UNIL | Université de Lausanne

Unicentre
CH-1015 Lausanne
http://serval.unil.ch

Year : 2024

OPTIMIZING THE TRADE-OFF BETWEEN RISK AND PROFITABILITY IN CRYPTOCURRENCY MINING AND UNDER MISCLASSIFICATION IN INSURANCE

Finger Dina

Finger Dina, 2024, OPTIMIZING THE TRADE-OFF BETWEEN RISK AND PROFITABILITY IN CRYPTOCURRENCY MINING AND UNDER MISCLASSIFICATION IN INSURANCE

Originally published at : Thesis, University of Lausanne

Posted at the University of Lausanne Open Archive http://serval.unil.ch
Document URN : urn:nbn:ch:serval-BIB_2D8AA750EE760

Droits d'auteur
L'Université de Lausanne attire expressément l'attention des utilisateurs sur le fait que tous les documents publiés dans l'Archive SERVAL sont protégés par le droit d'auteur, conformément à la loi fédérale sur le droit d'auteur et les droits voisins (LDA). A ce titre, il est indispensable d'obtenir le consentement préalable de l'auteur et/ou de l'éditeur avant toute utilisation d'une oeuvre ou d'une partie d'une oeuvre ne relevant pas d'une utilisation à des fins personnelles au sens de la LDA (art. 19, al. 1 lettre a). A défaut, tout contrevenant s'expose aux sanctions prévues par cette loi. Nous déclinons toute responsabilité en la matière.

Copyright
The University of Lausanne expressly draws the attention of users to the fact that all documents published in the SERVAL Archive are protected by copyright in accordance with federal law on copyright and similar rights (LDA). Accordingly it is indispensable to obtain prior consent from the author and/or publisher before any use of a work or part of a work for purposes other than personal use within the meaning of LDA (art. 19, para. 1 letter a). Failure to do so will expose offenders to the sanctions laid down by this law. We accept no liability in this respect.

UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES

DÉPARTEMENT DE SCIENCES ACTUARIELLES

OPTIMIZING THE TRADE-OFF BETWEEN RISK AND PROFITABILITY IN CRYPTOCURRENCY MINING AND UNDER MISCLASSIFICATION IN INSURANCE

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales
de l'Université de Lausanne

pour l'obtention du grade de
Doctorat en Sciences actuarielles

par

Dina FINGER

Directeur de thèse
Prof. Hansjoerg Albrecher

Co-directeur de thèse
Prof. Pierre-Olivier Goffard

Jury

Prof. Boris Nikolov, président
Prof. Joël Wagner, expert interne
Prof. Carole Bernard, experte externe

LAUSANNE
2024

UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES

DÉPARTEMENT DE SCIENCES ACTUARIELLES

OPTIMIZING THE TRADE-OFF BETWEEN RISK AND PROFITABILITY IN CRYPTOCURRENCY MINING AND UNDER MISCLASSIFICATION IN INSURANCE

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales
de l'Université de Lausanne

pour l'obtention du grade de
Doctorat en Sciences actuarielles

par

Dina FINGER

Directeur de thèse
Prof. Hansjoerg Albrecher

Co-directeur de thèse
Prof. Pierre-Olivier Goffard

Jury

Prof. Boris Nikolov, président
Prof. Joël Wagner, expert interne
Prof. Carole Bernard, experte externe

LAUSANNE
2024

UNIL | Université de Lausanne
Faculté des hautes études commerciales
Internef - 261
CH-1015 Lausanne

IMPRIMATUR

La Faculté des hautes études commerciales de l'Université de Lausanne autorise l'impression de la thèse de doctorat rédigée par

Dina FINGER

intitulée

Optimizing the Trade-off between Risk and Profitability in Cryptocurrency Mining and under Misclassification in Insurance

sans se prononcer sur les opinions exprimées dans cette thèse.

Lausanne, le 19.01.2024

Professeure Marianne Schmid Mast, Doyenne

Members of the Jury

Prof. **Boris Nikolov**
President of the jury,
Department of Finance, University of Lausanne, Switzerland.

Prof. **Hansjörg Albrecher**
Thesis supervisor,
Department of Actuarial Science, University of Lausanne, Switzerland.

Prof. **Pierre-Olivier Goffard**
Thesis co-supervisor,
Department of Mathematics and Computer Science, University of Strasbourg, France.

Prof. **Joël Wagner**
Internal expert,
Department of Actuarial Science, University of Lausanne, Switzerland.

Prof. **Carole Bernard**
External expert,
Department of Accounting, Law and Finance, Grenoble École de Management, France and Faculty of Economics, Vrije Universiteit, Belgium.

University of Lausanne
Faculty of Business and Economics

PhD in Actuarial Science

I hereby certify that I have examined the doctoral thesis of

Dina FINGER

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature:  Date: 22.12.2023

Prof. Hansjörg ALBRECHER
Thesis supervisor

University of Lausanne
Faculty of Business and Economics

PhD in Actuarial Science

I hereby certify that I have examined the doctoral thesis of

Dina FINGER

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature:  Date: 21/12/2023

Prof. Pierre-Olivier GOFFARD
Thesis co-supervisor

University of Lausanne
Faculty of Business and Economics

PhD in Actuarial Science

I hereby certify that I have examined the doctoral thesis of

Dina FINGER

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature:  Date: 21/12/2023

Prof. Joël WAGNER
Internal expert

University of Lausanne
Faculty of Business and Economics

PhD in Actuarial Science

I hereby certify that I have examined the doctoral thesis of

Dina FINGER

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature:  Date: December 21st, 2023

Prof. Carole BERNARD
External expert

Acknowledgements

First of all, I would like to thank my supervisors Prof. Hansjörg Albrecher and Prof. Pierre-Olivier Goffard for their scientific guidance and wisdom. Working with them was a life changing experience and a constant learning process thanks to their valuable advice. I was motivated to undertake this PhD journey when I first attended a lecture of Prof. Hansjörg Albrecher and this path lead me to the person I became now.

I would also like to thank my PhD committee members Prof. Joël Wagner and Prof. Carole Bernard. Their insightful comments and suggestions helped me to improve the quality of my thesis.

I gratefully acknowledge the funding received from Vaudoise Assurances and the employment opportunity that was given to me. Special thanks go to my Pricing & Reporting team that taught me important practical lessons.

Of course, I want to express my sincere acknowledgements to the Department of Actuarial Science and other residents of the Extranef building that made this adventure so special and full of memories. I will never forget our exciting trips and card games that might even open a new research area. Also, I want to express my gratitude to all my friends that followed my journey and encouraged me along the way.

Last but not least, I would like to thank my family that has always supported me throughout my life and particularly in this intense period. Without their constant care, all these accomplishments would not have been possible. Finally, I would like to thank Giulio for his never ending patience, love and support.

Summary

This thesis is composed of three articles addressing issues at the intersection of ruin theory, cryptocurrency mining, and insurance economics. A key focus is examining the quantitative trade-offs between risk and profitability across different contexts.

In the first article, we explore the risky nature of cryptocurrency mining under the Proof-of-Work consensus mechanism. We introduce the concept of ruin, where operational costs can outweigh mining rewards. Our research investigates the benefits and drawbacks of joining a mining pool, which reduces return variability at a specific cost. By applying ruin theory and risk-sharing principles from actuarial science, we derive explicit formulas for key metrics like expected value of the profit and ruin probability, demonstrated through numerical examples with parameters of practical relevance.

The second article builds upon these findings, shifting from theoretical to empirical analysis. Here, we scrutinize the validity of our theoretical models in real-world mining scenarios. A significant addition is the inclusion of transaction fees in block rewards. We introduce algorithms for fitting generalized hyperexponential distributions to actual data and conduct a sensitivity analysis to assess various factors in mining, particularly the effects of temporal dependencies and transaction fees. We conclude that despite the apparent time dependence in the rewards, approximations by combinations of exponentials with an i.i.d. assumption yields satisfying results.

In the third article, we shift our focus to the non-life insurance industry, specifically the challenges of risk classification. We address the issue of policyholder categorization into risk classes and the potential errors in this process, particularly relevant with the advent of automatic classification systems. We explore the consequences of misclassifying policyholders into two distinct risk types, providing a mean-variance framework to study the insurer's optimization problem in setting premiums. This analysis includes examining the cost-benefit trade-off when probabilities of classification errors are known. We develop a simple framework that can be further extended to include competition, multi-period dynamic games and other risk measures.

Résumé

Cette thèse se compose de trois articles abordant des problématiques à l'intersection de la théorie de la ruine, du minage de cryptomonnaies et de l'économie des assurances. L'accent est mis sur l'examen des compromis quantitatifs entre le risque et la rentabilité dans différents contextes.

Dans le premier article, nous explorons la nature risquée du minage de cryptomonnaies sous le mécanisme de consensus par preuve de travail. Nous introduisons le concept de ruine, où les coûts opérationnels peuvent dépasser les récompenses du minage. Notre recherche étudie les avantages et les inconvénients de rejoindre un pool de minage, qui réduit la variabilité des rendements à un coût spécifique. En appliquant la théorie de la ruine et les principes de partage des risques des sciences actuarielles, nous dérivons des formules explicites pour des mesures clés telles que la valeur espérée du profit et la probabilité de ruine, démontrées à travers des exemples numériques avec des paramètres pratiquement pertinents.

Le deuxième article s'appuie sur ces résultats, passant d'une analyse théorique à une analyse empirique. Ici, nous examinons la validité de nos modèles théoriques dans des scénarios de minage réels. Un ajout significatif est l'inclusion des frais de transaction dans les récompenses de blocs. Nous introduisons des algorithmes pour ajuster les distributions hyperexponentielles généralisées aux données réelles et réalisons une analyse de sensibilité pour évaluer divers facteurs dans le minage, en particulier les effets des dépendances temporelles et des frais de transaction. Nous concluons que malgré l'apparente dépendance temporelle dans les récompenses, les approximations par des combinaisons d'exponentielles avec une hypothèse i.i.d. procurent des résultats satisfaisants.

Dans le troisième article, nous orientons notre attention vers l'industrie des assurances non-vie, en particulier les défis de la classification des risques. Nous abordons la question de la catégorisation des preneurs d'assurance en classes de risque et les erreurs potentielles dans ce processus, particulièrement pertinentes avec l'avènement des systèmes de classification automatique. Nous explorons les conséquences d'une mauvaise classification des titulaires de polices en deux types de risques distincts, en fournissant un cadre moyenne-variance pour étudier le problème d'optimisation de l'assureur dans la fixation des primes. Cette analyse comprend l'examen du compromis coût-bénéfice lorsque les probabilités d'erreurs de classification sont connues. Nous développons un cadre simple qui peut être étendu pour inclure la concurrence, des jeux dynamiques sur plusieurs périodes et d'autres mesures de risque.

Contents

Chapter 1

Introduction

Throughout history, correctly understanding the drivers of a particular risk and accurately assessing its amplitude have been preoccupations of various members of the society in order to protect their peers from unexpected events. This thinking led - among others - to the development of actuarial science, which aims to determine the price of uncertainty for insurance applications.

Generally, in the non-life insurance sector, the pricing framework is based on specific characteristics of the insured individuals. A range of classification methods are employed to sort policyholders into risk groups, which are established by the insurance company. Hence, risk classification is a never-aging topic in insurance. Its accuracy ensures fair policy pricing and sustainable portfolio management. In a world where demographic shifts, new technologies, and changing consumer behaviour reshape the risk profiles, actuaries must recalibrate and refine classification techniques. However, with new techniques come new threats, such as erroneous risk classification. The consequences of misclassification are severe, from biased premiums to adverse selection and potential ruin.

In the domain of classification pricing strategies, a balance must be achieved between the accuracy of classification, the associated costs and risk, and the resulting expected profit. On one hand, the adoption of a more sophisticated classification system has the potential to attract better customer profiles, thus improving the quality of the underwritten risk. On the other hand, the implementation of such a system carries with it an increase in both operational costs and the potential for error, which in turn can reduce profitability. Moreover, this complexity in the classification system may have a cascading effect on various other risk metrics and key performance indicators. Thus, a tradeoff between various risk measures has to be found.

This search for balance for an insurer surprisingly finds a parallel in the world of cryptocurrency mining, where the stakes of managing risk and reward are equally high. Just as insurance companies classify risks to price policies accurately and ensure profitability, cryptocurrency miners must assess the cost-benefit dynamics of participating in mining activities. In particular, for a miner, this translates to the risk of operational costs surpassing mining rewards, leading to financial insolvency. This scenario is akin to a ruin event in insurance, where companies face the risk of

becoming unable to cover claims.

The area of cryptocurrencies is relatively new, having emerged just over the past fifteen years. This period has witnessed economic bubbles and financial crashes. Millions have engaged in mining or purchasing cryptocurrencies, often without a comprehensive grasp of the underlying mechanisms. Various reward systems and pooling schemes present a jungle of complex terminology. This thesis aims to offer an actuarial perspective on this topic and to draw connections to ruin theory.

This introductory chapter will lay out some foundations for the understanding of this present thesis. In Section 1.1, we will get an overview of classification techniques and criteria selection, which are a starting point for Chapter 4. Section 1.2 will present the classic results of ruin theory, which are further used in Chapters 2 and 3. Furthermore, Section 1.3 provides a short overview of the principles of blockchains and the emergence of cryptocurrencies, to increase the understanding of Chapters 2 and 3. Finally, we summarize the main contributions of this thesis in Section 1.4.

1.1 Risk classification in insurance pricing

Risk classification in insurance is a fundamental process used to categorize potential policyholders based on various risk factors. It involves assessing the likelihood and potential cost of a claim associated with an individual or group. This process is not only critical in determining premiums but also in ensuring the financial sustainability of insurance products. The underlying principle is that similar risks should be grouped together, and each group should be charged a premium proportional to the level of risk they represent. In a monograph by the American Academy of Actuaries [51] authors point out the importance of finding balance between maximizing the category size to increase historical data volume and minimizing the inhomogeneity within the class and thus the adverse selection.

1.1.1 Classification criteria

The rationale for risk classification has several drivers [63, 116]. Let us give an overview of the most important ones in the insurance realm.

At the heart of risk classification lies the actuarial and statistical basis, which forms the cornerstone of insurance underwriting. Actuaries employ statistical models to assess risk, drawing on historical data to forecast future claims. This approach, however, is not without its challenges. The accuracy of these models is contingent upon the quality and relevance of the data used. Errors can occur due to outdated data, incorrect assumptions, or oversimplified models that fail to capture the complexity of real-world scenarios. Such inaccuracies in risk assessment can lead to mispriced premiums and adverse selection from the customer's side.

Operational factors also play a crucial role in risk classification. These include the costs associated with the identification and processing of risk data, as well as the verifiability of the information collected. High operational costs may lead to shortcuts in risk assessment processes, potentially resulting in errors. Moreover, the

reliability of the data gathered is crucial: unverified or erroneous data can skew risk assessments, leading to misclassified risks.

The social dimension of risk classification covers the respect of sensible customer data and ensuring social acceptability and affordability for the insured. Insurance companies must navigate between thorough risk assessment, and the invasion of privacy and the use of unethical classification criteria, though they may be statistically relevant. Additionally, the affordability of insurance products is a social concern. If risk is misclassified, resulting in exorbitant premiums, insurance may become inaccessible to those who need it most.

Lastly, the legal framework within which insurance operates cannot be ignored. State regulations and equality legislation play a crucial role in shaping how risks are classified. As a prominent example, we mention here the European Court of Justice ruling on the prohibition of discrimination by gender in insurance pricing [112]. Compliance with these legal requirements is imperative to avoid penalties and maintain the insurer's reputation. However, legal constraints can also introduce challenges in risk classification, as they limit the factors that insurers can consider, potentially leading to less accurate risk assessments.

Recently, with the development of genetic testing and personalized medicine, new challenges have arisen with respect to data protection. By undergoing a genetic test and sharing data with the insurer, the latter gains knowledge about the individual's risk type which in turn influences the classification scheme. This triggers some concerns from the population with regard to privacy considerations for sensitive data sharing [50]. Also, the challenge of including personalized medicine into healthcare systems may cause inequalities at a very granular level, since some patients would receive very costly treatments that need to be covered by the insurer. A comprehensive review of those challenges can be found in [84].

One should not disregard the origins of insurance which lay in the principles of risk sharing among the undertakers. Insurance is not an exclusive instrument of protection against risk; an individual can alternatively turn to self-protection and a self-insurance technique to mitigate his risk. In this respect, risk classification arises as a means for the insurer to compete with alternatives [1]. With inaccurate risk classification, however, there is potential risk that the inefficiency is unevenly distributed and thus penalizes an already disfavoured group. Another aspect to note is the problem of adverse selection, mentioned above. From a public policy perspective, adverse selection may be desired to some extent, as it increases the coverage of particularly exposed individuals [139]. From this point of view, risk classification may be skewed in a suitable direction to achieve this objective via public policies.

In practice, multiple relevant criteria are used to classify the customer in categories. Those can be grouped in the following distinguishable categories. Some further examples of criteria and analysis of their fairness can be found in Charpentier [32].

Demographic characteristics This category includes criteria that are specific to the individual or entity seeking insurance. Among them we can cite the most notable examples relevant to both life and non-life insurance. Age and gender constitute for instance the main rating factors in a life insurance policy. The geographical position, such as the address, also constitutes a valid factor since the proximity to dangerous borders may favour theft, and the type of city among more rural or more urban areas may influence the driving behaviour. Further rating criteria in motor insurance are presented in Lemaire [93]. For working-related benefits, occupation is a relevant factor as some professions are deemed to have a more risky impact on the accident frequency.

Object characteristics In non-life insurance, the object of the contract is often not the individual, but a property or casualty (mainly, with exception of health insurance). For example, in motor insurance, the main rating factors, in addition to the driver's characteristics, would be the car's price, model, horsepower, type of gear. For a pet coverage, the insurer will be interested in the pet's breed and gender, whereas for household insurance, one will be questioned about the type of construction.

These two categories group criteria that are often observable and cannot be hidden, see e.g. Denuit et al. [48, chap. 1]. Mostly, they are verified upon the claim occurrence and the insurer reserves the right to decline a coverage in case of a falsely reported characteristic. For the next typology of criteria, the visibility of a factor is more complicated.

Risk behaviour This category involves criteria that directly relate to the history of the insured or the behaviour of the insured person or activity. Notably, the claims history enters this category. The latter often enters the ratemaking system via a bonus-malus framework, such as in the motor insurance. More recently, with the advances of technologies, possibilities of live-tracking customers with devices installed on the car engine raised the monitoring of the client's behaviour onto a new level [146]. The insurance based on telematics criteria is refereed to as *Pay-as-you-drive* in practice. This system can help preventing cheating on the behaviour declaration. On an individual level, lifestyle habits, such as smoking, can also lead to higher premiums. The type of usage of property can also categorize a client in a riskier category, such as the use of a house as primary or secondary residence, or the participation or not of an insured horse in exhausting sport competitions. Often of declarative nature, those criteria can be falsely reported to reduce the apparent risk.

Policy-related factors Finally, the insurance contract also differs in some policy-specific factors. This can include the coverage amount, deductibles, and policy limits, which directly affect the premium but are choices made by the policyholder. Analyzing these factors may uncover the self-selection by the insured people, as a policyholder who expects more claims might for instance take a lower deductible.

Even for verifiable criteria, errors are inevitable due to human mistakes or system-imposed restrictions such as limited input fields that cut away important

information. Furthermore, errors in classification might arise as the historical data on which estimations are made may be not sufficient and lead to wrong conclusions.

1.1.2 Classification techniques

In this subsection, we will give a brief overview of popular machine learning classification techniques. Generally, they fall into two broad classes of algorithms: supervised and unsupervised learning. In a survey on the use of artificial intelligence in actuarial science, Richman [118, 119] describes the difference between those two, and introduces the existence of a third one: reinforcement learning. The idea of the latter is learning what to do to maximize a numerical reward signal, without being told what actions to undertake [132]. Since very few applications of this class of artificial intelligence (AI) are used in insurance (see a rare occurrence in Choi et al. [35] for fraud detection in insurance claims), let us concentrate on the first two.

Supervised learning

This method involves training a model on a labelled dataset, where the outcomes are known. This method is highly valuable in actuarial tasks like claim prediction, risk assessment, and pricing models. The common point of these techniques is having an input set and an output set, which are somehow linked. The goal of these techniques is to use the inputs to predict the outputs [75]. Once a classifier has been "learned", it can be used to perform prediction on a novel dataset. The inputs in the classification task can consist of numerical variables, but categorical features are also admissible.

Bayesian classifier This very simple approach in classification is based on Bayesian decision theory, combining prior knowledge with new evidence [137]. Some techniques using Bayesian methods can be found in Wütrich and Merz [154, chap. 6]. Given a set of M classes $\omega_i,\ i = 1, 2, \ldots, M$ let us classify a novel input x to the class ω_i that maximizes the posterior probabilities, i.e.

$$\omega_i = \underset{\omega_j}{\operatorname{argmax}}\, \mathbb{P}\left(\omega_j \mid x\right).$$

After information on some observations x has been obtained, the classifier is updated following the Bayes theorem

$$\mathbb{P}\left(\omega_j \mid x\right) = \frac{f(x \mid \omega_j)\mathbb{P}\left(\omega_j\right)}{f(x)},\ j = 1, 2, \ldots, M,$$

where $f(x)$ is the p.d.f. of x. The classifier is updated using the posterior distributions

$$\omega_i = \underset{\omega_j}{\operatorname{argmax}}\, f(x \mid \omega_j)\mathbb{P}\left(\omega_j\right).$$

Both the conditional p.d.f. and the a priori probability can be estimated from the training dataset. An illustration of this classifier is shown in Figure 1.1. Also note that the Bayesian classifier minimizes the misclassification error on the training set [137].

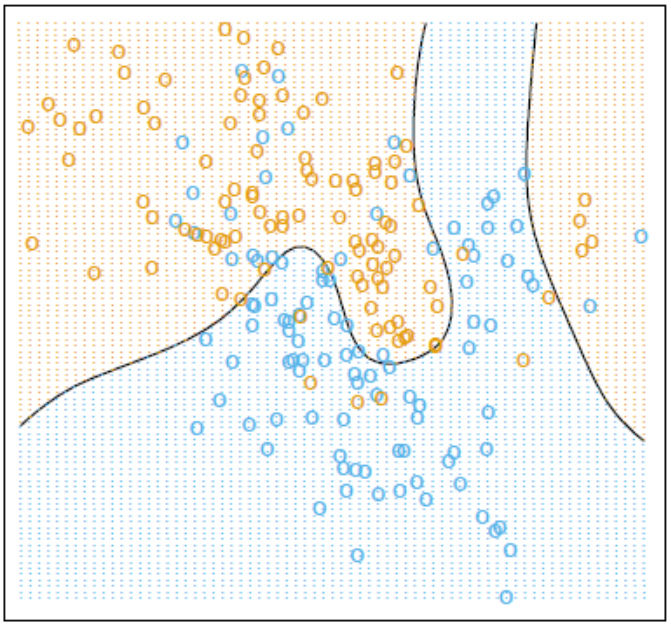

Figure 1.1: Bayesian classifier. Figure taken from [75].

Logistic regression Let us consider a dataset which needs to be classified in two categories. The logistic regression is particularly suited to model the probability that an observation falls into one class or another [81]. Assume we want to classify between two classes 0 and 1. Given a set of input (explicative) variables X, we observe the class Y (the output variable) for the training set. We aim to predict the probability of a novel observation to fall into one class, i.e.

$$p(X) = \mathbb{P}(Y = 1 \mid X).$$

As the predicted results must be a probability comprised between 0 and 1, the linear regression technique would not satisfy the conditions. Thus, we resort to the logistic function

$$p(X) = \frac{e^{\boldsymbol{\beta} X}}{1 + e^{\boldsymbol{\beta} X}}.$$

For the calibration of the vector $\boldsymbol{\beta}$ of parameters, a maximum likelihood maximisation can be performed. Naturally, $\boldsymbol{\beta}$ can admit an intercept. As for the results, the sensitivity of the prediction with respect to the change in X will not only depend on the underlying linear regression parameters, but also on the initial value of X itself.

k-nearest-neighbour In this method, one considers the observations in the training set closer in input space to the variable to classify as its *neighbours*, and the predicted output is then the same as those of the neighbours. Let x_i, $i = 1, 2, \ldots$ be the features of a risk i or the input variables. The output y_i is also observed for our training set. We denote the neighbourhood of x by $N_k(x)$ assuming k nearest neighbours (measured by a distance metric, like the Euclidean distance). The predicted output $\hat{Y}$ for a novel input set x will be given by

$$\hat{Y}(x) = \frac{1}{k} \sum_{x_i \in N_K(x)} y_i.$$

As we can observe, the prediction will change depending on the number of neighbours considered. With a smaller k, the separation between classified zones will be fuzzier, as every individual point will have a "zone of influence" around itself. For a higher k, the delimited predicted region frontier will be smoother, as the predicted class is chosen by a majority vote amongst the k nearest neighbours. One should pay attention to the metric used when choosing the unique k parameter: indeed, if $k = 1$, all the observations in the training set would be correctly classified, which leads to a 0 mean-square error on the training dataset. This does not ensure the same results in a independent dataset which results in an overfitting problem.

Unsupervised learning

In this approach, learning involves analyzing data without pre-labelled outcomes. It is useful for discovering underlying patterns, customer segmentation, and anomaly detection in insurance data [81]. We have access to a set of features or patterns, and the goal is to exploit the interesting information provided by those features. The prediction here is not the objective, we would rather group observations based on observed characteristics. For this purpose, an important tool is the principal component analysis, a tool often used for dimensionality reduction before supervised techniques are applied, and a range of clustering techniques for grouping into classes [75]. In insurance, it can be used for customer segmentation or identifying groups with similar risk profiles. Let us give an example of a clustering technique.

K-means One of the most popular methods in clustering, a technique for finding subgroups in a dataset [49]. In simple terms, we aim to maximize intraclass homogeneity and maximize interclass heterogeneity. In a K-means clustering method, one specifies in advance the desired number of clusters K. The algorithm then partitions the dataset into non-overlapping clusters. Define $C_1, \ldots, C_K$ as sets containing the observations in each cluster. Each observation belongs to at least one and not more than one cluster. Let $W(C_k)$ be a measure of within-cluster variation. Informally, we aim to minimize

$$\min_{C_1,\ldots,C_K} \sum_{k=1}^{K} W(C_k).$$

The common choice for the within-cluster variation measure is the squared Euclidean distance

$$W(C_k) = \frac{1}{|C_k|} \sum_{i,i' \in C_k} \sum_{j=1}^{p} (x_{ij} - x_{i'j})^2, \tag{1.1}$$

where $|C_k|$ is the number of observations in the k^{th} cluster and p the number of features of each observation. In other words, we compute the sum of the squared Euclidean distances of all pairwise observations in each cluster. Approaching this problem by brute force is hardly feasible - the possible number of combinations to check is equal to K^n. Therefore, a simple algorithm is used in the K-means clustering: First, randomly assign all observations to the K clusters. This is the initialization step. Next comes the iterative step. For each cluster, compute the centroid, i.e. the mean vector of the p features of the cluster observations. Assign each observation to the closest cluster centroid and repeat from the centroid calculation. An illustration of the K-means clustering is shown in Figure 1.2. The

results will depend on the initialization of the cluster, so it is important to repeat the algorithms multiple times with a different random starting position.

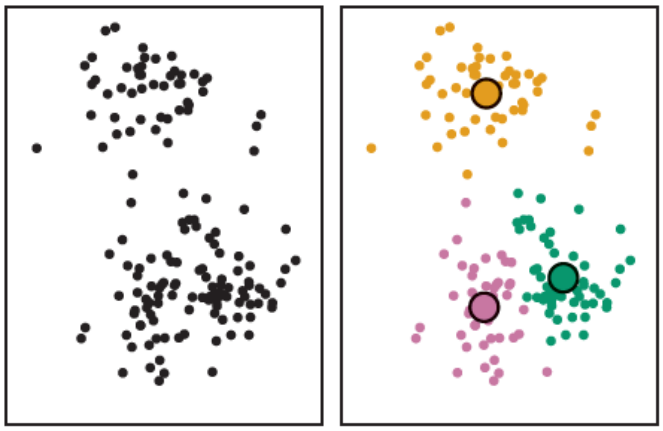

Figure 1.2: K-means classifier, example on the iris dataset[1].

1.2 Ruin theory

Let us define the classical risk reserve process of an insurance company $\{R_t\}_{t\geq 0}$, starting with initial capital $R_0 = u$. Further we assume that claims are defined by their claim count $\{N_t\}$ in the time interval $[0, t]$ and the size of the n^{th} claim is denoted by X_n. Additionally, the premium flow is set as p per time unit. Altogether, we then obtain

$$R_t = u + pt - \sum_{n=1}^{N_t} X_n.$$

Further, we define the ruin probability $\psi(u)$ in infinite time horizon as the probability that the reserve ever goes below zero given the initial capital u:

$$\psi(u) = \mathbb{P}\left(\inf_{t\geq 0} R_t < 0 \mid R_0 = u\right).$$

The time of ruin is then referred to as $\tau(u) = \inf\{t > 0 : R_t < 0\}$. An illustration of this risk reserve process is given in Figure 1.3. The study of the ruin probability, as well as the risk reserve process in general is referred to as the ruin theory. Surveys of general methods on collective risk theory can be found in Asmussen and Albrecher [16], Bühlmann [30], De Vylder [47], Dickson [52], Grandell [69], Kaas et al. [83], Rolski et al. [120] among others.

1.2.1 Cramér-Lundberg model

The origins of the Cramér-Lundberg model trace back to the early 20^{th} century, marking a significant evolution in actuarial thought. Filip Lundberg, a Swedish actuary and mathematician, first introduced the fundamental concepts of this model

[1] From `R` base datasets [117], first introduced by Fisher [64].

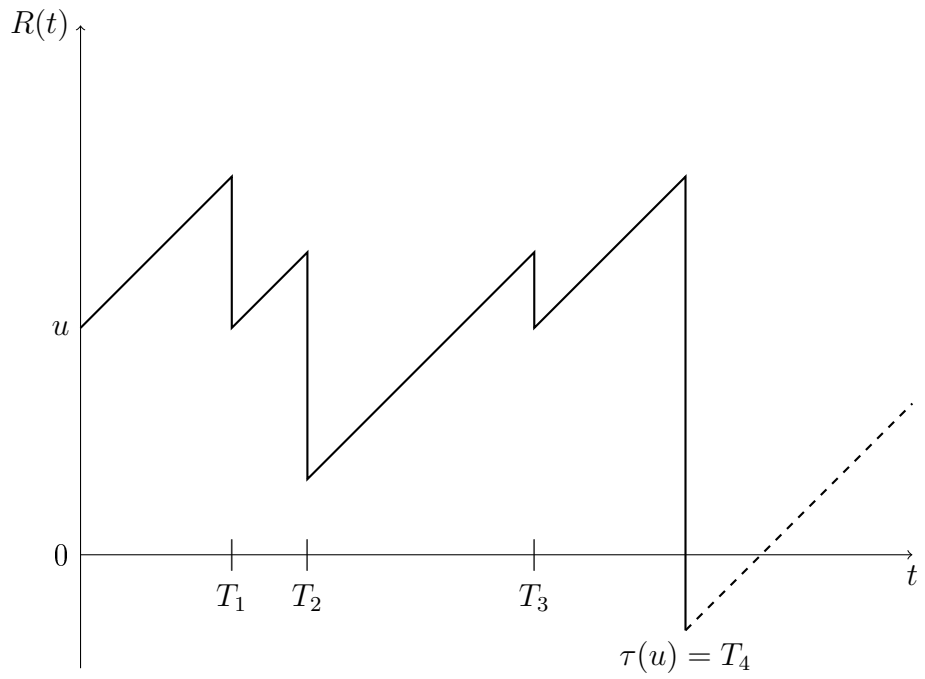

Figure 1.3: Risk reserve process $R(t)$. The claims arrival times are denoted by T_n.

in his thesis [100] as early as 1903, laying the foundations for a systematic approach to calculating probabilities of ruin in insurance companies. His work was further refined and popularized by Harald Cramér, a Swedish statistician, who expanded on Lundberg's ideas and integrated them within the developing field of stochastic processes [40, 41].

For the definition of the Cramér-Lundberg model, we need to respect assumptions 1.2.1-1.2.3.

Assumption 1.2.1. The claims count follows a Poisson process, i.e.

$$\mathbb{P}\left(N_t = k\right) = e^{-\lambda t}\frac{(\lambda t)^k}{k!},\ k = 0, 1, \ldots,$$

with $N_0 = 0$ and independent increments. In other words, N_t is Poisson distributed with mean λt. A comprehensive description of its various properties can be found in Tijms [142]. We can mention some notable properties such as the memoryless property. Define the inter-arrival times of the claims $W_n := T_n - T_{n-1},\ n = 1, 2, \ldots$ and $T_0 = 0$. Then, $\{W_n\}_{n\geq 1}$ are independent and exponentially distributed with mean $1/\lambda$. As we can see, the interarrival times do not depend on t, meaning that for any given moment, the time until the next event occurs is exponentially distributed, just like the initial time between events, and this remains true irrespective of the time passed since the last event. This unique feature of the Poisson process, known as being *memoryless*, means that the elapsed time since the last event does not influence the waiting time for the next one. The next important feature is that under thinning and merging operations, the resulting process remains a Poisson process, see e.g. [34]. For example, the process N_t can be separated into two independent processes N_t^1 and N_t^2 with respective probabilities p and $1 - p$. Both resulting processes will be again Poisson processes with rates $p\lambda$ and $(1 - p)\lambda$ respectively.

As for the merging, the sum of two independent Poisson processes with rates λ_1, λ_2 gives again a Poisson process with rate $\lambda_1 + \lambda_2$.

Assumption 1.2.2. The claims sizes $X_1, X_2, \ldots$ are positive i.i.d. with a common distribution F. Furthermore, claims sizes are independent of N_t the counting process.

Assumption 1.2.3. The premium rate p is constant.

In order for an insurance company to properly function, we also assume that the safety loading condition is satisfied, i.e. $p > \lambda \mathbb{E}(X)$. Now, to derive the ruin probability in infinite time, define $\phi(u) = 1 - \psi(u)$ and condition on the size of the first claim:

$$\begin{aligned} \phi(u) &= \int_0^\infty \lambda e^{-\lambda t} \int_0^{u+pt} \phi(u + pt - x) dF(x) dt \\ &= e^{-\lambda h} \phi(u + ph) + \int_0^h \lambda e^{-\lambda t} \int_0^{u+pt} \phi(u + pt - x) dF(x) dt. \end{aligned} \tag{1.2}$$

The first term in the equation (1.2) describes the case when there is no claim in the interval $[0, h]$. Next, by substitution and differentiation (see e.g. [120] for a proof of differentiability), we derive the following result:

$$\frac{d}{du}\phi(u) = \frac{\lambda}{p}\phi(u) - \frac{\lambda}{p}\int_0^u \phi(u - x) dF(x). \tag{1.3}$$

We define the Laplace transform of a function $h(x)$ as

$$\tilde{h}(s) = \int_0^\infty e^{-sx} h(x) dx.$$

Then, applying the Laplace transform to (1.3), and using $\lim_{u\to 0} \phi(u) = \lim_{s\to 0} s\tilde{\phi}(s) = 1$, we have

$$\tilde{\phi}(s) = \frac{p - \lambda \mathbb{E}(X)}{ps - \lambda(1 - \tilde{f}(s))},$$

where $\tilde{f}(s)$ is the Laplace transform of the p.d.f. of the claims. Note that $\phi(0) = 1 - \lambda \mathbb{E}(X)/p$. Inverting the Laplace transform yields the famous Pollaczek-Khintchine formula

$$\phi(u) = \left(1 - \frac{\lambda \mathbb{E}(X)}{p}\right) \sum_{n=0}^{\infty} \left(\frac{\lambda \mathbb{E}(X)}{p}\right)^n F_I^{*n}(u),$$

where F_I is the integrated tail of F, i.e. $F_I(x) = \frac{1}{\mathbb{E}(X)} \int_0^x (1 - F(y)) dy$ and F^{*n} denotes the n^{th} fold convolution of F. By applying for example a martingale approach, one can derive the Lundberg equation

$$\lambda \int_0^\infty e^{Rx} \, dF(x) = pR,$$

with $R > 0$ being the adjustment (Lundberg) coefficient. It exists whenever the claim distribution is light-tailed. In that case, an upper bound for the ruin probability is $\psi(u) \leq e^{-Ru}$, also known as Lundberg's inequality, see e.g. [69].

1.2.2 Dual risk model

We refer to the model with negative premium inflows and negative claims sizes as the *dual risk model*. This model is defined in (1.4).

$$R_t = u - ct + \sum_{n=1}^{N_t} Y_n. \tag{1.4}$$

The constant outflow can be interpreted as the operational cost and the random upward jumps as the present value of future innovations for example. Another example of this process is a life annuity [69]. An illustration of this risk reserve process is represented in Figure 1.4. Some first appearances can be traced back to Cramér [41]. More recently, developments around this model have been brought to light with respect to optimal dividend problem as in Avanzi et al. [18], Gerber and Smith [67], Ng [111] or with tax payments considerations in Albrecher et al. [4]. In that case, assuming Poisson arrival times, the ruin probability in infinite time

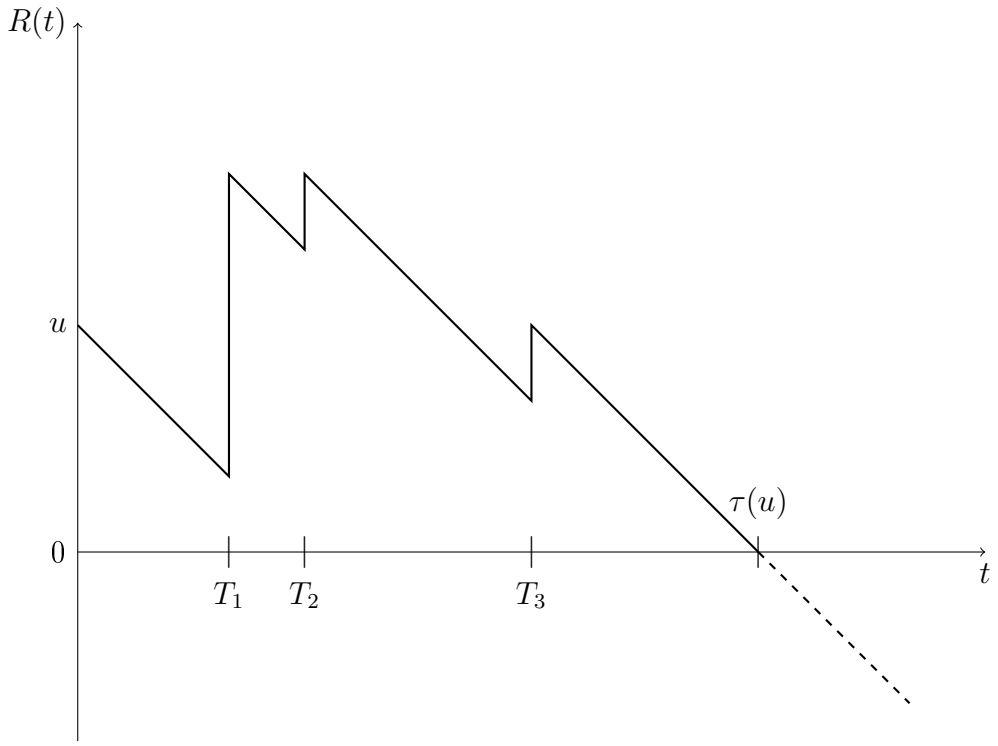

Figure 1.4: Risk reserve process $R(t)$. Ruin occurs at $\tau(u)$, which does not coincide with a jump.

horizon is given by

$$\psi(u) = e^{-Ru},$$

where R is the unique positive solution of the Lundberg equation

$$\lambda - cR = \lambda \tilde{f_Y}(R).$$

For proofs, see e.g. Asmussen and Albrecher [16] or Grandell [69, p.8]. This process is considered in Chapter 2, namely an individual miner's risk process can be interpreted as a dual risk model of the above kind.

1.2.3 Sparre-Anderson model

The Sparre-Anderson model represents a significant advancement in the field of ruin theory, generalizing the classic Cramér-Lundberg model. First ideas were introduced by Erik Sparre Andersen in 1957 [13] and further developed by Thorin [140, 141]. This model accommodates a wider range of claim arrival processes, allowing for other than Poisson distributions. From the Assumptions 1.2.1-1.2.3, Assumption 1.2.1 is relaxed and can be reformulated as Assumption 1.2.4. Assumptions 1.2.2 and 1.2.3 still hold.

Assumption 1.2.4. The arrival times form a renewal process. Let $\{W_n\}_{n\geq 1}$ be inter-arrival times of the claims. Then, $W_1, W_2, \ldots$ are i.i.d. with a common distribution F_W.

The corresponding claims count N_t is then a renewal counting process where $N_t = \sum_{n=1}^{\infty} \mathbb{1}\left(\sum_{k=1}^{n} W_k \leq t\right)$, where $\mathbb{1}(\cdot)$ is the indicator function [69], [120]. In recent literature, various adaptations of the model have been studied and extended for practical purposes. Here are some examples of possible applications. Willmot [152] studies the case of a class of renewal processes characterized by a rational Laplace–Stieltjes transform of the arrival inter-occurrence time distribution. Dufresne [55] derives the Laplace transform of the non-ruin probability for a wide class of claim inter-arrival times or severity distributions, when admitting a rational Laplace transform representation. The time value of ruin with Erlang(n) renewal process is explored in Li and Garrido [94]. A Sparre-Andersen dual model is studied in Yang and Sendova [156] to link this renewal process extension to the model described in the previous subsection.

1.3 Blockchain and cryptocurrencies

Finding a cryptocurrency section in an actuarial science thesis may seem unexpected. However, as Chapter 2 will demonstrate, actuarial tools find their use in a much broader area than just the insurance sector. In order to better understand the further chapters, this short section provides necessary background for further reading.

1.3.1 Blockchains

Blockchain technology, at its core, is a type of decentralized data ledger. It allows data to be stored globally on thousands of servers while letting anyone on the network see everyone else's entries in real-time. This makes it difficult for one user to gain control of the network or manipulate the data [110]. It is maintained collectively by a peer-to-peer network. The type of blockchain we consider in this thesis is public and therefore accessible by anyone. When transactions are realized, they are recorded into the blockchain according to a consensus protocol on which everyone agreed. Once a block entered the blockchain, it cannot be modified anymore. Each block consists of important information, which is immutable [163], which makes it a reliable storage of transactions. As minimal information, blocks usually contain their unique identifier, the *hash*, and the previous (parent) block identifier, which creates the link between them. The timestamp is also part of the block, which ensures that only the first block gets attached to the chain. Moreover, the block contains the

transactions details. Further components of a blockchain include its nodes. These nodes have the job of maintaining the blockchain's integrity. They store, spread, and preserve the blockchain data, thus making the system decentralized. Each node has a full record of the data that has been stored on the blockchain since its inception, allowing for a high level of transparency and security. If one node has an error in its data, it can use the thousands of other nodes as a reference to correct itself. This self-auditing system of a blockchain is what makes it reliable and secure. Acting as validators in this intricate system, there are miners. They use powerful computers to solve complex mathematical problems that validate and secure transactions. This process is known as mining. Mining involves finding a hash that matches the current transaction requirements. It is a trial and error process, and the probability of solving these problems is very low, making mining a competitive and resource-intensive process. In the case of cryptocurrency application, when a miner successfully solves the mathematical problem, they are allowed to add a block to the blockchain and are rewarded with the blockchain's native cryptocurrency. Finally, each blockchain relies on its consensus protocol used to achieve the necessary agreement on a single data value or a single state of the network among distributed processes [15].

Cryptocurrencies aside, there exist multiple applications of blockchains in other fields. Notably, let us mention its usage in smart contracts [162]. Smart contracts are contracts that automatically execute themselves upon meeting some conditions. To state an easy example, when you buy a chocolate in a vending machine, it is a form of smart contract, where the necessary condition is inserting the adequate amount of coins. The advantage of this system is not needing any intermediaries between parties, which is very useful when contracting parties do not trust each other. By using a blockchain, all the transactions within a contract are stored in distributed ledgers. Clauses and conditions of the service and counterpart are written in computer programs and automatically executed. Hence, smart contracts exclude fraud and are transparent.

Another example of application of a blockchain is in supply chain management. During the long process of creation of a good or service, the latter travels through various stakeholders, e.g. producers, retailers, manufacturers. By implementing blockchain technology on a good, via tracking devices, a company can track the movement of the merchandise through the whole process and, if necessary, record the surrounding conditions at each stage. For example, one could detect a failure of cooling systems during the transport of refrigerated food or medicine. Since no one can modify the data during the process as it is stored in a blockchain, it is trustable and transparent. With the help of this technology, the customer can be ensured of the legally certified provenance of the merchandise and reduce the incentive of companies to resort to unethical production sources. On a more operational side, adopting the tracking of goods via blockchain allows the data to be accessible by everyone in the process, and can be practical for inventory management. For example, whenever the inventory level is falling below a predetermined threshold, the order to the producer can be automatically triggered, and the delivery organized. For a review on this topic, the reader may refer to [59, 72].

In another sector, blockchain technology appears in the gaming industry [17]. For

example, let us examine the concept of NFTs (non-fungible tokens). They can be used as an in-game asset, traded or exchanged with other players. The ownership of such objects is registered in a blockchain, allowing transparency and enabling players to certify the ownership of their virtual assets. From a security perspective, the use of blockchain prevents games and their mechanisms from being copied or hacked.

1.3.2 Cryptocurrencies

Blockchain as a core technology for the emergence of a cryptocurrency was first proposed by Nakomoto in 2008 [109]. Cryptocurrencies are digital currencies, without a paper equivalent. They are decentralized and not emitted by a central authority (central bank), hence they are not controlled by any governments or official financial institutions. All the transactions executed by their means are anonymous and verified by a peer-to-peer network, and once they are validated, they cannot be unwritten, so they are irreversible. As a consequence, a user cannot betray its counterparty by sending virtual funds and cancelling the transaction once it has already entered the block.

The *crypto-market* is a fast evolving world. As of November 2023, the total market capitalization of all crytocurrencies is around $ 1'391 billion. As it can be seen in Figure 1.5, Bitcoin was - and still is - the dominant currency in this market, as the first one ever introduced [15], although in recent years, the development of alternatives, like Ethereum, Cardano, or Dogecoin has risen.

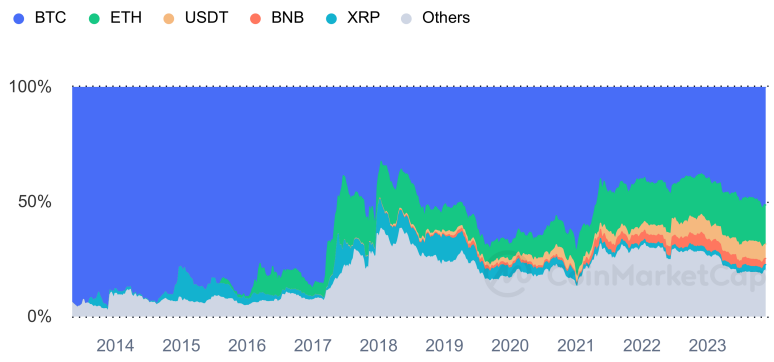

Figure 1.5: Percentage of total market capitalization of cryptocurrencies across time. Taken from `https://coinmarketcap.com/charts/`, last accessed on November 19, 2023.

From now on, let us focus on the Bitcoin currency. Its consensus protocol is the Proof-of-Work. It means that every person, called miner, must solve a cryptopuzzle. After someone succeeds in doing so, a block of transactions is appended to the blockchain. In exchange, the miner obtains a reward in Bitcoins, set at 6.25 Bitcoins as of November 2023. In other words, the winning miner gets the authorization from the community to choose which transactions will fill the new block. On average, one block is mined every ten minutes. To maintain this rate, the complexity of the cryptoproblem is adjusted every 2016 blocks (approximately every two weeks).

Blocks in the Bitcoin blockchain are public, which means that anyone can access the information about those blocks. In the same spirit, anyone can become a miner and therefore take part in the validation process. The usual information one can find in a block contains the list of transactions, the amounts, the timestamps and attached transaction fees. As an illustration, in Figure 1.6, the reader will observe for the block number 781'745 all information mentioned above and in particular the hash. Let us enter in more detail on this latter point.

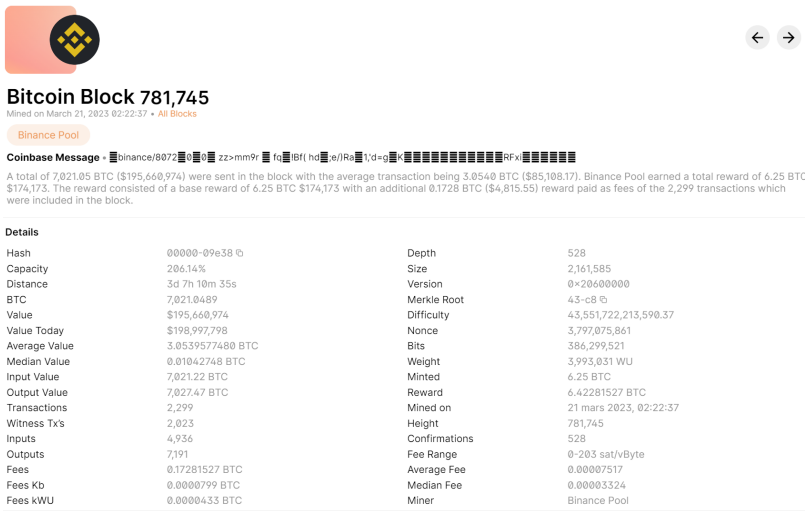

Figure 1.6: Overview of the Bitcoin block 781'745 as seen by the public. Taken from `https://www.blockchain.com/explorer/blocks/btc/781745`, last accessed on November 21, 2023.

A hash is a unique string of characters generated by a specific cryptographic function from input data of any size. In the Bitcoin context, this data typically includes transaction information. The hash function used in Bitcoin is known as SHA-256, which stands for Secure Hash Algorithm 256-bit, see e.g. [114]. This algorithm takes input data and produces a fixed-length, 256-bit hash, which appears as a sequence of 64 hexadecimal characters. To illustrate, consider the block 781'745 introduced in Figure 1.6. The pictured hash of this block is given by `00000000000000000003b630b4a574e1ded1af5495714aac7300d9e169909e38`. Despite the complexity and randomness of this string, it is deterministically generated: the same input data will always produce the same hash. However, even a minor change in the input, such as altering the amount of the transaction, or the timestamp, would result in a completely different hash.

In Bitcoin mining, the primary goal is to find a hash that is below a pre-set target. This target is what defines Bitcoin's difficulty level: the smaller the target, the rarer and more difficult it is to find a hash that meets this criterion. The process of finding such a hash involves a parameter called a nonce. A nonce, short for "number only used once", is a variable that miners alter repeatedly to change the hash outcome. Essentially, miners engage in a process of trial and error, altering the nonce and recalculating the hash until they find one that is less than the target [15]. The nonce is also part of a block information as can be seen in Figure 1.6. This mining process

is computationally intensive and consumes tremendous amounts of electricity. As a result, Bitcoin mining on a global scale has a profound energy footprint. To put it into perspective, the annual electricity consumption for Bitcoin mining worldwide is around 157.21 TWh[2] and this number is comparable to the annual electricity consumption of a country like Poland[3], a country with a population of 38 mio. inhabitants. This high energy demand raises environmental concerns and drives ongoing discussions about the sustainability of such cryptocurrencies. For instance, in [45], De Vries et al. investigate the Bitcoin's mining carbon footprint. In another work, De Vries and Stoll [46] explore the problem of electronic waste left by quickly obsolete mining equipment. Wendl et al. [151] provide an extensive review of the environmental impact of cryptocurrencies. The immense power consumption is an inherent aspect of ensuring the security and integrity of the Proof-of-work mining process, which remains for the moment the Bitcoin's consensus protocol.

1.3.3 Mining pools

Often, miners join together in mining pools. These pools are groups of miners who combine their computational resources to increase their chances of successfully mining Bitcoin blocks. Upon success, the rewards are distributed among pool participants in proportion to their contributed computational power, following predefined rules among the pool. As of December 2023, the Bitcoin mining landscape is dominated by several key players. According to data from `https://mempool.space/graphs/mining/pools`, seven major pools hold more than 85% of the pool market share.

The leader of the market with a market share of approximately 25% is AntPool. Founded in 2014 by Bitmain, a Chinese company, leading designer of chips for Bitcoin mining, it is one of the world's oldest mining pools. Besides Bitcoin, AntPool also supports a variety of other cryptocurrencies, including Bitcoin Cash, Litecoin, Ethereum Classic, Ethereum PoW, Dash, among others[4]. AntPool offers miners two types of reward methods: Full Pay Per Share (FPPS) and Pay Per Last N Shares (PPLNS). The FPPS method provides a payout for each share of computational power contributed by the miner, regardless of whether a block is mined or not. This method ensures a steady and predictable income stream for miners but includes a fee of 4% to cover the risk taken on by the pool, since the transaction fees are also distributed. On the other hand, the PPLNS method only pays out when a block is successfully mined, and the reward is divided among the miners who contributed to the last shares. This method can potentially be more profitable and no fees are retained by the pool but also carries more risk, as payouts are dependent on the pool's success in mining blocks[5].

The second biggest pool in the market is Foundry USA, with a market share of approximately 25% as of December 2023. It was founded in 2020 by a US

[2] `https://ccaf.io/cbeci/index`, last accessed on November 21, 2023.

[3] `https://yearbook.enerdata.net/electricity/electricity-domestic-consumption-data.html`, last accessed on November 21, 2023.

[4] `https://coinmarketcap.com/academy/glossary/antpool`, last accesses on 30/12/2023.

[5] `https://antpoolsupport-hc.zendesk.com/hc/en-us/articles/5983010227993-Miners-Settings-Fees`, last accessed on 30/12/2023.

based company active in cryptocurrency investment as a response to a largely Asian dominated market[6]. This pool focuses only on Bitcoin and Bitcoin Cash cryptocurrencies. Foundry operates under the FPPS rewards system. For the transaction fees, the payout rate is calculated as the mean percentage of transactions fees per block in the last elapsed mining day. The formulas for the calculation of the payouts can be found at `https://pool-faq.foundrydigital.com/what-is-foundry-usa-pools-payout-methodology`[7]. There are no pool management fees retained by the pool[8].

Very scarce information can be found on pools going bankrupt. Indeed, the visible information for a external observer is noting that pools stop mining new blocks, in which situation one can suspect a pool having gone out of business. The reasons of this are often unknown. In the early stages of Bitcoin history, we can cite the example of the Galaxy mining pool, launched in July 2013, we can see how it shut down 40 days later due to insufficient hash rate[9]. It is not reported that the pool has filed for official bankruptcy and one can infer from the aforementioned forum thread that the pool manager took preventive measures before collapsing.

In recent years, several notable examples of big mining companies going bankrupt can be found. For instance in 2022, three big mining companies operating in the US have filed for bankruptcy under Chapter 11[10]. Some explanations of these events may be the increase in electricity price due to an unstable geopolitical situation and a drop in the Bitcoin price in the second half of the year.

In the same year, Poolin, one of the top 10 largest pools in the world, announced facing liquidity problems[11]. It led to users experiencing issues with payout withdrawals. Reasons of these problems were not publicly announced but there are speculations about the pools involvement in yield farming, in other words lending cryptoholdings to earn returns on investment. As of December 2023, Pooling has lost more than 90% of its market share but still operates and is still part of the top 10 biggest pools. These events highlight the risks and challenges faced by Bitcoin mining pools, especially in volatile market conditions and when engaging in high-risk financial activities like yield farming.

[6] `https://foundrydigital.com/mining-service/foundry-usa-pool/`, last accessed on 30/12/2023.

[7] Last accessed on 30/12/2023.

[8] `https://www.hedgewithcrypto.com/best-bitcoin-mining-pools/`, last accessed on 30/12/2023.

[9] `https://bitcointalk.org/index.php?topic=248031.0`, last accessed on 30/12/2023.

[10] `https://ezblockchain.net/article/bitcoin-mining-companies-who-already-declared-bankruptcy-and-why`, last accessed on 30/12/2023.

[11] `https://www.pymnts.com/cryptocurrency/2022/crypto-insolvencies-spread-as-top-bitcoin-mining-pool-halts-withdrawals/`, last accessed on 30/12/2023.

1.4 Main contributions of this thesis

This thesis contains results in the area of ruin theory applied to cryptocurrency mining problems, as well as insurance economics with respect to risk classification. A recurrent and uniting theme throughout the thesis is to study a quantitative trade-off between risk and profitability. The results of Chapter 2 have already been published in *Insurance: Mathematics and Economics* and the content of Chapters 3 and 4 are submitted for publication.

In Chapter 2, we explore the resource-intensive process of mining blocks in blockchain technology, especially under a *Proof-of-Work* consensus mechanism. This method of mining can lead to a financial risk known as ruin, where the costs of mining operations surpass the gains obtained from mining rewards. Our focus is on the potential benefits for a miner of joining a mining pool, which can help reduce the variability of returns for them at a specific cost of joining. We are able to draw analogies to reinsurance in this context. Using concepts from ruin theory and risk-sharing principles in the field of insurance, we conduct a quantitative analysis to assess the impact of pooling in this domain. Via this approach, we derive a series of explicit formulas to quantify key aspects of interest, such as the expected value of the profit and the ruin probability. To demonstrate the practical applicability of our findings, we present numerical examples using parameters that hold in real-world scenarios.

After a more theoretically oriented view on the topic of ruin in blockchain mining activities, in Chapter 3 we turn to an empirical assessment of our previous findings. Specifically, we control whether the calculated probabilities of ruin and the expected future surplus, derived under certain model assumptions, hold true in practical scenarios. A key addition to our analysis is the inclusion of transaction fees within block rewards. Additionally, we introduce algorithms designed to fit generalized hyper-exponential distributions to real-world data. Furthermore, we conduct a sensitivity analysis to evaluate various factors of interest in the mining process. This includes assessing the impact of incorporating temporal dependencies and transaction fees into the models, in contrast to assuming i.i.d. mining rewards.

Finally, Chapter 4 shifts to a different aspect of the insurance world: risk classification. Here, pricing strategies often depend on individual criteria of policyholders. Insurers commonly employ various classification algorithms to sort policyholders into different risk categories. However, this method is not without its flaws, as errors in classification can occur. With the development of new data science technologies, the limitations of these risk classification systems are becoming more evident in practical applications. In this chapter, we explore the consequences of such misclassifications. We differentiate between two types of risks and establish a framework to quantitatively assess the problem of an insurer setting optimal premiums. This includes an analysis of the balance between costs and profits, when the probabilities of classification errors are known. In particular, we derive Markowitz-type efficient frontiers for identifying optimal trade-offs between risk and profitability. We also provide several extensions of the original model and give a number of numerical illustrations.

Chapter 2

Blockchain mining in pools: Analyzing the trade-off between profitability and ruin

This chapter is based on the following article:

H. Albrecher, D. Finger, and P.-O. Goffard. Blockchain mining in pools: Analyzing the trade-off between profitability and ruin. Published in *Insurance: Mathematics and Economics*, 105, 313-335 [8].

Abstract. The resource-consuming mining of blocks on a blockchain equipped with a *Proof-of-Work* consensus protocol bears the risk of ruin, namely when the operational costs for the mining exceed the received rewards. In this chapter we investigate to what extent it is of interest to join a mining pool that reduces the variance of the return of a miner for a specified cost for participation. Using methodology from ruin theory and risk sharing in insurance, we quantitatively study the effects of pooling in this context and derive several explicit formulas for quantities of interest. The results are illustrated in numerical examples for parameters of practical relevance.

2.1 Introduction

A blockchain is a decentralized data ledger maintained by a *Peer-to-Peer* network. Blockchain users issue transactions to the network peers who agree on those to be recorded by following a consensus protocol. In public and permissionless blockchains, such as the one for Bitcoin, the consensus protocol is called *Proof-of-Work* (PoW). The nodes, referred to as miners, compete to solve a challenging cryptographic puzzle using some brute force search algorithm. The first miner to come up with a solution includes the pending transactions in a block and is rewarded with newly minted crypto-coins. This reward compensates the operational cost of mining mainly induced by the consumption of electricity. The PoW protocol is designed to be incentive compatible in the sense that a miner is compensated proportionally to her computational effort. When the Peer-to-Peer network grows large, the share of the network computing power owned by a given miner shrinks, which in turn makes the gains infrequent. The constant operating costs therefore endanger the solvency of miners and has led them to join forces by forming mining pools.

Mining pools grant miners a steadier income, as block finding rewards are collected more often. The earnings are then fairly distributed to the pool participants based on their contribution to the computational effort. The simplest way to do so consists in splitting the reward whenever a block is found. This is the *proportional* reward system. More sophisticated reward schemes have been put together to increase the amount of risk transferred from the miners to the pool and to fill the gaps of the proportional system that we will discuss later. These more sophisticated systems require the supervision of a manager who undertakes part of the risk in exchange for a commission. An early work of Rosenfeld [121] provides a detailed overview on mining pool reward systems, see also the recent survey of Zhu et al. [164]. In practice, the individual contribution of a miner is measured through a *share* submission process. A *share* refers to an easier-to-find 'fake' solution to the crypto-puzzle that miners must send to the pool manager to prove their involvement (for instance, a solution to the crypto-puzzle with only m instead of the $n > m$ leading zeroes required for the successful mining of a block). In this work we provide a risk analysis of *Pay-per-Share* (PPS) reward systems in which the pool manager pays for each share submitted by the miners. In that way the manager takes on much of the randomness associated to the mining activity, which is therefore very appealing to the participant. Using utility theory, Schrijvers et al. [126] showed that such systems are incentive-compatible for risk-averse miners. Both Rosenfeld [121] and Zhu et al. [164] stressed that a scheme of this kind must go hand in hand with a proper capital allocation strategy on the part of the manager to avoid ruin. Reliable information on mining companies filing for bankruptcy are hard to come by, we are able to provide one. Bcause Mining, LLC, the mining affiliate operating in Virginia Beach, Virginia, filed voluntarily for bankruptcy court protection under Chapter 11 on April 11, 2019 in order to reorganize its debts. Unsecured creditors included a Virginia power company ($\$1,459,267.38$), the U.S. Customs and Border Patrol ($\$737,041.90$), landlords, and staff[1].

The aim of this chapter is to provide risk-analytic tools to inform the decision

[1]Source: https://www.theblockcrypto.com

making process of miners and pool managers. This is achieved by taking an approach inspired from insurance risk theory. The wealth of miners and pool managers is modelled via stochastic processes that take into account operational costs, pool participation fees and block finding rewards. The resulting processes are similar to those appearing in the surplus modelling for insurance companies which collect premiums continuously and have to pay loss reimbursements to policyholders in case a claim occurs. A standard risk measure in this context is the ruin probability defined as the probability that the wealth process falls below zero, see e.g. Asmussen and Albrecher [16] for an overview. This analogy was already used in Albrecher and Goffard [10], where the opportunity for miners to deviate from the prescribed protocol by withholding blocks was investigated. A first result was also obtained there in relation to the advantage of joining a mining pool which applies the proportional system. Our objective in this chapter is to considerably extend this line of thinking towards the *Pay-per-Share* redistribution systems that are more commonly used in practice. We will also consider a variant of the model in which the collected rewards are random variables. This assumption will enable the application of classical results from double-sided jumps in a risk reserve process for modelling insurance portfolios, see for example Albrecher at al. [9], Labbé and Sendova [91]. Incorporating random rewards allows us to account for the transactions fees and the exchange rate of cryptocurrencies to fiat ones. Transaction fees are included by blockchain users to entice the network to process their transactions, see Easley al. [60] and Kasahara and Kawahara [85]. The redistribution of the revenue generated by the transaction fees among the pool participants also varies from one mining pool to another. Closed-form expressions for the probability of ruin and the expected profit given that ruin has not occurred are provided up to an exponentially distributed time horizon. These formulas are amenable for a quick numerical evaluation to perform a sensitivity analysis of risk and reward indicators with respect to the model parameters. We believe that our results will be useful for miners and pool managers to make the right financial choices. Our indicators can also serve as the basis for a potential future regulatory framework for mining activity on blockchains equipped with the Proof-of-work consensus protocol.

A major concern associated to mining pool formation is the centralization of the network. Cong et al. [37] have explained that miners who direct their mining power to multiple small mining pools enjoy the same risk sharing benefits as miners that choose to join a single mining pool. Hence the intuition that a larger mining pool would grow even larger is misguided. Empirical data shows that the participation fees are greater in larger mining pools, which naturally slows down their growth. We aim at providing more insight on the risk of centralization in the light of our analysis.

The remainder of the chapter is organized as follows. Section 2.2 gives a brief description of the mining process in blockchains equipped with Proof-of-Work. Section 2.3 provides an overview of the existing reward systems and describes the *Pay-Per-Share* mechanism in more detail, as it will be the focus later on. Formulas for the ruin probability and expected surplus for the pool manager are derived for deterministic rewards in Section 2.4 and for randomized reward in Section 2.5. Section 2.6 provides formulas from the individual miner's perspective. Section 2.7 is devoted to numerical illustrations where the sensitivity of the risk and performance indicator

is analysed with respect to the model parameters. Section 2.8 concludes.

2.2 Mining blocks in a Proof-of-Work powered blockchain

A block consists of a header and a list of "transactions" that represents the information recorded through the blockchain. The header usually includes the date and time of creation of the block, the block height which is the index inside the blockchain, the hash of the block and the hash of the previous block. The hash of a block is obtained by concatenating the header and the transactions in a large character string thus forming a "message", to which a hash function is applied. A hash function is a function that can map data of arbitrary size to fixed-sized values. The hash functions used in blockchain applications must be cryptographic, i.e. quick to compute, one way and deterministic. It must be nearly infeasible to generate a message with a given hash value or to find two messages with the same hash value. A small change in the message should change dramatically the hash value so that the new hash value appears to be uncorrelated to the previous hash. We will not expand on how to build such a cryptographic hash function, we refer the interested reader to the work of Al-Kuwari et al. [2]. In the bitcoin blockchain as well as in many other applications, the standard is the SHA-256 function which converts any message into a hash value of 256 bits. The latter is usually translated into a hexadecimal digest, for instance the hash value of the title of the present chapter reads as

`98b1146926548f6b57c4347457713ff2f035beda9c93f12fbc9b202e9c512e80.`

The information recorded in a public blockchain may be retrieved by anyone and can be accessed through a blockchain explorer such as blockchain.com, the content of the block of height #724724 may be viewed through the following link block content. Mining a block means finding a block hash value lower than some target which can only be achieved by brute force search thanks to the properties of cryptographic hash functions. In practice, the search for an appropriate hash value, referred to as a solution, is done by appending a nonce to the block message before applying the hash function. A nonce is a 32 bits number, drawn at random by miners until a nonce resulting in a proper block hash value is found. For illustration, consider the block in Figure 2.1.

```
Block Hash: 1fc23a429aa5aaf04d17e9057e03371f59ac8823b1441798940837fa2e318aaa
Block Height: 0
Time:2022-02-25 12:42:04.560217
Nonce:0
Block data: [{'sender': 'Coinbase', 'recipient': 'Satoshi', 'amount': 100, 'fee': 0}, {'sen
der': 'Satoshi', 'recipient': 'Pierre-O', 'amount': 5, 'fee': 2}]
Previous block hash: 0
Mined: False
--------------
```

Figure 2.1: A block that has not been mined yet.

The hash value in decimal notation is $1.43e^{76}$ while the maximum value for a 256 bits number is $2^{256} - 1 \approx 1.16e^{77}$. We refer to the latter as the maximal target and denote it by $T_{\max}$. The Proof-of-Work protocol sets a target $T < T_{\max}$ and

ask miners to find a nonce such that the hash value of the block is smaller than T. Practitioners would rather talk about the *difficulty* which is defined as $D = T_{\max}/T$. If the difficulty is one, any hash value is acceptable. Increasing the difficulty reduces the set of allowable hash values, making the problem harder to solve. A hash value is then called *acceptable* if its hexadecimal digest starts with a given number of zeros. If we set the difficulty to 2^4, then the hexadecimal digest of the hash of the block must start with at least 1 leading zero, making the hash value of the block in Figure 2.1 not acceptable. After completing the nonce search we get the block in Figure 2.2. Note that it took 5 attempts to find this nonce. The number of needed

```
Block Hash: 0869032ad6b3e5b86a53f9dded5f7b09ab93b24cd5a79c1d8c81b0b3e748d226
Block Height: 0
Time:2022-02-25 13:41:48.039980
Nonce:2931734429
Block data: [{'sender': 'Coinbase', 'recipient': 'Satoshi', 'amount': 100, 'fee': 0}, {'sen
der': 'Satoshi', 'recipient': 'Pierre-O', 'amount': 5, 'fee': 2}]
Previous block hash: 0
Mined: True
--------------
```

Figure 2.2: A mined block with a hash value having on leading zero.

trials is geometrically distributed with parameter $1/D$, which means that with a difficulty of $D = 2^4$ it takes on average 16 trials. The protocol adjusts the difficulty automatically every $2,016$ block discoveries so as to (globally) maintain one block discovery every 10 minutes on average. The time between two block discoveries depends on the number of hash values computed by the network at a given instant. As of February 2022, the network computes 182.58 Exahashes per second and the difficulty is $27,967,152,532,434$.[2] For an exhaustive overview of the mining process in the bitcoin blockchain, we refer the reader to the book of Antonopoulos [15, Chapter 10]. As each trial (of the system) for mining a block is independent of the others and leads to a success with very small probability, the overall number of successes is binomially distributed and will be very well approximated by a Poisson random variable. This justifies the Poisson process assumption made in the sequel to model the block arrival and the reward collecting processes. Empirical studies of the block inter-arrival times data tend to confirm this hypothesis, see the work of Bowden et al. [27].

2.3 Risk models and reward systems

A risk model defines the wealth of some company or individual as a stochastic process

$$R_t = u - C_t + B_t,\ t \geq 0$$

which corresponds to the income $(B_t)_{t\geq 0}$ net of the expenses $(C_t)_{t\geq 0}$. The surplus process $(R_t)_{t\geq 0}$ starts at some initial level $R_0 = u > 0$. We take a continuous time approach where $t \in \mathbb{R}_+$, and $(C_t)_{t\geq 0}$ and $(B_t)_{t\geq 0}$ define increasing functions or stochastic processes. A risk analysis is relevant only if at least one of the model components is random. The activity of the company is profitable if on average

[2]Source: bitcoinblockhalf.com

the earnings exceed the expenses, namely $\mathbb{E}(B_t) > \mathbb{E}(C_t)$. Even if the net profit condition holds, the variability of the process $(R_t)_{t\geq 0}$ can lead to bankruptcy as it may become negative. Define the ruin time as

$$\tau_u = \inf\{t \geq 0 : \ R_t < 0\},$$

which corresponds to the first time at which the surplus goes below 0. The risk of bankruptcy is classically assessed by computing the ruin probability up to time $t \geq 0$ defined by

$$\psi(u,t) = \mathbb{P}(\tau_u \leq t). \tag{2.1}$$

It is sometimes more convenient from a mathematical point of view to consider the infinite-time horizon by letting $t \to \infty$, and in that case we write $\psi(u) := \lim_{t\to\infty} \psi(u,t)$. Following the rationale developed in [10], we also consider a performance indicator defined as

$$V(u,t) = \mathbb{E}(R_t \mathbb{I}_{\tau_u > t}), \tag{2.2}$$

which corresponds to the expected surplus at time $t \geq 0$ in case ruin did not occur until then.

Consider a network of n miners, where miner $i \in \{1, \ldots, n\}$ owns a share $p_i \in (0,1)$ of the network hashpower, i.e. $\sum_{i=1}^n p_i = 1$. If the number of blocks found by the network is governed by a homogeneous Poisson process $(N_t)_{t\geq 0}$ with intensity λ, then the number of blocks found by miner i is a (thinned) Poisson process $(N_t^i)_{t\geq 0}$ with intensity $p_i \cdot \lambda$. Denote by $b > 0$ the amount of the reward for finding a new block and assume that the cumulative operational cost is a linear function with slope $c_i > 0$ which depends on the price of the electricity and the computing power of miner i. The surplus process of miner i is then given by

$$R_t^i = u - c_i \cdot t + N_t^i \cdot b, \ t \geq 0. \tag{2.3}$$

Model (2.3) has been considered by Albrecher and Goffard [10], and formulas for both the finite-time ruin probability (2.1) and the expected surplus (2.2) were derived. To make the formulas more amenable for numerical evaluation, the authors then decided to approximate the fixed time horizon $t \geq 0$ by an exponential random variable $T \sim \text{Exp}(t)$ with mean $t \geq 0$, resulting in tractable expressions for

$$\widehat{\psi}(u,t) = \mathbb{E}[\psi(u,T)], \text{ and } \widehat{V}(u,t) := \mathbb{E}[V(u,T)], \tag{2.4}$$

which were then used to carry out a numerical analysis.

Remark 2.3.1. *Model* (2.3) *assumes that the block-finding reward is constant, while the bitcoin protocol stipulates a halving of the reward every* $210,000$ *blocks. However, first note that* $210,000$ *blocks take* 4 *years to be found which is greater than the time horizon we have in mind when defining* $\widehat{\psi}(u,t)$ *and* $\widehat{V}(u,t)$*. Second, since these halving dates are known to blockchain networks, the market automatically adjusts the cryptocurrency exchange rate before the halving occurs. Eventually, one could also replace* b *in Model* (2.3) *with a piecewise constant function to account for halving events. However, this would only be relevant for very long time horizons and will not be pursued in the present work.*

Consider now a situation where a subset of miners $I \subset \{1, \ldots, n\}$ decides to gather in a mining pool. The cumulated hashpower of this pool is then

$$p_I = \sum_{i \in I} p_i,$$

and the arrival rate of block rewards for a given miner i rises from $p_i \cdot \lambda$ to $p_I \cdot \lambda$. Because the reward is shared among the pool participants, the size of the reward collected by miner i decreases from b to $p_i \cdot b$. The expected surplus is the same when mining solo and mining for a pool, but the variance (and therefore the risk) is smaller when mining for a pool. The management of a mining pool relies heavily on the reward distribution mechanism set up by a pool manager. For the redistribution system to be fair, each miner must be remunerated in proportion to her calculation effort. Miner i must earn a share p_i/p_I of the mining pool total income. The pool manager has to find a way to estimate the contribution of each pool participant. This is done by submitting *shares* which are partial solutions to the cryptopuzzle easier to find than the actual solution. Recall from Section 2.2 that a proper solution corresponds to a hash value starting with a given number of zeros, so *shares* are hash values with a smaller number of leading zeros. If the current difficulty of the cryptopuzzle is D, then the difficulty for finding a *share* is set to $q \cdot D$ by the pool manager, where $q \in (0, 1)$. The manager's cut is a fraction $f \in (0, 1)$ of the block discovery reward b. We start by presenting the proportional reward system in Section 2.3.1.

2.3.1 The proportional reward system

The proportional reward system splits time in *rounds* which correspond to the time elapsed between two block discoveries. During these *rounds*, the miners submit *shares*. The ratio of the number of *shares* submitted by miner i over the total number of *shares* submitted by her fellow mining pool participants determines her share of the reward and should converge to her share of the mining pool computing power, that is p_i/p_I (for sufficiently low complexity of the shares, the latter limit will be a very good approximation for the actual situation indeed). The surplus of miner i is then given

$$R_t^i = u - c_i \cdot t + N_t^I \cdot (1 - f) \cdot \frac{p_i}{p_I} \cdot b, \ t \geq 0, \tag{2.5}$$

where (N_t^I) is a Poisson proccess of intensity $p_I \cdot \lambda$ that gives the number of blocks appended to the blockchain by the mining pool. The duration of a *round* is exponentially distributed $\mathrm{Exp}\left[(p_I\lambda)^{-1}\right]$. The uncertainty on the length of the round has undesirable consequences on the time value of the *shares* submitted by the miners. Indeed, if n shares are submitted during a round, then the value of a given *share* is $(1 - f) \cdot b/n$. The longer a *round* lasts, the greater the value of n is. The *shares* are worth less in longer rounds which triggers an exodus behavior of miners toward mining pools with shorter rounds. This phenomenon, called pool hopping, has been documented in the early work of Rosenfeld [121]. Yet another drawback is that a miner that has found a full solution may delay the submission until her ratio of *shares* submitted reflects her fraction of the mining pool computing power. The proportional system is not *incentive-compatible* using the terminology of Schrijvers

et al. [126]. A discounting factor may be applied to compensate the decreasing value of shares over time, see for instance the slush's method [130].

Our work is also concerned about the risk undertaken by pool managers. Within the frame of the proportional reward system, the surplus of the pool manager is given by

$$R_t^I = u + N_t^I \cdot f \cdot b,\ t \geq 0. \tag{2.6}$$

Model (2.6) does not account for any mining pool operating cost. The mining costs are entirely borne by miners and the mining pool manager only serves as coordinator. A proportional-type reward system should therefore lead to a low management fee f.

Although this system provides fairness, it has weaknesses that justify the introduction of a more sophisticated distribution mechanism. In particular, if miners seek to actually transfer some of the risk associated to the mining activity to the pool manager, then they should rather turn to a mining pool based on a *Pay-per-Share* system, which is the focus of this chapter and introduced in the next section.

2.3.2 The Pay-Per-Share reward system

In a *Pay-per-Share* reward system, the pool manager immediately rewards the miners for each *share* submitted. Let $(M_t)_{t\geq 0}$ be a Poisson process of intensity μ that counts the number of *shares* submitted by the entire network of miners up to time $t \geq 0$. Denote by $q \in (0,1)$ the relative difficulty of finding a block compared to finding a share. Let $0 < w < b$ be the reward for finding a *share*. The number of *shares* submitted by miner i is then a (thinned) Poisson process $(M_t^i)_{t\geq 0}$ with intensity $p_i \cdot \mu$, p_i being the share of the individual miner's network hashpower as defined above, and her surplus when joining a PPS mining pool becomes

$$R_t^i = u - c_i \cdot t + M_t^i \cdot w,\ t \geq 0. \tag{2.7}$$

The intensities of the processes $(N_t)_{t\geq 0}$ and $(M_t)_{t\geq 0}$ are linked through $\lambda = q \cdot \mu$. By setting $w = (1-f) \cdot b \cdot q$, we observe that the surplus (2.5) and (2.7) have the same expectation at time t, but the variance and therefore the risk associated to (2.7) is lower. This reward system has been shown to be resistant to pool hopping and is incentive compatible. It also entails a significant transfer of risk to the pool manager whose surplus process is now given by

$$R_t^I = u - M_t^I \cdot w + N_t^I \cdot b,\ t \geq 0, \tag{2.8}$$

making her subject to the risk of bankruptcy.

Remark 2.3.2. *Since the process $(M_t^I)_{t\geq 0}$ requires solving for a problem of lower complexity than $(N_t^I)_{t\geq 0}$, $(N_t^I)_{t\geq 0}$ is a subset of the path defined by the process $(M_t^I)_{t\geq 0}$. It means that both processes are not independent. Concretely, at the moment of the block reward payment b, at the same time there is a realisation of the miners' reward w. As we sometimes will need to isolate downward jumps without the simultaneous upward jump point, we define another process with a reduced intensity. We apply the superposition theorem (see e.g. [88]) to the Poisson process M_t^I by redefining the down jump process as $(M_t^{I,d})_{t\geq 0} \sim Poisson(\mu_d)$, where $\mu_d = \mu - \lambda$.*

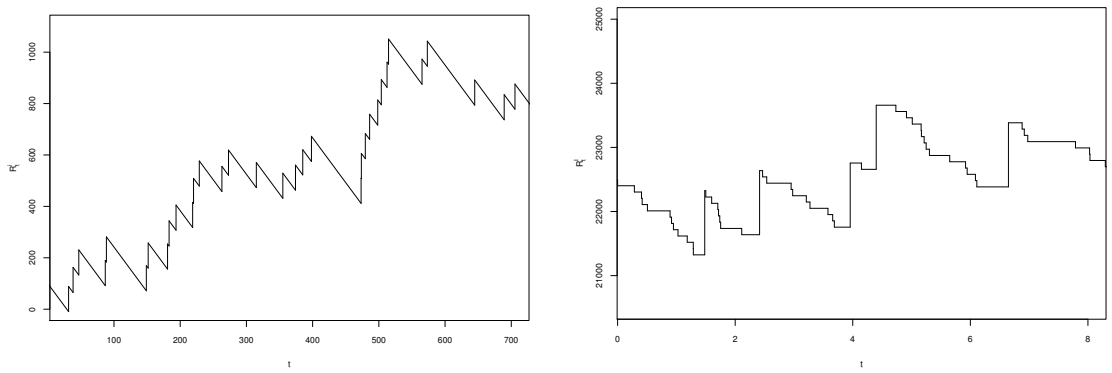

(a) Individual miner (b) Pool manager

Figure 2.3: Illustration of surplus paths for the pool members and the pool manager.

Figure 2.3 represents sample paths of the surplus processes for an individual miner and the pool manager.

In addition to the bounty for finding a new block, blockchain users usually include a small financial incentive for the network to process their transaction. These transaction fees (e.g. referred to as *gas* within the ETHEREUM blockchain), are known to be variable as they highly depend on the network congestion at a given time. Note also that since the operational cost is paid by miners using a fiat currency, it would be more accurate to account for the exchange rate of the cryptocurrency to some fiat currency. We can therefore model the successive rewards for *shares* and blocks as sequences of nonnegative random variables denoted by $(W_k)_{k\geq 1}$ and $(B_k)_{k\geq 1}$ respectively, which for simplicity we will both assume to be *i.i.d.* in this chapter. A reward system that features a *Pay-per-Share* mechanism and includes in the miners' reward the transaction fees is referred to as a *Full Pay-per-Share* reward system by practitioners. The surplus of miner i in a mining pool applying the FPPS system is given by

$$R_t^i = u - c_i \cdot t + \sum_{k=1}^{M_t^i} W_k, \ t \geq 0, \tag{2.9}$$

and the surplus of the pool manager then becomes

$$R_t^I = u - \sum_{k=1}^{M_t^I} W_k + \sum_{l=1}^{N_t^I} B_l, \ t \geq 0. \tag{2.10}$$

In the following sections, we will now derive formulas for the ruin probability and expected surplus in case ruin did not occur up to a given time horizon for the models discussed above.

2.4 Pool analysis with deterministic rewards

We start with a fixed time horizon. For simplicity, we drop the superscript I in the following developments.

2.4.1 Deterministic time horizon

For the pool manager's side, we first define some measures of interest. Let $\tau = \inf\{t \geq 0 : R_t < 0\}$ be the time of ruin of the pool manager, i.e. the first time his surplus reaches 0. The corresponding ruin probabilities in finite and infinite horizon respectively are given by

$$\psi(u,t) = \mathbb{P}(\tau \leq t), \text{ and } \psi(u) = \mathbb{P}(\tau < \infty). \tag{2.11}$$

The net profit condition in this case translates to $\lambda b > \mu w$. It implies from [16], that $\psi(u) < 1$. We also define the expected surplus at time t given that ruin has not occurred up to time t:

$$V(u,t) = \mathbb{E}(R_t \mathbb{I}_{\tau > t}). \tag{2.12}$$

In the sequel, we will use the process $(M_t^d)_{t\geq 0}$ defined in Remark 2.3.2 representing the pure downward jumps. Note that ruin can only occur at discrete times when the process $(M_t^d)_{t\geq 0}$ admits a jump. We can rewrite the ruin time τ as

$$\tau = \inf\{t \geq 0; M_t^d w > u + N_t(b-w)\} = \inf\{t \geq 0; M_t^d > u/w + N_t(b-w)/w\}. \tag{2.13}$$

Equivalently, we can rewrite it as

$$\tau = \inf\{t \geq 0; M_t^d w/(b-w) > u/(b-w) + N_t\} \tag{2.14}$$

to isolate the $(N_t)_{t\geq 0}$ process with unit jumps. The study of the p.d.f. f_τ of τ is analogous to the derivations in [68]. Since $(N_t)_{t\geq 0}$ and $(M_t^d)_{t\geq 0}$ are Poisson process, they enjoy the order statistic property. That is, given that $N_t = n$, the jump times $\{T_1, \ldots, T_n\}$ of the process N_t have the same distribution as the order statistics vector of a random variable having distribution $F_t(s) = s/t,\ 0 \leq s \leq t$. Further, let $\{S_n^d, n \in \mathbb{N}\}$ be the sequence of arrival times associated with the process $(M_t^d)_{t\geq 0}$. Its distribution function is denoted by $F_{S_n^d}(t)$ and its p.d.f. by $f_{S_n^d}(t)$. Denote by $\lceil x \rceil$ the ceiling function. Following Corollary 1 from [68], we proceed from Equation (2.14) and derive the next steps.

Theorem 2.4.1. *Let $(N_t,\ t \geq 0)$ and $(M_t^d,\ t \geq 0)$ be Poisson processes with intensities $\{\lambda, \mu_d\}$ respectively, and assume that the net profit condition $\lambda b > \mu w$ holds, then the p.d.f. of τ is given by*

$$f_\tau(t) = \sum_{n=0}^{+\infty} \mathbb{E}\left[\frac{(-1)^n}{t^n} G_n\left[0 | S_{v_0}, \ldots, S_{v_{n-1}}^d\right] | S_{v_n}^d = t\right] f_{S_{v_n}^d}(t) \mathbb{P}\left[N_t = n\right], \tag{2.15}$$

where $(v_n)_{n\geq 0}$ is a sequence of integers defined as $v_n = \lceil n(b-w)/w + u/w \rceil$, $n \geq 0$, and $(G_n(\cdot|\{\ldots\})_{n\in\mathbb{N}}$ is the sequence of Abel-Gontcharov polynomials defined in Appendix 2.A.

The proof is delegated to Appendix 2.B. The expression of the ruin time p.d.f. (2.15) is not convenient for numerical purposes. The infinite series in (2.15) must be truncated, possibly to a high order to reach an acceptable level of accuracy. Also, the evaluation of the Abel-Gontcharov polynomials relies on the recurrence relationships (2.68) which are known to suffer from numerical instabilities. Moreover, the conditional expectation with respect to $\{S_{v_n}^d = t\}$ itself requires the use of Monte Carlo simulations. Finally, a similar algebraic expression for $V(u,t)$ is out of sight. In view of all these difficulties, we therefore propose as in [10] a workaround which consists of replacing the deterministic time horizon by a random variable with exponential distribution.

2.4.2 Exponential time horizon

To obtain a nicer solution to the problem, we now randomize the time horizon T. The practical intuition suggests that the time horizon is never fixed in advance and is subject to various external factors, such as bitcoin price fluctuations, in-pool events etc. We choose the time horizon T to be exponentially distributed with rate $1/t$ (so that $\mathbb{E}(T) = t$). This leads to computable expressions having an intuitive justification due to the lack of memory property of the exponential distribution. Let $\hat{V}(u,t) := \mathbb{E}(R_T \mathbb{I}_{\tau>T})$ denote the expected value of the surplus at the exponential time horizon T.

Theorem 2.4.2. *Let b and w, $b > w$, be fixed positive integers and assume that the net profit condition $\lambda b > \mu w$ holds. Then the expected surplus at an exponential time horizon can be expressed in the form*

$$\hat{V}(u,t) = \sum_{i=1}^{w} c_i x_i^u + u + \lambda b\, t - (\lambda + \mu_d) w\, t,$$

where $x_1, \ldots, x_w$ are the w roots inside the unit disk of the equation

$$\lambda x^b - (\lambda + \mu_d + 1/t) x^w + \mu_d = 0, \tag{2.16}$$

and the constants $c_1, \ldots, c_w$ are the solution of the linear equation system

$$\begin{pmatrix} \lambda x_1^{b-w} - (\lambda+\mu_d+1/t) & \cdots & \lambda x_w^{b-w} - (\lambda+\mu_d+1/t) \\ \lambda x_1^{b-w+1} - (\lambda+\mu_d+1/t)x_1 & \cdots & \lambda x_w^{b-w+1} - (\lambda+\mu_d+1/t)x_w \\ \vdots & \ddots & \vdots \\ \lambda x_1^{b-1} - (\lambda+\mu_d+1/t)x_1^{w-1} & \cdots & \lambda x_w^{b-1} - (\lambda+\mu_d+1/t)x_w^{w-1} \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_w \end{pmatrix} = \begin{pmatrix} A_1 \\ A_2 \\ \vdots \\ A_w \end{pmatrix}, \tag{2.17}$$

with

$$A_i = (i-1)\mu_d + \mu_d t(\lambda b - (\lambda+\mu_d)w) - \mu_d w, \quad i = 1, \ldots, w.$$

Proof. Akin to the approach in [9], consider some small $h > 0$ and condition on the following scenarios during the time interval $(0, h)$:

1. no jump and $T > h$;
2. no jump and $T \le h$;
3. occurrence of an upward jump;
4. occurrence of a downward jump.

All other combinations of these events have negligible probability in the limit $h \to 0$ that we will pursue below. One then obtains

$$\begin{aligned} \hat{V}(u,t) = {} & e^{-(\frac{1}{t}+\lambda+\mu_d)h}\hat{V}(u,t) + \frac{1}{t}\int_0^h e^{-s/t} e^{-(\lambda+\mu_d)s} u\, ds \\ & + \lambda \int_0^h e^{-\lambda s} e^{-(1/t+\mu_d)s} \hat{V}(u+b-w,t)\, ds + \mu_d \int_0^h e^{-\mu_d s} e^{-(1/t+\lambda)s} \hat{V}(u-w,t)\, ds. \end{aligned} \tag{2.18}$$

We now take the derivative w.r.t. h and set $h = 0$ to obtain

$$\lambda \hat{V}(u+b-w,t) - (\lambda + \mu_d + 1/t)\hat{V}(u,t) + \mu_d \hat{V}(u-w,t) + u/t = 0, \quad u \geq 0. \tag{2.19}$$

By definition of $\hat{V}(u,t)$ we have the boundary conditions $\hat{V}(u,t) = 0$ for all $u < 0$ and the linear boundedness $0 \leq \hat{V}(u,t) \leq u + (\lambda b - \mu_d w)t$ in both u and t for all $u, t \geq 0$.

Equation (2.19) is an inhomogeneous difference equation with constant coefficients (see e.g. [82] for solution methods), which has the general solution

$$\hat{V}(u,t) = \sum_{i=1}^{b} c_i x_i^u + d_0 + d_1 u$$

with constants $\{c_i\}_{i=1}^b, \{x_i\}_{i=1}^b, d_0, d_1$ still to be determined.
Let us start with the inhomogeneous part: plugging the ansatz $d_0 + d_1 u$ into (2.19) gives

$$d_1 = 1, \; d_0 = \lambda b t - (\lambda + \mu_d) w t.$$

For the homogeneous part, we consider the characteristic equation (2.16), which by the Fundamental Theorem of Algebra has exactly b complex roots $x_1, \ldots, x_b$. The linear boundedness of $\hat{V}(u,t)$, however, excludes any solution with absolute value exceeding 1 (i.e., the corresponding constants c_i must be zero). In fact, it turns out that exactly w roots of the polynomial in (2.16) are located inside the unit disk in the complex plane. To see this, observe first that $(\lambda + \mu_d + 1/t)x^w + \mu_d$ has exactly w roots inside the unit disk (due to $\mu_d/(\lambda + \mu_d + 1/t) < 1$). Then Rouché's Theorem establishes that the same is true for the entire polynomial in (2.16), if

$$|\lambda z^b| < |-(\lambda + \mu_d + 1/t)z^w + \mu_d| \text{ on } |z| = 1,$$

which translates into the condition

$$|\mu_d - (\lambda + \mu_d + 1/t)z^w| > \lambda \text{ on } |z| = 1. \tag{2.20}$$

The reverse triangle inequality states for any complex $a, b \in \mathbb{C}$ that $|a-b| \geq \Big||a|-|b|\Big|$, which shows that for $|z| = 1$ the left-hand side of (2.20) is larger than $\lambda + 1/t$, so that (2.20) is indeed fulfilled.
It is now only left to determine the coefficients $c_1, \ldots, c_w$ corresponding to the w roots $x_1, \ldots, x_w \in \mathbb{C}$ with $|x_i| < 1$ of (2.16). To that end, note that (2.19) evaluated at $u = 0, \ldots, w-1$ gives the following system of equations:

$$\begin{aligned}
&\lambda \hat{V}(b-w,t) - (\lambda + \mu_d + 1/t)\hat{V}(0,t) = 0,\\
&\lambda \hat{V}(b-w+1,t) - (\lambda + \mu_d + 1/t)\hat{V}(1,t) + 1/t = 0,\\
&\cdots\\
&\lambda \hat{V}(b-1,t) - (\lambda + \mu_d + 1/t)\hat{V}(w-1,t) + (w-1)/t = 0.
\end{aligned}$$

Substituting the form

$$\hat{V}(u,t) = \sum_{i=1}^{w} c_i x_i^u + u + a_t$$

with $a_t = \lambda b\, t - (\lambda + \mu_d) w\, t$ into this system leads to

$$\lambda \sum_{i=1}^{w} c_i x_i^{b-w} + \lambda(b-w) + \lambda a_t - (\lambda + \mu_d + 1/t)\left(\sum_{i=1}^{w} c_i + a_t\right) = 0,$$
$$\lambda \sum_{i=1}^{w} c_i x_i^{b-w+1} + \lambda(b-w+1) + \lambda a_t - (\lambda + \mu_d + 1/t)\left(\sum_{i=1}^{w} c_i x_i + (1 + a_t)\right) + 1/t = 0,$$
$$\cdots$$
$$\lambda \sum_{i=1}^{w} c_i x_i^{b-1} + \lambda(b-1) + \lambda a_t - (\lambda + \mu_d + 1/t)\left(\sum_{i=1}^{w} c_i x_i^{w-1} + (w - 1 + a_t)\right) + (w-1)/t = 0.$$

But the latter can be rewritten in the form (2.17). □

Example 2.4.3. *Figure 2.4 depicts $\hat{V}(u,t)$ as a function of u for the parameters $b = 100, w = 9, t = 1, \lambda = 10, \mu_d = 90$. Note that for some capital levels u the increase of $\hat{V}(u,1)$ from u to $u+1$ is larger than for others. This is linked to how many down-jumps relative to up-jumps are needed to become negative, and due to the discrete nature of the problem such jumps in $\hat{V}(u,t)$ occur naturally.*

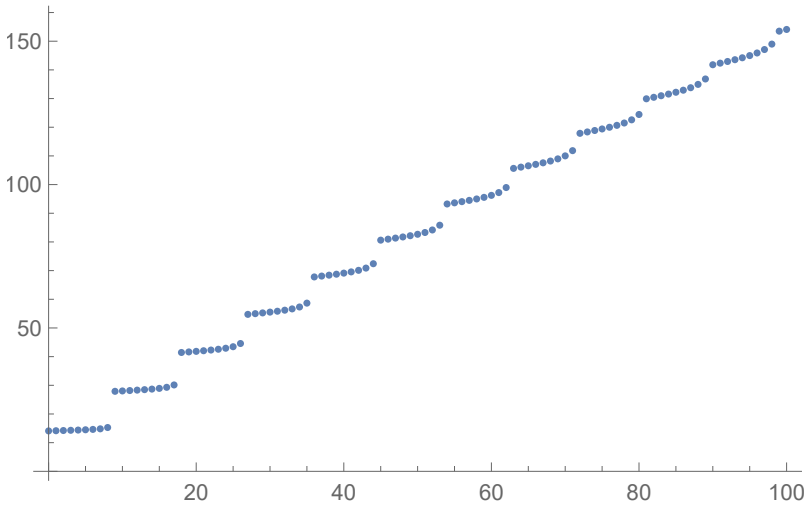

Figure 2.4: $\hat{V}(u,1)$ as a function of u

In an analogous way, an explicit formula for $\widehat{\psi}(u,t) = \mathbb{E}[\psi(u,T)]$ can be derived.

Theorem 2.4.4. *Let b and w, $b > w$, be fixed positive integers. Then the ruin probability up to an exponential time horizon with mean t is given by*

$$\widehat{\psi}(u,t) = \sum_{i=1}^{w} c_i x_i^u \tag{2.21}$$

where $x_1, \ldots, x_w$ are the w roots inside the unit disk of Equation (2.16) and the constants $c_1, \ldots, c_w$ are the solution of the linear equation system

$$\begin{pmatrix} \lambda x_1^{b-w} - (\lambda + \mu_d + 1/t) & \cdots & \lambda x_w^{b-w} - (\lambda + \mu_d + 1/t) \\ \lambda x_1^{b-w+1} - (\lambda + \mu_d + 1/t) x_1 & \cdots & \lambda x_w^{b-w+1} - (\lambda + \mu_d + 1/t) x_w \\ \vdots & \ddots & \vdots \\ \lambda x_1^{b-1} - (\lambda + \mu_d + 1/t) x_1^{w-1} & \cdots & \lambda x_w^{b-1} - (\lambda + \mu_d + 1/t) x_w^{w-1} \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_w \end{pmatrix} = \begin{pmatrix} -\mu_d \\ -\mu_d \\ \vdots \\ -\mu_d \end{pmatrix}. \tag{2.22}$$

Proof. We proceed in the same way as in the proof of Theorem 2.4.2. The analogue of (2.18) then is

$$\begin{aligned}\widehat{\psi}(u,t) = e^{-(\frac{1}{t}+\lambda+\mu_d)h}\widehat{\psi}(u,t) + \lambda\int_0^h e^{-\lambda s}e^{-(1/t+\mu_d)s}\widehat{\psi}(u+b-w,t)\,ds \\ + \mu_d\int_0^h e^{-\mu_d s}e^{-(1/t+\lambda)s}\widehat{\psi}(u-w,t)\,ds\end{aligned} \tag{2.23}$$

and (2.19) is replaced by

$$\lambda\widehat{\psi}(u+b-w,t) - (\lambda+\mu_d+1/t)\widehat{\psi}(u,t) + \mu_d\widehat{\psi}(u-w,t) = 0, \quad u \geq 0, \tag{2.24}$$

which is the homogeneous equation of the former. The boundary conditions here are given by $\widehat{\psi}(u,t) = 1$ for $u < 0$ as well as the obvious bound $\widehat{\psi}(u+b-w,t) \leq 1$ for all $u \geq 0$. Correspondingly, from the proof of the previous theorem we then know that

$$\widehat{\psi}(u,t) = \sum_{i=1}^{w} c_i x_i^u \tag{2.25}$$

with constants $c_1, \ldots, c_w$ still to be determined. Evaluating (2.24) at $u = 0, \ldots, w-1$ gives

$$\lambda\widehat{\psi}(b-w+j,t) - (\lambda+\mu_d+1/t)\widehat{\psi}(j,t) + \mu_d = 0, \quad j = 0, \ldots, w-1.$$

Substituting (2.25) into these leads to

$$\lambda\sum_{i=1}^{w} c_i x_i^{b-w+j} - (\lambda+\mu_d+1/t)\left(\sum_{i=1}^{w} c_i x_i^j\right) + \mu_d = 0, \quad j = 0, \ldots, w-1,$$

or equivalently (2.22). □

2.5 Pool analysis with stochastic rewards

Until this point, we considered deterministic rewards b and w for jump sizes of the surplus process. However, in practice, one may desire to incorporate variability in these quantities to account for instance for the incorporation of variable transaction fees attached to the block reward, or to capture the price volatilities to convert the reward to a fiat currency.

Let us therefore assume now that the up- and downward jumps in the dynamics of the pool manager's surplus are stochastic. Under certain assumptions on the nature of these jumps, this will allow us to still derive closed-form formulas for $\widehat{\psi}$ and $\hat{V}$ in the spirit of [9], see also [16, Ch. 4]. Equation (2.8) then is replaced by

$$R_t = u - \sum_{n=1}^{M_t^d} W_n + \sum_{n=1}^{N_t} B_{r,n}, \quad t \geq 0, \tag{2.26}$$

where we assume $W_n,\ n \in \mathbb{N}$ to be i.i.d. positive random variables with cumulative distribution function F_W and finite mean representing payments to the pool members, and $B_{r,n},\ n \in \mathbb{N}$ are assumed to be i.i.d. positive random variables with distribution function F_{B_r} and finite mean representing the remaining inflow of bounty

rewards diminished by the simultaneous payout to the respective pool member. If one wishes to draw a parallel to the previous case with deterministic rewards, the random variable B_r assumes the role of $b - w$, the fixed block reward minus the payout to the pool member.

Consider the expected surplus of the pool manager as defined previously with a random time horizon T. Concretely, T follows an exponential distribution with mean t. As in the previous section, we are interested in $\hat{V}(u,t)$.

Proposition 2.5.1. *The quantity $\hat{V}(u,t) = \mathbb{E}(R_T \mathbb{I}_{\tau>T})$ for the pool surplus process (2.26) is a solution of the integral equation*

$$\begin{aligned} &\lambda \int_0^\infty \hat{V}(u+b_r,t)\, dF_{B_r}(b_r) - (\lambda + \mu_d + 1/t)\hat{V}(u,t) \\ &+ \mu_d \int_0^u \hat{V}(u-w,t)\, dF_W(w) + u/t = 0, \quad u \geq 0, \end{aligned} \tag{2.27}$$

with boundary conditions $\hat{V}(u,t) = 0$ for all $u < 0$ and $0 \leq \hat{V}(u,t) \leq u + (\lambda\mathbb{E}[B_r] - \mu_d\mathbb{E}[W])t$ for all $u, t \geq 0$.

Proof. We extend the approach of the proof of Theorem 2.4.2 by conditioning on the size of the jump in case a jump occurs. For some small $h > 0$ we correspondingly get

$$\begin{aligned} \hat{V}(u,t) = e^{-(\frac{1}{t}+\lambda+\mu_d)h}\hat{V}(u,t) &+ \frac{1}{t}\int_0^h e^{-s/t}e^{-(\lambda+\mu_d)s} u\, ds \\ &+ \lambda \int_0^h e^{-\lambda s} e^{-(1/t+\mu_d)s} \int_0^\infty \hat{V}(u+b_r,t)\, dF_{B_r}(b_r)\, ds \\ &+ \mu_d \int_0^h e^{-\mu_d s} e^{-(1/t+\lambda)s} \int_0^u \hat{V}(u-w,t)\, dF_W(w)\, ds. \end{aligned} \tag{2.28}$$

Taking the derivative w.r.t. h and setting $h = 0$, one obtains (2.27). The property $\hat{V}(u,t) = 0$ for all $u < 0$ follows by definition and the linear upper bound in u and t is obtained from the inequality $\hat{V}(u,t) \leq \mathbb{E}(R_T)$.

□

Remark 2.5.2. *For degenerate F_{B_r} and F_W (i.e. for constant B_r and W), the integral equation (2.27) simplifies to (2.19) (and for integer constants, we get back to the setting of Theorem 2.4.2).*

For our purposes, it is very reasonable to assume (and will lead to simplified notation) that the generic random variables B_r and W are connected via

$$B_r = a\, W \tag{2.29}$$

for some constant $a > 1$ that depends on the number of miners in the pool. Indeed, W is the payment to the pool miner for solving a less complex puzzle, and B_r can be seen as the bounty reward when the more complex puzzle is solved minus the payment to the miner who solved it, and that latter payment will be a constant fraction, depending on the specification of the pool rules. Note that for most positive random variables, a scaled version of it belongs to the same class of

random variables with only the parameter(s) changed, and the latter is indeed the case for all distributional assumptions that we will pursue in this chapter. In any case, all results below can easily be adapted to the case when B_r and W follow distributions that are unrelated combinations of exponentials with different n and A_i's.

Let us now consider in more detail the case where both the up- and down-jumps are random variables whose distributions are combinations of exponentials. The latter class is dense in the class of all random variables on the positive half-line, so that the result is in fact quite general (see e.g. Dufresne [56]). Concretely, the density of downward jumps is then assumed to be of the form

$$f_W(w) = \sum_{i=1}^{n} A_i \alpha_i e^{-\alpha_i w}, \quad w > 0, \tag{2.30}$$

where $\alpha_1 < \alpha_2 < \ldots < \alpha_n$ and $A_1 + \cdots + A_n = 1$ (but the A_i are not necessarily positive). The Laplace transform of this density is given by

$$\tilde{f}_W(s) := \mathbb{E}(e^{-sW}) = \sum_{i=1}^{n} A_i \frac{\alpha_i}{\alpha_i + s}, \quad \mathrm{Re}(s) > -\alpha_1.$$

From (2.29), we then have

$$f_{B_r}(b_r) = \sum_{i=1}^{n} A_i \beta_i e^{-\beta_i b_r}, \quad b_r > 0 \tag{2.31}$$

with $\beta_i = \alpha_i / a$, $i = 1, \ldots, n$.

Theorem 2.5.3. *If W and B_r are random variables with densities given in* (2.30) *and* (2.31)*, then we have*

$$\hat{V}(u,t) = \sum_{k=1}^{n} C_k e^{-r_k u} + u + t \sum_{i=1}^{n} A_i \left(\frac{\lambda}{\beta_i} - \frac{\mu_d}{\alpha_i} \right), \tag{2.32}$$

where $r_1, \ldots, r_n$ are the solutions with positive real parts of

$$\lambda \sum_{i=1}^{n} A_i \frac{\beta_i}{\beta_i + r} + \mu_d \sum_{i=1}^{n} A_i \frac{\alpha_i}{\alpha_i - r} - (\lambda + \mu_d + 1/t) = 0 \tag{2.33}$$

and

$$C_k = \frac{\sum_{j=1}^{n} B_j \prod_{h=1}^{n} (\alpha_j - r_h) \prod_{i=1, i \neq j}^{n} \frac{r_k - \alpha_i}{\alpha_j - \alpha_i}}{\prod_{h=1, h \neq k}^{n} (r_k - r_h)}, \quad k = 1, \ldots, n \tag{2.34}$$

with

$$B_j = \frac{1}{\alpha_j^2} - \frac{t}{\alpha_j} \sum_{i=1}^{n} A_i \left(\frac{\lambda}{\beta_i} - \frac{\mu_d}{\alpha_i} \right), \quad j = 1, \ldots, n.$$

Proof. Substituting (2.30) and (2.31) into (2.27), we get

$$\lambda \sum_{i=1}^{n} A_i \beta_i \int_0^\infty \hat{V}(u+b_r,t) e^{-\beta_i b_r} db_r$$
$$-(\lambda+\mu_d+1/t)\hat{V}(u,t) + \mu_d \sum_{i=1}^{n} A_i \alpha_i \int_0^u \hat{V}(u-w,t) e^{-\alpha_i w} dw + u/t = 0, \quad u \geq 0.$$

The function $\hat{V}(u,t)$ then has the form

$$\hat{V}(u,t) = \sum_{k=1}^{n} C_k e^{-r_k u} + d_1 u + d_0,$$

for constants $C_1, \ldots, C_n, r_1, \ldots, r_n, d_0, d_1$ to be determined. In fact, plugging this ansatz into the above equation shows that comparing coefficients of $e^{-r_k u}$ exactly gives (2.33) (which is a generalized Lundberg equation in the terminology of ruin theory, cf. [16]). That equation has exactly n solutions with positive real part $r_1, \ldots, r_n$ and n solutions with negative real part (see e.g. [160]). The solutions with negative real part would enter $\hat{V}$ with positive real part and are correspondingly irrelevant for our purpose, as that would violate the linear boundedness of the resulting $\hat{V}$ (in other words, the coefficients in front of such terms need to be zero). Comparing coefficients of $e^{-\alpha_i u}$, $i = 1, \ldots, n$ gives

$$\sum_{k=1}^{n} \frac{C_k}{\alpha_i - r_k} = \frac{d_1}{\alpha_i^2} - \frac{d_0}{\alpha_i}, \quad i = 1, \ldots, n. \tag{2.35}$$

Coefficients in front of $u\, e^{-\alpha_i u}$, $i = 1, \ldots, n$ all cancel. After a little algebra, one sees that a comparison of coefficients of u in that equation establishes $d_1 = 1$ and a comparison of the constant coefficients gives

$$d_0 = t \sum_{i=1}^{n} A_i \left(\frac{\lambda}{\beta_i} - \frac{\mu_d}{\alpha_i} \right).$$

These obtained values of d_1 and d_0 can now be plugged into (2.35), and the resulting system of linear equations can be solved explicitly to give (2.34), either by realizing that the coefficient matrix is a Cauchy matrix or by using the trick of rational function representation developed in [9, Sec.4]. □

Example 2.5.4. *A particular simple example of the above is the case where W is exponentially distributed with parameter α and B_r is exponentially distributed with parameter β. In that case $n = 1$ in Theorem 2.5.3 and we obtain*

$$\hat{V}(u,t) = \left(\frac{1}{\alpha^2} - \frac{t}{\alpha} \left(\frac{\lambda}{\beta} - \frac{\mu_d}{\alpha} \right) \right) (\alpha - R) e^{-Ru} + u + t \left(\frac{\lambda}{\beta} - \frac{\mu_d}{\alpha} \right), \tag{2.36}$$

where R is the (unique) solution with positive real part of

$$\lambda \frac{\beta}{\beta + r} + \mu_d \frac{\alpha}{\alpha - r} - (\lambda + \mu_d + 1/t) = 0. \tag{2.37}$$

Let us now move on to study the ruin probability $\widehat{\psi}(u,t) = \mathbb{E}[\psi(u,T)]$ in the present context.

Theorem 2.5.5. *If W and B_r are random variables with densities given in* (2.30) *and* (2.31)*, then we have*

$$\widehat{\psi}(u,t) = \sum_{k=1}^{n} D_k e^{-r_k u}, \tag{2.38}$$

where $r_1, \dots, r_n$ are the n solutions with positive real parts of (2.33) *and*

$$D_k = \frac{\sum_{j=1}^{n} \frac{1}{\alpha_j} \prod_{h=1}^{n} (\alpha_j - r_h) \prod_{i=1, i\neq j}^{n} \frac{r_k - \alpha_i}{\alpha_j - \alpha_i}}{\prod_{h=1, h\neq k}^{n} (r_k - r_h)}, \quad k = 1, \dots, n. \tag{2.39}$$

Proof. We can proceed in the same way as in the proof of Proposition 2.5.1 to derive an integral equation for the ruin probability. The analogue of Equation (2.28) here is

$$\begin{aligned}\widehat{\psi}(u,t) = e^{-(\frac{1}{t}+\lambda+\mu_d)h}\widehat{\psi}(u,t) + \lambda \int_0^h e^{-\lambda s} e^{-(1/t+\mu_d)s} \int_0^\infty \widehat{\psi}(u+b_r,t)\, dF_{B_r}(b_r)\, ds \\ + \mu_d \int_0^h e^{-\mu_d s} e^{-(1/t+\lambda)s} \left(\int_0^u \widehat{\psi}(u-w,t)\, dF_W(w) + \int_u^\infty 1\, dF_W(w) \right) ds.\end{aligned} \tag{2.40}$$

Taking the derivative w.r.t. h and evaluating at $h = 0$ then gives

$$\begin{aligned}&\lambda \int_0^\infty \widehat{\psi}(u+b_r,t)\, dF_{B_r}(b_r) - (\lambda + \mu_d + 1/t)\widehat{\psi}(u,t) \\ &+ \mu_d \int_0^u \widehat{\psi}(u-w,t)\, dF_W(w) + \mu_d(1 - F_W(u)) = 0, \quad u \geq 0.\end{aligned} \tag{2.41}$$

Here the boundary conditions are $\widehat{\psi}(u,t) = 1$ for $u < 0$ and $\widehat{\psi}(u,t) \leq 1$ for $u \geq 0$ and arbitrary $t > 0$, and uniqueness of its solution follows analogously to Theorem 2.5.3. Under the assumptions on F_{B_r} and F_W this reads

$$\begin{aligned}\lambda \sum_{i=1}^{n} A_i \beta_i \int_0^\infty \widehat{\psi}(u+b_r,t)\, e^{-\beta_i b_r}\, db_r - (\lambda + \mu_d + 1/t)\widehat{\psi}(u,t) \\ + \mu_d \sum_{i=1}^{n} A_i \alpha_i \int_0^u \widehat{\psi}(u-w,t)\, e^{-\alpha_i w}\, dw + \mu_d \sum_{i=1}^{n} A_i e^{-\alpha_i u} = 0, \quad u \geq 0.\end{aligned} \tag{2.42}$$

In analogy to the proof of Theorem 2.5.3 we then see that the ruin probability must have the form

$$\widehat{\psi}(u,t) = \sum_{k=1}^{n} D_k e^{-r_k u}$$

for constants $D_1, \dots, D_n$ to be determined, and $r_1, \dots, r_n$ being the n positive solutions of (2.33). The constants D_k are now obtained by substituting the above expression into (2.42) and comparing coefficients of $e^{-\alpha_i u}$, $i = 1, \dots, n$. This gives

$$\sum_{k=1}^{n} \frac{D_k}{\alpha_i - r_k} = \frac{1}{\alpha_i}, \quad i = 1, \dots, n. \tag{2.43}$$

This system of linear equations is again of Cauchy matrix form with explicit solution (2.39), establishing the result. □

Example 2.5.6. *If W and B_r are exponentially distributed with parameter α and β, respectively, then* (2.38) *simplifies to*

$$\widehat{\psi}(u,t) = (1 - R/\alpha)e^{-Ru},\ u \geq 0, \tag{2.44}$$

where R is the (unique) solution with positive real part of (2.37).
Note that for $t \to \infty$ one obtains $R = (\lambda\alpha - \mu_d\beta)/(\lambda + \mu_d) > 0$, so that

$$\psi(u) = \frac{\mu_d(1+\beta/\alpha)}{\lambda+\mu_d} e^{-\frac{\lambda\alpha-\mu_d\beta}{\lambda+\mu_d}u},\ u \geq 0. \tag{2.45}$$

In particular, without initial capital in the pool, the infinite-time ruin probability amounts to

$$\psi(0) = \frac{\mu_d(1+\beta/\alpha)}{\lambda+\mu_d},$$

in accordance with Formula (8.1) in [9].

2.6 Individual miner analysis

2.6.1 Deterministic rewards

Comparing the formula describing the miner's surplus under the PPS pooling scheme (2.7) with the solo-mining surplus (2.3), one can see that they are in fact the same type of process, only distinguished by the reward amount and frequency. Correspondingly, the formulas obtained by Albrecher and Goffard [10] for the expected value of the surplus and the ruin probability of a honest miner apply in the PPS case with deterministic rewards. Adapted to the present context, we hence get:

Theorem 2.6.1. *[10] For the miner's surplus process $R_t^i = u - c_i \cdot t + M_t^i \cdot w,\ t \geq 0$, with $M_t^i \sim Poisson(p_i\mu t)$, the value function $\hat{V}(u,t)$ can be expressed as*

$$\hat{V}(u,t) = u + (p_i\mu w - c_i)t(1 - e^{\rho^* u}), \tag{2.46}$$

where ρ^* is the negative solution of the equation

$$-c_i\rho + p_i\mu(e^{w\rho} - 1) = 1/t. \tag{2.47}$$

Theorem 2.6.2. *[10] For the same surplus process, the ruin probability with exponential time horizon is given by $\hat{\psi}(u,t) = e^{\rho^* u}$, where ρ^* is the negative solution of* (2.47).

2.6.2 Stochastic rewards

Consider now the same surplus process as in the previous section, but with stochastic rewards. Let us define this process by

$$R_t^i = u - c_i \cdot t + \sum_{n=1}^{M_t^i} W_n,\ t \geq 0, \tag{2.48}$$

where we assume $W_n,\ n \in \mathbb{N}$ to be i.i.d. positive random variables with cumulative distribution function F_W and finite mean and $M_t^i \sim Poisson(p_i\mu t)$ as previously.

This type of process is denominated as the *dual problem* in the insurance context, see e.g. [18]. We assume that the net profit condition $p_i\mu\mathbb{E}[W_n] > c_i$ is satisfied. We are again interested in deriving the expected value of the surplus and the ruin probability for the miner. To simplify the computations, we consider again an exponential time horizon.

Theorem 2.6.3. *For exponential time horizon, the expected value of the miner's surplus $\hat{V}(u,t)$ can be expressed as the solution of the integro-differential equation*

$$c_i\hat{V}'(u,t) + (\frac{1}{t} + p_i\mu)\hat{V}(u,t) - p_i\mu\int_0^{+\infty}\hat{V}(u+w,t)dF_W(w) - u/t = 0, \quad (2.49)$$

with boundary conditions $\hat{V}(0,t) = 0$ and $0 \leq \hat{V}(u,t) \leq u - c_it + p_i\mu\mathbb{E}[W]$.

Proof. As in previous sections, by conditioning the occurrence of T to a small time interval $(0,h)$, we can write the value function as

$$\begin{aligned}\hat{V}(u,t) = e^{-h(\frac{1}{t}+p_i\mu)}\hat{V}(u-c_ih,t) + \int_0^h \frac{1}{t}e^{-s(\frac{1}{t}+p_i\mu)}(u-c_is)ds \\ + \int_0^h p_i\mu e^{-s(\frac{1}{t}+p_i\mu)}\int_0^{+\infty}\hat{V}(u-c_ih+w,t)dF_W(w)ds.\end{aligned} \quad (2.50)$$

Taking the derivative w.r.t. h and evaluating it at $h = 0$ gives us (2.49). The boundary condition follows from ruin considerations. □

For rewards whose distribution is a combination of exponentials (2.30), we can refine Theorem 2.6.3.

Theorem 2.6.4. *When W has density $f_W(w) = \sum_{j=1}^n A_j\alpha_j e^{-\alpha_j w}$, $w > 0$, then*

$$\hat{V}(u,t) = t\left(c_i - p_i\mu\sum_{j=1}^n\frac{A_j}{\alpha_j}\right)e^{-Ru} + u + t\left(p_i\mu\sum_{j=1}^n\frac{A_j}{\alpha_j} - c_i\right),\ u > 0, \quad (2.51)$$

where R is the unique solution with positive real part of the equation

$$c_iR + p_i\mu\sum_{j=1}^n\frac{A_j\alpha_j}{R+\alpha_j} - (\frac{1}{t} + p_i\mu) = 0.$$

Proof. Equation (2.49) translates into

$$c_i\hat{V}'(u,t) + (\frac{1}{t} + p_i\mu)\hat{V}(u,t) - p_i\mu\sum_{j=1}^n A_j\alpha_j\int_0^{+\infty}\hat{V}(u+w,t)e^{-\alpha_j w}dw - u/t = 0. \quad (2.52)$$

This equation has a solution of the form

$$\hat{V}(u,t) = Ce^{-Ru} + d_1u + d_0 \quad (2.53)$$

and we plug this ansatz into (2.52)

$$\begin{aligned}&c_i(-RCe^{-Ru} + d_1) + (\frac{1}{t} + p_i\mu)(Ce^{-Ru} + d_1u + d_0) \\ &- p_i\mu\sum_{j=1}^n A_j\alpha_j\int_0^{+\infty}(Ce^{-R(u+w)} + d_1(u+w) + d_0)e^{-\alpha_j w}dw - u/t = 0.\end{aligned} \quad (2.54)$$

Comparing coefficients, we obtain

$$d_1 = 1, \quad d_0 = t\left(p_i\mu \sum_{j=1}^{n} \frac{A_j}{\alpha_j} - c_i\right).$$

Further, a comparison of the coefficients in front of e^{-Ru} simplifies to the following equation:

$$c_i R + p_i\mu \sum_{j=1}^{n} \frac{A_j\alpha_j}{R+\alpha_j} - (\frac{1}{t} + p_i\mu) = 0. \tag{2.55}$$

Similarly to the Lundberg equation derived in [99], we note that there is one positive root R to this equation. To complete the proof, we consider the boundary condition $\hat{V}(0,t) = 0$ and substituting into the ansatz gives $C = -d_0$. □

Example 2.6.5. *When W is exponentially distributed, i.e. $f_W(w) = \alpha e^{-\alpha w}$, $w > 0$, Equation (2.51) simplifies to*

$$\hat{V}(u,t) = t\left(c_i - \frac{p_i\mu}{\alpha}\right)e^{-Ru} + u + t\left(\frac{p_i\mu}{\alpha} - c_i\right), \ u > 0, \tag{2.56}$$

where R is the solution with positive real part of

$$c_i R^2 + (\alpha c_i - \frac{1}{t} - p_i\mu)R - \alpha\frac{1}{t} = 0.$$

Theorem 2.6.6. *For exponential time horizon, the miner's ruin probability can be expressed as*

$$\hat{\psi}(u,t) = e^{-R\cdot u}, \tag{2.57}$$

where R is the unique positive root of

$$p_i\mu + \frac{1}{t} - c_i R = p_i\mu\mathbb{E}[e^{-RW_n}]. \tag{2.58}$$

Proof. The proof is adapted from Example 2 of Mazza and Rullière [103]. From the latter, we have that the Laplace transform of the ruin time τ in the dual problem is $\mathbb{E}[e^{-s\tau}] = e^{-R(s)\cdot u}$, with $R(s)$ being the unique positive root of $p_i\mu + s - c_i R = p_i\mu\mathbb{E}[e^{-RW_n}]$. Since the ruin probability up to an exponential time horizon can be rewritten as

$$\hat{\psi}(u,t) = \mathbb{E}[\mathbb{P}(T > \tau) \mid \tau], \tag{2.59}$$

with $T \sim Exp(1/t)$, it immediately follows that

$$\hat{\psi}(u,t) = \mathbb{E}[e^{\tau/t}] \tag{2.60}$$

which completes the proof. □

Example 2.6.7. *If W is an exponential random variable, i.e. $f_W(w) = \alpha e^{-\alpha w}$, $w > 0$, then the ruin probability reduces to*

$$\hat{\psi}(u,t) = e^{-R^* u}, \tag{2.61}$$

where

$$R^* = \frac{1/t + p_i\mu - c_i\alpha + \sqrt{\Delta}}{2c_i}, \ \Delta = (c_i\alpha - p_i\mu - 1/t)^2 + 4c_i\alpha/t. \tag{2.62}$$

Remark 2.6.8. *Results concerning the ruin probability can also be retrieved from the respective results for a more general renewal model considered in Alcoforado et al. [11].*

2.7 Numerical illustration

2.7.1 Pool manager

In this section, we will illustrate the pool dynamics in both the deterministic and stochastic setting. In addition, we will perform a sensitivity analysis on main decision variables from the pool's perspective.

First, let us define the set of parameters used in the following examples. For each illustration, we keep all the parameters fixed to these levels except the one that is varying : $t = 336,\ p_I = 0.1,\ q = 0.1,\ f = 0.02,\ b = 1000MU,\ w = (1-f)bq = 98MU,\ \lambda = 6p_I = 0.6,\ \mu_d = 6p_I(1/q - 1) = 5.4$. The units we use are hours (h) for the time parameters and monetary units (MU) for the value functions. The choice for the time horizon t is equal to 2 weeks because it is linked to the period of difficulty adjustment. The monetary units are related to bitcoin in this way : $1000MU = 6.25BTC$. The reason for this scaling is purely practical to solve the deterministic problem which involves integer constraints. As of May 28th 2021, $1BTC \approx \$35670.5$, so $1MU \approx \$231.85$.

Figure 2.5 compares the function $\hat{V}(u,t)$ defined in Theorem 2.4.2 with the Monte Carlo simulation of the mining process with deterministic and exponential time horizon fixed at the same mean parameter. The functions are reduced by u to isolate the expected gain realized by the pool manager. We can see that the exact formula falls nicely within the 95% confidence interval bounds of the MC simulations within fixed or exponential time horizon. The red line represents the upper limit of the function to which it converges as $u \to +\infty$, which is also the expected value of the gain in absence of ruin considerations. One can see that for small levels of initial capital potential ruin affects the resulting profit considerably, and for any given u the pool manager can quantify the undesirable effect of ruin.
Figure 2.6 exhibits the corresponding ruin probability $\hat{\psi}(u,t)$ for the mining pool. We can note that ruin is highly non-negligible for low levels of initial capital. Indeed, $\hat{\psi}(u, t = 336) < 5\%$ for $u > 22594$, which is equivalent to \$5238419. We also see how the exponential time horizon slightly underestimates the ruin probability for low capital levels, which is due to the skewness of the exponential distribution. This graph can also be interpreted sideways: if one fixes a threshold for the ruin probability on the vertical axis, the corresponding initial capital can be read off on the horizontal axis.

In Figure 2.7, we depict the sensitivity of the expected surplus and the ruin probability to the management fee f. Not surprisingly, the relationship between f and $\hat{\psi}(u,t)$ is decreasing, as the pool retains more reward for itself. The parameter f impacts the expected gain of the pool manager.

Remark 2.7.1. *Note that here we consider the interplay of all factors in a static set-up for a fixed number of participants in the pool. One may then go one step further to consider the fact that a higher fee f may deter some participants to join the pool, with respect to their willingness to pay and when comparing with fees of other competitive pools in the market. This outflow of miners would consequently impact negatively the expected profit of the pool. However, such considerations naturally ask*

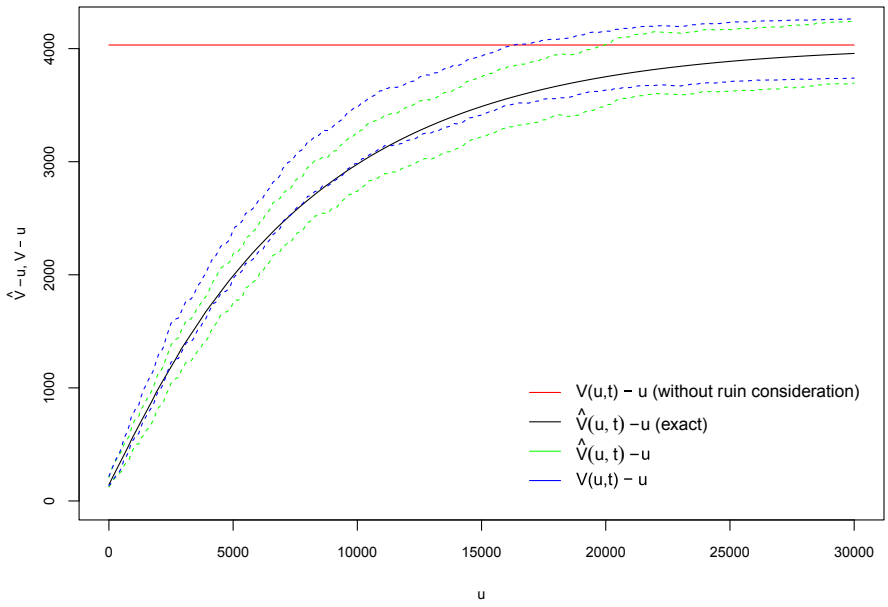

Figure 2.5: $\hat{V}(u,t) - u$ as a function of u and simulated $\hat{V}(u,t) - u$ and $V(u,t) - u$ with their 95% confidence interval bound in dashed with deterministic size jumps b and w.

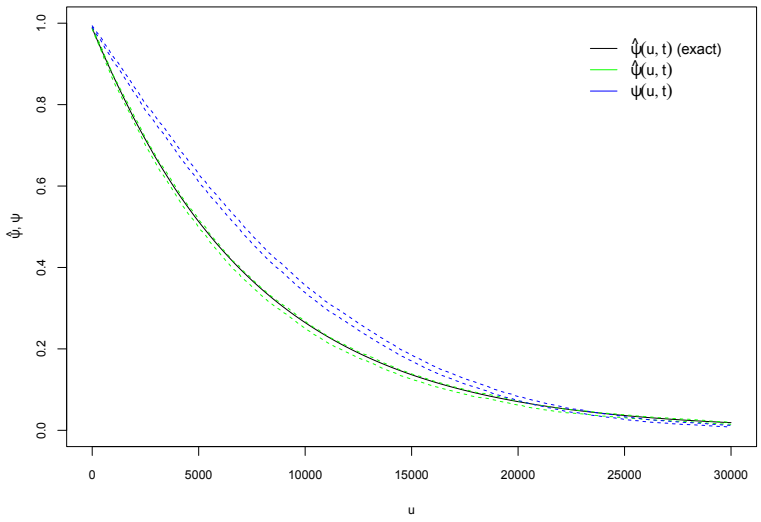

Figure 2.6: $\hat{\psi}(u,t)$ as a function of u and simulated $\hat{\psi}(u,t)$ and $\psi(u,t)$ with their 95% confidence interval bound in dashed with deterministic size jumps b and w.

for an analysis with competing pools, which is beyond the scope of this chapter.

In Figure 2.8, we explore the impact of the relative difficulty to find a share q on ruin and expected surplus.

It is worthwhile to note that increasing q is profitable to the pool manager. Indeed, as q increases, the payout of shares to the pool members is getting less

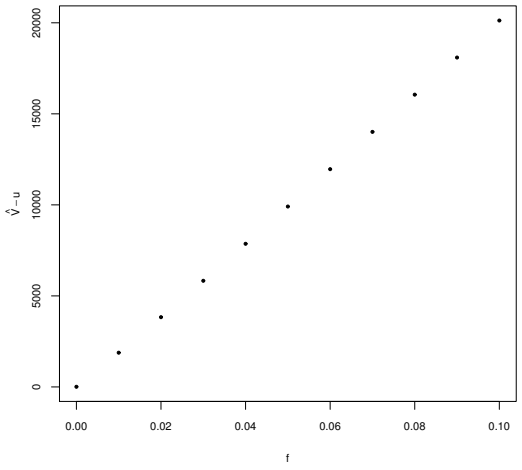

(a) $\hat{V}(u,t) - u$ as a function of f.

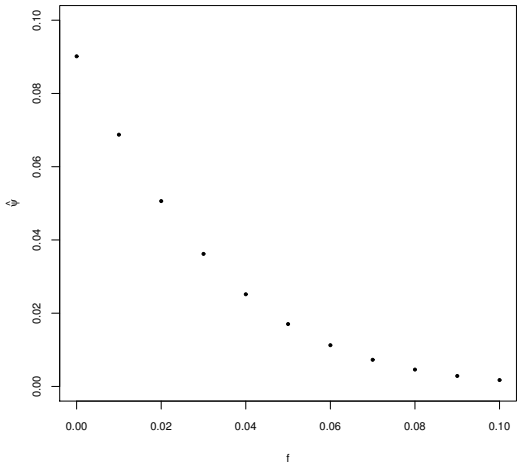

(b) $\hat{\psi}(u,t)$ as a function of f.

Figure 2.7: Sensitivity to f in case of deterministic rewards and exponential time horizon.

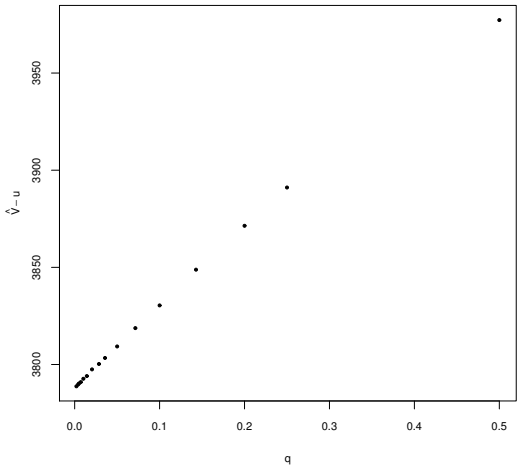

(a) $\hat{V}(u,t) - u$ as a function of q.

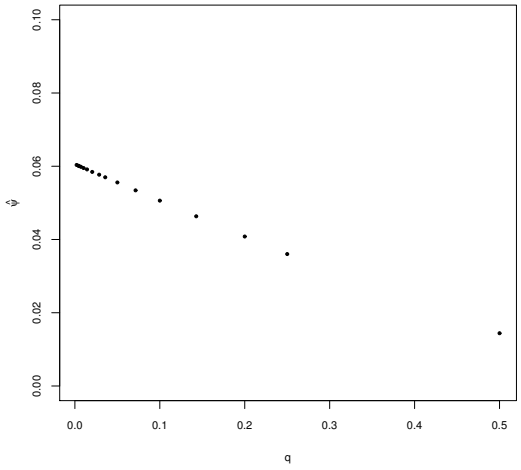

(b) $\hat{\psi}(u,t)$ as a function of q.

Figure 2.8: Sensitivity to q in case of deterministic rewards and exponential time horizon.

frequent, thus the pool manager retains more liquidity and controls his probability of ruin at lower levels. The parameter q adjusts the magnitude of the risk transfer between the miners and their manager.

Figures 2.9, 2.10, 2.11, 2.12 illustrate the same concepts with exponentially distributed rewards. For comparison, the parameters for the exponential distributions are chosen so that the resulting mean matches the deterministic jump sizes, i.e. $\alpha = 1/w = 1/98, \beta = 1/b = 1/1000$.

Figure 2.13 gives a two-way sensitivity analysis with respect to the pool size p_I and the pool fee f. The level curves indicate the expected profit for the pool manager for different pool sizes.

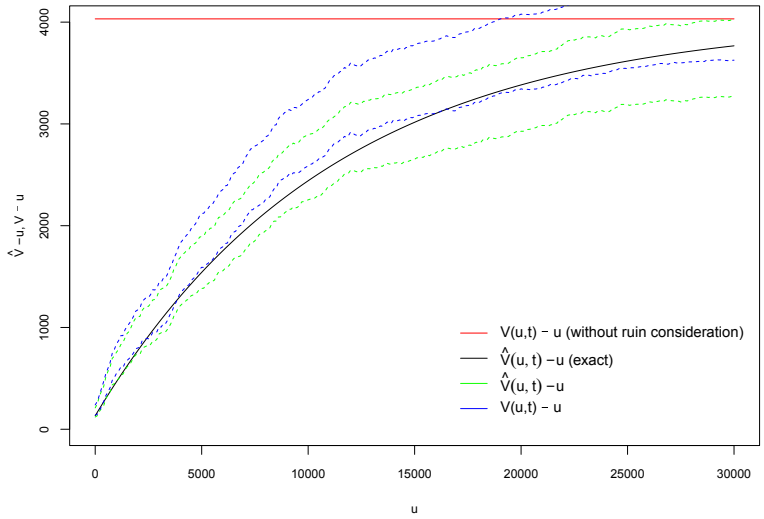

Figure 2.9: $\hat{V}(u,t) - u$ as a function of u and simulated $\hat{V}(u,t) - u$ and $V(u,t) - u$ with their 95% confidence interval bound in dashed. Both jumps are exponentially distributed.

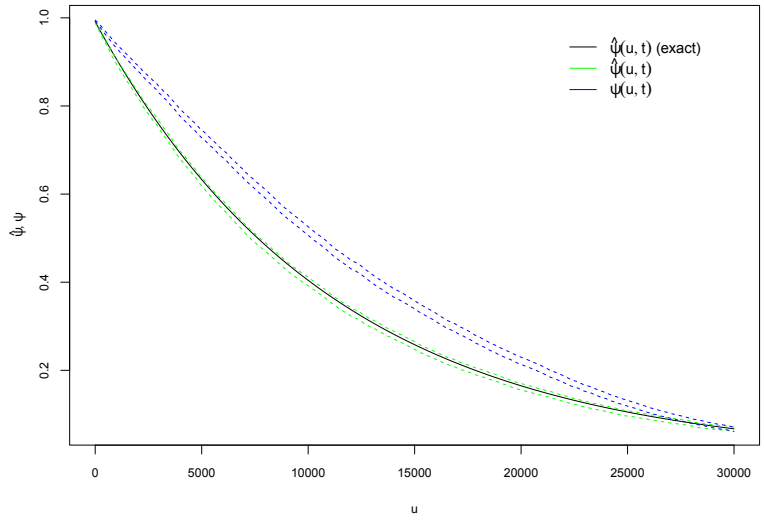

Figure 2.10: $\hat{\psi}(u,t)$ as a function of u and simulated $\hat{\psi}(u,t)$ and $\psi(u,t)$ with their 95% confidence interval bound in dashed. Both jumps are exponentially distributed.

For a bigger pool size p_I, in order to maintain the same level of expected profit, the pool manager can reduce the fee size. One can clearly see an inverse relationship between the pool size and the fee. Thus, a bigger pool can diminish its fees to attract more miners and thus to grow even more. This implies a threat on the decentralized nature of the consensus protocol. If a mining pool manager concentrates more than 50% of the total hashpower, then the blockchain is prone to 51%-type attacks such as double spending in the bitcoin context. How can a smaller mining pool tackle this problem? One solution consists in offering to take on more risk by decreasing the

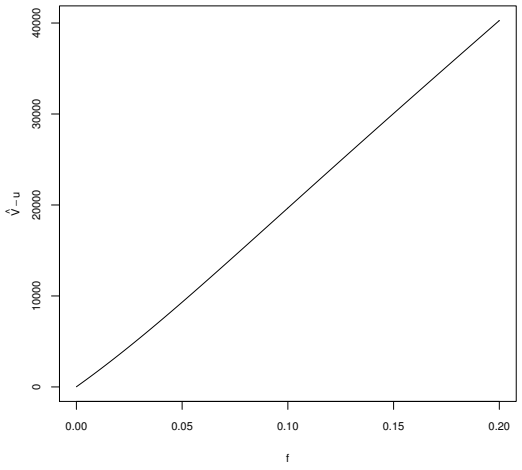

(a) $\hat{V}(u,t) - u$ as a function of f.

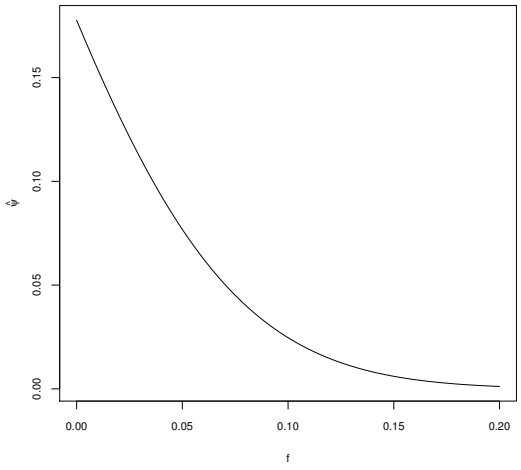

(b) $\hat{\psi}(u,t)$ as a function of f.

Figure 2.11: Sensitivity to f in case of exponentially distributed rewards and exponential time horizon.

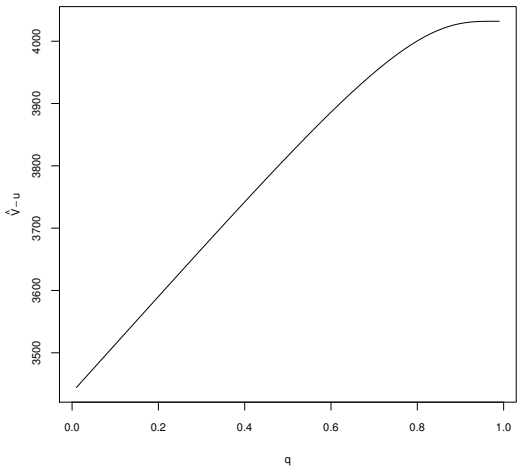

(a) $\hat{V}(u,t) - u$ as a function of q.

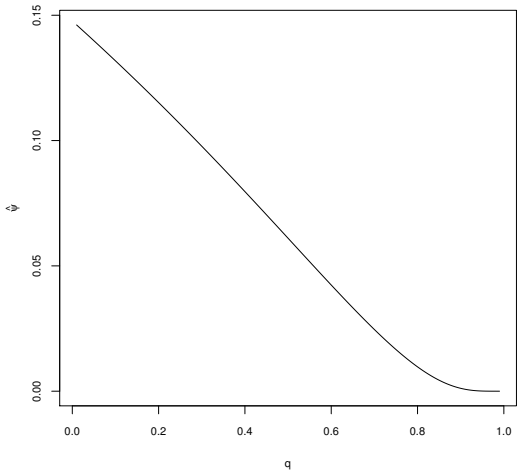

(b) $\hat{\psi}(u,t)$ as a function of q.

Figure 2.12: Sensitivity to q in case of exponentially distributed rewards and exponential time horizon.

difficulty of finding a share which reduces to decreasing the value of q. Figure 2.14 shows the expected profit of two mining pools, one for which $p_I = 0.1$ and a smaller one for which $p_I = 0.02$, both having an initial capital level $u = 22500$, for both the reward and the time horizon being exponentially distributed. The level curves indicate that in terms of expected profit a smaller miner may decrease q without increasing the pool fee f, while maintaining the same level of profitability. That is not the case for the larger mining pool whose expected profit turns out to be more sensitive to q.

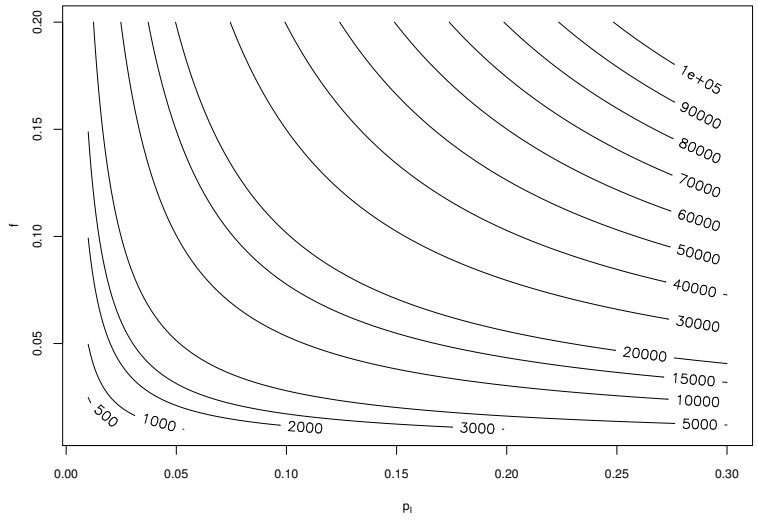

Figure 2.13: $\hat{V}(u,t)$ as a function of p_I and f for $u = 22500$. Both jumps are exponentially distributed.

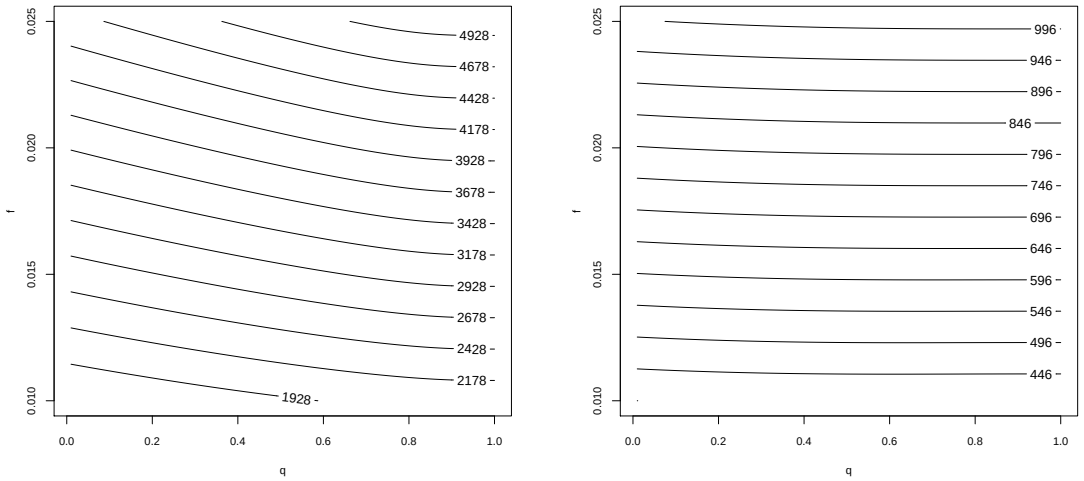

(a) $\hat{V}(u,t) - u$ as a function of q and f for a large mining pool ($p_I = 0.1$).

(b) $\hat{V}(u,t) - u$ as a function of q and f for a small mining pool ($p_I = 0.02$).

Figure 2.14: Sensitivity to q and f of the expected profit of two mining pools of different size over and exponentially distributed time horizon and reward for $u = 22500$.

2.7.2 Individual miner

Let us now compare the situation of an individual miner before and after joining the pool. We recall Figure 2.3 (left panel), which examplifies the pool members' surplus. Also, the surplus of the member is described by (2.9). Finally, we use the results presented in Sections 2.6.1 and 2.6.2 to assess the pool effect for the individual miner's surplus following the protocol. Consider a miner in a deterministic rewards environment. We assume a PPS pool and consider a pool member whose hashpower

is equal to 1% of the pool's total hashpower, i.e. $p_i = 0.001$. For the choice of other parameters, we assume that the cost of electricity c is given by

$$c = p_i \times e_W \times \pi_W, \tag{2.63}$$

where e_W is the electricity consumption of the network expressed in kWh, and π_W is the price of electricity per kWh. For the sake of our example, we take the estimate of e_W as $\frac{115.541\times10^9}{365.25\times24}$.[3] The price of electricity is taken to be \$0.06, then converted to our MU. Therefore, the net profit condition is satisfied both with and without joining the pool. Figures 2.15 and 2.16 illustrate the expected surplus and ruin probability with deterministic rewards and exponential time horizon. One can observe how effective the risk reduction in case of joining the pool is for the individual miner. Figure 2.16 particularly emphasizes the drastic decrease of ruin probability for low capital levels.

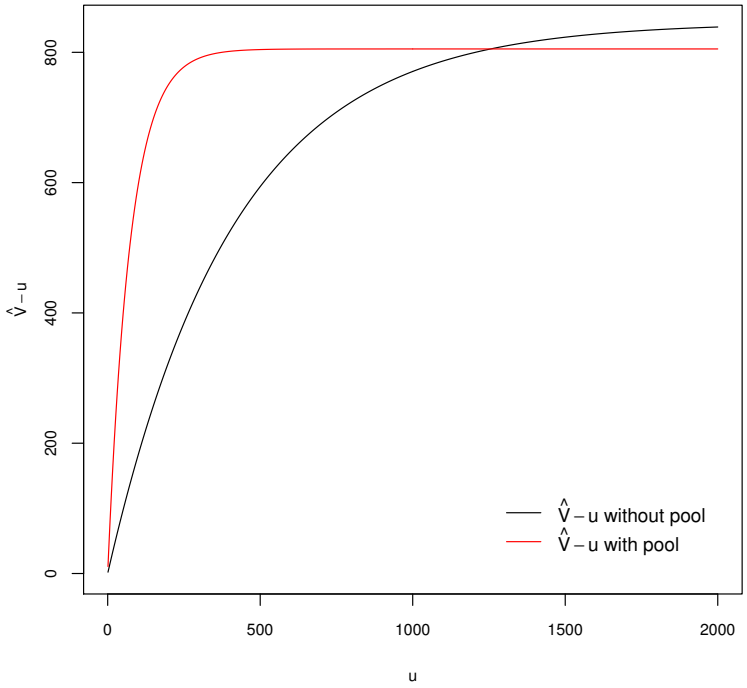

Figure 2.15: $\hat{V}(u,t) - u$ as a function of u for an individual pool miner alone in black and within the pool in red.

Up until a level of initial capital of $u = 1255$, it is more profitable for the miner to join the pool, whereas for higher levels of capital the pool fee becomes the main decision driver instead of the ruin considerations. Converted to USD, this amounts to approximately \$290,971. Recall that this is akin to the effects of reinsurance, as the miner cedes part of his risk to the pool in exchange of a fixed contractual payment (pool fee). In Figure 2.17, we can see how this indifference point evolves with respect to the proportion of the computational power p_i of an individual miner. One observes that this point grows with the size of the miner. This may seem counterintuitive, but one must note that in the formula giving us the result (2.51), the

[3] https://cbeci.org/, consulted on May 28th 2021.

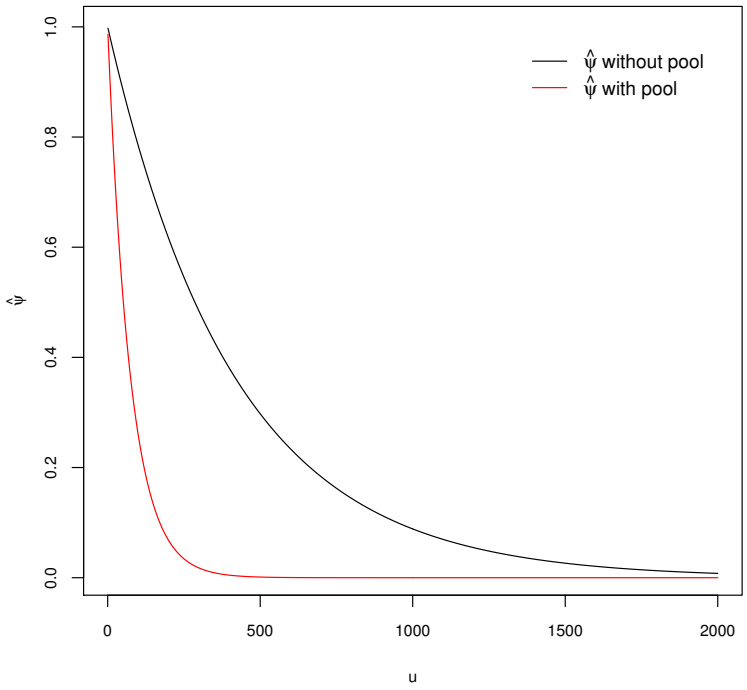

Figure 2.16: $\hat{\psi}(u,t)$ as a function of u for an individual pool miner alone in black and within the pool in red.

parameter p_i enters only as a scaling factor and is also indirectly incorporated in c_i via Equation (2.63). In practice, a miner with a very small computational power is unlikely to see a reward very fast, which may deter him from entering the mining market at all. Moreover, comparing the expected profit at the indifference point for a miner with $p_i = 0.001\%$ and $p_i = 0.1\%$, we obtain an expected profit of $8MU$ and $800MU$ respectively, which translates to $\$2,120$ and $\$212,000$ respectively, as converted on December 30, 2023. To give a practical example, let us consider a novice miner with one PC according to 2023 market standards. With such an equipment, this miner can mine at a rate of 50 megahashes per second. Using an online calculator[4], one can conclude that the expected profit is practically null. The miner in our previous example, amounting to even $p_i = 0.001\%$ of global computational power, is approximately $93,000,000$ times more powerful. Naturally, mining involves investing in better and specialized equipment, which makes the difference.

Finally, we investigate the sensitivity of the miner's expected surplus with respect to the key model parameters. In Figure 2.18, the miner can see for his level of initial capital u whether it is better to join the pool or not, depending on the employed fee f. As before, for higher levels of capital, the miner is less willing to accept high fees than a miner with less initial capital. We also observe that the two red lines (miner in the pool with different initial capital u) are much closer to each other than the two black lines (miner outside of the pool with different initial capital u). This is due to the risk reduction of the miner inside the pool, since he is transferring part

[4]For example: `https://bitcoinx.com/profit/`.

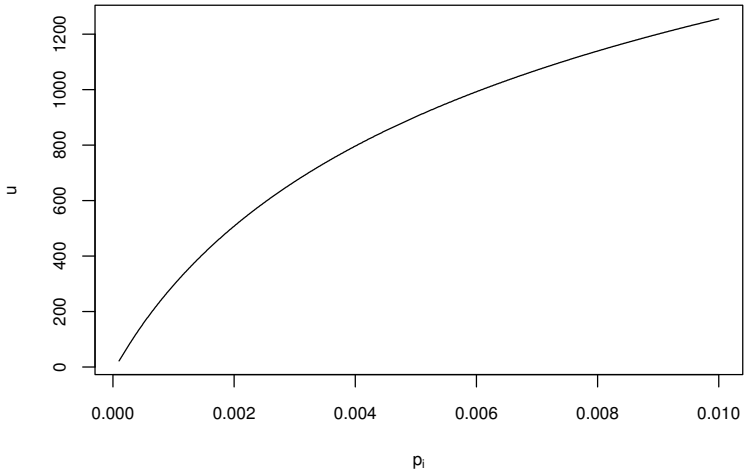

Figure 2.17: Level of initial capital u for which the individual miner obtains the same expected profit inside and outside of the pool, as a function of his computational power p_i.

of the risk to the pool and getting more frequent rewards.

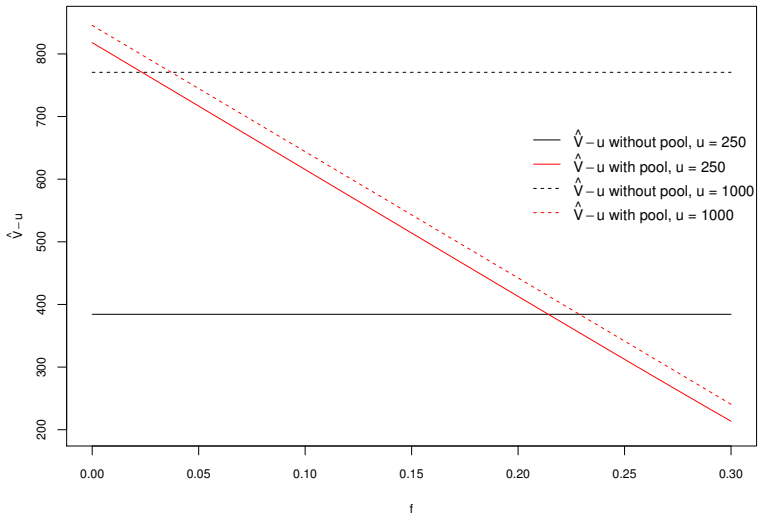

Figure 2.18: $\hat{V}(u,t) - u$ as a function of f for an individual pool miner alone in black and within the pool in red.

Figure 2.19 shows the level curves of $\hat{V}(u,t)$ with a varying difficulty for the miner's problem q and pool fee f. Note that not joining the pool is equivalent to setting the difficulty level equal to the block finding problem level and letting the pool fee be $f = 0$.

With such a two-way analysis, the pool can fix an appropriate fee and the miner can see whether he is better off joining the pool for his given level of capital u. In this graph, we depict the situation of one particular miner with his respective

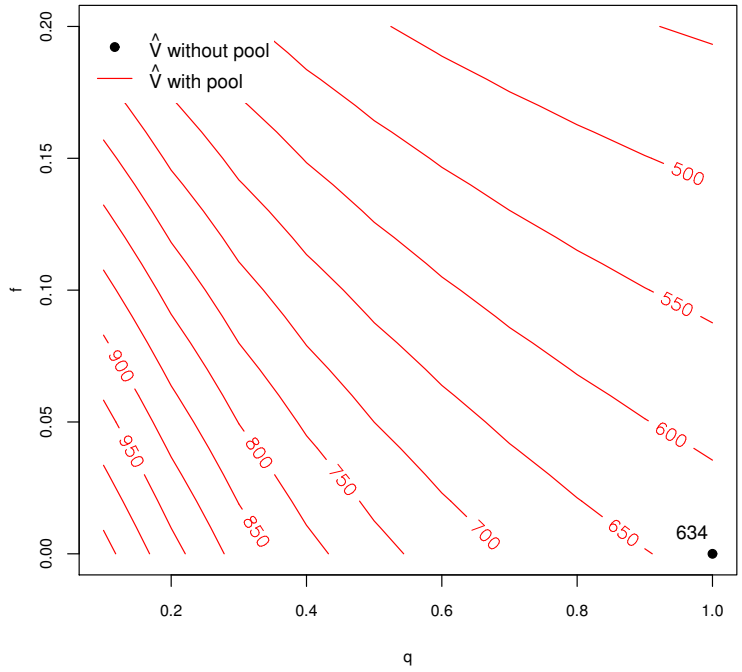

Figure 2.19: $\hat{V}(u,t)$ as a function of q and f for an individual pool miner alone in black and within the pool in red.

initial capital u. For another miner, the situation would be different and he might choose another more suitable pool. Therefore, multiple pools exist in the market to accommodate the various participants.

The miner's decision to join a *pay-per-share* mining pool does not depend on the size of the mining pool. Hence a miner will be indifferent whether to direct her hashpower towards a small or large pool. All that matters is the level of expected profit (decreasing in f) and the share of risk transferred to the mining pool (decreasing in q). Decentralization will prevail if the preferences, more specifically the risk aversion, of both the pool managers and the individual miners are sufficiently heterogeneous. Note that a situation where a mining pool would control most of the computing power is not desirable for anyone. The blockchain would then be prone to attacks and the associated cryptocurrency would no longer be of value.

2.8 Conclusion

In this chapter, we developed a framework for a bitcoin mining pool analysis from a risk and profitability perspective. Given a pay-per-share pooling scheme, we investigated the profitability of a pool under ruin probability considerations, which allows us to derive original results for the pool manager's expected profit. When describing the pool income process as a stochastic double-sided jump process, one can adapt techniques developed in the actuarial literature for applications in the blockchain universe. In addition, we also looked at the problem from the individual miner's

side, to identify conditions under which it is profitable for her to enter the pool or not.

We find that ignoring ruin considerations highly overestimates the expected gain for a pool for small values of initial capital and quantify the required capital level needed for which the ruin aspect becomes negligible. Moreover, we define a trade-off between the main pool defining parameters to set up conditions for optimizing the pool profit for different levels of capital. For an individual miner, pooling has similar effects as a reinsurance treaty for an insurer. We provide a sensitivity analysis that can be helpful for the miner to select the most appropriate pool given his initial parameters.

For a randomized time horizon, it was possible to obtain explicit formulas for all quantities of interest. The flexibility of our model enabled to consider deterministic as well as stochastic reward sizes. The established formulas for combinations of exponentials are in fact quite flexible, as any other distribution on the positive halfline can be approximated arbitrary well with such distributions (cf. [56]). Naturally, some restrictive assumptions were needed to enable the explicit mathematical treatment in this chapter, in particular the assumption of independent and identically distributed jump sizes. It will be interesting in future research to look into relaxing these assumptions.

The study of the formation of mining pools naturally raises the question of whether they pose a threat to the decentralized nature of blockchain-based applications. We find that the size of the mining pool does not interfere in a miner's decision making process. A miner chooses a mining pool according to the share of risk she wishes to cede and the profit she wishes to make. The preferences of miners and pool managers have been analysed using game theory in Cong et al. [37] and Li et al. [95]. The results of the present chapter may serve as concrete risk management tools for miners and pool managers that could also be integrated as value or cost functions within such a game-theoretic approach.

2.A Appendix A: Abel-Gontcharov polynomials

Let $U = \{u_i\ ,\ i \geq 1\}$ be a sequence of real non-decreasing numbers. The (unique) family $\{G_n(x|U)\ ,\ n \geq 0\}$ of *Abel-Gontcharov polynomials* of degree n in x attached to U is defined as follows. Starting with $G_0(x|U) = 1$, the polynomials $G_n(x|U)$ satisfy the differential equations

$$G_n^{(1)}(x|U) = n\, G_{n-1}(x|\mathcal{E}U), \tag{2.64}$$

where $\mathcal{E}U$ is the shifted family $\{u_{i+1}\ ,\ i \geq 1\}$, and with boundary conditions

$$G_n(u_1|U) = 0, \quad n \geq 1. \tag{2.65}$$

So, each G_n, $n \geq 1$, has the integral representation

$$G_n(x|U) = n! \int_{u_1}^{x} \left[\int_{u_2}^{y_1} \mathrm{d}y_2 \ldots \int_{u_n}^{y_{n-1}} \mathrm{d}y_n \right] \mathrm{d}y_1. \tag{2.66}$$

The polynomials G_n, $n \geq 1$, can be interpreted in terms of the joint distribution of the order statistics $(U_{1:n}, \ldots, U_{n:n})$ of a sample of n independent uniform random variables on $(0, 1)$. Indeed, for $0 \leq x \leq u_1 \leq \ldots \leq u_n \leq 1$, we have that

$$P[U_{1:n} \leq u_1, \ldots, U_{n:n} \leq u_n \text{ and } U_{1:n} \geq x] = (-1)^n\, G_n(x|u_1, \ldots, u_n).$$

This last identity is used inside the proof of Theorem 2.4.1 together with the following property

$$G_n(x|a + bU) = b^n G_n\left((x-a)/b\,|U\right), \quad n \geq 1. \tag{2.67}$$

Lastly, the numerical evaluation of (2.15) can rely on the recursive relations

$$G_n(x|U) = x^n - \sum_{k=0}^{n-1} \binom{n}{k} u_{k+1}^{n-k} G_k(x|U), \quad n \geq 1. \tag{2.68}$$

Formula (2.68) follows from an Abelian expansion of x^n based on (2.64), and (2.65).

2.B Appendix B: Proof of Theorem 4.1

The event $\{\tau \in (t, t+dt)\}$ can be viewed conditioned over the values of the process $(N_t)_{t \geq 0}$. In other terms,

$$\{\tau \in (t, t+dt)\} = \bigcup_{n=0}^{+\infty} \{\tau \in (t, t+dt)\} \cap \{N_t = n\}. \tag{2.69}$$

We distinguish according to the value of N_t. For $N_t = 0$, Equation (2.14) can be rewritten as

$$\tau = \inf\{t \geq 0; M_t^d > u/w\}, \tag{2.70}$$

which occurs when the $\lceil \frac{u}{w} \rceil^{th}$ jump of M_t^d occurs at t, where $\lceil x \rceil$ denotes the ceiling function. It follows that

$$\{\tau \in (t, t+dt)\} \cap \{N_t = 0\} = \{S^d_{\lceil \frac{u}{w} \rceil} \in (t, t+dt)\} \cap \{N_t = 0\} \tag{2.71}$$

and

$$f_{\tau|N_t=0}(t) = f_{S^d_{\lceil \frac{u}{w} \rceil}}(t),\ t \geq 0. \tag{2.72}$$

In case $N_t \geq 1$, one needs to constrain $\{M^d_t,\ t \geq 0\}$ so it does not reach $N_{u,s}w/(b-w) + u/(b-w)$ for any time $s < t$ but does so at t. Let $(v_n)_{n\geq 0}$ is a sequence of integers defined as $v_n = \lceil n(b-w)/w + u/w \rceil$, $n \geq 0$. We have

$$\{\tau \in (t, t+dt)\} \cap \{N_t \geq 1\} = \bigcup_{n=1}^{+\infty} \bigcap_{k=1}^{n} \{T_k \leq S^d_{v_{k-1}}\} \cap \{S^d_{v_n} \in (t, t+dt)\} \cap \{N_t = n\}, \tag{2.73}$$

as $M^d_t > \underbrace{N_t}_{=n}(b-w)/w + u/w$ at the time of the fatal jump (and before t, N_t reaches each step before the payout process surpasses it). Now

$$\begin{aligned} &\mathbb{P}\left[\{\tau \in (t, t+dt)\} \cap \{N_t \geq 1\}\right] \\ &= \sum_{n=1}^{+\infty} \mathbb{P}\left[\bigcap_{k=1}^{n} \{T_k \leq S^d_{v_{k-1}}\} \cap \{S^d_{v_n} \in (t, t+dt)\} \mid N_t = n\right] \mathbb{P}\left[N_t = n\right]. \end{aligned} \tag{2.74}$$

By the order statistic property, we get

$$\begin{aligned} &\mathbb{P}\left[\bigcap_{k=1}^{n} \{T_k \leq S^d_{v_{k-1}}\} \cap \{S^d_{v_n} \in (t, t+dt)\} \mid N_t = n\right] \\ &= \mathbb{P}\left[\bigcap_{k=1}^{n} \{U_{k:n} \leq F_t\left(S^d_{v_{k-1}}\right)\} \cap \{S^d_{v_n} \in (t, t+dt)\}\right] \\ &= \mathbb{P}\left[\bigcap_{k=1}^{n} \{U_{k:n} \leq F_t\left(S^d_{v_{k-1}}\right)\} \mid S^d_{v_n} \in (t, t+dt)\right] \mathbb{P}\left[S^d_{v_n} \in (t, t+dt)\right] \\ &= \mathbb{E}\left[(-1)^n G_n\left[0 \mid F_t\left(S^d_{v_0}\right), \ldots, F_t\left(S^d_{v_{n-1}}\right)\right] \mid S^d_{v_n} \in (t, t+dt)\right] \mathbb{P}\left[S^d_{v_n} \in (t, t+dt)\right], \end{aligned} \tag{2.75}$$

where $(U_{1:n}, \ldots, U_{n:n})$ denote the order statistics of n i.i.d. unit uniform r.v. and $G_n(.\mid .)$ denote the Abel-Gontcharov polynomials, see Appendix 2.A for a short presentation. Now take $F_t(s) = s/t,\ s \leq t$. In virtue of the property (2.67), we have

$$\begin{aligned} G_n\left[0 \mid F_t\left(S^d_{v_0}\right), \ldots, F_t\left(S^d_{v_{n-1}}\right)\right] &= G_n\left[0 \mid S^d_{v_0}/t, \ldots, S^d_{v_{n-1}}/t\right] \\ &= \frac{1}{t^n} G_n\left[0 \mid S^d_{v_0}, \ldots, S^d_{v_{n-1}}\right]. \end{aligned} \tag{2.76}$$

Inserting that last expression into (2.75) yields the announced result (2.15).

Chapter 3

Empirical risk analysis of mining a Proof-of-Work blockchain

This chapter is based on the following article:

H. Albrecher, D. Finger, and P.-O. Goffard. Empirical risk analysis of mining a Proof-of-Work blockchain. Submitted, [7].

Abstract. The process of mining blocks on a blockchain utilizing a Proof-of-Work consensus mechanism carries inherent risks, particularly when the operational expenses associated with mining exceed the rewards earned. Building on previous findings on mining in pools, this chapter delves into the question of whether the theoretical formulas for the ruin probability and the expected value of future surplus obtained under particular model assumptions are indeed validated empirically. In particular, we include the presence of transactions fees in the block rewards in our analysis. We also provide algorithms to fit the involved generalized hyperexponential distributions to actual data. Moreover, we perform a sensitivity analysis for different factors of interest, and we quantify the relevance of incorporating temporal dependence and transaction fees in the model.

3.1 Introduction

A blockchain is a data ledger which is maintained by a *Peer-to-Peer* network. The database entries, referred to as transactions, are recorded by batches called *blocks* resulting from the application of a consensus protocol. In the case of the bitcoin, the consensus protocol is the *Proof-of-Work*. The network participants, called *miners*, compete to solve a cryptoproblem via a trial and error approach. Computers are running 24/7 which consumes a lot of electricity. This operational cost is borne by the miners and compensated by a reward expressed in cryptocurrency units whenever a new block is found. The stability of blockchain systems relies heavily on this incentive mechanism. The balance between cost and reward is a stochastic process, which we model as

$$R_t = u - C_t + B_t,\ t \geq 0, \tag{3.1}$$

where the initial capital u is augmented by the income $(B_t)_{t\geq 0}$ net of the expenses $(C_t)_{t\geq 0}$. Such models were studied in [10, 8, 95] assuming that

$$C_t = c \cdot t,\ \text{for } c, t \geq 0,$$

and

$$B_t = \sum_{i=1}^{N_t} U_i,\ t \geq 0,$$

where $(N_t)_{t\geq 0}$ is a Poisson process and the block rewards B_i form a sequence of positive random variables. One goal of this chapter is to propose a more accurate model for these rewards.

Block rewards consist of a protocol-specified bounty augmented by transactions fees. When passing a transaction, users will typically attach a transaction fee to it. This fee mostly depends on the transaction volume, since each block has a fixed allotted space inside. The pending transactions are stored in the `memory pool` (often abbreviated as `mempool`), where they await confirmation, hereby forming a queue. The transaction fee level closely relates to the network congestion that may be tracked by looking at the `mempool` size. Algorithms have been developed to inform the users of the appropriate transaction fee levels, see for instance the book of Antonopoulos [15, Chapter 6]. Although common practice suggests that miners prioritize transactions with the highest transaction fee per byte rate, some deviations can be observed due to potential arbitrage opportunities linked to include specific transactions or ordering them in a given manner. Such considerations are beyond the scope of the present study, and we refer the reader to the work of Messias et al. [105]. In the bitcoin blockchain, the impact of the transaction fees is currently still minor when compared to the bounty for finding a new block. However, as the reward gets halved approximately every four years, the need to understand the underlying dynamics of transaction fees will become pivotal in the future. In Carlsten et al. [31], the authors envision the stability of the system when the block reward reduces to the transaction fees. Möser & Böhme [107] analyse the main drivers of the fees and conclude that higher fees lead to faster transaction processing. Easly et al. [60] link the proportion of zero-fee transactions to the bitcoin price and memory pool size through a linear model. Tedeschi et al. [136] build a neural net that outputs the

probability for a transaction to be included based on the transaction features. Finally, Rossi et al. [122] consider a queueing model to estimate the confirmation time.

The block reward is expressed in crypto-currency units, but the operational cost in (3.1) is likely to be expressed in fiat currency. The question of modelling the exchange rate of crypto against fiat currencies naturally arises when studying the profit and losses of blockchain miners via model (3.1). The evolution price of cryptocurrencies has been extensively studied in the literature. Ciaian et al. [36] study the bitcoin price formation incorporating market information, such as the Dow Jones stock market index or the oil price. Bouoiyour et al. [26] decompose the bitcoin price index using Empirical Mode Decomposition, which is similar to signal-processing techniques, but does not assume periodicity, see e.g. [158]. Many authors have applied neural network techniques to fit and predict bitcoin prices. For instance, Almeida et al. [12] use Artificial Neural Networks (ANN) and find that trading volumes are irrelevant. McNally et al. [104] use Recurrent Neural Networks (RNN) of Long Short Term Memory type to accomododate their three-year long dataset. Time series models like GARCH [86] and ARIMA [19, 153] have also been considered.

We will pursue two modelling strategies in this chapter. First, we assume that the block rewards are independent and identically distributed (i.i.d.) random variables with a generalized hyperexponential (GH) distribution, studied in Botta et al. [25], also referred to as the combination of exponentials model by Dufresne [56]. The probability density function of a GH distribution is a linear combination of exponentials which does not need to be convex (in contrast to mixture of exponential distributions). The GH class is a proper subset of matrix-exponential (ME) distributions which are probability distributions with rational Laplace transform, see already Cox [39] and Bladt & Nielsen [23] for a recent overview. The GH class leads to tractable calculations and is itself already dense in the set of probability distributions on the positive half-line. For applications in insurance risk theory, see e.g. Lin & Willmot [96, 97]. In that framework, (3.1) is called the dual model of the standard insurance risk model, as the wealth process performs upward jumps and decreases linearly in time. In our previous work [8], we found closed-form expressions for the ruin probability and the expected profit when the rewards are GH distributed. In this chapter we want to go one step further and fit such a distribution to actual bitcoin data.

Since the construction is not probabilistic, a priori one can not guarantee that a particular fitted set of parameters for a GH (and more generally ME) distribution represents a proper probability distribution. Conditions based on the roots and poles of the Laplace transform have been given in the works of Bean et al. [22] and Fackrell [62]. These conditions apply to ME distributions and were integrated in Fackrell [61] within a maximum likelihood estimation procedure. Dufresne [56] used GH distributions to approximate any distribution on the positive half-line via an expansion in terms of Jacobi polynomials. The non-negativity problem is addressed then by expanding the square root of the probability density. In this work, we explore the polynomial approach to fit data to a GH distribution. The non-negativity of the resulting probability density function is checked via a bisection procedure suggested in the work of Hanzon & Holland [73]. To the best of our knowledge, this chapter is

the first to fit combinations of exponentials to actual data in a non-parametric way instead of approximating a predefined distribution. The results may therefore also be applicable in other modelling contexts where such a distributional assumption is assumed, see e.g. [9] for an example in ruin theory.
The assumption of stationarity, a key premise in our first approach, might be considered too restrictive, as block rewards are influenced by transaction fees and fluctuations in cryptocurrency prices. Consequently, we adopt a second approach that views block reward data as an ARIMA time series. After fitting it to the data, we generate scenarios that enable us to estimate risk and performance indicators through Monte Carlo simulations. We then conduct a numerical comparison between the results obtained under this time series framework and the ones based on the i.i.d. assumption in our first approach, together with a series of further sensitivity results on various assumptions underlying the model for which a tractable formula for the key quantities is available.

The rest of the chapter is organized as follows. Section 3.2 provides a brief description of our data together with reminders about the way that Proof-of-Work blockchains operate. Section 3.3 describes the combination of exponentials model and our block reward distribution function estimators. The proposed estimation is first back-tested on synthetic data, before it is applied to the actual block reward data. Section 3.4 presents the result of our time series analysis. Section 3.5 compares the two modelling approaches, looking at their impact on the profitability and ruin of blockchain miners. It also contains sensitivity tests with respect to the inclusion of transaction fees in the modelling as well as the electricity price. Section 3.6 concludes.

3.2 Transaction fee concepts and descriptive data analysis

In this section, we will analyse empirical data on the transaction fees in the bitcoin cryptocurrency. Let us first give a short reminder on the definition of transaction fees and their importance in the mining process.

In the bitcoin *Proof-of-Work* verification algorithm, each miner solves a cryptoproblem in order to validate a block. Whenever a block is validated or *mined*, the miner receives the corresponding block reward set to 6.25 BTC as of November 2023. In addition, the sum of transaction fees attached to all the transactions included in this block are also given as an additional reward to the miner. As part of the validation system, the miner obtains the authorization from the community to choose which transactions from the *memory pool* ("waiting line") are entering the newly mined block. The user will typically attach a transaction fee to the required payment. This fee mostly depends on the transaction volume, since a block has a fixed allotted space inside. It also depends on the market congestion and the individual user's decisions [15]. There exist algorithms that help the user choose an appropriate fee. For example on `https://privacypros.io/tools/bitcoin-fee-estimator/`, one can estimate the expected confirmation time as a function of the chosen fee. Apart from the congestion, which can be studied by looking at the memory pool,

another determinant for the choice of transaction fees may be the electricity price and the bitcoin price. In fact, when the bitcoin price is high, miners may have more incentive to mine at a larger scale and thus speed up the inclusion of the transactions into the blocks, hence reducing the congestion. Another factor one may want to consider is the electricity price. When the electricity price is high, some miners will reduce their activity of even go out of business. This may raise the congestion and drive up the transaction fees. For this, including a stochastic cost process in the model (3.1) would be necessary, but is out of scope of this chapter. Moreover, this chapter focuses on very short (2 weeks on average) time horizons, in which electricity price variations can considered minor or being fixed via contracts with the provider. In the mining process, the miners have to decide which transactions to incorporate in the block, since it has limited memory space. They will consider the priority of the transaction based on the attached fee size per byte, but recent research shows that there may be other (selfish) interests in promoting transactions [105]. It is important to note again that with the scheduled halving of the fixed block rewards every four years, the share of the transactions fees in the total rewards will gain decisive importance over time. Thus, including and modelling this stochastic part will become even more relevant in the future.

For this analysis, we use publicly available data. Scrapping bitcoin-related data is possible by making calls to some APIs. This was done by looping through necessary blocks. For example, to access block number 650000 information, one can follow the address `https://chain.api.btc.com/v3/block/650000`[1]. It extracts information in JSON format, which can then be reformatted to our needs. With this method, we gather data for the period spanning from the last halving of the bitcoin reward on May 12, 2020 until September 16, 2021. For this time frame, we collect the following information:

- The transaction fees per block (Fees in BTC);
- The exchange rate BTC-USD per minute (Price in USD);
- The size of the memory pool of the transactions in bytes (daily) (Mempool size);
- The number of transactions in the memory pool (daily) (Mempool count);
- The current difficulty of the cryptoproblem (adjusts bi-weekly) (Difficulty).

Figure 3.1 depicts the extracted data and Table 3.1 contains some statistics of the dataset. One can observe a strong correlation between the transaction fees and the number of transactions in the memory pool, which is confirmed in Table 3.2.

For the analysis in this chapter, we consider a smaller sub-sample of the data containing 10'000 data points ranging from February 10, 2021 to April 21, 2021. Also, we merge the series of fees and prices by converting fees in USD, since we aim to model the latter. The covariate `"difficulty"` is not retained as it has a bi-weekly update and our focus will be on a two weeks horizon. For this chapter, we also do not further consider the strongly correlated variables `"mempool size"` and

[1] last accessed on 12/09/2023

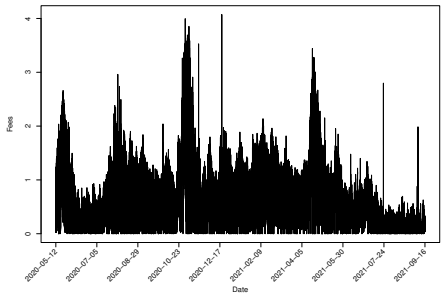

(a) Fees

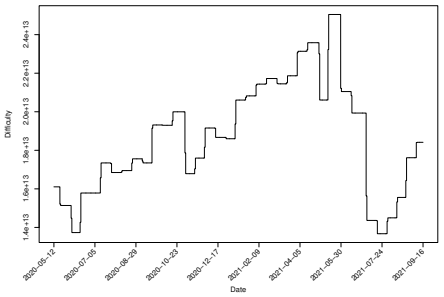

(b) Difficulty

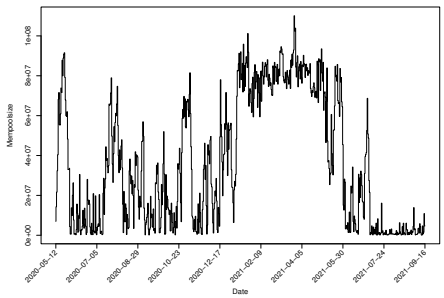

(c) Mempoolsize

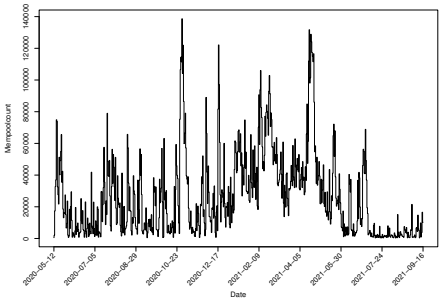

(d) Mempoolcount

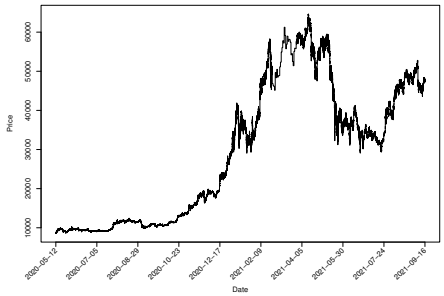

(e) BTC/USD price

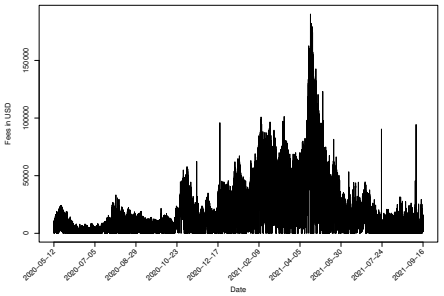

(f) Fees in USD

Figure 3.1: Illustration of the data.

Table 3.1: Main statistics of the dataset.

	Block	Date	Fees in BTC	Price in USD	Mempool size	Mempool count	Difficulty
Mean			0.594	29'535	36'090'588	25'610	1.86E+13
Min	630'014	12.05.2020	0	8'584	120'928	290	1.37E+13
25^{th} pct.			0.1731	11'324	3'889'528	4'757	1.68E+13
Median			0.4652	31'700	27'446'996	18'670	1.86E+13
75^{th} pct.			0.8872	46'331	70'698'937	38'942	2.08E+13
Max	700'851	16.09.2021	4.074	64'617	110'094'398	138'640	2.51E+13

"`mempool count`", because fitting a time series with a covariate implies modelling the dynamics of this covariate independently. For the purpose of this chapter, we opt for modelling only one time series only, the dynamics of which then aggregates all the dynamics playing a role for the output variable. As illustrated in Figure 3.2, the time series of fees exhibits strong autocorrelation, so that a stationarity assumption would not be appropriate. However, by taking the first difference, one

Table 3.2: Pearson correlation matrix.

	Fees in BTC	Mempool size	Mempool count	Difficulty	Price
Fees in BTC	1	0.4885764	0.6361694	0.3313047	-0.01781731
Mempool size	0.48857638	1	0.7462129	0.6715088	0.43744576
Mempool count	0.63616939	0.7462129	1	0.5342133	0.25523618
Difficulty	0.33130472	0.6715088	0.5342133	1	0.54365471
Price	-0.01781731	0.4374458	0.2552362	0.5436547	1

can find back stationarity to a large extent, see the right-hand side of Figure 3.2; only one lag is still significant for the fees. We will therefore consider only one lag when fitting a time series model in Section 3.4.

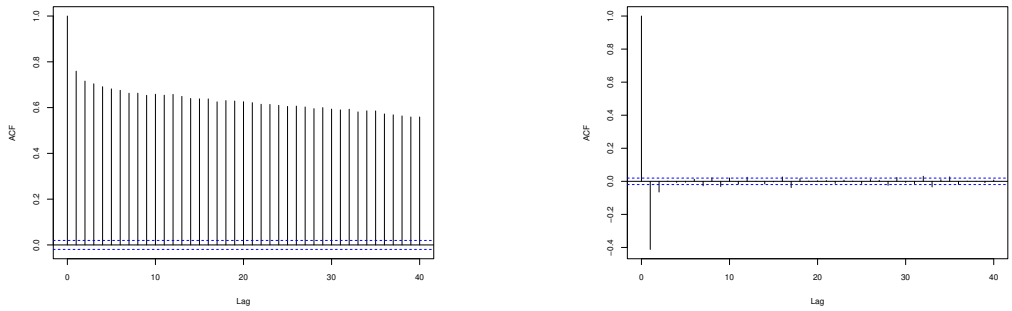

Figure 3.2: Autocorrelation function of fees in USD (left) and its first differences (right).

3.3 Block reward as a combination of exponentials

A random variable U has a generalized hyperexponential distribution if its cumulative distribution function (CDF) is given by

$$F_U(x) = 1 - \sum_{i=1}^{d} a_i e^{-\lambda_i x}, \text{ for } x \geq 0, \tag{3.2}$$

where $\lambda_1, \ldots, \lambda_d > 0$ and $a_1, \ldots, a_d \in \mathbb{R}$ with $\sum_{i=1}^{d} a_i = 1$. In the sequel, we assume that $\lambda_1 < \ldots < \lambda_d$. Define the vectors

$$\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_d), \boldsymbol{a} = (a_1, \ldots, a_d),$$

and the diagonal matrix $\Lambda = \text{diag}(\lambda)$, so that

$$F_U(x) = 1 - \boldsymbol{a} \cdot e^{-\Lambda x} \cdot \mathbf{1}_d, \text{ for } x \geq 0, \tag{3.3}$$

where $\mathbf{1}_d = (1, \ldots, 1)$. The probability density function (PDF) is given by

$$f_U(x) = \boldsymbol{b} \cdot e^{-\Lambda x} \cdot \mathbf{1}_d, \text{ for } x \geq 0, \tag{3.4}$$

with $\boldsymbol{b} = \boldsymbol{a} \cdot \Lambda$.

The Laplace transform of U

$$\mathbb{E}(e^{-\theta U}) = \boldsymbol{b} \cdot (\Lambda + \theta I_d) \cdot \mathbf{1}_d,\ \theta \geq 0, \tag{3.5}$$

where I_d is the identity matrix, is rational which implies that combinations of exponential distributions are instances of matrix-exponential distributions.

3.3.1 Non-negativity of GH probability distribution functions

Given a set of parameters $\boldsymbol{a}$ and $\boldsymbol{\lambda}$, there is no straightforward way to ensure that (3.3) is a proper CDF or that (3.4) is a proper PDF. Some characterization, based on the Laplace transform, have been provided by Bean et al. [22] and Fackrell [62] for ME distributions. We take a different road here, following up on the work of Hanzon and Holland [73]. Consider the function

$$f(x) = \boldsymbol{b} \cdot e^{-\Lambda x} \cdot \mathbf{1}_d, \quad x \geq 0. \tag{3.6}$$

In order for f to be a proper PDF, we need to ensure

$$f(x) > 0,\ \forall x > 0, \text{ and } \int_0^\infty f(x)\mathrm{d}x = 1. \tag{3.7}$$

Necessary conditions for (3.7) to hold include $b_1 > 0$ and $\boldsymbol{b} \cdot \Lambda^{-1} = \mathbf{1}_d$, but the latter are not sufficient. We therefore use a verification method via bisection, suggested in Hanzon and Holland [73]. We define the sequence

$$f_0(x) = b \cdot e^{-\Lambda x} \cdot \mathbf{1}_d = f_U(x) \text{ and } f_k(x) = b \cdot \prod_{i=1}^{k} (\lambda_j \cdot I_d - \Lambda) \cdot e^{-\Lambda x} \cdot \mathbf{1}_d, \text{ for } k = 1, \ldots, d.$$

It is characterized by the following property: for $x \in [0, M]$, the function f_k has at most one sign-changing zero between two sign-changing zeros or boundary points of f_{k+1} for $k = 0, 1, \ldots, d-1$. As $f_d(x) = 0$, one can recursively, starting from $f_d(x)$, check the presence of sign-changing points on a closed interval through a bisection procedure.

3.3.2 Fitting GH distributions to data via polynomial expansions

Dufresne [56] presents a method to approximate any PDF of a distribution on the positive half-line via a combination of exponentials. The approximation formula takes the form of an expansion in terms of the shifted Jacobi polynomials defined as

$$R_k^{(\alpha,\beta)}(x) = \sum_{j=0}^{k} \rho_{k,j} x^j, \tag{3.8}$$

where

$$\rho_{k,j} = \frac{(-1)^k (\beta+1)_k (-k)_j (k+\alpha+\beta+1)_j}{(\beta+1)_j k! j!}, \tag{3.9}$$

where $(z)_k = z \cdot (z+1) \cdot (z+2) \cdots (z+k-1)$ denotes the Pochhammer symbol. These polynomials are orthogonal on $[0, 1]$ w.r.t. the weight function $\phi(x) = (1-x)^\alpha x^\beta$.

Any function $g : (0,1) \mapsto \mathbb{R}$, square integrable w.r.t. $\phi(x)$ can be expanded as a shifted Jacobi polynomial expansion with

$$g(x) = \sum_{k=0}^{\infty} c_k R_k^{(\alpha,\beta)}(x), \tag{3.10}$$

where

$$c_k = \frac{1}{h_k} \int_0^1 g(x)(1-x)^\alpha x^\beta R_k^{(\alpha,\beta)}(x)\mathrm{d}x, \tag{3.11}$$

and

$$h_k = \int_0^1 (1-x)^\alpha x^\beta R_n^{(\alpha,\beta)}(x)^2 dx = \frac{\Gamma(n+\alpha+1)\Gamma(n+\beta+1)}{(2n+\alpha+\beta+1)n!\Gamma(n+\alpha+\beta+1)}. \tag{3.12}$$

The convergence in (3.10) takes place in the L^2 sense, see for instance the book of Nagy [134, Ch.7]. Our target is a PDF f on the positive half-line. Following Dufresne [56], we expand the function $f^*(x) = e^{prt} f(x)$ for some $p \in \mathbb{R}$ and $r > 0$ and use the change of variable

$$g(x) = f^* \left[-\frac{1}{r} \log(x) \right],$$

that maps the interval $(0,1)$ onto $(0,\infty)$. The expansion of f is then

$$f(x) = e^{-rpt} \sum_{k=0}^{\infty} c_k R_k^{(\alpha,\beta)}(e^{-rpt}), \tag{3.13}$$

and the square integrability condition on g translates directly to

$$\int_0^\infty e^{-(\beta+1-p)rt}(1-e^{-rt})^\alpha f^2(t)\mathrm{d}t < \infty.$$

The coefficients of the polynomial expansion can also be expressed as an integral in terms of f as

$$c_k = \frac{r}{h_k} \int_0^\infty e^{-(\beta-p+1)rt}(1-e^{-rt}) R_n^{(\alpha,\beta)}(e^{-rt}) f(t) dt. \tag{3.14}$$

A simple truncation of the infinite series in (3.13), followed by a normalization so that it integrates to 1, yields an aproximation

$$f(t) \approx e^{-rpt} \sum_{k=0}^{d-1} c_k R_k^{(\alpha,\beta)}(e^{-rt}), \tag{3.15}$$

which is consistent with what is referred to as Method A in Dufresne's work [56]. Replacing the polynomials $R_k^{(\alpha,\beta)}(x)$ by $\sum_{j=0}^k \rho_{k,j} x^j$ in (3.15) yields

$$f(t) \approx \sum_{j=1}^{d} \sum_{k=j-1}^{d-1} c_k \rho_{k,j-1} e^{-(j-1+p)rt},$$

which is a combination of exponentials as in (3.4) where

$$\mathbf{b} := \begin{pmatrix} \sum_{k=0}^{d-1} c_k\rho_{k,0} & \dots & c_{d-1}\rho_{d-1,d-1} \end{pmatrix} \text{ and } \lambda := \begin{pmatrix} pr & (1+p)r & \dots & (n+p)r \end{pmatrix}.$$

Remark 3.3.1. *The approximation method involves selecting parameters α, β, p, r, and determining the truncation order d. In the absence of established selection guidelines, we draw upon the parameter values utilized in Dufresne's work [56]. The choice of the truncation order should strike a balance between accuracy, computational efficiency, and numerical stability, favoring larger values wherever possible. When considering f^* instead of f, it ensures that $f(t)$ approaches zero as t tends towards infinity for any truncation order, provided that $0 < p < (\beta + 1)/2$.*

For our application, we do not have a known distribution function to approximate, but a dataset to fit. Assume that $\{x_1, \ldots, x_M\}$ form an i.i.d. sample of size M. We can replace the expansion coefficients defined in (3.14) by their empirical counterpart with

$$\widehat{c}_k = \frac{r}{h_k M} \sum_{m=1}^{M} e^{-(\beta-p+1)rx_m}(1 - e^{-rx_m}) R_k^{(\alpha,\beta)}(e^{-rx_m}), \text{ for } k = 0, \ldots, d-1. \quad (3.16)$$

An a posteriori control with the help of the bisection method from Section 3.3.1 can ensure the non-negativity of the estimated PDF

$$\widehat{f}_M(t) = e^{-rpt} \sum_{k=0}^{d-1} \widehat{c}_k R_k^{(\alpha,\beta)}(e^{-rt}). \quad (3.17)$$

The estimated PDF is a nonparametric density estimator relying on orthogonal functions, a method detailed in [148, Ch.8]. A recognized limitation of this approach is its susceptibility to occasional negative values stemming from sampling errors. In instances where our estimates exhibit negativity, we can employ what Dufresne [56] terms as 'Method B' as a corrective measure. Instead of expanding $e^{prt} f(t)$, consider expanding

$$\tilde{f}(t) = e^{prt} \sqrt{f(t)}.$$

We get an approximation formula of the form

$$\sqrt{f(t)} \approx e^{-prt} \sum_{k=0}^{d-1} c_k R_k^{(\alpha,\beta)}(e^{-rx}) = \sum_{j=1}^{d} b_j e^{-\lambda_j t}, \quad (3.18)$$

and finally squaring it yields an approximation of f which is a proper PDF after normalization:

$$f(t) \approx \sum_{j=1}^{d} \sum_{k=1}^{d} b_j b_k e^{-(\lambda_j+\lambda_k)t} = \sum_{m=1}^{2d-1} \tilde{b}_m e^{-(m-1+2p)rt},$$

with $\tilde{b}_m = \sum_{j=1}^{m} b_j b_{m+1-j}$ and $b_{j>d} := 0$, since $\lambda_j = (j-1+p)r,\ j = 1, 2, \ldots, d$. The coefficients of the polynomial expansion of $\sqrt{f}$ are given by

$$b_k = \frac{r}{h_k} \int_0^{\infty} e^{-(\beta-p+1)rt}(1 - e^{-rt}) R_k^{(\alpha,\beta)}(e^{-rt}) \sqrt{f(t)} dt. \quad (3.19)$$

To get a statistical estimation of the coefficients, we replace f in (3.19) by a kernel density estimator

$$\widehat{f}_h(x) = \frac{1}{M} \sum_{m=1}^{M} K_h(x - x_m) = \frac{1}{Mh} \sum_{m=1}^{M} K\Big(\frac{x - x_m}{h}\Big), \quad (3.20)$$

where $K(x)$ is the Gaussian kernel.

Remark 3.3.2. *In contrast to Method A, Method B guarantees a valid PDF. However, this advantage comes at the expense of significantly increased computational complexity due to doubling the number of terms (with a numerical integration of the kernel density estimator for each).*

3.3.3 Simulation study

We illustrate our fitting procedure for a combination of exponentials through a brief simulation study. Draw m samples $(x_1, \ldots, x_m)$ from a right-shifted gamma random variable $X = \gamma + Y$, where $\gamma > 0$ and Y has PDF

$$f(x) = \frac{\delta^r x^{r-1} e^{-\delta x}}{\Gamma(r)}, \ x > 0, \ r, \delta > 0. \tag{3.21}$$

The choice of this shift is motivated by the observed shape of the empirical distribution in our collected block reward data later. We set the parameters to $r = 3, \delta = 0.5$ and $\gamma = 5$. Figure 3.3 shows the fit (in red) of the combination of exponentials distribution to the data when using Method A with parameters $\alpha = 0, \beta = 0, r = 0.01, p = 0.9, d = 20$ for samples of sizes $n \in \{100, 1000, 10000, 100000\}$. The quality of the fit enhances as the sample size increases. However, note that also, for larger sample size, the density estimate occasionally exhibits negative values.
Figure 3.4 shows the fit (in red) of the combination of exponentials distribution to the data when using Method B with parameters $\alpha = 0, \beta = 0, r = 0.01, p = 0.9, d = 20$, for samples of sizes $n \in \{100, 1000, 10000, 100000\}$. The fit is less good when using Method B; however, it consistently results in a valid probability distribution. Due to this essential property, we have chosen to exclusively employ Method B in our application to the block reward data in the sequel. Method B is centred around the square root of the PDF; in the Appendix 3.A we perform an exploratory analysis on the approximations using higher order of roots.

3.3.4 Real data application

The block reward comprises two components: the reward for discovering a new block, which currently stands at BTC 6.25 as of November 2023, and the transaction fees detailed in Section 3.2. Given that miners typically operate within a fiat currency framework, such as USD, we apply the exchange rate applicable at the moment of block discovery. Figure 3.5 depicts the histogram of total rewards received by miners, presented both in BTC and USD. The data is indeed shifted away from zero due to the fixed block reward addition, which poses a challenge for parametric fitting methods, as the support of exponential random variables is the entire positive halfline. The non-parametric approach, utilizing a polynomial expansion, provides an advantage by capturing variations in the central mass, even in the absence of data points in the lower tail.

Figure 3.6 shows the fit of the combination of exponentials using Method B with two parameterizations.

As the number of terms increases, the central part of the density aligns more closely with the actual data, albeit with an increase in instability near the left end. Conversely, reducing the number of terms results in a smoother shape, but the original density's spikes are less pronounced. Renormalization alleviates this effect

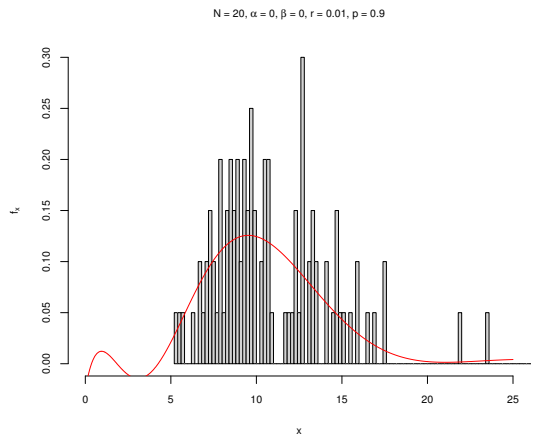

(a) Sample size: 100.

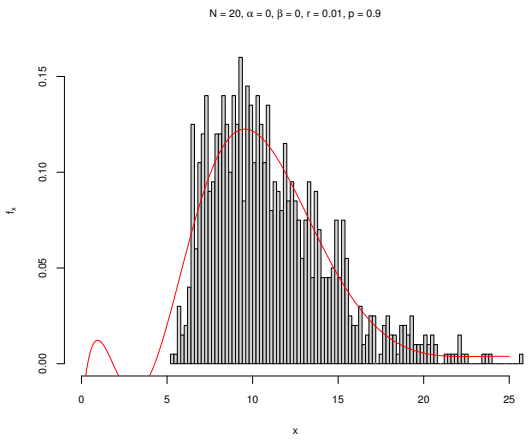

(b) Sample size: 1000.

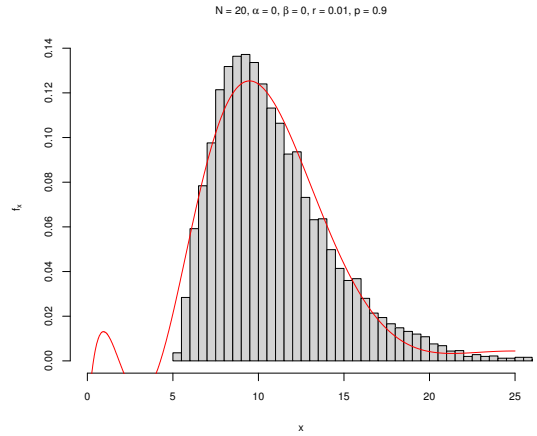

(c) Sample size: 10000.

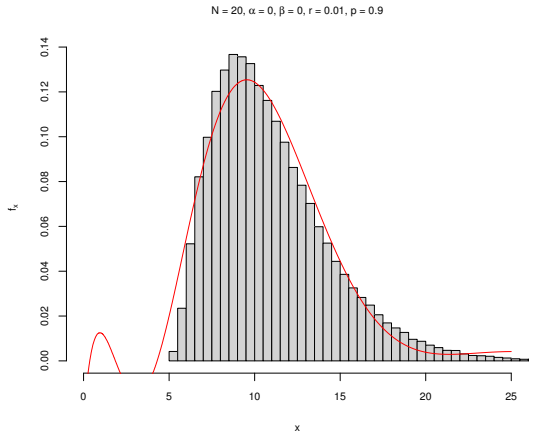

(d) Sample size: 100000.

Figure 3.3: Fitting of the shifted Gamma random sample by modified Method A.

to some extent (cf. the green line), and we choose to use the parametrization with $d = 10$ for the numerical analysis later.

3.4 Block rewards as time series

Since we also want to test the model for its sensitivity to non-stationarities, we fit the block reward data to an ARIMA model calibrated using the Box and Jenkins optimization method. Recall that a time series X_t is ARIMA(p, d, q) if $\nabla^d X_t$ is an ARMA(p, q) process, where ∇^d is the d^{th} difference operator. An ARMA(p, q) time series is a stationary process defined as

$$X_t = \phi_1 X_{t-1} + \cdots + \phi_p X_{t-p} + z_t + \theta_1 z_{t-1} + \cdots + \theta_q z_{t-q}, \tag{3.22}$$

with $\phi_p \neq 0$, $\theta_q \neq 0$ and z_t is white noise with mean 0 and variance σ_z^2, see e.g. [129]. Note that time series modeling is not a focus of this chapter, we refer to other papers for this purpose, see e.g. [19, 153]. Our choice of a simple ARIMA model is motivated by its suitability for generating plausible scenarios over a short time

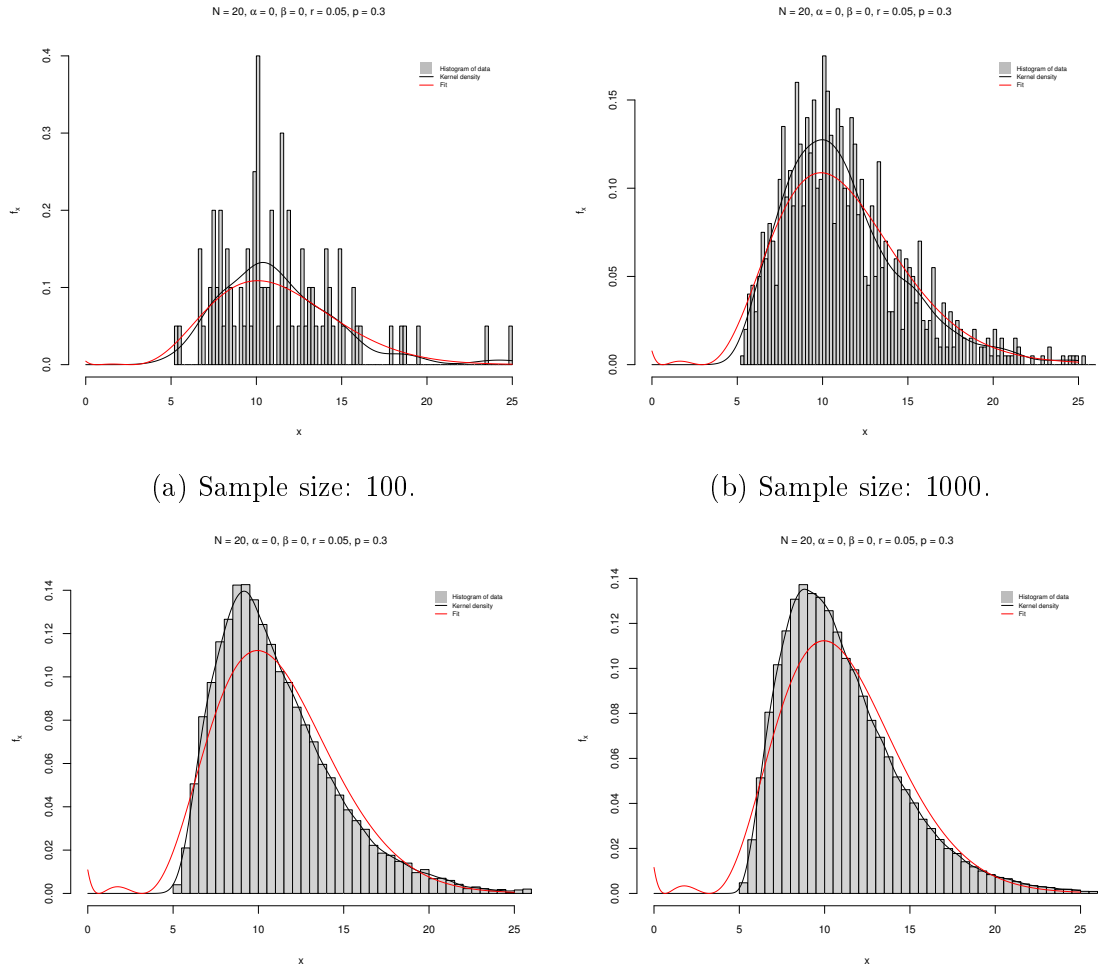

(a) Sample size: 100.

(b) Sample size: 1000.

(c) Sample size: 10000.

(d) Sample size: 100000.

Figure 3.4: Fitting of the shifted Gamma random sample by modified Method B.

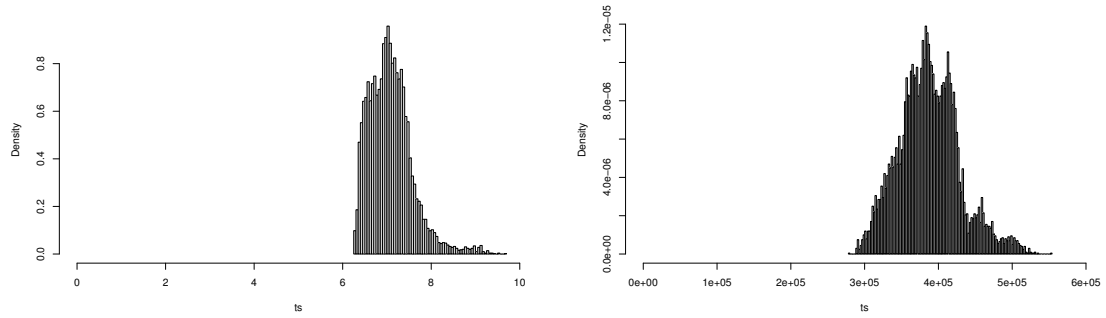

(a) Rewards expressed in BTC.

(b) Rewards expressed in USD.

Figure 3.5: Histogram of total rewards, period from February 10, 2021 to April 21, 2021.

horizon while departing from the i.i.d. assumption of Section 3.3. Also, we opted against more complex GARCH models due to their tendency to overfit the data

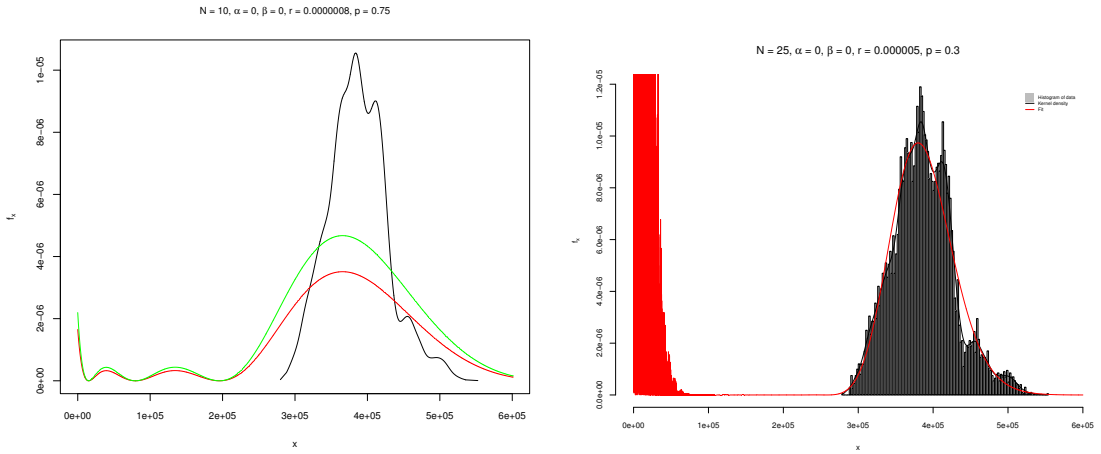

(a) $d = 10$ $(p = 0.75)$ (b) $d = 25$ $(p = 0.3)$

Figure 3.6: Approximation of the bitcoin rewards sample (Method B, $\alpha = 0, \beta = 0, r = 0.0000008$)

and their requirement for an extensive number of lags to ensure reliability. Hence, we deliberately select a simple ARIMA models as a pragmatic and reliable choice to conduct our risk analysis over a two weeks time horizon. For fitting our model, we consider a data set from a time frame between February 10, 2021 and April 6, 2021 that we split into a training set and a consecutive test set for checking the data prediction. In Table 3.3 we summarize the obtained results. The fitted model suggests indeed a once differentiated series.

Table 3.3: Summary of ARIMA model.

	ar1	ma1	ma2	ma3	ma4	ma5	ma6	ma7
	-0.7799	0.0989	-0.6543	-0.1146	-0.0337	-0.0330	-0.0075	-0.0267
s.e.	0.1122	0.1125	0.0771	0.0194	0.0136	0.0138	0.0123	0.0122
σ^2	219226983							
Log.Lik.	-88159.57							
AIC	176337.1							
BIC	176400							

In Figure 3.7 we show an illustration of the fitted model. On the left-hand side, one can see the fitted data points in comparison to the historical values of the training sample. On the right-hand side, we show simulated paths in different colors and the true historical values from our test sample in black. In addition to the statistical fit (cf. Figure 3.8 for a normal Q-Q plot for the remaining residuals), the fit also seems quite satisfactory visually, which is remarkable for the case of only a few free parameters.

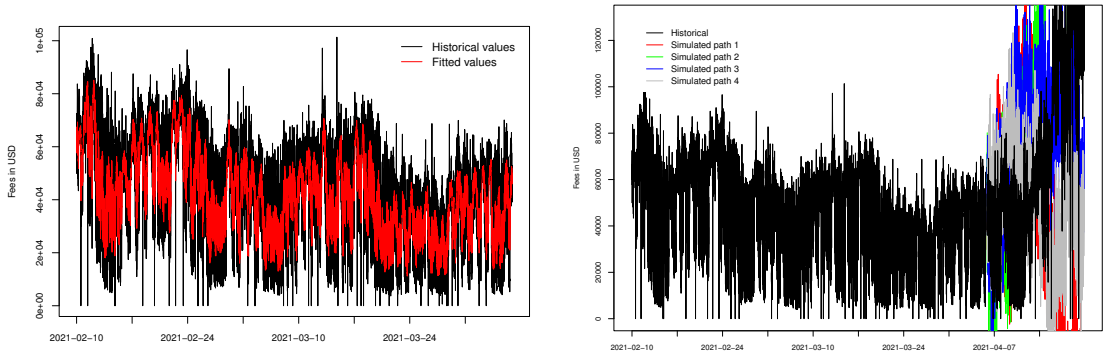

(a) Fitted values. (b) Simulated rewards per block in USD.

Figure 3.7: ARIMA(1,1,7) fit to block reward data.

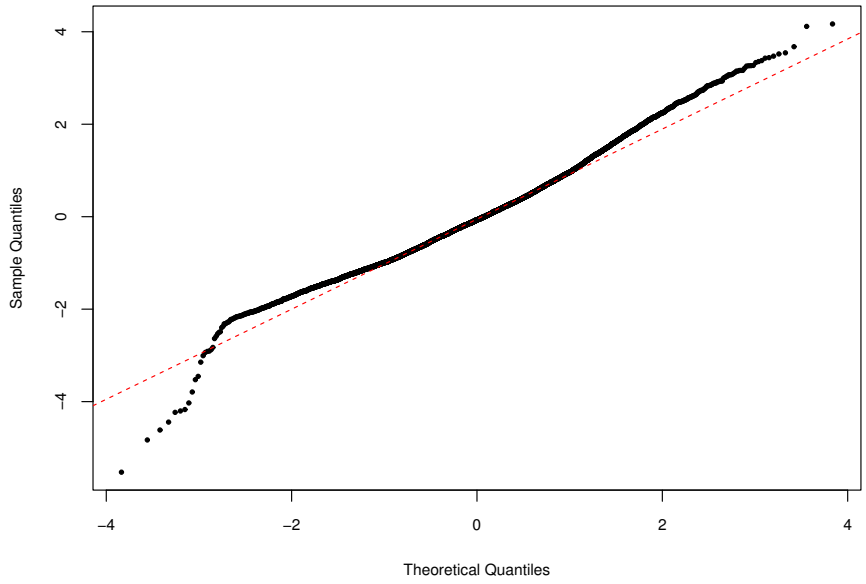

Figure 3.8: Normal Q-Q plot of ARIMA(1,1,7) residuals.

3.5 Comparison of the two modelling approaches in terms of risk

In [8, 10], explicit formulas were derived for the ruin probability and the expected surplus of a miner, taking into account the option of participating in a mining pool. Assume the miner's surplus has the form

$$R_t^i = u - c_i \cdot t + \sum_{i=1}^{N_t} U_i, \ t \geq 0,$$

where $N_t \sim Poisson(p_i \mu t)$ and the U_i's are i.i.d. with PDF $f_U(x) = \sum_{j=1}^{d} a_j \lambda_j e^{-\lambda_j x}$, $x > 0$. In [8], it was shown that under the assumption of i.i.d. rewards of GH

type, the ruin probability up to an exponential time horizon (with mean t) can be expressed as

$$\hat{\psi}(u,t) = e^{-R\cdot u}, \tag{3.23}$$

and the miner's expected value of the surplus at that exponential time horizon, given that it did not go negative until then, is given by

$$\hat{V}(u,t) = t\left(c_i - p_i\mu\sum_{j=1}^{d}\frac{a_j}{\lambda_j}\right)e^{-Ru} + u + t\left(p_i\mu\sum_{j=1}^{d}\frac{a_j}{\lambda_j} - c_i\right),\ u>0, \tag{3.24}$$

where R is the unique solution with positive real part of the equation

$$c_i R + p_i\mu\sum_{j=1}^{d}\frac{a_j\lambda_j}{R+\lambda_j} - \left(\frac{1}{t} + p_i\mu\right) = 0.$$

These metrics are of particular significance, given that the process of mining incurs considerable energy costs, and real-world miners may face the risk of financial ruin, which in turn affects their expected earnings.
Our goal in this section is to assess the sensitivity of the risk measures for which we have formulas (3.23) and (3.24) w.r.t. some of the model assumptions. First, in Section 3.5.1 we compare the formulas (using a GH fit to block rewards data from the time period February 10, 2021 to April 21, 2021) to the actual historical realization of the occurrence of blocks and the sizes of the rewards in the period from February 10, 2021 to a random time horizon in the future with a mean of 2 weeks. Note that there are still three random elements to be implemented in the historical path: the actual length of the random time horizon and the probability that a found block was found by a particular miner. Thirdly, we consider a 'Full Pay-per-Share' reward system (FPPS) for the pool, in which pool managers instantly reward miners for each share submitted, with the payout determined by the block reward and the entire estimated transaction fees associated with the block. That is, we have to simulate the arrival of shares in addition to the arrival of blocks. As the historical sample path both contains potential time-dependence of transaction fees and the actual block reward values (rather than the GH fit), but only one realization of the two latter effects, we subsequently compare in Section 3.5.2 the results of (3.23) and (3.24) with the simulated counterparts under an i.i.d. assumption, by bootstrapping block reward sizes from the empirical distribution function over the period February 10, 2021 to April 21, 2021. Subsequently, in Section 3.5.3 we provide a comparison of the formulas (3.23) and (3.24) with the simulated counterparts under the ARIMA assumption calibrated in Section 3.4. Finally, Sections 3.5.4 and 3.5.5 consider the sensitivity of the model output concerning the inclusion of transaction fees at all, as well as concerning the price of electricity.

3.5.1 Comparison with the historical path

With the collected data on the fees and prices, we can reconstitute the real surplus path of miners or pools in any specific past time period. Indeed, if we position ourselves at some starting date, we can replicate the outflows of mining costs and the inflows of block rewards for the individual miner as well as for the pool. For the following example, we select a two weeks time frame. We select the documented block

arrival in the entire system to the particular pool by simulating a Bernoulli random variable with probability equal to the pool's proportion of computational power in the global mining network. For the individual miner, we also simulate the more frequent share payouts by assuming Poisson distributed arrivals of their rewards. In Figures 3.9 and 3.10, one can see an illustration of the surplus path for the pool and the miner, respectively. Figures 3.9a and 3.9b are almost indistinguishable to the naked eye, but they are not identical. Indeed, for that short particular time horizon, the exchange rate was not very volatile. Here we choose the same parameters as in [8]: $t = 336\text{h} = 2$ weeks, $p_i = 0.001,\ q = 0.1,\ f = 0.02,\ \mu = 6p_i = 0.006,\ u = \$1M$, the cost of electricity c_i is given by $c = p_i \times e_W \times \pi_W$, where e_W is the electricity consumption of the network expressed in kWh, and $\pi_W = 0.04$ is the price of electricity per kWh. For e_W, we choose $\frac{115.541\times10^9}{365.25\times24}$.[2]

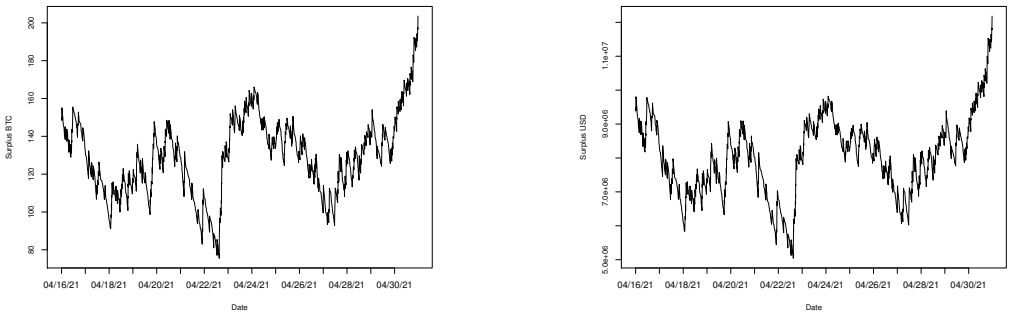

(a) BTC (b) USD

Figure 3.9: Pool surplus path.

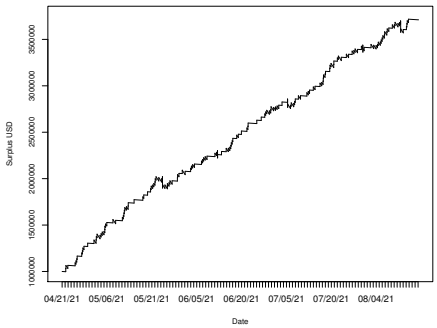

Figure 3.10: Miner surplus path USD.

Our analysis distinguishes between scenarios where the miner is engaging in solo mining or participating in a pool. More specifically, we implement the following approach:

1. We choose February 10, 2021 as the starting date.

2. We randomize the time horizon by simulating *nsim* durations following an exponential distribution with mean equal to 2 weeks.

[2] `https://cbeci.org/`

3. For each simulation run, we simulate $nsim$ share reward payment times.

4. For each block found by an individual miner, mining alone, he will receive the block reward and the transaction fees attached to this block. For a miner in a pool, the miner receives only a fraction of the block reward and of the transaction fees, assuming a FFPS pooling scheme. In practice, whereas the bitcoin block reward is known and fixed (at 6.25 BTC as of November 2023), the transaction fees included in a block can only be discovered after the block is appended to the blockchain. It means that the pool has to predict the future fee in order to pay the miners before the block appears. What often happens in practice, is that the fee is computed as the average fee over some short time frame, e.g. the last 24 hours [135]. This is the retained approach in our framework.

5. Combine all data to produce sample paths of the miner's surplus for different values of their initial capital u. Averaging over all iterations yields simulated values for $\hat{V}(u,t)$ and $\hat{\psi}(u,t)$.

On the other hand, we compute (3.23) and (3.24) under the i.i.d. assumption with a GH model for the block rewards as calibrated in Section 3.3, using the moderate truncation parameter $d = 10$ (a higher number of terms in the combinations of exponentials would significantly complicate the numerical evaluation). Figures 3.11 and 3.12 display the ruin probability and expected surplus based on the theoretical formulas and the historical path method, for both a stand-alone miner and one joining a pool. One can see that the results are rather close for the two cases, and the shape as a function of initial capital u is identified rather well. This suggests that the theoretical formulas are quite useful for practical use.

3.5.2 Sensitivity of the GH fit under the i.i.d. assumption

Under the assumption of i.i.d. block rewards, it is interesting to see how sensitive the results are w.r.t. to the fitted GH model. For that purpose, we compare the formulas (3.23) and (3.24) to a situation, where we resample block reward values from the empirical distribution function of rewards from the period February 10, 2021 to April 21, 2021. In a sense, this may also be considered a fairer comparison than the one in Section 3.5.1, since now the variability of block rewards is similar. Figures 3.13 and 3.14 depict the results, where the black lines represent the formulas using the GH fit, and the red lines are simulated values under the resampling together with a 95% confidence interval (also in all the remaining plots of this section, whenever we plot simulated values we do so together with its 95% confidence interval). Since this data sample served as the basis for the GH fit calibration, one expects a close similarity, which is indeed the case. Consequently, the GH approximation seems sufficient for purposes of drawing conclusions for our quantities of interest.

3.5.3 Sensitivity w.r.t. time dependence of rewards

We use the ARIMA model calibrated in Section 3.4 to simulate trajectories for the transaction fees, as well as another ARIMA model for the bitcoin price in USD relevant for the fixed 6.25 BTC contributions (as was suggested by Azari [19], for

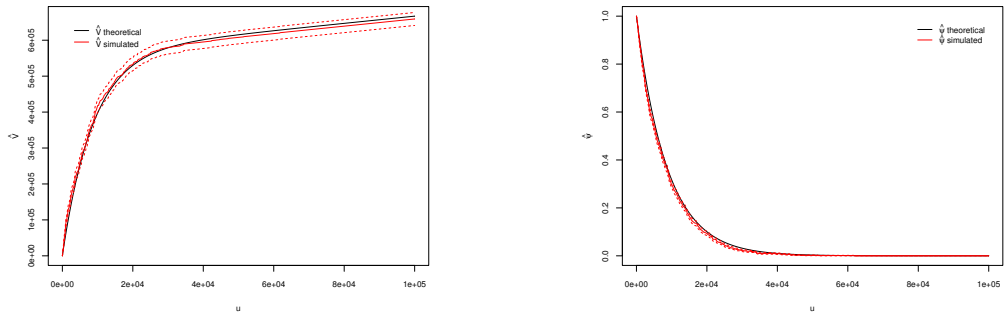

(a) Expected value of the surplus. (b) Ruin probability.

Figure 3.11: Simulation of a historical path vs i.i.d. GH fit, miner in a pool (starting date February 10, 2021).

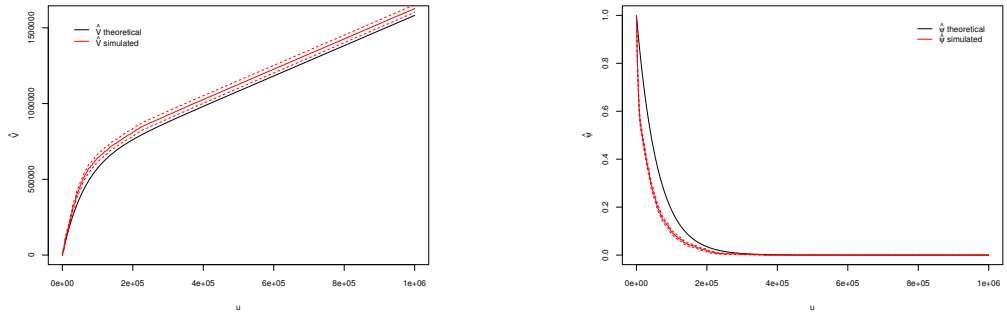

(a) Expected value of the surplus. (b) Ruin probability.

Figure 3.12: Simulation of a historical path vs i.i.d. GH fit, miner alone (starting date February 10, 2021).

more details see Appendix 3.B). For our time window, a calibration of the latter ARIMA model suggests an ARIMA(5,1,1) model. The sum of the two forms our block reward process, which feeds a Monte Carlo estimator of the ruin probability and expected surplus under that assumption. As previously, we fit the block rewards expressed in USD. For the miner, the price conversion seems interesting, as rewards enter less often, therefore the timing of the conversion is important and taken in account. From the pool perspective, it virtually operates in BTC for both inflows and outflows and therefore does not suffer conversion risk. Figure 3.15 and 3.16 compare the results (in red) with the formulas (3.23) and (3.24) (in black), which were derived under an i.i.d. assumption. One observes that for the ruin probability, the deviation from the i.i.d. assumption observed in the data is not relevant for the 2 weeks timeframe, and also for the expected value of the future surplus, the deviations are minor. The difference in the expected profit for large initial capital (where ruin becomes very unlikely) can be explained by the fact that for the calibration of the ARIMA model the time period is shorter (by 20%) compared to the one of the i.i.d. GH fit due to out-of-sample validation, cf. Section 3.4.

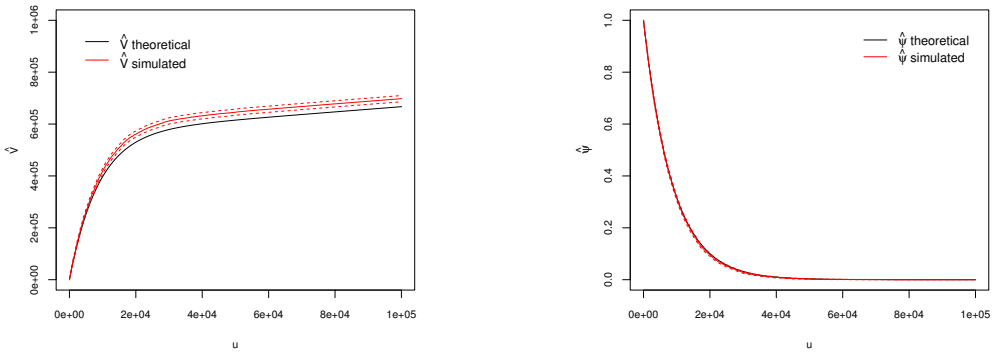

(a) Expected value of the surplus. (b) Ruin probability.

Figure 3.13: Empirical reward distribution vs i.i.d. GH fit, miner in a pool (starting date February 10, 2021).

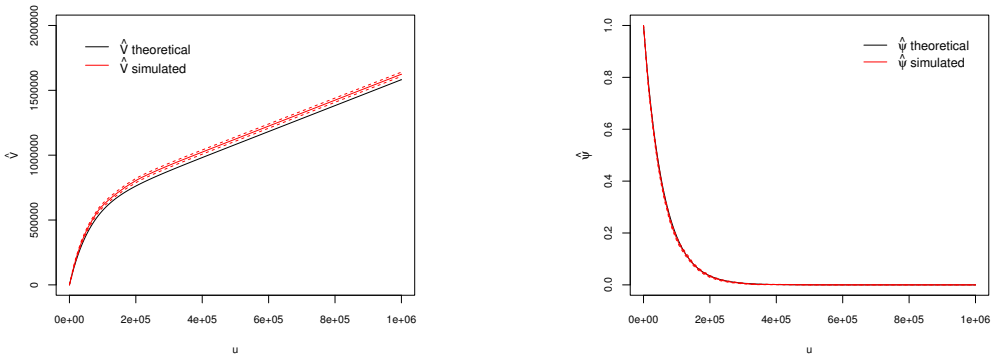

(a) Expected value of the surplus. (b) Ruin probability.

Figure 3.14: Empirical reward distribution vs i.i.d. GH fit, miner alone (starting date February 10, 2021).

3.5.4 Sensitivity w.r.t. transaction fees

Transaction fees represent an additional element of randomness in the block rewards, adding to the inherent volatility of the bitcoin (fiat) price. Our objective is to illustrate the effects of including or excluding transaction fees in the reward modelling process. Figures 3.17 and 3.18 depict the changes in the ruin probability and expected surplus for miners, both solo and within a pool, when transaction fees are inlcuded or left out (that is, only the ARIMA model for the bitcoin price is simulated in the latter case). One observes that the expected surplus converges to different limits as the initial capital u grows large. In that case, the impact of ruin on the miner's surplus diminishes, and only the positive effect of higher block rewards materializes (statistically, for the period May 12, 2020 to September 16, 2021, the transaction fees mounted to 8.6% of the fixed part of the block rewards 6.25 BTC per bounty). In terms of ruin probabilities, while the model including fees prevails over the one without fees, the difference is in fact very small. However, it is important to consider that the fixed block reward undergoes a scheduled halving approximately

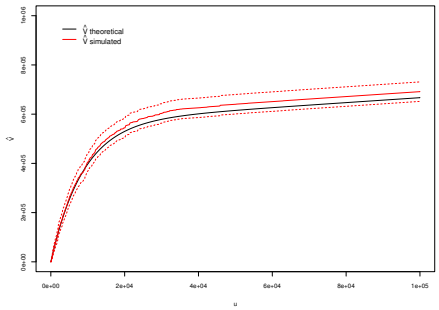

(a) Expected value of the surplus.

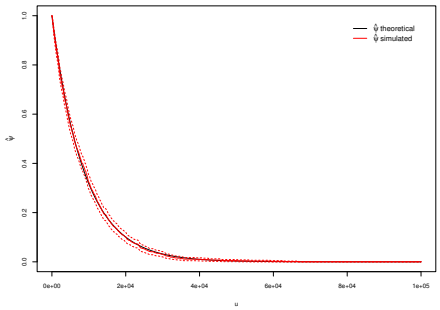

(b) Ruin probability.

Figure 3.15: ARIMA simulation vs i.i.d. GH fit, miner in a pool (starting date February 10, 2021).

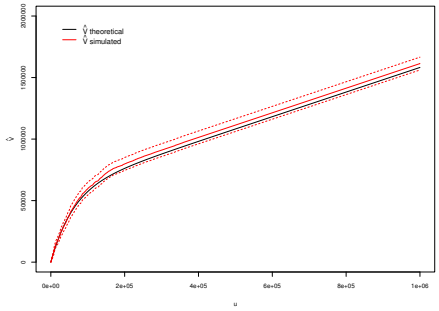

(a) Expected value of the surplus.

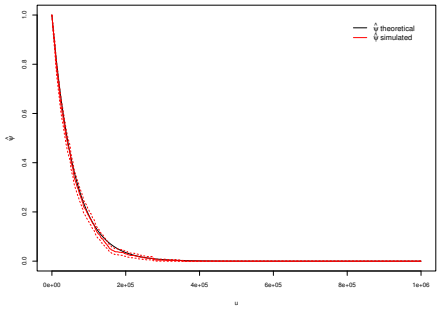

(b) Ruin probability.

Figure 3.16: ARIMA simulation vs i.i.d. GH fit, miner alone (starting date February 10, 2021).

every four years. According to the current countdown[3], the next halving is expected in April 2024, reducing the fixed reward to 3.125 BTC. In Figures 3.19 and 3.20, we therefore conduct the same analysis with 3.125 BTC for the fixed part, all other factors held constant (and assuming that the fee dynamics remain unchanged after that halving). In the latter case, the differences become much more pronounced.

3.5.5 Sensitivity w.r.t. electricity costs

Finally, let us consider the sensitivity of the results to electricity costs under the ARIMA model. Figures 3.21 and 3.22 show the ruin probability and expected surplus as a function of the price of electricity for a solo miner and a mining pool participant for a fixed initial capital of $u = 100'000$.

Such a graph can help to identify upper bounds for affordable electricity prices needed to ensure specific target levels of ruin probability or expected surplus for a given level of u.

Notably, the results reveal a substantial difference in the ruin probability between individual miners and those participating in a pool, emphasizing the risk mitigation benefits of pooling (already observed in [8]). Additionally, this quantifies the

[3]See, for example, `https://www.nicehash.com/countdown/btc-halving-2024-05-10-12-00`.

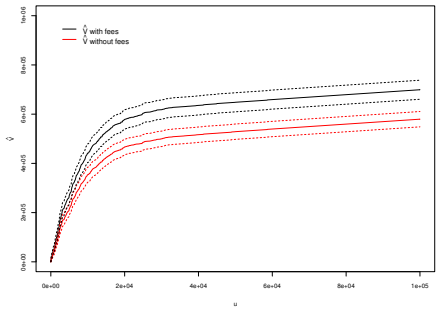

(a) Expected value of the surplus.

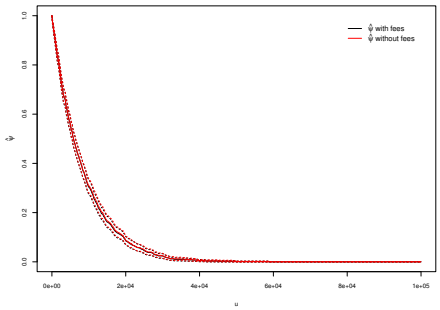

(b) Ruin probability.

Figure 3.17: Sensitivity to presence/absence of fees, miner in a pool.

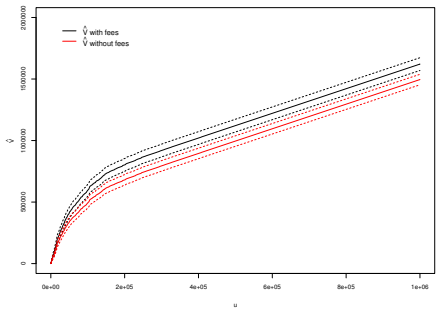

(a) Expected value of the surplus.

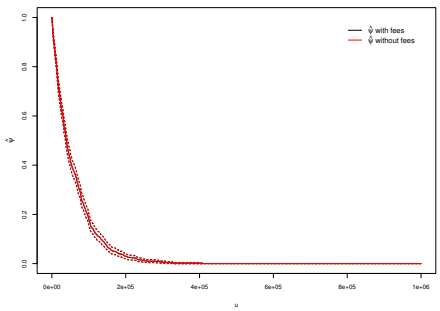

(b) Ruin probability.

Figure 3.18: Sensitivity to presence/absence of fees, miner alone.

influence of increased electricity costs, underlining the exposure and vulnerability of 'Proof-of-Work' cryptocurrencies to energy-related crises. Such considerations may (even economically) motivate the transition to alternative, less energy-intensive consensus protocols. For instance, Ethereum, the second-largest cryptocurrency, shifted from 'Proof-of-Work' to 'Proof-of-Stake' on September 15, 2022, reducing its energy consumption by a remarkable 99.95%.[4]

3.6 Conclusion

In this chapter, we have delved deeper into the framework for analyzing bitcoin mining from the perspective of risk and profitability. Previously, quantifying the choices available to miners, such as entering a mining pool and selecting the most suitable one, relied on formulas based on assumptions involving combinations of exponential distributions. In this work, we introduce a straightforward and efficient approach for fitting real-world data to these distributions, allowing us to apply theoretical results in practical mining scenarios. The remarkable flexibility of the resulting method enables us to explore shapes that are typically challenging to achieve using other approaches. The code for this work is available upon request.
Furthermore, we have explored the stochastic nature of block rewards, breaking down their variability into two key components: price volatility and the inclusion

[4]Source: `https://ethereum.org/en/upgrades/merge/` (last accessed on 13/09/2023).

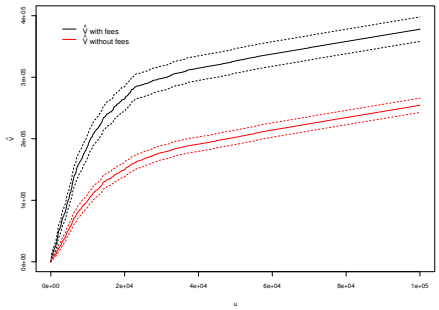

(a) Expected value of the surplus.

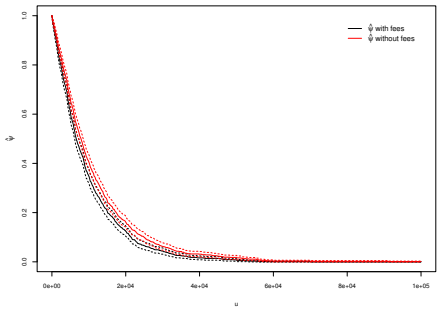

(b) Ruin probability.

Figure 3.19: Sensitivity to presence/absence of fees, miner in a pool, block reward 3.125 BTC.

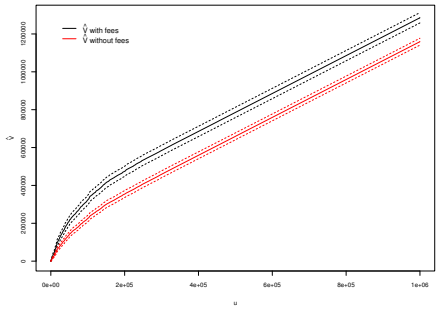

(a) Expected value of the surplus.

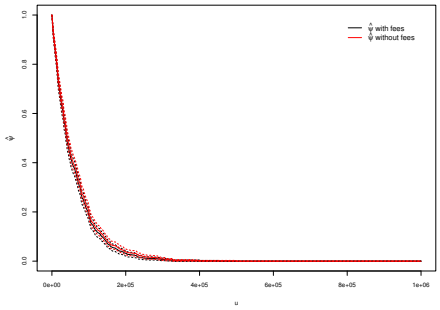

(b) Ruin probability.

Figure 3.20: Sensitivity to presence/absence of fees, miner alone, block reward 3.125 BTC.

of transaction fees. Our findings highlight the growing importance of incorporating fees in modeling, especially with the scheduled halving of the fixed block reward.
By expressing transaction fees as a time series, we can simulate our key metrics of interest—namely, the expected surplus and ruin probability—and compare them to theoretical results. This analysis confirms that the formulas derived in [8, 10], originally under an i.i.d. assumption, stand up well when compared to the outcomes derived from modeling strategies that incorporate time dependency or use empirical data more generally.
Finally, we emphasize the sensitivity of our key metrics to the inclusion of transaction fees. While the differences in magnitude are currently still small, we anticipate significant shifts in results after the next halving, making fees a more substantial factor in modeling. Our analysis also underlines the substantial impact of electricity costs on a miner's profitability and ruin probability. These insights contribute to the ongoing debate surrounding the viability of 'Proof-of-Work' consensus protocols in an increasingly energy-constrained world.

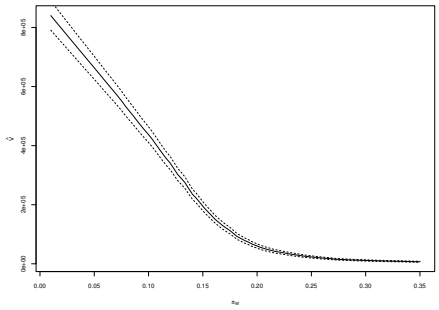

(a) Expected value of the surplus.

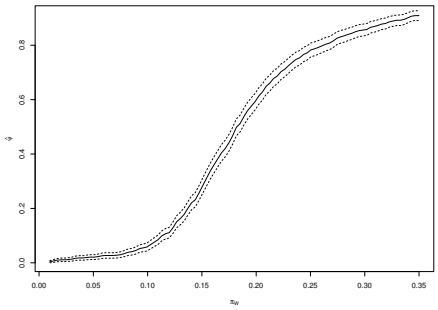

(b) Ruin probability.

Figure 3.21: Sensitivity to change of electricity price π_W, miner in a pool (starting date February 10, 2021).

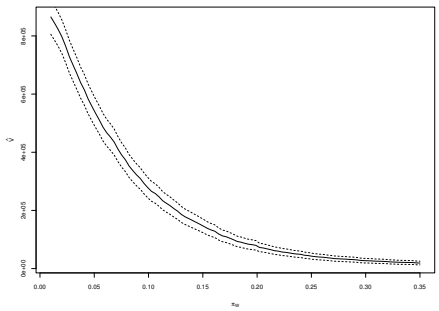

(a) Expected value of the surplus.

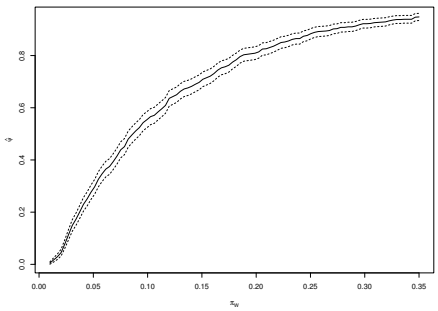

(b) Ruin probability.

Figure 3.22: Sensitivity to change of electricity price π_W, miner alone (starting date February 10, 2021).

3.A Appendix A: Variations of the root degree in the GH fitting procedure

In Section 3.3.2, one of the methods proposed to fit the empirical data by GH distributions performs a fit on the square root of the density function in order to guarantee non-negativity. One may wonder about the use of the squaring procedure. Indeed, the technique should also be applicable in a similar fashion to all n^{th} roots of a PDF, where n is even. In Figure 3.23, we illustrate the results obtained by taking $n = 2$, $n = 4$ and $n = 6$ to fit a toy sample drawn from a Gamma distribution. The main empirical observation is that the higher the root degree used, the fuzzier the approximated curves become, which is due to the increased number of terms in the GH distribution. With these considerations, we do not investigate higher order roots and proceed with the square root method in this chapter.

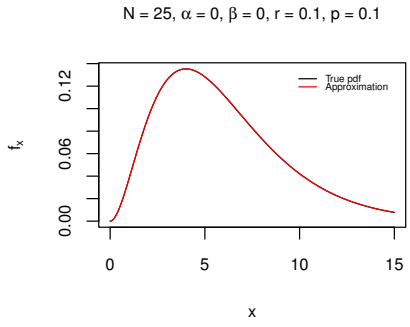

(a) $n = 2$

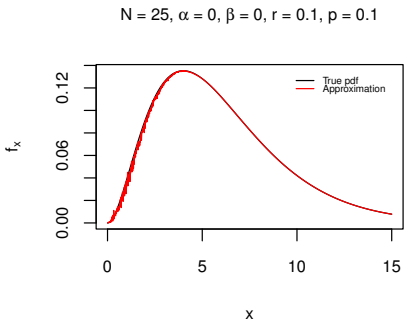

(b) $n = 4$

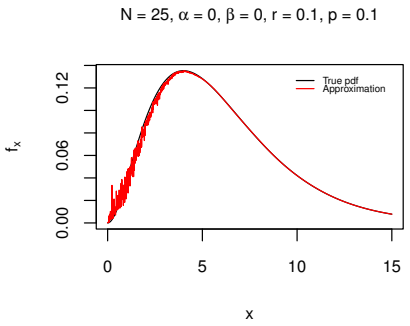

(c) $n = 6$

Figure 3.23: Approximation of the Gamma PDF by modified Method B.

3.B Appendix B: Additional time series models

3.B.1 ARIMA with covariates

In this subsection, we modify the model proposed in Section 3.4 in order to incorporate relevant covariates such as the mempoolcount in the form of a regression.

Regarding the relevance of the chosen variables, beside the linear correlation detected in Section 3.2, there is an intuitive explanation. The mempool count variable measures the number of transactions in the memory pool of the blockchain, awaiting validation by the miners. As this number grows, the congestion of the network increases. In order for the users to validate their transactions faster, there is an incentive for them to likewise increase the amount of fees attached to their transactions. In fact, most platforms incorporate calculators that help users to calculate necessary fees in order to almost surely ensure a transaction under a fixed time limit.

First of all, lets us present some additional descriptive statistics from the transaction fees dataset. In Figure 3.24, one can observe that by taking the first difference leads to stationarity. Indeed, as can be seen in Figure 3.25, the ACFs of the differenced variables show that only one lag for the fees is still significant, and none for the other covariates.

For the ARIMA with covariates regression, we follow the procedure from [79], where the times series is regressed on a vector of covariates and the residuals are forecasted with an ARIMA. In our case, the vector of covariates contains the mempool count variable, shifted in time for 100 periods, which gives a vector of 100 time-dependent variables. In other words, we define X_t^{mem}, the random variable describing the mempool count at time t. Let L be the lag operator, such as $L^k X_t^{mem} = X_{t-k}^{mem}$, for example $LX_t^{mem} = X_{t-1}^{mem}$. The covariates used in the regression would then be : $(LX_t^{mem}, L^2X_t^{mem}, \ldots, L^{99}X_t^{mem})$.

A potential problem arising from the choice of this regression period is that in practice, after 100 predicted values, one runs out of known covariates. So the predictions should be readjusted after this horizon. In practice, trend estimates can be taken from the covariates time series in order to incorporate them as regressors for the future time points in the model. Another problem again, is the absence of control on the negativity of the predicted future data points.

The summary of the result can be found in Table 3.4. In this table, one can observe the significance of each lag in the regression. The order of the model was chosen based on the AIC criterion from a search up to 20 lags. In the Figures 3.26, one can see on the left-hand side the fitted results by the model and on the right-hand side, the predicted values.

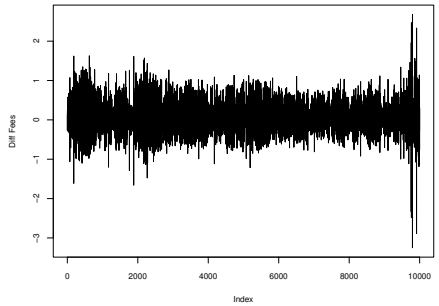

(a) Diff Fees

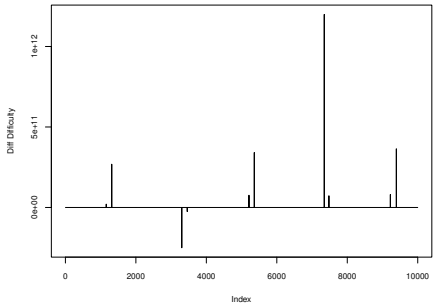

(b) Diff Difficulty

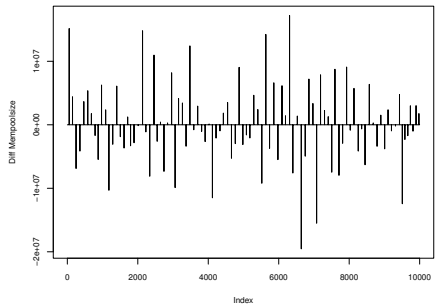

(c) Diff Mempoolsize

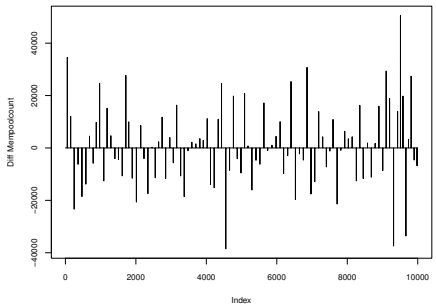

(d) Diff Mempoolcount

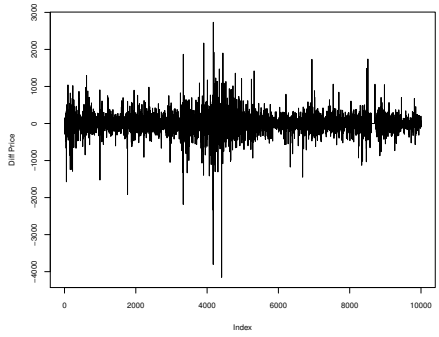

(e) Diff BTC/USD price

Figure 3.24: First difference of each covariate.

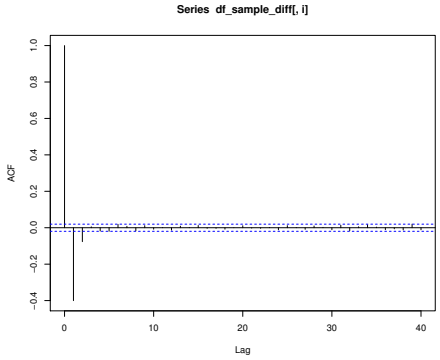

(a) ACF diff Fees

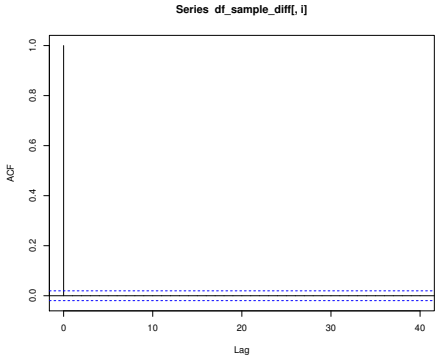

(b) ACF Log diff Difficulty

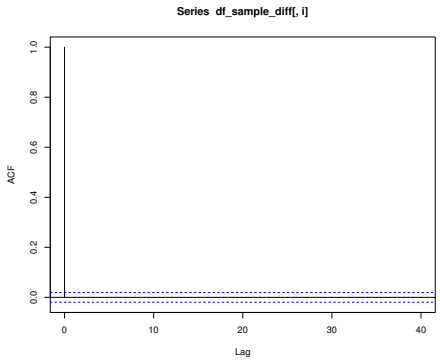

(c) ACF diff Mempoolsize

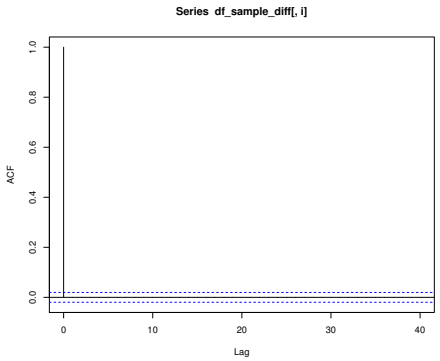

(d) ACF diff Mempoolcount

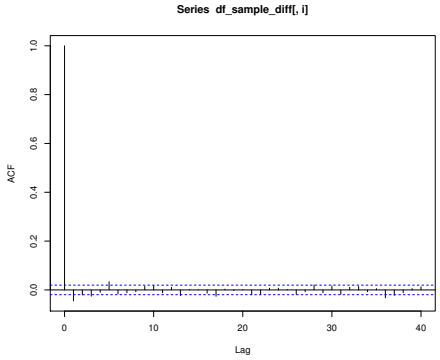

(e) ACF diff BTC/USD price

Figure 3.25: ACF difference of each covariate.

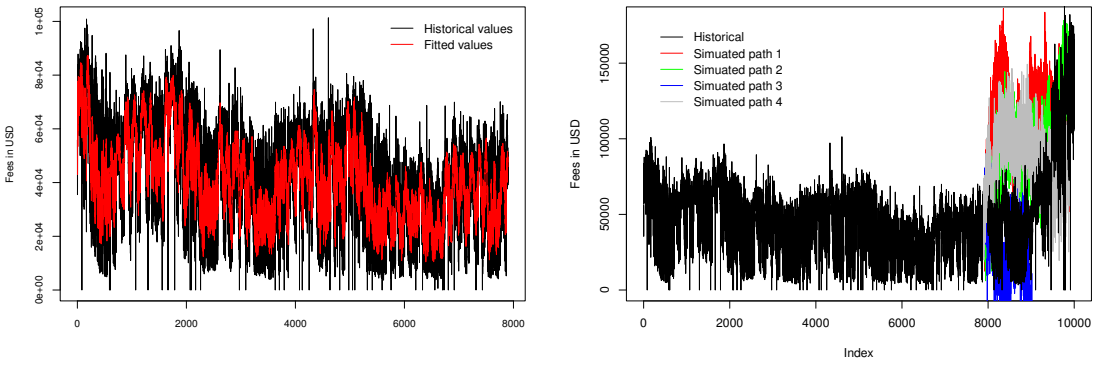

(a) Fitted values.

(b) Simulated fees per block in USD.

Figure 3.26: ARIMA(4,1,2) with linear regression on 100 lagged memory pool count covariates.

Table 3.4: Summary of ARIMA model with lagged regressors of memory pool count.

	ar1	ar2	ar3	ar4	ma1	ma2	xreg1	xreg2	xreg3
	-0.6308	0.2366	0.0624	0.0434	-0.0513	-0.7898	0.054	0.1093	-0.2924
s.e.	0.0679	0.0228	0.0167	0.0126	0.067	0.0603	0.1216	0.1472	0.1479
	xreg4	xreg5	xreg6	xreg7	xreg8	xreg9	xreg10	xreg11	xreg12
	0.1533	0.0656	-0.0312	0.2065	-0.1537	0.0489	0.0079	0.0657	0.0506
s.e.	0.1479	0.1479	0.1479	0.1479	0.1479	0.1479	0.1479	0.1479	0.1479
	xreg13	xreg14	xreg15	xreg16	xreg17	xreg18	xreg19	xreg20	xreg21
	-0.2628	0.0852	-0.123	0.0803	0.0948	0.089	-0.088	-0.1082	0.0212
s.e.	0.1479	0.1479	0.1479	0.1479	0.1479	0.1479	0.1479	0.1479	0.1479
	xreg22	xreg23	xreg24	xreg25	xreg26	xreg27	xreg28	xreg29	xreg30
	0.1854	0.0017	-0.1564	-0.0236	0.0297	0.0248	0.1464	-0.166	0.1341
s.e.	0.1479	0.1479	0.1479	0.1479	0.1479	0.1479	0.1479	0.1478	0.1477
	xreg31	xreg32	xreg33	xreg34	xreg35	xreg36	xreg37	xreg38	xreg39
	-0.1022	-0.0552	-0.1845	0.1116	-0.0922	0.2335	-0.1421	0.4361	-0.278
s.e.	0.1477	0.1477	0.1477	0.1477	0.1477	0.1477	0.1477	0.1477	0.1477
	xreg40	xreg41	xreg42	xreg43	xreg44	xreg45	xreg46	xreg47	xreg48
	-0.1486	0.3066	-0.0679	-0.2838	0.2205	-0.1769	0.1635	0.1928	-0.1592
s.e.	0.1477	0.1477	0.1477	0.1477	0.1477	0.1477	0.1477	0.1477	0.1477
	xreg49	xreg50	xreg51	xreg52	xreg53	xreg54	xreg55	xreg56	xreg57
	-0.0567	-0.1243	0.1907	0.0804	0.0154	-0.0461	-0.162	0.0289	0.1161
s.e.	0.1477	0.1477	0.1477	0.1477	0.1477	0.1477	0.152	0.1539	0.154
	xreg58	xreg59	xreg60	xreg61	xreg62	xreg63	xreg64	xreg65	xreg66
	-0.1808	0.0621	0.1966	-0.1887	0.0183	-0.1012	0.2699	-0.4113	0.1366
s.e.	0.154	0.154	0.154	0.154	0.154	0.154	0.154	0.154	0.154
	xreg67	xreg68	xreg69	xreg70	xreg71	xreg72	xreg73	xreg74	xreg75
	0.1899	0.077	-0.3807	0.1774	0.0485	-0.1507	0.1555	-0.0685	-0.162
s.e.	0.154	0.154	0.154	0.154	0.154	0.154	0.154	0.154	0.154
	xreg76	xreg77	xreg78	xreg79	xreg80	xreg81	xreg82	xreg83	xreg84
	-0.0186	-0.0585	0.312	-0.1822	0.0604	-0.144	0.0168	-0.0953	0.1755
s.e.	0.154	0.154	0.154	0.154	0.154	0.154	0.154	0.154	0.154
	xreg85	xreg86	xreg87	xreg88	xreg89	xreg90	xreg91	xreg92	xreg93
	-0.1648	-0.1744	0.1769	-0.1296	0.072	0.0338	0.1139	-0.0778	-0.0734
s.e.	0.154	0.154	0.154	0.154	0.154	0.154	0.154	0.154	0.154
	xreg94	xreg95	xreg96	xreg97	xreg98	xreg99	xreg100		
	0.0997	-0.2199	0.1156	-0.0932	-0.0331	0.054	0.2058		
s.e.	0.154	0.154	0.154	0.1541	0.1541	0.1533	0.1268		
σ^2	220137825								
Log.Lik.	-87024.43								
AIC	174262.9								
BIC	175009.1								

3.B.2 Vector Auto Regression

As with an ARIMA using a linear matrix of regressors we do not take in account the evolution of these regressors in time, a fit of a vector autoregressive model (VAR) would capture this. Again, we consider the memory pool count as the only regressor as including others did not significantly improve the fits based on previous observations. A multivariate time series $\mathbf{X_t}$ subject to a VAR(p) process can be expressed in the following form:

$$\mathbf{X_t} = \boldsymbol{\beta_1}\mathbf{X_{t-1}} + \cdots + \boldsymbol{\beta_p}\mathbf{X_{t-p}} + \mathbf{z_t}, \tag{3.25}$$

where $\boldsymbol{\beta_1}, \ldots, \boldsymbol{\beta_p}$ are non-zero vectors of coefficients and $\mathbf{z_t}$ is a multivariate white noise with mean 0 and covariance matrix Σ_z. For an overview of this model, see e.g. [129].

We choose the number of lags based on an AIC selection of models upto 200 lags. As it can be seen in Figure 3.27, the small number of lags let the forecasted values be strongly influenced by the trend of the last values and thus the predicted path may follow too closely the average of the last time points.

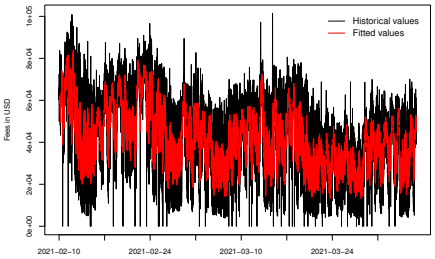

(a) Fitted values.

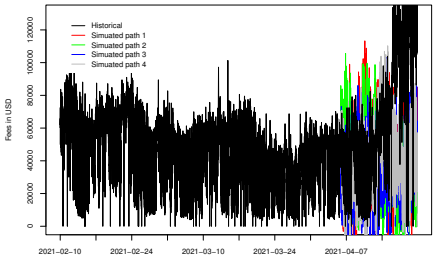

(b) Simulated fees per block in USD.

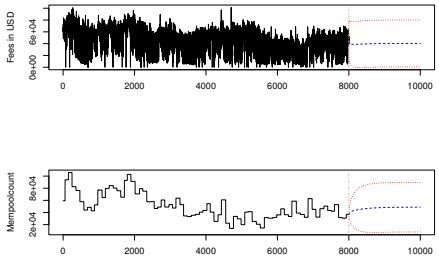

(c) Forecast confidence interval bounds.

Figure 3.27: VAR(16) forecast.

To compare, we also apply the model with 144 lags, which corresponds approximately to 1 day of data (see Figure 3.28) and 2016 lags (2 weeks) in Figure 3.29.

Figure 3.29 shows that with the increased number of lags, the confidence intervals for the prediction are getting larger, thus encompassing the true values in the 95% bounds.

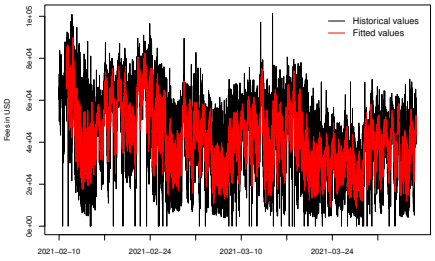

(a) Fitted values.

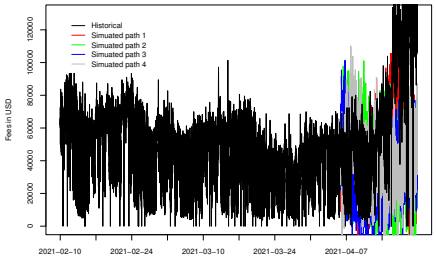

(b) Simulated fees per block in USD.

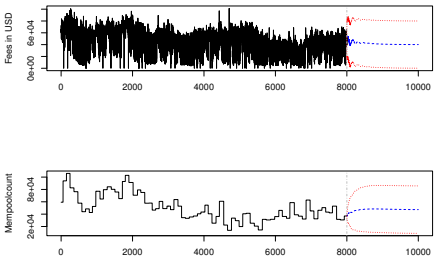

(c) Forecast confidence interval bounds.

Figure 3.28: VAR(144) forecast.

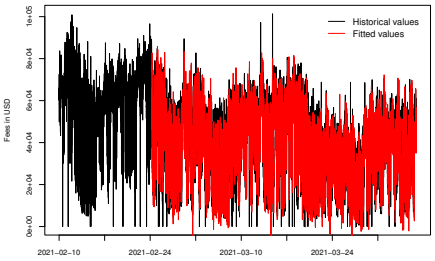

(a) Fitted values.

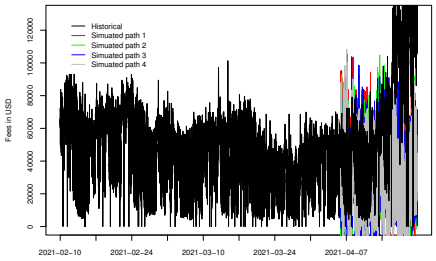

(b) Simulated fees per block in USD.

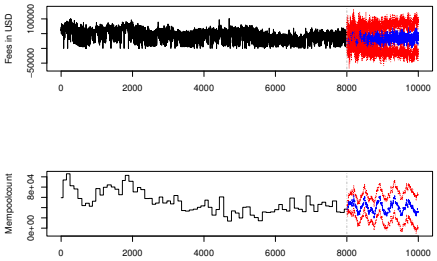

(c) Forecast confidence interval bounds.

Figure 3.29: VAR(2016) forecast.

3.B.3 ARIMA-GARCH

Recent literature already suggested modelling bitcoin prices through GARCH models (see e.g. [86]). GARCH type modelling is better suited to model the volatility observed in series. GARCH models were first introduced by Bollerslev in 1986 [24] and have the following representation: Let ϵ_t be a real-valued stochastic process following a GARCH(p,q) process. Then:

$$\epsilon_t \mid \epsilon_s \sim \mathbb{N}(0, h_t), \ \forall s < t,$$
$$h_t = \alpha_0 + \sum_{i=1}^{q} \alpha_i \epsilon_{t-i}^2 + \sum_{j=1}^{p} \beta_j h_{t-j},$$

where $p \geq 0,\ q \geq 0,\ \alpha_0 > 0,\ \alpha_i \geq 0,\ i = 1, \ldots, q,\ \beta_j \geq 0,\ i = 1, \ldots, p$, see also [65].

In order to fit our data, we will apply first the ARIMA model to our data (as above) and further introduce heteroskedasticity in the residuals via the GARCH model, as suggested in [155, 165]. In Figure 3.8, we showed the normal Q-Q plot of the residuals from the ARIMA model. As the reader can observe, there is room for improvement and standard statistical tests reject the Gaussian hypothesis. So in order to fill this gap, we fit a GARCH model to the residuals. For this, we base our choice of parameters on the AIC criterion. The best model out of the tested set is GARCH(17,2). In Figure 3.30a, one can see the improvement of the Q-Q plot of residuals after accounting for the variance. Finally in Figure 3.30b, one can observe the prediction with confidence interval bounds for the overall model. Up to 10 days, the realized path falls within the confidence interval bound, but this example is the worst case scenario since the training period strongly deviates from the test data. As in all other models, the main drawback is that the predicted results can be negative. This model is not retained in comparison to the ordinary ARIMA model as few improvements of the Q-Q plot are observable for many additional parameters.

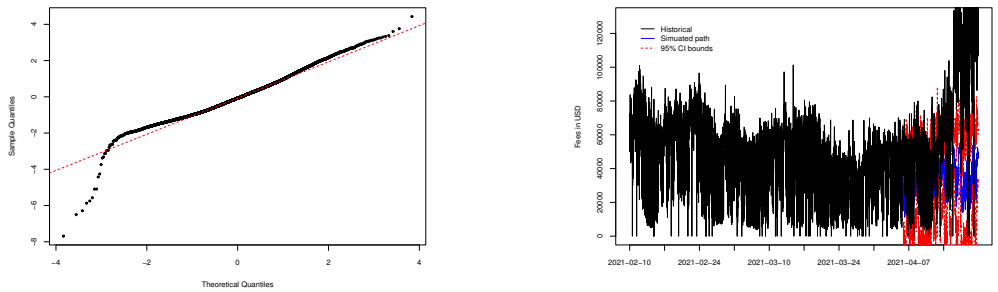

(a) Q-Q plot of the residuals. (b) Predicted path and confidence interval.

Figure 3.30: Results of the ARIMA(1,1,7)-GARCH(17,2) model.

3.B.4 Bitcoin price

In Section 3.5 we mentioned that the Bitcoin price was fitted using an ARIMA model inspired by Azari [19]. Let us present this model here. We fit the data on the

same period as the transaction fees, from February 10^{th}, 2021 to April 21^{th}, 2021. The corresponding price evolution is displayed in Figure 3.31a. On the right-hand side of Figure 3.31b, we can observe the ACF of the Bitcoin price, which exhibits autocorrelation.

(a) Evolution of the Bitcoin price.

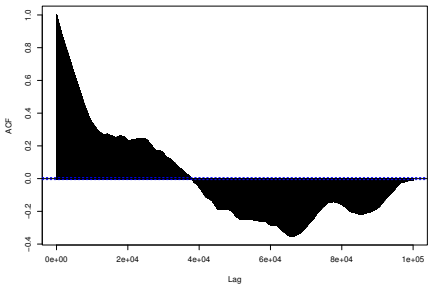

(b) ACF of the Bitcoin price.

Figure 3.31: BTC/USD price, 10/02/2021 to 21/04/2021.

A test for stationarity with an augmented Dickley-Fuller test and a Ljung-Box test is performed and rejects stationarity. In order to counter this problem, we transform the series by observing the log-differentiated series in Figure 3.32a. As can be seen in Figure 3.32b, the autocorrelation disappears. The Dickley-Fuller test startistic has a p-value of 0.01, which do not reject stationarity.

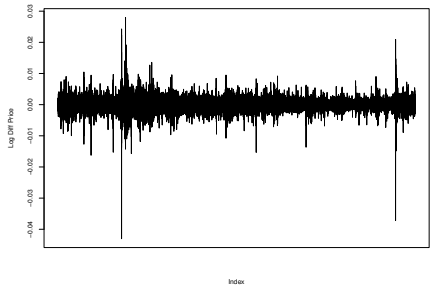

(a) Log difference of BTC/USD.

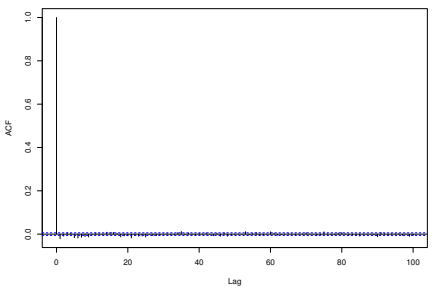

(b) ACF of the log difference of BTC/USD.

Figure 3.32: Log differentiated BTC/USD price, 10/02/2021 to 21/04/2021.

We can now confidently fit an ARIMA by minimizing AIC using *forecast* package from Hyndman and Khandakar [79]. The plot of observed vs fitted values almost overlap in Figure 3.33a. This fit is given by a ARIMA(5,1,1) model with parameters given in Table 3.5. Next, following [157, chap. 7], we can simulate paths of Bitcoin prices from the fitted ARIMA model. We can see simulated predicted paths in color in Figure 3.33b vs the true historical development in black.

3.C Appendix C: Other cryptocurrencies

In this Section, we aim to compare the dynamics of the transaction fees for other cryptocurrencies and see whether they can also be captured by simple models.

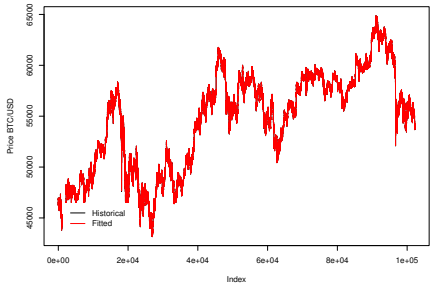

(a) Observed vs Fitted price.

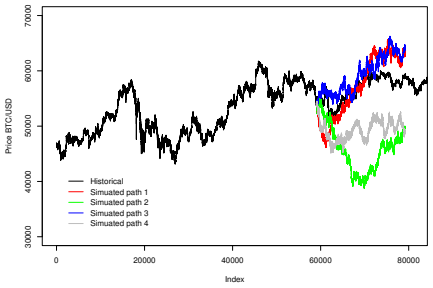

(b) Predicted paths in color vs historical path.

Figure 3.33: ARIMA(5,1,1) fit of the Bitoin price, 10/02/2021 to 21/04/2021.

Table 3.5: Coefficients of the ARIMA fit for the price model.

	ar1	ar2	ar3	ar4	ar5	ma1
	0.6991	0.0065	0.0126	-0.0015	-0.0251	-1.7217
s.e.	0.0370	0.0039	0.0038	0.0038	0.0033	003701
σ^2	4653					
Log.Lik.	-576682.5					
AIC	1153379					
BIC	1153446					

3.C.1 Cardano

Launched in 2012, Cardano[5] blockchain records transactions performed in Ada cryptocurrency. In contrast to the Bitcoin, Cardano adopted a *Proof-Of-Stake* consensus protocol. The cryptostaking does not require huge computational power to solve cryptoproblems. Instead, anyone can become a validator by staking their coins. Ada holder delegate their stake to stake pools, reliable server nodes run by a stake pool operator. With a higher amount of staked coins, the node gets a higher probability of becoming a validator, meaning he will construct the next block. Akin to the Bitcoin consensus protocol, the winner of this round will collect the reward. Since its conception, Cardano evolved through different *eras* of development. Currently, the block formation is structured as follows. Time is divided in *epochs*, which in their turn are subdivided in *slots*. Each slot last for one second and one epoch accommodates 432'000 slots. In every slot, there is a probability of electing one or more slot leaders. If one of these leaders produces a block, it is appended to the blockchain. See Zhang and Lee [159] for an overview of consensus protocols.

The rewards in the Cardano system are constituted by the sum of the transaction fees attached to the newly generated block as well as some newly issued Ada coins. The reward is distributed to the slot leader according to a predefined formula (see `https://docs.cardano.org/learn/pledging-rewards`). One of the

[5] `https://docs.cardano.org/`, `https://cardano.org/`

parameters in the formula is the saturation parameter, preventing the pools of staking to much and thus decentralizing the system. The transaction fees themselves obey a certain structure. The minimal fee is capped from below by a formula liner in the transaction size (see `https://docs.cardano.org/explore-cardano/fee-structure`, [29]). On the other hand, there is no upper bound on the fees.

For the purpose of our analysis, we collect data on the sum of the transaction fees per block during 2 periods of time. These periods match the beginning and the end of the extracted Bitcoin data series. Due to the high frequency of Cardano blocks, obtaining the full 16 months period data was not possible.

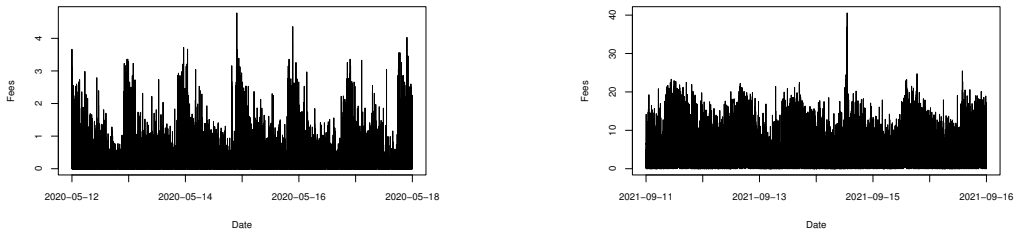

(a) From 12/05/2020 to 18/05/2020. (b) From 11/09/2021 to 16/09/2021.

Figure 3.34: Cardano fees.

The first observation that can be made in Figure 3.34 is a seemingly cyclical behaviour of the fees. This cyclicity may reflect the intensity of the transactions during different daytimes. For the most recent period, the peaks on the graphs consistently fall around 2-3 p.m. UTC. As for the lowest parts, there are on the opposite at 3 a.m.

The ACF in Figure 3.35 confirms the cyclical behaviour of the time series. The idea is then to fit a model accounting for seasonality (daily in our case). For this purpose, we apply the SARIMA model. Recall that a time series X_t is SARIMA$(p,d,q)(P,D,Q)_s$ defined as

$$\Phi_P(B^s)\phi(B)\nabla_s^D\nabla^d X_t = \Theta_Q(B^s)\theta(B)z_t, \tag{3.26}$$

where z_t is a white noise with mean 0 and variance σ_z^2, and ∇ denote backshift operators. For an overview of this model, see e.g. [129] and [78] for forecasting methods. The results are illustrated in Figure 3.37.

3.C.2 Ethereum

Ethereum, established as the second most prominent cryptocurrency following Bitcoin in terms of market capitalization, presents a interesting case in the world of digital currencies. This short subsection investigates its uniqueness and recent developments, and the implications of these changes for users and miners.

In the Ethereum blockchain, Ommer (previously known as Uncle) blocks represent an interesting phenomenon. These blocks are generated concurrently with the main blocks but do not form part of the main chain. Despite this, Ommer blocks are

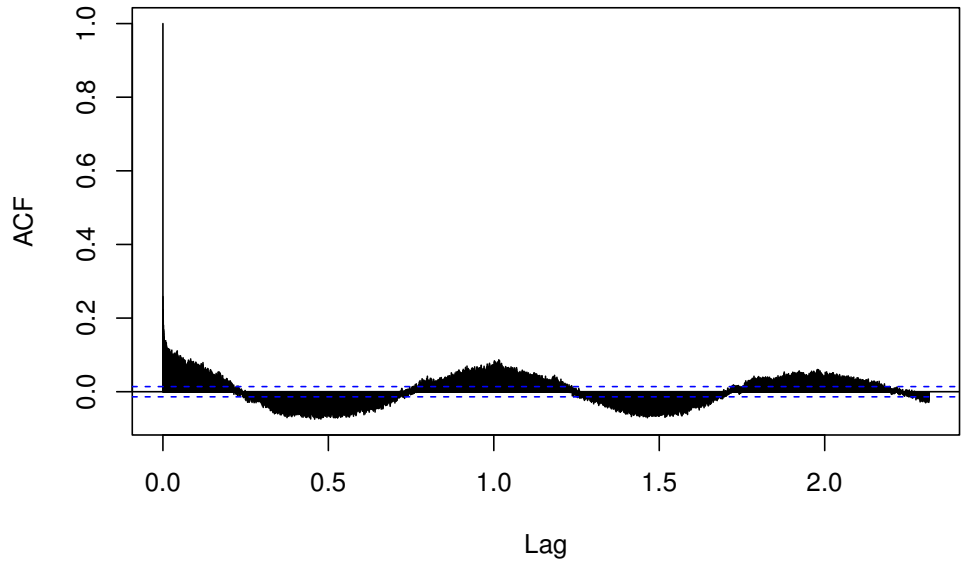

Figure 3.35: ACF of Ada transaction fees.

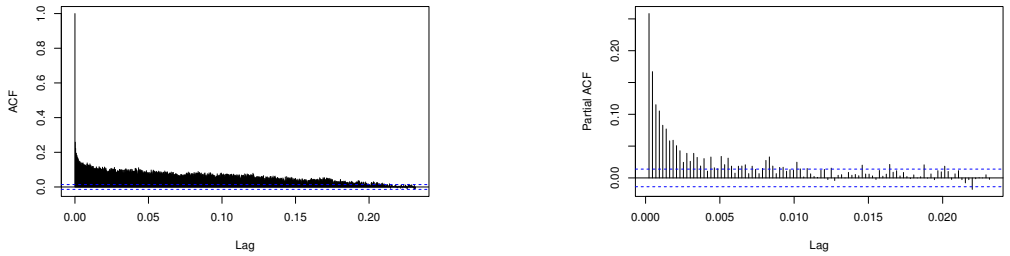

(a) ACF. (b) PACF.

Figure 3.36: ACF and PACF of Ada transaction fees.

integral to Ethereum's architecture, contributing to network security and rewarding miners for valid block production. However, with the transition of Ethereum from a Proof-of-Work (PoW) to a Proof-of-Stake (PoS) protocol, the role and dynamics of Ommer blocks have undergone significant changes, and are not part of the PoS protocol anymore[6], although the transaction fees mechanics remained unchanged for Ommer blocks after transition[7].

Ethereum transactions are powered by "gas," a unit that quantifies the computational effort required to execute operations. Each transaction necessitates a certain amount of gas, acting as a fee paid to miners for processing and validating transactions. Notably, if the gas limit for a transaction is reached without com-

[6] `https://ethereum.org/en/glossary/#ommer`, last accessed on 23/11/2023.

[7] `https://eips.ethereum.org/EIPS/eip-3675#block-and-ommer-rewards`, last accessed on 23/11/2023.

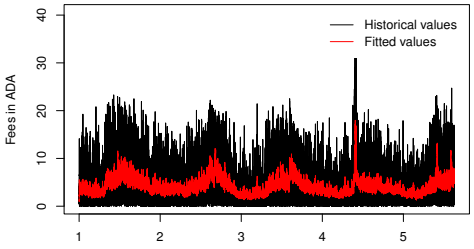

(a) Fitted values.

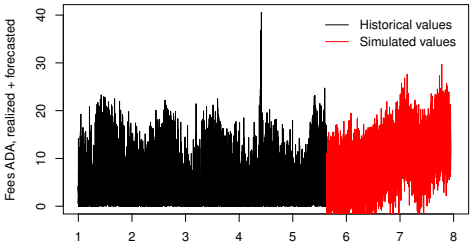

(b) Predicted values.

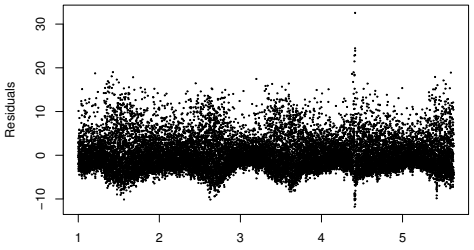

(c) Residuals.

Figure 3.37: SARIMA $(40, 1, 10)_{4320}$ fit of Ada transaction fees.

pletion, the transaction fails, emphasizing the importance of accurately estimating gas requirements [74]. The London Hard Fork[8], implemented on August 5, 2021, at block 12,965,000, introduced a pivotal change in the Ethereum fee structure. A portion of the transaction fees, previously awarded to miners, is now "burned" or permanently removed from circulation. This mechanism aims to reduce the overall supply of Ether, potentially influencing its value. Additionally, the update brought a new fee calculation model, incorporating a base fee per block (subject to fluctuation) and a priority fee, determined by users to expedite transaction processing[9].

A landmark event in Ethereum's history was "The Merge", completed on September 15, 2022, at block 15,537,393. This transition marked Ethereum's shift from the energy-intensive PoW consensus mechanism to the more environmentally friendly PoS protocol. This fundamental change not only alters the method by which new blocks are created and validated but also has significant implications for the overall security, scalability, and energy efficiency of the Ethereum network.

Figures 3.38a and 3.38b illustrate the total fees associated with Ethereum transactions over the period from 21/06/2022 to 07/10/2022. These charts provide valuable insights into the dynamics of Ethereum's network before and after the Merge, highlighting the fluctuating nature of transaction fees. Similarly, Figures 3.39a and 3.39b focus on the fees received by miners, accounting for the burned fees.

[8] `https://cointelegraph.com/news/ethereum-london-hard-fork-goes-live`, last accessed on 23/11/2023.

[9] `https://ethereum.org/en/developers/docs/gas/`, last accessed on 23/11/2023.

Ethereum's evolution, particularly through significant updates like the London Hard Fork and The Merge, shows its adaptive and innovative nature in the cryptocurrency domain. These developments not only demonstrate Ethereum's commitment to addressing scalability and environmental concerns but also show its potential to shape the future landscape of blockchain technology and digital currencies.

(a) From 21/06/2022 to 07/10/2022. (b) Zoom.

Figure 3.38: ETH total fees.

(a) From 21/06/2022 to 07/10/2022. (b) Zoom.

Figure 3.39: ETH fees received by miner (Fees - Burned Fees).

Chapter 4

On the cost of risk misspecification in insurance pricing

This chapter is based on the following article:

D. Finger, H. Albrecher, and L. Wilhelmy. On the cost of risk misspecification in insurance pricing: a mean-variance approach. Submitted.

Abstract. In the non-life insurance industry, pricing is often done relative to individual criteria of policyholders. Various classification algorithms are in use to categorize policyholders into risk classes defined by the insurer, but classification errors may result from this process. In the light of recent automatic classification practices, it becomes important to assess the risks caused by such errors. In this chapter we examine the impact of risk class misspecifications for a simple situation with two risk types. We provide a mean-variance framework for quantitatively studying the insurer's optimization problem of specifying premiums and we analyze the tradeoff of costs and benefits when classification error probabilities are known.

4.1 Introduction

Risk classification is a classical tool in actuarial practice. Indeed, the distribution of individual risks will often differ substantially, depending on personal characteristics, different exposure, environmental conditions etc. If an insurer applies the same premium across all such categories, this may lead to adverse selection, in the sense that individuals who face a lower premium than appropriate for their true risk will massively enter the contract, raising the price and squeezing out rational individuals with lower risk (see e.g. [14] for a general discussion). Classification is commonly used in insurance during the underwriting process and the tarification, cf. [113, 145] for health and life insurance and [149] for property lines. Different types of variables, such as quantitative or categorical can be used in various algorithms for actuarial classification purposes [128]. For an overview of classification methods we refer to [75]. In the literature, numerous authors discuss risk classification in view of adverse selection and efficiency in the Pareto sense, see [53] for a survey on this topic. In [42, 43, 77] authors study the efficiency of imperfect categorization in a utility setup. In [54, 116], authors consider the effect of bans on classification on the market efficiency. [138] deals with the customers' perspective on adverse selection and efficiency of risk classification. The contributions [42, 116] consider also costs of categorization. Indeed, risk classification may be costly both computationally (in terms of resource-consuming algorithms in the presence of huge datasets) and in monetary terms (e.g. involving data acquisition from official statistics or from competitors).

The insurer hence needs to select appropriate criteria for rating variables, which may take different forms and are not always motivated by actuarial drivers. For instance, [63] distinguishes actuarial, operational, social and legal criteria. In some cases, classification may be inaccurate when the main criterion is unobservable or there are legal constraints to use it, e.g. for social or political reasons. For instance, gender-based discrimination is nowadays forbidden in the European Union (see e.g. [124]), even if it is widely considered a relevant characteristic from the statistical point of view. In [125], authors discuss the perceptions of gender as pricing criteria by the customer and explain the complex interplay between anti-discrimination laws and actuarial principles within the insurance industry, see [98] for suggestions to deal with this issue. We refer to the recent book [32] for a rich source of information and ideas on this topic.

In recent years, advancements in genetic testing and personalized medicine have brought forth unique challenges in data protection. When individuals undergo genetic tests and share this information with their insurers, it gives the insurer insight into the individual's risk profile. This knowledge can affect how individuals are classified in insurance schemes, raising privacy concerns about the sharing of sensitive genetic data. Moreover, integrating personalized medicine into healthcare systems presents potential inequalities. Some patients might require expensive treatments, raising questions about coverage by insurance providers. Detailed discussions of these challenges are explored in various studies, including [50] and [84].

Although criteria to distinguish risk profiles will exist, they may often be un-

known to the insurer due to the asymmetry of information, see e.g. [3, 6]. At the same time, recent years have seen a dramatic increase in both the amount of risk information available and the ability to analyse it statistically [133]. In many situations, risks can be analysed on an almost individual basis, but in any case as elements of much smaller rating pools, i.e. pools of risks that share characteristics such that they can be assumed to have the same loss distribution. Offered insurance cover may also take into account the current risk situation, i.e. environmental variables (e.g. time, place) or even behavioural variables. The small size of rating pools leads to larger estimation errors in estimating the expected cost of insurance cover. In the presence of competing insurance providers and given the transparency of the prices of their insurance offers, customers may tend to choose the cheapest offer, i.e. an offer that is too low compared to the true but unknown production costs. From the perspective of the insurance company, this phenomenon is known as the "winner's curse". In order to avoid the negative economic consequences of the winner's curse, insurance companies like to apply tailor-made surcharges to the offered premium. It is likely that these surcharges, in addition to the higher cost of more granular risk assessment, will lead to higher overall costs for the entire insured portfolio of risks. Hence, the total welfare of the community is reduced. In addition, the overall coverage ratio may be reduced as some risk owners may find the increased cost of insurance too high. An overly granular approach may therefore be counterproductive and there may be an optimal level of rating granularity, together with legal boundaries, see for instance [63], [5] and [143].

Even when the resulting loss distributions are assumed to be known, each classification method will inherently contain classification errors, which may lead to unprofitable decision making. Note that error rates tend to be higher when one class is less represented in the population, cf. [87, 150]. For a classical reference to empirical evidence on concrete values for error probabilities arising from frequently used statistical methods, see e.g. [21].

In this chapter, we would like to quantitatively study the effects of such classification errors in the context of a simple model with only two possible risk classes. In that case one faces two types of errors: assigning a risk of the first class as one of the second and vice versa (a *false positive* and *false negative* in statistical terms, or also *sensitivity* vs. *specificity*, see e.g. [87, 145]). We compare three scenarios, one where the insurer does not classify the heterogeneous risks and two others where the risks are differentiated, but once with perfect knowledge and once with some probability of misspecification of policyholder types. While the introduction of the classification mechanism allows to price the insurance risks more efficiently, it entails certain fixed costs, and the potential misspecification can impact the insurer's profit further. In a similar spirit to [66], we develop a simple framework to assess the respective trade-off between costs of classification and profits. We consider linear and sigmoid-type demand functions of the premium for the probability that a customer accepts an offered policy (rather than deterministic demand functions as in [66]). Our main optimization target is the expected profit for the insurer (see e.g. [89]), and we consider several risk measures to assess the risk part in the analysis.

The remainder of the chapter is structured as follows: In Section 4.2, we start

with an insurer's expected profit approach and a piece-wise linear demand setting. We then consider different scenarios for this setup in Section 4.3. In Section 4.4, we then include the variance in our considerations, and establish mean-variance frontiers for the profit. We illustrate the results and its main drivers in Section 4.5. In order to assess the sensitivity of the results, Section 4.6 then develops a number of extensions, namely a sigmoid-type rather than piece-wise linear demand function, a lower semi-variance and a value-at-risk concept for replacing the variance in the risk assessment of the strategies, as well as a utility function approach to unite the consideration of profitability and variability in one function, together with numerical illustrations for each of these cases. Finally, Section 4.7 concludes.

4.2 Model Setting

Assume that there is only one insurer present in the market, so there is no competition. Let us further assume a population of n individuals, who are all willing to contract insurance, and independently from each other choose whether they enter the insurance contract for a given premium or not. The individuals fall into two types: low risk type with loss random variable L and high risk type with loss random variable H, with underlying cumulative loss distribution functions $F_L(x)$ and $F_H(x)$, $x \geq 0$ respectively (both L and H will typically have atoms at 0, signifying the case of no claim in the considered time period). Define the respective means and variances by

$$\mathbb{E}(L) = \mu_L, \ \mathbb{V}\mathrm{ar}(L) = \sigma_L^2, \tag{4.1}$$
$$\mathbb{E}(H) = \mu_H, \ \mathbb{V}\mathrm{ar}(H) = \sigma_H^2, \ \mu_H > \mu_L, \tag{4.2}$$

which are all assumed to be finite. All risks are assumed to be independent and identically distributed within each type. Let $\{p_L, p_H \geq 0\}$ be the actual proportion of the low- and high-risk type among the policyholders ($p_L + p_H = 1$).

Define an acceptance function f_i which for any proposed premium P gives the probability for an individual to enter the contract; the form of this function differs for each risk type i. Each individual is assumed to take the decision about entering independent of all the others. We distinguish the acceptance behaviour of each individual within its risk type, but this can be extended to distinguish between each individual at a personal level. If m individuals are offered a premium P and all use the same acceptance function f_i, we then expect $m \cdot f_i(P)$ individuals to enter the contract. Let us first assume that f_i is piece-wise linear

$$f_i(P) = \begin{cases} 1 - \dfrac{P}{P_i^{max}}, & 0 \leq P \leq P_i^{max}, \ i \in \{L, H\} \\ 0, & P_i^{max} < P, \end{cases} \tag{4.3}$$

where P_i^{max} is the so-called *reservation price*, cf. Figure 4.1. Clearly, f_i is non-increasing in P. We set the condition $\mu_i < P_i^{max}$ to ensure the possibility of positive expected profits for the company.

Our approach represents a stochastic setup for the acceptance of an offered contract, in contrast to [66] who work with (4.3) as a deterministic demand function for a given price. While for expected profits as dealt with in Section 4.3 this difference

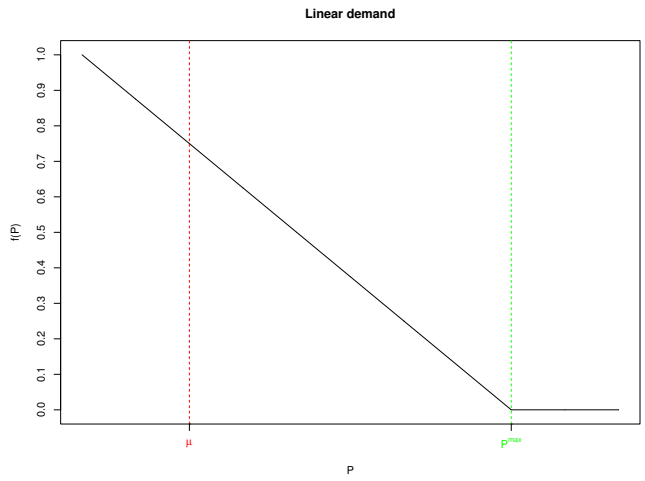

Figure 4.1: Acceptance probability $f_i(P)$ as a function of offered premium P.

does not matter, it will be important for the risk considerations in the subsequent sections. Despite the non-differentiability of the piece-wise linear f_i at P_i^{max}, this form will lead to simple local solutions. An extension of the results to a more complex, but analytically better tractable sigmoid form will be considered in Section 4.6.

Remark 4.2.1. *In the classical work of Rothschild and Stiglitz [123] as well as further models based on it, in a comparable setting of identifying contracts to sell, the authors find optimal solutions in terms of premiums and levels of coverage. For the present purpose and simplicity we, however, prefer to use the concept of the acceptance function, with individuals facing a binary choice of entering the contract or not. That is, we do not allow a partial coverage or deductibles here. Allowing partial coverages would change the model in some aspects. In Section 4.3, the difference between both settings would not be observable, as in expectation one may interpret the acceptance function of the individuals as their level of coverage instead of the probability to accept the contract. On the contrary, in Section 4.4, the variance would differ, as the partial coverages would not induce a variability in the number of insured individuals as the acceptance function does.*

In the following sections, we introduce three different scenarios that the insurer may face. We start with the case where the insurer can observe the risk type of each individual, which we refer to as the full information case. We then consider the situation where the insurer can not distinguish between the two types ex-ante at all. In that case the only possible method of pricing is to not differentiate individuals, and the profitability of the insurance business will then depend on the empirical fraction of each risk type in the population. Finally, the possibility to observe and measure a certain characteristic, which can be discrete or continuous, allows us to classify an observation, but with a certain error probability. This probability depends on the true class of the observation. The introduction of the classification mechanism has certain costs, but allows to better price according to the true risk class. We are interested to quantitatively assess the respective trade-off in this simple model setup.

4.3 Expected profit in three scenarios

In this section, we focus on the expected profit only. Let us introduce three different scenarios that the insurer may face, starting with the case where the insurer can actually observe the risk type of each individual.

4.3.1 Full information

The benchmark for our analysis is the situation with no asymmetry of information. Here the insurer can observe the risk type of each individual and therefore price according to the true type. Recall that we know the actual proportion $\{p_L, p_H\}$ of the population in each class. Therefore, we can maximize profit by differentiating between groups. If, for a given individual, the price is higher than its true risk premium, his/her willingness to accept the contract decreases. Let us denote by $X_j^{(L)}$ the j^{th} loss random variable of the low risk type (independent copies of L) and by $X_j^{(H)}$ the j^{th} high-risk loss variable (independent copies of H). Then, the (random) profit is given by

$$\Pi = \sum_{j=1}^{n\cdot p_L} I_j^L (P_L - X_j^{(L)}) + \sum_{j=1}^{n\cdot p_H} I_j^H (P_H - X_j^{(H)}), \tag{4.4}$$

where I_j^L and I_j^H are independent Bernoulli random variables with probabilities $f_L(P_L)$ and $f_H(P_H)$, respectively. In this case, an adaptation of [66] establishes that the optimal premiums are independent of n, p_L and p_H, and they are simply the average of the mean claim size and the maximum premium that the policyholders are willing to pay.

Theorem 4.3.1. *In the full information case, the expected profit of the insurer is maximized by the premium choice*

$$P_L = \frac{1}{2}(\mu_L + P_L^{max}) \tag{4.5}$$

and

$$P_H = \frac{1}{2}(\mu_H + P_H^{max}) \tag{4.6}$$

for the two risk classes.

Proof. The optimal premium choice is the solution of the following optimization problem:

$$\{P_L, P_H\} = \underset{x,y}{\operatorname{argmax}}\, \mathbb{E}(\Pi) = \underset{x,y}{\operatorname{argmax}}\, np_L f_L(x)(x-\mu_L) + np_H f_H(y)(y-\mu_H). \tag{4.7}$$

We notice that $\mathbb{E}(\Pi)$ is a continuous function of $\{x, y\}$ and that for $\{x < \mu_L,\ y < \mu_H\}$, $\mathbb{E}(\Pi) < 0$; $\{x = \mu_L,\ y = \mu_H\}$, $\mathbb{E}(\Pi) = 0$; $\{x > \mu_L,\ y > \mu_H\}$, $\mathbb{E}(\Pi) > 0$. Also, $\lim_{x\to+\infty, y\to+\infty} \mathbb{E}(\Pi) = 0$, which means that $\mathbb{E}(\Pi)$ admits a strictly positive maximum for some $\{x > \mu_L,\ y > \mu_H\}$ (and the optimization in x and y can in fact be separated). We can characterize this point by the following equations:

$$\frac{\partial \mathbb{E}(\Pi)}{\partial x} = np_L\left(f_L'(x)x + f_L(x)\right) - np_L f_L'(x)\mu_L \overset{!}{=} 0 \tag{4.8}$$

$$\frac{\partial \mathbb{E}(\Pi)}{\partial y} = np_H \left(f'_H(y)y + f_H(y)\right) - np_H f'_H(y)\mu_H \stackrel{!}{=} 0, \tag{4.9}$$

where the $\stackrel{!}{=}$ operator denotes a necessary condition. From (4.8) we have

$$\begin{aligned}\frac{\partial \mathbb{E}(\Pi)}{\partial x} =& np_L \left(f'_L(x)x + f_L(x)\right) - np_L f'_L(x)\mu_L = 0 \\ \iff & np_L \left(-\frac{1}{P_L^{max}}\right)(x - \mu_L) + np_L \left(1 - \frac{x}{P_L^{max}}\right) = 0 \\ \iff & x = \frac{1}{2}(\mu_L + P_L^{max}).\end{aligned}$$

Since the optimal solution P_L respects the condition $P_L \leq P_L^{max}$, as $\mu_L < P_L^{max}$, it is not necessary to distinguish cases of the piece-wise function. To prove it is indeed a local and global maximum, we can easily prove that the second derivative is negative for a linear f:

$$\begin{aligned}\frac{\partial^2 \mathbb{E}(\Pi)}{\partial x^2} =& np_L f''_L(x)(x - \mu_L) + 2np_L f'_L(x) \\ =& 2np_L \left(-\frac{1}{P_L^{max}}\right) < 0.\end{aligned}$$

The same reasoning holds for the first and second derivative w.r.t. y. □

Remark 4.3.2. *It is easy to check that if*

$$\mu_L + P_L^{max} < 2\mu_H, \tag{4.10}$$

then $P_L < \mu_H$, in which case charging a low-risk type premium to a high-risk type customer results in an expected loss on the individual level. □

In terms of sensitivities, we simply see from (4.5) and (4.6) that

$$\frac{\partial P_i}{\partial \mu_i} = \frac{\partial P_i}{\partial P_i^{max}} = \frac{1}{2}$$

for both risk types $i \in \{L, H\}$. That is, the reactivity of the optimal premium is constant for variation in the mean loss. Consequently, in case of increasing losses, the increase in premium will only cover half of the increase in losses, thus decreasing profits by the double effect of smaller margins and smaller acceptance rate. Similarly, shifting the endpoint P_i^{max} of the acceptance function (for invariant μ_i) also increases the optimal chargeable premium linearly with slope $1/2$.

4.3.2 No differentiation

Next, let us consider the situation where the insurer has no possibility to distinguish between risk types on the individual level, but still has an estimate for the fractions $\{p_L, p_H\}$ of the population in each class (e.g. through some historical figures). So we assume these numbers to be known ($p_L + p_H = 1$). In this scenario, the insurer proposes an identical premium P to every individual. This has the advantage that

one saves the cost of identification of risk types, and provides another benchmark for the sequel. The profit in this case is

$$\Pi = \sum_{j=1}^{np_L} I_j^L (P - X_j^{(L)}) + \sum_{j=1}^{np_H} I_j^H (P - X_j^{(H)}),$$

where I_j^L and I_j^H are Bernoulli random variables with probabilities $f_L(P)$ and $f_H(P)$ respectively, and the optimal premium then amounts to

$$P = \underset{z}{\operatorname{argmax}} \mathbb{E}(\Pi) = \underset{z}{\operatorname{argmax}}\, n(p_L f_L(z) + p_H f_H(z))z - np_L f_L(z)\mu_L - np_H f_H(z)\mu_H. \tag{4.11}$$

Comparing the resulting optimization problem

$$\max_z np_L f_L(z)(z - \mu_L) + np_H f_H(z)(z - \mu_H) \tag{4.12}$$

with (4.7), we see from $P_L \neq P_H$ there that the optimal solution P to (4.12) will now yield a smaller profit (which is intuitive, since we have less information available than in the setup of Section 4.3.1).

Theorem 4.3.3. *In the no differentiation case, the expected profit of the insurer is maximized by the premium choice*

$$P = a^* P_L + (1 - a^*) P_H, \tag{4.13}$$

where P_L and P_H are the optimal premiums of the full information case given in (4.5) *and* (4.6) *and* $a^* = \frac{p_L / P_L^{max}}{p_L / P_L^{max} + p_H / P_H^{max}}$.

Proof. Problem (4.12) can be solved using the first order condition

$$\begin{aligned}\frac{\partial \mathbb{E}(\Pi)}{\partial z} =& np_L \left(f_L'(z)z + f_L(z)\right) + np_H \left(f_H'(z)z + f_H(z)\right) \\ & - np_L f_L'(z)\mu_L - np_H f_H'(z)\mu_H \overset{!}{=} 0.\end{aligned}$$

Using $z < \min\{P_L^{max}, P_H^{max}\}$, plugging in the function f yields

$$\begin{aligned}& np_L \left(-\frac{1}{P_L^{max}}\right)(z - \mu_L) + np_L \left(1 - \frac{z}{P_L^{max}}\right) \\ & + np_H \left(-\frac{1}{P_H^{max}}\right)(z - \mu_H) + np_H \left(1 - \frac{z}{P_H^{max}}\right) = 0,\end{aligned}$$

which leads to

$$z = \frac{\frac{p_L}{P_L^{max}} \frac{\mu_L + P_L^{max}}{2} + \frac{p_H}{P_H^{max}} \frac{\mu_H + P_H^{max}}{2}}{\frac{p_L}{P_L^{max}} + \frac{p_H}{P_H^{max}}}$$

and finally (4.13).

The second order condition yields a strictly negative result, thus confirming the global maximum. □

Expression (4.13) shows that P is the average of the optimal premiums under full information, weighted by the proportions in the population and the maximum affordable premiums. Under the assumption (4.10), this also establishes

$$P_L \leq P \leq P_H.$$

Remark 4.3.4. *One should be careful to check whether $P > P_L^{max}$: in that case, L type customers do not enter the contract. This happens if $a^* < \frac{P_H - P_L^{max}}{P_H - P_L}$. Consequently, the optimal premium is that for higher risk types only, meaning $P = P_H$. If the expected profit for $P = P_H$ is greater than the one found above, then the optimal premium will be P_H and only H types will enter the contract.* □

The change of the expected profit when compared to the case of full information can now also be expressed as

$$\underbrace{np_L\left(f_L(P) - f_L(P_L)\right)(P_L - \mu_L)}_{\text{loss on } L \text{ not entering the contract}} + \overbrace{np_L f_L(P)(P - P_L)}^{\text{gain on extra margin on } L}$$

$$+ \underbrace{np_H\left(f_H(P) - f_H(P_H)\right)(P_H - \mu_H)}_{\text{gain on more } H \text{ entering the contract}} + \overbrace{np_H f_H(P)(P - P_H)}^{\text{loss on reduced margin on } H} < 0. \quad (4.14)$$

In particular, for low risk types the proposed premium P is higher than their appropriate optimal premium P_L under full information. Thus, with the decreasing shape of the acceptance function, on average the insurer loses low-risk type customers and the associated expected profit (negative first term in (4.14)). At the same time, those who remain bring higher profits (the second term in (4.14)). Correspondingly, due to the cheaper than appropriate premium P_H, more high-risk type customers join (positive third term in (4.14)), but they pay less premium now (negative fourth term).

4.3.3 Differentiation in two classes

Assume now that the insurer does not know the individuals' risk type, but has access to a mechanism that can assign (*classify*) the risk types correctly with a certain probability. Assume that the probability of misclassification is the same for each policyholder of the same type and given by

$$p_{H|L} := \mathbb{P}\left(i \text{ is classified as } H \mid i \in L\right),$$
$$p_{L|H} := \mathbb{P}\left(i \text{ is classified as } L \mid i \in H\right).$$

Remark 4.3.5. *If $p_{H|L} = p_{L|H} = 0$, we get back to the full information setting, as there is no classification error. If $p_{H|L} = p_{L|H} = 1$, all the true H end up in the L group and all the true L are classified in the H group (which would also result in knowing the true type of each one, but having to switch the categories). In the cases when $p_{H|L} = 1$ and $p_{L|H} = 0$ or $p_{H|L} = 0$ and $p_{L|H} = 1$, all individuals are classified in the same group. Typically, there is a tradeoff between the two error types: in an attempt to classify one risk more accurately, the precision on the other one will go down. For instance, in order to minimize $p_{L|H}$, we could simply attribute all observations to group H, which indeed gives $p_{L|H} = 0$, but $p_{H|L}$ would increase drastically as all L observations are then erroneously identified as H.*

The cost $c(n)$ of applying the classification algorithm will increase with population size n (the computational cost of different algorithms is increasing in the sample size (take for instance the simplest Bayesian classifier [161] with linear complexity),

the human time invested in analysing data and making decisions increases, and more powerful machines may be needed to run the algorithms, just to name a few reasons). At the same time, the marginal cost is likely to decrease in n (fixed costs in the process can be divided onto more policyholders, the insurer gains experience and recognizes patterns etc.). Hence, we define

$$c : \mathbb{R}^+ \mapsto \mathbb{R}^+,\ c'(n) \geq 0,\ c''(n) \leq 0.$$

A mathematically simple candidate for such a function is

$$c(n) = c_0 \log(\gamma n),$$

where γ offsets for the minimal cost amount and c_0 scales for the intensity of the effect of the population size.

The insurer will propose premiums, P_L^* and P_H^*, different from the ones in Section 4.3.1 under full information, and some customers receive 'wrong' offers, leading to a different customer behaviour with respect to accepting the contract. Figure 4.2 visualizes the pricing process. An initial population of n customers is subdivided into groups by their true risk type, rather than their identified risk type, and finally the insurer loses some customers because of the entailed acceptance patterns of policies.

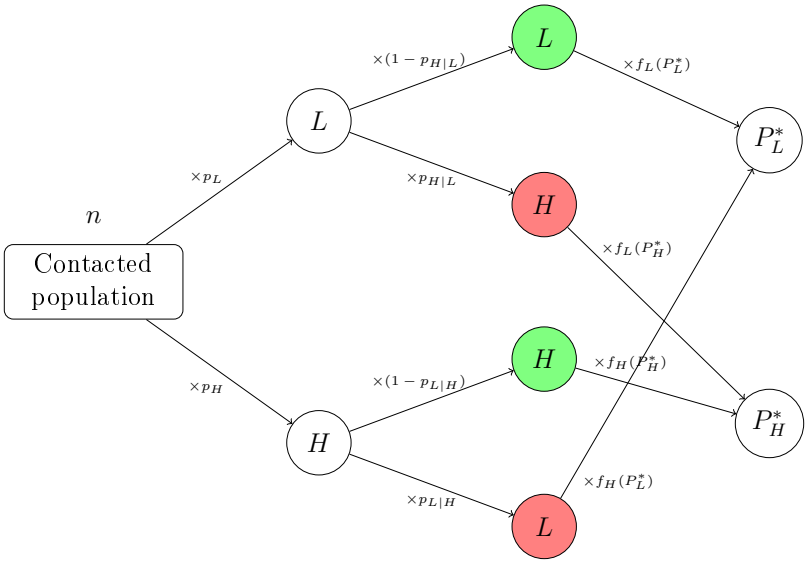

Figure 4.2: Visualisation of the pricing process.

In this situation, the profit is given by

$$\Pi = \sum_{j=1}^{n \cdot p_L \cdot (1-p_{H|L})} I_j^{L|L} \left(P_L^* - X_j^{(L)} \right) + \sum_{j=1}^{n \cdot p_L \cdot p_{H|L}} I_j^{H|L} \left(P_H^* - X_j^{(L)} \right)$$
$$+ \sum_{j=1}^{n \cdot p_H \cdot (1-p_{L|H})} I_j^{H|H} \left(P_H^* - X_j^{(H)} \right) + \sum_{j=1}^{n \cdot p_H \cdot p_{L|H}} I_j^{L|H} \left(P_L^* - X_j^{(H)} \right),$$

where $I_j^{L|L}$, $I_j^{H|L}$, $I_j^{H|H}$ and $I_j^{L|H}$ are Bernoulli random variables with parameters $f_L(P_L^*)$, $f_L(P_H^*)$, $f_H(P_H^*)$ and $f_H(P_L^*)$ respectively. The optimization procedure now

amounts to

$$\begin{aligned}\{P_L^*, P_H^*\} = &\underset{v,w}{\operatorname{argmax}}\, \mathbb{E}(\Pi) = \underset{v,w}{\operatorname{argmax}}\ np_L(1-p_{H|L})f_L(v)(v-\mu_L)\\ &+np_Lp_{H|L}f_L(w)(w-\mu_L) + np_H(1-p_{L|H})f_H(w)(w-\mu_H)\\ &+np_Hp_{L|H}f_H(v)(v-\mu_H) - c(n).\end{aligned} \tag{4.15}$$

Theorem 4.3.6. *In the differentiation case, the expected profit of the insurer is maximized by the premium choice*

$$P_L^* = b^*P_L + (1-b^*)P_H \tag{4.16}$$

and

$$P_H^* = c^*P_L + (1-c^*)P_H \tag{4.17}$$

for the two classified risk classes, where P_L and P_H are the optimal premiums (4.5) *and* (4.6) *of the full information case, $b^* = \frac{p_L(1-p_{H|L})/P_L^{max}}{p_L(1-p_{H|L})/P_L^{max}+p_Hp_{L|H}/P_H^{max}}$ and $c^* = \frac{p_Lp_{H|L}/P_L^{max}}{p_H(1-p_{L|H})/P_H^{max}+p_Lp_{H|L}/P_L^{max}}$.*

Proof. We make use of the following first order conditions from (4.15) to determine the optimal solution:

$$\begin{aligned}\frac{\partial\mathbb{E}(\Pi)}{\partial v} =&np_L(1-p_{H|L})\left(f_L'(v)v + f_L(v)\right) + np_Hp_{L|H}\left(f_H'(v)x + f_H(v)\right)\\ &- np_L(1-p_{H|L})f_L'(v)\mu_L - np_Hp_{L|H}f_H'(v)\mu_H \overset{!}{=} 0\end{aligned} \tag{4.18}$$

$$\begin{aligned}\frac{\partial\mathbb{E}(\Pi)}{\partial w} =&np_Lp_{H|L}\left(f_L'(w)w + f_L(w)\right) + np_H(1-p_{L|H})\left(f_H'(w)w + f_H(w)\right)\\ &- np_Lp_{H|L}f_L'(w)\mu_L - np_H(1-p_{L|H})f_H'(w)\mu_H \overset{!}{=} 0\end{aligned} \tag{4.19}$$

Equation (4.18) yields

$$\begin{aligned}&np_L(1-p_{H|L})\left(-\frac{1}{P_L^{max}}\right)(v-\mu_L) + np_L(1-p_{H|L})\left(1-\frac{v}{P_L^{max}}\right)\\ &+ np_Hp_{L|H}\left(-\frac{1}{P_H^{max}}\right)(v-\mu_H) + np_Hp_{L|H}\left(1-\frac{v}{P_H^{max}}\right) = 0,\end{aligned} \tag{4.20}$$

leading to (4.16). Formula (4.17) is obtained in a completely analogous way from (4.19).

The second order conditions are strictly negative and thus confirm the maximum. □

Like in the no differentiation case, the optimal premiums can again be expressed simply as a weighted average of the optimal premiums from the full information case, and the weights now involve the error probabilities.

Remark 4.3.7. *One needs to verify the limiting case of $P_H^* = P_L^{max}$ and $P_L^* = P_L^{max}$ to obtain the true maximum, since misclassified L individuals may not enter the contract after the limiting premium. This is the case when $c^* < \frac{P_H-P_L^{max}}{P_H-P_L}$ and*

$b^* < \frac{P_H - P_L^{max}}{P_H - P_L}$. *One can distinguish three cases. Firstly, if both* $P_H^* < P_L^{max}$ *and* $P_L^* < P_L^{max}$, *then the optimal solutions are given by Equations* (4.16) *and* (4.17). *Secondly, if only* $P_H^* > P_L^{max}$, *then the correct premium for the proposed* H *contract should be* $P_H^* = P_H$, *since we correctly price for only* H *types entering the group.* H *types will always enter the contract since their premium is a weighted average of* P_L *and* P_H, *and both are smaller than* P_H^{max} *from Section 4.3.1. Thirdly, if both* $P_H^* > P_L^{max}$ *and* $P_L^* > P_L^{max}$, *which could happen with a high proportion of misclassified* H *individuals, then the optimal solution would be to offer the contract only to* H *types by setting* $P_L^* = P_H$ *and* $P_H^* = P_H$. *It is worthwhile to notice that the second-order mixed partial derivatives* $\frac{\partial^2 \mathbb{E}(\Pi)}{\partial v \partial w} = \frac{\partial^2 \mathbb{E}(\Pi)}{\partial w \partial v} = 0$, *and therefore the optimal price for the low risk types does not depend on the optimal price for the high risk types and vice versa.* □

What is of particular interest is the situation where H individuals are wrongly classified as L. Indeed, since $P_i > \mu_i$, this is the only situation where the insurer makes losses, so it is important to maintain control over this group. The loss (presented here as a negative gain) compared to the benchmark of the situation of full information can be decomposed into

$$
\begin{aligned}
&\underbrace{np_L(1-p_{H|L})\left[(f_L(P_L^*)-f_L(P_L))\,(P_L-\mu_L)+f_L(P_L^*)(P_L^*-P_L)\right]}_{\text{True } L\text{: Loss on } L \text{ not entering the contract and gain on those who remain}}\\
&+\underbrace{np_Lp_{H|L}\left[(f_L(P_H^*)-f_L(P_L))\,(P_L-\mu_L)+f_L(P_H^*)(P_H^*-P_L)\right]}_{\text{False } H\text{: Loss on } L \text{ not entering the contract and gain on those who remain}}\\
&+\underbrace{np_H(1-p_{L|H})\left[(f_H(P_H^*)-f_H(P_H))\,(P_H-\mu_H)+f_H(P_H^*)(P_H^*-P_H)\right]}_{\text{True } H\text{: Gain on extra } H \text{ entering the contract and loss on them underpriced}}\\
&+\underbrace{np_Hp_{L|H}\left[(f_H(P_L^*)-f_H(P_H))\,(P_H-\mu_H)+f_H(P_L^*)(P_L^*-P_H)\right]}_{\text{False } L\text{: Gain on extra } H \text{ entering the contract and loss on them underpriced}}\\
&-\underbrace{c(n)}_{\text{Invested cost}}\quad .
\end{aligned}
$$

Recall that P was the optimal uniform premium for the case without differentiation. Differentiation of risk types only makes sense, if the resulting premiums P_L^*, P_H^* satisfy $P_L^* \leq P \leq P_H^*$ (cf. Figure 4.3). From Equations (4.13), (4.16) and

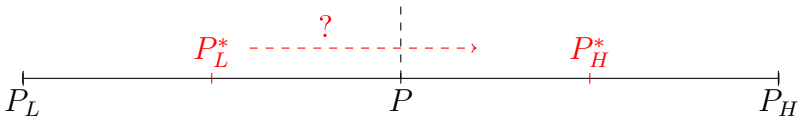

Figure 4.3: Illustration of premiums in different scenarios.

(4.17), this amounts to the condition $c^* \leq a^* \leq b^*$ which can easily translated to the following condition on the error probabilities:

$$p_{H|L} + p_{L|H} \leq 1.$$

This will always be fulfilled in practically relevant situations.

4.4 A mean-variance analysis

Proposing a unique premium P to both categories of risks attracts a higher relative proportion of H than the differentiating strategies. This heterogeneity in the portfolio composition generates a higher level of risk, which one should consider in the underwriting process. As the expected profit considered in Section 4.3 does not capture this aspect of the problem, we introduce the variance of the profit as a simple indicator that can be easily implemented in practical settings, as one only needs estimates for the first two moments of the underlying claim distributions for the analysis.

Define by N_L, N_H the (random) number of insured persons of risk type L and H, respectively, entering the contract. Their first two moments are summarized in Table 4.1, where P_L, P_H are given by Equations (4.5), (4.6), P by Equation (4.13) and P_L^*, P_H^* by Equations (4.16), (4.17).

Table 4.1: Expected value and variance of the number of insured for each scenario.

	Full information	No differentiation	Differentiation
$\mathbb{E}(N_L)$	$np_L f_L(P_L)$	$np_L f_L(P)$	$np_L(1-p_{H\|L})f_L(P_L^*)$ $+np_L p_{H\|L} f_L(P_H^*)$
$\mathbb{V}\mathrm{ar}(N_L)$	$np_L f_L(P_L)(1-f_L(P_L))$	$np_L f_L(P)(1-f_L(P))$	$np_L(1-p_{H\|L})f_L(P_L^*)(1-f_L(P_L^*))$ $+np_L p_{H\|L} f_L(P_H^*)(1-f_L(P_H^*))$
$\mathbb{E}(N_H)$	$np_H f_H(P_H)$	$np_H f_H(P)$	$np_H(1-p_{L\|H})f_H(P_H^*)$ $+np_H p_{L\|H} f_H(P_L^*)$
$\mathbb{V}\mathrm{ar}(N_H)$	$np_H f_H(P_H)(1-f_H(P_H))$	$np_H f_H(P)(1-f_H(P))$	$np_H(1-p_{L\|H})f_H(P_H^*)(1-f_H(P_H^*))$ $+np_H p_{L\|H} f_H(P_L^*)(1-f_H(P_L^*))$

For each risk type i, $N_i = \sum_{j=1}^{n} I_j^i$, where I_j^i are independent Bernoulli random variables with probability $f_i(P_i)$, so $\mathbb{E}(N_i) = nf_i(P_i)$ and variance $\mathbb{V}\mathrm{ar}(N_i) = nf_i(P_i)(1-f_i(P_i))$. The claim sizes are $X_j^{(i)}$ and the premium is P_i. From (4.4), we then get

$$\begin{aligned}\mathbb{V}\mathrm{ar}(\Pi_i) &= \mathbb{V}\mathrm{ar}\left(\sum_{j=1}^{N_i}\left(P_i - X_j^{(i)}\right)\right)\\ &= \mathbb{E}\left(\mathbb{V}\mathrm{ar}\left(\sum_{j=1}^{N_i}\left(P_i - X_j^{(i)}\right) \mid N_i\right)\right) + \mathbb{V}\mathrm{ar}\left(\mathbb{E}\left(\sum_{j=1}^{N_i}\left(P_i - X_j^{(i)}\right) \mid N_i\right)\right)\\ &= \mathbb{E}(N_i)\cdot\sigma_i^2 + (P_i-\mu_i)^2\cdot\mathbb{V}\mathrm{ar}(N_i).\end{aligned} \tag{4.21}$$

With this ingredient, we can now derive the variance of the profit in our three scenarios introduced in the previous section.

- Full information:

$$\mathbb{V}\mathrm{ar}(\Pi) = \mathbb{E}(N_L)\,\sigma_L^2 + (P_L-\mu_L)^2\mathbb{V}\mathrm{ar}(N_L) + \mathbb{E}(N_H)\,\sigma_H^2 + (P_H-\mu_H)^2\mathbb{V}\mathrm{ar}(N_H).$$

- No differentiation:

$$\mathbb{V}\mathrm{ar}(\Pi) = \mathbb{E}(N_L)\,\sigma_L^2 + (P-\mu_L)^2\mathbb{V}\mathrm{ar}(N_L) + \mathbb{E}(N_H)\,\sigma_H^2 + (P-\mu_H)^2\mathbb{V}\mathrm{ar}(N_H).$$

- Differentiation:

$$\begin{aligned}\mathbb{V}\mathrm{ar}\,(\Pi) =& np_L(1-p_{H|L})f_L(P_L^*)\sigma_L^2 + (P_L^*-\mu_L)^2 np_L(1-p_{H|L})f_L(P_L^*)(1-f_L(P_L^*))\\ &+np_Lp_{H|L}f_L(P_H^*)\sigma_L^2 + (P_H^*-\mu_L)^2 np_Lp_{H|L}f_L(P_H^*)(1-f_L(P_H^*))\\ &+np_H(1-p_{L|H})f_H(P_H^*)\sigma_H^2 + (P_H^*-\mu_H)^2 np_H(1-p_{L|H})f_H(P_H^*)(1-f_H(P_H^*))\\ &+np_Hp_{L|H}f_H(P_L^*)\sigma_H^2 + (P_H^*-\mu_L)^2 np_Hp_{L|H}f_H(P_L^*)(1-f_H(P_L^*)).\end{aligned}$$

Remark 4.4.1. *Note that in all of the above expressions a term containing the variance due to the randomness of claim sizes is followed by one with the variance due to the randomness of underwriting, i.e. the customer's probability to enter the contract or not. In Section 4.5, we will illustrate this decomposition with the help of a numerical example.*

With external pressures coming from different stakeholders, it may be preferable to not price-differentiate until a certain variance level. This threshold will be higher with increasing classification costs and error probabilities. Consequently, we introduce a variance constraint in the optimization problem, modifying the problem from Section 4.3 to

$$\begin{aligned}&\max \mathbb{E}\,(\Pi)\\ &\text{s.t.}\mathbb{V}\mathrm{ar}\,(\Pi) \leq \bar{\sigma}^2.\end{aligned}$$

Varying the value of $\bar{\sigma}$ will lead to a mean-variance efficient frontier in the spirit of Markowitz [102]. Introduce the Lagrange multipliers

$$\mathcal{L}(P_i,\lambda) = \mathbb{E}\,(\Pi(P_i)) + \lambda\left(\bar{\sigma}^2 - \mathbb{V}\mathrm{ar}\,(\Pi(P_i))\right) \tag{4.22}$$

for the premium P_i in any of the optimization programs (4.7), (4.11) and (4.15). The optimal premiums are then obtained by the first order conditions

$$\begin{aligned}\frac{\partial\mathcal{L}}{\partial P_i} &= \frac{\partial\mathbb{E}\,(\Pi)}{\partial P_i} - \lambda\frac{\partial\mathbb{V}\mathrm{ar}\,(\Pi)}{\partial P_i} \stackrel{!}{=} 0,\\ \frac{\partial\mathcal{L}}{\partial\lambda} &= \bar{\sigma}^2 - \mathbb{V}\mathrm{ar}\,(\Pi) \stackrel{!}{=} 0.\end{aligned}$$

We give here the corresponding equations for the full information case, the other cases follow in an analogous way. Equation (4.22) translates into

$$\begin{aligned}\mathcal{L}(P_L,P_H) =& n\left(p_L\left(1-\frac{1}{P_L^{max}}P_L\right)(P_L-\mu_L) + p_H\left(1-\frac{1}{P_H^{max}}P_H\right)(P_H-\mu_H)\right)\\ &+\lambda\left\{\bar{\sigma}^2 - \left[\left(np_L\left(1-\frac{1}{P_L^{max}}P_L\right)\right)^2 + np_L\left(1-\frac{1}{P_L^{max}}P_L\right)\right]\sigma_L^2\right.\\ &-(P_L-\mu_L)^2 np_L\left(1-\frac{1}{P_L^{max}}P_L\right)\frac{P_L}{P_L^{max}}\\ &-\left[\left(np_H\left(1-\frac{1}{P_H^{max}}P_H\right)\right)^2 + np_H\left(1-\frac{1}{P_H^{max}}P_H\right)\right]\sigma_H^2\\ &\left.-(P_H-\mu_H)^2 np_H\left(1-\frac{1}{P_H^{max}}P_H\right)\frac{P_H}{P_H^{max}}\right\}.\end{aligned}$$

The first order conditions are given by

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial P_L} =& np_L \left(-\frac{1}{P_L^{max}}\right)(P_L-\mu_L) + np_L\left(1-\frac{1}{P_L^{max}}P_L\right)\\
&-\lambda\left(np_L\left(-\frac{1}{P_L^{max}}\right)\sigma_L^2 + 2(P_L-\mu_L)np_L\left(1-\frac{1}{P_L^{max}}P_L\right)\frac{P_L}{P_L^{max}}\right.\\
&\left.+(P_L-\mu_L)^2 np_L \frac{1}{P_L^{max}}\left(1-2\frac{P_L}{P_L^{max}}\right)\right) \stackrel{!}{=} 0,\\
\frac{\partial \mathcal{L}}{\partial P_H} =& np_H \left(-\frac{1}{P_H^{max}}\right)(P_H-\mu_H) + np_H\left(1-\frac{1}{P_H^{max}}P_H\right)\\
&-\lambda\left(np_H\left(-\frac{1}{P_H^{max}}\right)\sigma_H^2 + 2(P_H-\mu_H)np_H\left(1-\frac{1}{P_H^{max}}P_H\right)\frac{P_H}{P_H^{max}}\right.\\
&\left.+(P_H-\mu_H)^2 np_H \frac{1}{P_H^{max}}\left(1-2\frac{P_H}{P_H^{max}}\right)\right) \stackrel{!}{=} 0,\\
\frac{\partial \mathcal{L}}{\partial \lambda} =& \bar{\sigma}^2 - \left(np_L\left(1-\frac{1}{P_L^{max}}P_L\right)\sigma_L^2 + (P_L-\mu_L)^2 np_L\left(1-\frac{1}{P_L^{max}}P_L\right)\frac{P_L}{P_L^{max}}\right.\\
&\left.+np_H\left(1-\frac{1}{P_H^{max}}P_H\right)\sigma_H^2 + (P_H-\mu_H)^2 np_H\left(1-\frac{1}{P_H^{max}}P_H\right)\frac{P_H}{P_H^{max}}\right) \stackrel{!}{=} 0.
\end{aligned}$$

This results in a system of three equations for the three unknowns P_L, P_H and λ which can be solved numerically for every choice of involved parameters.

4.5 Numerical illustrations

Let us now consider concrete numerical illustrations of the results of the previous sections. The following parametrization will be used throughout this section unless otherwise stated:

$$\mu_L = 1, \mu_H = 5, \sigma_L^2 = 1, \sigma_H^2 = 10, p_L = 0.9, p_{H|L} = p_{L|H} = 0.1. \tag{4.23}$$

For the shape of the cost function we assume $c(n) = 20\log n$ in the plots, but note that any other choice would be feasible as well.

4.5.1 Expected profit

Let $f_i(P)$ have the form (4.3) with $P_L^{max} = 4\mu_L = 4$ and $P_H^{max} = 4\mu_H = 20$. Then we get from the respective formulas of Section 4.3:

- Full information:

$$P_L = 2.5, \quad P_H = 12.5, \quad \mathbb{E}(\Pi) = 0.788\, n.$$

- No differentiation:

$$P \approx 2.717, \quad \mathbb{E}(\Pi) \approx 0.298\, n.$$

  Note that $\mu_L < P < \mu_H$. In this case, the insurer targets the low-risk L type customers because their proportion in the population is large enough to compensate for the losses on the H types.

- Differentiation in two classes:

$$P_L^* \approx 2.525, \quad P_H^* = 12.5, \quad \mathbb{E}(\Pi) \approx 0.687\, n - c(n).$$

  Note that $P_H^* > P_L^{max}$.

Applying the classification is hence only an advantage if

$$0.687\, n - c(n) \geq 0.298\, n.$$

Conversely, the maximum cost which the insurer will be willing to pay for the classification, given a population of size n, is

$$c(n) < 0.388\, n.$$

It is instructive to look into the sensitivity of the results. Let us first explore the variability of the profit under different error probabilities, which can be a helpful decision tool in case of limited investment resources. For each level of error probabilities, we recompute the optimal premiums. Figure 4.4 features the sensitivity of the expected profit with respect to both error probabilities. The classification cost still needs to be deducted here from the expected profit. If the insurer is given a choice of different classification algorithms or investment possibilities for improvement of precision with known resulting error probabilities, one can verify whether that investment is worthwhile. This figure may help the decision makers to judge

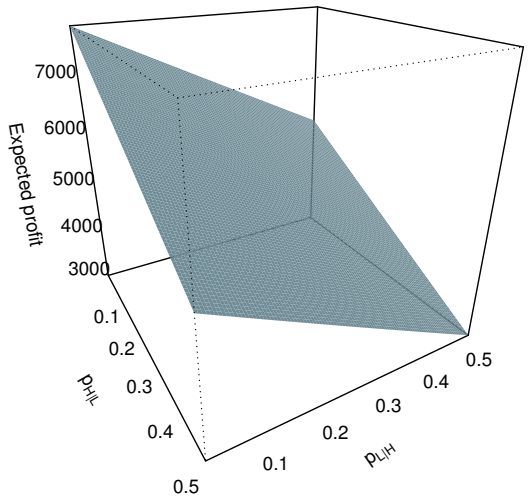

Figure 4.4: Expected profit as function of error probabilities.

whether with given error probabilities, a refinement in the classification may be of added value to the company.

4.5.2 Variance

Numerically, with the parameters defined in (4.23), we obtain the following results for the three cases of the linear acceptance function:

- Full information:
$$\mathbb{V}\mathrm{ar}\,(\Pi) = 2.505469n.$$
- No differentiation:
$$\mathbb{V}\mathrm{ar}\,(\Pi) = 1.79213n.$$
- Differentiation:
$$\mathbb{V}\mathrm{ar}\,(\Pi) = 2.355267n.$$

As the total variance is an increasing function of the number of policyholders, it will naturally be higher under a differentiation strategy, as the insurer gets more market share. But the structure of the variance will be different. Without differentiation, the variance inside the group is much higher than the average of internal group variances from the differentiation case, that difference being larger when the two distributions are further apart.

In Figure 4.5, we plot the variances in the three scenarios to illustrate their forms as a function of chosen premium. We split the variances according to the part stemming from the variability of claims (in red) and from the one of acceptance of contracts (in green). The humps indicate the region where the increase of variance due to the increasing deviation from the mean is compensated by the decrease in the number of underwritten policies. In Figures 4.5b and 4.5c, we can observe two humps, appearing because of the mixture of two risk types. With the help of this decomposition, we can clearly see that the humps in the plots come from the acceptance behaviour. In Figure 4.5a, we observe that if the premium becomes too high, the total variance decreases as the population does not enter the contract any more. In Figure 4.6 the same variance decompositions are depicted, but scaled with respect to the absolute value of the expected profit. In this way, one can compare the risk per unit of gained profit in each situation.

In Figures 4.7-4.9, we show the variances when varying one parameter at a time. In Figure 4.7, one can observe the variance shapes for a small range acceptance function. In this case, the variance is mostly defined by the claims behaviour, since the acceptance rate remains low and the humps are less pronounced. For an acceptance function ranging up to high premiums as shown in Figure 4.8, the acceptance variance dominates. The humps are more pronounced as policyholders exist in a broader range. Finally, with a bigger expected claim size difference between risk types, the relationship between the humps and the risk types becomes clearer as they are further apart, see Figure 4.9.

4.5.3 Mean-variance efficient frontier

To complete the numerical part, we now address the illustration of the mean-variance frontier as defined in Section 4.4 for a population size of $n = 10'000$. We see in Figure 4.10 that up to a certain variance level, the non-differentiation strategy dominates differentiation in terms of expected profit. This breaking point depends on the cost

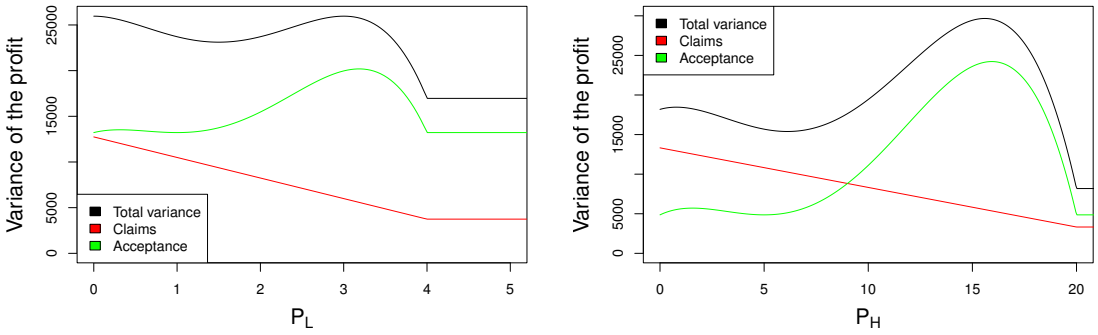

(a) Perfect information case. Left-hand side: $\mathbb{V}\mathrm{ar}(\Pi)$ as a function of P_L. Right-hand side: $\mathbb{V}\mathrm{ar}(\Pi)$ as a function of P_H.

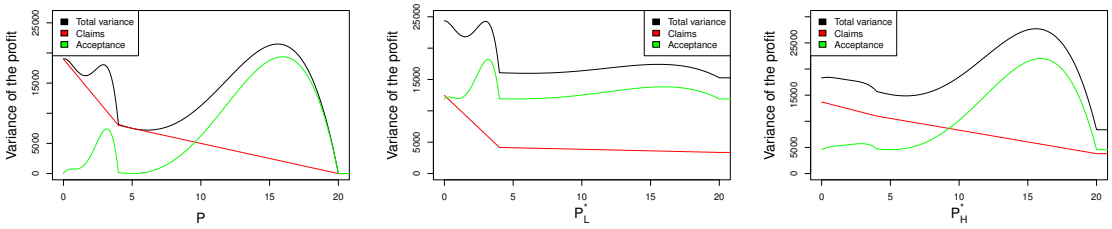

(b) No differentiation case. $\mathbb{V}\mathrm{ar}(\Pi)$ as a function of P.

(c) Differentiation case. Left-hand side: $\mathbb{V}\mathrm{ar}(\Pi)$ as a function of P_L^*. Right-hand side: $\mathbb{V}\mathrm{ar}(\Pi)$ as a function of P_H^*.

Figure 4.5: Decomposition of the variance.

function $c(n)$ and the error probabilities. One may also want to consider limiting constraints in practice such as regulatory constraints or the demands of stakeholders. The kinks in the frontier arise from the fact that for different variance limitations, a different portfolio composition becomes optimal. In other words, the optimal strategy switches in the points of the kinks by letting more of a lower or higher risk type entering the contract.

The mean-variance approach assumes variations to both sides as equally weighted since the variance is a symmetric risk measure. This framework can be extended to other risk measures, such as the lower semi-variance to take in account only one-sided deviations from the mean or the value-at-risk to consider minimal profit requirements. These adaptations are developed in Section 4.6, where we also present an alternative approach for the risk assessment based on utility functions.

4.6 Extensions

4.6.1 A sigmoid-type acceptance function

While the piece-wise linear acceptance functions used in this chapter allow for intuitive and transparent results, one may want to challenge this simplistic assumption. In this section we would like to extend the previous analysis to a possibly more realistic shape that still allows for an explicit treatment. Concretely, assume that

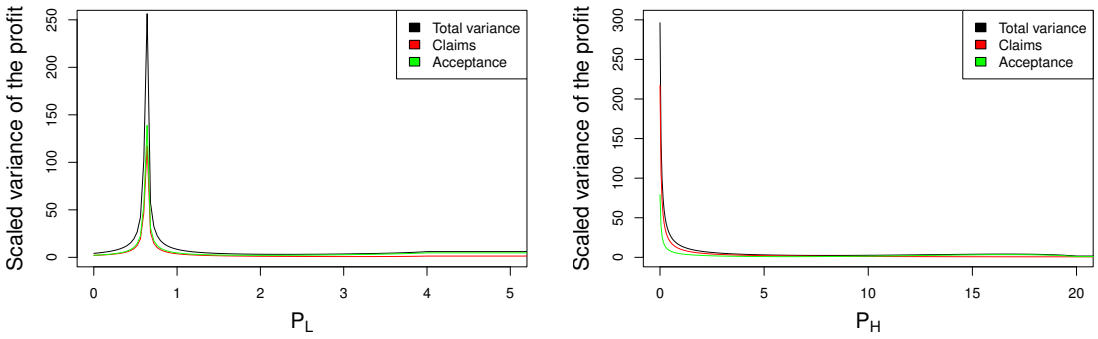

(a) Perfect information case. Left-hand side: $\mathbb{V}\mathrm{ar}\,(\Pi)$ divided by the absolute value of $\mathbb{E}\,(\Pi)$ as a function of P_L. Right-hand side: $\mathbb{V}\mathrm{ar}\,(\Pi)$ divided by the absolute value of $\mathbb{E}\,(\Pi)$ as a function of P_H.

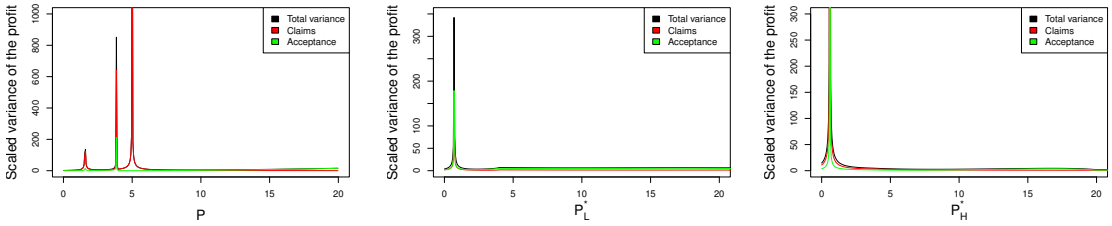

(b) No differentiation case. $\mathbb{V}\mathrm{ar}\,(\Pi)$ divided by the absolute value of $\mathbb{E}\,(\Pi)$ as a function of P.

(c) Differentiation case. Left-hand side: $\mathbb{V}\mathrm{ar}\,(\Pi)$ divided by the absolute value of $\mathbb{E}\,(\Pi)$ as a function of P_L^*. Right-hand side: $\mathbb{V}\mathrm{ar}\,(\Pi)$ divided by the absolute value of $\mathbb{E}\,(\Pi)$ as a function of P_H^*.

Figure 4.6: Decomposition of the variance scaled by the absolute value of the expectation.

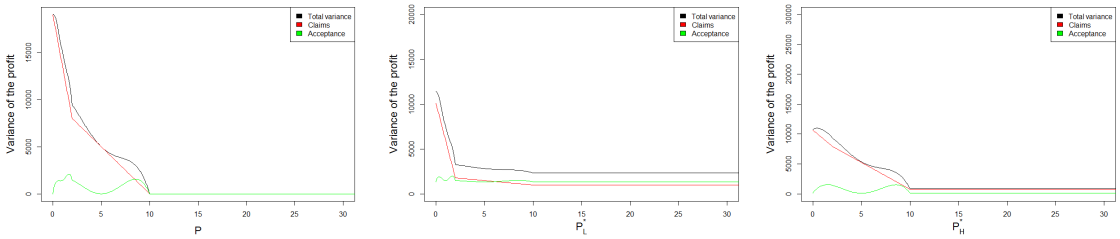

(a) No differentiation case.

(b) Differentiation case.

Figure 4.7: Decomposition of the variance, parameter $P_i^{max} = 2\mu_i$.

f belongs to the class of sigmoid functions, namely the logistic functions, which are smooth and monotone, thus suitable for our situation [90]. This form of function appears when applying a logit lapsing model with different risk factors, see e.g. [57, 70]. An example of a model using premiums as risk factors can be found in [28, 71] and particularly in [58]. Consider the following concrete shape of the acceptance function f_i of an individual of risk type i:

$$f_i(P) = \frac{1}{1 + e^{a_i(P - b_i)}},\ a_i \in \mathbb{R}^+,\ b_i \in \mathbb{R},\ i \in \{L, H\}, \tag{4.24}$$

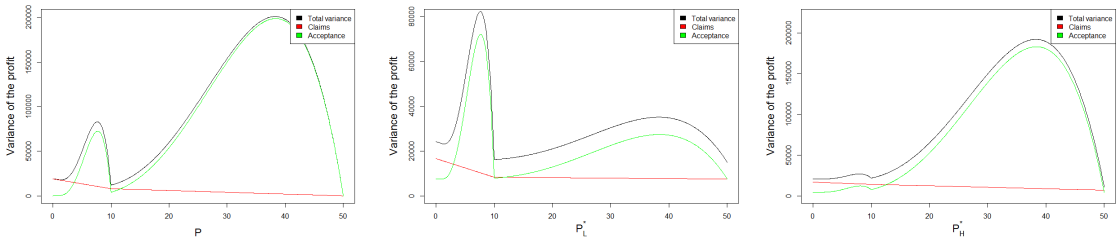

(a) No differentiation case. (b) Differentiation case.

Figure 4.8: Decomposition of the variance, parameter $P_i^{max} = 10\mu_i$.

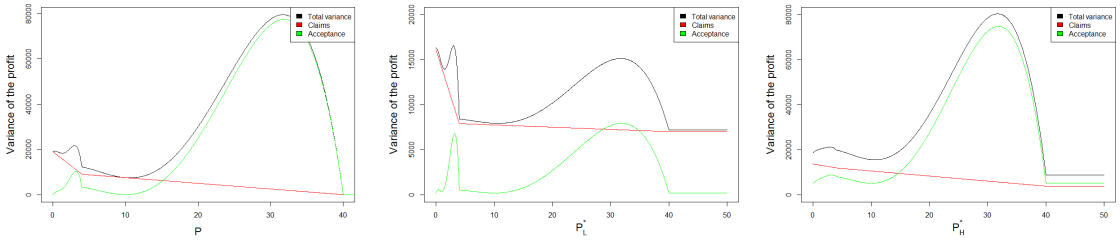

(a) No differentiation case. (b) Differentiation case.

Figure 4.9: Decomposition of the variance, parameter $\mu_H = 10$.

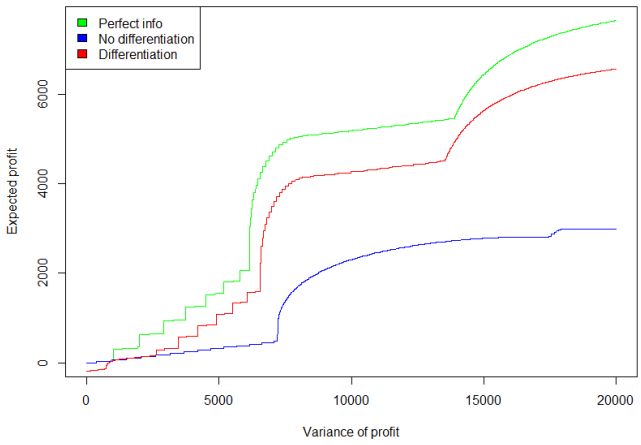

Figure 4.10: Mean-variance frontier with linear demand function.

where the parameters a_i and b_i need to be calibrated. We can suppose $b_i > \mu_i$, so that the function reaches value $1/2$ for premiums that are higher than the actuarially fair premium, cf. Figure 4.11. As a grows, the curve becomes steeper around the pivotal position determined by the parameter b (note that the choice of b also determines the value of f for $P = 0$ which will typically be smaller than 1). From an analytical point of view, the form (4.24) is more attractive than the piece-wise linear shape considered in the previous sections, as it is differentiable everywhere. Clearly, f_i is strictly decreasing in P:

$$\frac{\partial f_i(P)}{\partial P} = \frac{-a_i e^{a_i(P-b_i)}}{\left(1+e^{a_i(P-b_i)}\right)^2} < 0.$$

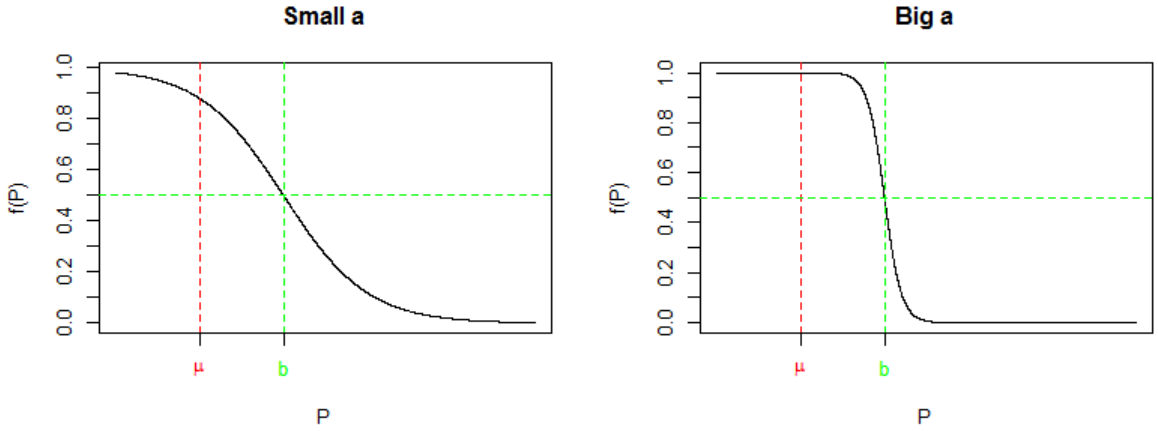

Figure 4.11: Acceptance function f for small (left) and large (right) parameter a.

Define further the *price elasticity* of a risk type as the change of the number of customers entering the contract with respect to the price variation:

$$E_P = \frac{\partial f_i(P)}{\partial P}\frac{P}{f_i} = \frac{-a_i e^{a_i(P-b_i)}P}{1+e^{a_i(P-b_i)}} = -a_i e^{a_i(P-b_i)} P f_i(P).$$

This measure illustrates the reactivity of the portfolio size to the variation of premium, cf. for instance [144, Ch.15].

Theoretical results

We first derive the analogous results to the ones in Sections 4.3, 4.4, under the sigmoid acceptance function. The full information case still leads to an explicit formula:

Theorem 4.6.1. *In the full information case, the expected profit of the insurer is maximized for the premium choice*

$$P_L = \mu_L + \frac{1}{a_L} + \frac{1}{a_L} W(e^{a_L b_L - a_L \mu_L - 1}) \tag{4.25}$$

and

$$P_H = \mu_H + \frac{1}{a_H} + \frac{1}{a_H} W(e^{a_H b_H - a_H \mu_H - 1}), \tag{4.26}$$

where $W(z)$ denotes the (principal branch of the) Lambert W function, which is the inverse function of $g(x) = xe^x$ (cf. [38]).

Proof. The optimal premium choice is the solution of the optimization problem

$$\{P_L, P_H\} = \underset{x,y}{\operatorname{argmax}}\, \mathbb{E}(\Pi) = \underset{x,y}{\operatorname{argmax}}\, n\left(p_L f_L(x)(x-\mu_L) + p_H f_H(y)(y-\mu_H)\right).$$

We can characterize the maxima by the equations

$$\frac{\partial \mathbb{E}(\Pi)}{\partial x} = np_L\left(f_L'(x)x + f_L(x)\right) - np_L f_L'(x)\mu_L \overset{!}{=} 0, \tag{4.27}$$

$$\frac{\partial \mathbb{E}(\Pi)}{\partial y} = np_H \left(f'_H(y)y + f_H(y)\right) - np_H f'_H(y)\mu_H \stackrel{!}{=} 0. \tag{4.28}$$

From (4.27) we have

$$\begin{aligned} & np_L \left(\frac{\partial \left(1 - \frac{1}{1+e^{-a_L(x-b_L)}}\right)}{\partial x}(x-\mu_L) + 1 - \frac{1}{1+e^{-a_L(x-b_L)}} \right) = 0 \\ \iff & a_L(x-\mu_L) = 1 + e^{-a_L(x-b_L)} \end{aligned}$$

leading to (4.25). Equation (4.26) for the high-risk individuals is then obtained in a completely analogous way.
To see that the extremal point is indeed a local maximum, one needs to verify

$$\frac{\partial^2 \mathbb{E}(\Pi)}{\partial x^2} = np_L f''_L(x)(x-\mu_L) + 2np_L f'_L(x) < 0.$$

Using $\{P_L > \mu_L,\ P_H > \mu_H\}$ and

$$\frac{\partial^2 f(P)}{\partial P^2} = \frac{a_i^2 e^{-a_i(P-b_i)}\left(1 - e^{-a_i(P-b_i)}\right)}{\left(1+e^{-a_i(P-b_i)}\right)^3},$$

we can rearrange the previous condition as

$$f''_L(x)(x-\mu_L) < -2f'_L(x).$$

Since $-2f'_L(x)$ is always positive, $f''_L(x) < 0$ for all $x < b_L$ and $x = P_L > \mu_L$. The same conclusion holds for the second derivative w.r.t. y. □

Remark 4.6.2. *Note that we can provide a necessary condition for P_L to be smaller than μ_H:*

$$e^{a_L(b_L-\mu_H)} < a_L(\mu_H - \mu_L) - 1. \tag{4.29}$$

This condition is of interest for analysing the case when the low-risk type premium yields losses in absolute terms if sold to a high-risk type. Also, one can easily obtain sensitivities of the premium with respect to the parameters by means of first derivatives.

$$\begin{aligned} \frac{\partial P_i}{\partial \mu_i} &= 1 + \frac{1}{a_i}\frac{\partial W}{\partial \mu_i} = 1 - \frac{W(e^{a_i b_i - a_i \mu_i - 1})}{1 + W(e^{a_i b_i - a_i \mu_i - 1})},\ \in [0,1]\ > 0, \\ \frac{\partial P_i}{\partial b_i} &= \frac{1}{a_i}\frac{\partial W}{\partial b_i} = \frac{W(e^{a_i b_i - a_i \mu_i - 1})}{1 + W(e^{a_i b_i - a_i \mu_i - 1})},\ \in [0,1]\ > 0, \\ \frac{\partial P_i}{\partial a_i} &= -\frac{1}{a_i^2} - \frac{1}{a_i^2}W + \frac{1}{a_i}\frac{\partial W}{\partial \mu_i} = \frac{b_i - \mu_i}{a_i}\frac{W(e^{a_i b_i - a_i \mu_i - 1})}{1 + W(e^{a_i b_i - a_i \mu_i - 1})} - \frac{1}{a_i^2}\left(1 + W(e^{a_i b_i - a_i \mu_i - 1})\right). \end{aligned}$$

Concerning the sign of the last term, under $b_i > \mu_i$ the first term is positive and the second is negative. The overall difference is negative for small values of a_i, but positive for larger a_i, that effect manifesting itself sooner if the difference $b_i - \mu_i$ is larger.

In case of no differentiation, we proceed as before by taking first order conditions of the expected profit defined above in Equation (4.11):

$$\begin{aligned}\frac{\partial \mathbb{E}(\Pi)}{\partial z} =& np_L\left(f_L'(z)z + f_L(z)\right) + np_H\left(f_H'(z)z + f_H(z)\right)\\ &- np_L f_L'(z)\mu_L - np_H f_H'(z)\mu_H \stackrel{!}{=} 0.\end{aligned}$$

Plugging in the sigmoid function f yields

$$\begin{aligned}&p_L\left(\frac{-a_L e^{-a_L(z-b_L)}}{\left(1+e^{-a_L(z-b_L)}\right)^2}(z-\mu_L) + 1 - \frac{1}{1+e^{-a_L(z-b_L)}}\right)\\ +&p_H\left(\frac{-a_H e^{-a_H(z-b_H)}}{\left(1+e^{-a_H(z-b_H)}\right)^2}(z-\mu_H) + 1 - \frac{1}{1+e^{-a_H(z-b_H)}}\right) = 0,\end{aligned}$$

which can be solved numerically. We establish that $P_L \leq P \leq P_H$, following the assumption in (4.29).

For the differentiation case, we have the problem defined in Equation (4.15) to solve. Once again, as the first order conditions are symmetric, we will detail only one of them.

$$\begin{aligned}\frac{\partial \mathbb{E}(\Pi)}{\partial v} =& np_L(1-p_{H|L})\left(f_L'(v)v + f_L(v)\right) + np_H p_{L|H}\left(f_H'(v)v + f_H(v)\right)\\ &- np_L(1-p_{H|L})f_L'(v)\mu_L - np_H p_{L|H} f_H'(v)\mu_H = 0\\ \iff& p_L(1-p_{H|L})\left(\frac{-a_L e^{-a_L(v-b_L)}}{\left(1+e^{-a_L(v-b_L)}\right)^2}(v-\mu_L) + 1 - \frac{1}{1+e^{-a_L(v-b_L)}}\right)\\ &+p_H p_{L|H}\left(\frac{-a_H e^{-a_H(v-b_H)}}{\left(1+e^{-a_H(v-b_H)}\right)^2}(v-\mu_H) + 1 - \frac{1}{1+e^{-a_H(v-b_H)}}\right) = 0.\end{aligned} \tag{4.30}$$

Similarly, we also get

$$\begin{aligned}\frac{\partial \mathbb{E}(\Pi)}{\partial w} = 0 \iff& p_L(p_{H|L}\left(\frac{-a_L e^{-a_L(w-b_L)}}{\left(1+e^{-a_L(w-b_L)}\right)^2}(w-\mu_L) + 1 - \frac{1}{1+e^{-a_L(w-b_L)}}\right)\\ &+p_H(1-p_{L|H})\left(\frac{-a_H e^{-a_H(w-b_H)}}{\left(1+e^{-a_H(w-b_H)}\right)^2}(w-\mu_H) + 1 - \frac{1}{1+e^{-a_H(w-b_H)}}\right) = 0.\end{aligned} \tag{4.31}$$

These conditions characterize the optimum, which is then solved numerically.

Numerical illustrations

Let us look into the case of a sigmoid acceptance function (4.24) with parameters

$$a_L = a_H = 1, b_i = 2\mu_i,\ i \in \{L, H\}.$$

All other parameters remaining identical to those from Section 4.5, we obtain the following results:

- Full information:

$$P_L \approx 2.567203,$$
$$P_H \approx 8.926367,$$
$$\mathbb{E}(\Pi) \approx 0.8030561n.$$

- No differentiation:

$$P \approx 8.836827,$$
$$\mathbb{E}(\Pi) \approx 0.2998947n.$$

  Note that $P > \mu_H > \mu_L$. In this case, the insurer targets the high risk type audience because even if its size is smaller, with this acceptance function form he can make higher margins on them, thus compensating their smaller size.

- Differentiation:

$$P_L^* \approx 2.601853,$$
$$P_H^* \approx 8.91731,$$
$$\mathbb{E}(\Pi) \approx 0.6993114n - c(n).$$

  Note that $\mu_L < P_L^* < \mu_H < P_H^*$.

Consequently, differentiation is only preferable here if the population size satisfies

$$0.6993114n - c(n) \geq 0.2998947n, \quad \text{that is} \quad n/\ln(\gamma n) \geq 2.503651c_0,$$

in our case $n \geq 282.617$. Conversely, the maximum cost the insurer is willing to pay given a population size is given by:

$$c(n) < 0.3994167n.$$

For the variances, the results are as follows:

- Full information:

$$\mathbb{V}\text{ar}(\Pi) = 1.874096n.$$

- No differentiation:

$$\mathbb{V}\text{ar}(\Pi) = 1.089133n.$$

- Differentiation:

$$\mathbb{V}\text{ar}(\Pi) = 1.800876n.$$

Figure 4.12 depicts the form of the variance as function of the proposed premiums in the different scenarios and its decomposition into the two parts (the variance arising from the acceptance function and the one from the claim size variability), showing again a hump pattern. Figures 4.13-4.15 illustrate the sensitivity of the variances of the profit with and without differentiation, when varying one of the parameters. In Figure 4.13, we observe that under small price elasticity, with more customers entering the contract, the hump behaviour disappears, since at the limit there is no gap in different risk types behaviour. The total variance is now mostly due to the

underwriting process via the acceptance function, and claim size variance has little effect on the total variance. In contrast, a high price elasticity pushes different risk types to stabilize around their pivotal point of their respective acceptance function, accepting contracts only below this point, see Figure 4.14. All the variance of the profit can then be explained by the claim size variance. Finally, Figure 4.16 gives the mean-variance frontier in case of this sigmoid acceptance function.

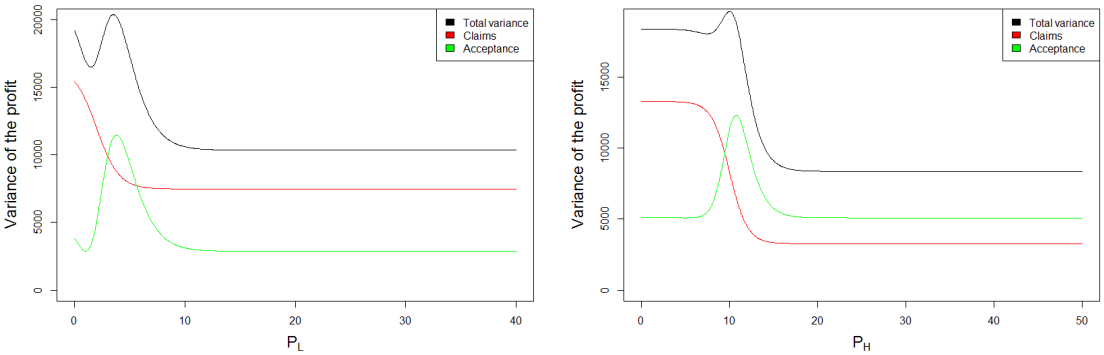

(a) Perfect information case. Left-hand side: $\mathbb{V}\mathrm{ar}(\Pi)$ as function of P_L. Right-hand side: $\mathbb{V}\mathrm{ar}(\Pi)$ as function of P_H.

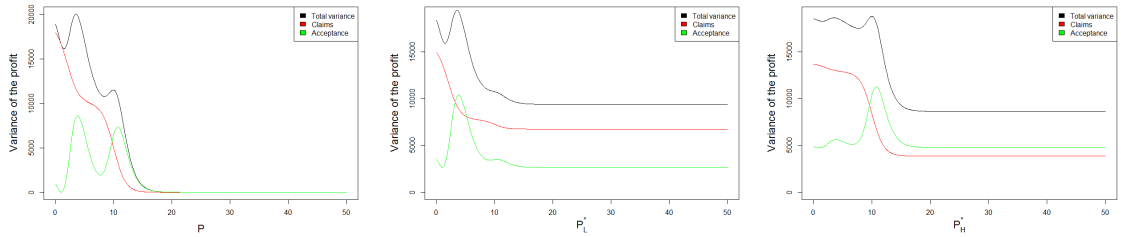

(b) No differentiation case. $\mathbb{V}\mathrm{ar}(\Pi)$ as function of P.

(c) Differentiation case. Left-hand side: $\mathbb{V}\mathrm{ar}(\Pi)$ as function of P_L^*. Right-hand side: $\mathbb{V}\mathrm{ar}(\Pi)$ as function of P_H^*.

Figure 4.12: Decomposition of the variance.

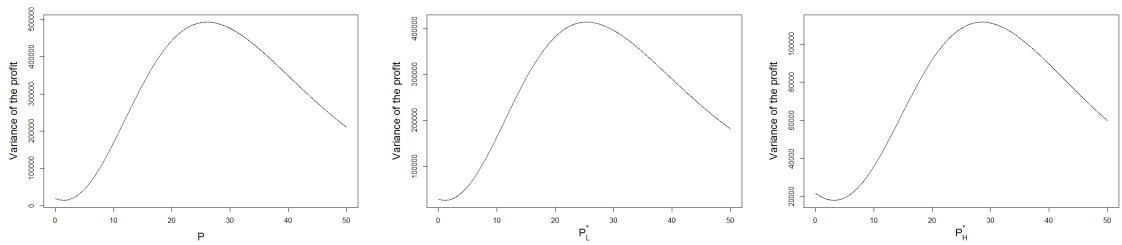

(a) No differentiation case.

(b) Differentiation case.

Figure 4.13: Parameter $a = 0.1$. Under small price-elasticity of demand, we observe higher levels of underwriting for both risk types, hence higher and smoother variance.

We observe that the no-differentiation strategy changes depending on the form of the acceptance function used in the analysis. This can be particularly relevant in the case when a company conducts a study using a simplified linear form instead of a more realistic logistic approximation.

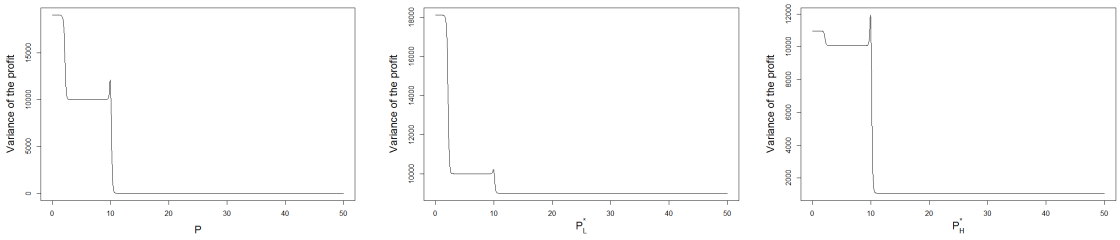

(a) No differentiation case. (b) Differentiation case.

Figure 4.14: Parameter $a = 10$. Under big price-elasticity of demand, demand concentrates around pivotal points, thus different risk types only enter contract until their pivotal point price, therefore steps are noticeable.

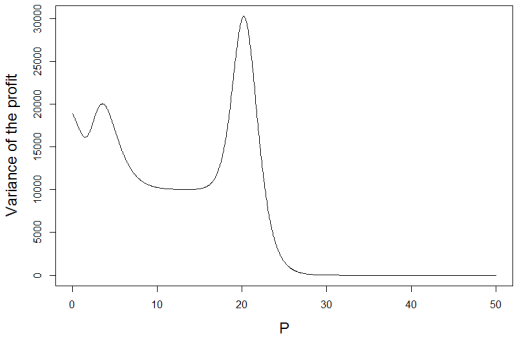

(a) No differentiation case.

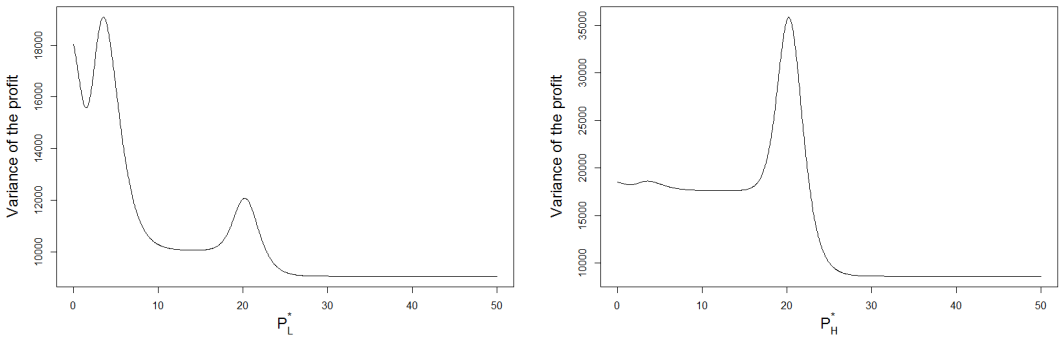

(b) Differentiation case.

Figure 4.15: Parameter $\mu_H = 10$. With a bigger expected claim size difference between risk types, the relationship between the humps and the risk types becomes clearer.

Sensitivities

Let us now investigate the sensitivity of the expected profit w.r.t. to each parameter. The change in expected profits will allow to compute the variation in the maximal cost for which the differentiation policy is still advantageous.

Parameter a: In the sigmoid curve of the acceptance function f, a represents the

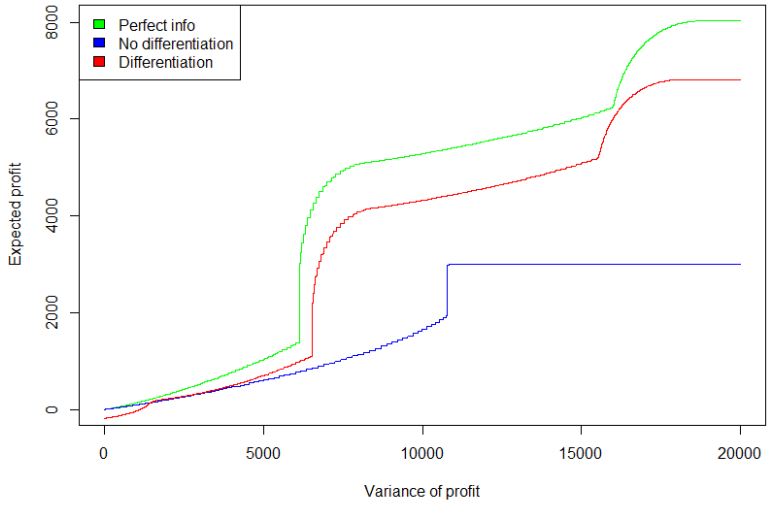

Figure 4.16: Efficient frontiers in the mean-variance setup for the three scenarios.

steepness, giving the speed at which the function changes around its central point (see Figure 4.11). We choose $a_L = a_H = a$, so we will vary it as a unique parameter. We observe in Figure 4.17 that for higher levels of a, the steepness of the twist increases, meaning that the values of the acceptance function grow closer to the points b_L and b_H. Thus, the price can be set closer to the twisting point, allowing a higher proportion of individuals to enter the contract. As in our initial parametrization $b_i > \mu_i$, we gain strictly positive profit when pricing around b_i. In the case of no differentiation, we cannot entirely benefit from this feature, as one of our types twisting point will end up far from the unique price P. Therefore, the maximum cost the insurer is willing to invest into the classification method is increasing in the parameter a. We can determine the limit:

$$\begin{aligned}
&\lim_{a\to+\infty} c(a) = \lim_{a\to+\infty} \mathbb{E}\left(\Pi(P_L^*(a), P_H^*(a))\right) - \lim_{a\to+\infty} \mathbb{E}\left(\Pi(P(a))\right) \\
&= \lim_{\{P_L^*\to b_L^-, P_H^*\to b_H^-\}} \mathbb{E}\left(\Pi(P_L^*(a), P_H^*(a))\right) - max\left(\lim_{P\to b_L^-} \mathbb{E}\left(\Pi(P(a))\right), \lim_{P\to b_H^-} \mathbb{E}\left(\Pi(P(a))\right)\right) \\
&= p_L(1-p_{H|L})(b_L-\mu_L)n + p_H(1-p_{L|H})(b_H-\mu_H)n + p_H p_{L|H}(b_L-\mu_H)n \\
&\quad - max\left(p_L(b_L-\mu_L)n + p_H(b_L-\mu_H)n, p_H(b_H-\mu_H)n\right),
\end{aligned}$$

which in our case gives $0.63n$.

Parameters b_L and b_H: Now we simultaneously vary the parameters b_H and b_L (the central points of the acceptance functions), see Figure 4.18. Naturally, the higher b_i, the higher will be the overall profit, as customers accept premiums until higher thresholds. Therefore, it becomes more and more attractive to differentiate customers to actually get this profit. Conversely, if b_H grows ceteris paribus, the profit increase becomes smaller with differentiation as the proportion of high-risk types is too low to strongly influence the non-differentiation premium.

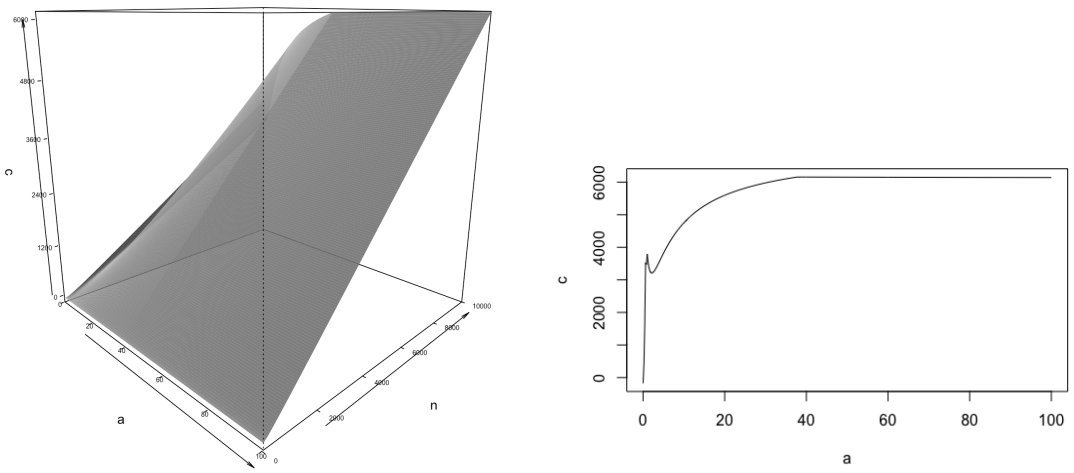

(a) Function of a and n.

(b) n fixed at 10000.

Figure 4.17: Maximum affordable investment cost for implementation of a differentiation mechanism as a function of a and n.

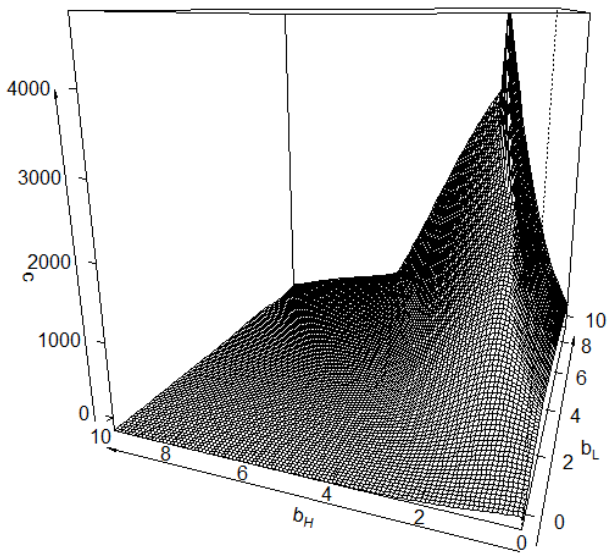

Figure 4.18: Maximum affordable cost as a function of b_H, b_L and $n = 10000$.

4.6.2 Other risk measures

We give a short comparative analysis for two other risk measures replacing the variance criterion (see Pflug and Römisch [115, Ch.5] for a more extensive list of possible alternatives in the context of efficient frontier studies in decision making).

Lower semi-variance

The main drawback of a variance risk constraint is that positive deviations from the mean are also penalized. In [101], Markowitz suggests the concept of the lower semi-variance (LSV)

$$\mathbb{V}\mathrm{ar}^{-}(\Pi) := \mathbb{E}\left((\min(0, \Pi - \mathbb{E}(\Pi))^2\right)$$

of the profit Π to account for asymmetry of positive and negative deviations from the profit target. In this case, analytical formulas are not feasible any more, but one can obtain similar results by Monte Carlo simulation, using 1000 simulation runs. For that purpose, rather than only specifying two moments, we need to take an assumption of the entire distribution of claim sizes. The left plot in Figure 4.19 depicts the resulting efficient frontiers for the three scenarios for an assumption of Gamma distributions for the individual claim sizes with an additional atom at 0 with probability 0.25 (parameters consistent with their first two moments from (4.23)) and all other parameters chosen as in (4.23). The right plot in Figure 4.19 shows

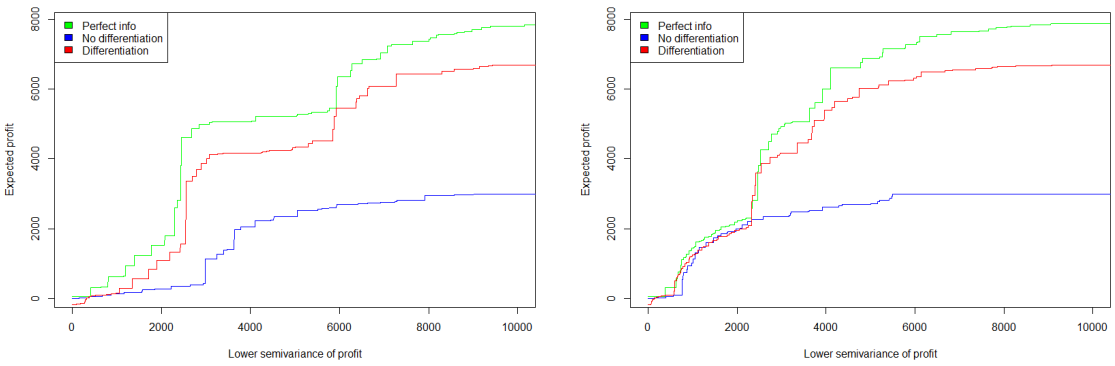

Figure 4.19: Efficient frontiers in the mean-LSV setup for the three scenarios under the assumption of Gamma-distributed H (left) and Log-Normal H (right).

the results for H being log-normally distributed risks (and again matching the first two moments). For instance, the intersection between the no-differentiation and differentiation scenario takes place at a much higher threshold.

Value-at-risk

Let us now instead consider the Value-at-risk

$$\mathrm{VaR}_\alpha(\Pi) := \inf\left\{x \in \mathbb{R} : F_\Pi(x) > \alpha\right\}$$

for some level $0 < \alpha < 1$. This measure is particularly focussing on the consequences of the tails of the loss (negative profit), when using small values of α. As for the LSV, we depict Monte Carlo results for the case of Gamma-distributed H and Log-normal H (Figure 4.20) risk types, where $\alpha = 0.025$. That is, the profit can be lower than the value of the abscissa in Figure 4.20 only with probability $\alpha = 0.025$, so that the more left in the abscissa one gets, the more risk-averse the strategy is. One observes that high values of $\mathrm{VaR}_{0.025}(\Pi)$ can only be obtained by the no-differentiation case. In regions where that VaR-value can be attained by all strategies, the differentiation

strategy always dominates the one without differentiation. Note that for this level of α, one virtually does not observe any difference between the case of light-tailed and heavy-tailed losses, which is also due to the size of the portfolio.

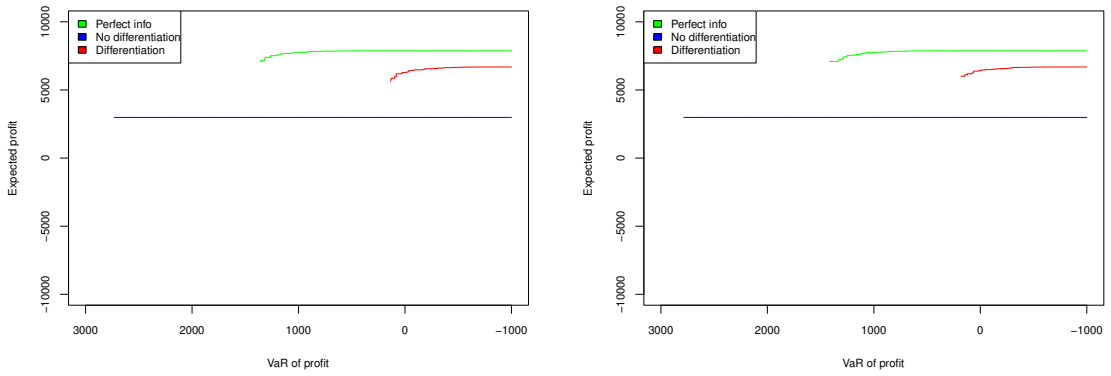

Figure 4.20: Efficient frontiers in the mean-VaR setup for the three scenarios for Gamma-distributed H (left) and Log-Normal H (right).

4.6.3 Utility functions

Utility theory is a classical tool to combine risk and profitability of an insurance undertaking in one function (see e.g. [123, 115]), so in this subsection we would like to briefly look at the problem posed in this chapter from the utility point of view. Note that in this case the knowledge of the full loss distribution is needed, and not only the first two moments as in Section 4.4. Assume that the insurer bases decisions on a risk-averse (i.e., increasing and concave) utility function $u(x)$. The insurer's optimization problem is then modified as follows:

$$\max_{P_L, P_H} \mathbb{E}\left(u(\Pi)\right), \tag{4.32}$$

where the profit Π is given by the (4.4), which we can also write as

$$\Pi = \sum_{j=1}^{N_L} \Pi_j^L + \sum_{j=1}^{N_H} \Pi_j^H. \tag{4.33}$$

Firstly, the moment-generating function of each Π_j^L is

$$M_{\Pi_j^L}(t) = \mathbb{E}\left(e^{t\Pi_j^L}\right) = \mathbb{E}\left(e^{t(P_L - L)}\right) = e^{tP_L} M_L(-t).$$

Analogously, $M_{\Pi_j^H}(t) = e^{tP_H} M_H(-t)$. By independence and classical collective risk theory calculations (cf. [83]), we can then determine the moment generating function of Π:

$$M_\Pi(t) = M_{N_L}(\log M_{\Pi_j^L}(t)) \cdot M_{N_H}(\log M_{\Pi_j^H}(t)).$$

The same reasoning applies to the non-differentiation case with setting $P_L = P_H = P$. Finally, for differentiating pricing, an analogous derivation gives

$$\begin{aligned} M_\Pi(t) =& M_{N_{L|L}}(\log M_{\Pi^{L|L}}(t)) \cdot M_{N_{H|L}}(\log M_{\Pi^{H|L}}(t)) \\ & \times M_{N_{H|H}}(\log M_{\Pi^{H|H}}(t)) \cdot M_{N_{L|H}}(\log M_{\Pi^{L|H}}(t)) e^{-tc(n)}. \end{aligned}$$

In each of the cases, $M_\Pi(t)$ can be inverted to obtain the c.d.f. $F_\Pi(x)$ of the profit, and the expected utility is then given by $\mathbb{E}\left(u(\Pi)\right) = \int_x u(\Pi(x))dF_\Pi(x)$.

For a numerical illustration, assume now that $L \sim \text{Exp}(\alpha_L)$ and $H \sim \Gamma(\alpha_H, \lambda_H)$. To be consistent with (4.23), we choose $\alpha_L = \mu_L = 1,\ \alpha_H = \mu_H^2/\sigma_H^2,\ \lambda_H = \mu_H/\sigma_H^2$. Since an explicit calculation of $\mathbb{E}\left(u(\Pi)\right)$ is not feasible, we add here numerical results from a Monte Carlo simulation, simulating its value for each choice of P_L, P_H (across a discrete grid of mesh size 0.05) using 1000 runs. For the sake of comparison, we use three popular utility functions:

- linear utility $u(x) = x$ (leading to simply the expected value of the profit);
- exponential utility $u(x) = -e^{-Ax}$ for some risk aversion coefficient $A > 0$;
- quadratic utility $u(x) = x - Bx^2,\ x \leq \frac{1}{2B}$.

The results in Figures 4.21, 4.22, 4.23 show the expected utility for each of the available premium combinations and each strategy for these three utility functions. Figure 4.21 serves as a reference point since it represents the simple expected profit as before. One observes that the optimal solution clearly depends on the chosen utility function. With the chosen parametrization of the exponential utility function, the difference in the expected utility between the differentiation and not differentiation case is less prominent than in the quadratic utility as the marginal utility of the quadratic function is greater in this region.

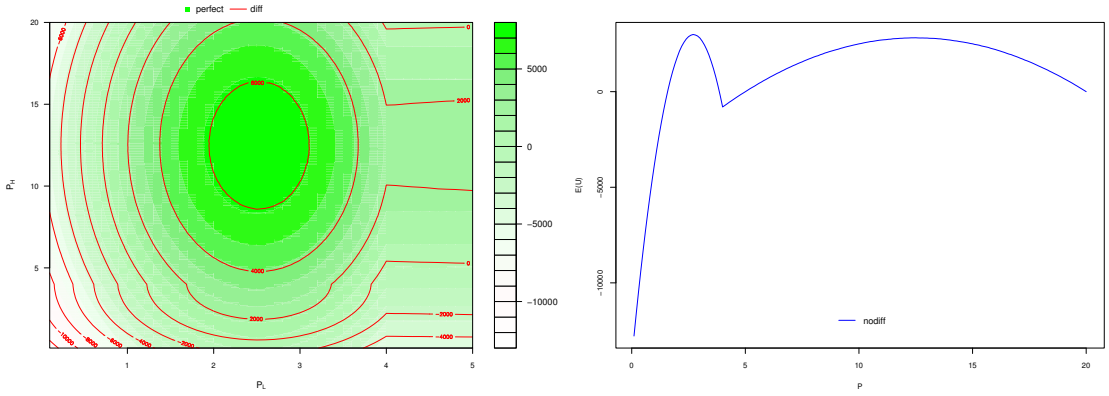

(a) Perfect information (green heatmap) and differentiation (red level curves).

(b) No differentiation.

Figure 4.21: Expected linear utility as a function of premiums.

4.7 Conclusion

In this chapter, we investigated the problem of risk categorization under the possibility of classification errors for an insurance company. We highlighted the impact of misspecification of risk classes on the company's profit, which is a relevant topic due to the growing use of black box techniques in classification. Resulting pricing errors may lead to adverse selection via a modified acceptance behaviour of individuals to

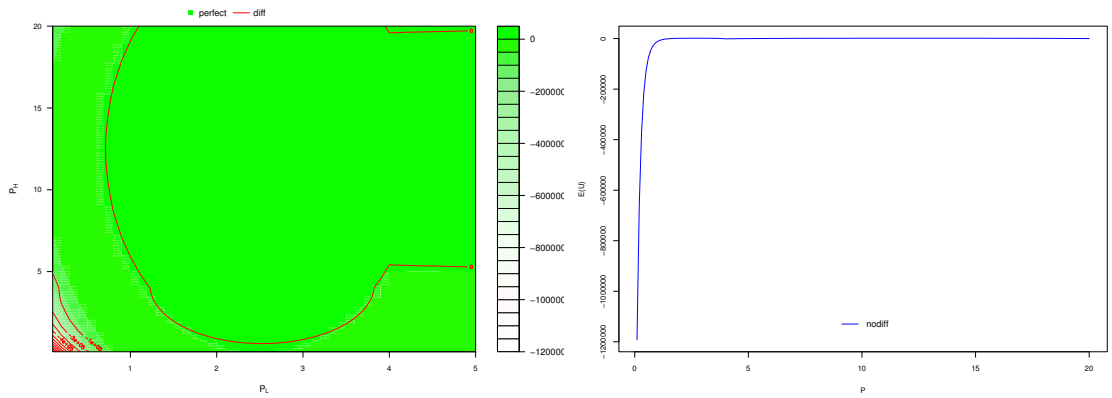

(a) Perfect information (green heatmap) and differentiation (red level curves).

(b) No differentiation.

Figure 4.22: Expected exponential utility as a function of premiums ($A = 0.0005$).

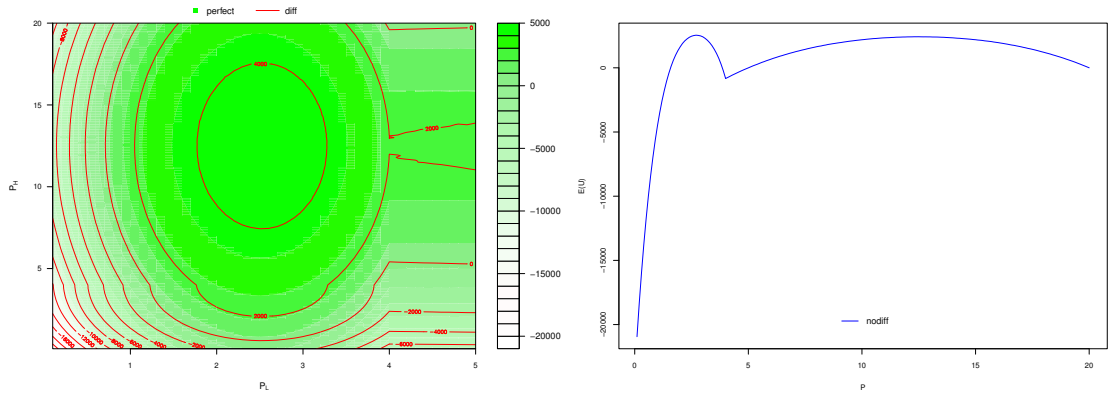

(a) Perfect information (green heatmap) and differentiation (red level curves).

(b) No differentiation.

Figure 4.23: Expected quadratic utility as a function of premiums ($B = 0.00005$).

enter a contract, potentially leading to extra costs due to a loss of premium inflow and the loss of the share of customers underwriting insurance. In a simple model with two risk types and piece-wise linear acceptance function, we distinguished three pricing scenarios: full information, undifferentiated pricing and costly price differentiation under error assumptions. In this framework, we studied the optimal solution for simply maximizing expected profit and more generally within a mean-variance framework, establishing efficient frontiers for the premium choices. The cost of the risk categorization as a function of population size will then eventually determine the optimal choice of premiums, and to what extent risk classification is profitable.

The simplicity of the introduced model allowed to quantify the effects and consequences of misspecification on the insurer's profit. Clearly, it will be of interest in future research to generalize the model assumptions in various directions. Beyond the extensions to more general acceptance functions and risk measures that we already address to a first extent in Section 4.6 of the chapter, it will be of interest to extend the study to more than two risk categories. Another important direction

will be to introduce market competition into this model (cf. [58, 108]), as well as the lapse behavior of policyholders between the different market players (see e.g. [20, 106]). For example, one could extend the setting to include a second competing insurer in a Stackelberg competition setting [92, 147], where we distinguish the "leader" who undertakes actions first and the "follower" who moves second. It is important for the leader to anticipate the followers moves in order to take them in account in the optimization problem. The first insurer obtains a market share that is further targeted by the second insurer. In [58], the authors define lapse functions for the customers to switch from the current insurer to another as a function of price-sensitivity. Also, customers incur some fixed cost for switching, so they need to observe a sufficiently high price difference to be motivated to change the insurer. This customer behaviour can be explained by administrative fees to terminate a contract, cost efforts to search a new insurer, but also by the customer's attachment to his first choice. These factors have been studied extensively in the literature, such as [80, 33, 131, 127] for consumer behaviour outside the insurance context. In [44], the authors consider the difference between the average market premium and the proposed insurer's premium as the driving factor of its market share. In [76], the author uses empirical data to quantify the searching cost and the switching cost for a customer in the US auto insurance market. This can be applied to our model by supposing costs are directly related to the underlying premium or claim distribution, so we can define an overall switching cost and apply it to the insurer as the cost of attracting customers, and to the customer as the minimal price difference for which they are willing to change insurers. The natural extension to this competition approach would be to allow for $n > 2$ insurers and thus reformulating the problem in a matrix form.

As we have explained above, a time dynamic setting is a natural extension for competition games. In the same spirit, one can introduce dynamism in the process of risk classification, as with every new period the insurer gains knowledge about the customer's risk type. This can allow him to refine his classification. Similar systems are used in bonus-malus frameworks. In this way, the insurer aims to compute an individual customer value.

Finally, while our probabilistic acceptance model already covers a certain degree of randomness in the choice of insurance policies, it could be interesting to more explicitly include bounded rationality as well as other elements of inertia of policyholders and the markets in the modelling framework.

References

[1] Kenneth S Abraham. Efficiency and fairness in insurance risk classification. *Virginia Law Review*, pages 403–451, 1985.

[2] Saif Al-Kuwari, James H. Davenport, and Russell J. Bradford. Cryptographic hash functions: Recent design trends and security notions. Cryptology ePrint Archive, Report 2011/565, 2011. `https://ia.cr/2011/565`.

[3] Hansjörg Albrecher. Asymmetric information and insurance. In *Cahiers de l'Institut Louis Bachelier*, volume 20, pages 12–15, 2016.

[4] Hansjörg Albrecher, Andrei Badescu, and David Landriault. On the dual risk model with tax payments. *Insurance: Mathematics and Economics*, 42(3):1086–1094, 2008.

[5] Hansjörg Albrecher, Antoine Bommier, Damir Filipović, Pablo Koch-Medina, Stéphane Loisel, and Hato Schmeiser. Insurance: models, digitalization, and data science. *European Actuarial Journal*, 9:349–360, 2019.

[6] Hansjörg Albrecher and Daily-Amir Dalit. On effects of asymmetric information on non-life insurance prices under competition. *International Journal of Data Analysis Techniques and Strategies*, 9(4):287–299, 2017.

[7] Hansjörg Albrecher, Dina Finger, and Pierre-O Goffard. Empirical risk analysis of mining a proof-of-work blockchain.

[8] Hansjörg Albrecher, Dina Finger, and Pierre-O Goffard. Blockchain mining in pools: Analyzing the trade-off between profitability and ruin. *Insurance: Mathematics and Economics*, 105:313–335, 2022.

[9] Hansjörg Albrecher, Hans U Gerber, and Hailiang Yang. A direct approach to the discounted penalty function. *North American Actuarial Journal*, 14(4):420–434, 2010.

[10] Hansjörg Albrecher and Pierre-Olivier Goffard. On the profitability of selfish blockchain mining under consideration of ruin. *Operations Research*, 70(1):179–200, 2022.

[11] Renata G. Alcoforado, Agnieszka I. Bergel, Rui M. R. Cardoso, Alfredo D. Egídio dos Reis, and Eugenio V. Rodríguez-Martínez. Ruin and dividend measures in the renewal dual risk model. *Methodology and Computing in Applied Probability*, jul 2021.

[12] João Almeida, Shravan Tata, Andreas Moser, and Vikko Smit. Bitcoin prediciton using ANN. *Neural Networks*, 7:1–12, 2015.

[13] E Sparre Andersen. On the collective theory of risk in case of contagion between claims. *Bulletin of the Institute of Mathematics and its Applications*, 12(2):275–279, 1957.

[14] Katrien Antonio and Emiliano A Valdez. Statistical concepts of a priori and a posteriori risk classification in insurance. *AStA Advances in Statistical Analysis*, 96:187–224, 2012.

[15] Andreas Antonopoulos. *Mastering Bitcoin*. O'Reilly UK Ltd., July 2017.

[16] Søren Asmussen and Hansjörg Albrecher. *Ruin probabilities*. World Scientific, Singapore, 2010.

[17] Mohsen Attaran and Angappa Gunasekaran. *Applications of blockchain technology in business: challenges and opportunities*. Springer Nature, 2019.

[18] Benjamin Avanzi, Hans U. Gerber, and Elias S.W. Shiu. Optimal dividends in the dual model. *Insurance: Mathematics and Economics*, 41(1):111–123, jul 2007.

[19] Amin Azari. Bitcoin price prediction: An ARIMA approach. *arXiv preprint arXiv:1904.05315*, 2019.

[20] Flavia Barsotti, Xavier Milhaud, and Yahia Salhi. Lapse risk in life insurance: Correlation and contagion effects among policyholders' behaviors. *Insurance: Mathematics and Economics*, 71:317–331, 2016.

[21] Eric Bauer and Ron Kohavi. An empirical comparison of voting classification algorithms: Bagging, boosting, and variants. *Machine learning*, 36:105–139, 1999.

[22] Nigel G. Bean, Mark Fackrell, and Peter Taylor. Characterization of matrix-exponential distributions. *Stochastic Models*, 24(3):339–363, aug 2008.

[23] Mogens Bladt and Bo Friis Nielsen. *Matrix-Exponential Distributions in Applied Probability*. Springer-Verlag GmbH, May 2017.

[24] Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986.

[25] Robert F. Botta, Carl M. Harris, and William G. Marchal. Characterizations of generalized hyperexponential distribution functions. *Communications in Statistics. Stochastic Models*, 3(1):115–148, jan 1987.

[26] Jamal Bouoiyour, Refk Selmi, Aviral Kumar Tiwari, and Olaolu Richard Olayeni. What drives Bitcoin price. *Economics Bulletin*, 36(2):843–850, 2016.

[27] Rhys Bowden, H. Paul Keeler, Anthony E. Krzesinski, and Peter G. Taylor. Modeling and analysis of block arrival times in the Bitcoin blockchain. *Stochastic Models*, 36(4):602–637, 2020.

[28] Patrick L Brockett, Linda L Golden, Montserrat Guillen, Jens Perch Nielsen, Jan Parner, and Ana Maria Perez-Marin. Survival analysis of a household portfolio of insurance policies: how much time do you have to stop total customer defection? *Journal of Risk and Insurance*, 75(3):713–737, 2008.

[29] Lars Brünjes and Murdoch J Gabbay. Utxo-vs account-based smart contract blockchain programming paradigms. In *International Symposium on Leveraging Applications of Formal Methods*, pages 73–88. Springer, 2020.

[30] Hans Bühlmann. Mathematical methods in risk theory. *Die Grundlehren der Mathematischen Wissenschaften in Einzeldarstellungen, Springer Verlag*, 1970.

[31] Miles Carlsten, Harry Kalodner, S Matthew Weinberg, and Arvind Narayanan. On the instability of bitcoin without the block reward. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 154–167, 2016.

[32] Arthur Charpentier. *Insurance, Biases, Discrimination and Fairness.* Springer, Heidelberg, 2023.

[33] Mei-Fang Chen and Ling-Huei Wang. The moderating role of switching barriers on customer loyalty in the life insurance industry. *The Service Industries Journal*, 29(8):1105–1123, 2009.

[34] Sung Nok Chiu, Dietrich Stoyan, Wilfrid S Kendall, and Joseph Mecke. *Stochastic geometry and its applications.* John Wiley & Sons, 2013.

[35] Jung-Moon Choi, Ji-Hyeok Kim, and Sung-Jun Kim. Application of reinforcement learning in detecting fraudulent insurance claims. *International Journal of Computer Science & Network Security*, 21(9):125–131, 2021.

[36] Pavel Ciaian, Miroslava Rajcaniova, and d'Artis Kancs. The economics of BitCoin price formation. *Applied Economics*, 48(19):1799–1815, 2016.

[37] Lin William Cong, Zhiguo He, and Jiasun Li. Decentralized Mining in Centralized Pools. *The Review of Financial Studies*, 34(3):1191–1235, 04 2020.

[38] Robert M Corless, Gaston H Gonnet, David EG Hare, David J Jeffrey, and Donald E Knuth. On the Lambert W function. *Advances in Computational mathematics*, 5:329–359, 1996.

[39] David Roxbee Cox. A use of complex probabilities in the theory of stochastic processes. *Mathematical Proceedings of the Cambridge Philosophical Society*, 51(2):313–319, apr 1955.

[40] Harald Cramér. *On the mathematical theory of risk.* Skandia Jubilee Volume, Stockholm, 1930.

[41] Harald Cramér. *Collective risk theory.* Skandia ins. Company, 1955.

[42] Keith J Crocker and Arthur Snow. The efficiency effects of categorical discrimination in the insurance industry. *Journal of Political Economy*, 94(2):321–344, 1986.

[43] Keith J Crocker and Arthur Snow. The theory of risk classification. *Handbook of insurance*, pages 281–313, 2013.

[44] Dalit Daily-Amir, Hansjörg Albrecher, Martin Bladt, and Joël Wagner. On market share drivers in the swiss mandatory health insurance sector. *Risks*, 7(4):114, 2019.

[45] Alex De Vries, Ulrich Gallersdörfer, Lena Klaaßen, and Christian Stoll. Revisiting Bitcoin's carbon footprint. *Joule*, 6(3):498–502, 2022.

[46] Alex De Vries and Christian Stoll. Bitcoin's growing e-waste problem. *Resources, Conservation and Recycling*, 175:105901, 2021.

[47] Florent Etienne De Vylder. *Advanced risk theory: a self-contained introduction*. Éd. de l'Univ. de Bruxelles, 1996.

[48] Michel Denuit, Donatien Hainaut, and Julien Trufin. *Effective Statistical Learning Methods for Actuaries I: GLMs and Extensions*. Springer Nature, 2019.

[49] Michel Denuit, Donatien Hainaut, and Julien Trufin. *Effective Statistical Learning Methods for Actuaries III: Neural Networks and Extensions*. Springer Actuarial Lecture Notes, 2019.

[50] Thibaud Deruelle, Veronika Kalouguina, Philipp Trein, and Joël Wagner. Designing privacy in personalized health: an empirical analysis. *Big Data & Society*, 10(1):20539517231158636, 2023.

[51] Arnold A. Dicke, Sam Gutterman, Burton Jay, and Mark Litow. On risk classification. *A public policy Monograph, American Academy of Actuaries, New York, NY*, 2011.

[52] David CM Dickson. *Insurance risk and ruin*. Cambridge University Press, 2016.

[53] Georges Dionne, Nathalie Fombaron, and Neil Doherty. *Adverse selection in insurance contracting*. Springer, 2013.

[54] Georges Dionne and Casey Rothschild. Economic effects of risk classification bans. *The Geneva Risk and Insurance Review*, 39:184–221, 2014.

[55] Daniel Dufresne. On a general class of risk models. *Australian Actuarial Journal*, 7:755–791, 2001.

[56] Daniel Dufresne. Fitting combinations of exponentials to probability distributions. *Applied Stochastic Models in Business and Industry*, 23(1):23–48, 2007.

[57] Christophe Dutang. The customer, the insurer and the market. *Bulletin Français d'Actuariat*, 2012.

[58] Christophe Dutang, Hansjoerg Albrecher, and Stéphane Loisel. Competition among non-life insurers under solvency constraints: A game-theoretic approach. *European Journal of Operational Research*, 231(3):702–711, 2013.

[59] Pankaj Dutta, Tsan-Ming Choi, Surabhi Somani, and Richa Butala. Blockchain technology in supply chain operations: Applications, challenges and research opportunities. *Transportation research part e: Logistics and transportation review*, 142:102067, 2020.

[60] David Easley, Maureen O'Hara, and Soumya Basu. From mining to markets: The evolution of bitcoin transaction fees. *Journal of Financial Economics*, 134(1):91–109, 2019.

[61] Mark Fackrell. Fitting with matrix-exponential distributions. *Stochastic Models*, 21(2-3):377–400, jan 2005.

[62] Mark Fackrell. An alternative characterization for matrix exponential distributions. *Advances in Applied Probability*, 41(4):1005–1022, dec 2009.

[63] Robert J Finger. Chapter 6 risk classification. *Foundations of Casualty Actuarial Science*, pages 287–341, 2001.

[64] Ronald A Fisher. The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7(2):179–188, 1936.

[65] Christian Francq and Jean-Michel Zakoian. *GARCH models: structure, statistical inference and financial applications*. John Wiley & Sons, 2019.

[66] Nadine Gatzert, Gudrun Schmitt-Hoermann, and Hato Schmeiser. Optimal risk classification with an application to substandard annuities. *North American Actuarial Journal*, 16(4):462–486, 2012.

[67] Hans U Gerber and Nathaniel Smith. Optimal dividends with incomplete information in the dual model. *Insurance: Mathematics and Economics*, 43(2):227–233, 2008.

[68] Pierre-Olivier Goffard. Fraud risk assessment within blockchain transactions. *Advances in Applied Probability*, 51(2):443–467, 2019.

[69] Jan Grandell. *Aspects of risk theory*. Springer Science & Business Media, 2012.

[70] Montserrat Guillén, Jan Parner, Chresten Densgsoe, and Ana M Perez-Marin. Using logistic regression models to predict and understand why customers leave an insurance company. In *Intelligent and other computational techniques in insurance: Theory and applications*, pages 465–490. World Scientific, 2003.

[71] Montserrat Guillen, Ana Maria Perez, and Manuela Alcañiz. A logistic regression approach to estimating customer profit loss due to lapses in insurance. *Document de Treball No. XREAP*, 13, 2011.

[72] Amulya Gurtu and Jestin Johny. Potential of blockchain technology in supply chain management: a literature review. *International Journal of Physical Distribution & Logistics Management*, 49(9):881–900, 2019.

[73] Bernard Hanzon and Finbarr Holland. Non-negativity analysis for exponential-polynomial-trigonometric functions on $[0,\infty)$. In *Spectral Theory, Mathematical System Theory, Evolution Equations, Differential and Difference Equations*, pages 399–412. Springer Basel, 2012.

[74] Thomas Hardjono, Alexander Lipton, and Alex Pentland. A contract service provider model for virtual assets. *The Journal of FinTech*, 1(02):2150004, 2021.

[75] Trevor Hastie, Robert Tibshirani, and Jerome H Friedman. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer, 2009.

[76] Elisabeth Honka. Quantifying search and switching costs in the us auto insurance industry. *The RAND Journal of Economics*, 45(4):847–884, 2014.

[77] Michael Hoy. Categorizing risks in the insurance industry. *The Quarterly Journal of Economics*, 97(2):321–336, 1982.

[78] Rob J Hyndman and George Athanasopoulos. *Forecasting: principles and practice*. OTexts, 2018.

[79] Rob J Hyndman and Yeasmin Khandakar. Automatic time series forecasting: the forecast package for R. *Journal of statistical software*, 27(1):1–22, 2008.

[80] Mark Israel. Tenure dependence in consumer-firm relationships: an empirical analysis of consumer departures from automobile insurance firms. *RAND Journal of Economics*, pages 165–192, 2005.

[81] Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. *An introduction to statistical learning*, volume 112. Springer, 2013.

[82] Abdul J Jerri. *Linear difference equations with discrete transform methods*, volume 363. Springer Science & Business Media, 2013.

[83] Rob Kaas, Marc Goovaerts, Jan Dhaene, and Michel Denuit. *Modern Actuarial Risk Theory*. 2008.

[84] Veronika Kalouguina and Joël Wagner. Challenges and solutions for integrating and financing personalized medicine in healthcare systems: A systematic literature review. *Journal of Risk and Financial Management*, 13(11):283, 2020.

[85] Shoji Kasahara and Jun Kawahara. Effect of bitcoin fee on transaction-confirmation process. *Journal of Industrial & Management Optimization*, 15(1):365–386, 2019.

[86] Paraskevi Katsiampa. Volatility estimation for bitcoin: A comparison of GARCH models. *Economics Letters*, 158:3–6, 2017.

[87] Mohammed Khalilia, Sounak Chakraborty, and Mihail Popescu. Predicting disease risks from highly imbalanced data using random forest. *BMC medical informatics and decision making*, 11:1–13, 2011.

[88] J.F.C. Kingman. *Poisson Processes*. Oxford Studies in Probability. Clarendon Press, 1992.

[89] Doron Kliger and Benny Levikson. Pricing insurance contracts—an economic viewpoint. *Insurance: Mathematics and Economics*, 22(3):243–249, 1998.

[90] Nikolay Kyurkchiev and Svetoslav Markov. Sigmoid functions: some approximation and modelling aspects. *LAP LAMBERT Academic Publishing, Saarbrucken*, 4, 2015.

[91] Chantal Labbé and Kristina P Sendova. The expected discounted penalty function under a risk model with stochastic income. *Applied Mathematics and Computation*, 215(5):1852–1867, 2009.

[92] Jean-Jacques Laffont and David Martimort. The theory of incentives: the principal-agent model. In *The theory of incentives*. Princeton university press, 2009.

[93] Jean Lemaire. *Automobile insurance: actuarial models*, volume 4. Springer Science & Business Media, 1985.

[94] Shuanming Li and Jose Garrido. On ruin for the erlang (n) risk process. *Insurance: Mathematics and Economics*, 34(3):391–408, 2004.

[95] Zongxi Li, A Max Reppen, and Ronnie Sircar. A mean field games model for cryptocurrency mining. *Management Science*, 2023.

[96] X.Sheldon Lin and Gordon E. Willmot. Analysis of a defective renewal equation arising in ruin theory. *Insurance: Mathematics and Economics*, 25(1):63–84, 1999.

[97] X.Sheldon Lin and Gordon E. Willmot. The moments of the time of ruin, the surplus before ruin, and the deficit at ruin. *Insurance: Mathematics and Economics*, 27(1):19–44, 2000.

[98] Mathias Lindholm, Ronald Richman, Andreas Tsanakas, and Mario V Wüthrich. Discrimination-free insurance pricing. *ASTIN Bulletin: The Journal of the IAA*, 52(1):55–89, 2022.

[99] Yi Lu. "A Direct Approach to the Discounted Penalty Function", Hansjörg Albrecher, Hans U. Gerber, and Hailiang Yang, Volume 14, No. 4, 2010. *North American Actuarial Journal*, 14(4):438–441, oct 2010.

[100] Fillip Lundberg. *Approximerad framställning av sannolikhetsfunktionen. Aterförsäkring av kollektivrisker. Akad.* PhD thesis, Almqvist och Wiksell Uppsala, Sweden, 1903.

[101] Harry M Markovitz. *Portfolio selection: Efficient diversification of investments*. John Wiley, 1959.

[102] Harry Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.

[103] Christian Mazza and Didier Rullière. A link between wave governed random motions and ruin processes. *Insurance: Mathematics and Economics*, 35(2):205–222, oct 2004.

[104] Sean McNally, Jason Roche, and Simon Caton. Predicting the price of bitcoin using machine learning. In *2018 26th euromicro international conference on parallel, distributed and network-based processing (PDP)*, pages 339–343. IEEE, 2018.

[105] Johnnatan Messias, Mohamed Alzayat, Balakrishnan Chandrasekaran, Krishna P Gummadi, Patrick Loiseau, and Alan Mislove. Selfish & opaque transaction ordering in the Bitcoin blockchain: the case for chain neutrality. In *Proceedings of the 21st ACM Internet Measurement Conference*, pages 320–335, 2021.

[106] Xavier Milhaud and Christophe Dutang. Lapse tables for lapse risk management in insurance: a competing risk approach. *European Actuarial Journal*, 8:97–126, 2018.

[107] Malte Möser and Rainer Böhme. Trends, tips, tolls: A longitudinal study of Bitcoin transaction fees. In *International Conference on Financial Cryptography and Data Security*, pages 19–33. Springer, 2015.

[108] Claire Mouminoux, Christophe Dutang, Stéphane Loisel, and Hansjoerg Albrecher. On a Markovian game model for competitive insurance pricing. *Methodology and Computing in Applied Probability*, 24(2):1061–1091, 2021.

[109] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Decentralized business review*, 2008.

[110] Arvind Narayanan, Joseph Bonneau, Edward Felten, Andrew Miller, and Steven Goldfeder. *Bitcoin and cryptocurrency technologies: a comprehensive introduction*. Princeton University Press, 2016.

[111] Andrew CY Ng. On a dual model with a dividend threshold. *Insurance: Mathematics and Economics*, 44(2):315–324, 2009.

[112] European Court of Justice. Verdict of March 1, 2011. *File Reference C-236/09*, 2011.

[113] Walter Olbricht. Tree-based methods: a useful tool for life insurance. *European Actuarial Journal*, 2:129–147, 2012.

[114] Wouter Penard and Tim van Werkhoven. On the secure hash algorithm family. *Cryptography in context*, pages 1–18, 2008.

[115] Georg Ch Pflug and Werner Romisch. *Modeling, measuring and managing risk*. World Scientific, 2007.

[116] Donatella Porrini. Risk classification efficiency and the insurance market regulation. *Risks*, 3(4):445–454, 2015.

[117] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2023.

[118] Ronald Richman. AI in actuarial science–a review of recent advances–part 1. *Annals of Actuarial Science*, 15(2):207–229, 2021.

[119] Ronald Richman. AI in actuarial science–a review of recent advances–part 2. *Annals of Actuarial Science*, 15(2):230–258, 2021.

[120] Tomasz Rolski, Hanspeter Schmidli, Volker Schmidt, and Jozef L Teugels. *Stochastic processes for insurance and finance*. John Wiley & Sons, 2009.

[121] Meni Rosenfeld. Analysis of bitcoin pooled mining reward systems. *CoRR*, abs/1112.4980, 2011.

[122] Sabina Rossi, Ivan Malakhov, and Andrea Marin. Analysis of the confirmation time in proof-of-work blockchains. *Available at SSRN 4031244*.

[123] Michael Rothschild and Joseph Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. In *Uncertainty in economics*, pages 257–280. Elsevier, 1978.

[124] Jörn Sass and Frank Thomas Seifried. Insurance markets and unisex tariffs: is the European Court of Justice improving or destroying welfare? *Scandinavian Actuarial Journal*, 2014(3):228–254, 2014.

[125] Hato Schmeiser, Tina Störmer, and Joël Wagner. Unisex insurance pricing: Consumers' perception and market implications. *The Geneva Papers on Risk and Insurance-Issues and Practice*, 39:322–350, 2014.

[126] Okke Schrijvers, Joseph Bonneau, Dan Boneh, and Tim Roughgarden. Incentive compatibility of Bitcoin mining pool reward functions. In *Financial Cryptography and Data Security*, pages 477–498. Springer Berlin Heidelberg, 2017.

[127] Stephan Seiler. The impact of search costs on consumer behavior: A dynamic approach. *Quantitative Marketing and Economics*, 11:155–203, 2013.

[128] Peng Shi and Kun Shi. Non-life insurance risk classification using categorical embedding. *North American Actuarial Journal*, 27(3):579–601, 2023.

[129] Robert H Shumway and David S Stoffer. *Time series analysis and its applications*, volume 3. Springer, 2000.

[130] slush pool. Reward system specifications. `https://help.slushpool.com/en/support/solutions/articles/77000426280-reward-system-specification`, 2021.

[131] Bruce A Strombom, Thomas C Buchmueller, and Paul J Feldstein. Switching costs, price sensitivity and health plan choice. *Journal of Health economics*, 21(1):89–116, 2002.

[132] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[133] Rick Swedloff. Risk classification's big data (r) evolution. *Conn. Ins. LJ*, 21:339, 2014.

[134] Béla Szőkefalvi-Nagy. *Introduction to real functions and orthogonal expansions*. Oxford University Press, 1965.

[135] Luxor Tech. Different Bitcoin mining pool payment methods PPS vs FPPS vs PPLNS vs PPS+), 2018. Accessed: 2022-06-17.

[136] Enrico Tedeschi, Tor-Arne S Nordmo, Dag Johansen, and Håvard D Johansen. On optimizing transaction fees in Bitcoin using AI: Investigation on miners inclusion pattern. *ACM Transactions on Internet Technology (TOIT)*, 2022.

[137] Sergios Theodoridis. *Machine Learning: A Bayesian and Optimization Perspective*. Academic Press, 2020.

[138] R Guy Thomas. Some novel perspectives on risk classification. *The Geneva Papers on Risk and Insurance-Issues and Practice*, 32:105–132, 2007.

[139] R Guy Thomas. Loss coverage as a public policy objective for risk classification schemes. *Journal of Risk and Insurance*, 75(4):997–1018, 2008.

[140] Olof Thorin. On the asymptotic behavior of the ruin probability for an infinite period when the epochs of claims form a renewal process. *Scandinavian Actuarial Journal*, 1974(2):81–99, 1974.

[141] Olof Thorin. Probabilities of ruin. *Scandinavian Actuarial Journal*, 1982(2):65–103, 1982.

[142] Henk C Tijms. *A first course in stochastic models*. John Wiley and sons, 2003.

[143] George Tzougas and Konstantin Kutzkov. Enhancing logistic regression using neural networks for classification in actuarial learning. *Algorithms*, 16(2):99, 2023.

[144] Hal R Varian. *Intermediate microeconomics with calculus: a modern approach*. WW norton & company, 2014.

[145] Francis Vekeman, Jesus Eric Piña-Garza, Wendy Y Cheng, Edward Tuttle, Philippe Giguère-Duval, Arman Oganisian, Joseph Damron, Mei Sheng Duh, Vivienne Shen, Timothy B Saurer, et al. Development of a classifier to identify patients with probable Lennox–Gastaut syndrome in health insurance claims databases via random forest methodology. *Current medical research and opinion*, 35(8):1415–1420, 2019.

[146] Roel Verbelen, Katrien Antonio, and Gerda Claeskens. Unravelling the predictive power of telematics data in car insurance pricing. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 67(5):1275–1304, 2018.

[147] Heinrich Von Stackelberg. *Market structure and equilibrium*. Springer Science & Business Media, 2010.

[148] Larry Wasserman. *All of nonparametric statistics*. Springer Science & Business Media, 2006.

[149] Wiltrud Weidner, Fabian WG Transchel, and Robert Weidner. Classification of scale-sensitive telematic observables for riskindividual pricing. *European Actuarial Journal*, 6:3–24, 2016.

[150] Gary M Weiss and Foster Provost. The effect of class distribution on classifier learning: an empirical study. Technical report, Rutgers University, 2001.

[151] Moritz Wendl, My Hanh Doan, and Remmer Sassen. The environmental impact of cryptocurrencies using proof of work and proof of stake consensus algorithms: A systematic review. *Journal of Environmental Management*, 326:116530, 2023.

[152] Gordon E Willmot. A laplace transform representation in a class of renewal queueing and risk processes. *Journal of Applied Probability*, 36(2):570–584, 1999.

[153] I Made Wirawan, Triyanna Widiyaningtyas, and Muchammad Maulana Hasan. Short term prediction on bitcoin price using ARIMA method. In *2019 International Seminar on Application for Technology of Information and Communication (iSemantic)*. IEEE, sep 2019.

[154] Mario V Wüthrich and Michael Merz. *Statistical foundations of actuarial learning and its applications*. Springer Nature, 2023.

[155] Ying Xiang. Using ARIMA-GARCH model to analyze fluctuation law of international oil price. *Mathematical Problems in Engineering*, 2022, 2022.

[156] Chen Yang and Kristina P Sendova. The ruin time under the Sparre-Andersen dual model. *Insurance: Mathematics and Economics*, 54:28–40, 2014.

[157] Aragon Yves. *Séries temporelles avec R: Méthodes et cas*. Springer-Verlag France–Pratique R–2011, 2011.

[158] Angela Zeiler, Rupert Faltermeier, Ingo R Keck, Ana Maria Tomé, Carlos García Puntonet, and Elmar Wolfgang Lang. Empirical mode decomposition-an introduction. In *The 2010 international joint conference on neural networks (IJCNN)*, pages 1–8. IEEE, 2010.

[159] Shijie Zhang and Jong-Hyouk Lee. Analysis of the main consensus protocols of blockchain. *ICT express*, 6(2):93–97, 2020.

[160] Zhimin Zhang, Hu Yang, and Shuanming Li. The perturbed compound Poisson risk model with two-sided jumps. *Journal of Computational and Applied Mathematics*, 233:1773–1784, 2010.

[161] Fei Zheng and Geoffrey I Webb. A comparative study of semi-naive Bayes methods in classification learning. *AUSDM05*, pages 141–155, 2005.

[162] Zibin Zheng, Shaoan Xie, Hong-Ning Dai, Weili Chen, Xiangping Chen, Jian Weng, and Muhammad Imran. An overview on smart contracts: Challenges, advances and platforms. *Future Generation Computer Systems*, 105:475–491, 2020.

[163] Zibin Zheng, Shaoan Xie, Hong-Ning Dai, Xiangping Chen, and Huaimin Wang. Blockchain challenges and opportunities: A survey. *International journal of Web and Grid services*, 14(4):352–375, 2018.

[164] Saide Zhu, Wei Li, Hong Li, Chunqiang Hu, and Zhipeng Cai. A survey: Reward distribution mechanisms and withholding attacks in Bitcoin pool mining. *Mathematical Foundations of Computing*, 1(4):393–414, 2018.

[165] Mehdi Zolfaghari and Samad Gholami. A hybrid approach of adaptive wavelet transform, long short-term memory and ARIMA-GARCH family models for the stock index prediction. *Expert Systems with Applications*, 182:115149, 2021.