Stockholm University

Large exposure asymptotics in insurance valuation and reserving, tree regularisation and stochastic control

Nils Engler

Doctoral Thesis in Mathematical Statistics at Stockholm University, Sweden 2026

Large exposure asymptotics in insurance valuation and reserving, tree regularisation and stochastic control

Nils Engler

Academic dissertation for the Degree of Doctor of Philosophy in Mathematical Statistics at Stockholm University to be publicly defended on Friday 29 May 2026 at 13.00 in Lärosal 4, Albano Hus 1, Vån 2, Albanovägen 28.

Abstract

This thesis investigates several topics in actuarial mathematics and applied probability, including insurance valuation and reserving, regularisation of regression trees, and stochastic optimisation in an extended dividend problem. The thesis is based on four papers.

Paper I provides a justification of the chain ladder predictor and Mack's estimator for the prediction error within a classical compound Poisson model under large exposure, that is, when the number of contracts tends to infinity. Although the model does not satisfy the assumptions of Mack's distribution-free chain ladder, both the predictor and the estimator are shown to arise in the large exposure limit.

Paper II studies the valuation of liability cashflows with capital requirements in a multi-period setting. Since explicit valuation is generally infeasible and Monte Carlo methods are often computationally challenging, an explicit and easily computable valuation formula is derived. The formula is obtained as a large exposure limit under a conditional weak convergence assumption on the liability cashflows.

Paper III introduces a regularisation method for regression trees based on node-wise statistical tests. At each node, a p-value is computed using a change point test, resulting in a regularised regression tree that is a deterministic function of the training data. Unlike cross-validation, the method avoids randomness from data splitting and ensures efficient use of the full dataset.

Paper IV revisits the classical dividend problem with ruin at zero by incorporating an additional default mechanism based on cumulative occupation time in a low-surplus region. This extension reflects realistic default triggers such as regulatory pressure or liquidity stress. The problem is solved explicitly, yielding closed-form expressions for both the optimal control and the value function.

Keywords: *claims reserving, valuation, regression trees, optimal dividends.*

Stockholm 2026
http://urn.kb.se/resolve?urn=urn:nbn:se:su:diva-253128

ISBN 978-91-8107-534-2
ISBN 978-91-8107-535-9

Stockholm University

Department of Mathematics

Stockholm University, 106 91 Stockholm

LARGE EXPOSURE ASYMPTOTICS IN INSURANCE VALUATION AND RESERVING, TREE REGULARISATION AND STOCHASTIC CONTROL

Nils Engler

Stockholm University

Large exposure asymptotics in insurance valuation and reserving, tree regularisation and stochastic control

Nils Engler

ISBN print 978-91-8107-534-2
ISBN PDF 978-91-8107-535-9

Printed in Sweden by Universitetsservice US-AB, Stockholm 2026

Für Mercedes.

List of Papers

This thesis is based on four papers, numbered I–IV, which are included in the following order.

I. Engler, Nils and Lindskog, Filip (2024). Mack's estimator motivated by large exposure asymptotics in a compound Poisson setting. *ASTIN Bulletin* 54, pp. 310–326

II. Engler, Nils and Lindskog, Filip (2025). Approximations of multi-period liability values by simple formulas. *Insurance: Mathematics and Economics* 123, 103112

III. Engler, Nils, Lindholm, Mathias, Lindskog, Filip and Nazar, Taariq (2025). Regularisation of regression trees by summation of p-values. arXiv: `2505.18769v2`

IV. Bodnariu, Andi, Engler, Nils and Rodosthenous, Neofytos (2026). Outrunning the Omega Clock: A Singular Control Problem for Dividend Optimisation with Ruin and Time-in-Distress Default. arXiv: `2601.21705`

Authors' contributions:

Paper I The choice of the model and the idea to examine its compatibility with the chain ladder technique originated from F. Lindskog. N. Engler contributed to all theoretical results and to the discussions leading to their formulation. In particular, N. Engler derived the limiting distributions of the estimators appearing in the analysis and implemented the numerical illustration.

Paper II The cost-of-capital valuation approach revisited in this paper was originally developed in earlier work by F. Lindskog, H. Engsner, and M. Lindholm, as was the valuation formula for the Gaussian cashflow case. The project idea originated from F. Lindskog. N. Engler derived the continuous convergence condition required for the convergence of the liability values and, together with F. Lindskog, formulated and proved all theoretical results presented in the paper.

Paper III N. Engler contributed to the idea of employing statistical tests at the nodes of a regression tree for regularisation purposes and to the derivation of the corresponding test statistic. Furthermore, N. Engler investigated its distribution under both the null and alternative hypotheses and contributed to the implementation of the associated programming code.

Paper IV The project idea originated from N. Rodosthenous. In joint work with A. Bodnariu, N. Engler formulated the theoretical results and developed their proofs. N. Engler also implemented the code used for the numerical illustrations.

Usage of AI:

The author did not use artificial intelligence in any of the above papers. Artificial intelligence (GPT-5.2) was employed to improve the writing style of some portions of Part I (Introduction) of this work. This use did not affect the intellectual content.

Acknowledgements

I would like to express my deep gratitude to my supervisor, Filip Lindskog, and my co-supervisor, Mathias Lindholm, for their constant availability, their support, and the inspiring exchanges we have shared over the past years. I am also very grateful to my collaborators and companions, Andi Bodnariu and Taariq Nazar, with whom I have spent countless hours exploring mathematical problems.

I would further like to thank my colleagues in the Department of Mathematics, in particular Jan McGarry Furriol, Benedetta Andina, Gabriel Saadia and Dario Giandinoto, for their friendship and for making my time at Stockholm University so enjoyable.

My curiosity about mathematics was significantly shaped by my two close friends and study companions, Fabian Bittl and Maximilian Kaske, whom I would like to thank for the memorable time we shared at Technische Universität Berlin.

Moreover, I want to deeply thank my parents for fostering my interest in mathematics at an early stage and for supporting me throughout my entire journey. Finally, and most importantly, I owe my greatest thanks to my family, Mercedes and Oskar, for reminding me every day of what truly matters in life.

Contents

Part I

Introduction

Preliminaries and notation

- $(\Omega, \mathcal{F}, \mathrm{P})$ is a probability space sufficiently rich to support all random variables that appear in this introduction.

- $L^0 = L^0(\Omega, \mathcal{F}, \mathrm{P})$ is the space of $\mathcal{F}$-measurable random variables from Ω to $\mathbb{R}$, modulo the equivalence relation of P-almost sure equality. For $1 \le p < \infty$, $L^p = L^p(\Omega, \mathcal{F}, \mathrm{P}) \subset L^0$ is the subspace of random variables whose p-th power of the absolute value is Lebesgue integrable with respect to $(\Omega, \mathcal{F}, \mathrm{P})$. Finally $L^\infty = L^\infty(\Omega, \mathcal{F}, \mathrm{P}) \subset L^p$ is the space of all bounded, $\mathcal{F}$-measurable random variables.

- We use the Euclidian norm $\|x\| := \sqrt{\sum_{k=1}^d x_k^2}$ for $x \in \mathbb{R}^d$.

- All inequalities and equalities between random variables are meant in the P-almost sure sense.

- $\mathbb{N} := \{1, 2, 3, \dots\}$, $\mathbb{N}_0 := \{0, 1, 2, 3, \dots\}$, $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$.

- For $d \in \mathbb{N}$, $C(\mathbb{R}^d)$ is the set of all continuous, real-valued functions on $\mathbb{R}^d$. $C_b(\mathbb{R}^d) \subset C(\mathbb{R}^d)$ is the subset of bounded continuous functions. For $m_1, \dots, m_d \in \mathbb{N}$, $C^{m_1, \dots, m_d}(\mathbb{R}^d)$ is the set of real-valued functions which are m_k-times continuously differentiable with respect to the kth component for $k = 1, \dots, d$.

- For $d \in \mathbb{N}$, $\mathcal{P}(\mathbb{R}^d)$ is the set of all probability measures on $\mathbb{R}^d$, equipped with its Borel sigma algebra $\mathcal{B}(\mathbb{R}^d)$.

- For $\mu_n, \mu \in \mathcal{P}(\mathbb{R}^d)$, we write $\mu_n \stackrel{w}{\to} \mu$, if $\int u d\mu_n \to \int u d\mu$ for all $u \in C_b(\mathbb{R}^d)$.

- If not stated otherwise, δ_x is the Dirac measure on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ at the point $x \in \mathbb{R}^d$. For $A \in \mathcal{B}(\mathbb{R}^d)$, $\delta_x(A) = 1$ if and only if $x \in A$.

- For two random variables X, Y, we write $X \stackrel{d}{=} Y$, if they have the same law.

- For $T \in \mathbb{N}$, we denote the set of indices $\mathcal{T} = \{1, \dots, T\}$.

- In this introduction, we do not distinguish between row and column vectors.
- A random variable Y taking values in $\overline{\mathbb{R}}$ is called the essential supremum of a family of random variables $(X_i)_{i\in I}$, if

  (i) $Y \geq X_i$ for all $i \in I$,

  (ii) $\tilde{Y} \geq Y$ for any other random variable $\tilde{Y}$ satisfying (i).

  Y is unique up to null sets and we denote $\operatorname{ess\,sup}_{i\in I} X_i := Y$. The essential infimum of $(X_i)_{i\in I}$ is defined correspondingly and denoted $\operatorname{ess\,inf}_{i\in I} X_i$.
- Let P, Q be probability measures on $(\Omega, \mathcal{F})$. The measure Q is absolutely continuous with respect to P if $\mathrm{P}(A) = 0$ implies that $\mathrm{Q}(A) = 0$ for all $A \in \mathcal{F}$. In that case we denote by $d\,\mathrm{Q}\,/d\,\mathrm{P}$ the Radon-Nikodym derivate of Q with respect to P.
- $X_{\leq t} := (X_1, \dots, X_t), \quad X_{>t} := (X_{t+1}, \dots, X_T)$, for a random variabel $X = (X_1, \dots, X_T)$ taking values in $\mathbb{R}^T$, analogous for non-random $x \in \mathbb{R}^T$.
- For a family of sets $A_i \in \mathcal{F}, i \in I$, where I is an index set, we denote by $\sigma(A_i \colon i \in I)$ the smallest (with respect to set inclusion) sigma algebra containing all of the sets $(A_i)_{i\in I}$.
- A function $f \colon \mathbb{R} \to \mathbb{R}$ is called *(strictly) increasing*, if $x_1 \leq x_2, (x_1 < x_2)$ implies $f(x_1) \leq f(x_2), (f(x_1) < f(x_2))$. A *(strictly) decreasing* function is defined correspondingly.
- $(f)^+ := \max(f, 0)$ for a real-valued function f.
- $\|f\|_\infty := \sup_{\{x\in A\}} |f(x)| \in \overline{\mathbb{R}}$ for a function $f \colon A \to \mathbb{R}$, $A \subset \mathbb{R}^d$.

Chapter 1

Motivation

This thesis covers a broad range of topics relevant to actuarial mathematics.

Reserving in insurance means ensuring sufficient capital is available to pay policyholders in the future. Therefore, and to reduce the risk of insurer insolvency, accurate reserving is fundamentally important. Among the reserving techniques used in practice, the chain ladder predictor for the ultimate claims amount is one of the most widely applied. Beyond producing a point prediction of outstanding claim payments, it is equally important for an insurance company to assess the uncertainty associated with this prediction. An estimator of the prediction error of outstanding claim reserves was introduced by Thomas Mack in 1993. We provide a theoretical justification of Mack's estimator by deriving it within the framework of one of the most common models for aggregate claims, the compound Poisson model. The derivation is carried out under the assumption of a large number of contracts in the insurance portfolio.

The insurance regulatory framework of the European Union (Solvency II) specifies rules governing the amount of capital that an insurance company must hold in order to reduce the risk of insolvency. Maintaining such capital entails costs, which must be taken into account for a correct valuation of insurance liability cashflows. Cost-of-capital based valuations are typically computed via backward recursion, starting from the latest payment date. For most stochastic cashflow models, the exact calculation of cost-of-capital valuations via backward recursion is neither analytically nor computationally tractable. The simple approximation formula proposed in this thesis therefore has substantial potential for practical application, both for regulators and for insurance companies.

A common class of models used, for example, in insurance pricing are regression trees. Such models are piecewise constant predictors which, in the pricing context, assign the same premium to all individuals within a given tariff cell, where the cell is determined by policyholder characteristics (e.g. age, driving experience in years). The most widely used regularisation technique for regression trees is cross-validation. This approach, however, has the disadvantage of introducing randomness into the model selection process. Moreover, it may effectively reduce the available training sample size and thereby negatively affect model fitting. This issue is particularly relevant in settings with limited data availability. To address these concerns, we propose an in-sample regularisation method that produces a regularised regression tree based on node-wise statistical tests for signal detection. We provide theoretical insights into the signal detection performance as the sample size tends to infinity and conclude with an example illustrating how the method can be used to reduce the complexity of a high-dimensional-parameter black-box predictor.

The dividend problem is a classical problem in stochastic optimal control. It provides a mathematical formalisation of the trade-off faced by a decision maker when determining dividend payments to shareholders. If dividends are paid excessively, the company risks default, resulting in the termination of all future dividend payments. Conversely, due to temporal discounting, there is an incentive to distribute dividends sooner rather than later. An optimal dividend strategy must therefore balance these two opposing effects. In the classical mathematical formulation, company default is modelled as the first time the surplus process reaches zero. In practice, however, default is rarely triggered by a zero surplus alone, but may instead result from factors such as loss of shareholder confidence or regulatory intervention following long periods of low surplus. We account for this by introducing a second default mechanism into the classical dividend problem, which is based on occupation time in low-surplus regions. We provide an explicit solution to the problem, thereby characterising the decision maker's optimal dividend strategy.

Taken together, these four projects address relevant problems in actuarial mathematics, spanning reserving, valuation, statistical learning, and stochastic control.

Chapter 2

Valuation and reserving for discrete-time cashflows

2.1 Claim count and claim amount distributions

This section presents some of the most common distributions for claim counts and claim amounts. Such models are considered both in Paper I and in Paper II. The outline of this section is based on Mikosch (2009).

We first introduce the renewal process as one of the most common classes of claim count distributions. Let (W_i) be an iid sequence of positive random variables. The sequence (T_k) defined by $T_0 = 0, T_k = \sum_{i=1}^{k} W_i, k \in \mathbb{N}$ is called a *renewal sequence* and the process $(M^\alpha)_{\alpha \geq 0}$ defined by $M^\alpha = \max\{k \in \mathbb{N}_0 \colon T_k \leq \alpha\}$ is called a *renewal process*. The sequence (T_k) models the arrival times of claims. When W_1 follows an $\mathrm{Exp}(\lambda)$ distribution with parameter $\lambda > 0$, then $(M^\alpha)_{\alpha \geq 0}$ is a *homogeneous Poisson process with intensity* λ. Renewal processes hence generalise homogeneous Poisson processes. From a modelling perspective, this is desirable since the distribution of the arrival of rare events like windstorms is heavy tailed and therefore not well described by the exponential distribution. In the following, assume that $0 < \mathrm{E}[W_1] < \infty$ and define $\lambda := 1/\mathrm{E}[W_1]$. Then, a strong law of large numbers applies for the renewal process, that is, $M^\alpha/\alpha \to \lambda$ almost surely. Under the additional assumption that $\mathrm{var}(W_1) < \infty$, a central limit theorem holds for M^α,

$$\alpha^{-1/2}(M^\alpha - \alpha\lambda) \xrightarrow{d} \mathrm{N}(0, \sigma^2), \tag{2.1}$$

where $\sigma^2 = \mathrm{var}(W_1)\lambda^3$. While the homogeneous Poisson process admits explicit pre-limit formulas for expectation and variance, that is, $\mathrm{E}[M^\alpha]/\alpha = \lambda, \mathrm{var}(M^\alpha)/\alpha = \mathrm{var}(W_1)\lambda^3$, the renewal process satisfies these equalities only asymptotically as $\alpha \to \infty$. In the context of this work, the parameter α is interpreted mostly as being proportional to the number of policyholders. That is why we decided to use the variable name α instead of t.

The next class of claim count distributions we wish to present are the mixed Poisson processes. Let $\widetilde{M}^\alpha$ be a homogeneous Poisson process with $\lambda = 1$ and let $\mu\colon [0,\infty) \to [0,\infty)$ be increasing, right-continuous and satisfy $\mu(0) = 0$. A function μ satisfying these properties is denoted a *mean value function*. Suppose that θ is a positive random variable independent of $(\widetilde{M}^\alpha)$ and μ a mean value function. The process $M^\alpha := \widetilde{M}^{(\theta\mu(\alpha))}$ is denoted a *mixed Poisson process with mixing variable* θ. When $\theta \equiv 1$, $M^\alpha := \widetilde{M}^{\mu(\alpha)}$ is called a *(general) Poisson process*. In this case, for any $0 < \alpha_1 < \alpha_2 < \infty$, it holds that $M^{\alpha_2} - M^{\alpha_1}$ follows a Poisson distribution with parameter $\mu(\alpha_2) - \mu(\alpha_1)$. The idea of randomising the mean value function also has an interpretation in claim count modelling. In the example of a portfolio of car insurance policies, a realisation $\theta(\omega)$ can represent properties of an insured individual, such as driving skills, age or driving experience. This way, a mixed Poisson process can capture more variability in the path space than a general Poisson process. An important example of mixed Poisson processes is given by setting $\mu(\alpha) = \alpha$ and letting θ be Gamma distributed with positive parameters (γ, β) (in shape–rate parametrisation). That is, θ has a pdf given by

$$f(x) = \frac{\beta^\gamma}{\Gamma(\gamma)} x^{\gamma-1} e^{-\beta x}$$

for $x > 0$, where Γ is the gamma function. It can be shown that M^α then follows a negative binomial distribution with parameters $(\beta/(\alpha+\beta), \gamma)$, i.e.

$$\mathrm{P}(M^\alpha = k) = \frac{\Gamma(k+\gamma)}{k!\Gamma(\gamma)} \left(\frac{\beta}{\alpha+\beta}\right)^\gamma \left(1 - \frac{\beta}{\alpha+\beta}\right)^k$$

for $k \in \mathbb{N}_0$. In claim count modelling, the negative binomial distribution is often the preferred choice compared to the Poisson distribution, because it relaxes the rather restrictive property $\mathrm{E}[M^\alpha] = \mathrm{var}[M^\alpha]$ of a Poisson distributed M^α. As a matter of fact, if $\mathrm{var}(\theta) > 0$ and $\mu(\alpha) > 0$, using the variance decomposition formula applied to M^α and θ, it can be shown that the mixed Poisson process always satisfies $\mathrm{var}(M^\alpha) > \mathrm{E}[M^\alpha]$. This property is called *overdispersion* and is one of the major differences between general Poisson processes and mixed Poisson processes.

We now introduce models for the total claim amount. Let (M^α) be any of the above defined claim count processes and assume that (Z_k) is an iid sequence of positive random variables, independent of (M^α). Consider

$$C^\alpha = \sum_{k=1}^{M^\alpha} Z_k.$$

Here, Z_k models the size of the kth claim payment. Different choices of claim count processes M^α and laws of Z_1 give rise to most of the commonly used models for the total claim amount C^α. When M^α is a homogeneous Poisson process, C^α is called a *compound Poisson model (or Cramér-Lundberg model).* When M^α is a renewal process, the resulting C^α is called a *renewal model (or Sparre-Anderson model).*

We will now outline some of the properties, given C^α is a renewal model with $\mathrm{var}(W_1) < \infty, \mathrm{var}(Z_1) < \infty$. Using tower property, variance decomposition and the independence between M^α and (Z_k), one can derive the following formulas for the expectation and variance of C^α, which are given by

$$\mathrm{E}[C^\alpha] = \mathrm{E}[M^\alpha]\,\mathrm{E}[Z_1], \quad \mathrm{var}(C^\alpha) = \mathrm{E}[M^\alpha]\,\mathrm{var}(Z_1) + \mathrm{var}(M^\alpha)\,\mathrm{E}[Z_1]^2.$$

Moreover, similar to the renewal process for claim counts, a strong law of large numbers and a central limit theorem also apply for C^α. Indeed, it holds that $C^\alpha/\alpha \to \lambda\,\mathrm{E}[Z_1]$ almost surely as well as

$$\alpha^{-1/2}\left(C^\alpha - \alpha\lambda\,\mathrm{E}[Z_1]\right) \xrightarrow{d} \mathrm{N}(0, \sigma^2), \tag{2.2}$$

where $\sigma^2 = \lambda\left(\mathrm{E}[Z_1^2] + (\lambda^2\,\mathrm{var}(W_1) - 1)\,\mathrm{E}[Z_1]^2\right)$.

Let us now consider $T \in \mathbb{N}$ (non-random) time periods and an iid sequence (Z_k, D_k), independent of (M^α), where D_1 takes values in $\{1, \dots, T\}$ and $\mathrm{P}(D_1 = t) = q_t > 0$. As above, Z_k is positive and models the size of the kth claim payment, while D_k is the delay between accident event and the time of claim payment. To be more precise, $\{D_k = t\}$ is the event that the kth accident event causes a claim payment of size Z_k in the time period $(t-1, t]$. Consider the random vector with values in $\mathbb{R}^T$, given by

$$C^\alpha = \sum_{k=1}^{M^\alpha} \left(\mathbb{1}_{\{D_k=1\}}, \dots, \mathbb{1}_{\{D_k=T\}}\right) Z_k. \tag{2.3}$$

If M^α is a renewal process and $\mathrm{var}(W_1) < \infty, \mathrm{var}(Z_1) < \infty$, C^α satisfies a multivariate central limit theorem,

$$\alpha^{-1/2}(C^\alpha - \mathrm{E}[C^\alpha]) \xrightarrow{d} \mathrm{N}_T\left(0, \Sigma\right), \tag{2.4}$$

where

$$\begin{aligned}\Sigma_{s,t} =\lambda\Big(&\mathrm{E}[\mathbb{1}_{\{D_1=s\}}\mathbb{1}_{\{D_1=t\}}Z_1^2] \\ &+ (\lambda^2\,\mathrm{var}(W_1) - 1)\,\mathrm{E}[\mathbb{1}_{\{D_1=s\}}Z_1]\,\mathrm{E}[\mathbb{1}_{\{D_1=t\}}Z_1]\Big).\end{aligned} \tag{2.5}$$

Note that the expression (2.5) generalises the limiting variance in the central limit theorem of the renewal model (2.2) with $T = 1$ and also generalises the limiting variance in the central limit theorem of the renewal process for claim counts (2.1) when choosing $T = 1$ and $Z_1 = 1$. In the special case where M^α is a homogeneous Poisson process with parameter $\lambda > 0$, it holds that $\mathrm{var}(W_1) = 1/\lambda^2$, such that Σ in (2.5) is a diagonal matrix. Therefore, in the limit, there is no dependence between the components of C^α (centred and scaled). It turns out that this property even holds in the pre-limit. More concretely, as a consequence of Mikosch (2009, Theorem 3.3.6), when M^α is a homogeneous Poisson process, then the vector C^α from (2.3) has independent components.

The (asymptotic) dependence between the components of general models C^α satisfying a central limit theorem will play a central role in the study of convergence of valuations (cf. Paper II). In Paper I, we consider both a general model as in (2.3), where M^α is assumed to satisfy a strong law of large numbers, as well as a special model where the claim counts follow the distribution of a homogeneous Poisson process.

2.2 Risk measures

In this section, we let $X \in \mathcal{X} = L^\infty(\mathrm{P})$ describe the discounted net worth of an (actuarial or financial) position at the end of a fixed period of time. In many cases, the occurring definitions can be extended to $\mathcal{X} = L^0(\mathrm{P})$. Negative values of X are net losses while positive values are net gains. The goal is to assign a real number to each $X \in \mathcal{X}$ which measures the risk in connection with the position X. Moreover, risk measures provides an interpretation of risk as a capital requirement, that is, the additional amount needed to make the position X acceptable from the perspective of a regulator. This viewpoint will play a role in the cost-of-capital valuation framework that Paper II deals with (cf. Section 2.3). The outline given here is adapted from Föllmer and Schied (2016).

A mapping $\rho\colon \mathcal{X} \to \mathbb{R}$ is a *(monetary) risk measure*, if for all $X, Y \in \mathcal{X}, m \in \mathbb{R}$,

$$\begin{aligned} &\text{(i) } \rho(0) = 0 \textit{ (normalisation)},\\ &\text{(ii) } X \le Y \text{ implies } \rho(X) \ge \rho(Y) \textit{ (monotonicity)},\\ &\text{(iii) } \rho(X+m) = \rho(X) - m \textit{ (cash additivity)}. \end{aligned} \tag{2.6}$$

Moreover, ρ is called

$$\begin{aligned} &\text{(iv) } \textit{positively homogeneous}, \text{ if } \rho(\lambda X) = \lambda\rho(X), \text{ for all } \lambda \ge 0,\\ &\text{(v) } \textit{convex, if } \rho(\lambda X + (1-\lambda)Y) \le \lambda\rho(X) + (1-\lambda)\rho(Y) \text{ for all } \lambda \in [0,1]. \end{aligned} \tag{2.7}$$

Note that (iv) implies (i) choosing $\lambda = 0$. A risk measure satisfying (i)–(v) is called *coherent*. Coherent risk measures have the property that they admit a so-called robust representation of the form

$$\rho(X) = \sup_{\mathrm{Q} \in \mathcal{Q}} \mathrm{E}_Q[-X], \tag{2.8}$$

where $\mathcal{Q}$ is some set of probability measures on Ω. The *acceptance set* of a risk measure ρ is given by

$$\mathcal{A}_\rho := \{X \in \mathcal{X}\colon \rho(X) \le 0\}.$$

It turns out that the acceptance set entirely characterises the risk measure via $\rho(X) = \inf\{m \in \mathbb{R} \colon m + X \in \mathcal{A}_\rho\}$. This allows to interpret $\rho(X)$ as a capital requirement (cf. Section 2.3). In this context, $\rho(X)$ is seen as the minimal additional amount needed in order to turn X into an acceptable position from the perspective of a supervising agency, since by property (iii), $\rho(X + \rho(X)) = 0$. Moreover, properties (iv) and (v) are equivalent to $\mathcal{A}_\rho$ being conic and convex, respectively. As a result, there exists a one-to-one relationship between coherent risk measures and convex cones in $\mathcal{X}$.

We now turn to a model with $T \in \mathbb{N}$ periods. Let $\{\emptyset, \Omega\} = \mathcal{F}_0 \subset \mathcal{F}_1 \subset \cdots \subset \mathcal{F}_T \subset \mathcal{F}$ and $L_t^\infty = L^\infty(\Omega, \mathcal{F}_t, \mathrm{P})$. Informally, similarly to the above, a dynamic risk measure will be interpreted as the minimal capital requirement needed to make a position at time $t \in \mathcal{T}$ acceptable. This quantity will therefore be an $\mathcal{F}_t$-measurable random variable. More concretely, in analogy to (2.6) and (2.7), a mapping $\rho(\cdot \mid \mathcal{F}_t) \colon \mathcal{X} \to L_t^\infty$ is denoted a *conditional risk measure*, if for all $X, Y \in \mathcal{X}, X_t \in L_t^\infty$,

(i) $\rho(0 \mid \mathcal{F}_t) = 0$ *(normalisation)*,
(ii) $X \le Y$ implies $\rho(X \mid \mathcal{F}_t) \ge \rho(Y \mid \mathcal{F}_t)$ *(monotonicity)*, (2.9)
(iii) $\rho(X + X_t \mid \mathcal{F}_t) = \rho(X \mid \mathcal{F}_t) - X_t$ *(conditional cash additivity)*.

Moreover, ρ is called

(iv) *conditionally positively homogeneous*, if $\rho(\lambda X \mid \mathcal{F}_t) = \lambda \rho(X \mid \mathcal{F}_t)$, for all $\lambda \in L_t^\infty, \lambda \ge 0$,
(v) *conditionally convex, if* $\rho(\lambda X + (1-\lambda)Y) \le \lambda \rho(X \mid \mathcal{F}_t) + (1-\lambda)\rho(Y \mid \mathcal{F}_t)$ for all $\lambda \in L_t^\infty, \lambda \in [0,1]$.
(2.10)

Again, $\rho(\cdot \mid \mathcal{F}_t)$ is uniquely determined by its (now random) acceptance set

$$\mathcal{A}_t = \{X \in L^\infty \colon \rho(X \mid \mathcal{F}_t) \le 0\}$$

via the relation

$$\rho(X \mid \mathcal{F}_t) = \operatorname{ess\,inf}\{Y_t \in L_t^\infty \colon X + Y_t \in \mathcal{A}_t\}.$$

We next provide representations of two conditional risk measures commonly used in practice. We only state definitions of the conditional risk measures noting that unconditional versions can be obtained by conditioning on the sigma algebra $\mathcal{F}_0 = \{\emptyset, \Omega\}$. Fix $u \in (0,1)$ and consider the acceptance set

$$\mathcal{A}_t = \{X \in L^\infty \colon \mathrm{P}(X < 0 \mid \mathcal{F}_t) \le u\}.$$

The conditional risk measure corresponding to $\mathcal{A}_t$ is called *conditional Value at Risk at level* u and admits the representations

$$\begin{aligned} \mathrm{V@R}_u(X \mid \mathcal{F}_t) &= \operatorname{ess\,inf}\{m_t \in L_t^\infty \colon \mathrm{P}(X + m_t < 0 \mid \mathcal{F}_t) \le u\} \\ &= F^{-1}_{-X|\mathcal{F}_t}(1-u), \end{aligned}$$

where the latter is the lower conditional quantile function as introduced in (3.5). Conditional on $\mathcal{F}_t$, the conditional Value at Risk at level u can be interpreted as the minimal amount of capital which, if added to X, bounds the probability of a negative outcome by u. A second interpretation is that, conditional on $\mathcal{F}_t$, it is the minimal loss value $-x_u$ (typically positive), such that the probability of a larger loss $\{-X > -x_u\}$ is bounded by u. While loss probabilities are dealt with, the size of losses is not taken into account by the Value at Risk. The conditional Value at Risk satisfies properties (i)–(iv) in (2.9)–(2.10). However, it fails to satisfy the conditional convexity property (v) and does hence not belong to the class of coherent conditional risk measures. In financial terms, this has the effect that the Value at Risk might penalise diversification instead of encouraging it.

The *conditional average Value at Risk at level* u admits the representations

$$\begin{aligned}
\mathrm{AV@R}_u(X \mid \mathcal{F}_t) &= \operatorname*{ess\,sup}_{\mathrm{Q} \in \mathcal{Q}_t^u} \mathrm{E}_{\mathrm{Q}}[-X \mid \mathcal{F}_t] \\
&= \frac{1}{u} \int_0^u \mathrm{V@R}_s(X \mid \mathcal{F}_t) ds \\
&\geq \mathrm{V@R}_u(X \mid \mathcal{F}_t).
\end{aligned}$$

Here, $\mathcal{Q}_t^u = \{\mathrm{Q} \in \mathcal{M}_1(\mathrm{P})\colon\ \mathrm{Q} = \mathrm{P} \text{ on } \mathcal{F}_t, d\,\mathrm{Q}/d\,\mathrm{P} \leq u\}$, $\mathcal{M}_1(\mathrm{P})$ is the set of probability measures on $(\Omega, \mathcal{F})$ which are absolutely continuous with respect to P and $d\,\mathrm{Q}/d\,\mathrm{P}$ is the Radon-Nikodym derivative of Q with respect to P. Since the conditional average Value at Risk is coherent, it admits a robust representation as in the first line (cf. (2.8) in the unconditional case). If the distribution $\mathcal{L}(X \mid \mathcal{F}_t)(\omega)$ is continuous for almost all $\omega \in \Omega$, one can furthermore express the conditional average Value at Risk as a conditional expected loss, conditional on the u-fraction of worst outcomes.

2.3 Cost-of-capital valuation

We give a brief outline of the valuation procedure of multi-period liability cash flows under capital requirements dealt with in Paper II. The approach is based on Engsner et al. (2017). A more detailed description can be found in (Paper II, Sections 1 – 4).

For illustration purposes, we first consider a one-period model. Let C_1 denote a random discounted payment at the future point in time $T = 1$. We assume that no further payments occur after that time. Regulation obliges the company to set aside a buffer capital R_0 to ensure that C_1 can be paid with a high probability. The required availability of R_0 influences the value V_0 of the liability C_1. The idea of the valuation procedure lies in a transfer of the liability to an external agent which does not have any other assets or liabilities (an empty company). The value V_0 should then be thought of as the minimal amount that the insurance company needs to pay to the external agent such that the latter accepts to compensate for C_1 and to have the buffer capital R_0 available at time zero. The amount $R_0 - V_0$ needed to comply with regulatory capital requirements is provided by a capital provider. This capital provider demands an expected return of $\eta_0 > 0$ on her investment, giving rise to the equation

$$\frac{\mathrm{E}[R_0 - C_1]}{R_0 - V_0} = 1 + \eta_0. \tag{2.11}$$

Note that the term inside the expectation in the numerator is the net worth of the empty company at time one, i.e. $(V_0 - C_1) + (R_0 - V_0) = R_0 - C_1$. Solving (2.11) for V_0 gives

$$V_0 = \frac{1}{1+\eta_0}\,\mathrm{E}[C_1] + \frac{\eta_0}{1+\eta_0} R_0,$$

a convex combination of the expected payment and the capital requirement. Alternatives to (2.11) taking utility functions or limited liability of the capital provider into account, are discussed in (Paper II, Section 4.1). The capital requirement R_0 is assumed to be given by a monetary risk measure of the negative discounted payment, i.e. $R_0 = \rho(-C_1)$. Here, ρ is supposed to be positively homogeneous, but not necessarily convex (cf. Section 2.2).

We continue to look at a multi-period case, where $T \in \mathbb{N}$ is a finite time horizon and C is a random discounted cashflow vector taking values in $\mathbb{R}^T$. Let $\mathcal{F}_0 = \{\Omega, \emptyset\} \subset \mathcal{F}_1 \subset \cdots \subset \mathcal{F}_T \subset \mathcal{F}$ a filtration such that $\sigma(C_{\leq t}) \subset \mathcal{F}_t$. In this case, the value of the liability cash flow will be modelled as an $(\mathcal{F}_t)$-adapted stochastic process and denoted by $V_t = V_t(C)$, $t = 0, \ldots, T$. As in the one-period case, this value should be thought of as the minimal amount of capital that the insurance company needs to pay to an external agent such that the latter accepts to manage $C_{>t}$. Moreover, due to regulation, the future cashflow $C_{>t}$ requires an available capital at time t, which is given by a conditional risk measure $R_t = \rho\left(-C_{t+1} - V_{t+1} \mid \mathcal{F}_t\right)$ of the negative sum of claim payments and cash flow value at time $t+1$. Since the external agent has already received the amount V_t from the insurance company, in order to comply with regulation constraints, it only requires the remaining amount of $R_t - V_t$, which is made available by a capital provider expecting a return $\eta_t > 0$ on the investment. Altogether, this motivates the equality

$$\mathrm{E}[R_t - C_{t+1} - V_{t+1} \mid \mathcal{F}_t] = (1+\eta_t)(R_t - V_t).$$

Note that, in contrast to equation in a single period (cf. (2.11)), the net worth of the empty company at time $t+1$ is given by $V_t - C_{t+1} + (R_t - V_t) - V_{t+1} = R_t - C_{t+1} - V_{t+1}$, where V_{t+1} is the amount needed to be reserved for future payments. Setting $Y = C_{t+1} + V_{t+1}$ and solving for V_t yields

$$V_t = \varphi_t(Y) := \frac{1}{1+\eta_t}\,\mathrm{E}[Y \mid \mathcal{F}_t] + \frac{\eta_t}{1+\eta_t}\rho(-Y \mid \mathcal{F}_t). \tag{2.12}$$

It is sensible to set $V_T = 0$, since no more payments occur after time T. The remaining values can then be determined by means of a backward recursion

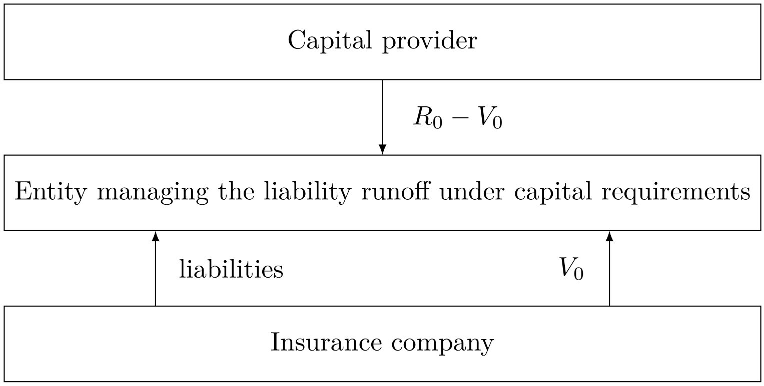

Figure 2.1: Transfer of liabilities and cash amounts at time $t = 0$

following (2.12). More concretely, for $t < T$, we let

$$V_T(C) = 0, \quad V_t(C) = \varphi_t(C_{t+1} + V_{t+1}(C)). \tag{2.13}$$

Assuming that ρ is a positively homogeneous conditional monetary risk measure (cf. Section 2.2), the mappings $\varphi_t\colon L^1(\mathcal{F}_T) \to L^1(\mathcal{F}_t)$ satisfy

$$\varphi_t(aY + \widetilde{Y}) = a\varphi_t(Y) + \widetilde{Y} \tag{2.14}$$

for $a \geq 0, \widetilde{Y} \in L^1(\mathcal{F}_t)$. As a consequence, alternatively to (2.13), we can express the value process via

$$V_T(C) = 0, \quad V_t(C) = \varphi_t \circ \cdots \circ \varphi_{T-1}\Bigg(\sum_{s=t+1}^{T} C_s\Bigg), \quad t < T,$$

where $\circ$ denotes composition. In particular, by (2.14), for any constants $a \geq 0$, $b \in \mathbb{R}^T$, it holds that

$$V_t(aC + b) = aV_t(C) + \sum_{s=t+1}^{T} b_s. \tag{2.15}$$

2.4 Convergence of valuations

Having presented the valuation method for a single cash flow in Section 2.3, let us now consider sequences of cashflows $C^1, C^2, C^3, \dots$ with values in $\mathbb{R}^T$. We assume that there exist non-random sequences $a_n > 0, b_n \in \mathbb{R}^T$, such that the scaled and centred vectors $X^n := a_n^{-1}(C^n - b_n)$ satisfy a central limit theorem, i.e. $\mathcal{L}(X^n) \xrightarrow{w} \mathcal{L}(X)$, where the latter is the law of a non-degenerate multivariate Gaussian. For the outline in this introduction, we let $(\mathcal{F}_t^n)$ be the filtration generated by X^n, that is, $\mathcal{F}_t^n = \sigma(X_{\leq t}^n) \subset \mathcal{F}$, $\mathcal{F}_t = \sigma(X_{\leq t}) \subset \mathcal{F}$ for $t = 1, \dots, T$ and $\mathcal{F}_0 = \mathcal{F}_0^n = \{\Omega, \emptyset\}$. As an example, consider the vector of discounted incremental claim payments

$$C_t^n = \sum_{k=1}^{M^n} \mathbb{1}_{\{D_k = t\}} Z_k, \tag{2.16}$$

for $t = 1, \dots, T$, where M^n denotes the total number of claim payments, $M^n \to \infty$ almost surely as $n \to \infty$, $D_k \in \{1, \dots, T\}$ denotes the time and Z_k denotes the size of the discounted kth claim's payment. Under certain assumptions (cf. for example (2.3) in Section 2.1 with $\alpha = n$), the scaled and centred vector

$$X^n = \mathrm{E}[M^n]^{-1/2} \left(C^n - \mathrm{E}[C^n]\right)$$

satisfies a central limit theorem, i.e. X^n converges in distribution to a multivariate Gaussian X. Here, the index n can be seen as the number of insurance contracts. Let the mappings $\varphi_t^n = 0, \dots, T-1$ and $V_t^n, t = 0, \dots, T$ be defined as in (2.12), with $\mathcal{F}_t$ replaced by $\mathcal{F}_t^n$. The goal is to obtain the value of the cash flow at time zero, that is, $V_0^n(C^n)$. Computing this quantity arithmetically is rarely possible (a simple example for $T = 2$ periods where the calculation is possible is presented in (Paper II, Section 6) and applying Monte Carlo methods to approximate conditional distributions gives rise to computational challenges. However, applying properties of the multivariate Gaussian distribution, the value of the limit X can be calculated via the formula

$$\begin{aligned} V_0(X) &= \mathrm{E}\left[\sum_{t=1}^{T} X_t\right] \\ &+ \sum_{t=1}^{T} \varphi_{t-1}(\varepsilon_t) \Bigg(\operatorname{var}\left(\sum_{u=t}^{T} X_u \mid X_{\leq t-1}\right) - \operatorname{var}\left(\sum_{u=t}^{T} X_u \mid X_{\leq t}\right)\Bigg)^{1/2}, \end{aligned} \tag{2.17}$$

where each ε_t is standard normally distributed and independent of $\mathcal{F}_{t-1}$. Note that the conditional variances in (2.17) are non-random. We remark that (2.17) resembles the standard deviation principle in classical premium calculation (cf. Mikosch (2009, Section 3.1.3)) with the difference that here, differences of conditional variances appear. If it was known that $V_0^n(X^n) \to V_0(X)$ as $n \to \infty$, then, applying (2.15), the approximation

$$\begin{aligned} V_0^n(C^n) = V_0^n(a_n X^n + b_n) &= a_n V_0^n(X^n) + \sum_{s=1}^{T} b_{n,s} \\ &\approx a_n V_0(X) + \sum_{s=1}^{T} b_{n,s}. \end{aligned}$$

would be justified.

It turns out that the weak convergence $\mathcal{L}(X^n) \xrightarrow{w} \mathcal{L}(X)$ is not sufficient to guarantee $V_0^n(X^n) \to V_0(X)$. Assume that, additionally, for each $t = 1, \dots, T-$

1, and each convergent sequence $x^n_{\leq t} \to x_{\leq t} \in \mathbb{R}^t$,

$$\mathcal{L}\left(X^n \mid X^n_{\leq t} = x^n_{\leq t}\right) \to \mathcal{L}\left(X \mid X_{\leq t} = x_{\leq t}\right) \tag{2.18}$$

in Wasserstein-1-distance. Then, $V^n_0(X^n) \to V_0(X)$ holds true. In Chapter 3 we provide more details about property (2.18).

2.5 Chain ladder reserving

Reserving in insurance is the process of setting aside funds to meet future claim obligations arising from insurance contracts already written. One of the most popular methods for predicting these obligations is the chain ladder method whose core assumption is that the conditional expectation of future claim payments is proportional to the obligations of the present date. In order to estimate uncertainty, Mack (1993) suggests an estimator of the mean squared error of the chain ladder predictor under three model assumptions. The following outline is based on Mack (1993).

Fix $T \in \mathbb{N}$ and let $i, t \in \{1, \dots, T\}$. In accordance with classical reserving approaches that structure claims data in run-off triangles, we assume T accident and T development years. We let $C_{i,t}$ denote the cumulative total claims amount from accidents in accident year i paid up to and including development year t. The sigma algebra of the past is given by $\mathcal{D} = \sigma(C_{i,t} \colon i+t \leq T+1)$ and corresponds to the available observations at the present time. The goal of the chain ladder method is to estimate the ultimate claims amount $C_{i,T}$ and the outstanding claims reserve $R_i = C_{i,T} - C_{i,T-i+1}$ for accident years $i = 2, \dots, T$.

	1	2	3	4	5
1	$C_{1,1}$	$C_{1,2}$	$C_{1,3}$	$C_{1,4}$	$C_{1,5}$
2	$C_{2,1}$	$C_{2,2}$	$C_{2,3}$	$C_{2,4}$	
3	$C_{3,1}$	$C_{3,2}$	$C_{3,3}$		
4	$C_{4,1}$	$C_{4,2}$			
5	$C_{5,1}$				

Table 2.1: Run-off triangle with $T = 5$. The chain ladder method provides estimators for future cumulative claims amounts (blank cells).

The *chain ladder method* consists of calculating

$$\widehat{f}_t = \frac{\sum_{i=1}^{T-t} C_{i,t+1}}{\sum_{i=1}^{T-t} C_{i,t}} \tag{2.19}$$

for $1 \leq t \leq T-1$ and then estimating $C_{i,T}$ and R_i by

$$\widehat{C}_{i,T} = C_{i,T-i+1} \widehat{f}_{T-i+1} \cdot \ldots \cdot \widehat{f}_{T-1}, \quad \widehat{R}_i = C_{i,T-i+1} \left(\widehat{f}_{T-i+1} \cdot \ldots \cdot \widehat{f}_{T-1} - 1 \right). \tag{2.20}$$

Mack (1993) suggests the following three conditions on a model in order to be in line with the chain ladder method.

$$\begin{aligned}
&\text{(i) For each } t = 1, \dots, T-1, \text{there exists } f_t > 0, \text{ such that} \\
&\qquad \mathrm{E}[C_{i,t+1} \,|\, C_{i,1}, \dots, C_{i,t}] = f_t C_{i,t}, \\
&\text{(ii) for each } t = 1, \dots, T-1, \text{there exists } \sigma_t^2 > 0, \text{ such that} \\
&\qquad \mathrm{var}(C_{i,t+1} \,|\, C_{i,1}, \dots, C_{i,t}) = \sigma_t^2 C_{i,t}, \\
&\text{(iii) } (C_{1,1}, \dots, C_{1,T}), \dots, (C_{T,1}, \dots, C_{T,T}) \text{ are independent.}
\end{aligned} \tag{2.21}$$

The conditions (i)–(iii) in (2.21) are referred to as *Mack's distribution-free chain ladder model (MCL)*. Under the assumptions (i) and (iii), it is shown in Mack (1993) that $\widehat{C}_{i,T}$ is an unbiased estimator of $\mathrm{E}[C_{i,T} \mid \mathcal{D}]$ in the sense that $\mathrm{E}[\widehat{C}_{i,T}] = \mathrm{E}[\mathrm{E}[C_{i,T} \mid \mathcal{D}]]$ and analogously for the estimator $\widehat{R}_i$. Moreover, under (i)–(iii), an estimator for the mean squared error of prediction

$$\mathrm{E}\left[(C_{i,T} - \widehat{C}_{i,T})^2 \,\middle|\, \mathcal{D}\right]$$

is derived, which is given by

$$(\widehat{C}_{i,T})^2 \sum_{t=T-i+1}^{T-1} \frac{\widehat{\sigma}_t^2}{\widehat{f}_t^2} \left(\frac{1}{\widehat{C}_{i,t}} + \frac{1}{\sum_{j=1}^{T-t} C_{j,t}} \right), \tag{2.22}$$

where $\widehat{C}_{i,T-i+1} = C_{i,T-i+1}$, $\widehat{C}_{i,t} = C_{i,T-i+1} \widehat{f}_{T-i+1} \cdot \ldots \cdot \widehat{f}_{t-1}$ for $t > T-i+1$ and

$$\widehat{\sigma}_t^2 = \frac{1}{T-t-1} \sum_{i=1}^{T-t} C_{i,t} \left(\frac{C_{i,t+1}}{C_{i,t}} - \widehat{f}_t \right)^2 \tag{2.23}$$

can be shown to be unbiased estimators of σ_t^2 (cf. (ii) in (2.21)) for $t = 1, \dots, T-2$. Note that (2.23) is not well-defined for $t = T-1$, so that σ_{T-1}^2 needs to be estimated in a different way. Mack (1993) suggests to set $\widehat{\sigma}_{T-1}^2 = \min(\widehat{\sigma}_{T-2}^4 / \widehat{\sigma}_{T-3}^2, \min(\hat{\sigma}_{T-3}^2, \widehat{\sigma}_{T-2}^2))$ which is based on an extrapolation of the sequence $\widehat{\sigma}_1^2, \cdots, \widehat{\sigma}_{T-2}^2$.

For most models, the conditions in (MCL) in (2.21) are not fulfilled. The question arises whether the use of the estimators in (2.20) and (2.22) can still be justified in standard models for multi-period claims amounts. This question is addressed in Paper I.

Chapter 3

Convergence of conditional distributions

This chapter provides the mathematical background on the convergence of valuations of liability cashflows discussed in Section 2.4.

3.1 Wasserstein and adapted Wasserstein spaces

Let $p \in [1, \infty)$ and $T \in \mathbb{N}$. The *Wasserstein-p-space* is given by

$$\mathcal{P}_p(\mathbb{R}^T) := \{\mu \in \mathcal{P}(\mathbb{R}^T) \colon \, \mathrm{E}_\mu[\|X\|^p] < \infty\}. \tag{3.1}$$

The notation $\mathrm{E}_\mu[X]$ means that $\mathcal{L}(X) = \mu$. For two elements μ, ν of $\mathcal{P}_p(\mathbb{R}^T)$, we let $\Pi(\mu, \nu)$ denote the set of all couplings of μ and ν, that is, the set of all $\pi \in \mathcal{P}(\mathbb{R}^{2T})$ having μ and ν as marginals. On the Wasserstein-p-space, one can define a metric via

$$\mathcal{W}_p(\mu, \nu) := \inf_{\pi \in \Pi(\mu,\nu)} \mathrm{E}_\pi[\|X - Y\|^p]^{\frac{1}{p}} \tag{3.2}$$

denoted the *Wasserstein-p-distance between μ and ν*. Similar to before, the notation $\mathrm{E}_\pi[\|X - Y\|^p]$ means that $\mathcal{L}((X, Y)) = \pi$. In the given set-up, the infimum in (3.2) is always attained. This metric is inspired by the idea of an optimal transport of mass from μ to ν. In this context, an element $\pi \in \Pi(\mu, \nu)$ is also referred to as a *transference plan*. An important role in the study of (3.2) is played by the set of *Monge couplings*. A Monge coupling of μ and ν is an element $\pi \in \Pi(\mu, \nu)$, such that $\nu = \mu \circ H^{-1}$ for some measurable $H \colon \mathbb{R}^T \to \mathbb{R}^T$. For example, when $p = 2$ and μ, ν are absolutely continuous with respect to the Lebesgue measure, the unique minimiser in (3.2) is given by a Monge coupling Villani (2009, Theorem 9.4).

In the following, we make the connection to weak convergence. For this purpose, consider $\mu_n, \mu \in \mathcal{P}_p(\mathbb{R}^T), n \in \mathbb{N}$. The following assertions are equivalent.

(i) $\mathcal{W}_p(\mu_n, \mu) \to 0$,

(ii) $\mu_n \overset{w}{\to} \mu$ and $\int \|x\|^p d\mu_n \to \int \|x\|^p d\mu(x)$,

(iii) $\mu_n \overset{w}{\to} \mu$ and $\lim_{R\to\infty} \sup_{n\in\mathbb{N}} \int_{\|x\|\geq R} \|x\|^p d\mu_n(x) = 0$.

Hence, the Wasserstein-p-topology (induced from the Wasserstein-p-distance) is finer than the weak topology. Characterisation (iii) can be rephrased as: For $\mathbb{R}^T$-valued random variables X, X_n with $\mathcal{L}(X) = \mu, \mathcal{L}(X_n) = \mu_n$, we have $X_n \overset{d}{\to} X$ and $(\|X_n\|^p)$ is uniformly integrable. This coincides with the formulation we employ in (Paper II, Theorem 2).

When $t = 1, \dots, T$ models time such that the vectors $X, Y \in \mathbb{R}^T$ describe stochastic processes, it turns out that neither the weak topology nor the Wasserstein-p-topology provide reasonable notions of distance between $\mathcal{L}(X)$ and $\mathcal{L}(Y)$. The reason lies in the fact that the available information at time t, typically modelled by a sequence of sigma algebras $\mathcal{F}_t$, is increasing in t. The information structure of stochastic processes is therefore mostly modelled in terms of filtrations, which have the property that $\mathcal{F}_s \subset \mathcal{F}_t$ whenever $s < t$. Given a point in time t, it is desirable that the laws of two stochastic processes are "close" only if their future laws, conditional on the information available at time t, are close. This property however is obviously not satisfied for the weak and Wasserstein-p-topologies, since "future" and "past" are not taken into account in their definitions.

The following example illustrates this point (cf. Figure 3.1). Let X^n, X take values in $\mathbb{R}^2$ with

$$\begin{aligned} &\mathcal{L}(X_1^n) = \frac{1}{2}\delta_{-\frac{1}{n}} + \frac{1}{2}\delta_{\frac{1}{n}}, \quad \mathcal{L}(X_2^n \mid X_1^n) = \mathbb{1}_{\{X_1^n = -\frac{1}{n}\}}\delta_{-1} + \mathbb{1}_{\{X_1^n = \frac{1}{n}\}}\delta_1, \\ &\mathcal{L}(X_1) = \delta_0, \quad \mathcal{L}(X_2 \mid X_1) = \mathcal{L}(X_2) = \frac{1}{2}\delta_{-1} + \frac{1}{2}\delta_1. \end{aligned} \tag{3.3}$$

In this example, it holds that $\mathcal{L}(X^n) \overset{w}{\to} \mathcal{L}(X)$: Indeed, for $u \in C_b(\mathbb{R}^2)$, we see that

$$\begin{aligned} \mathrm{E}[u(X_1^n, X_2^n)] =& \frac{1}{2}u\left(-\frac{1}{n}, -1\right) + \frac{1}{2}u\left(\frac{1}{n}, 1\right) \\ \to& \frac{1}{2}u(0,-1) + \frac{1}{2}u(0,1) = \mathrm{E}[u(X_1, X_2)]. \end{aligned}$$

Also convergence in Wasserstein-p-distance holds (for any p): Consider the

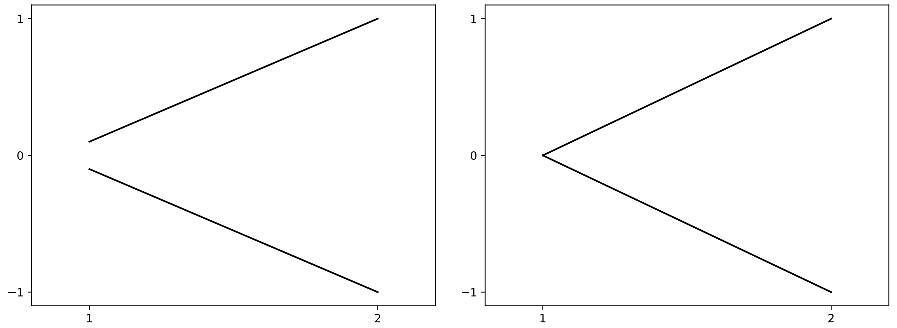

Figure 3.1: The paths (seeing X, X^n as stochastic processes at time points $t_1 = 1$, $t_2 = 2$) of X^n (left) and X (right). The described coupling transports the upper path in the left panel to the upper path in the right panel and analogously for the lower paths.

coupling (X^n, X) specified by setting $X_2^n = X_2$. It then follows that

$$\begin{aligned}\mathcal{W}_p(\mathcal{L}(X^n), \mathcal{L}(X)) &\leq \mathrm{E}\left[\left(\sqrt{(X_1 - X_1^n)^2 + (X_2 - X_2^n)^2}\right)^p\right]^{1/p} \\ &\leq \mathrm{E}\left[\left(\sqrt{\left(\frac{1}{n}\right)^2 + 0^2}\right)^p\right]^{1/p} = \frac{1}{n} \to 0.\end{aligned}$$

The processes in (3.3) have a distinct information structure. Still, as the example shows, in the weak or Wasserstein-p-topology, they are considered close.

To motivate the definition of a distance which takes the available information in a stochastic process indexed by time into account, we turn back to the above mentioned Monge couplings, where for two $\mathbb{R}^T$-valued random variables X, Y, there exists a measurable function $H = (H_1, \ldots, H_T) \colon \mathbb{R}^T \to \mathbb{R}^T$, such that

$$\mathcal{L}\left((Y_1, \ldots, Y_T)\right) = \mathcal{L}\left((H_1(X_1, \ldots, X_T), \ldots, H_T(X_1, \ldots, X_T))\right).$$

Informally, we would like to restrict the set of Monge couplings to those functions H satisfying that

$$\mathcal{L}\left((Y_1, Y_2 \ldots, Y_T)\right) = \mathcal{L}\left((H_1(X_1), H_2(X_1, X_2)) \ldots, H_T(X_1, \ldots, X_T))\right),$$

so that the law of Y_t does not depend on the future $(X_{t+1}, \ldots, X_T)$. This informal idea leads to the definition of the set of *bicausal couplings*

$$\begin{aligned}\Pi_{\mathrm{bc}}(\mu, \nu) := \{\pi \in \Pi(\mu, \nu) \colon & \pi(X_1, \ldots, X_t \mid Y) = \pi(X_1, \ldots, X_t \mid (Y_1, \ldots, Y_t)) \\ & \pi(Y_1, \ldots, Y_t \mid X) = \pi(Y_1, \ldots, Y_t \mid (X_1, \ldots, X_t))\}.\end{aligned}$$

Similar to the above, the notation here is to be understood as $\pi = \mathcal{L}((X, Y))$, $\mathcal{L}(X) = \mu$ and $\mathcal{L}(Y) = \nu$. The *adapted Wasserstein-p-distance* is then defined as

$$\mathcal{AW}_p(\mu,\nu) := \inf_{\pi\in\Pi_{\mathrm{bc}}(\mu,\nu)} \mathrm{E}_\pi[\|X-Y\|^p]^{\frac{1}{p}}. \tag{3.4}$$

Clearly, $\mathcal{AW}_p(\mu,\nu) \geq \mathcal{W}_p(\mu,\nu)$ as $\Pi_{\mathrm{bc}}(\mu,\nu) \subset \Pi(\mu,\nu)$. Hence, the adapted Wasserstein-p-topology is finer than the Wasserstein-p-topology which is again finer than the weak topology (cf. (3.1)).

Returning to the example in (3.3), we have that $(X^n, X) \notin \Pi_{\mathrm{bc}}(\mathcal{L}(X^n), \mathcal{L}(X))$ and adapted Wasserstein-p-convergence of X^n to X fails to apply.

3.2 Regular conditional distributions

In this section, we introduce regular conditional distributions, as they play a central role in the convergence of valuations of Paper II.

Let $\mathcal{G} \subset \mathcal{F}$ a sub-sigma-algebra and $X \in L^1(\Omega, \mathcal{F}, \mathrm{P})$. A real-valued random variable Z is called *conditional expectation of X given $\mathcal{G}$* if

(i) Z is $\mathcal{G}$-measurable,

(ii) for each $A \in \mathcal{G}$, $\mathrm{E}[X\mathbb{1}_A] = \mathrm{E}[Z\mathbb{1}_A]$.

The random variable Z exists and is almost surely unique which justifies to set $\mathrm{E}[X \mid \mathcal{G}] := Z$. For a random variable Y on $(\Omega, \mathcal{F}, \mathrm{P})$ taking values in some Polish space $(E, \mathcal{E})$, one defines $\mathrm{E}[X \mid Y] := \mathrm{E}[X \mid \sigma(Y)]$.

Let us consider a fixed event $B \in \mathcal{B}(\mathbb{R})$, the Borel sigma algebra on $\mathbb{R}$. We can then define the random variable

$$\mu(B \mid \mathcal{G}) := \mathrm{E}[\mathbb{1}_{\{X\in B\}} \mid \mathcal{G}].$$

Since by (i), this definition only holds up to a P-nullset which depends on the set B, it is unclear whether it also leads to a probability measure for almost all $\omega \in \Omega$. The reason is that properties of a probability measure like additivity for disjoint events $(B_i)_{i=1}^\infty \subset \mathcal{B}(\mathbb{R})$ must hold for all ω except for a P-nullset, which is independent of the sequence (B_i). Since there are uncountably many such sequences, it is unclear whether the union of corresponding nullsets forms again a P-nullset.

A *regular conditional distribution of X given $\mathcal{G}$* is a mapping $\kappa\colon (\Omega, \mathcal{B}(\mathbb{R})) \to [0,1]$ which satisfies

(i) for each $B \in \mathcal{B}(\mathbb{R})$, it holds that $\kappa(\omega, B) = \mathrm{E}[\mathbb{1}_{\{X\in B\}} \mid \mathcal{G}](\omega)$ for P-almost all $\omega \in \Omega$,

(ii) for each $\omega \in \Omega$, $B \mapsto \kappa(\omega, B)$ is a probability measure on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$.

In our case where X is real-valued, a regular conditional distribution exists. The same holds if X takes values in a Polish space. In the following, we will use the notation $\mathcal{L}(X \mid Y)(\omega)$ to denote the $\mathcal{P}(\mathbb{R})$-valued random variable $\omega \mapsto \kappa(\omega, \cdot)$, which is the regular conditional distribution of X given $\sigma(Y)$. The given outline is easily extended to X, Y taking values in $\mathbb{R}^{d_1} \times \mathbb{R}^{d_2}$, $d_1, d_2 \in \mathbb{N}$. Finally, for $p \in (0,1)$, we define the conditional lower quantile function $F^{-1}_{X|\mathcal{G}}(p)$ as

$$\omega \mapsto \min\{m \in \mathbb{R} : \kappa(\omega, (-\infty, m]) \geq p\}. \tag{3.5}$$

3.3 Extended weak convergence

In his monograph, Aldous (1981) defines the prediction process of a continuous-time stochastic process $(X_t)_{t\geq 0}$ as the measure-valued process assigning to each t the regular conditional distribution of X given $\mathcal{F}_t$. We will transform his formulation into a discrete-time setting with finite time-horizon where the filtration is generated by the process itself. Let $T \in \mathbb{N}$ and X a random vector in $\mathbb{R}^T$. Define the prediction process $(Z_t)_{t=0,\dots,T}$ by $Z_t := \mathcal{L}(X \mid X_{\leq t})$ taking values in $\mathcal{P}(\mathbb{R}^T)$. Here, $Z_0 = \mathcal{L}(X)$ is a constant while $Z_T = \delta_X$ is a random point mass. Aldous defines *extended weak convergence of X^n to X* as

$$\mathcal{L}(X^n, Z^n) \xrightarrow{w} \mathcal{L}(X, Z)$$

in $\mathcal{P}\left(\mathbb{R}^T \times \mathcal{P}(\mathbb{R}^T)^{T+1}\right)$, where Z^n denotes the prediction process of the $\mathbb{R}^T$-valued random vector X^n. Backhoff-Veraguas et al. (2020) show that when replacing the Euclidean metric by a bounded one in the definition of (3.4), the notion of extended weak convergence is equivalent to $\mathcal{AW}_1(\mathcal{L}(X^n), \mathcal{L}(X)) \to 0$. We apply this result in (Paper II, Theorem 5) to show that adapted Wasserstein-p-convergence is a consequence of the continuous convergence of conditional distributions (cf. Section 3.4).

3.4 Continuous convergence

Continuous convergence of conditional distributions is the central property in Paper II to ensure convergence of valuations. It is a stronger condition than adapted Wasserstein-p-convergence as shown in (Paper II, Theorem 5). When explicit pre-limit laws are available, it is also a condition which is verified in a straightforward manner, as we demonstrate in (Paper II, Theorem 6).

Let $d \in \mathbb{N}$ and let $\mathcal{X}$ be a Polish space. A sequence of mappings $f_n \colon \mathbb{R}^d \to \mathcal{X}$ is *continuously convergent in $\mathcal{X}$*, if for each converging sequence $x_n \to x \in \mathbb{R}^d$, $f_n(x_n)$ converges in $\mathcal{X}$. Taking constant sequences, it follows that f_n converges pointwise to a limit $f \colon \mathbb{R}^d \to \mathcal{X}$. If two sequences x_n, y_n converge to the same limit $x \in \mathbb{R}^d$, then the sequences $f_n(x_n)$, $f_n(y_n)$ have the same limit in $\mathcal{X}$, which is given by $f(x)$. In addition, f is continuous. It turns out that

$f_n \to f$ continuously if and only if $f_n \to f$ uniformly on each compact $K \subset \mathbb{R}^d$. The notion of continuous convergence was first introduced for real-valued functions by Hahn (1921). Even though being more practical to handle than the equivalent uniform convergence on compacts as remarked by Carathéodory (1929), the latter has received far more attention.

Assume that the sequence of measurable functions $g_n \colon \mathbb{R} \to \mathbb{R}$ converges continuously towards $g \colon \mathbb{R} \to \mathbb{R}$. Moreover, suppose that $X^n \xrightarrow{d} X$, where X^n, X are real-valued. By a generalised version of the continuous mapping theorem as stated in Kallenberg (2002, Theorem 4.27), it then follows that $g_n(X^n) \xrightarrow{d} g(X)$. Taking $g \in C(\mathbb{R})$ and the continuously converging constant sequence $g_n = g$ for all n, it can be seen that this indeed generalises the classical continuous mapping theorem.

For $T \in \mathbb{N}$, consider random vectors Z^n, Z taking values in $\mathbb{R}^T$ whose regular conditional distributions, seen as mappings from $\mathbb{R}^t$, $t = 0, \dots, T-1$, into the Polish space $\mathcal{P}(\mathbb{R}^T)$, equipped with the weak topology, converge continuously. In other words, we have *continuous convergence of conditional distributions of* Z^n *to* Z, if

$$\mathcal{L}(Z^n \mid Z^n_{\leq t} = \cdot) \to \mathcal{L}(Z \mid Z_{\leq t} = \cdot) \text{ continuously in } \mathcal{P}(\mathbb{R}^T) \text{ for } t = 0, \dots, T-1.$$

Here, the convergence for $t = 0$ is to be interpreted as $\mathcal{L}(Z^n) \xrightarrow{w} \mathcal{L}(Z)$.

Sweeting (1989) introduced and analysed continuous convergence in the context of conditional distributions. Its importance in the present work is in the induction step in the proof of (Paper II, Theorem 2). There, the generalised continuous mapping theorem as introduced in the previous paragraph is applied to the continuous convergence of conditional distributions of a discrete-time stochastic liability cash flow with finite time horizon. This, together with a uniform integrability assumption, allows to prove convergence of the values of the insurance liability cashflows (cf. Section 2.4).

Chapter 4

Regression trees

This chapter provides the theoretical frameworks of regression trees, change point testing and explains how these concepts can be used to construct node-wise tests for signal detection for the purpose of regularisation. The theory presented in this chapter is employed in Paper III.

4.1 CART regression trees

In this section we introduce the *Classification and Regression Tree (CART)* as presented in Breiman et al. (1984). A CART regression tree is a piecewise constant function $\mu\colon \mathbb{R}^d \to \mathbb{R}$,

$$\mu(x) := \sum_{k=1}^{m} \zeta_k \mathbb{1}_{\{x \in A_k\}}, \tag{4.1}$$

where $(A_k)_{k=1}^m$ is a partition of $\mathbb{R}^d$ (of a particular shape being made precise below) and values $\zeta_k \in \mathbb{R}$. The complexity of a CART regression tree is usually measured by the number of leaves or regions, $m \in \mathbb{N}$. Let $(X^{(i)}, Y^{(i)})_{i=1}^n$ be an iid sequence of observed covariate and response pairs taking values in $\mathbb{R}^d \times \mathbb{R}$. Using this data, the CART method uses a recursive greedy binary splitting approach in order to obtain a partition of the covariate space. This results in a piecewise constant predictor $\mu(x) = \mu(x; (X^{(i)}, Y^{(i)})_{i=1}^n)$ as in (4.1), depending on the data $(X^{(i)}, Y^{(i)})_{i=1}^n$. When using the L^2 loss, this is achieved by recursively computing a minimiser (j^*, ξ^*) of

$$\sum_{i:X^{(i)} \in R_{\text{left}}(j,\xi)} (Y^{(i)} - \overline{Y}_{\text{left}}(j,\xi))^2 + \sum_{i:X^{(i)} \in R_{\text{right}}(j,\xi)} (Y^{(i)} - \overline{Y}_{\text{right}}(j,\xi))^2. \tag{4.2}$$

Here, $\overline{Y}_{\text{left}}(j,\xi)$ is the average of all $Y^{(i)}$ for which $X^{(i)} \in R_{\text{left}}(j,\xi)$, and correspondingly for $\overline{Y}_{\text{right}}(j,\xi)$, where

$$R_{\text{left}}(j,\xi) = \{x \in \mathbb{R}^d : x_j \le \xi\}, \quad R_{\text{right}}(j,\xi) = \{x \in \mathbb{R}^d : x_j > \xi\}$$

are the two disjoint half-spaces corresponding to the split point (j, ξ). For each region A constructed by the CART method, the corresponding value ζ is computed as the average of all responses $Y^{(i)}$ for which $X_i \in A$. In the first step of the algorithm, the tree only consists of a root node, so that $\mu^{(1)}(x) = \overline{Y}$ is a constant. The CART regression tree after one split is given by

$$\mu^{(2)}(x) = \overline{Y}_{\text{left}}(j^*, \xi^*) \mathbb{1}_{\{x \in R_{\text{left}}(j^*, \xi^*)\}} + \overline{Y}_{\text{right}}(j^*, \xi^*) \mathbb{1}_{\{x \in R_{\text{right}}(j^*, \xi^*)\}}.$$

The procedure is then repeated restricted to the covariate subspace $R_{\text{left}}(j^*, \xi^*)$ using only the data $((X^{(i)}, Y^{(i)})_{i=1}^n : X^{(i)} \in R_{\text{left}}(j^*, \xi^*))$ and analogously in the right half-space. This way, $m-1$ splits (corresponding to internal nodes) yield a regression tree as in (4.1) with m disjoint regions forming a partition of $\mathbb{R}^d$ (corresponding to leaves). The method is stopped by e.g. specifying a maximal tree depth or a minimal number of data points in each leaf. By construction of the CART method, each of the sets $A_1, \ldots, A_m$ in (4.1) is given by an intersection of half-spaces whose normal vectors are standard unit vectors of $\mathbb{R}^d$.

The CART regression tree μ, as defined in (4.1), does not contain any information about the internal nodes or the associated regions that arise in the CART tree-growing procedure described above. Strictly speaking, one should therefore distinguish between two objects: the piecewise constant function μ and the tree viewed as a graph, which contains the information on all split points at all nodes. In what follows, we will be slightly imprecise, sometimes referring to the former and sometimes to the latter when discussing a tree.

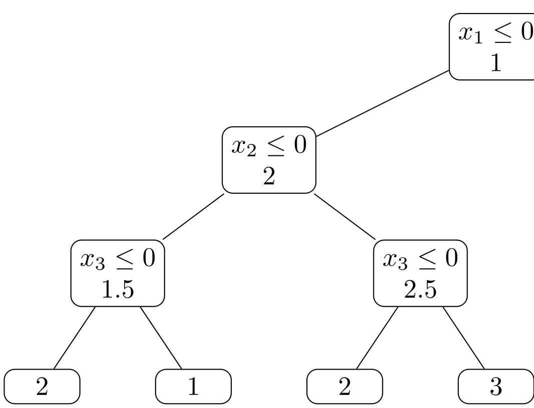

Figure 4.1: Example of a regression tree with $d = 3$ and $m = 5$. Each left child answers the inequality with "true". First row (if existent): split point (j, ξ), second row: value ζ.

4.2 Change point testing

In this section, which is based on Yao and Davis (1986), we introduce a test statistic to test for a change point in mean in a given sequence of independent, normally distributed random variables.

4.2.1 Variance known

Let $\sigma^2 > 0$ be known and $\mu_1 \neq \mu_2$ be unknown real parameters. Consider a sequence of independent observations $Y_1, \ldots, Y_n$, such that $Y_1, \ldots, Y_r$ follow $\mathrm{N}(\mu_1, \sigma^2)$ while $Y_{r+1}, \ldots, Y_n$ follow $\mathrm{N}(\mu_2, \sigma^2)$, where the change point $1 \leq r \leq n$ is unknown. We define a null and an alternative hypotheses by

$$\begin{aligned} &H_0 \colon r = n \\ &H_1 \colon r < n. \end{aligned} \tag{4.3}$$

H_0 means that no change of mean occurs while H_1 corresponds to a change of mean at r from μ_1 to μ_2. The likelihood function for the unknown parameters r, μ_1, μ_2 is given by

$$L(r, \mu_1, \mu_2) := (2\pi\sigma^2)^{-n/2} \exp\left(-\frac{1}{2\sigma^2} \left(\sum_{i=1}^{r} (Y_i - \mu_1)^2 + \sum_{i=r+1}^{n} (Y_i - \mu_2)^2 \right) \right). \tag{4.4}$$

For each fixed r, (4.4) is maximised by $\mu_1 = \overline{Y}_{\leq r}, \mu_2 = \overline{Y}_{>r}$, denoting the averages of the first r, and of the last $n-r$ responses, respectively. As a result, the generalised likelihood ratio of H_0 against H_1 is

$$\frac{\sup_{1 \leq r < n, \mu_1, \mu_2} L(r, \mu_1, \mu_2)}{\sup_{\mu_1} L(n, \mu_1, 0)} \tag{4.5}$$

$$= \frac{\max_{1 \leq r < n} \exp\left(-\frac{1}{2\sigma^2} \left(\sum_{i=1}^{r} (Y_i - \overline{Y}_{\leq r})^2 + \sum_{i=r+1}^{n} (Y_i - \overline{Y}_{>r})^2 \right) \right)}{\exp\left(-\frac{1}{2\sigma^2} \left(\sum_{i=1}^{n} (Y_i - \overline{Y}_{\leq n})^2 \right) \right)}. \tag{4.6}$$

Note that the choice of the value zero in the denominator of (4.5) is of no relevance here. Denoting by r^* a maximiser in the numerator of (4.6), taking the logarithm and multiplying by the constant 2, yields the test statistic

$$T_n^2 := \frac{1}{\sigma^2} \left(\sum_{i=1}^{n} (Y_i - \overline{Y}_{\leq n})^2 - \sum_{i=1}^{r^*} (Y_i - \overline{Y}_{\leq r^*})^2 - \sum_{i=r^*+1}^{n} (Y_i - \overline{Y}_{>r^*})^2 \right) \geq 0, \tag{4.7}$$

which is twice the generalised log-likelihood ratio of H_0 against H_1. Large values of T_n^2 mean significance. In order to construct a valid ε-level test for a given significance level $\varepsilon > 0$, an analysis of the distribution of (4.7) under H_0 is required. The exact pdf of T_n was derived in Hawkins (1977, Theorem 1).

However, its computation gives rise to runtime challenges for large values of n. Using the property that under H_0,

$$\sqrt{T_n^2} =: T_n \stackrel{d}{=} \max_{t=\frac{1}{n},\ldots,\frac{n-1}{n}} \frac{|W_t - tW_1|}{\sqrt{t(1-t)}},$$

where $(W_t)_{t\geq 0}$ is a standard Brownian motion, as well as results from extreme value theory, Yao and Davis (1986) show that a shifted and scaled version of T_n converges to a Gumbel distribution as n tends to infinity. More concretely, denoting by P_0 a probability measure under which H_0 holds true, it is shown that

$$\mathrm{P}_0\left(\frac{T_n - b_n}{a_n} \leq x\right) \to \exp\left(-\frac{2}{\sqrt{\pi}} e^{-x}\right) \tag{4.8}$$

for all $x \in \mathbb{R}$. Here, $a_n = (2\ln_2(n))^{-1/2}$, $b_n = a_n^{-1} + a_n \ln_3(n)/2$, and $\ln_k$ is the k-times iterated natural logarithm, e.g. $\ln_2(\cdot) = \ln(\ln(\cdot))$. The convergence (4.8) also holds uniformly, such that a natural approximation of the law of T_n under H_0 is given by

$$\mathrm{P}_0\left(T_n \leq x\right) \approx \exp\left(-\frac{2}{\sqrt{\pi}} e^{-(a_n x + b_n)}\right). \tag{4.9}$$

However, since (4.8) is very slow, (4.9) provides a very inaccurate approximation for any value of n relevant in practice. As argued in Yao and Davis (1986, Remark 2.3) and based on results from Hawkins (1977),

$$\mathrm{P}_0\left(\frac{T_n - b_n}{a_n} \leq x\right) - \left(\Phi\left(a_n x + \tilde{b}_n\right)\right)^{2\ln(n/2)} \to 0, \tag{4.10}$$

where $\tilde{b}_n = b_n - a_n(\ln_3(n) + \ln(2))$ and Φ is the standard normal cdf. Compared to (4.8), the convergence in (4.10) is much faster. Since the subtrahend in (4.10) also converges uniformly to the Gumbel distribution in (4.8), this motivates the approximation

$$\mathrm{P}_0\left(T_n \leq x\right) \approx \left(\Phi\left(x - a_n(\ln_3(n) + \ln(2))\right)\right)^{2\ln(n/2)}. \tag{4.11}$$

As discussed in Yao and Davis (1986, Remark 2.3), (4.11) provides reasonable approximations even for small sample sizes $20 \leq n \leq 50$.

We next turn to the distribution of T_n under the alternative. Let $c_n = b_n - a_n \ln\left(\frac{\sqrt{\pi}}{2}\ln(\frac{1}{1-\varepsilon})\right)$, which is the $(1-\varepsilon)$-quantile of the distribution corresponding to the right-hand side in (4.9). Let $t_0 \in (0,1)$ and H_n a sequence of alternative hypotheses such that the change in mean between $\lfloor nt_0 \rfloor$ and $\lfloor nt_0 \rfloor + 1$ given by $\theta_n := |\mu_1^n - \mu_2^n|$, satisfies

$$\theta_n \sqrt{nt_0(1-t_0)} - \sqrt{2\ln_2(n)} \to \infty. \tag{4.12}$$

Then, denoting by P_n a probability measure such that H_n holds, it follows that

$$\mathrm{P}_n(T_n > c_n) \to 1 \tag{4.13}$$

as n tends to infinity. This is a direct consequence of Yao and Davis (1986, Theorem 3.1). Examples of sequences satisfying (4.12) are $\theta_n = \alpha n^{-\gamma}$ for $\gamma < 1/2, \alpha > 0$. In particular, the result holds for a constant change in mean ($\gamma = 0$). The convergence (4.13) is the key ingredient for the proof of Proposition 1 in Paper III.

4.2.2 Variance unknown

Including $\sigma^2 > 0$ as an unknown parameter in the generalised log likelihood ratio approach (cf. (4.4) - (4.7)) leads to the test statistic

$$U_n^2 := \frac{1}{\hat{\sigma}^2}\left(\sum_{i=1}^{n}(Y_i - \overline{Y}_{\leq n})^2 - \sum_{i=1}^{r^*}(Y_i - \overline{Y}_{\leq r^*})^2 - \sum_{i=r^*+1}^{n}(Y_i - \overline{Y}_{>r^*})^2\right) \geq 0, \tag{4.14}$$

where $\hat{\sigma}^2 = \frac{1}{n}\sum_{i=1}^{n}(Y_i - \overline{Y}_{\leq n})^2$. Since

$$\sqrt{U_n^2} =: U_n \stackrel{d}{=} T_n \frac{\sigma}{\hat{\sigma}},$$

and $\frac{\sigma}{\hat{\sigma}} \to 1$ P_0-almost surely, the results on the asymptotic distribution of T_n under H_0 and its approximations (cf. (4.8) - (4.11)) also hold for U_n, which follows from the Slutzky theorem.

Consider the sequence of alternative hypotheses H_n as defined above equation (4.12). To analyse the asymptotic type II error as n tends to infinity of the change point test statistic U_n, we need a convergence result as in (4.13), but for U_n instead of T_n. With the additional assumption that $\limsup_n \theta_n < \infty$, we show that indeed $\mathrm{P}_n(U_n > c_n) \to 1$ holds true. This is then used to prove Proposition 1 in Paper III.

4.3 Regularisation in a single split based on p-values

We will now explain the p-value-based regularisation approach employed in Paper III. Consider iid data as in Section 4.1 and a fully grown CART regression tree μ built from this data. The complexity m (number of leaves) of μ could for example be determined by specifying a maximal tree depth k. If k is chosen large, μ will in general overfit to the data. The idea to avoid this is to conduct a statistical test for signal in each internal node of the tree. In the following, we fix an internal node of μ, which corresponds to a subset $A \subset \mathbb{R}^d$ of the covariate space and we assume that $(X^{(i)}, Y^{(i)})_{i=1}^n$ are all the data points with $X_i \in A$. Let us further fix a covariate $j \in \{1, \dots, d\}$ and assume that the data $(X^{(i)}, Y^{(i)})_{i=1}^n$ is ordered by covariate j, that is, $X_j^{(1)} \leq \cdots \leq X_j^{(n)}$ and that the responses appear in the same order. Then, the idea is to test for a change point in the sequence $Y^{(1)}, \dots, Y^{(n)}$ of responses. Referring to (4.14) in Section 4.2, under a normality assumption, a test statistic based on a generalised log likelihood ratio is given by

$$U_{n,j}^2 := \frac{1}{\hat{\sigma}^2}\left(\sum_{i=1}^{n}(Y^{(i)} - \overline{Y}_{\le n})^2 - \sum_{i=1}^{r^*}(Y^{(i)} - \overline{Y}_{\le r^*})^2 - \sum_{i=r^*+1}^{n}(Y^{(i)} - \overline{Y}_{>r^*})^2\right),$$

where $\hat{\sigma}^2 = \frac{1}{n}\sum_{i=1}^{n}(Y^{(i)} - \overline{Y}_{\le n})^2$. Note that, in contrast to (4.14), we have added an index j to highlight the implicit dependence of $U_{n,j}^2$ on j via the ordering of the responses. As outlined in Section 4.2, the null distribution of $U_{n,j}^2$ can be approximated well via (4.11).

Since our aim is to test for signal in any covariate dimension $j = 1, \dots, d$, we extend the definitions of the null and the alternative hypotheses from (4.3). The null hypothesis H_0 will be defined as: there exists $\mu \in \mathbb{R}$, $\sigma^2 > 0$, such that $Y^{(1)}, \dots, Y^{(n)}$ follow $\mathrm{N}(\mu, \sigma^2)$. This means that there is no change point in mean in the sequence $Y^{(1)}, \dots, Y^{(n)}$ (for any ordering corresponding to some covariate j). In contrast, the alternative hypothesis H_1 says that there exist $\mu_1 \neq \mu_2, \sigma^2 > 0$, and a covariate dimension $j \in \{1, \dots, d\}$ such that if the responses $Y^{(1)}, \dots, Y^{(n)}$ are ordered by covariate j as described in the previous paragraph, there exists a change point $r \in \{1, \dots, n-1\}$, such that $Y^{(1)}, \dots, Y^{(r)}$ follow $\mathrm{N}(\mu_1, \sigma^2)$ and $Y^{(r+1)}, \dots, Y^{(n)}$ follow $\mathrm{N}(\mu_2, \sigma^2)$.

A statistic to test H_0 against H_1 is then given by

$$U_{n,\max}^2 := \max_{j=1,\dots,d} U_{n,j}^2,$$

where large values indicate significance, that is, a change point in mean in some covariate ordering. Recalling (4.2), the CART method with L^2 loss computes an optimal split point (j^*, ξ^*) which makes the computation of $U_{n,\max}^2$ immediate. Under the null, the upper tail probability of $U_{n,\max}^2$ can be bounded by

$$\begin{aligned}
\mathrm{P}_0(U_{n,\max}^2 > u) &= \mathrm{P}_0\left(\bigcup_{j=1}^{d}\{U_{n,j}^2 > u\}\right) && (4.15)\\
&\le \sum_{j=1}^{d}\mathrm{P}_0(U_{n,j}^2 > u) && (4.16)\\
&\approx\ d\left(1 - \Phi\left(u^{1/2} - \frac{\ln_3(n) + \ln(2)}{(2\ln_2(n))^{1/2}}\right)^{2\ln(n/2)}\right) && (4.17)\\
&=: dp_n(u). && (4.18)
\end{aligned}$$

In (4.17), the approximation (4.11) is employed, as well as the fact that under the null, $U_{n,1}^2, \dots, U_{n,d}^2$ are identically distributed. Formula (4.17) enables the simple computation of a p-value approximation $dp_n(U_{n,\max}^2)$, where, as mentioned above, $U_{n,\max}^2$ naturally appears as a byproduct in the optimisation procedure of the CART method with L^2 loss. Since under the null, (4.18) is an upper bound of the upper tail probability of $U^2_{n,\max}$, the test $\varphi = \mathbb{1}_{\{dp_n(U^2_{n,\max})<\varepsilon\}}$ provides a valid ε-level test in the sense that

$$\mathrm{E}_0(\varphi) \leq \varepsilon,$$

where $\varepsilon > 0$ is some fixed significance level. The convergence result (4.13) allows for an asymptotic analysis of the type II error of φ when n tends to infinity (cf. Paper III, Proposition 1).

Chapter 5

Controlled diffusions

Let $(\mathcal{F}_t)_{t\geq 0}$ be a filtration with $\mathcal{F}_\infty = \bigcup_{t\geq 0} \mathcal{F}_t \subset \mathcal{F}$. We assume that $(\mathcal{F}_t)$ satisfies the usual conditions, that is, $\mathcal{F}_0$ contains all P-null sets and for all $t \geq 0$, $\bigcap_{s>t} \mathcal{F}_s = \mathcal{F}_t$ (right continuity). Moreover, we assume that $(\mathcal{F}_t)$ supports a standard Brownian motion $(W_t)_{t\geq 0}$ which is adapted to it. In this chapter, we will employ several notions from stochastic calculus, amongst others stochastic integration in the Itô-sense, the Itô-formula, stochastic differential equations (SDEs), strong solutions, path-wise uniqueness and the infinitesimal generator of a diffusion. For a theoretical introduction to these concepts, we refer to Klenke (2013, chapters 25 and 26).

We will first introduce the classical stochastic control problem and then continue to the singular type which Paper IV deals with.

5.1 Classical stochastic control

This section is based on the outline in Touzi (2004). A process $(v_t)_{t\geq 0}$ is *progressively measurable*, if for each $t \geq 0$, the mapping $[0,t] \times \Omega \to \mathbb{R}$, $(s,\omega) \mapsto v_s(\omega)$ is $\mathcal{B}([0,t]) \otimes \mathcal{F}_t$-measurable. Progressively measurable processes are adapted. Conversely, adapted processes with càdlàg paths are progressively measurable. For a measurable $U \subset \mathbb{R}$, we denote by $\mathcal{A}_0$ the set of all progressively measurable processes with values in U and denote by $v \in \mathcal{A}_0$ a *control*.

Let $b\colon \mathbb{R} \times U \to \mathbb{R}$ and $\sigma\colon \mathbb{R} \times U \to \mathbb{R}$ satisfy a uniform Lipschitz condition, that is, there exists $K > 0$, such that

$$|b(x,u) - b(y,u)| + |\sigma(x,u) - \sigma(y,u)| \leq K|x-y| \tag{5.1}$$

for all $x, y \in \mathbb{R}$, $u \in U$.

For a given control $v \in \mathcal{A}_0$ and some initial condition $\xi \in L^2(\mathcal{F}_0)$, the stochastic differential equation

$$dX_t^v = b(X_t^v, v_t)dt + \sigma(X_t^v, v_t)dW_t, \quad X_0^v = \xi, \tag{5.2}$$

is called the *controlled* SDE. Note that (5.2) may be a path-dependent SDE if (v_t) is non-Markovian (cf. definition further down). Fix a time horizon $T > 0$. Let $\mathcal{A} \subset \mathcal{A}_0$, the set of *admissible controls*, be given by those $v \in \mathcal{A}_0$ satisfying

$$\mathrm{E}\left[\int_0^T \left(|b(x, v_t)| + |\sigma(x, v_t)|^2\right) dt\right] < \infty \tag{5.3}$$

for all $x \in \mathbb{R}$. Given some initial condition $\xi \in L^2(\mathcal{F}_0)$, (5.3) together with (5.1) ensure the existence of a path-wise unique strong solution X_t^v solving (5.2), denoted *controlled process*. Let $\mathrm{E}_{t,x}$ denote the expectation with respect to a probability measure $\mathrm{P}_{t,x}$ satisfying $\mathrm{P}_{t,x}(X_t = x) = 1$. We define the expected reward as

$$J(t, x; v) = \mathrm{E}_{t,x}\left[\int_t^T e^{-\int_t^s r(X_l^v, v_l)dl} f(X_s^v, v_s)ds + e^{-\int_t^T r(X_l^v, v_l)dl} g(X_T^v)\right]. \tag{5.4}$$

Here, $f\colon \mathbb{R} \times U \to \mathbb{R}$, $g\colon \mathbb{R} \to \mathbb{R}$ are measurable and satisfy the quadratic growth condition

$$|f(x, u)| + |g(x)| \leq C(1 + |x|^2) \tag{5.5}$$

for some constant $C > 0$ and all $u \in U$ and the *discount rate* $r\colon \mathbb{R} \times U \to \mathbb{R}$ is bounded from below. The process X^v in (5.4) is the path-wise unique strong solution to (5.2). The quadratic growth condition (5.5) and the fact that r is bounded from below ensure that $J(t, x; v)$ is well-defined for each admissible control $v \in \mathcal{A}$.

We remark that the framework can be extended to time-dependent functions b, σ, f and r. However, since the corresponding coefficients in the problem we consider in Paper IV have no time-dependence, and in order to simplify the presentation, we focus on the time-independent case. Moreover, note that even though the above mentioned functions have no time-dependence, the expected reward (5.4) does depend on t. The search for a maximiser

$$V(t, x) = \sup_{v \in \mathcal{A}} J(t, x; v) \tag{5.6}$$

for $(t, x) \in [0, T) \times \mathbb{R}$ is called a *stochastic control problem* and a function V fulfilling (5.6) is denoted the *value function.* A maximiser v^* in (5.6) is called an *optimal control.* One important class of controls are *Markovian controls* which can be written in the form $v_t = h(t, X_t^D)$ for some measurable $h\colon [0, T) \times \mathbb{R} \to U$. We remark that unless v is Markovian, the expected reward (5.4) may depend on information prior to time t. However, it can be shown that the value function (5.6) only depends on the present time and state (t, x).

5.1.1 The Hamilton-Jacobi-Bellman equation

The idea of constructing a candidate value function is to set up necessary conditions on its local behaviour. This leads to the *dynamic programming principle* which is one of the main tools of stochastic control. Let $(t,x) \in [0,T) \times \mathbb{R}$. Then, for any stopping time $t \le \tau \le T$, it can be shown that

$$V(t,x) = \sup_{v \in \mathcal{A}} \mathrm{E}_{t,x} \left[\int_t^\tau e^{-\int_t^s r(X_l^v, v_l)dl} f(X_s^v, v_s)ds + e^{-\int_t^\tau r(X_l^v, v_l)dl} V(\tau, X_\tau^v) \right], \tag{5.7}$$

where V is the value function defined in (5.6). In general terms, the dynamic programming principle states that the optimality of all subproblems is necessary for the optimality of the original problem. As another example, the problem of finding a shortest path in a graph follows the dynamic programming principle, since subpaths of shortest paths need to be shortest paths themselves. In the context of stochastic control, the subproblems are given by sub time intervals $[t,\tau] \subset [t,T]$, noting that $V(T,x) = g(x)$ in (5.4).

While the proof that the right-hand side in (5.7) dominates the value function only applies the tower property of conditional expectation as well as the multiplicative structure of the exponential discounting, the reverse inequality is more delicate.

In the following, we will motivate the *Hamilton-Jacobi-Bellman equation (HJB)* from the dynamic programming principle (5.7). The HJB equation characterises the local behaviour of the value function under the assumption that $V \in C^{1,2}([0,T), \mathbb{R})$. Let $(X_s^u)_{t \le s \le T}$ be the unique strong solution to (5.2) with initial condition $X_t^v = x \in \mathbb{R}$, where $v = u \in U$ is a constant control. In the following, the dynamics of the process

$$\left(e^{-\int_0^s r(X_l^u, u)dl} V(s, X_s^u) \right)_{0 \le s \le T}$$

play an important role. In order to analyse these dynamics, we apply the Itô formula with $F \in C^{1,1,2}([0,T), (0,\infty), \mathbb{R})$ given by $F(t,\beta,x) = \beta V(t,x)$ to the augmented process $(Y_s)_{t \le s \le T}$ given by

$$Y_s = \begin{bmatrix} s \\ e^{-\int_0^s r(X_l^u, u)dl} \\ X_s^u \end{bmatrix},$$

started at $Y_t = y = (t, \beta, x)$. Let $t \le \tau_n \le T$ a sequence of stopping times given by

$$\tau_n := \inf\{s \ge t \colon (s - t, X_s^u - x) \notin [0, 1/n) \times (-\alpha, \alpha)\}$$

for some constant $\alpha > 0$. Then,

$$
\begin{aligned}
&F(Y_{\tau_n}) - F(y) \\
&\quad = e^{-\int_0^{\tau_n} r(X_s^u,u)ds} V(\tau_n, X_{\tau_n}^u) - \beta V(t,x) \\
&\quad = \int_t^{\tau_n} e^{-\int_0^{s} r(X_l^u,u)dl} \Bigg(\frac{d}{ds} V(s,X_s^u) - r(X_s^u,u)V(s,X_s^u) \\
&\qquad\qquad + b(X_s^u,u)\frac{d}{dx}V(s,X_s^u) + \frac{1}{2}\sigma^2(X_s^u,u)\frac{d^2}{dx^2}V(s,X_s^u)\Bigg) ds \\
&\qquad + \int_t^{\tau_n} e^{-\int_0^{s} r(X_l^u,u)dl} \frac{d}{dx}V(s,X_s^u)\sigma(X_s^u,u)dW_s.
\end{aligned}
\tag{5.8}
$$

Define the differential operator

$$
(L^u h)(t,x) := -r(x,u)h(t,x) + b(x,u)\frac{d}{dx}h(t,x) + \frac{1}{2}\sigma^2(x,u)\frac{d^2}{dx^2}h(t,x)
$$

for $h \in C^{1,2}$. The last integral in (5.8) is a martingale with mean zero. Therefore, after rearranging, canceling β, taking expectations and substituting L^u, we arrive at

$$
\begin{aligned}
V(t,x) = \mathrm{E}_{t,x}\Bigg[&\int_t^{\tau_n} e^{-\int_t^{s} r(X_l^u,u)dl} \Big(-\frac{d}{ds}V(s,X_s^u) - (L^u V)(s,X_s^u)\Big) ds \\
&+ e^{-\int_t^{\tau_n} r(X_s^u,u)ds} V(\tau_n, X_{\tau_n}^u)\Bigg].
\end{aligned}
\tag{5.9}
$$

Comparing (5.9) to the dynamic programming principle (5.7), this motivates that the value function should satisfy

$$
-\frac{d}{dt}V(t,x) - (L^u V)(t,x) \geq f(x,u). \tag{5.10}
$$

This can be shown precisely by letting n tend to infinity and arguing that a dominated convergence theorem applies. Define the function $H\colon \mathbb{R}^4 \to \overline{\mathbb{R}}$ by

$$
H(x,y,z,w) = \sup_{u\in U} \left(f(x,u) - r(x,u)y + b(x,u)z + \frac{1}{2}\sigma^2(x,u)w \right).
$$

If $H \in C(\mathbb{R}^4)$, $r(\cdot,u), f(\cdot,u) \in C(\mathbb{R})$ for each $u \in U$ as well as $\|r\|_\infty < \infty$, it can be shown that V solves the *Hamilton-Jacobi-Bellman equation (HJB)*, i.e.

$$
\begin{aligned}
&\frac{d}{dt}V(t,x) + H\left(x, V(t,x), \frac{d}{dx}V(t,x), \frac{d^2}{dx^2}V(t,x)\right) && (5.11)\\
=&\frac{d}{dt}V(t,x) + \sup_{u\in U}\left(f(x,u) + (L^u V)(t,x)\right) && (5.12)\\
=&\,0
\end{aligned}
$$

for all $(t,x) \in [0,T) \times \mathbb{R}$ (cf. Touzi (2004, Theorem 1.3)). In particular, a maximiser $u^* \in U$ of (5.12) makes (5.10) an equality. In other words, the HJB equation states that, if v^* is an optimal control, then the value function satisfies that the process

$$\left(e^{-\int_0^t r(X_s^{v^*}, v^*)ds} V(t, X_t^{v^*}) - \int_0^t e^{-\int_0^s r(X_l^{v^*}, v^*)dl} f(X_s^{v^*}, v^*) ds \right)_{0 \le t \le T}$$

is a local martingale.

The crucial assumption of this section that $V \in C^{1,2}$ is a-priori usually unknown. The general strategy to solve a stochastic control problem is therefore to derive a candidate value function using the Hamilton-Jacobi-Bellman equation and, in a second step, to prove that the given candidate indeed solves (5.6). This second step is referred to as proving a *verification theorem*. For more details on the verification argument, we refer to Touzi (2004, Section 1.4).

5.1.2 Example: The dividend problem with restricted dividend rate

Let X_t^v model the reserve of an insurance company at time t, which is governed by the controlled SDE

$$dX_t^v = (\mu - v_t)dt + \sigma dW_t, \quad X_0^v = x > 0, \tag{5.13}$$

where $\mu, \sigma > 0$, $v_t \in [0, a_0]$ for some $a_0 > 0$ is the rate of dividend payment at time t. We define the reward function as the total discounted dividend payment, i.e.

$$J(x; v) = \mathrm{E}_x \left[\int_0^\tau e^{-ct} v_t dt \right], \tag{5.14}$$

where $\tau = \inf\{t \ge 0 \colon X_t^v = 0\}$ is the first hitting time of zero, $c > 0$ is the discount rate and let the value function be given by

$$V(x) = \sup_{v \in \mathcal{A}} J(x; v).$$

Here, the set of admissible control processes $\mathcal{A}$ is given by all $(\mathcal{F}_t)$-adapted càdlàg-processes $(v_t)_{t \ge 0}$ taking values in $[0, a_0]$. Note that in contrast to the presentation in Section 5.1.1 where a fixed time horizon T was given, here, neither the reward nor the value function depend on time. This is due to the time-independent horizon τ. The presented model was analysed and solved in Asmussen and Taksar (1997). Comparing to (5.11), the Hamilton-Jacobi-Bellman equation of the problem is

$$\sup_{0 \le u \le a_0} \left(\frac{\sigma^2}{2} V''(x) + (\mu - u)V'(x) - cV(x) + u \right) = 0, \tag{5.15}$$

$$V(0) = 0, \tag{5.16}$$

where the boundary condition in the second line is a direct consequence of the formulation of the reward in (5.14). It can be seen that the maximiser u^* of (5.15) is given by

$$u^*(x) = \begin{cases} 0, & V'(x) > 1 \\ a_0, & V'(x) \leq 1. \end{cases} \tag{5.17}$$

This leads to the two linear ordinary differential equations (ODEs)

$$\frac{\sigma^2}{2} V''(x) + \mu V'(x) - cV(x) = 0, \quad V'(x) > 1, \tag{5.18}$$

$$\frac{\sigma^2}{2} V''(x) + (\mu - a_0) V'(x) - cV(x) + a_0 = 0, \quad V'(x) \leq 1. \tag{5.19}$$

Using an exponential ansatz $V(x) = e^{\gamma x}$, the linear space of solutions to (5.18) is given by

$$\{C_1 e^{\gamma_1(\mu)x} + C_2 e^{\gamma_2(\mu)x} \colon C_1, C_2 \in \mathbb{R}\}, \tag{5.20}$$

where $\gamma_1(\mu) > 0 > \gamma_2(\mu)$ are the two solutions to $\frac{\sigma^2}{2}\gamma^2 + \mu\gamma - c = 0$. Meanwhile, since $\frac{a_0}{c}$ is a particular solution to the non-homogeneous, linear ODE (5.19), it follows that

$$\left\{ \frac{a_0}{c} + C_3 e^{\gamma_1(\mu - a_0)x} + C_4 e^{\gamma_2(\mu - a_0)x} \colon C_3, C_4 \in \mathbb{R} \right\}$$

is the affine-linear solution space to (5.19), where $\gamma_1(\mu - a_0) > 0 > \gamma_2(\mu - a_0)$ are the two solutions to $\frac{\sigma^2}{2}\gamma^2 + (\mu - a_0)\gamma - c = 0$. The constants $C_1, \ldots, C_4$ are determined by the initial condition $V(0) = 0$ and by demanding that the solution is twice continuously differentiable. Altogether, it can be shown that there exists a unique function $h \in C^2(\mathbb{R}_{\geq 0})$ solving (5.15) – (5.16). When $\frac{a_0}{c} + \frac{1}{\gamma_2(\mu - a_0)} \leq 0$,

$$h(x) = \frac{a_0}{c}(1 - e^{\gamma_2(\mu - a_0)x}).$$

When $\frac{a_0}{c} + \frac{1}{\gamma_2(\mu - a_0)} > 0$,

$$h(x) = \begin{cases} C(e^{\gamma_1(\mu)x} - e^{\gamma_2(\mu)}x), & 0 \leq x \leq b^* \\ \frac{a_0}{c} - de^{\gamma_2(\mu - a_0)x}, & b^* < x, \end{cases}$$

where $C, d, b^* \in \mathbb{R}$ are uniquely determined. Having found a solution h to the Hamilton-Jacobi-Bellman equation, it remains to show that indeed $h = V$. Here, since $h \in C^2(\mathbb{R}_{\geq 0})$, this can be achieved by means of Itô's formula. The optimal control can be shown to be given by $v_t^* = a_0$ when $\frac{a_0}{c} + \frac{1}{\gamma_2(\mu - a_0)} \leq 0$, while $v_t^* = a_0 \mathbb{1}_{\{X_t^{v^*} > b^*\}}$ when $\frac{a_0}{c} + \frac{1}{\gamma_2(\mu - a_0)} > 0$. Moreover, $J(x; v^*) = h(x) = V(x)$.

We would like to point out that the condition (5.17) implies that, whenever $V'(x) \leq 1$, it is optimal to pay out as much dividends as possible. This gives some insight into the problem solution in the case where $\mathcal{A}$ is not restricted to bounded controls. In Section 5.2.3, we will continue the example discussing an adaptation to non-restricted controls.

5.2 Singular stochastic control

In the example in Section 5.1.2, the cumulative displacement of the controlled process (5.13) caused by the control is given by the process $D_t := \int_0^t v_s ds$, where (v_s) is càdlàg. As a consequence, $t \mapsto D_t$ is absolutely continuous. We now consider a class of stochastic control problems where this absolute continuity fails to apply. We can hence not express the cumulative displacement as an integral of a càdlàg process. The word "singular" originates from relaxing the set of admissible controls to allow the process (D_t) to be singular, that is, $t \mapsto D_t$ is continuous, increasing, and has a derivative of zero for Lebesgue-almost all t. The standard example of a singular function is the Cantor function. However, one can even further relax the set of admissible controls to allow for discontinuities of (D_t).

We will now introduce some of the main concepts of singular stochastic control and then elaborate on the example in Section 5.1.2.

5.2.1 Reflected stochastic differential equations

We first introduce the purely non-probabilitstic Skorokhod problem following Pilipenko (2014). Given $T > 0$, $b^* > 0$, $\varphi \in C([0,T])$ with $\varphi(0) \leq b^*$, a pair of functions $\psi, l \in C([0,T])$ is called a *solution to the Skorokhod problem for* φ, if

$$\begin{aligned} &\text{(i) } \psi(t) \leq b^* \text{ for all } t \in [0,T], \\ &\text{(ii) } \psi(t) = \varphi(t) - l(t) \text{ for all } t \in [0,T], \\ &\text{(iii) } l(0) = 0 \text{ and } l \text{ is increasing}, \\ &\text{(iv) } \int_0^T \mathbb{1}_{\{\psi(s) < b^*\}}(s) dl(s) = 0. \end{aligned} \tag{5.21}$$

This problem formalises a downward reflection at b^*. The function ψ is the reflection of φ while l, informally speaking, is the minimal downwards push needed to ensure that $\psi \leq b^*$. The Skorokhod problem's unique solution is given by

$$l(t) = \max_{s \in [0,t]} \left(\max(\varphi(s) - b^*, 0) \right), \quad \psi(t) = \varphi(t) - l(t).$$

As an example, in Figure 5.1, we plot a non-differentiable but continuous function φ as well as the unique pair (ψ, l) solving the Skorokhod problem for φ and with $b^* = 1$. It can be seen that the continuous function l only increases when ψ hits the barrier b^*, corresponding to property (iv) in (5.21).

Define the Skorokhod mapping

$$\Gamma : C([0,T]) \to C([0,T]), \quad \varphi \mapsto \varphi - \max_{s \in [0,\cdot]} \left(\max(\varphi(s) - b^*, 0) \right).$$

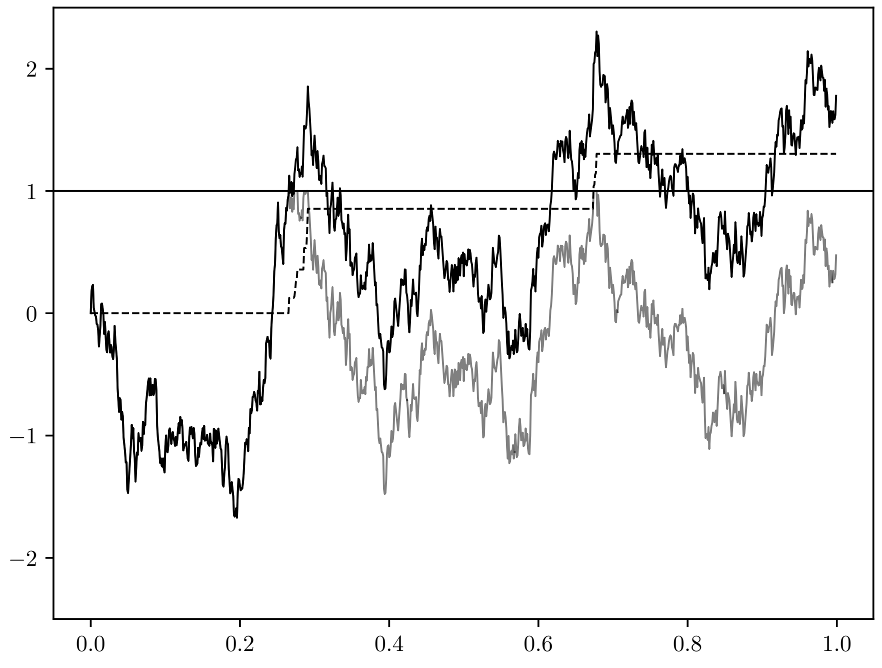

Figure 5.1: Function $\varphi \in C([0,1])$ with $\varphi(0) = 0$ (black) and unique solution to the Skorokhod problem (ψ, l) (grey, dashed) with reflection at $b^* = 1$.

Now, consider the stochastic differential equation

$$dX_t = b(X_t)dt + \sigma(X_t)dW_t, \quad X_0 = x_0 \leq b^*, \tag{5.22}$$

whose coefficients are assumed to satisfy a Lipschitz condition to ensure existence and uniqueness of a strong solution. The aim is to construct a process X_t^D which behaves like (5.22) when $X_t^D < b^*$ and continuously reflects downwards when $X_t^D = b^*$. Inspired by (5.21), we call a pair of continuous $(\mathcal{F}_t)$-adapted processes (X_t^D, D_t) a *solution to the stochastic differential equation*

$$dX_t^D = b(X_t^D)dt + \sigma(X_t^D)dW_t - dD_t \tag{5.23}$$

with downward reflection at b^* *and initial condition* $X_0^D = x_0 \leq b^*$, if, almost surely,

$$\begin{aligned}
&\text{(i) } X_t^D \leq b^* \text{ for all } t \in [0,T],\\
&\text{(ii) } X_t^D = x_0 + \int_0^t b(X_s^D)ds + \int_0^t \sigma(X_s^D)dW_s - D_t,\\
&\text{(iii) } D_0 = 0 \text{ and } D \text{ is increasing},\\
&\text{(iv) } \int_0^T \mathbb{1}_{\{X_s^D < b^*\}}(s)dD_s = 0
\end{aligned} \tag{5.24}$$

and all integrals are well-defined.

In order to make the connection between (5.21) and (5.24), assume that we are given a solution (X_t^D, D_t) to (5.23) with downward reflection at b^*. Let

$$Y_t = x_0 + \int_0^t b(X_s^D)ds + \int_0^t \sigma(X_s^D)dW_s$$

and pick $\omega \in \Omega$ such that all conditions in (5.24) are satisfied. Then, the conditions in (5.24) coincide with those in (5.21) where $\varphi = Y(\omega)$. As a consequence, the pair $(X^D(\omega), D(\omega))$ is the unique solution to the Skorokhod problem for $Y(\omega)$ and it holds that

$$X^D(\omega) = \Gamma Y(\omega).$$

Hence, Y_t solves the path-dependent SDE

$$Y_t = x_0 + \int_0^t b((\Gamma Y)_s)ds + \int_0^t \sigma((\Gamma Y)_s)dW_s. \tag{5.25}$$

It should be remarked here that if Y is continuous and $(\mathcal{F}_t)$ adapted, then the same holds for ΓY by definition of the Skorokhod map. Given a solution Y_t to (5.25), $X^D := \Gamma Y$ and $D = Y - X^D$ yield a solution to the reflected SDE. As a consequence, finding a solution to the reflected SDE can be reduced to solving (5.25). Given a Lipschitz- and a linear growth condition on the coefficients, i.e. given there exist constants $L > 0, C > 0$, such that for all $t \geq 0$, $x_1, x_2 \in \mathbb{R}$,

$$|b(x_1) - b(x_2)| + |\sigma(x_1) - \sigma(x_2)| \leq L|x_1 - x_2| \tag{5.26}$$
$$|b(x_1)| + |\sigma(x_1)| \leq C(1 + |x_1|), \tag{5.27}$$

it can be shown that (5.25) has a unique strong solution (cf. Liptser and Shiryaev (2013, Theorem 4.6)), such that there exists a unique solution (X^D, D) to the reflected SDE.

5.2.2 Relation to local time

For a continuous semimartingale (X_t), *the (symmetric) local time of* (X_t) *at* $a \in \mathbb{R}$ is given by the limit

$$L_t^a(X) = \lim_{\varepsilon \downarrow 0} \frac{1}{2\varepsilon} \int_0^t \mathbb{1}_{[a-\varepsilon, a+\varepsilon]}(X_s)d\langle X\rangle_s. \tag{5.28}$$

Here, the process $(\langle X\rangle_t)$ is the quadratic variation of (X_t). The limit in (5.28) exists almost surely. An alternative representation is given by the Tanaka formula (cf. Revuz and Yor (1999, Chapter 6))

$$L_t^a(X) = |X_t - a| - |X_0 - a| - \int_0^t \operatorname{sgn}^*(X_s - a)dX_s,$$

where sgn^* is the symmetric sign function given by

$$\mathrm{sgn}^*(x) = \begin{cases} 1, & x > 0 \\ 0, & x = 0 \\ -1, & x < 0. \end{cases}$$

Let (X^D, D) be the unique solution to the reflected SDE derived in Section 5.2.1. Then, applying Tanaka's representation to the process X_t^D, as well as the properties (5.24) of (X^D, D), one can show that

$$L_t^{b^*}(X^D) = D_t. \tag{5.29}$$

5.2.3 Example: The dividend problem with non-restricted dividend rate

We continue with the example in 5.1.2, with the adaptation of, informally, relaxing the set of admissible controls to allow v_t to take values in $[0, \infty]$. A control $(D_t) \in \mathcal{A}$ will now model the cumulative dividend payments. The controlled process is governed by the SDE

$$dX_t^D = \mu dt + \sigma dW_t - dD_t \tag{5.30}$$

$$X_{0-}^D = x > 0, \tag{5.31}$$

and the set of admissible controls given by

$$\mathcal{A} := \{(D_t)_{t\geq 0} \text{ is increasing, right-continuous, } (\mathcal{F}_t)\text{-adapted,} \\ \text{with } D_{0-} = 0, \text{ such that } D_t - D_{t-} \leq X_{t-}^D, \text{ for all } \; t \geq 0\}.$$

In contrast to the restricted case, the processes D_t and X_t^D are in general neither absolutely continuous nor continuous. The dividend payment at time t is given by $D_t - D_{t-}$. Since $X_0^D = x - D_0$, in case of $D_0 > 0$, there is an instantaneous dividend payout and X_0^D decreases from x to $x - D_0$. Hence, E_x now refers to a probability measure satisfying $\mathrm{P}_x(X_{0-}^D = x) = 1$.

The reward function will be adapted to the cumulative notation as

$$J(x; D) = \mathrm{E}_x\left[\int_0^\tau e^{-ct} dD_t\right],$$

where the integral is of ordinary Lebesgue-Stieltjes type (D_t is locally of bounded variation) and τ is the first hitting time of zero, as before.

We recall equation (5.17) from Section 5.1.2, which says that, depending on the first derivative of the value function, the optimal control consists of paying out nothing or the maximal amount a_0. This motivates the idea that in the unrestricted case, the optimal dividend rate $u^*(x)$ should be zero or infinite. Indeed, it can be shown that, if the value function is twice continuously differentiable, it satisfies

$$\frac{\sigma^2}{2}V''(x) + \mu V'(x) - cV(x) \leq 0 \tag{5.32}$$

$$V'(x) \geq 1. \tag{5.33}$$

Note that (5.32) coincides with the corresponding no-payout equation in the bounded case (cf. (5.15), (5.17)). The proof idea for the equation $V'(x) \geq 1$ is to show that the strategy of an instantaneous payout of $\delta > 0$ units will reduce the value by δ units. Hence, any optimal strategy should be better and thus satisfy $V(x) - V(x-\delta) \geq \delta$ (value increase larger than invested capital). As a matter of fact, it can be shown that one of the equations (5.32), (5.33), is always tight, so that the following Hamilton-Jacobi-Bellman equation holds

$$\begin{aligned} \max\left(\frac{\sigma^2}{2}V''(x) + \mu V'(x) - cV(x), 1 - V'(x)\right) &= 0, \\ V(0) &= 0. \end{aligned} \tag{5.34}$$

A solution $h \in C^2(\mathbb{R})$ to (5.34) is given by

$$h(x) = \begin{cases} C\left(e^{\gamma_1(\mu)x} - e^{\gamma_2(\mu)x}\right), & x \leq b^* \\ C\left(e^{\gamma_1(\mu)b^*} - e^{\gamma_2(\mu)b^*}\right) + x - b^*, & x > b^*, \end{cases} \tag{5.35}$$

where

$$C = \frac{1}{\gamma_1(\mu)e^{\gamma_1(\mu)b^*} - \gamma_2(\mu)e^{\gamma_2(\mu)b^*}}, \quad b^* = \frac{\log\left(\frac{\gamma_2^2(\mu)}{\gamma_1^2(\mu)}\right)}{\gamma_1(\mu) - \gamma_2(\mu)}. \tag{5.36}$$

Here, $\gamma_1(\mu) > 0 > \gamma_2(\mu)$ are defined as in (5.20) and below, and the constants b^*, C have been determined by the initial condition $h(0) = 0$ as well as by ensuring that $h \in C^2(\mathbb{R})$.

It remains to prove a verification theorem stating that the constructed candidate h indeed coincides with the value function. This is achieved in two steps. First, it is shown that our candidate dominates the value function, i.e. $h \geq V$. Since $h \in C^2(\mathbb{R})$, this can be done by applying a generalised Itô formula for semimartingales (possibly with jumps) (cf. Peskir (2007, Theorem 3.2)) to the candidate h.

Second, one needs to prove that h is attained by some control. For this purpose, we recall from (5.34) and (5.35) that our candidate h satisfies

$$\frac{\sigma^2}{2}h''(x) + \mu h'(x) - ch(x) = 0, \quad 0 \leq x \leq b^*, \tag{5.37}$$

$$h'(x) = 1, \quad x > b^*. \tag{5.38}$$

From the bounded case with $a_0 > 0$ large, in particular from (5.17), we recall that the optimal control partitions the state space $[0, \infty)$ into a waiting region (no payout) $\mathcal{W} = \{x > 0\colon V'(x) > 1\}$ and an action region $\mathcal{D} = \{x > 0\colon V'(x) \leq 1\}$ where the maximal dividend rate a_0 is payed out. In the present case, in particular by (5.37)–(5.38), this motivates to set up an optimal control candidate which is defined by a waiting region $\mathcal{W} = (0, b^*)$ and an action region $\mathcal{D} = [b^*, \infty)$. When the controlled process is started in $x \in \mathcal{W}$, it behaves like an uncontrolled Brownian motion with drift according to (5.30) with $dD_t = 0$. Upon hitting b^*, a minimal downward control is exerted to keep the process within the closure of the waiting region $\overline{\mathcal{W}}$. On the other hand, when the process is started in the action region $x \in \mathcal{D}$, there is an instantaneous dividend payout $D_0 - D_{0-} = x - b^*$ which pushes the process downward into the closure of the waiting region.

When $x \in \overline{\mathcal{W}}$, the above described dynamics correspond to an SDE with downward reflection at b^* as defined in (5.23). Since the constant coefficients $\mu, \sigma > 0$ satisfy the Lipschitz- and linear growth conditions (5.26) – (5.27), we know that there exists a path-wise unique solution $(X_t^{D^{b^*}}, D^{b^*})$ to the SDE reflected at b^* which satisfies

$$\begin{aligned}
&\text{(i) } X_t^{D^{b*}} \leq b^*, \quad t \geq 0, \\
&\text{(ii) } X_t^{D^{b^*}} = x + \int_0^t \mu dt + \int_0^t \sigma dW_t - D_t^{b^*}, \\
&\text{(iii) } D_0^{b^*} = 0 \text{ and } D_t^{b^*} \text{ is increasing,} \\
&\text{(iv) } \int_0^t \mathbb{1}_{\{X_s^{D^{b*}} < b^*\}}(s) dD_s^{b^*} = 0, \quad t \geq 0
\end{aligned} \tag{5.39}$$

and moreover, by (5.29), $D_t^{b*} = L_t^{b^*}(X^{D^{b*}})$, the local time of $X^{D^{b*}}$ at b^*. Incorporating the dynamics if the process is started in the interior of the action region, we arrive at our candidate optimal control

$$D_t^{b^*} = (x - b^*)^+ + L_t^{b^*}(X^{D^{b*}}), \quad t \geq 0, \quad D_{0-}^{b^*} = 0.$$

Note that $D^{b*} \in \mathcal{A}$ by (5.39), (iii), and since $(\mathcal{F}_t)$-adapted by the Skorokhod construction. Finally, using (5.39) and (5.37) – (5.38), the Itô formula applied to h can be used to show that $h(x) = \mathrm{E}_x\left[\int_0^\tau e^{-ct} dD_t^{b^*}\right] = J(x; D^*)$. This finishes the verification and as a consequence, $h = V$.

Chapter 6

Overview of papers

6.1 Paper I

We suppose to be given insurance data of cumulative total claims amounts in the form of a so-called run-off triangle. For a fixed time horizon $T \in \mathbb{N}$ and $i, t \in \mathcal{T} = \{1, \dots, T\}$, let $C_{i,t}$ denote the sum of all claim payments from events in accident year i paid up to and including development year t. We assume to have an observation of $C_{i,t}$, if $i + t \leq T + 1$ and define the sigma algebra of the past observations accordingly by $\mathcal{D} = \sigma\left(\{C_{i,t} \colon i + t \leq T + 1\}\right)$. The aim of the chain ladder method is to provide an estimator of the ultimate claims amount $C_{i,T}$ and of the outstanding claims reserve $R_i = C_{i,T} - C_{i,T-i+1}$ for accident years $i = 2, \dots, T$.

The core idea is to estimate so-called development factors, $\widehat{f}_t$, $t = 1 \dots, T-1$, by (2.19), and then to calculate predictors of $C_{i,T}$ and R_i by

$$\widehat{C}_{i,T} = C_{i,T-i+1} \widehat{f}_{T-i+1} \cdot \ldots \cdot \widehat{f}_{T-1}, \quad \widehat{R}_i = C_{i,T-i+1} \left(\widehat{f}_{T-i+1} \cdot \ldots \cdot \widehat{f}_{T-1} - 1 \right). \tag{6.1}$$

Since its development, the question has been investigated which model assumptions on $(C_{i,t})_{i,t\in\mathcal{T}}$ underly the chain ladder method. In an answer to this question, in his pioneering contribution, Mack (1993) states three conditions, known as Mack's distribution-free chain ladder (MCL), given by (2.21). Mack shows that under (MCL), conditionally on $\mathcal{D}$, (6.1) provide unbiased predictors and derives an estimator of the conditional mean squared error of prediction $\mathrm{E}\left[(C_{i,T} - \widehat{C}_{i,T})^2 \,\middle|\, \mathcal{D}\right]$, as introduced in (2.22). However, it turns out that most models for cumulative total claims amounts fail to satisfy (MCL).

In Paper I, we consider two models, a general and a special one. For ease of exposition of this overview, we only present the latter, denoting it by (SM). For $i, t \in \mathcal{T}$, assume that the incremental total claims amount $X^{\alpha}_{i,1} = C^{\alpha}_{i,1}, X^{\alpha}_{i,t} = C^{\alpha}_{i,t} - C^{\alpha}_{i,t-1}$ for $t \geq 2$, is given by a compound Poisson model (Cramér-

Lundberg) model,

$$X_{i,t}^{\alpha} = \sum_{k=1}^{N_{i,t}^{\alpha}} Z_{i,t,k}, \tag{6.2}$$

where $(N_{i,t}^{\alpha})_{\alpha \geq 0}$ is a homogeneous Poisson process with intensity $\lambda_i q_t$, $\sum_{t=1}^{T} q_t = 1$, $q_t \in (0,1)$ for all t, and $\lambda_1, \ldots, \lambda_T > 0$. Here, $(N_{i,t}^{\alpha})$ is independent of $(Z_{i,t,k})_k$ and all appearing random variables on the right hand side of (6.2) are independent of those from a different cell $(\widetilde{i}, \widetilde{t}) \neq (i,t)$. In particular, there is no dependence between incremental claims amounts between development years. This model is a special case of the total claims amount model introduced in (2.3) in Section 2.1.

(SM) does not satisfy (MCL). However, we show that the chain ladder predictors in (6.1) can be derived for (SM) under large-exposure asymptotics, i.e. as $\alpha \to \infty$. Moreover, we derive Mack's estimator for the conditional mean squared error of prediction (2.22) for (SM) as α tends to infinity. Our results thus justify the use of (6.1) and (2.22) in a standard model for the total claims amount under large exposure, even though the conditions of (MCL) are not fulfilled.

The most challenging part of Paper I was to establish the weak convergence of $\widehat{\sigma}_t^2$ in (2.23) to a scaled chi-squared distribution as $\alpha \to \infty$ (cf. (4.3) in (Paper I, Theorem 1)). To this end, it was necessary to determine the eigenvalues of the matrix B, including their geometric multiplicities, in the proof of Theorem 1. This required conjecturing the eigenvectors and subsequently verifying their correctness.

For model classes more general than (SM), it appears unlikely that Mack's estimator can be derived under large exposure, since the derivation relies on the special property of Poisson counting variables that expectation and variance coincide.

Our results do not show that the distributions of the random variables

$$\mathrm{E}\left[(C_{i,T}^{\alpha} - \widehat{C}_{i,T}^{\alpha})^2 \,\middle|\, \mathcal{D}\right]$$

and (2.22) are similar under (SM) (cf. (Paper I, Figure 1) for an illustration). Rather, they show that, for large α, the natural estimator of $\mathrm{E}\left[(C_{i,T}^{\alpha} - \widehat{C}_{i,T}^{\alpha})^2 \,\middle|\, \mathcal{D}\right]$ coincides with Mack's estimator (2.22). Moreover, the expectations of the respective limits of $\mathrm{E}\left[(C_{i,T}^{\alpha} - \widehat{C}_{i,T}^{\alpha})^2 \,\middle|\, \mathcal{D}\right]$ and (2.22), standardised by $(C_{T-i+1}^{\alpha})^{-1}$, coincide (asymptotic unbiasedness).

6.2 Paper II

This paper is motivated by computational challenges in the valuation of multi-period liability cashflows with capital requirements. We consider the set-up as presented in Sections 2.3 and 2.4 with a slight generalisation of a potentially larger filtration $(\mathcal{F}_t)$. Let $T \in \mathbb{N}$ a finite time horizon, $X, X^1, X^2, \dots$ be random vectors taking values in $\mathbb{R}^T$ and $Y, Y^1, Y^2, \dots$ be random vectors in $(\mathbb{R}^d)^T$, where $d \in \mathbb{N}$. The vector $X^n = (X_1^n, \dots, X_T^n)$ denotes a discounted incremental cashflow (centred and scaled) in a model with T periods and $Y^n = (Y_1^n, \dots, Y_T^n)$ is a random vector providing additional information (cf. (Paper II, Section 4) for examples of processes (X^n, Y^n)). We make the assumption that a central limit theorem applies to (X^n, Y^n), such that $\mathcal{L}(X^n, Y^n) \xrightarrow{w} \mathcal{L}(X, Y)$, where the latter is the law of a multivariate Gaussian vector. This holds, for example, when X^n is given by a multi-period model for the total claim amount (cf. (2.4) in Section 2.1, with n playing the role of α as an index of exposure). For example, n may be proportional to the number of policyholders. We let $\mathcal{F}_t^n = \sigma(X_{\le t}^n, Y_{\le t}^n), \mathcal{F}_t = \sigma(X_{\le t}, Y_{\le t})$, $\mathcal{F}_0 = \mathcal{F}_0^n = \{\emptyset, \Omega\}$.

Consider the values at time zero of the liability cashflows $V_0^n(X^n)$, $V_0(X)$ defined in terms of a backward recursion and depending on the filtrations $(\mathcal{F}_t^n), (\mathcal{F}_t)$ as in (2.12) in Section 2.3. Note that we consider a slightly more general set-up here, since $\mathcal{F}_t^n, \mathcal{F}_t$ may hold more information than contained the cashflows $X_{\le t}^n, X_{\le t}$, respectively.

The exact calculation of the values V_0^n for a general random vector (X^n, Y^n) is in general not possible arithmetically. Also Monte Carlo approaches are computationally challenging, in particular when the time horizon T is large. However, when (X, Y) is a multivariate Gaussian vector, the value $V_0(X)$ can be calculated explicitly via the formula (2.17) in Section 2.4. This brings up the question whether convergence of the real sequence

$$V_0^n(X^n) \to V_0(X) \tag{6.3}$$

applies. The ultimate goal is an approximation of the value at time zero of the original non-centred and non-scaled liability cashflow $C^n = a_n X^n + b_n$, where $a_n > 0, b_n \in \mathbb{R}^T$ are non-random sequences (for more details and an example consider Section 2.4 and (2.16) therein). Employing (2.15), the convergence (6.3) would justify to approximate

$$V_0^n(C^n) = V_0^n(a_n X^n + b_n) = a_n V_0^n(X^n) + \sum_{s=1}^{T} b_{n,s} \tag{6.4}$$

$$\approx a_n V_0(X) + \sum_{s=1}^{T} b_{n,s}. \tag{6.5}$$

The main contribution of Paper II is to set up conditions under which (6.3) holds. It turns out that only the assumption of weak convergence $\mathcal{L}(X^n, Y^n) \xrightarrow{w} \mathcal{L}(X, Y)$ is not sufficient for (6.3). Suppose additionally that for each t and each convergent sequence $(x^n, y^n)_{\leq t} \to (x, y)_{\leq t}$,

$$\mathcal{L}\Big((X^n, Y^n) \mid (X^n, Y^n)_{\leq t} = (x^n, y^n)_{\leq t}\Big) \xrightarrow{w} \mathcal{L}\Big((X, Y) \mid (X, Y)_{\leq t} = (x, y)_{\leq t}\Big), \tag{6.6}$$

$$\Big(\mathcal{L}\big(\|X^n\| \mid (X^n, Y^n)_{\leq t} = (x^n, y^n)_{\leq t}\big)\Big)_{n\in\mathbb{N}} \quad \text{is uniformly integrable.} \tag{6.7}$$

Then, (6.3) holds and (6.5) provides a simple approximation formula for the value of the liability cashflow. In the example of a negative multinomial cashflow model, we show that both (6.6) and (6.7) are satisfied. Moreover, we compute exact pre-limit values $V_0^n(X^n)$ in $T = 2$ periods and demonstrate their convergence to $V_0(X)$ as n tends to infinity.

The most technically demanding aspect of Paper II was the proof of Theorem 2, as the induction step involves several delicate arguments concerning weak convergence, leading to rather technical expressions. In addition, recognising that the domain of the convex optimisation problem in the proof of Theorem 4 can be characterised as a convex cone was a key structural observation.

A beneficial extension of the results would be to establish that the convergence of (X^n, Y^n) to (X, Y) in the adapted Wasserstein-1-distance already suffices for the convergence of valuations. Such a result would bring the theoretical framework closer to the existing literature, where adapted Wasserstein distances are more prevalent than continuous convergence of conditional distributions.

Finally, although we prove convergence $V_0^n(X^n) \to V_0(X)$, no convergence rates are derived. Consequently, the accuracy of the approximation (6.5) cannot be guaranteed from a theoretical standpoint. Nevertheless, the numerical example presented in (Paper II, Section 6) suggests that the approximation performs satisfactorily in practice.

6.3 Paper III

This paper deals with the regularisation of CART regression trees. The standard regularisation technique determining the complexity of CART trees is based on cross-validation in order to obtain predictors that generalise well to unseen data. Due to the random split into folds, this implies that the resulting predictor is not a deterministic function of the data. Moreover, cross-validation is computationally expensive and may result in an inefficient use of data. The main contribution of Paper III is to propose an in-sample method to stop the growing of a CART regression tree based on node-wise statistical tests. These tests are derived from the theory of change point detection, where the null hypothesis corresponds to “no signal”.

Suppose we are given a fully grown CART regression tree μ with m leaves and internal nodes (non-leaves, splits) $k = 1, \dots, m-1$. Here, we regard μ as a graph, not as a function as in (4.1). μ was built by greedy recursive binary splitting of the covariate space using iid data $(X^{(i)}, Y^{(i)})_{i=1}^n$ of covariate-response pairs taking values in $\mathbb{R}^d \times \mathbb{R}$, where $d \in \mathbb{N}$ is the covariate space dimension. For more details, we refer to Section 4.1. We would like to select a pruned subtree of μ which does not overfit to the data and with good prediction performance. For this purpose, for a given internal node $k \in \{1, \dots, m-1\}$ of μ, we derive a test statistic $U_{\max,k}$, where large values indicate the existence of signal in node k. The test statistic is based on a generalised log likelihood ratio in the framework of change point testing. Approximating the distribution of $U_{\max,k}$ under the null hypothesis of no signal enables to compute an approximate p-value p_k. Low values of p_k indicate that the null hypothesis of no signal in node k is unlikely. In this light, p_k can be interpreted as a "no-signal penalty".

Let us suppose to be given a sequence of nested trees $\mu^{(1)} \subset \mu^{(2)} \subset \cdots \subset \mu^{(r)} = \mu$, where the inclusion $\widetilde{\mu} \subset \mu$ means that $\widetilde{\mu}$ can be obtained from μ by iteratively removing pairs of sibling leaves (pruning). $\mu^{(1)}$ consists of only a root node while $\mu^{(r)} = \mu$ is the fully grown CART regression tree with m leaves. Our goal is to select a tree μ^* from the sequence of subtrees with good prediction and regularisation performance.

To this end, fix a tolerance level $\delta > 0$. The idea is to look at one subtree at a time and to sum up all the p-values of its internal nodes. The sum of p-values can be regarded as a term penalising "no signal". Under the null, each node-p-value follows a uniform distribution on $[0, 1]$. Therefore, if the sum of p-values is bounded by δ, we reject the null hypotheses and we continue to the next larger subtree. Otherwise, we accept the null hypothesis and the method stops.

For a fixed node k, we prove that if there exists a step signal with strictly positive step size (denoted by $\mathrm{P}_{\mathcal{A}}$), then our p-value approximation converges to zero in $\mathrm{P}_{\mathcal{A}}$-probability as the number of data points n tends to infinity (cf. (Paper III, Proposition 1)). In addition, this holds true if the step size decreases to zero at a slow rate. Hence, our test detects any positive step signal with a probability tending to one as n tends to infinity.

The central contribution of this paper is the establishment of a link between CART split-point selection and the theoretical framework of change point testing. In particular, the test statistic $U_{\max,k}$ admits a natural motivation as the maximum of generalised log-likelihood ratio statistics for testing the presence of a change point in a covariate dimension. At the same time, $U_{\max,k}$ emerges as a byproduct of the CART optimisation procedure under an L^2 loss.

In the present work we restrict attention to continuous covariates and assume that the responses, conditional on a covariate realisation, are normally distributed with constant (homoscedastic) variance. One may argue that, deeper in the tree, the distribution of $U_{\max,k}$ should be conditioned on the sequence of rejections of the null hypothesis along the path from the root to node k. In our approach, such a conditioning is neglected.

Possible extensions of the analysis include deriving finite-sample power-bounds for the proposed test. Such results could provide quantitative guidance on sample sizes required to attain a prescribed type II error level.

6.4 Paper IV

In Paper IV, we consider a dividend problem with ruin at zero surplus and an additional random time horizon. This additional random time horizon is based on the occupation time of the surplus process below a given distress threshold. Depending on this threshold, three different solution structures for the value function and the corresponding optimal control are derived. Allowing for default triggered by a persistently low surplus constitutes a realistic extension of the classical dividend problem. Examples include capital providers withdrawing funds or regulatory intervention following a prolonged period of low surplus. Formally, consider the surplus process governed by the controlled SDE

$$dX_t^D = \mu dt + \sigma dW_t - dD_t, \quad t \geq 0, \quad X_{0-}^D = x \geq 0, \tag{6.8}$$

where $\mu, \sigma > 0$ and $(W_t)_{t\geq 0}$ is a Brownian motion. The control process D_t describes the cumulative dividends up to time t and belongs to the set

$$\begin{aligned}\mathcal{A} := \{&(D_t)_{t\geq 0} \text{ is increasing, right-continuous, } (\mathcal{F}_t)\text{-adapted,}\\ &\text{with } D_{0-} = 0, \text{ such that } D_t - D_{t-} \leq X_{t-}^D, \text{ for all } t \geq 0\}.\end{aligned} \tag{6.9}$$

The occupation time that the surplus process spends below a given distress threshold $y > 0$ is defined as

$$\omega_t^y = q \int_0^t I_{\{X_s^D < y\}} ds, \quad t \geq 0, \tag{6.10}$$

where $q > 0$ can be interpreted as a penalisation rate, determining how likely a low surplus leads to default. We further define the expected reward by

$$J(x; y, D) = \mathrm{E}_x \left[\int_0^{\tau_0^D} e^{-rt} \mathbb{1}_{\{\omega_t^y < e_1\}} dD_t \right], \tag{6.11}$$

where $r > 0$ is the discount rate, τ_0^D is the first hitting time of zero and e_1 is an exponentially distributed random variable with unit mean which is independent of X^D and D. We are interested in

$$V(x, y) := \sup_{D \in \mathcal{A}} J(x; y, D), \tag{6.12}$$

denoted the value function of the problem as well as in finding maximisers $D^* \in \mathcal{A}$ attaining $V(x, y)$. An important observation is that the expected reward (6.11) can be rewritten as

$$J(x; y, D) = \mathrm{E}_x \left[\int_0^{\tau_0^D} e^{-rt - \omega_t^y} dD_t \right].$$

This formulation shows that the problem can be interpreted as a classical dividend problem with state-dependent discounting.

We show that for any fixed distress level $y > 0$, the optimal control can be characterised by a waiting region $\mathcal{W}^y$ and an action region $\mathcal{D}^y$, such that $\mathcal{D}^y \cap \mathcal{W}^y = \emptyset$ and $\mathcal{D}^y \cup \mathcal{W}^y = (0, \infty)$. When the surplus process starts at a point $x \in \mathcal{W}^y$, no dividends are paid, and the surplus evolves like an uncontrolled Brownian motion with drift. On the other hand, when the surplus process is started at a point $x \in \mathcal{D}^y$, there is an instantaneous minimal dividend payout $D_0^* - D_{0-}^* = D_0^* \geq 0$, such that $x - D_0^* \in \overline{\mathcal{W}}^y$, the closure of the waiting region. We show that there exist three regimes for the distress threshold y giving rise to different solution structures for the optimal control. More formally, there exist $0 < y_l < y_u < \infty$ depending only on the four positive exogenous model parameters μ, σ, q and r, such that

(i) *(low distress threshold)* if $0 < y \leq y_l$, the optimal control is given by a single-barrier strategy with $\mathcal{W}^y = \big(0, b^*(y)\big), \mathcal{D}^y = \big[b^*(y), \infty\big)$, where $b^*(y) > y$ admits an explicit form,

(ii) *(intermediate distress threshold)* if $y_l < y < y_u$, the optimal control is given by a double barrier strategy characterised by $\mathcal{W}^y = \big(0, b^*_{r+q}\big) \cup \big(\underline{b}^*(y), \overline{b}^*(y)\big)$, $\mathcal{D}^y = \big[b^*_{r+q}, \underline{b}^*(y)\big] \cup \big[\overline{b}^*(y), \infty\big)$ with $b^*_{r+q} < \underline{b}^*(y) < y < \overline{b}^*(y)$,

(iii) *(large distress threshold)* if $y_u \leq y$, the optimal control is again given by a single-barrier strategy and $\mathcal{W}^y = \big(0, b^*_{r+q}\big), \mathcal{D}^y = \big[b^*_{r+q}, \infty\big)$, where $b^*_{r+q} < y$ and the reflecting barrier does not depend on y.

Here, b^*_{r+q} denotes the reflecting barrier of the classical dividend problem with $(r+q)$-discounting and admits an explicit expression given in (5.36). By contrast, the existence and uniqueness of $\underline{b}^*(y) < y < \overline{b}^*(y)$ are established via the intermediate value theorem, and only implicit characterisations are available. For each of the regimes (i) – (iii), we provide explicit expressions for the value function $V(x; y)$.

A possible extension of this problem is to consider state-dependent coefficients μ and σ. This would, however, break the exponential solution structure of the value function (cf. (5.20)) and thus substantially complicate the analytical approach adopted in this work.

Sammanfattning på svenska

Denna avhandling behandlar flera ämnen inom försäkringsmatematik och tillämpad sannolikhetsteori, däribland försäkringsvärdering och reservsättning, regularisering av regressions- och klassificeringsträd samt stokastisk optimering i ett utvidgat utdelningsproblem. Avhandlingen bygger på fyra artiklar.

Artikel I ger en motivering av chain ladder-prediktorn och Macks estimator för prediktionsfelet inom en klassisk sammansatt Poissonmodell under ett antagande som stor exponering (t.ex. stort antal kunder). Även om modellen inte uppfyller antagandena för Macks fördelningsfria chain ladder-metod visas att både prediktorn och estimatorn uppstår naturligt asymptotiskt då vi låter exponeringen växa.

Artikel II studerar värdering av skuldkassaflöden med kapitalkrav i en flerperiodsmodell. Eftersom explicit värdering i allmänhet är ogenomförbar och Monte Carlo-metoder ofta är beräkningsmässigt krävande, härleds en explicit och lättberäknad värderingsformel. Formeln erhålls som ett gränsfall vid stor exponering under ett antagande om betingad svag konvergens för skuldkassaflödena.

Artikel III introducerar en regulariseringsmetod för CART-regressionsträd baserad på nodvisa statistiska tester. I varje nod beräknas ett p-värde med hjälp av ett statistiskt test för en ändring av väntevärde för en följd av stokastiska variabler, vilket resulterar i ett regulariserat CART-regressionsträd som är en deterministisk funktion av träningsdatan. Till skillnad från korsvalidering undviker metoden slumpmässighet från uppdelning av datan i delmängder och säkerställer ett effektivt nyttjande av hela datamängden.

Artikel IV återbesöker det klassiska utdelningsproblemet med ruin vid nivån noll genom att införa en ytterligare defaultmekanism baserad på kumulativ uppehållstid i ett låg-överskottsområde. Denna utvidgning återspeglar realistiska defaultutlösare såsom regulatoriskt tryck eller likviditetsstress. Problemet löses explicit, vilket ger slutna uttryck för både den optimala styrningen och värdefunktionen.

Bibliography

Aldous, David (1981). "Weak convergence and general theory of processes". Unpublished draft of monograph; Department of Statistics, University of California, Berkeley, CA 94720.

Asmussen, Søren and Taksar, Michael (1997). Controlled diffusion models for optimal dividend pay-out. *Insurance: Mathematics and Economics* 20.1, pp. 1–15.

Backhoff-Veraguas, Juan, Bartl, Daniel, Beiglböck, Mathias and Eder, Michael (2020). All adapted topologies are equal. *Probability Theory and Related Fields* 178, pp. 1125–1172.

Bodnariu, Andi, Engler, Nils and Rodosthenous, Neofytos (2026). Outrunning the Omega Clock: A Singular Control Problem for Dividend Optimisation with Ruin and Time-in-Distress Default. arXiv: `2601.21705`.

Breiman, Leo, Friedman, Jerome H., Olshen, Richard A. and Stone, Charles J. (1984). *Classification and Regression Trees*. Belmont (California): Wadsworth.

Carathéodory, Constantin (1929). Stetige Konvergenz und normale Familien von Funktionen. *Mathematische Annalen* 101, pp. 515–533.

Engler, Nils, Lindholm, Mathias, Lindskog, Filip and Nazar, Taariq (2025). Regularisation of regression trees by summation of p-values. arXiv: `2505.18769v2`.

Engler, Nils and Lindskog, Filip (2024). Mack's estimator motivated by large exposure asymptotics in a compound Poisson setting. *ASTIN Bulletin* 54, pp. 310–326.

Engler, Nils and Lindskog, Filip (2025). Approximations of multi-period liability values by simple formulas. *Insurance: Mathematics and Economics* 123, 103112.

Engsner, Hampus, Lindholm, Mathias and Lindskog, Filip (2017). Insurance valuation: A computable multi-period cost-of-capital approach. *Insurance: Mathematics and Economics* 72, pp. 250–264.

Föllmer, Hans and Schied, Alexander (2016). *Stochastic Finance: An Introduction in Discrete Time*. Berlin, Boston: De Gruyter.

Hahn, Hans (1921). *Theorie der reellen Funktionen*. Berlin, Heidelberg: Springer.

Hawkins, Douglas M. (1977). Testing a sequence of observations for a shift in location. *Journal of the American Statistical Association* 72.357, pp. 180–186.

Kallenberg, Olav (2002). *Foundations of Modern Probability*. 2nd ed. New York: Springer.

Klenke, Achim (2013). *Wahrscheinlichkeitstheorie*. 3rd ed. Berlin, Heidelberg: Springer Spektrum.

Liptser, Robert and Shiryaev, Albert (2013). *Statistics of Random Processes: I General Theory*. Berlin, Heidelberg: Springer.

Mack, Thomas (1993). Distribution-free calculation of the standard error of chain ladder reserve estimates. *ASTIN Bulletin* 23.2, pp. 213–225.

Mikosch, Thomas (2009). *Non-life Insurance Mathematics*. Berlin, Heidelberg: Springer.

Peskir, Goran (2007). "A Change-of-Variable Formula with Local Time on Surfaces". *Séminaire de Probabilités XL*. Ed. by Donati-Martin, Catherine, Émery, Michel, Rouault, Alain and Stricker, Christophe. Vol. 1899. Lecture Notes in Mathematics. Berlin, Heidelberg: Springer, pp. 70–96.

Pilipenko, Andrey (2014). *An introduction to stochastic differential equations with reflection*. Vol. 1. Lectures in pure and applied mathematics. Potsdam: Potsdam University Press.

Revuz, Daniel and Yor, Marc (1999). *Continuous Martingales and Brownian Motion*. 3rd ed. Berlin, Heidelberg: Springer.

Sweeting, Trevor J. (1989). On conditional weak convergence. *Journal of Theoretical Probability* 2.4, pp. 461–474.

Touzi, Nizar (2004). *Stochastic Control Problems, Viscosity Solutions and Application to Finance*. Pisa: Scuola Normale Superiore.

Villani, Cédric (2009). *Optimal Transport: Old and New*. Berlin, Heidelberg: Springer.

Yao, Yi-Ching and Davis, Richard A. (1986). The asymptotic behavior of the likelihood ratio statistic for testing a shift in mean in a sequence of independent normal variates. *Sankhyā: The Indian Journal of Statistics, Series A* 48.3, pp. 339–353.

Part II

Papers