties utilize this specific choice, they can easily be adapted to a wider range of models.

We now consider estimation of $\mathbf{f}$. Disregarding incomplete event adjudication for a brief moment, Assumption 5.2.2 implies that the log-likelihood for the reporting delays becomes

$$\ell_U(\mathbf{f}) = \sum_{m=1}^{N^c_{\cdot\cdot}(\eta)} \log\left[\mathrm{pr}_U\{\mathrm{d}U_m^c; \mathbf{J}^c(T_m^c), \mathbf{f}\}/\mathrm{pr}_U\{\eta - T_m^c; \mathbf{J}^c(T_m^c), \mathbf{f}\}\right].$$

Define the reverse time hazard

$$\alpha\{s; \mathbf{J}^c(T_m^c), \mathbf{f}\}\,\mathrm{d}s = \mathrm{pr}_U\{\mathrm{d}s; \mathbf{J}^c(T_m^c), \mathbf{f}\}/\mathrm{pr}_U\{s; \mathbf{J}^c(T_m^c), \mathbf{f}\}$$

similarly to Kalbfleisch & Lawless (1991) and write the distribution function in terms of the reverse time hazard $\mathrm{pr}_U\{t; \mathbf{J}^c(T_m^c), \mathbf{f}\} = \exp\left[-\int_t^\infty \alpha\{s; \mathbf{J}^c(T_m^c), \mathbf{f}\}\,\mathrm{d}s\right]$.

Then

$$\ell_U(\mathbf{f}) = \sum_{m=1}^{N_{\cdot\cdot}^c(\eta)} \log \alpha\{U_m^c; \mathbf{J}^c(T_m^c), \mathbf{f}\} - \int_{U_m^c}^{\eta - T_m^c} \alpha\{s; \mathbf{J}^c(T_m^c), \mathbf{f}\}\, \mathrm{d}s.$$

Parameterizing in terms of reverse time hazards maintains the form of the likelihood under right-truncation similarly to how hazards maintain the form of the likelihood under right-censoring. Returning to incomplete event adjudication, we analogously to $\ell(\boldsymbol{\theta}; \mathbf{f}, \xi_{\langle\eta\rangle})$ write $\ell_U(\mathbf{f}) = \ell_U(\mathbf{f}; \xi_{\langle\eta\rangle})$ and let the objective function for $\mathbf{f}$ be the imputed likelihood $\ell_U(\mathbf{Z}, \mathbf{f}; \mathbf{g}) = \mathbb{E}_{\mathbf{g}}\{\ell_U(\mathbf{f}) \mid \mathbf{Z}\} = w(0, \mathbf{Z}; \mathbf{g})\ell_U(\mathbf{f}; 0) + w(1, \mathbf{Z}; \mathbf{g})\ell_U(\mathbf{f}; 1)$. For a given estimator $\hat{\mathbf{g}}_n$, we let $\hat{\mathbf{f}}_n$ be the maximizer of $\sum_{i=1}^n \ell_U(\mathbf{Z}_i, \mathbf{f}; \hat{\mathbf{g}}_n)$ with respect to $\mathbf{f}$ over a suitable parametric class.

5.3.3 Estimation of $\boldsymbol{\theta}$

The estimator $\hat{\boldsymbol{\theta}}_n$ was described in Section 5.3.1. Evaluating $\ell(\boldsymbol{\theta}; \mathbf{f})$ is costly due to the repeated numerical integration required owing to γ_j. This would be further compounded when evaluating the score, and we hence recommend derivative-free maximization for example via quasi-Newton methods. For these algorithms to perform well, one needs a good starting value for $\boldsymbol{\theta}$. We here describe one approach to overcoming this problem.

Define $\check{\gamma}_j$ by replacing $\mathrm{pr}_U\{\eta - (t - V_{t-}); \mathbf{J}(t - V_{t-}), \mathbf{f}\}$ with one in the expression for γ_j, effectively ignoring the reporting delay of the previous jump, and let $\check{\mu}_{jk} = \check{\gamma}_j \times \nu_{jk}$. Note that the approximation is exact if there is never a reporting delay for the prior jump or if there has been no prior jump. Note $I_j(t)P^*\{t; \mathbf{J}(t-), \boldsymbol{\theta}\} \leq \check{\gamma}_j(t) \leq 1$ so $\check{\mu}_{jk} \approx \nu_{jk}$. Furthermore, under additional smoothness, we show in Section B of the Supplementary material that $\check{\mu}_{jk}\{t; \mathbf{J}(t-), \boldsymbol{\theta}, \mathbf{f}\} = \nu_{jk}\{t; \mathbf{J}(t-), \boldsymbol{\theta}, \mathbf{f}\} + \mathrm{err}(t)$ with $|\mathrm{err}(t)| \leq V_{t-} \times \sup_{s\in[t-V_{t-}, t]} \mu_{j\bullet}^*\{s; \mathbf{J}(s-), \boldsymbol{\theta}\}^2$, so the error is a second order term. We call the approximation $\mu_{jk} \approx \nu_{jk}$ a Poisson approximation, since the approximate hazard ν_{jk} has a similar form to the one found in the Poisson process setup of Kalbfleisch & Lawless (1989). The corresponding approximate likelihood $\ell^{\mathrm{app}}(\boldsymbol{\theta}; \mathbf{f})$ reads

$$\begin{aligned}\ell^{\mathrm{app}}(\boldsymbol{\theta}; \mathbf{f}) = \sum_{j,k=1}^{K} & \int_V^C \log \mu_{jk}^*\{t; \mathbf{J}(t-), \boldsymbol{\theta}\}\, \mathrm{d}N_{jk}(t) \\ & - \int_V^C I_j(t)\mu_{jk}^*\{t; \mathbf{J}(t-), \boldsymbol{\theta}\}\mathrm{pr}_U[\eta - t; \{\mathbf{J}(t-), t, k\}, \mathbf{f}]\, \mathrm{d}t.\end{aligned}$$

Analogously to Equation (5.3.1), we introduce the imputed version $\ell^{\mathrm{app}}(\mathbf{Z}, \boldsymbol{\theta}; \hat{\mathbf{f}}_n, \hat{\mathbf{g}}_n)$. Maximization of the approximate likelihood is significantly simpler than the true likelihood; computationally, the approximate case is the same as a standard multi-state likelihood, but where the contribution of a given path is weighted with a scalar and where the exposure is reduced based on the closeness to the time of analysis via the reporting delay distribution. Practical aspects of the implementation are discussed in Section D of the Supplementary material.

Remark 5.3.2. The framework could be extended to allow for time-dependent covariates by using the associated partial likelihood. Note, however, that the covariate paths then have to be included in $\mathcal{H}$, meaning that the adjudication model also needs to depend on these covariates. Since the adjudication model has to be used prospectively, one would need a model for how the covariates evolve over time. This is about as complicated as including the covariates in the marked point process itself, which is also more natural when the goal is prediction rather than inference.

5.4 Asymptotic Properties

The large sample properties of $\hat{\boldsymbol{\theta}}_n$ largely follow from the two-step M- and Z-estimation results of Newey & McFadden (1994), Hahn (1996), and Delsol & Van Keilegom (2020). We also show weak consistency and asymptotic normality of a bootstrap procedure analogous to Efron's simple nonparametric bootstrap as studied in Gross & Lai (1996), where $(\mathbf{Z}_{ni}^{\text{obs}})_{i=1}^n$ are drawn with replacement from $(\mathbf{Z}_i^{\text{obs}})_{i=1}^n$. The estimators based on the bootstrap sample are denoted $\hat{\mathbf{g}}_n^{\text{boot}}$, $\hat{\mathbf{f}}_n^{\text{boot}}$, and $\hat{\boldsymbol{\theta}}_n^{\text{boot}}$. The main result is:

Theorem 5.4.1.

(i) Consistency: Under Assumptions 1-3 and Assumptions 5-7 from Section E of the Supplementary material, $\hat{\mathbf{g}}_n$, $\hat{\mathbf{f}}_n$, and $\hat{\boldsymbol{\theta}}_n$ as well as $\hat{\mathbf{g}}_n^{\text{boot}}$, $\hat{\mathbf{f}}_n^{\text{boot}}$, and $\hat{\boldsymbol{\theta}}_n^{\text{boot}}$ are weakly consistent.

(ii) Asymptotic normality: Under Assumptions 1-3 and Assumptions 5-8 from Section E of the Supplementary material, $\hat{\mathbf{g}}_n$, $\hat{\mathbf{f}}_n$, and $\hat{\boldsymbol{\theta}}_n$ are asymptotically normal. Furthermore $n^{1/2}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) \to N(0, \mathbf{V})$ and $n^{1/2}(\hat{\boldsymbol{\theta}}_n^{\text{boot}} - \hat{\boldsymbol{\theta}}_n) \to N(0, \mathbf{V})$ in distribution.

The proof and the form of $\mathbf{V}$ are given in Section E of the Supplementary material. Implementing confidence intervals via estimators for $\mathbf{V}$ is challenging since there is no closed form for the score. However, Theorem 5.4.1 implies that a percentile confidence interval based on a simple bootstrap procedure is valid. One may also use an empirical bootstrap variance estimator for $\mathbf{V}$ if $n^{1/2}(\hat{\boldsymbol{\theta}}_n^{\text{boot}} - \hat{\boldsymbol{\theta}}_n)$ is uniformly integrable.

Remark 5.4.2. As may be seen from the proof of Theorem 5.4.1, the proposed estimator can, under smoothness conditions, be cast as a generalized method of moment estimator with weighting matrix equal to the identity by using score equations as moment conditions. In the class of regular asymptotically linear estimators that only restrict the statistical model through moment conditions, it is known from Hansen (1982), Chamberlain (1987), and Section 5 of Newey & McFadden (1994) that generalized method of moment estimators satisfying Equation (5.4) of Newey & McFadden (1994) are efficient. A sufficient condition for this is that the weighting matrix is equal to the inverse of the variance of the score, while for the maximum likelihood estimator, the identity weight matrix is also efficient due to Bartlett's identities. Despite using score-type equations as moment conditions, our proposed estimator is not efficient, since an identity of Barlett type is not generally satisfied. Hwang & Sun (2018), however, find that the finite sample performance of the efficient estimator is often inferior to an a priori chosen weighting matrix, since the uncertainty in estimating the weighting matrix is ignored in the asymptotics. If the cost-benefit trade-off is unclear, they recommend sticking to an a priori chosen weighting matrix, in line with our proposal of using $\hat{\boldsymbol{\theta}}_n$. Compared to the efficient estimator, our proposal also has the computational advantage that the score need not be computed.

5.5 Numerical Study

5.5.1 Data-generating process

We use the multistate event model and three stage adjudication model depicted in Figure 5.1. A jump is confirmed if and when state 3 of the adjudication model is hit. Multistate models are simulated using Lewis' thinning algorithm from Ogata (1981). We refer to Section F of the Supplementary material for additional figures, tables, and details regarding the numerical study. We compare our approach with oracle methods that use the filtration $\mathcal{F}^*$, which is unavailable to the statistician, and naive methods that either use all reported events without back-censoring (Naive 1) or back-censor by one year and delete reported events if they have been under adjudication for over two years at time η (Naive 2).

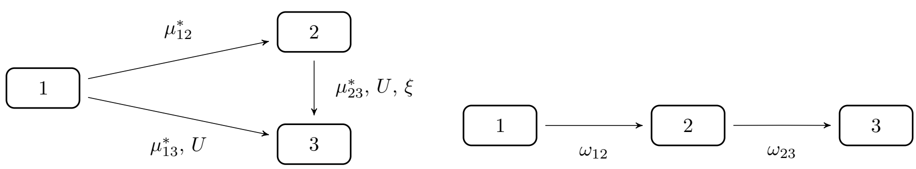

Figure 5.1: *Event history model (left) and adjudication model (right). Symbols U and ξ indicate the presence of reporting delays and adjudication processes, respectively.*

A total of 400 samples of size $n = 1500$ with time-horizon $\eta = 5$ years are considered. For a given subject, the data is generated as follows: $V \sim \text{Uniform}(0, 1)$ and $C \mid V \sim \text{Uniform}(V, \eta)$. Subjects enter in state 1 with no additional left-truncation information. The baseline covariate is $X \sim \text{Uniform}(-4, 4)$. We let $\boldsymbol{\theta} = (\log 0.15, 0.1, 0.4, \log 0.1, 0.03, -0.3, -0.3)$ and generate events using the rates

$\tilde{\mu}_{12}(t;X) = \exp\{\theta_1 + \theta_2 \times (t+X) + \theta_3 \times \sin(0.5\pi X)\}$, $\tilde{\mu}_{13}(t;X) = \exp\{\theta_4 + \theta_5 \times t^2 + \theta_6 \times \cos(0.5\pi X)\}$, and $\tilde{\mu}_{23}(t;\tilde{T}_1, X) = \exp\{\theta_7 \times (t - \tilde{T}_1)X^2\}$, which upon deletion of unreported jumps gives $(\tilde{T}_m, \tilde{Y}_m)_{m\geq 1}$ or upon deletion of rejected jumps gives $(T^*_m, Y^*_m)_{m\geq 1}$. The rates are relatively large, and on average around 260 transitions from 1 to 3, 415 transitions from 1 to 2, and 180 transitions from 2 to 3 are generated.

To model reporting delays, impose a Weibull distribution with proportional reverse time hazards, that is $\alpha(t;X) = \alpha_0(t)\exp(X\beta)$ for $\alpha_0(t) = \frac{k\lambda^k t^{k-1}}{\exp\{(\lambda t)^k\}-1}$. This results in a power model for the distribution function $t \mapsto [1 - \exp\{-(\lambda t)^k\}]^{\exp(X\beta)}$ which we denote Weibull(λ, k, β). Let $\mathbf{f} = (2, 0.5, 0.1, 1, 1.5, 0.2)$. Reporting delays are distributed as Weibull(f_1, f_2, f_3) when coming from state 1 and as Weibull(f_4, f_5, f_6) when coming from state 2. The generated reporting delays have means close to one and standard deviations around 0.15 and 0.05, respectively. The adjudication rates are $\omega_{12}(t;\mathcal{H}_{t-}) = g_1 \times \{X/(t - \tilde{T}_2 - \tilde{U}_2 + 2)\}^2$ and $\omega_{23}(t;\mathcal{H}_{t-}) = \exp\{g_2 \times (t - \tau_{1,1})\}$ with $\mathbf{g} = (0.8, -1.2)$, leading to a long-term confirmation rate of 37%. To go from $\tilde{\mu}_{jk}$ to μ^*_{jk}, note that $\mu^*_{12}(t;X) = \tilde{\mu}_{12}(t;X)$ and $\mu^*_{13}(t;X) = \tilde{\mu}_{13}(t;X)$, while an expression for $\mu^*_{23}(t, V_t; X)$ may be found in Section F of the Supplementary material.

5.5.2 Finite sample performance

The regularity assumptions from Theorem 5.4.1 are easily seen to be satisfied, and we hence expect the proposed estimators to be unbiased and asymptotically normal. In Table 5.1 we report the bias and the empirical standard deviation (SD) of the parameter estimators of $\boldsymbol{\theta}$ using the different methods. A corresponding table for the estimators $\hat{\mathbf{g}}_n$ and $\hat{\mathbf{f}}_n$ can be found in Section F of the Supplementary material. The results are generally as expected. We find comparable bias for the proposed method and oracle methods, with a higher SD for the proposed method. The bias is slightly higher for the Poisson approximation with SD comparable to the proposed method. The naive methods are slightly worse than the proposed method for $(\theta_1, \theta_2, \theta_3)$, moderately worse for $(\theta_4, \theta_5, \theta_6)$, and substantially worse for θ_7. The naive methods generally have comparable performance except that Naive 1 is superior for θ_4 and Naive 2 is superior for θ_7. We conclude that our proposed method works well.

The Poisson approximation is also expected to perform well since there is no jump after the transition affected by reporting delays, confer with Section 5.3.3. Since the hazards are relatively large, the decrease in bias when going from the Poisson approximation to our proposed method is non-negligible. Depending on the application, this reduction in bias might be worth the significantly longer computation times. With our implementation and hardware setup, the Poisson approximation for θ_ℓ ($\ell = 1, \ldots, 6$) took only 7 seconds, while the proposed method took about 300 seconds.

We also seek to verify the bootstrap procedure outlined in Section 5.4. We here focus on θ_7 since its estimation is affected by both reporting delays and incomplete event adjudication. Based on $k = 400$ estimates of θ_7 with 1000 bootstrap resamples and confidence level $1-\alpha = 0.90, 0.95, 0.99$, we obtain coverage rates of $89.5, 95.0, 98.7$. We conclude that the bootstrap works well as the coverage rates are close to the nominal level.

Table 5.1: *Bias and empirical standard deviation (SD) of the estimator $\hat{\boldsymbol{\theta}}_n$ based on 400 simulations of size $n = 1500$.*

	Proposed method		Oracle		Poisson approx.		Naive 1		Naive 2	
Parameter	Bias	SD	Bias	SD	Bias	SD	Bias	SD	Bias	SD
$\theta_1 = \log 0.15$	-.004	.067	-.008	.031	-.010	.067	-.011	.066	-.010	.066
$\theta_2 = 0.1$	-.000	.020	-.001	.020	-.006	.020	-.006	.020	-.006	.021
$\theta_3 = 0.4$	.003	.078	.003	.078	-.002	.078	-.002	.078	-.000	.079
$\theta_4 = \log 0.1$	.003	.084	.001	.083	.012	.091	-.041	.077	-.051	.082
$\theta_5 = 0.03$	.000	.012	-.000	.013	-.006	.016	-.018	.011	-.015	.014
$\theta_6 = -0.3$	-.000	.094	-.001	.088	.007	.094	-.009	.087	-.007	.090
$\theta_7 = -0.3$	-.011	.066	-.011	.054	-.012	.066	.157	.023	.148	.069

5.5.3 Robustness against misspecification

To investigate the performance of the proposed method under misspecification, we use the same simulations but consider parameters estimated under a misspecified parametric family for $(T_m^*, Y_m^*)_{m\geq 1}$ characterized by the hazard rates

$$\begin{aligned}
\mu_{12}^{\text{miss}}(t; X) &= \exp\{\theta_1 + \theta_2 \times t + \theta_3 \times X\}, \\
\mu_{13}^{\text{miss}}(t; X) &= \exp\{\theta_4 + \theta_5 \times t + \theta_6 \times X\}, \\
\mu_{23}^{\text{miss}}(t; \tilde{T}_1, X) &= \exp\{\theta_7 + \theta_8 \times (t - \tilde{T}_1) + \theta_9 \times X\}.
\end{aligned}$$

This misspecification is substantial since it disregards the non-linearities and interactions present in the true hazards. We both consider the case where the reporting delay and adjudication models are correctly specified, and a doubly misspecified case where these are also misspecified via the following parametric families:

$$\begin{aligned}
\omega_{12}^{\text{miss}}(t; \mathcal{H}_{t-}) &= \exp\{g_1 \times X + g_2 \times (t - \tilde{T}_2 - \tilde{U}_2)\}, \\
\omega_{23}^{\text{miss}}(t; \mathcal{H}_{t-}) &= \exp\{g_3 \times X + g_4 \times (t - \tau_{1,1})\},
\end{aligned}$$

and $\alpha(t; X) = \alpha_0(t)\exp(X\beta)$ with α_0 being the reverse time hazard for a Gamma distribution with shape parameter k and rate parameter λ.

Since it is no longer meaningful to compare parameter values, we instead compare the performance of the different estimation approaches for the expected duration

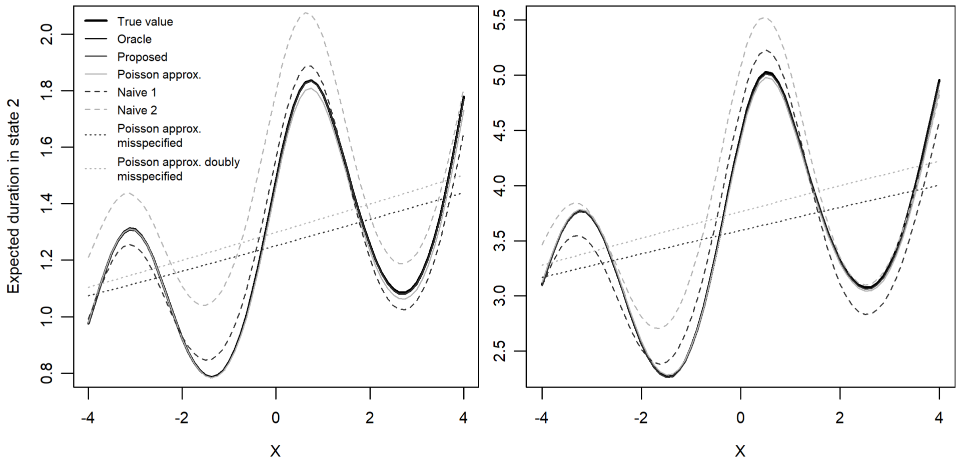

Figure 5.2: *Plot of plug-in estimates of $V_a(t;x)$ averaged over 400 samples of size $n = 1500$ with $t = \eta$ on the left and $t = 2\eta$ on the right.*

spent in state 2 before time t, which is denoted $V_a(t;X) = \mathbb{E}\{\int_{[0,t]} I_2^*(s)\,\mathrm{d}s \mid X\}$, with $t \in \{\eta, 2\eta\}$. In actuarial contexts, this may be recognized as an active reserve for a unit disability annuity. Note that the case $t = 2\eta$ requires extrapolation which is not inherently an issue for parametric models. The estimand V_a is calculated by solving the differential equation from Corollary 7.2 in Adékambi & Christiansen (2017) using the fourth-order Runge-Kutta method. Since it was observed that the Poisson approximation performed very well for this estimand, as illustrated in Figure 5.2, we have chosen to only consider the misspecified model used in conjunction with the Poisson approximation. The results are illustrated in Figure 5.2.

The oracle, proposed, and Poisson approximation approaches are all very close to the true value of V_a, with the Poisson approximation showing slight deviations around $X = 0$ and $X = 2$. Naive 1 performs reasonably albeit with some noticeable bias, while Naive 2 seems to be systematically biased upwards except for $t = 2\eta$ and $X \geq 2$ where it is very close to the true value. The misspecified parametric model captures the local x-dependence of $V_a(t;x)$ the worst, but identifies the correct overall level for the estimand. This behavior is similar to how maximum likelihood estimators behave under misspecification, see Halbert (1982). Including polynomial or sinusoidal transformations of X in the parametric model would likely have led to a better fit for the local behavior of V_a, but including too many non-linear terms could lead to high variance estimators due to overfitting. The performance is comparable for $t = \eta$ and $t = 2\eta$, showing that our method is able to capture the general trends reasonably well even under substantial model misspecification.

5.6 Data Application

5.6.1 Data characteristics

We introduce a new data set LEC-DK19 (Loss of Earning Capacity – Denmark 2019) collected by a large Danish insurer in the period 31/01/2015 to 01/09/2019 with time of analysis $\eta = 01/09/2019$ and apply our proposed estimation procedure to this data. Observations are available on a monthly grid. Subjects enter into the data when they enter the portfolio and are censored when they leave the portfolio. Claims may be reported after censoring, meaning that information on reporting delays and adjudication may arrive after censoring. A total of 416,483 insured are included across five tables concerning disabilities, reactivations, disability delays, disability adjudications, and reactivation adjudications. Of the 413,139 insured in the disability data, 1,773 (0.43%) have an unrefuted disability; disabilities are thus rare. Of the 3,011 insured in the reactivation data, 1,133 (37.63%) have an unrefuted reactivation. The data is based on raw data that has been anonymized and slightly altered so as not to reveal any confidential information about the individual subjects or the insurance portfolio. Available baseline covariates are gender and date of birth.

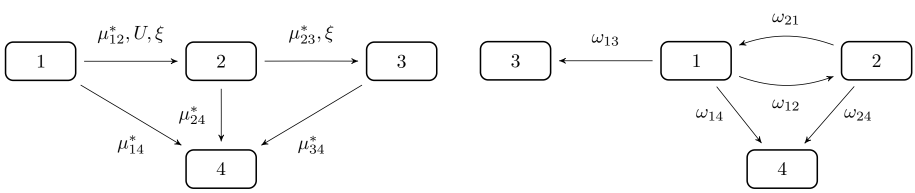

Figure 5.3: *Event history model (left) and adjudication model (right). For events, active is 1, disabled is 2, reactivated is 3, and dead is 4. For adjudications, active report is 1, inactive report is 2, adjudicated is 3, and dead is 4.*

5.6.2 Model specification

We model the data with the multistate model depicted in Figure 5.3. Subjects can become disabled, reactivate from disablement, and die. A disability starts in the adjudication state 1, is confirmed if state 3 is hit, and otherwise rejected. A reactivation starts in adjudication state 2 and the reactivation is annulled if state 3 is hit. Only disablements seem to exhibit reporting delay, likely because the reactivations result from the insurer terminating payments.

We employ the Poisson approximation, which implies that the hazards can be estimated separately. We here only estimate μ^*_{12} and μ^*_{23}, partly because these are of primary interest, but also because only deaths recorded during the adjudication period are included in the data, impeding estimation of the full event history model.

Via back-censoring, it is seen that the hazard μ^*_{12} likely has order of magnitude 10^{-2}, which is one to two orders of magnitudes less than the hazards in the numerical study of Section 5.5, so the incurred approximation error for μ^*_{12} is expected to be negligible. The lack of mortality data makes it difficult to assess the approximation error for μ^*_{23}, however an alternative is to approximate the disabled mortality hazard with zero.

The hazards are regressed on the baseline covariates and the time elapsed since 31/01/2015, and μ^*_{23} is additionally regressed on the duration spent in the disabled state. An adjudication hazard is estimated for each transition. Following the form of $\mathcal{H}_t$ suggested in Section 5.2, the adjudication hazards for disabilities are regressed on the duration since the disability occurred and the associated reporting delay, and the adjudication hazards for reactivations are regressed on the time since the disability and reactivation events. In addition, for disabilities we regress on whether state 2 has previously been hit. For each adjudication and event hazard, the covariates enter in a linear predictor with log link. The log link implies that the hazards are always positive, so a disability rejection or reactivation confirmation is only certain once the insured dies, but the probability of a confirmed jump decreases towards zero as the duration of adjudication increases towards infinity. For reporting delays, we employ the Weibull proportional reverse time hazards of Section 5.5 with age entering as age at disability onset. Since the data set contains monthly records, we employ the Poisson regression approach for implementation as described in Section D of the Supplementary material.

5.6.3 Empirical results

In Table 5.2, we present the non-nuisance parameter estimates and percentile bootstrap confidence intervals computed using 400 bootstrap samples. For comparison, we additionally employ a naive estimation procedure consisting of using all reported unrefuted events with no back-censoring. A corresponding table for the nuisance parameters is provided in Section G of the Supplementary material. For this data, we see that adjusting for reporting delays and incomplete event adjudication has a substantial effect on the estimated level and calendar-time dependence of the disability hazard. The effects on the parameter estimates of the reactivation hazard are also noticeable, but contrary to the case for the disability hazard, the naive estimates stay firmly within the confidence intervals except for the calendar-time dependence. In Figure 5.4, predicted rates across the data set, using both the proposed and naive procedure, are shown against the corresponding empirical rates aggregated by yearly tenths. Since the naive empirical rates use all unrefuted events, they are higher than the adjusted empirical rates for reactivations. For disabilities, they are higher in the first years and lower in the later years, where the influence of reporting delays is more substantial. Both rates show large calendar-time dependence, which is a distinct feature of the data set that implies an especially poor forecasting performance of back-censoring and underscores the importance of our proposed methods for disability insurance applications.

Table 5.2: *Parameter estimates (Est.), Naive parameter estimates (Naive), and 95% bootstrap percentile confidence interval (CI) for* $\boldsymbol{\theta}$ *using the proposed method and 400 bootstrap resamples.*

	μ^*_{12}			μ^*_{23}		
Parameter	Est.	Naive	CI	Est.	Naive	CI
Age	.023	.022	(0.018, 0.027)	-.012	-.012	(-0.018, -0.007)
Male	-7.46	-7.11	(-7.65, -7.23)	.334	.262	(0.073, 0.624)
Female	-8.66	-8.25	(-8.87, -8.40)	.756	.674	(0.456, 1.07)
Time	.304	.098	(0.188, 0.357)	-.125	-.068	(-0.168, -0.073)
Duration	-	-	-	-1.04	-1.03	(-1.15, -0.942)

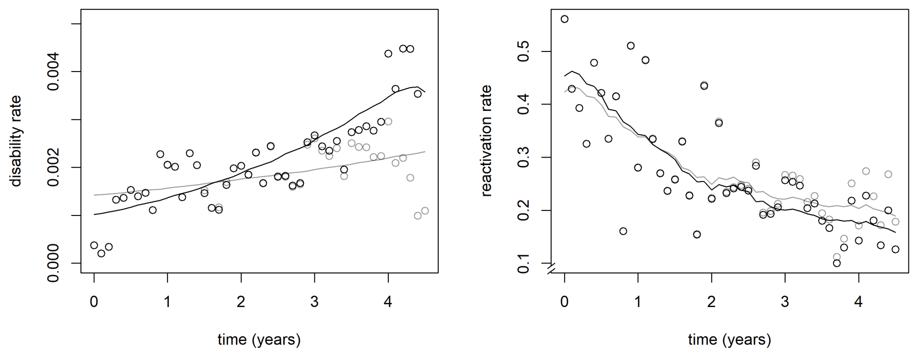

Figure 5.4: *Fitted rates (lines) and occurrence-exposure rates (points) for the proposed method (black) and the naive method (gray). Disability rates are shown on the left and reactivation rates on the right.*

The alternative approach, where the Poisson approximation is not employed but the disabled mortality is approximated with zero, leads to the estimated regressors of μ^*_{23} being $(-0.011, 0.373, 0.840, -0.260, -0.907)$. Compared with the regressors obtained when using the Poisson approximation, the main difference is that a portion of the duration dependence is shifted to calendar-time dependence. This would have a noticeable influence on the forecasted reactivation rates. For real applications, we therefore recommend using the full estimation procedure with access to mortality data when estimating the reactivation hazard. We also recommend validating the models by comparing the original predictions with predictions obtained using the same data but with different time of analyses. However, this is not possible with the current data, as the disability and reactivation occurrences, along with exposures, are available only for a single valuation date. Moreover, their values at different valuation dates cannot be inferred from the adjudication data, since the same ID does not correspond to the same insured individual across the separate data tables. Note also that usual out-of-sample validation is not applicable since the data itself is biased.

In Figure 5.5, we use the link between reverse time hazard estimation and Poisson regression, which is described in Section D of the Supplementary material, to plot the fitted reverse time hazard rates against empirical occurrence-exposure rates. The model shows no obvious lack of goodness-of-fit. Figure 5.6 contains similar plots for the adjudication hazards as a function of the duration since the unadjudicated event occurred and which also do not show no systematic deviations.

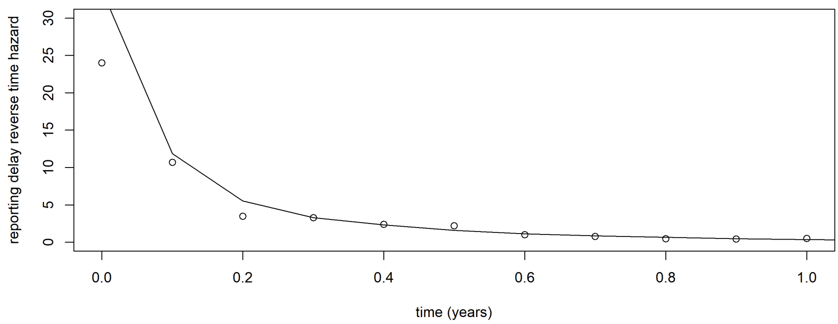

Figure 5.5: *Fitted rate (line) and occurrence-exposure rate (points) for the reverse time hazard of the reporting delay distribution.*

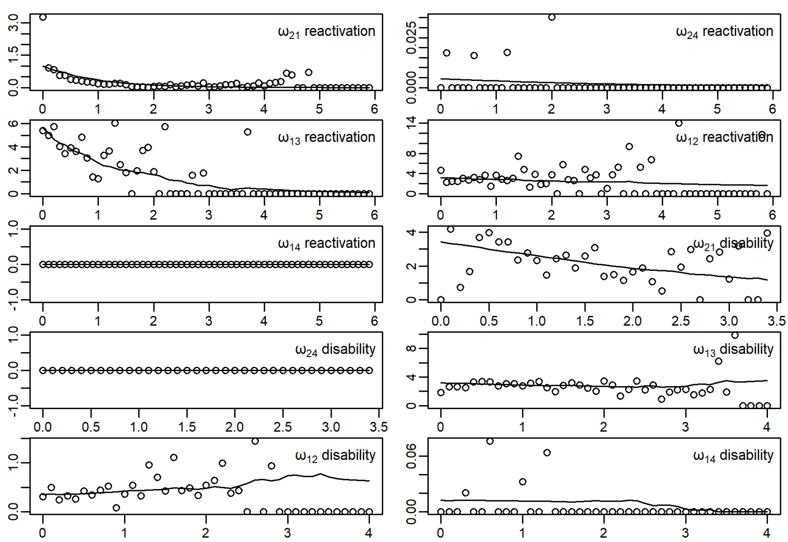

Figure 5.6: *Fitted rates (lines) and occurrence-exposure rates (points) for the adjudication hazards graphed as a function of the duration since the unadjudicated event occurred.*

5.7 Closing Remarks

We have proposed a parametric two-step method to estimate hazards when data is left-truncated, right-censored, and contaminated by reporting delays and incomplete event adjudication. We proved that the estimators and their bootstrapped versions are weakly consistent and asymptotically normal. The numerical study showed favorable performance of the proposed method compared to other alternatives, but also highlighted that for small hazards the Poisson approximation performs reasonably while being considerably less computationally demanding. Our approach overcomes the need for back-censoring, meaning that the most recent data may be used. This feature is particularly useful in monitoring the emergence of new trends on a population level, confer with the data application in Section 5.6.

The Supplementary material contains Section A–G which includes Tables and Figures referenced in Sections 5.2-5.6. The R code and the LEC-DK19 data are available on GitHub: `https://github.com/oliversandqvist/Web-appendix-estimation-contamination`.

Acknowledgments and declarations of interest

Oliver Lunding Sandqvist is an Industrial PhD student at PFA Pension and Kristian Buchardt is an Affiliated Professor of Actuarial Mathematics at the University of Copenhagen. Significant parts of the research presented in this paper were carried out while Kristian Buchardt was employed at PFA Pension. The authors thank Munir Hiabu for insightful comments that helped improve the manuscript. Oliver Lunding Sandqvist's research has partly been funded by the Innovation Fund Denmark (IFD) under File No. 1044-00144B.

Supplementary material

The supplementary material is organized as follows. Section 5.A provides a discussion of the link between the proposed estimation procedure and imputation. In Section 5.B, we prove Lemma 5.3.1 and derive an error-bound for the Poisson approximation. In Section 5.C, we discuss the estimation procedure and its efficiency under endogenous reporting delays. Section 5.D provides a way to implement parts of the proposed estimation procedure via Poisson regression and explains how this link may be used to construct goodness-of-fit plots. In Section 5.E, we derive the large-sample behavior of our estimators and prove Theorem 5.4.1. Section 5.F contains further details regarding the data-generating mechanism used in the numerical study as well as additional figures and tables. Section 5.G contains additional figures and tables for the data application.

5.A Relation to Imputation

The proposed estimation procedure corresponds to frequentist imputation when the number of imputations b tends to infinity. With finite b, one may form a single 'clustered data' objective function $\frac{1}{b}\sum_{j=1}^{b}\frac{1}{n}\sum_{i=1}^{n}\ell(\boldsymbol{\theta};\mathbf{f},k_{ij})$, where k_{ij} is the imputed value of the adjudication outcome for individual i in the imputed dataset j. This procedure is described and analyzed in Wang & Robins (1998), Robins & Wang (2000), and Kim (2011) for both finite b and in the limit $b\to\infty$. An approximate maximizer of $\sum_{i=1}^{n}\ell(\mathbf{Z}_i,\boldsymbol{\theta};\mathbf{f},\mathbf{g})$ could thus alternatively have been based on a large but finite number of imputed datasets (simulating k_{ij} based on its distribution under $\mathbf{g}$ conditional on $\mathbf{Z}_i$), where the user inputs the $b\times n$ observations as 'independent' observations into a standard algorithm that is able to maximize $\ell(\boldsymbol{\theta};\mathbf{f})$. For finite b, one may alternatively use the multiple imputation approach described on p. 76 of Rubin (1987), where an estimate $\hat{\boldsymbol{\theta}}_{nj}$ $(j=1,\ldots,b)$ is computed for each of the b imputed datasets by generating the imputed values under the posterior distribution of the missing values given the observed values under some specified Bayesian model. A final estimate is then formed by $\hat{\boldsymbol{\theta}}_n=\frac{1}{b}\sum_{j=1}^{b}\hat{\boldsymbol{\theta}}_{nj}$. Under some regularity assumptions, p. 117 of Robins & Wang (2000) shows these procedures to be n-asymptotically equivalent for $b=\infty$, while for finite b they have the same asymptotic mean but the frequentist method has a strictly smaller asymptotic variance.

5.B Distribution of the Thinned Marked Point Process

5.B.1 Proof of Lemma 5.3.1

We derive an operational expression for the hazard μ_{jk}, which in turn gives an expression for the partial log-likelihood $\ell(\boldsymbol{\theta};\mathbf{f})$. The calculations become straightforward by focusing on the jump time and jump mark distribution of the next jump in $(T_m,Y_m)_{m\geq 1}$ given the previous jump times and jump marks. Let $\mathbf{J}_m^*=(T_\ell^*,Y_\ell^*)_{1\leq\ell\leq m}$ and $\mathbf{J}_m=(T_\ell,Y_\ell)_{1\leq\ell\leq m}$ denote the first m jump times and jump marks. The baseline covariates are suppressed throughout.

Proof. Let P be the distribution under $\boldsymbol{\theta}$ and $\mathbf{f}$. All calculations hold on the event $(T_m<\infty)$. For the jump time distribution, Assumption 5.2.1 implies that

$$
\begin{aligned}
P(T_{m+1}\leq t\mid\mathbf{J}_m)&=P(T_{m+1}^*\leq t,T_{m+1}^*+U_{m+1}^*\leq\eta\mid\mathbf{J}_m^*,T_m^*+U_m^*\leq\eta)\\
&=\int_0^t\int_0^{\eta-s}\frac{P(U_{m+1}^*\in\mathrm{d}u\mid\mathbf{J}_m^*,T_{m+1}^*=s)}{P(T_m^*+U_m^*\leq\eta\mid\mathbf{J}_m^*)}P(T_{m+1}^*\in\mathrm{d}s\mid\mathbf{J}_m^*).
\end{aligned}
$$

For $P(U_{m+1}^*\in\mathrm{d}u\mid\mathbf{J}_m^*,T_{m+1}^*=s)$, one may use the law of total probability with respect to Y^*_{m+1} to get

$$P(U^*_{m+1} \in \mathrm{d}u \mid \mathbf{J}^*_m, T^*_{m+1} = s) = \sum_{k=1}^{J} \mathrm{pr}_U\{\mathrm{d}u; (\mathbf{J}^*_m, s, k), \mathbf{f}\} \frac{\mu^*_{Y^*_m k}(s; \mathbf{J}^*_m, \boldsymbol{\theta})}{\mu^*_{Y^*_m \bullet}(s; \mathbf{J}^*_m, \boldsymbol{\theta})}.$$

Furthermore,

$$P(T^*_{m+1} \in \mathrm{d}s \mid \mathbf{J}^*_m) = 1_{(s \geq T^*_m)} \exp\left\{-\int_{T_m}^{s} \mu^*_{Y^*_m \bullet}(v; \mathbf{J}^*_m, \boldsymbol{\theta})\,\mathrm{d}v\right\} \mu^*_{Y^*_m \bullet}(s; \mathbf{J}^*_m, \boldsymbol{\theta})\,\mathrm{d}s$$

by Assumption 5.2.1. Inserting these results gives

$$\begin{aligned} P(T_{m+1} \leq t \mid \mathbf{J}_m) = \mathrm{pr}_U(\eta - T^*_m; \mathbf{J}^*_m, \mathbf{f})^{-1} \times \int_{T^*_m}^{t} \exp\left\{-\int_{T_m}^{s} \mu^*_{Y^*_m \bullet}(v; \mathbf{J}^*_m, \boldsymbol{\theta})\,\mathrm{d}v\right\} \\ \times \sum_{k=1}^{J} \mathrm{pr}_U\{\eta - s; (\mathbf{J}^*_m, s, k), \mathbf{f}\} \mu^*_{Y^*_m k}(s; \mathbf{J}^*_m, \boldsymbol{\theta})\,\mathrm{d}s. \end{aligned}$$

We now move on to the jump mark distribution. For $t \leq \eta$,

$$\begin{aligned} P(Y_{m+1} = k \mid \mathbf{J}_m, T_{m+1} = t) &= P(Y^*_{m+1} = k \mid \mathbf{J}^*_m, T^*_{m+1} = t, T^*_{m+1} + U^*_{m+1} \leq \eta) \\ &= \frac{P(Y^*_{m+1} = k, U^*_{m+1} \leq \eta - t \mid \mathbf{J}^*_m, T^*_{m+1} = t)}{P(U^*_{m+1} \leq \eta - t \mid \mathbf{J}^*_m, T^*_{m+1} = t)}. \end{aligned}$$

Similar factorizations to before lead to

$$\begin{aligned} P(Y_{m+1} = k \mid \mathbf{J}_m, T_{m+1} = t) &= \frac{\mathrm{pr}_U\{\eta - t; (\mathbf{J}^*_m, t, k), \mathbf{f}\} \mu^*_{Y^*_m k}(t; \mathbf{J}^*_m, \boldsymbol{\theta})}{\sum_{\ell=1}^{J} \mathrm{pr}_U\{\eta - t; (\mathbf{J}^*_m, t, \ell), \mathbf{f}\} \mu^*_{Y^*_m \ell}(t; \mathbf{J}^*_m, \boldsymbol{\theta})} \\ &= \frac{\nu_{Y^*_m k}(t; \mathbf{J}^*_m, \boldsymbol{\theta}, \mathbf{f})}{\nu_{Y^*_m \bullet}(t; \mathbf{J}^*_m, \boldsymbol{\theta}, \mathbf{f})}. \end{aligned}$$

Therefore,

$$\begin{aligned} &\frac{P(T_{m+1} \in \mathrm{d}t, Y_{m+1} = k \mid \mathbf{J}_m)}{P(T_{m+1} \geq t \mid \mathbf{J}_m)} \\ &= \frac{\exp\left\{-\int_{T_m}^{t} \mu^*_{Y_m \bullet}(s; \mathbf{J}_m, \boldsymbol{\theta})\,\mathrm{d}s\right\} \nu_{Y_m k}(t; \mathbf{J}_m, \boldsymbol{\theta}, \mathbf{f})}{\mathrm{pr}_U(\eta - T_m; \mathbf{J}_m, \mathbf{f}) - \int_{T_m}^{t} \exp\left\{-\int_{T_m}^{s} \mu^*_{Y_m \bullet}(v; \mathbf{J}_m, \boldsymbol{\theta})\,\mathrm{d}v\right\} \nu_{Y_m \bullet}(s; \mathbf{J}_m, \boldsymbol{\theta}, \mathbf{f})\,\mathrm{d}s}\,\mathrm{d}t. \end{aligned}$$

By Jacod's formula for the intensity, see Proposition (3.1) of Jacod (1975), we immediately get

$$\mu_{jk}\{t; \mathbf{J}(t-), \boldsymbol{\theta}, \mathbf{f}\} = \gamma(t) \times \nu_{jk}\{t; \mathbf{J}(t-), \boldsymbol{\theta}, \mathbf{f}\}$$

on $\{I_j(t) = 1\}$ using

$$\begin{aligned} &\exp\left(-\int_{T_m}^{t} \mu^*_{Y_m \bullet}(s; \mathbf{J}_m, \boldsymbol{\theta})\,\mathrm{d}s\right) \\ &\quad = 1 - \int_{T_m}^{t} \exp\left(-\int_{T_m}^{s} \mu^*_{Y_m \bullet}(v; \mathbf{J}_m, \boldsymbol{\theta})\,\mathrm{d}v\right) \mu^*_{Y_m \bullet}(s; \mathbf{J}_m, \boldsymbol{\theta})\,\mathrm{d}s. \end{aligned}$$

□

Remark 5.B.1. Expressing the thinned hazard in terms of the original hazard allows one to use observations in the thinned process to estimate the original hazard. Usually one has some expert knowledge about the form of the underlying hazard, while positing a reasonable model for the thinned hazard can be more challenging. Therefore, it becomes natural to express the thinned hazard in terms of the original hazard, as is our approach. If one had instead sought to use a data-adaptive method to estimate the thinned hazard μ directly, it would have been more natural to invert the above formula to obtain an expression for the original hazard μ^* as a function of μ. Under sufficient smoothness conditions, one can obtain a coupled system of differential equations for the hazards $\mu^*_{jk}\{s;\mathbf{J}(s-),\boldsymbol{\theta}\}$ in terms of $\mu_{jk}\{s;\mathbf{J}(s-),\boldsymbol{\theta}\}$. We briefly sketch the argument. Note that, on $(T_m<\infty, Y_m=j)$,

$$\begin{aligned}
&\exp\left\{-\int_{T_m}^t \mu_{j\bullet}(s;\mathbf{J}_m,\boldsymbol{\theta},\mathbf{f})\,\mathrm{d}s\right\}\mu_{jk}(t;\mathbf{J}_m,\boldsymbol{\theta},\mathbf{f}) \\
&= \exp\left\{-\int_{T_m}^t \mu^*_{j\bullet}(s;\mathbf{J}_m,\boldsymbol{\theta})\,\mathrm{d}s\right\}\mu^*_{jk}(t;\mathbf{J}_m,\boldsymbol{\theta}) \times \frac{\mathrm{pr}_U\{\eta-t;(\mathbf{J}_m,t,k),\mathbf{f}\}}{\mathrm{pr}_U(\eta-T_m;\mathbf{J}_m,\mathbf{f})}.
\end{aligned}$$

This follows by casting $\mathbb{P}(T_{m+1}\in\mathrm{d}t\mid\mathbf{J}_m)\mathbb{P}(Y_{m+1}=k\mid\mathbf{J}_m,T_{m+1}=t)$ first in terms of μ, invoking the definition of μ_{jk}, and secondly using the expressions obtained above. Applying the logarithm and differentiating with respect to t gives a differential equation for μ^*_{jk}. Together, these differential equations form a system of nonlinear quadratic differential equations which may be solved numerically.

Remark 5.B.2. If instead of imposing a parametric model, one had assumed the generalized Cox model of Dabrowska (1997), which allows for time-dependent covariates, the derivation of the hazard in the thinned marked point process would be unchanged, implying that μ_{jk} would no longer be on the proportional hazards form. However, ν_{jk} stays on the proportional hazards form since the baseline hazard absorbs the reporting delay distribution. When using the smoothed profile likelihood of Dabrowska (1997) and the Poisson approximation, one may thus ignore reporting delays for inference about the parametric part of a Cox model.

5.B.2 Approximating the hazard

Recall

$$\check{\gamma}(t)=\frac{1-\int_{t-V_{t-}}^t P^*\{s;\mathbf{J}(s-),\boldsymbol{\theta}\}\mu^*_{j\bullet}\{s;\mathbf{J}(s-),\boldsymbol{\theta}\}\,\mathrm{d}s}{1-\int_{t-V_{t-}}^t P^*\{s;\mathbf{J}(s-),\boldsymbol{\theta}\}\nu_{j\bullet}\{s;\mathbf{J}(s-),\boldsymbol{\theta},\mathbf{f}\}\,\mathrm{d}s}.$$

Since $\nu_{jk}\{s;\mathbf{J}(s-),\boldsymbol{\theta},\mathbf{f}\}\le\mu^*_{jk}\{s;\mathbf{J}(s-),\boldsymbol{\theta}\}$, it holds that $\check{\gamma}(t)\le 1$. Conversely, it holds that $\nu_{j\bullet}\{s;\mathbf{J}(s-),\boldsymbol{\theta},\mathbf{f}\}\ge 0$ so $\check{\gamma}(t)\ge P^*\{t;\mathbf{J}(t-),\boldsymbol{\theta}\}$. These are the bounds reported in the main text. To show that $\check{\mu}_{jk}$ and ν_{jk} only differ by a second order term, impose Assumption 5.B.3.

Assumption 5.B.3. *As a function of s, the hazards $\mu_{jk}(s;\mathbf{J},\boldsymbol{\theta},\mathbf{f})$ and $\nu_{jk}(s;\mathbf{J},\boldsymbol{\theta},\mathbf{f})$ are continuous and almost everywhere continuously differentiable.*

Assumption 5.B.3 is not used elsewhere in the paper, but similar assumptions are made when deriving asymptotic properties of the estimators. Assumption 5.B.3 implies that the hazards are almost everywhere continuously differentiable so Taylor's theorem with Lagrange remainder at the point $t - V_{t-}$ and some straightforward calculations yield that

$$\begin{aligned}
\check{\gamma}(t) &= 1 + V_{t-} \times \mu^*_{j\bullet}\{\zeta; \mathbf{J}(\zeta-), \boldsymbol{\theta}\} \times P^*\{\zeta; \mathbf{J}(\zeta-), \boldsymbol{\theta}\} \\
&\quad \times \Bigg(\frac{\nu_{j\bullet}\{\zeta; \mathbf{J}(\zeta-), \boldsymbol{\theta}, \mathbf{f}\}}{\mu^*_{j\bullet}\{\zeta; \mathbf{J}(\zeta-), \boldsymbol{\theta}\}} \left[1 - \int_{t-V_{t-}}^{\zeta} P^*\{s; \mathbf{J}(s-), \boldsymbol{\theta}\} \mu^*_{j\bullet}\{s; \mathbf{J}(s-), \boldsymbol{\theta}\}\, \mathrm{d}s\right] \\
&\qquad + \int_{t-V_{t-}}^{\zeta} P^*\{s; \mathbf{J}(s-), \boldsymbol{\theta}\} \nu_{j\bullet}\{s; \mathbf{J}(s-), \boldsymbol{\theta}, \mathbf{f}\}\, \mathrm{d}s - 1 \Bigg) \\
&\quad \times \left[1 - \int_{t-V_{t-}}^{\zeta} P^*\{s; \mathbf{J}(s-), \boldsymbol{\theta}\} \nu_{j\bullet}\{s; \mathbf{J}(s-), \boldsymbol{\theta}, \mathbf{f}\}\, \mathrm{d}s\right]^{-2} \\
&= 1 + R(t)
\end{aligned}$$

for some ζ between $t - V_{t-}$ and t. Since $\check{\gamma}$ is bounded above by one, the term $R(t)$ is negative. The term $\nu_{j\bullet}\{\zeta; \mathbf{J}(\zeta-), \boldsymbol{\theta}, \mathbf{f}\} / \mu^*_{j\bullet}\{\zeta; \mathbf{J}(\zeta-), \boldsymbol{\theta}\} \times \Big[1 - \int_{t-V_{t-}}^{\zeta} P^*\{s; \mathbf{J}(s-), \boldsymbol{\theta}\}$ $\times \mu^*_{j\bullet}\{s; \mathbf{J}(s-), \boldsymbol{\theta}\}\, \mathrm{d}s\Big]$ is positive and $R(t)$ thus necomes more negative by removing this term. We thus get the following crude bound:

$$\begin{aligned}
|R(t)| &\leq V_{t-} \times \mu^*_{j\bullet}\{\zeta; \mathbf{J}(\zeta-), \boldsymbol{\theta}\} \times P^*\{\zeta; \mathbf{J}(\zeta-), \boldsymbol{\theta}\} \\
&\quad \times \frac{1 - \int_{t-V_{t-}}^{\zeta} P^*\{s; \mathbf{J}(s-), \boldsymbol{\theta}\} \nu_{j\bullet}\{s; \mathbf{J}(s-), \boldsymbol{\theta}, \mathbf{f}\}\, \mathrm{d}s}{[1 - \int_{t-V_{t-}}^{\zeta} P^*\{s; \mathbf{J}(s-), \boldsymbol{\theta}\} \nu_{j\bullet}\{s; \mathbf{J}(s-), \boldsymbol{\theta}, \mathbf{f}\}\, \mathrm{d}s]^2} \\
&= V_{t-} \times \mu^*_{j\bullet}\{\zeta; \mathbf{J}(\zeta-), \boldsymbol{\theta}\} \times \check{\gamma}(\zeta) \\
&\leq V_{t-} \times \mu^*_{j\bullet}\{\zeta; \mathbf{J}(\zeta-), \boldsymbol{\theta}\}.
\end{aligned}$$

The bound in the main text now follows by using that

$$\nu_{jk}\{t; \mathbf{J}(t-), \boldsymbol{\theta}, \mathbf{f}\} \leq \mu^*_{j\bullet}\{t; \mathbf{J}(t-), \boldsymbol{\theta}\} \leq \sup_{s \in [t-V_{t-}, t]} \mu^*_{j\bullet}\{s; \mathbf{J}(s-), \boldsymbol{\theta}\}.$$

5.C Fully Endogenous Reporting Delays

5.C.1 Estimation procedure

In the proposed method, adjudication enters exogenously while reporting delays enter partially exogenously and partially endogenously with respect to the marked point process of interest. Included jumps are on the form $(T^*_m, Y^*_m) \times 1_{(T^*_m + U^*_m \leq \eta)}$ instead of the fully endogenous form $(T^*_m, Y^*_m, U^*_m) \times 1_{(T^*_m + U^*_m \leq \eta)}$, while the adjudication information $\mathcal{A}_{\langle \eta \rangle, \eta}$ only affects the imputation of $\xi_{\langle \eta \rangle}$. Here, we investigate what happens if reporting delays instead are made fully endogenous. A result of this is that one must describe how subsampling by prior reporting delays affects the intensity of future jumps. However, if one is able to specify this correctly, for example through independence assumptions, there would be an efficiency gain in using the resulting estimator. If $\mathcal{A}_{m,t}$ is generated by a marked point process, one could similarly make the adjudication information endogenous by extending with these jump times and jump marks. However, it seems reasonable to retain adjudication as exogeneous, as opposed to reporting delays, since the latter often stem from the subject's conduct, while adjudication likely results from the internal processes of the data collector. This leads to a marked point process $(T_m^*, Y_m^*, U_m^*)_{m\geq 1}$. To relate the hazards of this model to the hazards of $(T_m^*, Y_m^*)_{m\geq 1}$, one has to posit a model for how the hazards subsampled on reporting delays compare to the ones where one does not condition on prior reporting delays. An example of such a model could be the full independence $(T_{m+1}^*, Y_{m+1}^*) \perp\!\!\!\perp U_1^*, \dots, U_m^* \mid \mathbf{J}_m^*$, meaning that the hazards are unchanged by the additional information. Define $\mathbf{J}_U(t)$ as the time t event history of $(T_m, Y_m, U_m)_{m\geq 1}$. Redoing the calculations of the distribution of the thinned marked point process leads to the partial likelihood

$$\begin{aligned}\ell(\boldsymbol{\theta}; \mathbf{f}) = \sum_{j,k=1}^{J} \int_V^C &\log \mu_{jk}\{t; \mathbf{X}, \mathbf{J}_U(t-), \boldsymbol{\theta}, \mathbf{f}\}\, \mathrm{d}N_{jk}(t) \\ &- \int_V^C I_j(t)\mu_{jk}\{t; \mathbf{X}, \mathbf{J}_U(t-), \boldsymbol{\theta}, \mathbf{f}\}\, \mathrm{d}t \\ &+ \int_V^C \log\left(\frac{\mathrm{pr}_U[\mathrm{d}U_{N_{\bullet\bullet}(t)}; \{\mathbf{J}_U(t-), t, k\}, \mathbf{f}]}{\mathrm{pr}_U[\eta - t; \{\mathbf{J}_U(t-), t, k\}, \mathbf{f}]}\right) \mathrm{d}N_{jk}(t).\end{aligned}$$

In a slight abuse of notation, we have reused the previous notations $\ell(\boldsymbol{\theta}; \mathbf{f})$ and μ_{jk}, where the latter now reads $\mu_{jk}\{t; \mathbf{J}_U(t-), \boldsymbol{\theta}, \mathbf{f}\} = \gamma(t) \times \nu_{jk}\{t; \mathbf{J}_U(t-), \boldsymbol{\theta}, \mathbf{f}\}$ for $\gamma(t)$ being

$$\frac{1 - \int_{t-V_{t-}}^{t} P^*\{s; \mathbf{J}(s-), \boldsymbol{\theta}\}\mu_{j\bullet}^*\{s; \mathbf{J}(s-), \boldsymbol{\theta}\}\, \mathrm{d}s}{\mathrm{pr}_U\{\eta - (t - V_{t-}); \mathbf{J}(t - V_{t-}), \mathbf{f}\} - \int_{t-V_{t-}}^{t} P^*\{s; \mathbf{J}(s-), \boldsymbol{\theta}\}\nu_{j\bullet}\{s; \mathbf{J}_U(s-), \boldsymbol{\theta}, \mathbf{f}\}\, \mathrm{d}s}$$

and $\nu_{jk}\{t; \mathbf{J}_U(t-), \boldsymbol{\theta}, \mathbf{f}\} = \mu_{jk}^*\{t; \mathbf{J}(t-), \boldsymbol{\theta}\}\mathrm{pr}_U\big[\eta - t; \{\mathbf{J}_U(t-), t, k\}, \mathbf{f}\big]$. Adjusting for adjudication then leads us to introduce $\ell(\mathbf{Z}, \boldsymbol{\theta}; \mathbf{f}, \mathbf{g}) = \mathbb{E}_{\mathbf{g}}\{\ell(\boldsymbol{\theta}; \mathbf{f}) \mid \mathbf{Z}\} = w(0, \mathbf{Z}; \mathbf{g})\ell(\boldsymbol{\theta}; \mathbf{f}, 0) + w(1, \mathbf{Z}; \mathbf{g})\ell(\boldsymbol{\theta}; \mathbf{f}, 1)$. The estimation algorithm becomes more involved, even under the Poisson approximation, since $\mathbf{f}$ in the optimizing step now additionally enters into the first two terms of $\ell(\boldsymbol{\theta}, \mathbf{f})$. For fixed $\mathbf{f}$, the estimation of $\boldsymbol{\theta}$ proceeds as in Section 5.3.3. This suggests to optimize $\ell(\boldsymbol{\theta}; \mathbf{f})$ by defining $\hat{\boldsymbol{\theta}}_n(\mathbf{f})$ as the estimated value of $\boldsymbol{\theta}$ for a fixed value $\mathbf{f}$, and then optimizing $\ell\{\hat{\boldsymbol{\theta}}_n(\mathbf{f}); \mathbf{f}\}$ over $\mathbf{f}$ using some suitable numerical optimizer.

5.C.2 Efficiency

To discuss efficiency, we assume as in Section 5.4 that scores exist and we work in a neighborhood of the true parameter where the score has a unique zero. We also impose certain regularity conditions such that $\hat{\boldsymbol{\theta}}_n$ is asymptotically normal in both cases. In principle, the efficiency gain or loss by making the reporting delays endogenous can be inferred by comparing the expressions for the asymptotic variance of $\hat{\boldsymbol{\theta}}_n$ in the two cases. The relation between the covariance matrices, however, appears unclear in the general case. Since the space of distributions satisfying the moment conditions in the case where reporting delays are endogenous is strictly contained in the space of distributions satisfying the moment conditions in the partially endogenous case, the semiparametric efficiency bound in the fully endogenous case is smaller or equal to the one in the partially endogenous case. Consequently, when using the optimal weighting estimator as discussed in Section 5.4, the estimator in the fully endogenous case is at least as efficient as the one in the partially endogeneous case.

5.C.3 Admissibility in data application

The assumption of full independence is not admissible for the disability semi-Markov application of Section 5.6. This is due to the fact that it is assumed that death has no reporting delay, while monotone reporting remains imposed; this has the effect that knowing (T^*_m, U^*_m) implies that a death cannot have occurred in $[T^*_m, T^*_m + U^*_m]$. It might then instead be natural to say that the only extra information gained by conditioning on reporting delays is that the subject has not died before $T^*_m + U^*_m$ In Section 6 of Hoem (1972), the author obtains an expression for the transition probabilities and rates of a semi-Markov process conditional on the time of absorption not having occurred before a given time. Let Y be a semi-Markov process with duration process V, and let the transition probabilities be denoted $p^*_{jk}(s,t,u,z) = \mathbb{P}(Y_t = k, V_t \le z \mid Y_s = j, V_s = u)$. Assume the state space $1, \dots, J$ can be partitioned into two parts $\mathbf{R}_1$ and $\mathbf{R}_2$, where $\mathbf{R}_2$ is absorbing in the sense that $\mathbf{R}_1$ cannot be reached from $\mathbf{R}_2$. The transition rates conditional on being in $\mathbf{R}_1$ at time t are given by

$$\mu^*_{jk}(s,u;t) = \mu^*_{jk}(s,u) \frac{\sum_{i \in \mathbf{R}_1} p^*_{ki}(s,t,0,\infty)}{\sum_{i \in \mathbf{R}_1} p^*_{ji}(s,t,u,\infty)}$$

for $s < t$ and $j \in \mathbf{R}_1$, while $\mu^*_{jk}(s,u;t) = \mu^*_{jk}(s,u)$ for $s \ge t$. So in principle, this situation may also be handled by our approach. However, it would make the estimation procedure considerably more computationally demanding since the inclusion of the transition probabilities p^*_{jk} adds another layer of numerical integration that has to be performed separately for each candidate value of $\boldsymbol{\theta}$.

5.D Implementation of Estimation Procedure

A general approach to implementation of the estimation procedure for our parametric models is to discretize the integrals and use Poisson regression. For counting process likelihoods written in terms of the usual hazards, this is well-known and has been noted in, for instance, Lindsey (1995), but we see that this is also true for the reporting delay distribution parameterized in terms of reverse time hazards, as well as for the Poisson approximation of the likelihood. Contrary to the other supplements, here the dependence on baseline covariates is made explicit.

Select a partition $0 = t_0 < \ldots < t_A = \eta$ and discretize the likelihood according to this partition. Then

$$\ell_\xi(\mathbf{Z}^{\text{obs}}, \mathbf{g}) = \sum_{m=1}^{\langle\eta\rangle} \sum_{j,k=1}^{K} \sum_{a=0}^{A-1} \log \omega_{jk}(t_a; \mathcal{H}_{t_a}, \mathbf{g}) O_{jk}^{\xi}(a, m) \\ - \omega_{jk}(t_a; \mathcal{H}_{t_a}, \mathbf{g}) E_j^{\xi}(a, m) + \varepsilon_\xi(A)$$

where $O_{jk}^{\xi}(a,m) = N_{m,jk}(t_{a+1}) - N_{m,jk}(t_a)$ and $E_j^{\xi}(a,m) = \int_{t_a}^{t_{a+1}} I_{m,j}(s)\,\mathrm{d}s$ are the adjudication occurrences and exposures, respectively, while $\varepsilon_\xi(A) \to 0$ for $A \to \infty$. Multiplying with the exposures inside the logarithm term, we recognize this as the likelihood corresponding to the situation where $O_{jk}^{\xi}(a,m)$, $(a = 0, \ldots, A-1; m = 1, \ldots, \langle\eta\rangle; j, k = 1, \ldots, K)$, are mutually independent Poisson random variables with mean $\log\{\omega_{jk}(t_a; \mathcal{H}_{t_a}, \mathbf{g}) E_j^{\xi}(a,m)\}$. This further corresponds to a generalized linear model with the Poisson family, log link, mean $\log \omega_{jk}(t_a; \mathcal{H}_{t_a}, \mathbf{g})$, and offset $\log E_j^{\xi}(a,m)$. The discretization may either be seen as an approximation which converges to the true criterion function for $A \to \infty$, or it may be seen as the true likelihood if the partition is chosen to be equal to the precision of the time measurement.

Similarly for the reporting delays, let $\bar{t}_a = (t_{a+1} + t_a)/2$ and note that

$$\ell_U(\mathbf{f}) = \sum_{m=1}^{N_{\cdot\cdot}^c(\eta)} \sum_{a=0}^{A-1} \log \alpha\{\bar{t}_a; \mathbf{J}^c(T_m^c), \mathbf{X}, \mathbf{f}\} O^U(a, m) \\ - \alpha\{\bar{t}_a; \mathbf{J}^c(T_m^c), \mathbf{X}, \mathbf{f}\} E^U(a, m) + \varepsilon_U(A),$$

where $O^U(a,m) = 1_{(U_m^c \in (t_a, t_{a+1}])}$ and $E^U(a,m) = \int_{t_a}^{t_{a+1}} 1_{(U_m^c \le s \le \eta - T_m^c)}\,\mathrm{d}s$ are the reporting delay occurrences and exposures, while $\varepsilon_U(A) \to 0$ for $A \to \infty$. Note that the partition here is allowed to differ from the one of $\ell_\xi(\mathbf{Z}^{\text{obs}}, \mathbf{g})$ despite using the same notation. Here, the midpoints are chosen to take into account that the reverse time hazard takes the value $+\infty$ at time 0.

For the approximate likelihood of $\boldsymbol{\theta}$, we get

$$\ell^{\text{app}}(\boldsymbol{\theta}; \mathbf{f}) = \sum_{j,k=1}^{J} \sum_{a=0}^{A-1} \log \mu_{jk}^{*}\{t_a; \mathbf{J}(t_a), \mathbf{X}, \boldsymbol{\theta}\} O_{jk}(a) \\ - \mu_{jk}^{*}\{t_a; \mathbf{J}(t_a), \mathbf{X}, \boldsymbol{\theta}\} \mathbb{P}_U[\eta - t_a; \{\mathbf{J}(t_a), t_a, k\}, \mathbf{X}, \mathbf{f}] E_j(a) + \varepsilon(A)$$

where $O_{jk}(a) = N_{jk}^c(t_{a+1}) - N_{jk}^c(t_a)$ and $E_j(a) = \int_{t_a}^{t_{a+1}} 1_{(V,C]}(s) I_j(s)\,\mathrm{d}s$ are the occurrences and exposures, respectively, while $\varepsilon(A) \to 0$ for $A \to \infty$. When computing $\ell^{\mathrm{app}}(\mathbf{Z}, \boldsymbol{\theta}; \mathbf{f}, \mathbf{g})$, the occurrence and exposure contributions have to be weighted with the adjudication probabilities, but otherwise nothing changes. Most generalized linear model software packages allow the user to specify such weights.

The link to Poisson regression simplifies practical implementation significantly. Another useful consequence is that the link provides a natural way to construct plots in which the fitted rates can be compared with empirical rates. If the hazards are assumed to be piecewise constant, the maximum likelihood estimates become the well-known occurrence-exposure rates. It is these occurrence-exposure rates which we refer to as the empirical rates. The piecewise constant hazard assumption can be imposed for any continuous covariate, for example calendar time, age, or the duration since the occurrence of given event. By comparing the empirical rates with the predicted rates from the fitted parametric model, one obtains a goodness-of-fit plot.

In our numerical study, that data is on a jump time and jump mark format. There we forego the Poisson regression, opting instead to utilize the specific choice of parametric model to obtain a direct implementation; this includes the use of a standard optimizer to maximize the resulting expressions. In our data application, the data is on a monthly grid format, and there is considerably more data, which leads us to use the Poisson regression approach outlined above, since estimation might then exactly be based on aggregated data. When the number of observations is significantly larger than the covariate space, which is the case in our data application, this saves considerable amounts of storage and computation time.

5.E Asymptotics for the Parameters

Consistency of the nuisance parameter estimators $\hat{\mathbf{f}}_n$ and $\hat{\mathbf{g}}_n$ is needed in order to achieve consistency of $\hat{\boldsymbol{\theta}}_n$ and these estimators generally need a convergence rate of a higher order than $n^{1/4}$ in order to obtain asymptotic normality of $\hat{\boldsymbol{\theta}}_n$, see e.g. Theorem 2 of Chen et al. (2003). In this framework, the convergence rate is obtained with plenty to spare since parametric models under regularity conditions and correct specifications admit $n^{1/2}$-rates.

5.E.1 Consistency

Before showing asymptotic normality, we establish consistency. We focus on weak consistency, as this suffices for asymptotic distribution theory. Showing strong consistency of $\hat{\mathbf{g}}_n$, $\hat{\mathbf{f}}_n$, and $\hat{\boldsymbol{\theta}}_n$ would, however, require almost no additional effort. For Theorem 2.1 in Newey & McFadden (1994), one would change (iv) to hold almost surely in order to obtain strong consistency. Similarly, for Theorem 1 in Delsol & Van Keilegom (2020), one could change (A1) such that the estimator is required to be the actual maximizer of the objective function, which is satisfied for our estimators, and further impose that (A3) and (A4) hold almost surely instead of in probability. Since we show (iv) and (A4) by a Glivenko–Cantelli class argument, which implies almost sure convergence, we could equally well claim strong consistency.

Our approach to showing consistency and asymptotic normality is not to find minimal assumptions, but rather to formulate sufficient assumptions that are easy to check in a given application and, in particular, are not too restrictive for our applications.

Assumption 5.E.1.

(i) Different $\mathbf{g}$ *give rise to* ω_{jk} *that are not almost surely equal.*

(ii) The $\mathbf{g}$ *parameter space* $\mathbb{G}$ *is a compact subset of a Euclidean space, and* $\mathbf{g}_0$ *belongs to the interior.*

(iii) The hazards $\omega_{jk}(t;\mathbf{H},\mathbf{g})$ *are continuous in* $\mathbf{g}$*.*

(iv) $\sup_{\mathbf{g}} |\ell_\xi(\mathbf{Z}^{\mathrm{obs}},\mathbf{g})|$ *has finite expectation and* $\sup_{\mathbf{g},\mathbf{H}} \omega_{jk}(t;\mathbf{H},\mathbf{g})$ *are Lebesgue-integrable on* $[0,\eta]$*.*

In the following, we use the usual stochastic order notation o_P, that is we write $x_n = o_P(r_n)$ if $|x_n|/r_n \to 0$ in probability.

Proposition 5.E.2. *Under Assumptions 5.2.1-5.2.3 and Assumption 5.E.1, it holds that* $\|\hat{\mathbf{g}}_n - \mathbf{g}_0\| = o_P(1)$*.*

Proof. We verify the conditions of Theorem 2.1 in Newey & McFadden (1994). Condition (i) is satisfied by non-negativity of the Kullback-Leibler divergence, also known as Gibbs' inequality or the information inequality, and the identifiability of the model implied by Assumption 5.E.1(i); each contribution to the sum in ℓ_ξ is a usual log-likelihood of a marked point process which is maximized in $\mathbf{g}_0$ by Gibbs' inequality, and the identifiability ensures that a different $\mathbf{g}$ would lead to a smaller value for at least one of the contributions using Gibbs' inequality and the one-to-one correspondence between compensators and distributions. Condition (ii) is satisfied by Assumption 5.E.1(ii). We now consider Condition (iii). For a given $\mathbf{g} \in \mathbb{G}$, take $(\mathbf{g}_j)_{j\geq 1}$ with $\mathbf{g}_j \to \mathbf{g}$. By Assumption 5.E.1(iv), we may use dominated convergence which implies $\lim_{j\to\infty} \mathbb{E}\{\ell_\xi(\mathbf{Z}^{\mathrm{obs}},\mathbf{g}_j)\} = \mathbb{E}\{\lim_{j\to\infty} \ell_\xi(\mathbf{Z}^{\mathrm{obs}},\mathbf{g}_j)\}$. Again by Assumption 5.E.1(iv), dominated convergence and Assumption 5.E.1(iii) gives $\lim_{j\to\infty} \ell_\xi(\mathbf{Z}^{\mathrm{obs}},\mathbf{g}_j) = \ell_\xi(\mathbf{Z}^{\mathrm{obs}},\mathbf{g})$. Thus, Condition (iii) holds. Condition (iv) corresponds to showing that $\{\ell_\xi(\cdot,\mathbf{g}) : \mathbf{g} \in \mathbb{G}\}$ is Glivenko–Cantelli. We will use Example 19.8 of Van der Vaart (1998). We note that $\mathbf{g} \mapsto \ell_\xi(\mathbf{Z}^{\mathrm{obs}},\mathbf{g})$ is continuous for any $\mathbf{Z}^{\mathrm{obs}}$ by previous observations and has an integrable envelope function by

Assumption 5.E.1(iv). Hence the class is Glivenko–Cantelli and Condition (iv) is satisfied. □

Assumption 5.E.3.

(i) Different $\mathbf{f}$ *gives rise to* α *that are not almost surely equal.*

(ii) The $\mathbf{f}$ *parameter set* $\mathbb{F}$ *is a compact subset of a Euclidean space, and* $\mathbf{f}_0$ *belongs to the interior.*

(iii) The reverse time hazard $\alpha(t; \mathbf{J}, \mathbf{X}, \mathbf{f})$ *is continuous in* $\mathbf{f}$*.*

(iv) $\sup_{\mathbf{f}} |\ell_U(\mathbf{f}; k)|$ *for* $k = 0, 1$ *has finite expectation and* $\sup_{\mathbf{f},\mathbf{J},\mathbf{X}} \alpha(t; \mathbf{J}, \mathbf{X}, \mathbf{f})$ *is Lebesgue-integrable on* $[0, \eta]$*.*

(v) The adjudication probabilities $w(1, \mathbf{Z}; \mathbf{g})$ *are continuous in* $\mathbf{g}$*.*

Proposition 5.E.4. *Under Assumptions 5.2.1-5.2.3, Assumption 5.E.3, and if* $\|\hat{\mathbf{g}}_n - \mathbf{g}_0\| = o_P(1)$*, it holds that* $\|\hat{\mathbf{f}}_n - \mathbf{f}_0\| = o_P(1)$*.*

Proof. We verify the conditions in Theorem 1 of Delsol & Van Keilegom (2020). Condition (A1) is satisfied by the specification of $\hat{\mathbf{f}}_n$ as the maximizer of the objective function. By standard arguments, Condition (A2) is satisfied if $\mathbb{E}\{\ell_U(\mathbf{Z}, \mathbf{f}; \mathbf{g}_0)\}$ is continuous in $\mathbf{f}$ since $\mathbb{F}$ is compact. This is because $\mathbb{F}_\varepsilon = \{\mathbf{f} \in \mathbb{F} : \|\mathbf{f} - \mathbf{f}_0\| \geq \varepsilon\}$ is closed and hence compact, so in this case there exists $\mathbf{f}_\varepsilon$ which satisfies $\sup_{\mathbf{f}\in\mathbb{F}_\varepsilon} \mathbb{E}\{\ell_U(\mathbf{Z}, \mathbf{f}; \mathbf{g}_0)\} = \mathbb{E}\{\ell_U(\mathbf{Z}, \mathbf{f}_\varepsilon; \mathbf{g}_0)\} < \mathbb{E}\{\ell_U(\mathbf{Z}, \mathbf{f}_0; \mathbf{g}_0)\}$ by uniqueness of the maximum implied by Gibbs' inequality and the identifiability implied by Assumption 5.E.3(i). We show simultaneous continuity in $\mathbf{f}$ and $\mathbf{g}$, as this is useful for checking other regularity conditions. The continuity is a direct consequence of Assumptions 5.E.3(iii)-(v) by two applications of dominated convergence since $\sup_{\mathbf{f},\mathbf{g}} |\ell_U(\mathbf{Z}, \mathbf{f}; \mathbf{g})| \leq \sup_{\mathbf{f}} |\ell_U(\mathbf{f}; 1)| + \sup_{\mathbf{f}} |\ell_U(\mathbf{f}; 0)|$. Condition (A3) follows by the assumption that $\hat{\mathbf{g}}_n$ is weakly consistent. Condition (A4) follows if the class of functions $\{\ell_U(\cdot, \mathbf{f}; \mathbf{g}) : \mathbf{f} \in \mathbb{F}, \mathbf{g} \in \mathbb{G}\}$ is Glivenko–Cantelli. We use Example 19.8 in Van der Vaart (1998). Continuity and the existence of an integrable envelope function for the class follows by previous observations. Hence Condition (A4) is satisfied. For Condition (A5), we take $\mathbf{g}_j \to \mathbf{g}_0$ for $j \to \infty$. Note that by the triangle inequality, $|\mathbb{E}\{\ell_U(\mathbf{Z}, \mathbf{f}; \mathbf{g}_j)\} - \mathbb{E}\{\ell_U(\mathbf{Z}, \mathbf{f}; \mathbf{g}_0)\}| \leq \mathbb{E}\{|w(0, \mathbf{Z}; \mathbf{g}_j) - w(0, \mathbf{Z}; \mathbf{g}_0)| \times \sup_{\mathbf{f}} |\ell_U(\mathbf{f}; 0)| + |w(1, \mathbf{Z}; \mathbf{g}_j) - w(1, \mathbf{Z}; \mathbf{g}_0)| \times \sup_{\mathbf{f}} |\ell_U(\mathbf{f}; 1)|\}$. The right-hand side does not depend on $\mathbf{f}$ and converges to 0 for $\mathbf{g}_j \to \mathbf{g}_0$ by dominated convergence and Assumption 5.E.3(v), hence Condition (A5) holds. □

Assumption 5.E.5.

(i) Different $\boldsymbol{\theta}$ *gives rise to* μ^*_{jk} *that are not almost surely equal.*

(ii) *The* $\boldsymbol{\theta}$ *parameter space* Θ *is compact, and* $\boldsymbol{\theta}_0$ *belongs to the interior.*

(iii) $\sup_{\boldsymbol{\theta},\mathbf{f},\mathbf{g}} |\ell(\mathbf{Z},\boldsymbol{\theta};\mathbf{f},\mathbf{g})|$ *has finite expectation.*

(iv) $\sup_{\boldsymbol{\theta},\mathbf{J},\mathbf{X}} \mu^*_{jk}(t;\mathbf{J},\mathbf{X},\boldsymbol{\theta})$ *is Lebesgue-integrable on* $[0,\eta]$.

(v) *The hazards* $\mu^*_{jk}(t;\mathbf{J},\mathbf{X},\boldsymbol{\theta})$ *are continuous in* $\boldsymbol{\theta}$.

(vi) *The partial derivatives of* $\ell(\mathbf{Z},\boldsymbol{\theta};\mathbf{f},\mathbf{g})$ *with respect to the coordinates of* $\boldsymbol{\theta}$, $\mathbf{f}$, *and* $\mathbf{g}$ *all exist and are continuous. Furthermore,* $\sup_{\boldsymbol{\theta},\mathbf{f},\mathbf{g}} \|\nabla_{(\boldsymbol{\theta},\mathbf{f},\mathbf{g})}\ell(\mathbf{Z},\boldsymbol{\theta};\mathbf{f},\mathbf{g})\|$ *has finite expectation.*

Proposition 5.E.6. *Under Assumptions 5.2.1-5.2.3, Assumption 5.E.5, and if* $\|\hat{\mathbf{g}}_n - \mathbf{g}_0\| = o_P(1)$ *and* $\|\hat{\mathbf{f}}_n - \mathbf{f}_0\| = o_P(1)$, *it holds that* $\|\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0\| = o_P(1)$.

Proof. We verify the conditions in Theorem 1 of Delsol & Van Keilegom (2020). Condition (A1) is satisfied by the specification of $\hat{\theta}_n$ as the maximizer of the objective function. Similarly to Proposition 5.E.4, Condition (A2) is satisfied by standard arguments if $E\{\ell(\mathbf{Z},\boldsymbol{\theta};\mathbf{g}_0,\mathbf{f}_0)\}$ is continuous in $\boldsymbol{\theta}$ since Θ is compact by Assumption 5.E.5(ii) and the model is identifiable by Assumption 5.E.5(i). We start by showing the Lipschitz-type property $|\ell(\mathbf{Z},\boldsymbol{\theta}_1;\mathbf{f}_1,\mathbf{g}_1) - \ell(\mathbf{Z},\boldsymbol{\theta}_2;\mathbf{f}_2,\mathbf{g}_2)| \leq b(\mathbf{Z})\,\|(\boldsymbol{\theta}_1,\mathbf{f}_1,\mathbf{g}_1) - (\boldsymbol{\theta}_2,\mathbf{f}_2,\mathbf{g}_2)\|$ with b being an integrable function, as this is useful for showing other regularity conditions. If this holds, we see that

$$\begin{aligned}\Big|\mathbb{E}\{\ell(\mathbf{Z},\boldsymbol{\theta}_1;\mathbf{f}_1,\mathbf{g}_1)\} - \mathbb{E}\{\ell(\mathbf{Z},\boldsymbol{\theta}_2;\mathbf{f}_2,\mathbf{g}_2)\}\Big| &\leq \mathbb{E}\{|\ell(\mathbf{Z},\boldsymbol{\theta}_1;\mathbf{f}_1,\mathbf{g}_1) - \ell(\mathbf{Z},\boldsymbol{\theta}_2;\mathbf{f}_2,\mathbf{g}_2)|\} \\ &\leq \mathbb{E}\{b(\mathbf{Z})\,\|(\boldsymbol{\theta}_1,\mathbf{f}_1,\mathbf{g}_1) - (\boldsymbol{\theta}_2,\mathbf{f}_2,\mathbf{g}_2)\|\}\end{aligned}$$

and since b is integrable, dominated convergence implies that the right-hand side converges to 0 for $(\boldsymbol{\theta}_1,\mathbf{f}_1,\mathbf{g}_1) \to (\boldsymbol{\theta}_2,\mathbf{f}_2,\mathbf{g}_2)$, and hence $\mathbb{E}\{\ell(\mathbf{Z},\boldsymbol{\theta};\mathbf{f},\mathbf{g})\}$ is continuous in $\boldsymbol{\theta}$, $\mathbf{f}$, and $\mathbf{g}$. To show the Lipschitz-type property, note that Assumption 5.E.5(vi) implies that $\ell(\mathbf{Z},\boldsymbol{\theta};\mathbf{f},\mathbf{g})$ is differentiable in $(\boldsymbol{\theta},\mathbf{f},\mathbf{g})$ for any fixed $\mathbf{Z}$. The mean-value theorem thus implies

$$\begin{aligned}&\ell(\mathbf{Z},\boldsymbol{\theta}_1;\mathbf{f}_1,\mathbf{g}_1) - \ell(\mathbf{Z},\boldsymbol{\theta}_2;\mathbf{f}_2,\mathbf{g}_2) \\ &= \nabla_{(\boldsymbol{\theta},\mathbf{f},\mathbf{g})}\ell[\mathbf{Z},\{1-k(\mathbf{Z})\}(\boldsymbol{\theta}_1;\mathbf{f}_1,\mathbf{g}_1) + k(\mathbf{Z})(\boldsymbol{\theta}_2;\mathbf{f}_2,\mathbf{g}_2)]^\top\{(\boldsymbol{\theta}_1,\mathbf{f}_1,\mathbf{g}_1) - (\boldsymbol{\theta}_2,\mathbf{f}_2,\mathbf{g}_2)\}\end{aligned}$$

for some $k(\mathbf{Z}) \in (0,1)$. The Cauchy-Schwarz inequality now gives the Lipschitz-type property, with the integrability condition holding due to Assumption 5.E.5(vi). Hence Condition (A2) holds. Condition (A3) follows by the assumption that $\hat{\mathbf{g}}_n$ and $\hat{\mathbf{f}}_n$ are weakly consistent. Condition (A4) follows if $\{\ell(\cdot,\boldsymbol{\theta};\mathbf{f},\mathbf{g}) : \mathbf{f} \in \mathbb{F}, \mathbf{g} \in \mathbb{G}, \boldsymbol{\theta} \in \Theta\}$ is Glivenko–Cantelli. This is the case by Example 19.7 and Theorem 19.4 in Van der Vaart (1998) due to the Lipchitz-type property. Hence Condition (A4) follows. Condition (A5) also follows from the Lipchitz-type property and Assumption 5.E.5(ii). $\square$

We now also consider consistency of Efron's 'simple' nonparametric bootstrap introduced in Efron (1979) and described in Section 5.4. Based on the previous results, we immediately get weak consistency of the bootstrap estimator. The proposition is comparable with Proposition 1 in Hahn (1996).

Proposition 5.E.7. *Under Assumptions 5.2.1-5.2.3 and Assumptions 5.E.1-5.E.5, one has* $\|\hat{\mathbf{g}}_n^{\text{boot}} - \mathbf{g}_0\| = o_P(1)$, $\|\hat{\mathbf{f}}_n^{\text{boot}} - \mathbf{f}_0\| = o_P(1)$, *and* $\|\hat{\boldsymbol{\theta}}_n^{\text{boot}} - \boldsymbol{\theta}_0\| = o_P(1)$.

Proof. By the Glivenko–Cantelli property shown in Propositions 5.E.2-5.E.6, each of the function classes $\{\ell_\xi(\cdot, \mathbf{g}) : \mathbf{g} \in \mathbb{G}\}$, $\{\ell_U(\cdot, \mathbf{f}; \mathbf{g}) : \mathbf{f} \in \mathbb{F}, \mathbf{g} \in \mathbb{G}\}$, and $\{\ell(\cdot, \boldsymbol{\theta}; \mathbf{f}, \mathbf{g}) : \mathbf{f} \in \mathbb{F}, \mathbf{g} \in \mathbb{G}, \boldsymbol{\theta} \in \Theta\}$ converge uniformly $\sup_q \left\|n^{-1} \sum_{i=1}^n q(\mathbf{Z}_i^{\text{obs}}) - E[q(\mathbf{Z}^{\text{obs}})]\right\| = o_P(1)$, where q runs over the function class in question. Furthermore, we have shown that there exists an integrable envelope function for each of the classes. The bootstrap Glivenko–Cantelli theorem, Theorem 2.6 in Giné & Zinn (1990), then implies that

$$\sup_q \left\| n^{-1} \sum_{i=1}^n q(\mathbf{Z}_i^{\text{obs}}) - n^{-1} \sum_{i=1}^n q(\mathbf{Z}_{ni}^{\text{obs}}) \right\| = o_P(1)$$

almost surely. Consequently, the bootstrap estimators maximize their respective non-bootstrap empirical objective functions up to an $o_P(1)$-term, which is sufficient to invoke both Theorem 2.1 of Newey & McFadden (1994) and Theorem 1 of Delsol & Van Keilegom (2020). □

A different approach to proving Proposition 5.E.7 would be to verify the conditions of Theorem 3.5 in Arcones & Giné (1992).

5.E.2 Asymptotic Normality

We now consider the stacked estimating equation

$$n^{-1} \sum_{i=1}^n \mathbf{h}(\mathbf{Z}_i^{\text{obs}}, \boldsymbol{\theta}, \mathbf{f}, \mathbf{g}) = \mathbf{0}$$

for $\mathbf{h}(\mathbf{Z}^{\text{obs}}, \boldsymbol{\theta}, \mathbf{f}, \mathbf{g}) = \{\frac{\mathrm{d}}{\mathrm{d}\mathbf{g}} \ell_\xi(\mathbf{Z}^{\text{obs}}, \mathbf{g}), \frac{\mathrm{d}}{\mathrm{d}\mathbf{f}} \ell_U(\mathbf{Z}, \mathbf{f}; \mathbf{g}), \frac{\mathrm{d}}{\mathrm{d}\boldsymbol{\theta}} \ell(\mathbf{Z}, \boldsymbol{\theta}; \mathbf{f}, \mathbf{g})\}^\top$. For the nuisance parameters, define $\mathbf{k}(\mathbf{Z}^{\text{obs}}, \mathbf{f}, \mathbf{g}) = \{\frac{\mathrm{d}}{\mathrm{d}\mathbf{g}} \ell_\xi(\mathbf{Z}^{\text{obs}}, \mathbf{g}), \frac{\mathrm{d}}{\mathrm{d}\mathbf{f}} \ell_U(\mathbf{Z}, \mathbf{f}; \mathbf{g})\}^\top$. The existence of the scores is ensured by Assumption 5.E.5(vi) and Assumption 5.E.8(i). Introduce also

$$\begin{aligned} \mathbf{H}_1 &= \mathbb{E}\{\nabla_{\boldsymbol{\theta}} \mathbf{h}(\mathbf{Z}^{\text{obs}}, \boldsymbol{\theta}_0, \mathbf{f}_0, \mathbf{g}_0)\}, & \mathbf{H}_2 &= \mathbb{E}\{\nabla_{(\mathbf{f},\mathbf{g})} \mathbf{h}(\mathbf{Z}^{\text{obs}}, \boldsymbol{\theta}_0, \mathbf{f}_0, \mathbf{g}_0)\}, \\ \mathbf{K} &= \mathbb{E}\{\nabla_{(\mathbf{f},\mathbf{g})} \mathbf{k}(\mathbf{Z}^{\text{obs}}, \mathbf{f}_0, \mathbf{g}_0)\}, & \boldsymbol{\psi}(\mathbf{Z}^{\text{obs}}) &= -\mathbf{K}^{-1} \mathbf{k}(\mathbf{Z}^{\text{obs}}, \mathbf{f}_0, \mathbf{g}_0). \end{aligned}$$

Assumption 5.E.8.

(i) There exists a neighborhood of the true parameter denoted $(\Theta \times \mathbb{F} \times \mathbb{G})_\delta = \{(\boldsymbol{\theta}, \mathbf{f}, \mathbf{g}) : \|(\boldsymbol{\theta}, \mathbf{f}, \mathbf{g}) - (\boldsymbol{\theta}_0, \mathbf{f}_0, \mathbf{g}_0)\| \leq \delta\}$ *such that both* $\frac{\mathrm{d}}{\mathrm{d}\mathbf{g}} \ell_\xi(\mathbf{Z}^{\text{obs}}, \mathbf{g})$ *and* $\frac{\mathrm{d}}{\mathrm{d}\mathbf{f}} \ell_U(\mathbf{Z}, \mathbf{f}; \mathbf{g})$ *exist and such that* $\mathbb{E}\{\mathbf{h}(\mathbf{Z}^{\text{obs}}, \boldsymbol{\theta}, \mathbf{f}, \mathbf{g})\} = \mathbf{0}$ *has a unique solution in this neighborhood.*

(ii) The score $\mathbf{h}(\mathbf{Z}^{\text{obs}}, \boldsymbol{\theta}, \mathbf{f}, \mathbf{g})$ *is dominated by a square-integrable function of* $\mathbf{Z}^{\text{obs}}$.

(iii) The gradient of the score $\nabla_{(\boldsymbol{\theta}, \mathbf{f}, \mathbf{g})} \mathbf{h}(\mathbf{Z}^{\text{obs}}, \boldsymbol{\theta}, \mathbf{f}, \mathbf{g})$ *exists and is continuous in the parameters. Furthermore, it is dominated by a square-integrable function of* $\mathbf{Z}^{\text{obs}}$.

(iv) $(\mathbf{H}_1, \mathbf{H}_2)^T (\mathbf{H}_1, \mathbf{H}_2)$ *is nonsingular.*

We may now show asymptotic normality. For the ordinary estimator, one could proceed via Theorem 6.1 in Newey & McFadden (1994), but we instead use Theorem 1 of Hahn (1996) as this also gives asymptotic normality of the bootstrap estimator. Similar conditions to those of Hahn (1996) are found in Arcones & Giné (1992).

Proposition 5.E.9. *Under Assumptions 5.2.1-5.2.3 and Assumptions 5.E.1-5.E.8, it holds that*

(i) The ordinary estimators $\hat{\mathbf{g}}_n$, $\hat{\mathbf{f}}_n$, *and* $\hat{\boldsymbol{\theta}}_n$ *are asymptotically normal.*

(ii) $\sqrt{n}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) \to N(0, \mathbf{V})$ *in distribution for*

$$V = \mathbf{H}_1^{-1} \mathbb{E}[\{\mathbf{h}(\mathbf{Z}^{\text{obs}}, \boldsymbol{\theta}_0, \mathbf{f}_0, \mathbf{g}_0) + \mathbf{H}_2 \boldsymbol{\psi}(\mathbf{Z}^{\text{obs}})\} \\ \times \{\mathbf{h}(\mathbf{Z}^{\text{obs}}, \boldsymbol{\theta}_0, \mathbf{f}_0, \mathbf{g}_0) + \mathbf{H}_2 \boldsymbol{\psi}(\mathbf{Z}^{\text{obs}})\}^\top](\mathbf{H}_1^{-1})^\top.$$

(iii) The bootstrap estimators $\hat{\mathbf{g}}_n^{\text{boot}}$, $\hat{\mathbf{f}}_n^{\text{boot}}$, *and* $\hat{\boldsymbol{\theta}}_n^{\text{boot}}$ *are asymptotically normal and, in particular,* $\sqrt{n}(\hat{\boldsymbol{\theta}}_n^{\text{boot}} - \hat{\boldsymbol{\theta}}_n) \to N(0, \mathbf{V})$ *in distribution.*

Proof. Due to Assumption 5.E.8(i), our estimator can be formulated as a generalized method of moment estimator with identity weighting matrix and parameter space $(\Theta \times \mathbb{F} \times \mathbb{G})_\delta$. Theorem 1 of Hahn (1996) therefore becomes applicable. Condition (i) is satisfied since a) the score equation has the true parameters as its unique solution by Assumption 5.E.8(i), b) our data $(\mathbf{Z}_i^{\text{obs}})_{i=1}^n$ is assumed to be independent and identically distributed, c) the solution is well-separated by standard arguments since $(\Theta \times \mathbb{F} \times \mathbb{G})_\delta$ is compact and $\mathbb{E}\{\mathbf{h}(\mathbf{Z}^{\text{obs}}, \boldsymbol{\theta}, \mathbf{f}, \mathbf{g})\}$ is continuous in the parameters by dominated convergence and Assumptions 5.E.8(ii)-(iii), d) Glivenko–Cantelli holds for the class $\{\mathbf{h}(\cdot, \boldsymbol{\theta}; \mathbf{f}, \mathbf{g}) : (\boldsymbol{\theta}, \mathbf{f}, \mathbf{g}) \in (\Theta \times \mathbb{F} \times \mathbb{G})_\delta\}$ by Example 19.8 of Van der Vaart (1998) since $\mathbf{h}(\mathbf{Z}^{\text{obs}}, \boldsymbol{\theta}, \mathbf{f}, \mathbf{g})$ is continuous in the parameters by Assumption 5.E.8(iii) and has an integrable envelope function by Assumption 5.E.8(ii), e) there is an integrable envelope function for the class as noted before, and f) the weighting matrix is equal to the identity weighting matrix. Condition (ii) holds by specification of the estimators. Condition (iii) holds by Assumptions 5.E.8(ii)-(iii) via dominated convergence, since these assumptions imply that $\mathbf{h}$ is continuous in the parameters and that the integrand is dominated by an integrable function, just use $(a-b)^2 \leq 2(a^2+b^2)$. Note that an exponent of 2 seems to be missing in the Hahn (1996) paper; for comparison, see Theorem 21 (functional central limit theorem) in Chapter VII of Pollard (1984). Condition (iv) holds if the class $\{\mathbf{h}(\cdot,\boldsymbol{\theta},\mathbf{f},\mathbf{g}) : (\boldsymbol{\theta},\mathbf{f},\mathbf{g}) \in (\Theta \times \mathbb{F} \times \mathbb{G})_\delta\}$ is Donsker since Condition (iii) holds at any parameter value, not just the true value. This can be shown via Example 19.7 and Theorem 19.5 in Van der Vaart (1998) in the same way as how the Lipschitz-type property was shown in the proof of Proposition 5.E.6: It is sufficient that $\mathbf{h}$ is differentiable in the parameters and that the norm of the gradient is dominated by a square-integrable function, and these properties are ensured by Assumption 5.E.8(iii). Condition (v) holds by Assumptions 5.E.8(iii)-(iv). The first part of Condition (vi) holds by Assumption 5.E.8(ii) and the second part holds by Theorem 3.1 of Andersen & Dobric (1987) and their following remark when setting $p=2$ and with their function f being the score $\mathbf{h}$. This is due to the following observations. Their equation (3.1.2) holds by Assumption 5.E.8(ii) and the triangle inequality, their equation (3.1.3) holds since the scores are both bounded and uniformly continuous with probability one by Assumptions 5.E.8(ii)-(iii) combined with the Heine-Cantor theorem, and the parameter set is compact (and hence totally bounded) with respect to the original Euclidean metric.

The variance matrix calculation is now identical to that of Theorem 6.1 in Newey & McFadden (1994), which leads to the desired result. □

Theorem 5.4.1 is now a consequence of Propositions 5.E.2-5.E.9. Note that the proof can be adapted to the case where one uses the optimal weight matrix.

5.F Additional Details for Numerical Study

5.F.1 Data-generating Process

Let the marked point process generated by the rates $\tilde{\mu}_{jk}$ be denoted $(T^\dagger_m, Y^\dagger_m)_{m\geq 1}$. Deletion of unreported events from $(T^\dagger_m, Y^\dagger_m)_{m\geq 1}$ leads to $(\tilde{T}_m, \tilde{Y}_m)_{m\geq 1}$, while deleting events that will be rejected leads to $(T^*_m, Y^*_m)_{m\geq 1}$. The data is generated according to the following algorithm:

(i) Simulate X, V, and C.

(ii) Simulate events in $(V, C]$ using the rates $\tilde{\mu}_{jk}$ and initial state 1.

(iii) For jumps to state 3, simulate corresponding reporting delays with a distribution that depends on whether or not a transition to state 2 has occurred previously.

(iv) For transitions from state 2 to state 3, simulate adjudication events in $(T_2^\dagger + U_2^\dagger, \eta]$ if this interval is nonempty using the rates ω_{jk} and initial state 1.

Events that are reported after η are flagged and deleted except for when implementing oracle methods. The adjudication outcome $\xi_{\langle \eta \rangle}$ is simulated from a Bernoulli distribution with success probability equal to its expectation given $\mathcal{H}_\eta$, and events with adjudication outcome zero are deleted when implementing oracle methods. For simulation of the reporting delays, we note that if $U \sim \mathrm{Weibull}(\lambda, k, \beta)$, then $[1 - \exp\{-(\lambda U)^k\}]^{\exp(X\beta)} = W$ in distribution, where $W \sim \mathrm{Unif}(0,1)$. Trivial calculations then show that $U = \lambda^{-1} \times [-\log\{1 - W^{\exp(-X\beta)}\}]^{1/k}$ in distribution. Contour plots of the hazards $\tilde{\mu}_{jk}$, reporting delay reverse time hazards on log-scale, and adjudication hazards are depicted in Figure 5.7.

To obtain an expression for μ_{23}^*, we first let $p(X)$ be the probability that a newly reported jump becomes confirmed. We have

$$\begin{aligned} p(X) &= \left(1 - \exp\left[-\int_0^\infty g_1 \times \{X/(v+2)\}^2 \,\mathrm{d}v\right]\right) \\ &\quad \times \left[1 - \exp\left\{-\int_0^\infty \exp(g_2 \times v)\,\mathrm{d}v\right\}\right] \\ &= \left\{1 - \exp\left(-g_1 \times X^2/2\right)\right\}\{1 - \exp(1/g_2)\}. \end{aligned}$$

Note that $\tilde{\mu}_{23}(t; \tilde{T}_1, X)$ only depends on t and $\tilde{T}_1$ through $t - \tilde{T}_1$, which equals V_t due to the lack of reporting delays for the transition to state 2, so we write $\tilde{\mu}_{23}(t; \tilde{T}_1, X) = \tilde{\mu}_{23}(V_t; X)$. Using that thinning $(T_m^\dagger, Y_m^\dagger)_{m \geq 1}$ with the adjudication outcomes leads to $(T_m^*, Y_m^*)_{m\geq 1}$, similar calculations to the ones found in Section 5.B of the Supplementary material yield

$$\mathbb{P}(T_2^* \geq t \mid T_1^*, Y_1^* = 2) = 1 - p(X)\left[1 - \exp\left\{\frac{1 - \exp(\theta_7 V_t X^2)}{\theta_7 X^2}\right\}\right]$$

and hence by Jacod's formula for the intensity, see Proposition (3.1) of Jacod (1975),

$$\mu_{23}^*(t, V_t; X, \boldsymbol{\theta}) = \frac{p(X)\exp[\{1 - \exp(\theta_7 V_t X^2)\}/(\theta_7 X^2)]\tilde{\mu}_{23}(V_t; X)}{1 - p(X)(1 - \exp[\{1 - \exp(\theta_7 V_t X^2)\}/(\theta_7 X^2)])}.$$

In order to speed up the computation time when calculating $\hat{\boldsymbol{\theta}}_n$, we use that

$$\begin{aligned} \int \mu_{1\bullet}^*(s; \emptyset, \boldsymbol{\theta})\,\mathrm{d}s &= \{\pi/(4\theta_5)\}^{1/2} \exp\{\theta_4 + \theta_6 \times \cos(0.5\pi X)\} \times \mathrm{erfi}(\theta_5^{1/2} \times s) \\ &\quad + \theta_2^{-1} \times \exp\{\theta_1 + \theta_2 \times X + \theta_3 \times \sin(0.5\pi X)\} \times \exp(\theta_2 \times s) \end{aligned}$$

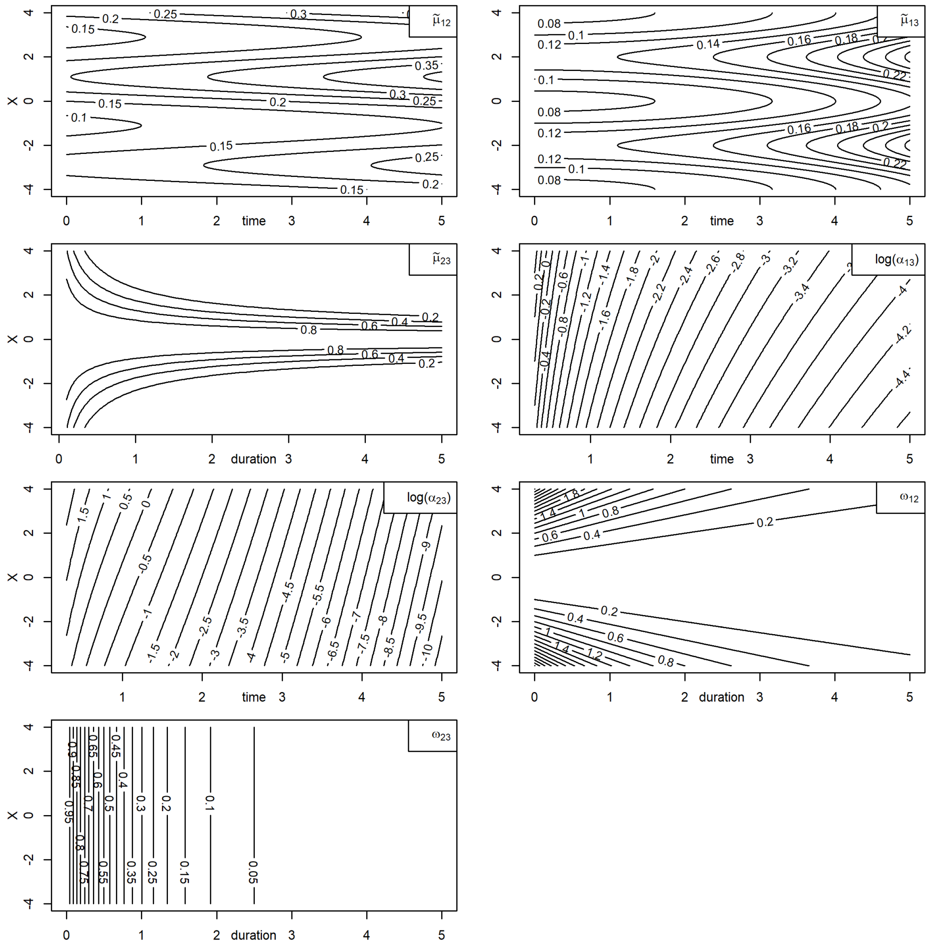

Figure 5.7: *Contour plots of the hazards from the numerical study. The first axis is time or duration depending on whether the hazard depends on the time or the duration in the current state; none of the hazards depend on both.*

to calculate $Y_1(t)P^*\{t; \mathbf{J}(t-), \boldsymbol{\theta}\}$ efficiently, where erfi is the imaginary error function. This function is implemented in many software packages, but we instead use the following 15'th order Taylor expansion around 0 to further speed up calculations: $\mathrm{erfi}(z) = \pi^{-1/2} \times (2z + \frac{2}{3}z^3 + \frac{1}{5}z^5 + \frac{1}{21}z^7 + \frac{1}{108}z^9 + \frac{1}{660}z^{11} + \frac{1}{4680}z^{13} + \frac{1}{37800}z^{15}) + O(z^{17})$. The approximation error for $Y_1(t)P^*\{t; \mathbf{J}(t-), \boldsymbol{\theta}\}$ was found to be of the order 10^{-7} for parameter choices near the true values, hence negligible. We further note that simple calculations gave analytical expressions for the exposure-terms of the log-likelihoods for $\mathbf{g}$ and $\mathbf{f}$ as well as for the Poisson approximation to the transition from state 1 to state 2.

Table 5.3: *Bias, empirical standard deviation (SD), and root mean squared error (RMSE) of the estimators* $\hat{\mathbf{g}}_n$ *and* $\hat{\mathbf{f}}_n$ *based on 400 simulations of size* $n = 1500$.

Parameter	Bias	SD	RMSE
$g_1 = 0.8$	-.008	.104	.105
$g_2 = -1.2$	-.035	.294	.296
$f_1 = 2$	.010	.465	.465
$f_2 = 0.5$	.004	.037	.037
$f_3 = 0.1$	-.000	.032	.032
$f_4 = 1$	-.001	.089	.089
$f_5 = 1.5$	.040	.170	.174
$f_6 = 0.2$	.004	.060	.060

Table 5.4: *Bias and root mean squared error (RMSE) of the estimator* $\hat{\boldsymbol{\theta}}_n$ *based on 400 simulations of size* $n = 1500$.

	Proposed method		Oracle		Poisson approx.		Naive 1		Naive 2	
Parameter	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE
$\theta_1 = \log 0.15$	-.004	.067	-.008	.033	-.010	.067	-.011	.067	-.010	.067
$\theta_2 = 0.1$	-.000	.020	-.001	.020	-.006	.021	-.006	.021	-.006	.021
$\theta_3 = 0.4$	.003	.078	.003	.078	-.002	.078	-.002	.078	-.000	.079
$\theta_4 = \log 0.1$	.003	.083	.001	.083	.012	.092	-.041	.087	-.051	.096
$\theta_5 = 0.03$	.000	.012	-.000	.013	-.006	.017	-.018	.022	-.015	.021
$\theta_6 = -0.3$	-.000	.094	-.001	.088	.007	.094	-.009	.087	-.007	.090
$\theta_7 = -0.3$	-.011	.067	-.011	.055	-.012	.067	.157	.158	.148	.163

5.F.2 Results

In Table 5.3 we report the bias, the SD, and the RMSE of the parameter estimators of $\mathbf{g}$ and $\mathbf{f}$. In Table 5.4 we report the bias and the RMSE of the estimators of $\boldsymbol{\theta}$. The findings for SD and RMSE are highly comparable. Histograms of the obtained estimators are shown in Figure 5.8 and generally show approximate Gaussianity with varying degrees of skewness. The bootstrap results are shown in Table 5.5, and are based on $k = 50, 100, 200, 300, 400$ estimates of θ_7 and 1000 bootstrap resamples. We report the coverage of the confidence bands for several values of the number of bootstrap runs k and the confidence level $1 - \alpha$.

5.G Additional Details for Data Application

In Table 5.7 and Table 5.6, we present the parameter estimates and percentile bootstrap confidence intervals for $\mathbf{f}$ and $\mathbf{g}$, respectively. The confidence intervals are computed using 400 bootstrap resamples. The variable 'rejected before' is an indicator of the event that adjudication state 2 has been entered previously. This variable is trivial in adjudication state 2 and hence not included in that

Table 5.5: *Coverage (%) of percentile bootstrap confidence bands for θ_7 with $k = 50, 100, 200, 300, 400$ as well as $1-\alpha = 0.90, 0.95, 0.99$.*

k	Coverage (%)		
	90	95	99
50	91.8	93.9	95.9
100	87.9	91.9	97.0
200	87.9	93.0	98.0
300	90.0	95.0	98.3
400	89.5	95.0	98.7

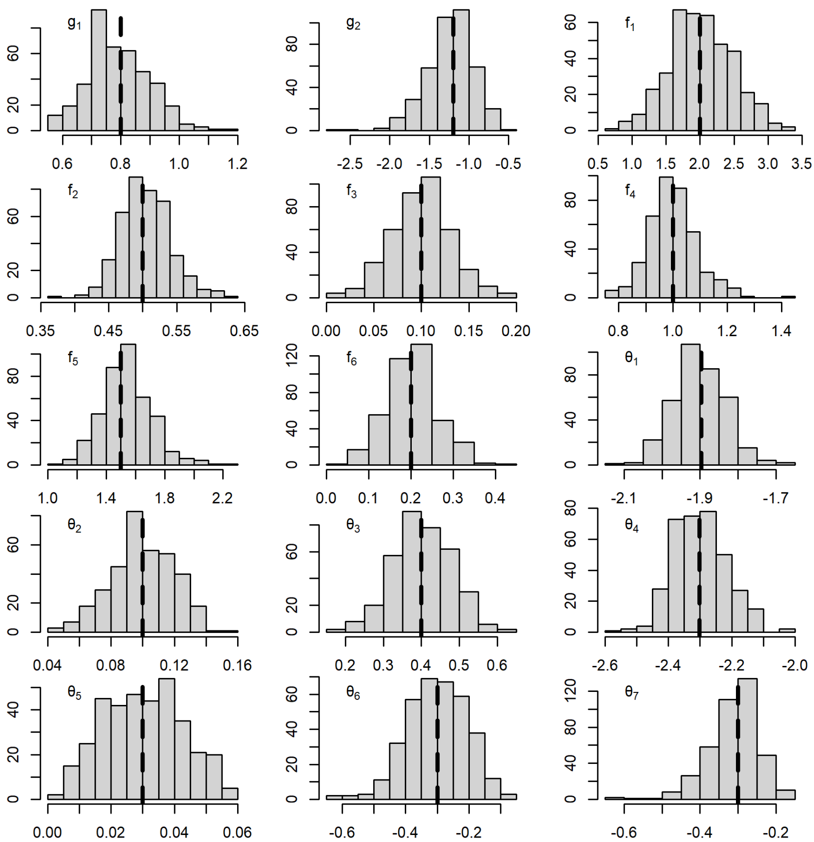

Figure 5.8: *Histograms of $\hat{\mathbf{g}}_n$, $\hat{\mathbf{f}}_n$, and $\hat{\boldsymbol{\theta}}_n$ based on 400 samples of size $n = 1500$ with the true values indicated by dashed lines.*

regression. There are no deaths observed from adjudication state 2 for disablements, and we hence estimate no parameters for this transition, letting the hazard be identically zero. Similarly, only one death is observed from adjudication state 1 for reactivations, causing the Poisson regression algorithm to diverge, and we hence also let this hazard be identically zero. One could of course have changed the regression model to a simple intercept model and obtained a small positive hazard, but its effect on the adjudication probabilities would be negligible and we hence forego this modification.

		ω_{12}		ω_{13}		ω_{14}		ω_{21}		ω_{24}	
Parameter	Type	Est.	CI	Est.	CI	Est.	CI	Est.	CI	Est.	CI
Age	Reactivation	.005	(-0.010, 0.018)	.015	(0.007, 0.026)	-	-	.000	(-0.011, 0.010)	.023	(-0.059, 0.105)
Male	Reactivation	.973	(0.295, 1.63)	1.29	(0.752, 1.76)	-	-	-.208	(-0.703, 0.295)	-5.94	(-10.45, -2.87)
Female	Reactivation	1.18	(0.524, 1.84)	.771	(0.159, 1.31)	-	-	-.053	(-0.793, 0.626)	-21.10	(-26.24, -18.17)
Disability duration	Reactivation	-.065	(-0.247, 0.153)	-.198	(-0.380, 0.005)	-	-	.206	(0.101, 0.325)	-.464	(-1.86, 0.136)
Reactivation duration	Reactivation	-.042	(-0.273, 0.205)	-.490	(-0.803, -0.200)	-	-	-1.25	(-1.49, -1.05)	.148	(-1.29, 1.66)
Age	Disability	-.005	(-0.015, 0.005)	.001	(-0.003, 0.004)	.013	(-0.054, 0.101)	.003	(-0.008, 0.014)	-	-
Male	Disability	-.800	(-1.27, -0.310)	1.06	(0.870, 1.25)	-4.85	(-10.04, 4.52)	1.16	(0.578, 1.74)	-	-
Female	Disability	-.806	(-1.30, -0.316)	1.19	(0.988, 1.41)	-5.29	(-22.06, 3.17)	1.08	(0.353, 1.71)	-	-
Disability duration	Disability	.077	(-0.190, 0.315)	-.234	(-0.343, -0.129)	.169	(-1.45, 1.23)	-.255	(-0.610, 0.063)	-	-
Report duration	Disability	-.250	(-0.802, 0.280)	.470	(0.307, 0.680)	-.532	(-211.85, 2.12)	-.090	(-0.525, 0.413)	-	-
Rejected before	Disability	.982	(0.572, 1.38)	-.118	(-0.360, 0.072)	-14.68	(-16.70, -11.75)	-	-	-	-

Table 5.6: *Parameter estimates (Est.) and 95% bootstrap percentile confidence interval (CI) for the parameter* $\mathbf{g}$ *using the proposed method and 400 bootstrap resamples.*

Table 5.7: *Parameter estimates (Est.) and 95% bootstrap percentile confidence interval (CI) for the parameter* $\mathbf{f}$ *using the proposed method and 400 bootstrap resamples.*

Parameter	Est.	CI
λ	2.23	(0.964, 2.78)
k	1.05	(0.931, 2.00)
Male	.356	(-0.518, 0.682)
Female	.113	(-0.672, 0.783)
Age at disability	.001	(-0.027, 0.006)

Chapter 6

Doubly robust inference with censoring unbiased transformations

This chapter contains the manuscript *Sandqvist (2024)*.

Abstract

This paper extends doubly robust censoring unbiased transformations to a broad class of censored data structures under the assumption of coarsening at random and positivity. This includes the classic survival and competing risks setting, but also encompasses multiple events. A doubly robust representation for the conditional bias of the transformed data is derived. This leads to rate double robustness and oracle efficiency properties for estimating conditional expectations when combined with cross-fitting and linear smoothers. Simulation studies demonstrate favourable performance of the proposed method relative to existing approaches. An application of the methods to a regression discontinuity design with censored data illustrates its practical utility.

Keywords: Censored data; Conditional effects; Machine learning; Nonparametric regression; Pseudo-values; Regression discontinuity design

6.1 Introduction

In many situations, one is interested in modeling the effect of covariates W on a variable Y. Powerful regression methods $\hat{\mathbb{E}}_n[Y \mid W = w]$ based on i.i.d. observations $(W_1, Y_1), \ldots, (W_n, Y_n)$ allows for flexible ways to estimate such effects without imposing strong parametric assumptions. Examples include local polynomial regression, neural networks, and tree-based methods. If $Y_1, \ldots, Y_n$ are not fully observed due to censoring or other coarsening mechanisms, it is not possible to form the regression estimator $\hat{\mathbb{E}}_n[Y \mid W = w]$ directly, and sub-sampling on complete-case data might lead to substantial biases in the estimates. In this paper, the observations are incomplete due to censoring, so one is interested in an outcome on the form $Y = Y(X)$ for $X = \{X(t)\}_{t\geq 0}$ but only $X^C = \{X(C \wedge t)\}_{t\geq 0}$ and the censoring time C are observed.

Instead of tailoring regression methods to censored data, one may construct pseudo-outcomes $Y^* = Y^*(C, X^C)$ by transforming (C, X^C) using a censoring unbiased transformation (CUT) and then run the regression method on the pseudo-outcomes; see Rubin & Van der Laan (2007) for an overview of different censoring unbiased transformations. The term "unbiased" refers to the fact that the pseudo-outcomes should satisfy $\mathbb{E}[Y^* \mid W = w] = \mathbb{E}[Y \mid W = w]$. A doubly robust censoring unbiased transformation (DRCUT) was first introduced in Rubin & Van der Laan (2007) for survival data $X(t) = \{W, 1_{(T\leq t)}\}$ with $Y(X) = T$ under independent censoring $C \perp\!\!\!\perp T \mid W$. This transformation depends on both the conditional mean outcome $\mathbb{E}[T \mid W = w, T > t]$ and the conditional censoring distribution $\mathbb{P}(C \leq t \mid W = w)$ but has the upside that it gives the correct conditional mean if just one of these is correctly specified. A generalization of this transformation has recently been introduced in Steingrimsson et al. (2016). They generalize the DRCUT of Rubin & Van der Laan (2007) to arbitrary $Y(X)$ for survival data. DRCUTs for censored data have to the author's knowledge not been explored outside of the survival data setting.

The main contributions of this paper are as follows. The DRCUT and regression discontinuity design (RDD) methodologies are generalized to any censored data satisfying coarsening at random and positivity, extending results from the survival setting. A doubly robust representation of the conditional bias of the DRCUT is obtained. Using this representation and building on the framework of Kennedy (2023), large sample properties of DRCUT-based estimators are established, including rate double robustness and oracle efficiency results. In passing, the analysis of Kennedy (2023) is extended from sample-splitting to cross-fitting and results on cross-fitting are extended to estimands that converge slower than $\sqrt{n}$.

Double robustness was initially discovered as a property of efficient-influence-function-based estimators for nonparametric and semiparametric models in cases where the data-generating process is affected by a missingness mechanism. Such estimators have been explored extensively for censored data, where the missingness mechanism stems from the fact that subjects are unobserved after a random censoring time, and in causal inference, where the missingness mechanism stems from the fact that not all potential outcomes are observed, confer with Section 6.6 in Bickel et al. (1998), Van der Laan & Robins (2003), Bang & Robins (2005), and Van der Laan & Rose (2011). For estimation in nonparametric models via efficient influence functions more broadly, see the recent reviews Kennedy (2022) and Hines et al. (2022).

As noted in Section 5.3 of Kennedy (2022), when the estimand is not an expectation but rather a regression function i.e. a conditional expectation, the efficient influence function does generally not exist, and existing theory on efficient estimation hence cannot be applied. One way to overcome this limitation is explored in Kennedy (2023) for the conditional average treatment effect (CATE) in a potential outcome setting with no censoring. The idea is that semiparametrically efficient estimators estimate a marginal estimand $\mathbb{E}[Y]$ by averaging over the uncentered efficient influence function of $\mathbb{E}[Y]$, so a conditional estimand $\mathbb{E}[Y \mid W = w]$ may be estimated by regressing the uncentered efficient influence function of $\mathbb{E}[Y]$ on covariates. One may recognize the DRCUTs of Rubin & Van der Laan (2007) and Steingrimsson et al. (2016) as uncentered efficient influence functions, making the efficient influence function of $\mathbb{E}[Y]$ a natural candidate for a DRCUT in the situation where only coarsening at random and positivity are imposed, and this paper proves that this is indeed a DRCUT.

To facilitate desirable large sample properties, Kennedy (2023) suggested a sample splitting approach, first estimating nuisance parameters needed for the transformation on one sample and then regressing pseudo-outcomes on covariates in the second sample. A motivation for this approach is that the conditional bias of the pseudo-outcomes has a product structure. Sample splitting allows one to exploit this product structure to show that estimators based on these pseudo-outcomes have a rate double robustness property, meaning that the convergence rate depends on the product of the convergence rate of each of the nuisance estimators. A main result of this paper is to show that a similar product structure emerges for DRCUTs and to exploit this to derive desirable large sample properties of the DRCUT estimators. Large sample properties of DRCUT estimators have not been investigated in the literature, but they are essential when the goal is inference rather than prediction.

Pseudo-outcomes which are based on the jack-knife rather than the efficient influence function have also been explored for survival and competing risks data. The jack-knife pseudo-outcomes were introduced in Andersen et al. (2003) and is an area of continued study, see for example Jacobsen & Martinussen (2016), Andersen et al. (2017), Overgaard et al. (2017), and Parner et al. (2023). The latter paper proposes to use so-called infinitesimal jack-knife pseudo-outcomes, which turn out to be exactly the uncentered efficient influence function for $\mathbb{E}[Y]$. The motivation in that paper is that the jack-knife pseudo-outcomes were observed to be asymptotically equivalent to the infinitesimal jack-knife pseudo-outcomes under suitable regularity conditions, and the latter were found to be much faster to compute. Infinitesimal jack-knife pseudo-outcomes and DRCUT pseudo-outcomes are thus identical, which seems to have been overlooked by Parner et al. (2023).

It is fruitful to make this connection; an assumption that has been persistent in all literature on jack-knife pseudo-outcomes is that the censoring is completely random $C \perp\!\!\!\perp (T, W)$. This has been highlighted as a key assumption, but also as perhaps the most restrictive one, see e.g. Section 4 in Overgaard et al. (2017). Instead of taking the jack-knife pseudo-outcomes as a starting point, the DRCUT, or equivalently the infinitesimal jack-knife pseudo-outcomes, is here taken as the starting point. With this point of view, it is not difficult to allow the censoring distribution to depend on covariates and the past trajectory of X since the efficient influence function is still well-known in this situation. Thus, this key independence assumption is relaxed substantially.

The nuisance parameters become more complicated when they are allowed to depend on covariates and the past trajectory for X. When they are simple e.g. belonging to a Donsker class, they may usually be estimated in-sample without adding asymptotic variance which is the case that has been explored in the jack-knife pseudo-outcome literature so far, see e.g. Overgaard et al. (2017) and Parner et al. (2023). This paper proposes a K-fold cross-fitting approach to allow for inference in the presence of flexible nuisance estimators e.g. depending on hyperparameters selected in data-adaptive ways.

The paper is structured as follows. In Section 6.2, the DRCUT is proposed and a doubly robust representation for its conditional bias is presented. In Section 6.3, the asymptotics of DRCUT-based estimators are studied. In Section 6.4, simulation experiments are conducted which show favorable performance of the proposed method relative to competing approaches. In Section 6.5, the approach is illustrated via an application to the Longitudinal Study of Young People in England (LSYPE) in the form of a regression discontinuity design (RDD) to infer a CATE. All implementations are available on GitHub at `https://github.com/oliversandqvist/Web-appendix-drcut`.

6.2 Doubly robust censoring unbiased transformation

Let the background probability space be $(\Omega, \mathcal{F}, \mathbb{P})$. The full data is the stochastic process $X = \{X(t)\}_{t \geq 0}$ indexed by a time-variable $t \in [0, \infty)$, which can refer to calendar time, time since study entry etc. Assume that X is a stochastic process on a metric space $\mathbb{D}$ equipped with its Borel σ-algebra $\mathcal{B}(\mathbb{D})$ and that the sample paths $t \mapsto X(t)(\omega)$ are càdlàg. One can thus alternatively consider X as a map from Ω into $\mathcal{X} = D([0, \infty), \mathbb{D})$ being the set of all functions $z : [0, \infty) \mapsto \mathbb{D}$ that are càdlàg. This space is equipped with the projection σ-algebra generated by the projection maps $z \mapsto z(t)$ from $\mathcal{X}$ to $\mathbb{D}$ which makes X a measurable map from Ω into $\mathcal{X}$.

The process X stopped at a time $u \geq 0$ is denoted $X^u = \{X(u \wedge t)\}_{t \geq 0}$ and the censoring variable is denoted $C : \Omega \mapsto (0, \infty)$, so the observed data is (C, X^C). It is assumed that $X(0)$ contains some baseline covariates denoted W. The outcome of interest is $Y(X)$ with $Y : \mathcal{X} \mapsto \mathbb{R}$. Vector-valued outcomes can be accommodated by applying each result coordinate-wise. It is assumed throughout that Y is square integrable. In order for Y to be nonparametrically identifiable from the observed data, it is assumed that $Y(X) = \tilde{Y}(X^\eta)$ for a suitable function $\tilde{Y}$. The approach in this paper fails when Y is not nonparametrically identifiable since the DRCUT utilizes inverse probability of censoring weighting. In such cases, one might be able to use the proposed methods to estimate the identifiable part of the estimand and use other methods to estimate what remains.

6.2.1 Transformation

Let $r(c \mid x)$ denote the conditional density of C given X with respect to a fixed reference measure μ, which in most applications will be the Lebesgue measure λ. For identifiability, it is standard to impose the assumption that the observed data are a *coarsening at random* (CAR) of the full data. CAR extends the concept of missing at random to situations where a many-to-one function of the complete data is observed. For right-censored counting processes, CAR is closely related to independent censoring defined as in Andersen et al. (1993), see Gill et al. (1997) or Lemma 1 in Munch et al. (2023). Here, the CAR formulation from Van der Vaart (2004) is used.

Assumption 6.2.1. *(Coarsening at random.)*
There is a measurable function $\tilde{r} : [0,\infty) \times \mathcal{X} \mapsto [0,\infty)$ such that

$$r(c \mid x) = \tilde{r}(c, x^c)$$

for $x^c = \{x(c \wedge t)\}_{t \geq 0}$ with $x \in \mathcal{X}$. $\diamond$

Assume that there is a positive probability of observing the full data for any realization of X. This is a standard assumption for nonparametric estimation with censored data.

Assumption 6.2.2. *(Positivity.)*
It holds that $\mathbb{P}(C \geq \eta \mid X) \geq \epsilon > 0$ for some deterministic ϵ. $\diamond$

In the Supplementary material, it is shown that the efficient influence function of $\mathbb{E}[Y(X)]$ is

$$\begin{aligned}\mathrm{IF}(C, X^C) = &-\mathbb{E}[Y(X)] + \frac{Y(X)1_{(C\geq\eta)}}{\mathbb{P}(C \geq \eta \mid X)} \\ &+ \int_{[0,\eta)} \frac{\mathbb{E}[Y(X) \mid X^u]}{\mathbb{P}(C > u \mid X)} \left\{ \mathrm{d}1_{(C\leq u)} - 1_{(C\geq u)} \frac{\mathbb{P}(C \in \mathrm{d}u \mid X)}{\mathbb{P}(C \geq u \mid X)} \right\}.\end{aligned}$$

Note that $\mathbb{E}[Y(X)]$ is the population version of the estimand of interest $\mathbb{E}[Y(X) \mid W]$. This motivates the DRCUT in Theorem 6.2.3. Define $\gamma(u \mid X) = r(u \mid X)/\mathbb{P}(C \geq u \mid X)$ which is the Radon-Nykodym derivative of the hazard measure for $C \mid X$ with respect to μ under $\mathbb{P}$, and define γ_1 similarly but under a different measure $\mathbb{P}_1$.

Theorem 6.2.3. *(Doubly robust censoring unbiased transformation.)*
Let $\mathbb{P}_1$ and $\mathbb{P}_2$ be two probability measures, which may be thought of as candidate measures for $\mathbb{P}$. Let

$$
\begin{aligned}
Y^*_{\mathbb{P}_1,\mathbb{P}_2}(C, X^C) = &\frac{Y(X)1_{(C\geq\eta)}}{\mathbb{P}_1(C \geq \eta \mid X)} \\
&+ \int_{[0,\eta)} \frac{\mathbb{E}_2[Y(X) \mid X^u]}{\mathbb{P}_1(C > u \mid X)} \left\{ \mathrm{d}1_{(C\leq u)} - 1_{(C\geq u)} \frac{\mathbb{P}_1(C \in \mathrm{d}u \mid X)}{\mathbb{P}_1(C \geq u \mid X)} \right\}.
\end{aligned}
$$

When $\mathbb{P}$ and $\mathbb{P}_1$ satisfy Assumption 6.2.1 and 6.2.2 it holds that

$$
\begin{aligned}
&\mathbb{E}[Y^*_{\mathbb{P}_1,\mathbb{P}_2}(C, X^C) - Y(X) \mid W] \\
&= \mathbb{E}\Big[\int_{[0,\eta)} \{\mathbb{E}[Y(X) \mid X^u] - \mathbb{E}_2[Y(X) \mid X^u]\} \\
&\qquad \times \{\gamma_1(u \mid X) - \gamma(u \mid X)\} \frac{\mathbb{P}(C \geq u \mid X)}{\mathbb{P}_1(C > u \mid X)} \,\mathrm{d}\mu(u) \mid W \Big].
\end{aligned}
$$

Furthermore, it holds that $\mathrm{Var}[Y^*_{\mathbb{P},\mathbb{P}}(C, X^C) \mid W] \geq \mathrm{Var}[Y \mid W]$ *with equality only in the degenerate case where* $\mathrm{Var}[Y \mid X^u] = 0$ *almost surely for all* $u \in [0, \eta)$ *where* $\mathbb{P}(C \in \mathrm{d}u \mid X)$ *is non-zero.*

The proof of Theorem 6.2.3 is deferred to the Supplementary material. The notation $\mathbb{E}_2$ denotes expectation under $\mathbb{P}_2$. The outcome $Y(X)$ is observed on $(C \geq \eta)$ since $Y(X) = \tilde{Y}(X^\eta)$, so the first part of the transformation only uses complete case data but corrects for the incurred bias by reweighing with the *inverse probability of censoring weights* (IPCW). This is in itself a censoring unbiased transformation and can hence be used to construct so-called IPCW pseudo-outcomes. The second term includes the contributions for the partial observations thus making more efficient use of the data. Note that the expression in Theorem 6.2.3 is not immediately well-defined due to the uncountably many null sets associated with the conditional expectations $\mathbb{E}_2[Y(X) \mid X^u]$, and one should hence take fixed regular conditional expectation throughout as in Van der Vaart (2004).

The transformation in Theorem 6.2.3 is a generalization of the ones found in Rubin & Van der Laan (2007) and Steingrimsson et al. (2016), which only apply to survival settings where C and the survival time T are continuously distributed. That it simplifies to the known transformation in the survival setting may be seen by noting that both $\mathbb{P}_1(C \in \mathrm{d}u \mid X)$ and $\mathrm{d}1_{(C\leq u)}$ are zero on $(T \leq u)$. The variance result appears to be new and shows that pseudo-outcomes have increased variance even if the true nuisance parameters are used.

In Section 6.3, it is seen that $\mathbb{E}[Y^*_{\mathbb{P}_1,\mathbb{P}_2}(C, X^C) - Y^*_{\mathbb{P},\mathbb{P}}(C, X^C) \mid W]$, which is the conditional bias of the pseudo-outcomes, is important for determining the asymptotic behavior of DRCUT-based regression estimators. Theorem 6.2.3 immediately implies a double robustness property as stated in Corollary 6.2.6, which in turn implies that Theorem 6.2.3 gives a doubly robust representation for the conditional bias. The usefulness of Theorem 6.2.3 comes from these observations.

Remark 6.2.4. (Estimands with dynamical conditioning information.)
It is possible to use the proof of Theorem 6.2.3 to show that the transformation

$$\begin{aligned}
&Y^*_{\mathbb{P}_1,\mathbb{P}_2}(t, C, X^C) \\
&= \frac{Y(X)1_{(C\geq\eta)}}{\mathbb{P}_1(C \geq \eta \mid X)} + \int_{[0,\eta)} \frac{\mathbb{E}_2[Y(X) \mid X^{u\vee t}]}{\mathbb{P}_1(C > u \mid X)} \left\{ \mathrm{d}1_{(C\leq u)} - 1_{(C\geq u)} \frac{\mathbb{P}_1(C \in \mathrm{d}u \mid X)}{\mathbb{P}_1(C \geq u \mid X)} \right\}
\end{aligned}$$

for any $t \geq 0$ satisfies

$$\begin{aligned}
&\mathbb{E}[Y^*_{\mathbb{P}_1,\mathbb{P}_2}(t, C, X^C) - Y(X) \mid X^t] \\
&= \mathbb{E}\left[\int_{[0,\eta)} \mathbb{E}[Y(X) \mid X^{u\vee t}] - \mathbb{E}_2[Y(X) \mid X^{u\vee t}] \;\mathrm{d}\left\{ \frac{\mathbb{P}(C > u \mid X)}{\mathbb{P}_1(C > u \mid X)} \right\} \mid X^t \right].
\end{aligned}$$

This extends the use case of doubly robust transformations to parameters on the form $\mathbb{E}[Y(X) \mid X^t]$ which are often of interest. An example could be the total duration spent as disabled in an illness-death model given the state and duration at time t. The methods and results of this paper are straightforward to generalize to such estimands. $\triangledown$

Remark 6.2.5. (Estimands with interventions.)
Let A be a coordinate of X denoting the observed treatment which for simplicity is assumed to be binary. Let $X^{(a)}$ be the potential outcome corresponding to what would have happened if A had been $a \in \{0, 1\}$ and set $X = X^{(A)}$. Assume no unmeasured confounding $A \perp\!\!\!\perp X^{(a)} \mid W$ which leads to the following identification formula for the treatment-specific conditional mean

$$\mathbb{E}[Y(X^{(a)}) \mid W] = \mathbb{E}[Y(X^{(a)}) \mid W, A = a] = \mathbb{E}[Y(X) \mid W, A = a].$$

Impose treatment positivity $\mathbb{P}(A = a \mid W) \geq \epsilon > 0$. In this case, the uncentered efficient influence function motivates the transformation

$$\begin{aligned}
&Y^*_{\mathbb{P}_1,\mathbb{P}_2,\mathbb{P}_3}(a, C, X^C) \\
&= \frac{Y(X)1_{(C\geq\eta)}1_{(A=a)}}{\mathbb{P}_1(C \geq \eta \mid X)\mathbb{P}_3(A = a \mid W)} - \frac{1_{(A=a)} - \mathbb{P}_3(A = a \mid W)}{\mathbb{P}_3(A = a \mid W)} \mathbb{E}_2[Y(X) \mid W, A = a] \\
&\quad + \frac{1_{(A=a)}}{\mathbb{P}_3(A = a \mid W)} \int_{[0,\eta)} \frac{\mathbb{E}_2[Y(X) \mid X^u]}{\mathbb{P}_1(C > u \mid X)} \left\{ \mathrm{d}1_{(C\leq u)} - 1_{(C\geq u)} \frac{\mathbb{P}_1(C \in \mathrm{d}u \mid X)}{\mathbb{P}_1(C \geq u \mid X)} \right\}.
\end{aligned}$$

See the Supplementary material for details. Related results are Theorem 6.1 in Van der Laan & Robins (2003) for the discrete-time case and Theorem 1 in Rytgaard et al. (2022) for the continuous-time case where treatment starts strictly after time 0. By calculations similar to those in the proof of Theorem 6.2.3, one can show

$\mathbb{E}[Y^*_{\mathbb{P},\mathbb{P},\mathbb{P}}(a, C, X^C) \mid W] = \mathbb{E}[Y(X) \mid W, A = a]$. This implies

$$\begin{aligned}
&\mathbb{E}[Y^*_{\mathbb{P}_1,\mathbb{P}_2,\mathbb{P}_3}(a, C, X^C) - Y^*_{\mathbb{P},\mathbb{P},\mathbb{P}}(a, C, X^C) \mid W] \\
&= \mathbb{E}\Bigg[Y^*_{\mathbb{P}_1,\mathbb{P}_2,\mathbb{P}_3}(a, C, X^C) - \frac{Y(X)1_{(A=a)}}{\mathbb{P}_3(A = a \mid W)} \\
&\qquad - \left(\frac{1_{(A=a)}}{\mathbb{P}(A = a \mid W)} - \frac{1_{(A=a)}}{\mathbb{P}_3(A = a \mid W)}\right) Y(X) \mid W\Bigg] \\
&= \mathbb{E}\left[\frac{1_{(A=a)}}{\mathbb{P}_3(A = a \mid W)} \int_{[0,\eta)} \mathbb{E}[Y(X) \mid X^u] - \mathbb{E}_2[Y(X) \mid X^u] \, \mathrm{d}\left\{\frac{\mathbb{P}(C > u \mid X)}{\mathbb{P}_1(C > u \mid X)}\right\} \mid W\right] \\
&\quad + \frac{\mathbb{P}(A = a \mid W) - \mathbb{P}_3(A = a \mid W)}{\mathbb{P}_3(A = a \mid W)} \left(\mathbb{E}[Y(X) \mid W, A = a] - \mathbb{E}_2[Y(X) \mid W, A = a]\right)
\end{aligned}$$

where the last equality follows from factoring out $1_{(A=a)}/\mathbb{P}_3(A = a \mid W)$ in the first two terms of the transformation and then proceeding as in Theorem 6.2.3. This is analogous to Theorem 6.2.3 and implies that the conditional mean is correct if either $\mathbb{P}_2 = \mathbb{P}$ or $\mathbb{P}_1 = \mathbb{P}_3 = \mathbb{P}$. The methods and results of this paper generalize straightforwardly to such estimands. $\triangledown$

Corollary 6.2.6. *(Double robustness.)*
Under the same assumptions as in Theorem 6.2.3, $\mathbb{P}_1 = \mathbb{P}$ or $\mathbb{P}_2 = \mathbb{P}$ implies that

$$\mathbb{E}[Y^*_{\mathbb{P}_1,\mathbb{P}_2}(C, X^C) \mid W] = \mathbb{E}[Y(X) \mid W].$$

Corollary 6.2.6 follows immediately from Theorem 6.2.3. In the case where $\mathbb{P}_1 = \mathbb{P}_2 = \mathbb{P}$, the pseudo-outcomes are referred to as *oracle pseudo-outcomes*. Note that in this paper, the oracle knows the correct transformation but not the uncensored data.

6.3 Asymptotics and inference

This section considers estimation and large sample properties of the proposed DRCUT. A sample-splitting approach similar to Kennedy (2023) is used since this allows one to exploit the product structure from Theorem 6.2.3 to prove a rate double robustness property, permitting fast convergence rates for the estimand of interest even in settings where estimating the nuisance parameters is hard such that their individual convergence rates are slower. An example could be high dimensional settings, where regularization is used to keep the variance of the estimator from blowing up, which however makes the bias of the estimator decrease slower than it otherwise would have. In addition, sample-splitting removes the need for Donsker conditions. If Donsker conditions hold, meaning that the nuisance estimators belong to sufficiently simple function classes, then the convergence happens uniformly over that class, and the bias introduced by overfitting, i.e. by estimating the nuisance estimators in-sample, vanishes asymptotically. Donsker conditions may fail when the function class becomes too big such as for models where the dimension of the W is modeled as increasing with the sample size. For a more detailed discussion of these points, confer with Chernozhukov et al. (2018).

Sample splitting also makes the proofs simple and model agnostic, allowing for flexible nuisance estimators whose exact statistical properties may be difficult to determine e.g. estimators that depend on hyperparameters selected in a data-adaptive way. Finally, when using sample-splitting, many established asymptotic results for the second-step regression method can be immediately lifted to asymptotic results about the proposed two-step estimators. This in turn means that one can leverage existing software packages for estimation and inference since the two-step estimator then asymptotically behaves as the second-step regression but where the pseudo-outcomes enter as if it was unmodified observed data, see Proposition 6.3.5 below. Sample-splitting however has the downside that only a subset of the data is used for estimating the estimand of interest, but full-sample efficiency can be regained using cross-fitting as shown in Section 6.3.2.

The proposed estimation algorithm is analogous to Algorithm 1 in Kennedy (2023) and is described in Algorithm 1. Assume that the available data D^{2n} consists of $2n$ i.i.d. observations. Randomly partition the data into D_1^n and D_2^n of size n each. Denote by (C, X^C) a generic outcome not from D_1^n and $(C_i, X_i^{C_i})$ the outcome for the i'th subject in D_2^n.

Algorithm 1 Pseudo-algorithm for doubly robust learning with censored data.

Input: Data D^{2n} split into D_1^n and D_2^n.

1: **Nuisance estimation:** Construct estimators $\hat{\mathbb{P}}_{1,n}$ and $\hat{\mathbb{P}}_{2,n}$ of $\mathbb{P}$ using D_1^n.

2: **Pseudo-outcome regression:** In the sample D_2^n, construct the pseudo-outcomes

$$\hat{Y}^*_{\hat{\mathbb{P}}_{1,n},\hat{\mathbb{P}}_{2,n}}(C, X^C) = \frac{Y(X)1_{(C\geq\eta)}}{\hat{\mathbb{P}}_{1,n}(C \geq \eta \mid X)} + \int_{[0,\eta)} \frac{\hat{\mathbb{E}}_{2,n}[Y(X) \mid X^u]}{\hat{\mathbb{P}}_{1,n}(C > u \mid X^u)} \times \left\{ \mathrm{d}1_{(C\leq u)} - 1_{(C\geq u)} \frac{\hat{\mathbb{P}}_{1,n}(C \in \mathrm{d}u \mid X^u)}{\hat{\mathbb{P}}_{1,n}(C \geq u \mid X^u)} \right\}$$

and regress them on covariates W, which results in a regression function

$$\hat{\mathbb{E}}_n[\hat{Y}^*_{\hat{\mathbb{P}}_{1,n},\hat{\mathbb{P}}_{2,n}}(C, X^C) \mid D_1^n, W = w].$$

3: **Cross-fitting (optional):** Repeat steps 1 and 2, swapping the roles of D_1^n and D_2^n. Average over the results as a the final estimate. K-fold cross-fitting is also possible.

Output: Estimator of $\mathbb{E}[Y(X) \mid W = w]$.

Remark 6.3.1. (Doubly robust random forests.)
Steingrimsson et al. (2016) and Steingrimsson et al. (2019) use a DRCUT for the composite outcome $Y(X) = L(T, W)$ where L is a loss function. This is sufficient for their purposes since the second-step estimator $\hat{\mathbb{E}}_n$ is restricted to regression trees and random forests, see Algorithm 1 and 2 of Steingrimsson et al. (2019), which only depend on data through the loss of the individual observations. Differently from Algorithm 1, it is proposed to estimate the nuisance parameters using all the data, and then to fit the regression model on the full data set of pseudo-outcomes.

In Section 2.4 of Steingrimsson et al. (2016), the possibility of estimating nuisance parameters in parallel with fitting the regression trees is discussed and it is stated to impair performance. The possibility of using sample splitting is not discussed. On the contrary, Steingrimsson et al. (2016) states that "...it is not obvious why using pre-computed estimators of these functions derived from the entire dataset should lead to overly optimistic risk estimators". Even if the approach does not lead to overfitting when the goal is prediction, a nuisance estimation performed in-sample can affect inference due to the added variability induced by estimating the nuisance parameters. For valid inference, one would need to quantify this added variability and adjust the standard errors coming from the second-step regression accordingly. When nuisance estimation is performed using sample-splitting, Proposition 6.3.5 implies that the effect on inference is simple; the estimator behaves as if one had access to the oracle pseudo-outcomes. ▽

6.3.1 Sample splitting estimator

For notational convenience, attention is initially restricted to the sample splitting estimator, consisting of step 1 and 2 from Algorithm 1. The extension to cross-fitting is given in Section 6.3.2. In order to use the results of Kennedy (2023), some extra notation is introduced. Introduce the shorthand $\hat{Y}^* = Y^*_{\hat{\mathbb{P}}_{1,n},\hat{\mathbb{P}}_{2,n}}$ and write $Y^* = Y^*_{\mathbb{P},\mathbb{P}}$ for the oracle pseudo-outcomes. Introduce the conditional bias

$$\hat{b}(w; D_1^n) = \mathbb{E}[\hat{Y}^*(C, X^C) - Y^*(C, X^C) \mid D_1^n, W = w].$$

The effect of conditioning on D_1^n is that $\hat{\mathbb{P}}_{1,n}$ and $\hat{\mathbb{P}}_{2,n}$ are fixed in the conditional expectation. Define:

$$\begin{aligned}
m(w) &= \mathbb{E}[Y^*(C, X^C) \mid W = w], \\
\hat{m}(w) &= \hat{\mathbb{E}}_n[\hat{Y}^*(C, X^C) \mid D_1^n, W = w], \\
\tilde{m}(w) &= \hat{\mathbb{E}}_n[Y^*(C, X^C) \mid W = w].
\end{aligned}$$

Thus, $m(w)$ is the oracle conditional expectation which also equals $\mathbb{E}[Y(X) \mid W = w]$, $\hat{m}(w)$ is the regression estimator obtained from regressing $\hat{Y}^*(C, X^C)$ on W in the sample D_2^n using a given regression estimator $\hat{\mathbb{E}}_n$, and $\tilde{m}(w)$ is the oracle regression estimator. To infer the asymptotics of $\hat{m}(w)$, decompose $\hat{m}(w) - m(w)$ into the sum of $\hat{m}(w) - \tilde{m}(w)$ and $\tilde{m}(w) - m(w)$. The asymptotics of $\tilde{m}(w) - m(w)$ can often be inferred from known asymptotic theory for the chosen regression estimator $\hat{\mathbb{E}}_n$. From hereon, it is assumed that the convergence rate is α.

Assumption 6.3.2. *(Convergence rate of second-step regression method.)*
It holds that $\tilde{m}(w) - m(w) = O_{\mathbb{P}}(n^{-\alpha})$. $\diamond$

Remark 6.3.3. (Oracle mean squared error and convergence rates.)
In some situations, it might be more natural to take the convergence rate of the oracle mean squared error $R_n^*(w) = \mathbb{E}[\{\tilde{m}(w) - m(w)\}^2]^{1/2}$ as a starting point rather than that of $\tilde{m}(w) - m(w)$. Straightforward calculations show that if $n^\alpha\{\tilde{m}(w) - m(w)\}$ converges in distribution to some distribution with mean μ and variance σ^2 and the first and second moments also converge then $n^\alpha R_n^*(w) \to (\sigma^2 + \mu^2)^{1/2}$ for $n \to \infty$ implying $R_n^*(w) = O_{\mathbb{P}}(n^{-\alpha})$. The convergence rates are hence the same whenever one has sufficient integrability. $\triangledown$

Remark 6.3.4. (Pointwise estimation.)
In this paper, the focus is the pointwise problem of estimating $\mathbb{E}[Y(X) \mid W = w]$ for any given w. This is the relevant estimand in the data application of Section 6.5 since RDDs utilize conditional expectations at a specific boundary value to estimate a local causal effect. It is seen from Remark 6.3.3 that studying the pointwise convergence rates is equivalent to studying the convergence rate of the mean squared error. Hence, if the integrated MSE rather than the pointwise MSE is the relevant performance metric, the approach in Rambachan et al. (2022) might be more suitable. $\triangledown$

For the remaining terms, write

$$\begin{aligned}\hat{m}(w) - \tilde{m}(w) &= \hat{\mathbb{E}}_n[\hat{b}(W; D_1^n) \mid D_1^n, W = w] \\ &\quad + \left(\hat{m}(w) - \tilde{m}(w) - \hat{\mathbb{E}}_n[\hat{b}(W; D_1^n) \mid D_1^n, W = w]\right).\end{aligned}$$

Theorem 6.2.3 gives a doubly robust representation for $\hat{b}(W; D_1^n)$, so the first term can usually be made $o_{\mathbb{P}}(n^{-\alpha})$ by having the nuisance estimators converge sufficiently fast. Following Definition 1 of Kennedy (2023), the regression method $\hat{\mathbb{E}}_n$ is said to be *stable* if the second term is $o_{\mathbb{P}}(n^{-\alpha})$ whenever $d(\hat{Y}^*, Y^*) = o_{\mathbb{P}}(1)$ for a suitable stochastic distance d.

By Theorem 1 in Kennedy (2023), the class of linear smoothers

$$\hat{\mathbb{E}}_n[f(C, X^C; D_1^n) \mid D_1^n, W = w] = \sum_{i=1}^n p_i(w; W^n) f(C_i, X_i^{C_i}; D_1^n)$$

for $W^n = (W_k)_{1 \leq k \leq n}$ is stable under suitable regularity conditions. Some prominent methods that belong to this class are listed in Kennedy (2023). Importantly, local linear regression is a linear smoother, which is the de facto method used in RDDs and which is also employed in Section 6.4 and 6.5. It is also possible to force the output of more flexible methods to be on this form to keep inference tractable and enhance interpretability. This can for example be done following Verdinelli and Wasserman (2021) which also relies on sample splitting, first fitting a random forest and then using the resulting estimator to define a kernel used for local linear regression in the second split.

To define the distance under which linear smoothers are stable, first introduce the conditional $L^2(\mathbb{P})$-norm

$$\|f(Z)\|_{w,D_1^n} = \mathbb{E}[f(Z)^2 \mid D_1^n, W = w]^{1/2}.$$

Theorem 1 in Kennedy (2023) then implies that linear smoothers are stable at $W = w$ with respect to the stochastic distance $d_{w,D^{2n}}$ given by

$$d_{w,D^{2n}}(g,f) = \sum_{i=1}^{n} \left\{ \frac{p_i(w; W^n)^2}{\sum_{j=1}^{n} p_j(w; W^n)^2} \|g(C, X^C; D_1^n) - f(C, X^C)\|_{W_i, D_1^n}^2 \right\}$$

whenever $d_{w,D^{2n}}(0, \mathrm{Var}[Y^*(C, X^C) \mid W = \bullet])^{-1} = O_{\mathbb{P}}(1)$.

Thus, if $d_{w,D^{2n}}(\hat{Y}^*, Y^*) = o_{\mathbb{P}}(1)$ and $d_{w,D^{2n}}(0, \mathrm{Var}[Y^*(C, X^C) \mid W = \bullet])^{-1} = O_{\mathbb{P}}(1)$ then stability gives

$$\hat{m}(w) - \tilde{m}(w) = \hat{\mathbb{E}}_n[\hat{b}(W; D_1^n) \mid D_1^n, W = w] + o_{\mathbb{P}}(n^{-\alpha}).$$

One can thus focus on the asymptotics of the conditional bias. Due to its product structure, one can obtain rate double robustness results like the one in Proposition 6.3.5. Define the stochastic norm $\|\cdot\|_{2,w,D^{2n}}$ by

$$\|f(u, X; D_1^n)\|_{2,w,D^{2n}}^2 = \sum_{i=1}^{n} \frac{|p_i(w; W^n)|}{\sum_{j=1}^{n} |p_j(w; W^n)|} \int_{[0,\eta)} \|f(u, X; D_1^n)\|_{W_i, D_1^n}^2 \, \mathrm{d}\mu(u).$$

Proposition 6.3.5. *(Rate double robustness under weighted L^2-rates.)*
Impose the assumptions from Theorem 6.2.3, Assumption 6.3.2, and for a fixed w

(a) $\inf_z \mathrm{Var}[Y(X) \mid W = z] > 0$;

(b) $d_{w,D^{2n}}(\hat{Y}^*, Y^*) = o_{\mathbb{P}}(1)$;

(c) $\sum_{i=1}^{n} |p_i(w; W^n)| = O_{\mathbb{P}}(1)$;

(d) $\|\mathbb{E}[Y(X) \mid X^u] - \hat{\mathbb{E}}_{2,n}[Y(X) \mid X^u]\|_{2,w,D^{2n}} = O_{\mathbb{P}}(n^{-\alpha_1})$;

(e) $\|\hat{\gamma}_{1,n}(u \mid X) - \gamma(u \mid X)\|_{2,w,D^{2n}} = O_{\mathbb{P}}(n^{-\alpha_2})$;

(f) $\alpha_1 + \alpha_2 > \alpha$.

Then $\hat{m}(w) - m(w) = \tilde{m}(w) - m(w) + o_{\mathbb{P}}(n^{-\alpha})$ *i.e. oracle efficiency is obtained.*

The proof of Proposition 6.3.5 is deferred to the Supplementary material. As discussed in Remark 5 of Kennedy (2023), results like Proposition 6.3.5 are important for inference since they imply that the asymptotic distribution when using estimated and oracle pseudo-outcomes are identical. Confidence intervals for $m(w)$ may hence be constructed by treating the estimated pseudo-outcomes as if they were observed outcomes and employing the usual asymptotic distributional approximation. Standard implementations can therefore be used.

Local polynomial regression achieves the minimax optimal convergence rate cf. Stone (1980, 1982), so the proposed methods do as well when Proposition 6.3.5 applies. The sample size used in the asymptotic approximation however becomes n, which is half the number of observations that was originally available, leading to an inferior constant in the minimax risk of the estimator. As shown in the next section, cross-fitting may be used to regain full sample efficiency.

Remark 6.3.6. (On the regularity conditions in Proposition 6.3.5.)
Condition (a) is used to prove $d_{w,D^{2n}}(0, \mathrm{Var}[Y^*(C, X^C) \mid W = \bullet])^{-1} = O_{\mathbb{P}}(1)$ which together with (b) implies stability of linear smoothers. As noted in Kennedy (2023), many linear smoothers satisfy that $\sum_{i=1}^n |p_i(w; W^n)|$ is bounded by a fixed constant with probability one (which is also implied by condition (1) in Theorem 1 of Stone (1977) regarding universal weak consistency of linear smoothers). This would imply (c). To obtain reasonable convergence rates in (d) and (e), one likely needs to assume that the dependence on X^u can be captured by a d-dimensional stochastic process $Z(u) = f(u, X^u)$ taking values in a compact subset of $\mathbb{R}^d$. Assuming that the $L^q(\mathbb{P})$ convergence rates results from Stone (1980, 1982) carry over to the weighted $L^2(\mathbb{P})$-norms $\|\cdot\|_{2,w,D^{2n}}$ and that $z_u \mapsto \mathbb{E}[Y(X) \mid Z(u) = z_u]$ is s-times continuously differentiable, Theorem 1 in Stone (1980, 1982) implies that the optimal convergence rate in a minimax sense is $O_{\mathbb{P}}(n^{-r})$ for $r = 1/(2 + d/s)$ under some regularity conditions. This rate can be obtained using e.g. series or local polynomial estimators. Other structured assumptions such as sparsity or additivity are popular alternatives when they are applicable, see for instance Yang & Tokdar (2015). ▽

Remark 6.3.7. (Extension to vector-valued outcomes.)
The results and proofs of Section 6.2.1 are unchanged if Y takes values in $\mathbb{R}^p$ for some $p \geq 1$ rather than $p = 1$. Similarly, if Proposition 6.3.5 holds for each coordinate of Y, then also $\hat{m}(w) - m(w) = \tilde{m}(w) - m(w) + o_{\mathbb{P}}(n^{-\alpha})$ as random vectors so in this case the joint asymptotic distribution are the same by Slutsky's lemma. ▽

Remark 6.3.8. (Oracle efficiency without sample splitting.)
If nuisance estimators converge sufficiently fast and uniformly, the added variance from estimating nuisance parameters in-sample may become asymptotically negligible, see e.g. Lemma 19.24 in Van der Vaart (1998) and Lemma 2 in Cui et al. (2023). In this case, sample-splitting is not necessary and the conditional bias is less

relevant. Furthermore, that approach might generalize more easily to estimators that are not linear smoothers. The present approach is chosen to allow for weaker conditions on the convergence rates which exploit the product structure of the conditional bias. $\triangledown$

6.3.2 Cross-fitted estimator

In this section, the arguments are extended to K-fold cross-fitting. Assume that one has access to n observations and deterministically partition the data into folds of size n/K. For notational simplicity, assume that $\hat{\mathbb{E}}_n$ is asymptotically Gaussian such that $n^\alpha\{\tilde{m}(w) - m(w)\} \to \mathcal{N}(\mu, \sigma^2)$. Let $\tilde{m}_k(w)$ be the oracle estimator when only data from fold k ($k = 1, \dots, K$) is used. Let D^{-k} be the data not in fold k. Write $\hat{Y}^*_{-k}$ for the pseudo-outcomes with estimates based on D^{-k}. Similarly, $\hat{\mathbb{E}}_k$ is the regression estimator based on the data in fold k. Denote by

$$\hat{m}_k(w) = \hat{\mathbb{E}}_k[\hat{Y}^*_{-k}(C, X^C) \mid D^{-k}, W = w]$$

the estimator obtained from estimating the nuisance parameters using D^{-k} and then regressing over the pseudo-outcomes from fold k. The proposed cross-fitted estimator is then $\hat{m}^{\mathrm{CF}}(w) = 1/K \sum_{k=1}^K \hat{m}_k(w)$. This cross-fitting scheme is similar to DML1 in Chernozhukov et al. (2018). An alternative not explored here could be like DML2 to first compute all the pseudo-outcomes $\hat{Y}^*_{-k}(C, X^C)$ and then input them simultaneously into $\hat{\mathbb{E}}_n$. The following proposition shows that the cross-fitted estimator regains full-sample efficiency.

Proposition 6.3.9. *(Asymptotic distribution of cross-fitting estimator.) Under the assumptions from Proposition 6.3.5 it holds that*

$$\hat{m}_k(w) - m(w) = \tilde{m}_k(w) - m(w) + o_{\mathbb{P}}(n^{-\alpha})$$

and

$$n^\alpha\{\hat{m}^{\mathrm{CF}}(w) - m(w)\} \to \mathcal{N}(K^\alpha \mu, K^{2\alpha-1}\sigma^2)$$

in distribution.

The proof of Proposition 6.3.9 is deferred to the Supplementary material. Since $\alpha \leq 1/2$, the asymptotic variance of the cross-fitted estimator is no larger than that of the full-sample oracle estimator $\tilde{m}(w)$ and is strictly less when $\alpha < 1/2$, while the bias is increased by a factor of K^α. In the special case where $\mu = 0$ and $\alpha = 1/2$, the asymptotic distribution of $\tilde{m}(w)$ and $\hat{m}^{\mathrm{CF}}(w)$ are identical which agrees with previous results in the literature, see e.g. Remark 3.1 and Theorem 3.1 in Chernozhukov et al. (2018). Proposition 6.3.9 implies that the standard error of $\hat{m}^{\mathrm{CF}}(w)$ can be estimated by averaging the estimated standard errors $\hat{\sigma}_k/(n/K)^\alpha$ of $\hat{m}_k(w)$ ($k = 1, \dots, K$) and scaling by $K^{-1/2}$.

Remark 6.3.10. (Bias-variance tradeoff with cross-fitting.)
It can be seen from the proof of Proposition 6.3.9 that the asymptotic distribution comes from averaging $\tilde{m}_k(w)$ $(k = 1, \dots, K)$. We therefore conjecture that a DML2-variant would, under similar regularity assumptions, have the same asymptotic distribution as $\tilde{m}(w)$. Thus, when $\alpha < 1/2$, the estimator $\hat{m}^{\mathrm{CF}}(w)$ trades a decrease in variance for an increase in bias compared to a DML2-variant. This phenomenon seems to imply that an efficient estimator converging at a rate slower than $\sqrt{n}$ cannot have a non-zero asymptotic bias since splitting and averaging would then decrease the asymptotic variance without a corresponding increase in bias. For an asymptotically unbiased estimator converging at a sub-optimal rate e.g. univariate local linear regression with undersmoothing, the variance reduction due to averaging is analogous to increasing the proportionality constant in the bandwidth and can hence be thought of as smoothing. $\triangledown$

6.4 Simulations

6.4.1 Data-generating process

To examine the finite sample predictive and inference performance of the proposed estimator and to demonstrate the double robustness property, a numerical study is conducted. The complete-case data is specified as $X = (Z, W)$ where Z follows the irreversible illness-death model depicted in Figure 6.1 with a time-horizon of $\eta = 5$ and initial state $Z(0) = 1$ and the baseline covariate is $W \sim \mathrm{Uniform}(-4, 4)$. The outcome of interest is the duration spent in the illness state before the end of the observation window meaning $Y(X) = \int_{[0,\eta)} 1\{Z(s) = 2\} \,\mathrm{d}s$ and $\mathbb{E}[Y(X) \mid W] = \int_{[0,\eta)} p_2(s, W) \,\mathrm{d}s$ for the state-occupation probability $p_2(s, W) = \mathbb{P}\{Z(s) = 2 \mid W\}$. The censored outcome (C, X^C) is simulated by first simulating W and then simulating $(C, Z^C) \mid W$ using Lewis' thinning algorithm from Ogata (1981). The `R` implementation (R Development Core Team, 2023) is available on GitHub (`https://github.com/oliversandqvist/Web-appendix-drcut`). A total of

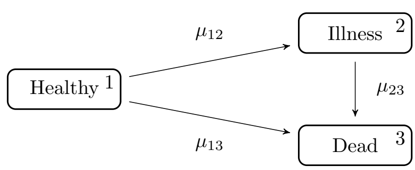

Figure 6.1: *The irreversible illness-death model for the process Z. Transitions from state j to state k has the transition hazard μ_{jk}.*

500 samples of sizes $n \in \{1000, 5000, 10000, 30000\}$ are considered. For a given subject, the data is generated as follows: The hazard γ of $C \mid X$ is set to $\gamma(t, W) = 1_{\{Z(t)=1\}} \exp\{\log(0.2) + 0.6 \times 1_{(-2 \le W < 2)}\}$, which results in a substantial amount of right-censoring as well as highly state-dependent censoring. Subjects not censored

before time η are administratively censored. Events are simulated according to the transition hazards

$$\mu_{12}(t,W) = \exp\big\{\log(0.3) + 0.15\times\cos(\pi W/2) + 0.15\times 1_{(t>2.5)} - 0.05\times W\big\},$$
$$\mu_{13}(t,W) = \exp\big\{\log(0.1) + 0.3\times\sin(\pi W/2) + 0.05\times t\big\},$$
$$\mu_{23}(t,S(t),W) = \exp\big[-0.75\times\min\{t-S(t),3\}\times(1.07+0.09\times\bar{W} - 0.024\times\bar{W}^2 - 0.014\times\bar{W}^3 + 0.001\times\bar{W}^4 + 0.00065\times\bar{W}^5)\big],$$

where $S(t)=\sup\{s\le t : Z(s)\neq Z(t)\}$ is the latest jump time and $\bar{W}=\min(W,3)$. With these specifications, Assumption 6.2.1 holds because the censoring intensity is adapted to the filtration generated by X and Assumption 6.2.2 holds because the largest probability of becoming censored before time η is obtained by remaining in the Healthy state, and this leads to a censoring probability that is strictly less than one. In addition, Y clearly has finite expectation. The required assumptions for use of IPCW and doubly robust pseudo-outcomes are hence satisfied.

For computation of Y^* it is convenient to introduce the prospective illness duration $Y(X,t)=\int_{(t,\eta)} 1_{\{Z(s)=2\}}\,\mathrm{d}s$ and its conditional expectation $V_{Z(t)}\{t,S(t),W\} = \mathbb{E}[Y(X,t)\mid X^t]$. Then

$$\mathbb{E}[Y(X)\mid X^t] = \int_{(0,t]} 1_{\{Z(s)=2\}}\,\mathrm{d}s + V_{Z(t)}\{t,S(t),W\}$$

and V may be calculated using differential equations whenever transition hazards exist, see Corollary 7.2 in Adékambi & Christiansen (2017), giving

$$\frac{\mathrm{d}}{\mathrm{d}t}V_1(t,s,w) = \{\mu_{12}(t,w)+\mu_{13}(t,w)\}\times V_1(t,s,w) - \mu_{12}(t,s,w)\times V_2(t,t,w),$$
$$\frac{\mathrm{d}}{\mathrm{d}t}V_2(t,s,w) = -1 + \mu_{23}(t,s,w)\times V_2(t,s,w),$$

with $V_1(\eta,s,w)=V_2(\eta,s,w)=0$. The fourth-order Runge-Kutta method is used to solve these differential equations. Note that computing $V_j(t,s,w)$ via this approach also yields $V_j(u,s,w)$ for all $u\ge t$.

Remark 6.4.1. (Relevance of the estimand.)
This setup is motivated by disability insurance applications, where the length of an illness is a key driver of expenses since insureds often receive disability benefits as long as they are disabled to make up for lost wages. Since they receive large benefits, subjects do not leave the portfolio while disabled, so there is no censoring hazard while in the illness state. This could also be a relevant estimand in medical applications where the length of an illness or a hospital stay could be an important aspect to predict or to make inferences about. ▽

6.4.2 Estimators

The following estimators are considered:

(a) A plug-in estimator with p_2 estimated by the Conditional Aalen-Johansen (CAJ) of Bladt & Furrer (2024) using the R-package `AalenJohansen`;

(b) A plug-in estimator with transition hazards μ_{12}, μ_{13}, and μ_{23} estimated by the Highly Adaptive Lasso (HAL) of Benkeser & Van Der Laan (2016) and Munch et al. (2024) using a custom implementation relying on the R-package `glmnet`;

(c) Two-fold cross-fitted doubly robust pseudo-outcomes with transition and censoring hazards estimated by HAL;

(d) Two-fold cross-fitted doubly robust pseudo-outcomes with transition hazards estimated by HAL and censoring hazards estimated by a misspecified parametric model;

(e) Two-fold cross-fitted doubly robust oracle pseudo-outcomes;

(f) Estimators (c), (d), and (e) but with IPCW pseudo-outcomes.

For all pseudo-outcome-based methods, the second-step regression method is chosen as a local linear regression using `lprobust` from the R-package `nprobust` with accompanying paper Calonico et al. (2019) using default parameters except for the bandwidth. With a one-dimensional covariate, the MSE and MISE optimal bandwidths satisfy $h \propto n^{-1/5}$ and lead to reasonable finite sample performance. However, to obtain a non-vanishing bias in the asymptotic Gaussian distribution, one needs $nh^5 \to 0$ (undersmoothing) or an explicit bias correction. For further details and discussions, see Calonico et al. (2019) and the references therein. We proceed via undersmoothing, first using a separate simulation to find a bandwidth with good performance when $n = 5\,000$ and then letting $h \propto n^{-1/4.5}$. In this case, the convergence rate of $\hat{\mathbb{E}}_n$ is $\sqrt{nh} \propto n^{7/18}$ so $\alpha = 7/18$.

The custom implementation of HAL for hazard estimation is a modification of the code from Rytgaard et al. (2022) and Rytgaard et al. (2023) allowing for higher-order interactions than second-order which is needed for estimation of μ_{23}. HAL is chosen since it is a general-purpose estimator that in Benkeser & Van Der Laan (2016) is demonstrated to have reasonable empirical performance both in smooth and discontinuous settings and is shown to have desirable asymptotic properties in Munch et al. (2024) whenever the true function is multivariate càdlàg. With this choice of hazard rates, HAL is expected to estimate the censoring hazard very closely since the true hazard is piecewise constant, while it is expected to have a harder time estimating the transition hazards of the illness-death model since these are more complicated.

The misspecified parametric family for the censoring hazard is chosen as the parametric family where $X \mid C$ has hazard $\gamma(t, W; \beta) = 1_{\{Z(t)=1\}} \exp(\beta_1 + \beta_2 \times t + \beta_3 \times W)$. This is expected to have poor performance due to the form of $\gamma(t, W)$.

The CAJ estimator is expected to be biased since the model is non-Markovian and the censoring is state-dependent, confer with Assumption 2 and Remark 2.1 in Bladt & Furrer (2024). Similarly to Munch et al. (2023) and Gunnes et al. (2007), it was observed that highly non-Markovian behaviour as well as high degrees of state-dependent censoring were required for the bias to be sizeable, and it further seems that the effect of covariates has to be small compared to the non-Markovianity of Z and the state-dependence of $C \mid X$.

Remark 6.4.2. (Marginal estimands for the irreversible illness-death model.)
The paper Munch et al. (2023) also uses efficient-influence-function-based estimators of estimands formulated using multi-state models and employs HAL to estimate nuisance parameters. Their results are however specialized to marginal state-occupation probabilities in the illness state for an illness-death model. This paper can be viewed as an extension of their approach to any square-integrable real-valued outcome and, more importantly, an extension to estimands that may depend on baseline covariates. Remark 2.8 further allows for conditioning on the history of the multi-state process. ▽

6.4.3 Results

The results for the first simulation are depicted in Figure 6.2. The estimand is a lower dimensional and smoother object than the individual hazards which makes it possible to nonparametrically estimate at a faster rate than the hazards. Since the pseudo-outcome methods use local linear regression as the second step estimator, this additional structure is exploited and these methods are therefore expected to perform well as long as the pseudo-outcomes are close to their oracle counterparts.

As seen on the left part of Figure 6.2, HAL captures the general shape of the transition hazards reasonably well and the censoring hazard extremely well as was expected. However, the right plot shows that the plug-in estimator based on HAL-estimated transition hazards performs poorly. HAL employs regularization to balance bias and variance to be optimal for the individual hazards, but this bias-variance trade-off is seen to be suboptimal for the estimand of interest as the estimate becomes too biased. The CAJ estimator is biased in this setting, and this bias carries over to the plug-in estimator, but the estimator still performs better than the plug-in HAL estimator. The HAL-based IPCW estimator performs well and is almost indistinguishable from the oracle IPCW estimator which is unsurprising since the estimated censoring hazard is very close to the true value. As expected, the misspecified IPCW estimator performs poorly. The doubly robust pseudo-outcomes perform well, resulting in values similar to those of the non-misspecified IPCW pseudo-outcomes. For the HAL-based doubly robust pseudo-values, one might have suspected that this was solely a consequence of the good performance of the IPCW term, but then the doubly robust pseudo-values with a misspecified censoring hazard should have performed poorly which is not the case. For those pseudo-values, one sees the remarkable phenomenon that although both the HAL-estimated transitions hazards and the misspecified censoring hazard gave poor estimates by themselves, the doubly robust property of the pseudo-values makes them perform well when used jointly. In the Supplementary material, one sees that the true curve is always contained in the pointwise 95% confidence bands albeit barely around $W = 0.5$.

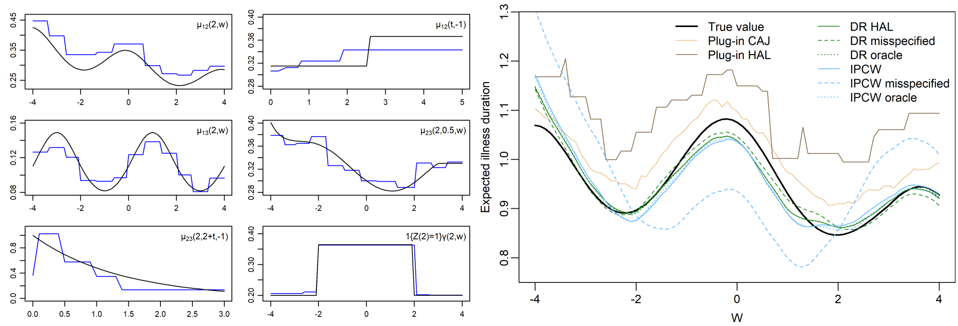

Figure 6.2: ***Left Panel:*** *Fitted HAL estimates and actual hazards at specific input values indicated at the top right corner for a single simulation.* ***Right Panel:*** *Estimators and true value of* $\mathbb{E}[Y(X) \mid W]$ *as a function of* W *for a single simulation.*

Similar patterns emerge across the 500 simulations. The reported performance metrics are the $L^2([-4, 4], \lambda)$ error which is relevant for prediction, and the empirical coverages of the confidence intervals for methods (c) and (e) which is relevant for inference. Additional performance metrics for prediction were computed, but their results were qualitatively highly similar and are hence not reported. Figure 6.3 leads to many of the same qualitative conclusions as Figure 6.2 regarding which estimators perform well. Additionally, one can see that the average performance of the plug-in CAJ, plug-in HAL, and misspecified IPCW estimator does not improve noticeably after $n = 5\,000$ although the variability decreases. For the remaining estimators, both the average performance and variability improves as n increases, and their densities are similar.

Although the performance of the doubly robust pseudo-outcomes with a misspecified and HAL-estimated censoring hazard appear similar in terms of predictive performance, it can be seen from the left plot in Figure 6.4 that the one using HAL agrees better with the Gaussian distributional approximation obtained from the oracle values and also with the true value of the estimand. The right plot shows that the empirical coverages of the confidence intervals deviate somewhat from their nominal values, but more importantly for this study is that the confidence intervals for the oracle and estimated doubly robust pseudo-values are highly similar especially for $n \geq 5\,000$ indicating that oracle efficiency is obtained. The coverages are close to their nominal value when W is away from -2, 0, and 2, where the curvature of the true estimand is the greatest, suggesting that the chosen bandwidth has led to too much smoothing for these values of W.

The choice of bandwidth greatly affects the validity of inference based on kernel estimators, see for example Table I in Calonico et al. (2014), making bandwidth selection important. In a setting resembling this numerical study, it would hence be natural to select different bandwidths for different values of W. It would be desirable to do this in some data-adaptive way, but then the resulting regression estimator might fall outside the class of linear smoothers and hence also outside of Proposition 6.3.5.

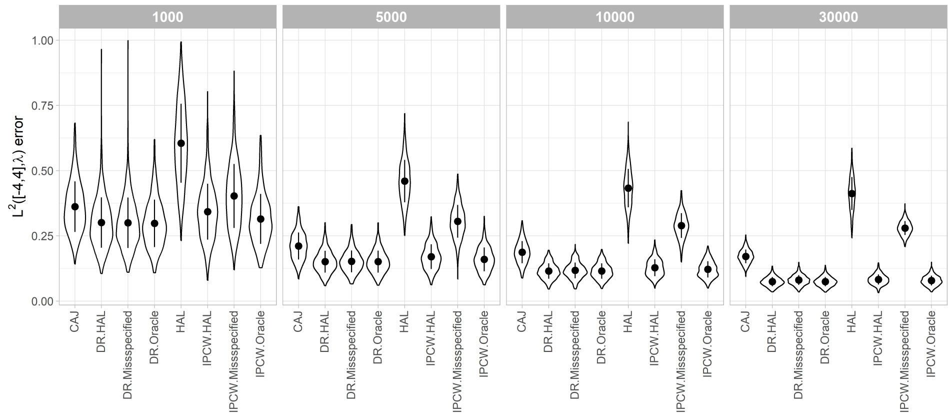

Figure 6.3: *Violin plot of the $L^2([-4,4],\lambda)$ error for different estimators and values of n with Mean $\pm$ Standard deviation indicated as a point range using 500 simulations.*

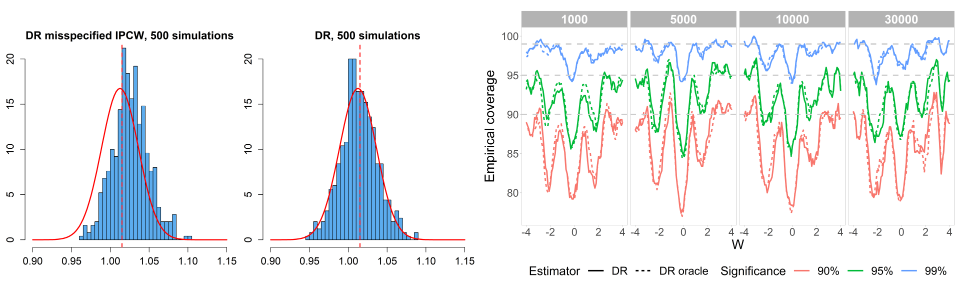

Figure 6.4: ***Left Panel:*** *Histogram of estimates at the point $W=-1$ for the doubly robust pseudo-outcomes with censoring estimated by HAL and a misspecified parametric family. The Gaussian approximation is obtained from the oracle pseudo-values and the dashed line is the true value of the estimand. Based on 500 simulations of size $n = 30\,000$.* ***Right Panel:*** *The empirical coverages of the 99%, 95%, and 90% confidence intervals using a Gaussian approximation with standard errors obtained from* `lprobust` *using HAL-estimated and oracle doubly robust pseudo-outcomes. Nominal values are shown with dashed lines.*

6.5 Data application

The proposed method is demonstrated by an application to data from LSYPE, Waves 1 to 5 (Centre for Longitudinal Studies (2024), Calderwood & Sanchez (2016)). LSYPE is a panel survey of initially around 16 000 young people (YP) born between September 1989 and August 1990 in England. Data was collected starting in 2004 and the first five Waves consisted of annual interviews with the YP and their carers. Thus, YP were in Year 9 during Wave 1. YP were allowed to leave school after Year 11 with post-compulsory schooling consisting of Years 12 and 13. We aim to estimate the impact of the Education Maintenance Allowance (EMA), a conditional cash transfer program, on time spent in full-time education. This is achieved by using the proposed methods to construct an RDD in the presence of censored data.

6.5.1 Background

The EMA program was established in England to encourage YP to continue their education after Year 11. It was piloted in September 1999, rolled out nationally in 2004, and abolished in September 2010. YP could apply for EMA during Year 11 and if EMA was awarded, YP would receive a weekly cash transfer during Year 12 and Year 13 provided they stayed in further education. EMA was awarded based on a household's annual income for the previous year submitted to the EMA administration via a bank statement. YP in households with annual incomes below £20,817 received £30, those between £20,818 and £25,521 received £20, and those between £25,522 and £30,810 received £10. No EMA was given for incomes over £30,810. The presence of these thresholds suggests that the causal effect of EMA can be estimated using an RDD. An RDD estimates a local causal effect by comparing groups just above and below a treatment threshold mimicking a (local) randomized controlled trial, see e.g. Hahn et al. (2001), Imbens & Lemieux (2008), and Cattaneo & Titiunik (2022). An RDD is therefore able to infer causal effects under relatively weak assumptions, avoiding no unmeasured confounding and similar graphical causal model based criteria, confer with Pearl (2009) and Hernán and Robins (2020).

We restrict our attention to measuring the effect of receiving high EMA since its effect is expected to be the highest and since 80% of YP receiving EMA were paid the highest rate of £30, see Bolton (2011). Note that those not receiving high EMA could still be receiving moderate or low rates, and the causal effect estimated here is therefore only valid in environments where these rates are also present. Under assumptions about how CATE changes as a function of salary, e.g. linear dependence, one could exploit the multiple thresholds to infer the causal effect of high EMA versus no EMA but this is not pursued here. For simplicity, we similarly restrict attention to whether high EMA is received in Wave 4 making treatment binary. It would be of interest to extend this approach to dynamical treatments, using the treatment status from both Wave 4 and Wave 5.

Studies based on self-reports indicated that ”only” 12% of recipients stayed in education because of EMA which the government used as a key reason for abolishing EMA, see Bolton (2011). This highlights the importance of statistical analyses in evaluating the effectiveness of such programs to guide informed policymaking. These numbers were consistent with other studies that used matching between the pilot and control groups, see Maguire et al. (2001), Middleton et al. (2005), and Dearden et al. (2009). An issue that was identified, but not controlled for, was that students staying in full-time education seemed more likely to remain in the survey, see e.g. Chapter 2.5.3 of Middleton et al. (2005). Additionally, the effect of EMA in the pilot might have been different than the national effect.

Not many studies have explored the effect of EMA after it was rolled out nationally. The only studies identified on the subject were Holford (2015), the working paper McKendrick (2022), and the unpublished PhD Rahman (2014) that employ panel regression, augmented inverse propensity weighted linear regression and Causal Forests, and an RDD, respectively. Except for the RDD, all previous studies hence rely on the assumption of no unmeasured confounding. The RDD in Rahman (2014) had some methodological weaknesses which are improved upon in this analysis. Firstly, observations in Wave 4 and Wave 5 were pooled such that a YP interviewed in Wave 4 and Wave 5 would contribute with two observations. Censored observations were discarded. This can create confounding over time e.g. if YP that responded positively to EMA and stayed in education were more likely to respond to the survey as was found in Middleton et al. (2005). Secondly, polynomial regression was used to estimate the relevant conditional expectations and to perform inference. As noted in Hahn et al. (1999), this is fragile to misspecification so local linear regression might be preferred since it is nonparametric and has good boundary properties.

Consequently, we find that an RDD based on observational data and utilizing the proposed methods can be a valuable complementary study for measuring the effect of EMA since it does not rely on no unmeasured confounding, allows the censoring distribution to depend on whether YP stays in education or not, uses the cohort is the one that emerged when EMA was well-established on a national level, and allows for the use of flexible nonparametric estimators for inference. This leads to both higher internal and external validity of the estimates.

6.5.2 Model and results

The present RDD is fuzzy since not all eligible YP apply for EMA and since the income information in LSYPE could deviate from the one submitted to the EMA administration. Additionally, the exact income is only available in Wave 1 and

Wave 2 and in banded form in Wave 3 which was the year where EMA application were submitted. The income in Wave 3 is thus estimated by taking the income from Wave 2 if this is within the band and otherwise simulate uniformly over the band. Fortunately, the bands align well with the EMA thresholds, so the risk of moving an observation across a threshold is very low. A handful of seeds were tested for the simulation and they all gave quantitatively similar results in terms of the final estimate.

Let time 0 be Wave 3, and the outcome Y be the amount of years spent in full-time education during Wave 4 and Wave 5. Censoring C takes the value 1 if YP becomes censored in Wave 4, 2 if YP becomes censored in Wave 5, and 3 if not censored in Wave 4 and 5. Let X be baseline covariates from Waves 1-3 as well as a time-dependent coordinate which at the end of the year increases by 1 if YP was in full-time education during that school year so that X^C is observable from the data and $Y = Y(X)$. Similarly, let the treatment outcome be denoted A and $Z = \{Z(t)\}_{t\geq 0}$ be as X but where the time-dependent coordinate is 1 if YP receives high EMA at the end of the year and 0 otherwise such that $A = A(Z)$. Assume (C, X^C) is a CAR of X and (C, Z^C) is a CAR of Z and that positivity holds. Let W be income in Wave 3 and $w_0 = £20{,}817$. Let $Y^{(a)}$ be the potential outcome corresponding to treatment $a \in \{0, 1\}$ and specify the causal estimand of interest as the CATE

$$\tau = \mathbb{E}[Y^{(1)} - Y^{(0)} \mid W = w_0].$$

For identification, assumptions analogous to those in Theorem 2 of Hahn et al. (2001) are imposed.

Assumption 6.5.1. *(RDD identification.)*

(i) $a^+ = \lim_{w\downarrow w_0} \mathbb{P}(A = 1 \mid W = w)$ *and* $a^- = \lim_{w\uparrow w_0} \mathbb{P}(A = 1 \mid W = w)$ *exist and* $a^+ \neq a^-$.

(ii) $\mathbb{E}[Y^{(1)} \mid W = w]$ *and* $\mathbb{E}[Y^{(0)} \mid W = w]$ *are continuous in* w *at* w_0.

(iii) $A \perp\!\!\!\perp (Y^{(1)} - Y^{(0)}) \mid W = w$ *in the limit for* $w \to w_0$. ◇

Theorem 2 in Hahn et al. (2001) then implies

$$\tau = \frac{y^+ - y^-}{a^+ - a^-}$$

where $y^+ = \lim_{w\downarrow w_0} \mathbb{E}[Y \mid W = w]$ and $y^- = \lim_{w\uparrow w_0} \mathbb{E}[Y \mid W = w]$. Assume oracle efficiency is obtained for each of y^+, y^-, a^+, and a^- when using cross-fitted doubly robust pseudo-outcomes and local linear regression, which holds under conditions given in Proposition 6.3.5 and 6.3.9. Then by Remark 6.3.7, oracle efficiency is obtained for τ since this is determined by the asymptotic distribution of $(\hat{y}^+, \hat{y}^-, \hat{a}^+, \hat{a}^-)$, confer with Hahn et al. (1999). Standard implementations for inference based on asymptotic approximations may thus be used, treating the pseudo-values as ordinary outcomes.

Estimation proceeds via Algorithm 1. Computation of the pseudo-outcomes can be formulated as sequential classification problems. Estimation is performed using the `R`-package `xgboost` where five-fold cross-validation with negative log-likelihood and AUC loss functions were used to determine suitable hyperparameters. Since `xgboost` is tree-based, it should be able to capture discontinuities caused by the EMA thresholds well. For predicting censoring, even with a moderate amount of hyperparameter optimization, it is hard to improve the performance of the model using only the covariates identified as predictors of non-response in Section 4.4 of Collingwood et al. (2010) compared to using all available covariates. A high-dimensional X may sometimes be desirable to make CAR more plausible, but since this does not seem to be needed here, we proceed with the lower dimensional model containing 14 covariates, even though the performance of the higher dimensional model is substantially better for predicting education outcomes. The analysis was also performed for the high-dimensional choice of X but is not reported as it lead to highly similar results. The second-step estimator is a local-linear-regression-based RDD implemented using the `R`-package `rdrobust` with standard parameters except for the bandwidth which is set to $h = 3\,500$. This leads to around 860 and 720 observations to the left and right of the threshold, respectively. It is an attractive feature of the approach that X can be made high-dimensional to make CAR more plausible while keeping the final estimation as a low-dimensional local linear regression which has desirable properties for inference.

The relationship between the estimated pseudo-outcomes and the income in Wave 3 is depicted in the right panel of Figure 6.5. This shows a clear discontinuity in treatment probability at w_0 indicating that an RDD is indeed applicable. In the left-panel, one sees that the expected outcome seems to increase with salary until around £60,000 after which it appears constant. The level also appears constant until around £30,000 which could be an indication that the level below £20,000 is artificially high due to EMA. The middle panel focuses on a neighborhood of w_0, and also seems to indicate a discontinuity for the education outcomes although its statistical significance is less clear. A desirable feature of using pseudo-outcomes for RDD is that the regression discontinuity can still be plotted when data is censored. Such graphical tools are important for RDD analyses, see Imbens & Lemieux (2008).

Algorithm 1 leads to the estimated value and standard error

$$\hat{\tau} = 0.703, \qquad \mathrm{SE}(\hat{\tau}) = 0.614,$$

resulting in a p-value of around 0.25 and thus not reaching statistical significance at conventional significance levels.

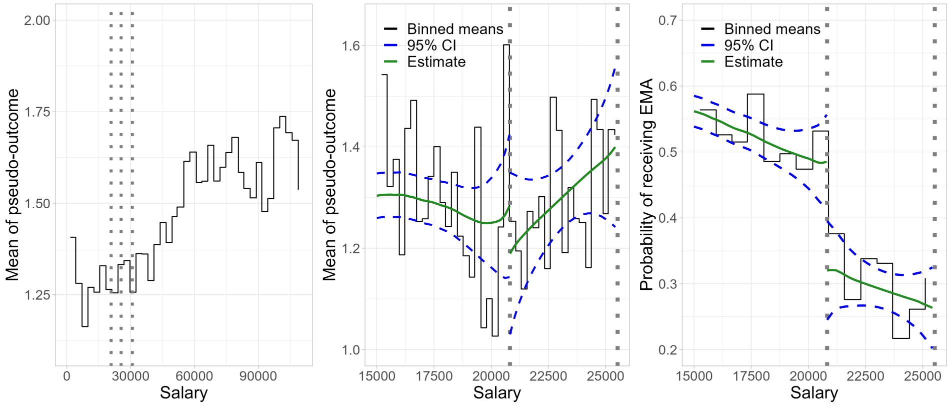

Figure 6.5: ***Left Panel:*** *Binned-means over the pseudo-outcomes for Y and EMA thresholds for income in Wave 3.* ***Middle Panel:*** *Binned-means of the pseudo-outcomes for Y, estimate of the conditional mean, 95% confidence intervals, and EMA thresholds for Wave 3 incomes around the threshold for receiving high EMA.* ***Right Panel:*** *The same as the middle panel with pseudo-outcomes for A.*

To get a feeling for the sensitivity of the result with respect to the bandwidth, the estimation was repeated with $b = 4\,704$ and $b = 2\,352$ which is the distance to the next EMA threshold and half that distance, respectively. The estimated values are in this case $\{\hat{\tau}, \mathrm{SE}(\hat{\tau})\} = (0.391, 0.588)$ and $\{\hat{\tau}, \mathrm{SE}(\hat{\tau})\} = (1.197, 0.697)$, respectively. Thus, the absolute size of the estimate changes considerably, but the effect remains large and positive. Note that the 90% significance level is reached for the smaller bandwidth. Making Figure 6.5 with these alternative bandwidths (not shown) indicates over- and undersmooth, respectively, and the original estimate hence seems to be the most reliable.

The proposed methods yield wider confidence intervals than those in Rahman (2014), leading to statistically insignificant results. This likely better reflects the uncertainty in the estimate since the assumptions imposed here are substantially weaker. If the effect is genuine, an additional 0.7 years of education would be a large effect and would be another piece of evidence that EMA was successful in its initial aim of keeping YP in education. Transparency regarding uncertainty is important for policymakers when assessing findings and deciding if more data is needed before making decisions. The p-value suggests a larger sample could be beneficial in clarifying the effect, or statistical power could be increased by using data from the multiple EMA thresholds, as discussed, though at the cost of some internal validity. This is left for future work.

Remark 6.5.2. (Model extensions.)
A slightly more sophisticated model would have accommodated the fact that interviews took place over a few months rather than simultaneously, using that the interview month is available from the data to model C on a monthly rather than

yearly grid. The effect of this is however expected to be minor in the present study. Additionally, one could have weakened the assumption that (C, X^C) is a CAR of X by including more outcomes from Wave 4 in X, but nuisance estimators would then have to model the entire distribution of X at Wave 4 given X^0. $\triangledown$

Remark 6.5.3. (RDD with survival data.)
The use of RDD for survival data has been studied in Adeleke et al. (2022) for an accelerated failure time model. The methods proposed in this paper seem to be the first that allow for nonparametric inference for an RDD when data is censored even for the survival setting. Note that the outcome Y specified above cannot be represented as survival data since some leave school in Wave 4 but return in Wave 5. $\triangledown$

Our approach generalizes the class of problems where an RDD is applicable. This could be a valuable tool in exploring long-ranging consequences of policies in cases where a longitudinal no unmeasured confounding assumption might be unsuitable but coarsening at random for the censoring mechanism is believable. Many existing datasets could likely be analyzed using methods similar to those employed in this section, and the availability of the methods might also incentivize more studies to be on a form where a longitudinal RDD could be applied. The LSYPE data for example allows one to explore several long-term consequences of EMA. Waves 6-8 enable examination of university attendance and choice of subjects, and Wave 8 contains information on labour market outcomes. Other linked administrative data are also available, though under stricter access requirements. Similar datasets are however available during Covid-19 years (2020-2021), allowing one to explore the long-term effect of EMA on self-reported health, amount of hours worked, trust in the government etc. Here it might be natural to let time 0 be Wave 4 such that treatment is a baseline covariate and Remark 6.2.5 may be used. This is left to future work.

Acknowledgments and declarations of interest

No competing interest is declared. This research has partly been funded by the Innovation Fund Denmark (IFD) under File No. 1044-00144B. Significant parts of the research were conducted during a visit to the Department of Mathematical Statistics at Stockholm University. The author gratefully acknowledges the hospitality of Filip Lindskog and Mathias Lindholm and thanks them for many fruitful discussions. The author also thanks their supervisor Christian Furrer for general feedback and Niels Richard Hansen for helpful discussions on the smoothing effect of averaging estimators.

6.A Derivation of the efficient influence function

Sample means of the IPCW pseudo-outcomes

$$Y^{\circ}(C, X^C) = \frac{Y(X)1_{(C\geq\eta)}}{\mathbb{P}(C \geq \eta \mid X)}$$

provide an estimator of $\mathbb{E}[Y(X)]$. Furthermore, this estimator is a regular and linear (hence also asymptotically linear) estimator of $\mathbb{E}[Y(X)]$ with influence function $Y^{\circ}(C, X^C) - \mathbb{E}[Y(X)]$, see e.g. Section 3 in Tsiatis (2006). The efficient influence function is any influence function subtracted its projection in $L^2(\Omega, \mathcal{F}, \mathbb{P})$ onto the CAR tangent space which may be found using Van der Vaart (2004) or Section 3.4 of Van der Laan & Robins (2003). Let $M(\mathrm{d}u) = \mathrm{d}1_{(C\leq u)} - 1_{(C\geq u)}\mathbb{P}_1(C \in \mathrm{d}u \mid X)/\mathbb{P}_1(C \geq u \mid X)$. The aforementioned projection is then

$$\begin{aligned}\int & \mathbb{E}[Y^{\circ}(u, X^u) - \mathbb{E}[Y(X)] \mid C > u, X^u] \\ & - \mathbb{E}[Y^{\circ}(C, X^C) - \mathbb{E}[Y(X)] \mid C > u, X^u]\, M(\mathrm{d}u)\end{aligned}$$

when the conditional expectations are taken to be 0 if $\mathbb{P}(C > u \mid X^u) = 0$. Thus, there are only contributions on $[0, \eta)$. The marginals expectations cancel and since $1_{(u\geq\eta)} = 0$ for $u \in [0, \eta)$ the only remaining term is $-\int_{[0,\eta)} \mathbb{E}[Y^{\circ}(C, X^C) \mid C > u, X^u]\, M(\mathrm{d}u)$. Note

$$\mathbb{E}[Y^{\circ}(C, X^C) \mid C > u, X^u] = \frac{\mathbb{E}[Y(X)1_{(C\geq\eta)}/\mathbb{P}(C \geq \eta \mid X) \mid X^u]}{\mathbb{P}(C > u \mid X^u)}$$

since $1_{(C\geq\eta)}1_{(C>u)} = 1_{(C\geq\eta)}$ for $u \in [0, \eta)$. Using the tower-property when conditioning on X in the numerator gives $\mathbb{E}[Y(X) \mid X^u]/\mathbb{P}(C > u \mid X^u)$.

6.B Proof of Theorem 6.2.3

The proof proceeds in two parts.

6.B.1 Proof of conditional expectation result

Proof. Write

$$\begin{aligned}Y^*_{\mathbb{P}_1,\mathbb{P}_2}(C, X^C) = & \frac{Y(X)1_{(C\geq\eta)}}{\mathbb{P}_1(C \geq \eta \mid X)} + \frac{\mathbb{E}_2[Y(X) \mid X^u]}{\mathbb{P}_1(C > u \mid X)}\Big|_{u=C} \times 1_{(C<\eta)} \\ & - \int_{[0,\eta)} 1_{(C\geq u)} \frac{\mathbb{E}_2[Y(X) \mid X^u]}{\mathbb{P}_1(C > u \mid X)} \frac{\mathbb{P}_1(C \in \mathrm{d}u \mid X)}{\mathbb{P}_1(C \geq u \mid X)}.\end{aligned}$$

The conditional expectation given W of each term is treated separately. The strategy is to write the expectation in terms of $X \mid W$ and $C \mid X$.

$$\mathbb{E}\left[\frac{Y(X)1_{(C\geq\eta)}}{\mathbb{P}_1(C \geq \eta \mid X)} \mid W\right] = \mathbb{E}\left[\frac{Y(X)\mathbb{P}(C \geq \eta \mid X)}{\mathbb{P}_1(C \geq \eta \mid X)} \mid W\right]$$

by the tower-property when conditioning on X. Similarly,

$$\mathbb{E}\left[\frac{\mathbb{E}_2[Y(X) \mid X^u]}{\mathbb{P}_1(C > u \mid X)}\Big|_{u=C} \times 1_{(C<\eta)} \mid W\right]$$
$$= \mathbb{E}\left[\int_{[0,\eta)} \frac{\mathbb{E}_2[Y(X) \mid X^u]}{\mathbb{P}_1(C > u \mid X)} \mathbb{P}(C \in \mathrm{d}u \mid X) \mid W\right]$$

and

$$\mathbb{E}\left[\int_{[0,\eta)} 1_{(C \geq u)} \frac{\mathbb{E}_2[Y(X) \mid X^u]}{\mathbb{P}_1(C > u \mid X)} \frac{\mathbb{P}_1(C \in \mathrm{d}u \mid X)}{\mathbb{P}_1(C \geq u \mid X)} \mid W\right]$$
$$= \mathbb{E}\left[\int_{[0,\eta)} \frac{\mathbb{P}(C \geq u \mid X)\mathbb{E}_2[Y(X) \mid X^u]}{\mathbb{P}_1(C > u \mid X)} \frac{\mathbb{P}_1(C \in \mathrm{d}u \mid X)}{\mathbb{P}_1(C \geq u \mid X)} \mid W\right].$$

Note that

$$\frac{\mathbb{P}(C \in \mathrm{d}u \mid X)}{\mathbb{P}_1(C > u \mid X)} = \gamma(u \mid X) \frac{\mathbb{P}(C \geq u \mid X)}{\mathbb{P}_1(C > u \mid X)} \,\mathrm{d}\mu(u)$$

and

$$\frac{\mathbb{P}(C \geq u \mid X)}{\mathbb{P}_1(C > u \mid X)} \frac{\mathbb{P}_1(C \in \mathrm{d}u \mid X)}{\mathbb{P}_1(C \geq u \mid X)} = \gamma_1(u \mid X) \frac{\mathbb{P}(C \geq u \mid X)}{\mathbb{P}_1(C > u \mid X)} \,\mathrm{d}\mu(u).$$

Thus,

$$\mathbb{E}[Y^*_{\mathbb{P}_1,\mathbb{P}_2}(C, X^C) - Y(X) \mid W]$$
$$= \mathbb{E}\Bigg[\frac{\mathbb{P}(C \geq \eta \mid X)}{\mathbb{P}_1(C \geq \eta \mid X)} Y(X) - Y(X)$$
$$+ \int_{[0,\eta)} \mathbb{E}_2[Y(X) \mid X^u]\{\gamma(u \mid X) - \gamma_1(u \mid X)\} \frac{\mathbb{P}(C \geq u \mid X)}{\mathbb{P}_1(C > u \mid X)} \,\mathrm{d}\mu(u) \mid W\Bigg].$$

Using p. 868 of Shorack & Wellner (1986), write

$$\frac{\mathbb{P}(C \geq \eta \mid X)}{\mathbb{P}_1(C \geq \eta \mid X)} Y(X) - Y(X) = \int_{[0,\eta)} Y(X) \,\mathrm{d}\left\{\frac{\mathbb{P}(C > u \mid X)}{\mathbb{P}_1(C > u \mid X)}\right\}.$$

Integration by parts for finite variation functions, see p. 868 of Shorack & Wellner (1986), implies

$$\mathrm{d}\left\{\frac{\mathbb{P}(C > u \mid X)}{\mathbb{P}_1(C > u \mid X)}\right\} = -\frac{\mathbb{P}(C \in \mathrm{d}u \mid X)}{\mathbb{P}_1(C > u \mid X)}$$
$$+ \frac{\mathbb{P}(C \geq u \mid X)}{\mathbb{P}_1(C \geq u \mid X)\mathbb{P}_1(C > u \mid X)} \mathbb{P}_1(C \in \mathrm{d}u \mid X)$$

using Assumption 6.2.2 for $\mathbb{P}_1$. By the previous calculations, one therefore obtains

$$\int_{[0,\eta)} Y(X) \,\mathrm{d}\left\{\frac{\mathbb{P}(C > u \mid X)}{\mathbb{P}_1(C > u \mid X)}\right\}$$
$$= -\int_{[0,\eta)} Y(X)\{\gamma(u \mid X) - \gamma_1(u \mid X)\} \frac{\mathbb{P}(C \geq u \mid X)}{\mathbb{P}_1(C > u \mid X)} \,\mathrm{d}\mu(u)$$

Because of CAR, it holds that

$$\mathbb{P}_1(C > u \mid X) = 1 - \int_{[0,u]} r_1(s \mid X)\,\mathrm{d}\mu(s) = 1 - \int_{[0,u]} \tilde{r}_1(s, X^s)\,\mathrm{d}\mu(s)$$

so $\mathbb{P}_1(C > u \mid X) = \mathbb{P}_1(C > u \mid X^u)$ by the tower property. Similar calculations hold for $\mathbb{P}_1(C \geq u \mid X)$ and $\mathbb{P}(C \geq u \mid X)$. Hence,

$$\begin{aligned}
&\mathbb{E}\left[\int_{[0,\eta)} Y(X)\{\gamma(u \mid X) - \gamma_1(u \mid X)\}\frac{\mathbb{P}(C \geq u \mid X)}{\mathbb{P}_1(C > u \mid X)}\,\mathrm{d}\mu(u) \mid W\right] \\
&= \mathbb{E}\left[\int_{[0,\eta)} \mathbb{E}[Y(X) \mid X^u]\{\gamma(u \mid X) - \gamma_1(u \mid X)\}\frac{\mathbb{P}(C \geq u \mid X)}{\mathbb{P}_1(C > u \mid X)}\,\mathrm{d}\mu(u) \mid W\right]
\end{aligned}$$

using Fubini to take the expectation inside the integral, then tower with X^u and use Fubini to take the expectation outside again. Collecting the results leads to the desired expression. □

6.B.2 Proof of conditional variance result

Proof. For shorthand, write $Y^* = Y^*_{\mathbb{P},\mathbb{P}}(C, X^C)$ and $Y^\circ = Y^\circ(C, X^C)$ recalling the IPCW notation from Section 6.A of the Supplementary material. The first part of the proof generalizes the calculations from Proposition 5 of Suzukawa (2004) and S.5.3 in the Supplementary Material of Steingrimsson et al. (2019). Note

$$\mathrm{Var}[Y^* \mid W] = \mathbb{E}[(Y^*)^2 \mid W] - \mathbb{E}[Y^* \mid W]^2.$$

By the first part of Theorem 6.2.3, it holds that $\mathbb{E}[Y^* \mid W] = \mathbb{E}[Y \mid W]$. For the other term, expanding the square gives $(Y^*)^2 = R^{(1)} + R^{(2)} + R^{(3)}$ where

$$\begin{aligned}
R^{(1)} &= (Y^\circ)^2, \\
R^{(2)} &= \left[\int_{[0,\eta)} \frac{\mathbb{E}[Y \mid X^u]}{\mathbb{P}(C > u \mid X)}\left\{\mathrm{d}1_{(C \leq u)} - 1_{(C \geq u)}\frac{\mathbb{P}(C \in \mathrm{d}u \mid X)}{\mathbb{P}(C \geq u \mid X)}\right\}\right]^2, \\
R^{(3)} &= 2Y^\circ \int_{[0,\eta)} \frac{\mathbb{E}[Y \mid X^u]}{\mathbb{P}(C > u \mid X)}\left\{\mathrm{d}1_{(C \leq u)} - 1_{(C \geq u)}\frac{\mathbb{P}(C \in \mathrm{d}u \mid X)}{\mathbb{P}(C \geq u \mid X)}\right\}.
\end{aligned}$$

Straigtforward calculations give

$$\begin{aligned}
\mathbb{E}[R^{(1)} \mid W] &= \mathbb{E}\left[\frac{Y^2}{\mathbb{P}(C \geq \eta \mid X)} \mid W\right], \\
\mathbb{E}[R^{(3)} \mid W] &= -2\mathbb{E}\left[Y \int_{[0,\eta)} \frac{\mathbb{E}[Y \mid X^u]}{\mathbb{P}(C > u \mid X)}\frac{\mathbb{P}(C \in \mathrm{d}u \mid X)}{\mathbb{P}(C \geq u \mid X)} \mid W\right].
\end{aligned}$$

Expanding the square gives $R^{(2)} = R^{(2.1)} + R^{(2.2)} + R^{(2.3)}$ for

$$\begin{aligned} R^{(2.1)} &= \left\{ \int_{[0,\eta)} \frac{\mathbb{E}[Y \mid X^u]}{\mathbb{P}(C > u \mid X)} \, \mathrm{d}1_{(C \leq u)} \right\}^2, \\ R^{(2.2)} &= \left\{ \int_{[0,\eta)} \frac{\mathbb{E}[Y \mid X^u]}{\mathbb{P}(C > u \mid X)} 1_{(C \geq u)} \frac{\mathbb{P}(C \in \mathrm{d}u \mid X)}{\mathbb{P}(C \geq u \mid X)} \right\}^2, \\ R^{(2.3)} &= -2 \int_{[0,\eta)} \frac{\mathbb{E}[Y \mid X^u]}{\mathbb{P}(C > u \mid X)} \, \mathrm{d}1_{(C \leq u)} \\ &\quad \times \int_{[0,\eta)} \frac{\mathbb{E}[Y \mid X^u]}{\mathbb{P}(C > u \mid X)} 1_{(C \geq u)} \frac{\mathbb{P}(C \in \mathrm{d}u \mid X)}{\mathbb{P}(C \geq u \mid X)}. \end{aligned}$$

Note

$$\begin{aligned} \mathbb{E}[R^{(2.2)} \mid W] = \mathbb{E}\bigg[\int_{[0,\eta)^2} & \frac{\mathbb{E}[Y \mid X^u]}{\mathbb{P}(C > u \mid X)} \frac{\mathbb{E}[Y \mid X^v]}{\mathbb{P}(C > v \mid X)} 1_{(C \geq u \vee v)} \\ & \times \frac{\mathbb{P}(C \in \mathrm{d}u \mid X)}{\mathbb{P}(C \geq u \mid X)} \frac{\mathbb{P}(C \in \mathrm{d}v \mid X)}{\mathbb{P}(C \geq v \mid X)} \mid W \bigg] \end{aligned}$$

by Fubini's theorem. By symmetry, this is twice the contribution where the indicator $1_{(C \geq v)}$ is replaced by $1_{(u \geq v)}$. Inserting this and towering on X gives

$$\begin{aligned} \mathbb{E}[R^{(2.2)} \mid W] = 2\mathbb{E}\bigg[& \int_{[0,\eta)} \frac{\mathbb{E}[Y \mid X^u]}{\mathbb{P}(C > u \mid X)} \\ & \times \left\{ \int_{[0,u]} \frac{\mathbb{E}[Y \mid X^v]}{\mathbb{P}(C > v \mid X)} \frac{\mathbb{P}(C \in \mathrm{d}v \mid X)}{\mathbb{P}(C \geq v \mid X)} \right\} \mathbb{P}(C \in \mathrm{d}u \mid X) \mid W \bigg]. \end{aligned}$$

Straightforward calculations thus imply $\mathbb{E}[R^{(2.2)} \mid W] = -\mathbb{E}[R^{(2.3)} \mid W]$ so these terms cancel. Finally, note

$$\mathbb{E}[R^{(2.1)} \mid W] = \mathbb{E}\left[\int_{[0,\eta)} \frac{\mathbb{E}[Y \mid X^u]^2}{\mathbb{P}(C > u \mid X)^2} \mathbb{P}(C \in \mathrm{d}u \mid X) \mid W \right]$$

Collecting the results gives

$$\begin{aligned} \mathrm{Var}[Y^* \mid W] = \mathbb{E}\bigg[& \frac{Y^2}{\mathbb{P}(C \geq \eta \mid X)} + \int_{[0,\eta)} \frac{\mathbb{E}[Y \mid X^u]^2}{\mathbb{P}(C > u \mid X)^2} \mathbb{P}(C \in \mathrm{d}u \mid X) \\ & - 2Y \int_{[0,\eta)} \frac{\mathbb{E}[Y \mid X^u]}{\mathbb{P}(C > u \mid X)} \frac{\mathbb{P}(C \in \mathrm{d}u \mid X)}{\mathbb{P}(C \geq u \mid X)} \mid W \bigg] - \mathbb{E}[Y \mid W]^2. \end{aligned}$$

Note that $Y^2 / \mathbb{P}(C \geq \eta \mid X) = \int_{[0,\eta)} Y^2 \, \mathrm{d}\{1/\mathbb{P}(C > u \mid X)\} + Y^2$ and integration by parts implies

$$\mathrm{d}\left\{ \frac{1}{\mathbb{P}(C > u \mid X)} \right\} = \frac{\mathbb{P}(C \in \mathrm{d}u \mid X)}{\mathbb{P}(C \geq u \mid X)\mathbb{P}(C > u \mid X)}.$$

Inserting this and collecting the integral terms implies

$$\mathrm{Var}[Y^* \mid W] - \mathrm{Var}[Y \mid W] = \mathbb{E}\Bigg[\int_{[0,\eta)} \bigg\{ \frac{Y^2}{\mathbb{P}(C \geq u \mid X)} + \frac{\mathbb{E}[Y \mid X^u]^2}{\mathbb{P}(C > u \mid X)} - \frac{2\mathbb{E}[Y \mid X^u]^2}{\mathbb{P}(C \geq u \mid X)} \bigg\} \frac{\mathbb{P}(C \in \mathrm{d}u \mid X)}{\mathbb{P}(C > u \mid X)} \mid W \Bigg].$$

By writing $\mathbb{P}(C \in \mathrm{d}u \mid X) = r(u \mid X)\,\mathrm{d}\mu(u)$, one may take the expectation inside the integral and can then tower on X^u and then take the expectation outside the integral again, leading to Y^2 being replaced by $\mathbb{E}[Y^2 \mid X^u]$. By bounding $\mathbb{E}[Y \mid X^u]^2/\mathbb{P}(C > u \mid X) \geq \mathbb{E}[Y \mid X^u]^2/\mathbb{P}(C \geq u \mid X)$ and using Jensen's inequality for conditional expectations to bound $\mathbb{E}[Y^2 \mid X^u] \geq \mathbb{E}[Y \mid X^u]^2$ gives the desired conclusion. □

Similarly to Theorem 3.1 in Steingrimsson et al. (2019), one could further have shown that $\mathrm{Var}[Y^*_{\mathbb{P},\mathbb{P}_2}(C, X^C) \mid W] \geq \mathrm{Var}[Y^* \mid W]$ so using a misspecified outcome distribution leads to larger variance of the pseudo-outcomes. This result is however not directly useful for our purposes and is hence omitted.

6.C Efficient influence function in Remark 6.2.5

The efficient influence function can be derived using similar arguments to those in Section 6.A of the Supplementary material. Define the inverse probability weighted pseudo-outcomes for treatment a as

$$Y^\circ(a, C, X^C) = \frac{Y(X)1_{(C \geq \eta)}1_{(A=a)}}{\mathbb{P}(C \geq \eta \mid X)\mathbb{P}(A = a \mid W)}.$$

Sample means of these pseudo-outcomes provide an estimator of $\mathbb{E}[\mathbb{E}[Y(X) \mid W, A = a]]$, which is the population version of the estimand of interest $\mathbb{E}[Y(X) \mid W, A = a]$. For these estimands, $(C, X^{(a)}, A)$ is the complete data and (C, X^C) is the observed data. Following the arguments in Rytgaard et al. (2022), the projection onto the relevant tangent space is given by

$$\begin{aligned}
&\int_{[0,\eta)} \mathbb{E}[Y^\circ(a, u, X^C) \mid C > u, X^u] - \mathbb{E}[Y^\circ(a, C, X^C) \mid C > u, X^u]\, M(\mathrm{d}u) \\
&+ \sum_{k \in \{0,1\}} \big(\mathbb{E}[Y^\circ(a, C, X^C) \mid W, A = k] - \mathbb{E}[Y^\circ(a, C, X^C) \mid W]\big)\, 1_{(A=k)} \\
&= -\frac{1_{(A=a)}}{\mathbb{P}(A = a \mid W)} \int_{[0,\eta)} \frac{\mathbb{E}[Y(X) \mid X^u]}{\mathbb{P}(C > u \mid X)} M(\mathrm{d}u) \\
&\quad + \frac{1_{(A=a)} - \mathbb{P}(A = a \mid W)}{\mathbb{P}(A = a \mid W)} \mathbb{E}[Y(X) \mid W, A = a]
\end{aligned}$$

with M defined as in Section 6.A of the Supplementary material. This result also appears in Section 6.4.3 of Van der Laan & Robins (2003) when $Y(X) = 1_{(T \leq t)}$ for a survival time T.

6.D Proof of Proposition 6.3.5

Proof. Write

$$\hat{m}(w) - m(w) = \hat{m}(w) - \tilde{m}(w) + \tilde{m}(w) - m(w).$$

To show stability of linear smoothers, note that

$$\begin{aligned}
&d_{w,D^{2n}}(0, \mathrm{Var}[Y^*(C, X^C) \mid W = \bullet\,]) \\
&= \sum_{i=1}^{n} \left\{ \frac{p_i(w; W^n)^2}{\sum_{j=1}^{n} p_j(w; W^n)^2} \mathrm{Var}[Y^*(C, X^C) \mid W = W_i]^2 \right\} \\
&\geq \inf_{z \in \{W_1, \ldots, W_n\}} \mathrm{Var}[Y^*(C, X^C) \mid W = z]^2 \\
&\geq \inf_{z} \mathrm{Var}[Y(X) \mid W = z]^2.
\end{aligned}$$

where the last inequality follows from the second part of Theorem 6.2.3. It therefore holds that $d_{w,D^{2n}}(0, \mathrm{Var}[Y^*(C, X^C) \mid W = \bullet\,])^{-1}$ is bounded and thus also $O_{\mathbb{P}}(1)$. This result combined with condition (ii) gives stability of the linear smoother. Therefore $\hat{m}(w) - \tilde{m}(w) = \hat{\mathbb{E}}_n[\hat{b}(W; D_1^n) \mid D_1^n, W = w] + o_{\mathbb{P}}(n^{-\alpha})$.

Introduce the stochastic norm

$$\|f(u, X; D_1^n)\|_{3,z,D_1^n} = \left\{ \int_{[0,\eta)} \|f(u, X; D_1^n)\|_{z,D_1^n}^2 \, \mathrm{d}\mu(u) \right\}^{1/2}.$$

By the first part of Theorem 6.2.3,

$$\begin{aligned}
\hat{b}(z; D_1^n) = \int_{\mathcal{X} \times [0,\eta)} &\Big\{ \mathbb{E}[Y(X) \mid x^u] - \hat{\mathbb{E}}_{2,n}[Y(X) \mid x^u] \Big\} \{\hat{\gamma}_{1,n}(u \mid x) - \gamma(u \mid x)\} \\
&\frac{\mathbb{P}(C \geq u \mid x)}{\hat{\mathbb{P}}_{1,n}(C > u \mid x)} \mathbb{P}(X \in \mathrm{d}x \mid W = z) \otimes \mathrm{d}\mu(u)
\end{aligned}$$

so

$$\begin{aligned}
|\hat{b}(z; D_1^n)| \leq \varepsilon^{-1} &\|\mathbb{E}[Y(X) \mid X^u] - \hat{\mathbb{E}}_{2,n}[Y(X) \mid X^u]\|_{3,z,D_1^n} \\
&\times \|\hat{\gamma}_{1,n}(u \mid X) - \gamma(u \mid X)\|_{3,z,D_1^n}
\end{aligned}$$

by taking the absolute value onto the integrand, using positivity, and then employing the Cauchy-Schwarz inequality. Note

$$\begin{aligned}
&\left|\hat{\mathbb{E}}_n[\hat{b}(W; D_1^n) \mid D_1^n, W = w]\right| \\
&\leq \sum_{i=1}^{n} |p_i(w; W^n)| \times |\hat{b}(W_i; D_1^n)| \\
&\leq \varepsilon^{-1} \sum_{i=1}^{n} |p_i(w; W^n)|^{1/2} \|\mathbb{E}[Y(X) \mid X^u] - \hat{\mathbb{E}}_{2,n}[Y(X) \mid X^u]\|_{3,W_i,D_1^n} \\
&\qquad |p_i(w; W^n)|^{1/2} \|\hat{\gamma}_{1,n}(u \mid X) - \gamma(u \mid X)\|_{3,W_i,D_1^n}.
\end{aligned}$$

By the Cauchy-Schwarz inequality

$$
\begin{aligned}
&\left|\hat{\mathbb{E}}_n[\hat{b}(W; D_1^n) \mid D_1^n, W = w]\right| \\
&\leq \varepsilon^{-1} \left\{ \sum_{i=1}^{n} |p_i(w; W^n)| \times \|\mathbb{E}[Y(X) \mid X^u] - \hat{\mathbb{E}}_{2,n}[Y(X) \mid X^u]\|_{3,W_i,D_1^n}^2 \right\}^{1/2} \\
&\qquad \times \left\{ \sum_{i=1}^{n} |p_i(w; W^n)| \times \|\hat{\gamma}_{1,n}(u \mid X) - \gamma(u \mid X)\|_{3,W_i,D_1^n}^2 \right\}^{1/2} \\
&= \sum_{i=1}^{n} \frac{|p_i(w; W^n)|}{\varepsilon} \times \|\mathbb{E}[Y(X) \mid X^u] - \hat{\mathbb{E}}_{2,n}[Y(X) \mid X^u]\|_{2,w,D^{2n}} \\
&\quad \times \|\hat{\gamma}_{1,n}(u \mid X) - \gamma(u \mid X)\|_{2,w,D^{2n}}
\end{aligned}
$$

where the final equality follows from the definition of the norm.

Note that the sum of the absolute weights is $O_{\mathbb{P}}(1)$ by (iii). By Assumption (iv) and (v), the right hand side is thus $O_{\mathbb{P}}(n^{-\alpha_1-\alpha_2})$. To obtain the oracle rate, this term should be $o_{\mathbb{P}}(n^{-\alpha})$. This is satisfied if $n^{\alpha_1+\alpha_2} > n^{\alpha}$ or equivalently $\alpha_1+\alpha_2 > \alpha$, which holds by (vi). It thus holds that $\hat{m}(w) - \tilde{m}(w) = o_{\mathbb{P}}(n^{-\alpha})$ which implies $\hat{m}(w) - m(w) = O_{\mathbb{P}}(n^{-\alpha})$ as desired. □

6.E Proof of Proposition 6.3.9

Proof. Note that

$$
\hat{m}^{\mathrm{CF}}(w) - m(w) = \frac{1}{K}\sum_{k=1}^{K}\{\tilde{m}_k(w) - m(w)\} + \frac{1}{K}\sum_{k=1}^{K}\{\hat{m}_k(w) - \tilde{m}_k(w)\}.
$$

Each $\tilde{m}_k(w) - m(w)$ can be analyzed analogously to $\tilde{m}(w) - m(w)$ but just using n/K observations instead of n. Therefore $n^{\alpha}\{\tilde{m}_k(w) - m(w)\} \to \mathcal{N}(K^{\alpha}\mu, K^{2\alpha}\sigma^2)$ in distribution since $n^{\alpha} = K^{\alpha}(n/K)^{\alpha}$. Furthermore, since $\tilde{m}_k(w) - m(w)$ for different values of k are independent, one obtains $n^{\alpha}[K^{-1}\sum_{k=1}^{K}\{\tilde{m}_k(w)-m(w)\}] \to \mathcal{N}(K^{\alpha}\mu, K^{2\alpha-1}\sigma^2)$ in distribution. For the second sum, note that each term $\hat{m}_k(w)-\tilde{m}_k(w)$ can be analyzed analogously to the sample split version $\hat{m}(w)-\tilde{m}(w)$. Under the assumptions from Proposition 6.3.5, it thus holds that $\hat{m}_k(w) - \tilde{m}_k(w) = o_{\mathbb{P}}(n^{-\alpha})$ so also $1/K\sum_{k=1}^{K}\{\hat{m}_k(w) - \tilde{m}_k(w)\} = o_{\mathbb{P}}(n^{-\alpha})$. Slutsky's lemma then implies, still under the assumptions from Proposition 6.3.5, that

$$
n^{\alpha}\{\hat{m}^{\mathrm{CF}}(w) - m(w)\} \to \mathcal{N}(K^{\alpha}\mu, K^{2\alpha-1}\sigma^2)
$$

in distribution. □

6.F Figures

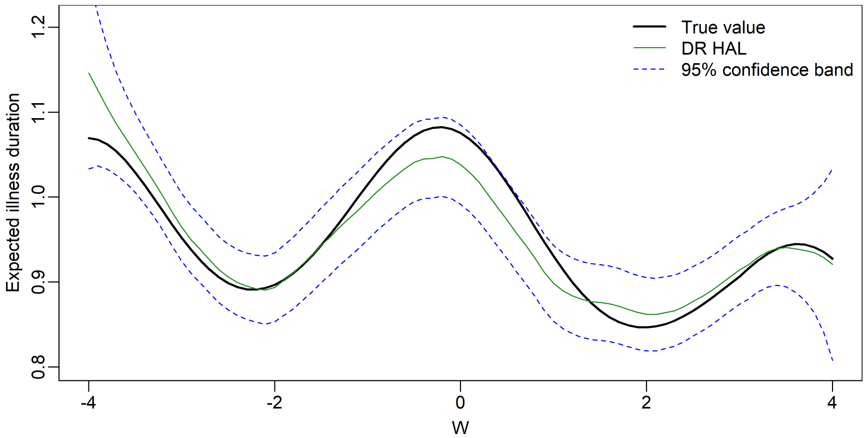

Figure 6.6: *Estimates using two-fold cross-fitted doubly robust pseudo-outcomes with nuisance parameters estimated by HAL, the true curve, and pointwise 95% confidence bands outputted by* `lprobust` *for a single simulation.*

Bibliography

Adékambi, F. and Christiansen, M. C. (2017). Integral and differential equations for the moments of multistate models in health insurance. *Scandinavian Actuarial Journal* **2017**, 29–50. doi: 10.1080/03461238.2015.1058854.

Adeleke, M. O., Baio, G., and O'Keeffe, A. G. (2022). Regression Discontinuity Designs for Time-to-Event Outcomes: An Approach using Accelerated Failure Time Models. *Journal of the Royal Statistical Society Series A: Statistics in Society* **185**(3), 1216–1246. doi: 10.1111/rssa.12812.

Andersen, P.K., Borgan, O., Gill, R.D., & Keiding, N. (1993). *Statistical Models Based on Counting Processes*. New York, United States: Springer. doi: 10.1007/978-1-4612-4348-9.

Andersen, N. T. and Dobric, V. (1987). The Central Limit Theorem for Stochastic Processes. *The Annals of Probability* **15**, 164–177. doi: 10.1214/aop/1176992262.

Andersen, P. K., Klein, J. P., and Rosthøj, S. (2003). Generalised linear models for correlated pseudo-observations, with applications to multi-state models. *Biometrika* **90**, 15–27. doi: 10.1093/biomet/90.1.15.

Andersen, P. K., Syriopoulou, E., and Parner, E. T. (2017). Causal inference in survival analysis using pseudo-observations. *Statistics in Medicine* **36**(17), 2669–2681. doi: 10.1002/sim.7297.

Antonio, K. and Plat, R. (2014). Micro-level stochastic loss reserving for general insurance. *Scandinavian Actuarial Journal* **2014**(7), 649–669. doi: 10.1080/03461238.2012.755938.

Arcones, M. A. and Giné, E. (1992). On the bootstrap of M-estimators and other statistical functionals. In: LePage, L. and Billard, L. Eds. *Exploring the Limits of Bootstrap*, 13–47. New York, United States: Wiley.

Aro, H., Djehiche, B., and Löfdahl, B. (2015). Stochastic modelling of disability insurance in a multi-period framework. *Scandinavian Actuarial Journal* **2015**, 88–106. DOI: 10.1080/03461238.2013.779594.

BADESCU, A. L., LIN, X. S., and TANG, D. (2016). A marked Cox model for the number of IBNR claims: Theory. *Insurance: Mathematics and Economics* **69**, 29–37. DOI: 10.1016/j.insmatheco.2016.03.016.

BADESCU, A. L., LIN, X. S., and TANG, D. (2019). A MARKED COX MODEL FOR THE NUMBER OF IBNR CLAIMS: ESTIMATION AND APPLICATION. *ASTIN Bulletin* **49**(3), 709–739. DOI: 10.1017/asb.2019.15.

BANG, H. and ROBINS, J. M. (2005). Doubly robust estimation in missing data and causal inference models. *Biometrics* **61**(4), 962–973. DOI: 10.1111/j.1541-0420.2005.00377.x.

BARABAS, C., VIRZA, M., DINAKAR, K., ITO, J., and ZITTRAIN, J. (2018). Interventions over Predictions: Reframing the Ethical Debate for Actuarial Risk Assessment. *Proceedings of the 1st Conference on Fairness, Accountability and Transparency, PMLR* **81**, 62–76.

BETTONVILLE, C., D'OULTREMONT, L., DENUIT, M., TRUFIN, J., and VAN OIRBEEK, R. (2021). Matrix calculation for ultimate and 1-year risk in the Semi-Markov individual loss reserving model. *Scandinavian Actuarial Journal* **2021**(5), 380–407. DOI: 10.1080/03461238.2020.1848912.

BECKER, N. G. and CUI, J. S. (1997). Estimating a delay distribution from incomplete data, with application to reporting lags for AIDS cases. *Statistics in Medicine* **16**(20), 2339–2347. DOI: 10.1002/(SICI)1097-0258(19971030)16:20¡2339::AID-SIM648¿3.0.CO;2-E.

BENKESER, D. and VAN DER LAAN, M. (2016). The Highly Adaptive Lasso Estimator. *2016 IEEE international conference on data science and advanced analytics (DSAA)*, 689–696. DOI: 10.1109/DSAA.2016.93.

BICKEL, P. J., KLAASSEN, C. A. J., RITOV, Y., and WELLNER, J. A. (1998). *Efficient and Adaptive Estimation for Semiparametric Models*. Berlin, Germany: Springer.

BISCHOFBERGER, S. M., HIABU, M., and ISAKSON, A. (2020). Continuous chain-ladder with paid data. *Scandinavian Actuarial Journal* **2020**(6), 477–502. DOI: 10.1080/03461238.2019.1694973.

BLADT, M., ASMUSSEN, S., and STEFFENSEN, M. (2020). Matrix representations of life insurance payments. *European Actuarial Journal* **10**, 29–67. doi: 10.1007/s13385-019-00222-0.

Bladt, M. and Furrer, C. (2024). Expert Kaplan–Meier estimation. *Scandinavian Actuarial Journal* **2024**, 1–27. doi: 10.1080/03461238.2023.2197442.

Bolton, P. (2011). Education Maintenance Allowance (EMA) Statistics. Report. Available online at `https://commonslibrary.parliament.uk/research-briefings/sn05778/`.

Botzen, W. J. W., Kunreuther, H., and Michel-Kerjan, E. (2019). Protecting against disaster risks: Why insurance and prevention may be complements. *Journal of Risk and Uncertainty* **59**, 151–169. doi: 10.1007/s11166-019-09312-6.

Bratsberg, B., Fevang, E., and Røed, K. (2013). Job loss and disability insurance. *Labour Economics* **24**, 137–150. doi: 10.1016/j.labeco.2013.08.004.

Brennan, M. J. and Schwartz, E. S. (1979). Alternative Investment Strategies for the Issuers of Equity Linked Life Insurance Policies with an Asset Value Guarantee. *The Journal of Business* **52**, 63–93.

Brennan, M. J. and Schwartz, E. S. (1979). *Pricing and Investment Strategies for Guaranteed Equity-Linked Life Insurance*. S.S. Huebner Foundation for Insurance Education, Wharton School, University of Pennsylvania, Philadelphia.

Buchardt, K., Møller, T., and Schmidt, K. B. (2015). Cash flows and policyholder behaviour in the semi-Markov life insurance setup. *Scandinavian Actuarial Journal* **2015**(8), 660–688. doi: 10.1080/03461238.2013.879919.

Buchardt, K., Furrer, C., and Sandqvist, O. L. (2023). Transaction time models in multi-state life insurance. *Scandinavian Actuarial Journal* **2023**(10), 974–999. doi: 10.1080/03461238.2023.2181708.

Buchardt, K., Furrer, C., and Sandqvist, O. L. (2025). Estimation for multistate models subject to reporting delays and incomplete event adjudication. Preprint. Available online at `https://arxiv.org/abs/2311.04318`.

Bücher, A. and Rosenstock, A. (2024). Combined modelling of micro-level outstanding claim counts and individual claim frequencies in non-life insurance. *European Actuarial Journal* **14**(2), 623–655. doi: 10.1007/s13385-024-00383-7.

Calderwood, L. and Sanchez, C. (2016). Next Steps (formerly known as the Longitudinal Study of Young People in England). *Journal of Open Health Data*

4, 1–3. DOI: doi.org/10.5334/ohd.16.

CALONICO, S., CATTANEO, M. D., and FARRELL, M. H. (2019). nprobust: Nonparametric Kernel-Based Estimation and Robust Bias-Corrected Inference. *Journal of Statistical Software* **91**(8), 1–33. DOI: 10.18637/jss.v091.i08.

CALONICO, S., CATTANEO, M. D., and TITIUNIK, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica* **82**(6), 2295–2326. DOI: 10.3982/ECTA11757.

CASPER, T. C. and COOK, T. D. (2012). Estimation of the Mean Frequency Function for Recurrent Events when Ascertainment of Events is Delayed. *The International Journal of Biostatistics* **8**, 1–20. DOI: 10.1515/1557-4679.1303.

CATTANEO, M. D. and TITIUNIK, R. (2022). Regression Discontinuity Designs. *Annual Review of Economics* **14**, 821–851. DOI: 10.1146/annurev-economics-051520-021409.

CHAMBERLAIN, G. (1987). Asymptotic efficiency in estimation with conditional moment restrictions. *Journal of Econometrics* **34**(3), 305–334. DOI: 10.1016/0304-4076(87)90015-7.

CHEN, C., LINTON, O., and VAN KEILEGOM, I. (2003). Estimation of Semiparametric Models when the Criterion Function Is Not Smooth. *Econometrica* **71**(5), 1591–1608. DOI: 10.1111/1468-0262.00461.

CHERNOZHUKOV, V., CHETVERIKOV, D., DEMIRER, M., DUFLO, E., HANSEN, C., NEWEY, W., and ROBINS, J. M. (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal* **21**, 1–68. DOI: 10.1111/ectj.12097.

CHRISTIANSEN, M. C. (2012). Multistate models in health insurance. *AStA Advances in Statistical Analysis* **96**(2), 155–186. DOI: 10.1007/s10182-012-0189-2.

CHRISTIANSEN, M. C. and DJEHICHE, B. (2020). Nonlinear reserving and multiple contract modifications in life insurance. *Insurance: Mathematics and Economics* **93**, 187–195. DOI: 10.1016/j.insmatheco.2020.05.004.

CHRISTIANSEN, M. C. (2021a). On the calculation of prospective and retrospective reserves in non-Markov models. *European Actuarial Journal* **11**(2), 441–462. DOI: 10.1007/s13385-021-00277-y.

CHRISTIANSEN, M. C. (2021b). Time-dynamic evaluations under non-monotone information generated by marked point processes. *Finance and Stochastics* **25**(3), 563–596. DOI: 10.1007/s00780-021-00456-5.

CHRISTIANSEN, M. C. and FURRER, C. (2021). Dynamics of state-wise prospective reserves in the presence of non-monotone information. *Insurance: Mathematics and Economics* **97**, 81–98. DOI: 10.1016/j.insmatheco.2021.01.005.

CENTRE FOR LONGITUDINAL STUDIES (2024). Next Steps (formerly the Longitudinal Study of Young People in England). *12th Release*, 2004–2024. Available online at `https://beta.ukdataservice.ac.uk/datacatalogue/doi/?id=5545#!#7`.

COLLINGWOOD, A., CHESHIRE, H., NICOLAAS, G., D'SOUZA, J., ROSS, A., HALL, J., et al. (2010). A review of the Longitudinal Study of Young People in England (LSYPE): recommendations for a second cohort. Report. Available online at `https://assets.publishing.service.gov.uk/media/5a7a4a7340f0b66eab99b215/DFE-RR048.pdf`.

COOK, T. D. (2000). Adjusting survival analysis for the presence of unadjudicated study events. *Controlled Clinical Trials* **21**(3), 208–222. DOI: 10.1016/s0197-2456(00)00053-2.

COOK, T. D. and KOSOROK, M. R. (2004). Analysis of Time-to-Event Data With Incomplete Event Adjudication. *Journal of the American Statistical Association* **99**(468), 1140–1152. DOI: 10.1198/016214504000000566.

CREVECOEUR, J., ANTONIO, K., and VERBELEN, R. (2019). Modeling the number of hidden events subject to observation delay. *European Journal of Operational Research* **277**(3), 930–944. DOI: 10.1016/j.ejor.2019.02.044.

CREVECOEUR, J., ROBBEN, J., and ANTONIO, K. (2022). A hierarchical reserving model for reported non-life insurance claims. *Insurance: Mathematics and Economics* **104**, 158–184. DOI: 10.1016/j.insmatheco.2022.02.005.

CREVECOEUR, J., ANTONIO, K., DESMEDT, S., and MASQUELEIN, A. (2022). Bridging the gap between pricing and reserving with an occurrence and development model for non-life insurance claims. *ASTIN Bulletin* **53**(2), 185–212. DOI: 10.1017/asb.2023.14.

CUI, Y., KOSOROK, M. R., SVERDRUP, E., WAGER, S., and ZHU, R. (2023). Estimating heterogeneous treatment effects with right-censored data via causal survival forests. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **85**(2), 179–211. DOI: 10.1093/jrsssb/qkac001.

DABROWSKA, D. (1995). Estimation of transition probabilities and bootstrap in a semiparametric markov renewal model. *Journal of Nonparametric Statistics* **5**(3), 237–259. DOI: 10.1080/10485259508832646.

DABROWSKA, D. M. (1997). Smoothed Cox regression. *The Annals of Statistics* **25**(4), 1510–1540. DOI: 10.1214/aos/1031594730.

DEARDEN, L., EMMERSON, C., FRAYNE, C., and MEGHIR, C. (2009). Conditional Cash Transfers and School Dropout Rates. *The Journal of Human Resources* **44**(4), 827–857.

DELONG, Ł., LINDHOLM, M., and WÜTHRICH, M. V. (2021). Collective reserving using individual claims data. *Scandinavian Actuarial Journal* **2022**, 1–18. DOI: 10.1080/03461238.2021.1921836.

DELSOL, L. and VAN KEILEGOM, I. (2020). Semiparametric M-estimation with non-smooth criterion functions. *Annals of the Institute of Statistical Mathematics* **72**, 577–605. DOI: 10.1007/s10463-018-0700-y.

DJEHICHE, B. and LÖFDAHL, B. (2014). Risk aggregation and stochastic claims reserving in disability insurance. *Insurance: Mathematics and Economics* **59**, 100–108. DOI: 10.1016/j.insmatheco.2014.09.001.

DJEHICHE, B. and LÖFDAHL, B. (2016). Nonlinear reserving in life insurance: Aggregation and mean-field approximation. *Insurance: Mathematics and Economics* **69**, 1–13. DOI: 10.1016/j.insmatheco.2016.04.002.

DUBOIS, M. (2011). Insurance and prevention: Ethical aspects. *The Journal of Primary Prevention* **32**, 3–15. DOI: 10.1007/s10935-011-0234-z.

EFRON, B. (1979). Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics* **7**, 1–26. DOI: 10.1214/aos/1176344552.

EHRLICH, I. and BECKER, G. S. (1972). Market Insurance, Self-Insurance, and Self-Protection. *Journal of Political Economy* **80**(4), 623–648.

EKSPERTGRUPPEN FOR FREMTIDENS BESKÆFTIGELSESINDSATS (2024). Anbefalinger til fremtidens beskæftigelsesindsats. Report. Available online at `https://xn--fremtidensbeskftigelsesindsats-0uc.dk/afrapportering/`.

ERHVERVSMINISTERIET (2015). Bekendtgørelse om finansielle rapporter for forsikringsselskaber og tværgående pensionskasser. Available online at `https://www.retsinformation.dk/eli/lta/2015/937`.

Esbjerg, S., Keiding, N., and Koch-Henriksen, N. (1999). Reporting delay and corrected incidence of multiple sclerosis. *Statistics in Medicine* **18**(13), 1691–1706. DOI: 10.1002/(SICI)1097-0258(19990715)18:13¡1691::AID-SIM160¿3.0.CO;2-D.

European Insurance and Occupational Pensions Authority (2009). Solvency II Directive (Directive 2009/138/EC). Available online at `https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:02009L0138-20140523&from=EN`.

Fischer, J. B. and Kvist, J. (2023). *Frihed, lighed og privat velfærd: På vej mod en ny samfundsmodel*. Copenhagen, Denmark: Gyldendal.

Friedman, M. (1982). Piecewise Exponential Models for Survival Data with Covariates. *The Annals of Statistics* **10**, 101–113. DOI: 10.1214/aos/1176345693.

Furrer, C. (2020). *Multi-state modeling in the mathematics of life insurance: meditations and applications*. PhD thesis, University of Copenhagen.

Furrer, C. and Sandqvist, O. L. (2025). Loss of earning capacity in Denmark – an actuarial perspective. Preprint. Available online at `https://arxiv.org/abs/2501.11578`.

Gauchon, R., Loisel, S., and Rullière, J. (2020). Health policyholder clustering using medical consumption. *European Actuarial Journal* **10**(2), 599—626. DOI: 10.1007/s13385-020-00244-z.

Gauchon, R., Loisel, S., Rullière, J., and Trufin, J. (2020). Optimal prevention strategies in the classical risk model. *Insurance: Mathematics and Economics* **91**, 202–208. DOI: 10.1016/j.insmatheco.2020.02.003.

Gauchon, R., Loisel, S., Rullière, J., and Trufin, J. (2021). Optimal prevention of large risks with two types of claims. *Scandinavian Actuarial Journal* **2021**(4), 323–334. DOI: 10.1080/03461238.2020.1844791.

Giné, E. and Zinn, J. (1990). Bootstrapping General Empirical Measures. *The Annals of Probability* **18**(2), 851–869. DOI: 10.1214/aop/1176990862.

Gill, R. D., Van Der Laan, M. J., and Robins, J. M. (1997). Coarsening at Random: Characterizations, Conjectures, Counter-Examples. In: Lin, D.Y. and Fleming, T.R. Eds. *Proceedings of the First Seattle Symposium in Biostatistics: Survival Analysis*, 255–294. New York, United States: Springer. DOI: 10.1007/978-1-4684-6316-3_14.

GROSS, S. T. and LAI, T. L. (1996). BOOTSTRAP METHODS FOR TRUNCATED AND CENSORED DATA. *Statistica Sinica* **6**(3), 509–530.

GUNNES, N., BORGAN, Ø., and AALEN, O. O. (2007). Estimating stage occupation probabilities in non-Markov models. *Lifetime Data Analysis* **13**(2), 211–240. DOI: 10.1007/s10985-007-9034-4.

HAASTRUP, S. and ARJAS, E. (1996). Claims Reserving in Continuous Time; A Nonparametric Bayesian Approach. *ASTIN Bulletin* **26**(2), 139–164. DOI: 10.2143/AST.26.2.563216.

HABERMAN, S. and PITACCO, E. (1998). *Actuarial models for disability insurance*. Boca Raton, United States: Chapman & Hall.

HAHN, J. (1996). A Note on Bootstrapping Generalized Method of Moments Estimators. *Econometric Theory* **12**, 187–197. DOI: 10.1017/S0266466600006496.

HALBERT, W. (1982). Maximum Likelihood Estimation of Misspecified Models. *Econometrica* **50**, 1–25. DOI: 10.2307/1912526.

HAHN, J., TODD, P., and VAN DER KLAAUW, W. (1999). Evaluating the Effect of an Antidiscrimination Law Using a Regression-Discontinuity Design. *NBER Working Papers* **7131**, 1–22. DOI: 10.3386/w7131.

HAHN, J., TODD, P., and VAN DER KLAAUW, W. (2001). Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design. *Econometrica* **69**, 201–209.

HANSEN, L. P. (1982). Large Sample Properties of Generalized Method of Moments Estimators. *Econometrica* **50**(4), 1029–1054. DOI: 10.2307/1912775.

HELWICH, M. (2008). *Durational effects and non-smooth semi-Markov models in life insurance*. PhD thesis, University of Rostock.

HERNÁN, M. A. and ROBINS, J. M. (2020). *Causal Inference: What If*. Boca Raton, United States: Chapman & Hall/CRC.

HINES, O., DUKES, O., DIAZ-ORDAZ, K., and VANSTEELANDT, S. (2022). Demystifying Statistical Learning Based on Efficient Influence Functions. *The American Statistician* **76**(3), 292–304. DOI: 10.1080/00031305.2021.2021984.

HOEM, J. M. (1969). Markov Chain Models in Life Insurance. *Blätter der DGVFM* **9**, 91–107. DOI: 10.1007/BF02810082.

HOEM, J. M. (1972). Inhomogeneous Semi-Markov Processes, Select Actuarial Tables, and Duration-Dependence in Demography. In: Greville, T.N.E. Ed. *Population dynamics*, 251–296. New York, United States: Academic Press. DOI: 10.1016/B978-1-4832-2868-6.50013-8.

HOEM, J. M. and AALEN, O. O. (1978). Actuarial values of payment streams. *Scandinavian Actuarial Journal* **1978**, 38–47. DOI: 10.1080/03461238.1978.10414317.

HOLFORD, A. (2015). The labour supply effect of Education Maintenance Allowance and its implications for parental altruism. *Review of Economics of the Household* **13**(3), 531–568. DOI: 10.1007/s11150-015-9288-7.

HOUGAARD, P. (2000). *Analysis of Multivariate Survival Data*. New York, United States: Springer. DOI: 10.1007/978-1-4612-1304-8.

HU, P. and TSIATIS, A. A. (1996). Estimating the survival distribution when ascertainment of vital status is subject to delay. *Biometrika* **83**(2), 371–380. DOI: 10.1093/biomet/83.2.371.

HWANG, J. and SUN, Y. (2018). Should we go one step further? An accurate comparison of one-step and two-step procedures in a generalized method of moments framework. *Journal of Econometrics* **207**(2), 381–405. DOI: 10.1016/j.jeconom.2018.07.006.

ICHIMURA, H. and LEE, S. (2010). Characterization of the asymptotic distribution of semiparametric M-estimators. *Journal of Econometrics* **159**(2), 252–266. DOI: 10.1016/j.jeconom.2010.05.005.

IMBENS, G. W. and LEMIEUX, T. (2008). Regression discontinuity designs: A guide to practice. *Journal of Econometrics* **142**(2), 615–635. DOI: 10.1016/j.jeconom.2007.05.001.

JACOBSEN, M. (2006). *Point Process Theory and Applications: Marked Point and Piecewise Deterministic Processes*. Boston, United States: Birkhauser. DOI: 10.1007/0-8176-4463-6.

JACOBSEN, M. and MARTINUSSEN, T. (2016). A Note on the Large Sample Properties of Estimators Based on Generalized Linear Models for Correlated Pseudo-observations. *Scandinavian Journal of Statistics* **43**(3), 845–862. DOI: 10.1111/sjos.12212.

JACOD, J. (1975). Multivariate point processes: predictable projection, Radon-Nikodym derivatives, representation of martingales. *Zeitschrift für Wahrschein-*

lichkeitstheorie und verwandte Gebiete **31**(3), 235–253. DOI: 10.1007/BF00536010.

JANSSEN, J. (1966). Application des processus semi-markoviens à un probléme d'invalidité. *Bulletin de l'Association Royale des Actuaries Belges* **63**, 35–52.

JENSEN, S. E. H., KALLESTRUP-LAMB, M., KIERKEGAARD, L.M., STEFFENSEN, M., and TANGGAARD, C. (2019). Forbrugernes udfordringer med pensionsbeslutninger. Report. Available online at `https://www.raadtilpenge.dk/-/media/PPP/Rapporter/Rapport-om-forbrugerens-udfordringer-med-pensionsbeslutninger.pdf`

KALBFLEISCH, J. D. and LAWLESS, J. F. (1989). Inference Based on Retrospective Ascertainment: An Analysis of the Data on Transfusion-Related AIDS. *Journal of the American Statistical Association* **84**(406), 360–372. DOI: 10.1080/01621459.1989.10478780.

KALBFLEISCH, J. D. and LAWLESS, J. F. (1991). REGRESSION MODELS FOR RIGHT TRUNCATED DATA WITH APPLICATIONS TO AIDS INCUBATION TIMES AND REPORTING LAGS. *Statistica Sinica* **1**, 19–32.

KAMINSKY, K. S. (1987). Prediction of IBNR claim counts by modelling the distribution of report lags. *Insurance: Mathematics and Economics* **6**(2), 151–159. DOI: 10.1016/0167-6687(87)90024-2.

KENKEL, D. S. (2000). Prevention. In: Culyer, A.J., and Newhouse, J.P. Eds. *Handbook of Health Economics* (Vol. 1B), 1675–1722. Oxford, United Kingdom: North Holland. DOI: 10.1016/S1574-0064(00)80044-X.

KENNEDY, E. H. (2022). Semiparametric doubly robust targeted double machine learning: a review. Preprint. Available online at `https://arxiv.org/pdf/2203.06469`.

KENNEDY, E. H. (2023). Towards optimal doubly robust estimation of heterogeneous causal effects. *Electronic Journal of Statistics* **17**(2), 3008–3049. DOI: 10.1214/23-EJS2157.

KIM, J. K. (2011). Parametric fractional imputation for missing data analysis. *Biometrika* **98**, 119–132. DOI: 10.1093/biomet/asq073.

KOMMISSIONEN OM TILBAGETRÆKNING OG NEDSLIDNING (2022). Fremtidssikring af et stærkt pensionssytem. Report. Available online at `https://bm.dk/media/20703/fremtidssikring-af-et-staerkt-pensionssystem.pdf`.

König, B., Weber, F., and Wüthrich, M. V. (2011). Prediction of disability frequencies in life insurance. *Zavarovalniski horizonti* **7**(3), 51–69.

Kristensen, D. and Salanié, B. (2017). Higher-order properties of approximate estimators. *Journal of Econometrics* **198**(2), 189–208. doi: 10.1016/j.jeconom.2016.10.008.

Lagakos, S. W., Barraj, L. M., and De Gruttola, V. (1988). Nonparametric Analysis of Truncated Survival Data, with Application to AIDS. *Biometrika* **75**(3), 515–523. doi: 10.2307/2336602.

Lagakos, S. W., Sommer, C. J., and Zelen, M. (1978). Semi-Markov Models for Partially Censored Data. *Biometrika* **65**(2), 311–317. doi: 10.2307/2335209.

Lindsey, J. K. (1995). Fitting Parametric Counting Processes by Using Log-Linear Models. *Journal of the Royal Statistical Society: Series C* **44**(2), 201–212. doi: 10.2307/2986345.

Little, R. J. (2021). Missing Data Assumptions. *Annual Review of Statistics and Its Application* **8**, 89–107. doi: 10.1146/annurev-statistics-040720-031104.

Lopez, O., Milhaud, X., and Thérond, P. E. (2018). Micro-level VS macro-level reserving in non-life insurance: why and when? Preprint. Available online at `https://hal.science/hal-01868744`.

Lopez, O., Xavier, M., and Thérond, P. (2019). A TREE-BASED ALGORITHM ADAPTED TO MICROLEVEL RESERVING AND LONG DEVELOPMENT CLAIMS. *ASTIN Bulletin* **49**(3), 741–762. doi: 10.1017/asb.2019.12.

Mack, T. (1993). Distribution-free Calculation of the Standard Error of Chain Ladder Reserve Estimates. *ASTIN Bulletin* **23**(2), 213–225. doi: 10.2143/AST.23.2.2005092.

Mack, T. (1999). The Standard Error of Chain Ladder Reserve Estimates: Recursive Calculation and Inclusion of a Tail Factor. *ASTIN Bulletin* **29**(2), 361–366. doi: 10.2143/AST.29.2.504622.

Maguire, M. J., Maguire, S., and Vincent, J. (2001). Implementation of the Education Maintenance Allowance Pilots: The First Year. Report. Available online at `https://dera.ioe.ac.uk/id/eprint/4680/1/RR255.pdf`.

McKendrick, A. (2022). Paying Students to Stay in School: Short- and Long-term Effects of a Conditional Cash Transfer in England. Preprint. Available online at https://www.lancaster.ac.uk/media/lancaster-university/content-assets/documents/lums/economics/working-papers/LancasterWP2022_002.pdf.

Middleton, S., Perren, K., Maguire, S., Rennison, J., Battistin, E., Emmerson, C., and Fitzsimons, E. (2005). Evaluation of Education Maintenance Allowance Pilots: Young People Aged 16 to 19 Years - Final Report of the Quantitative Evaluation. Report. Available online at https://dera.ioe.ac.uk/id/eprint/5734/1/RR678.pdf.

Milbrodt, M. & Röhrs, V. (2016). *Aktuarielle Methoden der deutschen Privaten Krankenversicherung* (2nd ed.). Karlsruhe, Germany: VVW GmbH.

Munch, A., Breum, M. S., Martinussen, T., and Gerds, T. A. (2023). Targeted estimation of state occupation probabilities for the non-Markov illness-death model. *Scandinavian Journal of Statistics* **50**(3), 1532–1551. DOI: 10.1111/sjos.12644.

Munch, A., Gerds, T. A., van der Laan, M. J., and Rytgaard, H. C. (2024). Estimating conditional hazard functions and densities with the highly-adaptive lasso. Preprint. Available online at https://arxiv.org/abs/2404.11083.

Murphy, K. M. and Topel, R. H. (1985). Estimation and Inference in Two-Step Econometric Models. *Journal of Business & Economic Statistics* **3**(4), 370–379. DOI: 10.2307/1391724.

Møller, C. M. (1993). A stochastic version of Thiele's differential equation. *Scandinavian Actuarial Journal* **1993**, 1–16. DOI: 10.1080/03461238.1993.10413910.

Møller, T. and Steffensen, M. (2007). *Market-Valuation Methods in Life and Pension Insurance*. Cambridge, United Kingdom: Cambridge University Press. DOI: 10.1017/CBO9780511543289.

National Association of Insurance Commissioners (2023). U.S. Health Insurance Industry Analysis Report. Available online at https://content.naic.org/sites/default/files/topics-industry-snapshot-analysis-reports-2023-annual-report-health.pdf.

Newey, W. K. and McFadden, D. L. (1994). Large sample estimation and hypothesis testing. In: Engle, R.F. and McFadden, D.L. Eds. *Handbook of Econometrics*, 2111–2245. Amsterdam, Netherlands: Elsevier. DOI: 10.1016/S1573-4412(05)80005-4.

Norberg, R. (1990). Payment Measures, Interest, and Discounting: An Axiomatic Approach with Applications to Insurance. *Scandinavian Actuarial Journal* **1990**, 14–33. doi: 10.1080/03461238.1990.10413870.

Norberg, R. (1991). Reserves in Life and Pension Insurance. *Scandinavian Actuarial Journal* **1991**, 3–24. doi: 10.1080/03461238.1991.10557357.

Norberg, R. (1992). Hattendorff's theorem and Thiele's differential equation generalized. *Scandinavian Actuarial Journal* **1992**, 2–14. doi: 10.1080/03461238.1992.10413894.

Norberg, R. (1993). Prediction of Outstanding Liabilities in Non-Life Insurance. *ASTIN Bulletin* **23**, 95–115. doi: 10.2143/AST.23.1.2005103.

Norberg, R. (1999). Prediction of Outstanding Liabilities II. Model Variations and Extensions. *ASTIN Bulletin* **29**, 5–25. doi: 10.2143/AST.29.1.504603.

Norberg, R. (2005). Anomalous PDEs in Markov chains: Domains of validity and numerical solutions. *Finance and Stochastics* **9**(4), 519–537. doi: 10.1007/s00780-005-0157-8.

Norris, J. R. (1998). *Markov Chains*. Cambridge, United Kingdom: Cambridge University Press. doi: 10.1017/CBO9780511810633.

Noufaily, A., Farrington, P., Garthwaite, P., Enki, D. G., Andrews, N., and Charlett, A. (2016). Detection of Infectious Disease Outbreaks From Laboratory Data With Reporting Delays. *Journal of the American Statistical Association* **111**(514), 488–499. doi: 10.1080/01621459.2015.1119047.

Ogata, Y. (1981). On Lewis' simulation method for point processes. *IEEE* **27**, 23–31. doi: 10.1109/TIT.1981.1056305.

Okine, A. N. A., Frees, E. W., and Shi, P. (2022). JOINT MODEL PREDICTION AND APPLICATION TO INDIVIDUAL-LEVEL LOSS RESERVING. *ASTIN Bulletin* **2022**, 91–116. doi: 10.1017/asb.2021.28.

Overgaard, M., Parner, E. T., and Pedersen, J. (2017). Asymptotic theory of generalized estimating equations based on jack-knife pseudo-observations. *The Annals of Statistics* **45**(5), 1988–2015. doi: 10.1214/16-AOS1516.

Pagano, M., Tu, X.M., De Gruttola, V., and MaWhinney, S. (1994). Regression Analysis of Censored and Truncated Data: Estimating Reporting-Delay Distributions and AIDS Incidence from Surveillance Data. *Biometrics*

50(4), 1203–1214. DOI: 10.2307/2533459.

PARNER, E. T., ANDERSEN, P. K., and OVERGAARD, M. (2023). Regression models for censored time-to-event data using infinitesimal jack-knife pseudo-observations, with applications to left-truncation. *Lifetime Data Analysis* **29**(3), 654–671. DOI: 10.1007/s10985-023-09597-5.

PEARL, J. (2009). *Causality: Models, Reasoning, and Inference* (2nd ed.). Cambridge, United Kingdom: Cambridge University Press.

PEDERSEN, A. W. and KUHNLE, S. (2017). The Nordic welfare state models. In: Knutsen, O. Ed. *The Nordic models in political science: Challenged, but still viable?*, 219–238. Bergen, Norway: Fagbokforlaget.

POLLARD, D. (1984). *Convergence of Stochastic Processes*. New York, United States: Springer. DOI: 10.1007/978-1-4612-5254-2.

PSOTKA, M. A., ABRAHAM, W. T., FIUZAT, M., FILIPPATOS, G., LINDENFELD, J., AHMAD, T., et al. (2020). Conduct of Clinical Trials in the Era of COVID-19: JACC Scientific Expert Panel. *Journal of the American College of Cardiology* **76**(20), 2368–2378. DOI: 10.1016/j.jacc.2020.09.544.

R DEVELOPMENT CORE TEAM (2023). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

RAHMAN, M. M. (2014). *Estimation of Treatment Effects Using Regression Discontinuity Design*. PhD thesis, University of Manchester.

RAMBACHAN, A., COSTON, A., and KENNEDY, E. (2022). Counterfactual risk assessments under unmeasured confounding. Preprint. Available online at `https://arxiv.org/abs/2212.09844`.

ROBINS, J. M. and WANG, N. (2000). Inference for Imputation Estimators. *Biometrika* **87**, 113–124.

RUBIN, D.B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* **66**(5), 688–701. DOI: 10.1037/h0037350.

RUBIN, D. B. (1987). *Multiple Imputation for Nonresponse in Surveys*. New York, United States: Wiley. DOI: 10.1002/9780470316696.

RUBIN, D. and VAN DER LAAN, M. J. (2007). A doubly robust censoring unbiased transformation. *The International Journal of Biostatistics* **3**, 1–19. DOI: 10.2202/1557-4679.1052.

RYTGAARD, H. C., ERIKSSON, F., and VAN DER LAAN, M. J. (2023). Estimation of time-specific intervention effects on continuously distributed time-to-event outcomes by targeted maximum likelihood estimation. *Biometrics* **79**(4), 3038–3049. DOI: 10.1111/biom.13856.

RYTGAARD, H. C., GERDS, T., and VAN DER LAAN, M. J. (2022). Continuous-time targeted minimum loss-based estimation of intervention-specific mean outcomes. *The Annals of Statistics* **50**(5), 2469–2491. DOI: 10.1214/21-AOS2114.

SANDQVIST, O. L. (2025). A multistate approach to disability insurance reserving with information delays. Preprint. Available online at `https://arxiv.org/abs/2312.14324`.

SANDQVIST, O. L. (2024). Doubly robust inference with censoring unbiased transformations. Preprint. Available online at `https://arxiv.org/pdf/2411.04909`.

SCHERVISH, M. J. (1995). *Theory of Statistics.* New York, United States: Springer. DOI: 10.1007/978-1-4612-4250-5.

SEGERER, G. (1993). The actuarial treatment of the disability risk in Germany, Austria and Switzerland. *Insurance: Mathematics and Economics* **13**(2), 131–140. DOI: 10.1016/0167-6687(93)90835-D.

SHORACK, G. R. and WELLNER, J. A. (1986). *Empirical Processes with Applications to Statistics.* New York, United States: Wiley.

SNODGRASS, R. and AHN, I. (1985). A taxonomy of time databases. *ACM Sigmod Record* **14**(4), 236–246. DOI: 10.1145/971699.318921.

SPENDER, A., BULLEN, C., ALTMANN-RICHER, L., CRIPPS, J., DUFFY, R., FALKOUS, C., et al. (2019). Wearables and the Internet of Things: Considerations for the Life and Health Insurance Industry. *British Actuarial Journal* **24**, 1–31. DOI: 10.1017/S1357321719000072.

SPITONI, C., VERDUIJN, M., and PUTTER, H. (2012). Estimation and asymptotic theory for transition probabilities in Markov renewal multi-state models. *The International Journal of Biostatistics* **8**, 23. DOI: 10.1515/1557-4679.1375.

STEINGRIMSSON, J. A., DIAO, J. A., and STRAWDERMAN, R. L. (2016). Doubly robust survival trees. *Statistics in Medicine* **35**(20), 3595–3612. DOI: 10.1002/sim.6949.

STEINGRIMSSON, J. A., DIAO, J. A., and STRAWDERMAN, R. L. (2019). Censoring Unbiased Regression Trees and Ensembles. *Journal of the American Statistical Association* **114**(525), 370–383. DOI: 10.1080/01621459.2017.1407775.

STONE, C. J. (1977). Consistent Nonparametric Regression. *The Annals of Statistics* **5**(4), 595–620. DOI: 10.1214/aos/1176343886.

STONE, C. J. (1980). Optimal Rates of Convergence for Nonparametric Estimators. *The Annals of Statistics* **8**(6), 1348–1360. DOI: 10.1214/aos/1176345206.

STONE, C. J. (1982). Optimal Global Rates of Convergence for Nonparametric Regression. *The Annals of Statistics* **10**(4), 1040–1053. DOI: 10.1214/aos/1176345969.

STONER, O., HALLIDAY, A., and ECONOMOU, T. (2023). Correcting delayed reporting of COVID-19 using the generalized-Dirichlet-multinomial method. *Biometrics* **79**(3), 2537–2550. DOI: 10.1111/biom.13810.

SUZUKAWA, A. (2004). Unbiased Estimation of Functionals Under Random Censorship. *Journal of the Japan Statistical Society* **34**(2), 153–172. DOI: 10.14490/jjss.34.153.

SWISS RE (2022). The Legacy of COVID-19: The ongoing mental health impact on Australia's community. White paper. Available online at `https://www.swissre.com/dam/jcr:affd6238-0c29-4ebd-9808-ed32159c09b0/white-paper-the-legacy-of-covid-19-mental-health-impact-on-australias-community.pdf`.

TAYLOR, G. C. (1971). SICKNESS: A STOCHASTIC PROCESS. *Journal of the Institute of Actuaries* **97**, 69–83.

TSIATIS, A. A. (2006). *Semiparametric Theory and Missing Data*. New York, United States: Springer. DOI: 10.1007/0-387-37345-4.

VAN DER LAAN, M. J. and HUBBARD, A. E. (1998). Locally Efficient Estimation of the Survival Distribution with Right-Censored Data and Covariates when Collection of Data is Delayed. *Biometrika* **85**(4), 771–783.

VAN DER LAAN, M. J. and ROBINS, J. M. (2003). *Unified Methods for Censored Longitudinal Data and Causality*. New York, United States: Springer. DOI:

10.1007/978-0-387-21700-0.

Van der Laan, M. J. and Rose, S. (2011). *Targeted Learning: Causal Inference for Observational and Experimental Data*. New York, United States: Springer. doi: 10.1007/978-1-4419-9782-1.

Van der Vaart, A. W. (1998). *Asymptotic Statistics*. Cambridge, United Kingdom: Cambridge University Press. doi: 10.1017/CBO9780511802256.

van der Vaart, A. (2004). On Robins' formula. *Statistics & Decisions* **22**(3), 171–200. doi: 10.1524/stnd.22.3.171.57065.

Verbelen, R., Antonio, K., Claeskens, G., and Crevecoeur, J. (2022). Modeling the Occurrence of Events Subject to a Reporting Delay via an EM Algorithm. *Statistical Science* **37**(3), 394–410. doi: 10.1214/21-STS831.

Verdinelli, I. and Wasserman, L. (2021). Forest Guided Smoothing. Preprint. Available online at `https://arxiv.org/pdf/2103.05092`.

Wang, N. and Robins, J. M. (1998). Large-Sample Theory for Parametric Multiple Imputation Procedures. *Biometrika* **85**(4), 935–948.

Wüthrich, M. V. (2018). Neural networks applied to chain–ladder reserving. *European Actuarial Journal* **8**(2), 407–436. doi: 10.1007/s13385-018-0184-4.

Yackel, J. (1968). A Random Time Change Relating Semi-Markov and Markov Processes. *The Annals of Mathematical Statistics* **39**(2), 358–364. doi: 10.1214/aoms/1177698396.

Yang, L., Shi, P., and Huang, S. (2024). A copula model for marked point process with a terminal event: An application in dynamic prediction of insurance claims. *The Annals of Applied Statistics* **18**(4), 2679–2704. doi: 10.1214/24-AOAS1902.

Yang, Y. and Tokdar, S. T. (2015). MINIMAX-OPTIMAL NONPARAMETRIC REGRESSION IN HIGH DIMENSIONS. *The Annals of Statistics* **43**(2), 652–674.