Index-linked catastrophe instrument valuation

Mario Nicoló Giuricich

A dissertation submitted to the Faculty of Commerce, University of Cape Town, in fulfilment of the requirements for the degree of Doctor of Philosophy specialising in Quantitative Finance.

October 26, 2018

University of Cape Town

Supervisory committee:
Associate Professor Krzysztof Burnecki
Associate Professor Peter Ouwehand
Honorary Professor Eckhard Platen

Plagiarism declaration

I declare that this dissertation is my own, unaided work. It is being submitted for the degree of Doctor of Philosophy specialising in Quantitative Finance at the University of the Cape Town. It has not been submitted before for any degree or examination in any other University.

Signed by candidate

October 26, 2018

Declaration: inclusion of publication

I confirm that I have been granted permission by the University of Cape Town's Doctoral Degrees Board to include the following publication in my PhD thesis, and where co-authorships are involved, my co-authors have agreed that I may include the publication (which comprises most of Chapter 5):

- Burnecki, K. and Giuricich, M.N., 2017. Stable weak approximation at work in index-linked catastrophe bond pricing. *Risks*, **5**(4), p.64.

October 26, 2018

*To my mother, **Rosetta**, and my late Nonna, **Maria***

Thank you for contributing selflessly to my upbringing

Abstract

This thesis proposes four contributions to the literature on index-linked catastrophe instrument valuation. Invariably, any exercise to find index-linked catastrophe instrument prices involves three key steps: construct a suitable arbitrage-free valuation model, estimate the parameters for the underlying loss process and simulate the instrument prices. Chapters 3 to 5 of this thesis loosely follow this process.

In Chapter 3 we propose an index-linked catastrophe bond pricing model, which pervades in subsequent chapters. We furthermore show how, under certain assumptions, our model can use real-world catastrophe loss-data to find arbitrage-free index-linked catastrophe bond prices.

Chapter 4 demonstrates how we estimate parameters for the catastrophe-related insurance-loss process on which our pricing model relies. In practice, data from such insurance-loss processes is both left-truncated and heavy-tailed. We build on an existing procedure for modelling left-truncated data via a compound non-homogeneous Poisson process, and modify their fitting process so that it becomes systematically applicable in the context of heavy-tailed data. We close this chapter by presenting an importance sampling technique for simulating index-linked catastrophe bond prices.

Chapter 5 treats the new problem of finding simple, closed-form expressions for index-linked catastrophe bond prices. By using the weak convergence of compound renewal processes to α-stable Lévy motion, we derive weak approximations to these catastrophe bond prices. Their applicability is then highlighted in the context of our catastrophe-bond pricing model from Chapter 3.

Chapter 6 deviates from the ambit of catastrophe bond pricing and considers a new type of insurance-linked security, namely the contingent convertible catastrophe bond. Our foremost contribution is that we comprehensively formalise the design and features of this instrument. Subsequently, we derive analytical valuation formulae for index-linked contingent-convertible catastrophe bonds. Using selected parameter values in line with earlier research, we empirically analyse our valuation formulae for index-linked contingent-convertible catastrophe bonds.

Acknowledgements

I would like to begin by expressing my deepest appreciation and gratitude to my three supervisors, **Associate Professor Krzysztof Burnecki**, **Associate Professor Peter Ouwehand** and **Honorary Professor Eckhard Platen**. I would also like to include **Professor Zbigniew Palmowski** in this note of thanks, and acknowledge his selfless inputs which assisted me in solving the many problems which arose in Chapter 6. Professor Burnecki and Professor Ouwehand, thank you very much for your insightful guidance, patient assistance and very wise direction during the entire process. It is also to Professor Burnecki and his wife, **Ewa**, that I am deeply indebted to for welcoming me to Poland and immersing me into their family during my four visits to Wroclaw. And last but not least, even though it did set me back on quite a few occasions, I am exceptionally grateful for Professor Ouwehand's continual role as the devil's advocate when working on ideas for my PhD thesis. Without his generous help, time spent and energy expended, the completion of my thesis would not have been possible. But please do note that all remaining errors are my responsibility.

I would also like to express my heartfelt and warmest appreciation to **Professor David Taylor**, the head of the African Institute of Financial Markets and Risk Management. I thank you for believing in me from the day I first met you, and for allowing me to return to your department and pursue my PhD. I am moreover grateful for your wise advice which you persistently planted in my head over this four-year journey. And finally, I wish to thank you for all the generous funding you provided me with during this time.

I wish to thank my close and dear friend, **Melusi Mavuso**, for all the selfless help and guidance you offered me during my PhD journey. I have learned so much from you, and I will always cherish the love for mathematics which you instilled into me. Please continue spreading your contagious love for mathematics to others. I also thank my other friends who helped and advised me along the way, **Alex Backwell**, **Jonas Becker** and **Werner Opperman**. Thank you for always being ready to assist me with many of the academic- and practitioner-related issues I unearthed in my PhD work.

I also would like to express my deepest gratitude to **Irene Petrony**, who has always stood by me and supported me along the paths of my actuarial and quantitative careers. Irene, I also thank you for always taking the time to chat with me, no matter how taxing my concerns were and for continuously reminding me to "keep the PhD in sight".

I additionally wish to thank **Lameez Fakier** for quietly working behind the scenes and making the administrative parts of my MPhil and PhD journey seamless and worthwhile. May you continue to help other students in the future just like you helped me.

Going back into my past, I wish to express my appreciation for being connected with two amazing teachers, **Mrs. Karen Taylor** and **Mrs. Lelys Maddock**, who always believed in me and, furthermore perpetually forced to give of my best in whatever I wrote. The many hours we worked together and the constant chats about believing in myself, I feel, have paid off.

Finally, words alone cannot express the thanks I owe to my loving parents, **Leonard** and **Rosetta**, as well as my brother, **Alessio**. Thank you for your constant inspiration, kindness, support and unconditional love. Thank you also for consistently giving me the space to speak and rant when things did not go well during my studies, and for listening to me long and hard when I needed someone to speak to. Thank you for believing in me during times of difficulty, and for all the energy expended on words and encouragement and faith. But most importantly, thank you for giving me invaluable opportunities in my life, for allowing me to make my own mistakes and for consistently encouraging me to give of my best.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The aim of this chapter is to provide the reader with a brief overview of our work. We will, moreover, expand on the contributions that this thesis makes, in the spheres of (1) pricing of catastrophe-linked instruments, (2) modelling left-truncated heavy-tailed natural catastrophe data and (3) the design of novel natural catastrophe-linked instruments. This chapter will then be concluded by offering an outline of the chapters which follow.

1.1 Background and problem statements

Underinsurance of property-related risks[1] remains a challenge for short-term insurers and reinsurers. In the insurance markets, the threat of natural catastrophes[2] poses large amounts of property-related risks, and (re-)insurers usually finance claims arising from these risks by building up internal reserves from (re-)insurance premiums. Because of the unpredictable and extreme nature of the claims caused by catastrophic events, classical insurance mechanisms, such as, *inter-alia*, prudent reserving, reinsurance, coinsurance, sidecars and insurance loss-warranties, may each individually be unsuitable for addressing the extreme nature of the losses caused by catastrophic events. In fact, many of these catastrophe-related property risks still remain uninsured.

Natural catastrophes can give rise to large-scale property damage, and the risk arising from these natural events is called catastrophe risk. Catastrophe risks can be viewed as high-severity, low-probability risks that do not typically conform to the law of large numbers. In fact, let us consider some statistics. The 10-year inflation-adjusted average of total economic losses from both natural catastrophes and man-

[1] From an insurer's perspective, the term property risk specifically refers to risks relating to buildings and content, as well as business interruption risk. We will, however, adopt a less general understanding of the term "property risk" at a later stage.

[2] In the realm of property insurance, "catastrophe" means a natural or man-made disaster, which is unexpectedly severe and affects a very large number of insurers and policyholders simultaneously.

made disasters was US $192 billion in 2015; up by US $2 billion since 2014 (SwissRe, 2016). In order to protect against such risks, insurers typically make use of reinsurance, but have lately also begun resorting to insurance-linked securities (a form of alternative risk transfer). Although the former is more frequently used, the latter's associated markets has been argued to have a larger appetite for such risks. Therefore, a convergence of the capital and insurance markets has occurred over the last decade (Cummins and Weiss, 2009). But not all of the non-reinsured catastrophe risk is passed on to the capital markets: according to Willis (2017), the end of the third quarter of 2017 only had US $25.5 billion outstanding in non-life insurance-linked securities, a 12% increase from that in the fourth quarter of 2016.

It is because of this inherent nature of catastrophe risk, and also because of large portions of it remaining uninsured, that (re-)insurers are seeking more ways to offer protection from it, specifically property risk arising from natural catastrophes. Now , the development of the insurance-linked securities (a form of alternative risk transfer) market has offered (re-)insurers such a manner to do so.

In light of the above, we begin, in this study, by considering one of the largest sub-categories of insurance-linked securities, namely catastrophe (CAT) bonds. CAT bonds are a type of insurance-linked security, which transfer a pre-specified set of catastrophe-related risks from the sponsor (typically a (re-)insurer) to investors in the capital markets. We focus on a specific type of CAT bond, namely index-linked CAT bonds (see Chapter 2 for further details on this). This leads to our first aim - to robustly model the indices underlying these CAT bonds.

Our second and third aims involve developing a pricing model to value index-linked CAT bonds primarily at their issue-date, and we would like it to be as simple as possible (even if prices have to be approximated). The fourth and final aim is to consider the structure of the CAT bond in further detail, and consider how it can be altered and made more attractive to investors. This is primarily done in order to provide further avenues which can address the problem that many property-related risks remain underinsured.

In working towards these aims, a number of contributions were made in the sphere of heavy-tailed data modelling, catastrophe bond pricing and catastrophe bond contract design. These will now be discussed.

1.2 Scope, limitations and contribution of the study

In line with the aims of the doctorate outlined in Section 1.1, the following main points will be addressed in this thesis.

1. We will introduce a rigorous methodology for modelling heavy-tailed, left-

truncated[3], catastrophe data in the compound non-homogeneous Poisson process framework. But first, we shall begin by showing that our data belong to the domain of attraction of an α-stable distribution. Then, we will go on to compare two procedures for estimating the severity and frequency distributions of the Property Claim Services (PCS) loss index[4], an index exhibiting non-random left-truncation. Of the two procedures we compare, one does not account for the left-truncation while the other does. The performance of the two approaches will then be compared for several heavy-tailed distributions. To achieve this, we introduce a modification of the maximum-product-of-spacings parameter estimation technique, and we will also, for the purposes of simulation, propose a Monte Carlo importance-sampling algorithm to ensure that large losses are satisfactorily simulated from the fitted heavy-tailed distributions. Finally, we will illustrate the potential usage of both the fitting and simulation methodologies, by presenting index-linked catastrophe bond prices with the trigger specified in relation to the behaviour of the PCS loss index.

2. We will attempt to provide a CAT instrument pricing model, primarily for pricing these instruments at their issue-date. In particular, we will work towards having a model that does not assume, as is done in the classical literature of CAT bond pricing, that catastrophe-risk variables retain their distributional characteristics when moving from the risk-neutral probability measure to the real-world probability measure (see Cummins and Geman (1995a); Cox and Pedersen (2000); Lee and Yu (2002); Vaugirard (2003a,b); Ma and Ma (2013) and Nowak and Romaniuk (2013)). Note that these arguments all stem from the framework considered by Merton (1976). In these models, natural catastrophes have been treated as idiosyncratic risks that can be (almost) fully diversified. But take note that our model will assume that catastrophe risk is independent of financial markets risk. Finally, we consider applying our CAT instrument pricing model to pricing CAT bonds.

3. We will also consider the subject of approximating tail probabilities in the general compound renewal process framework, where severity data are assumed to follow a heavy-tailed law (in that only the first moment is assumed to exist). By using weak convergence of compound renewal processes to α-stable Lévy motion, we will derive such approximations. Their applicabil-

[3] The reason for considering left-truncated data is because of the left-truncation feature prevalent in the data which we later work with.

[4] This is the insurance-loss index upon which our index-linked CAT bonds will be assumed to be based, and it is used extensively in the papers by Chernobai *et al.* (2006) and Ma and Ma (2013).

ity will then be highlighted in the context of our CAT bond pricing model, which is indeed based on a compound renewal processes. The behaviour of our weak approximations will then be compared to both Monte Carlo simulations and first-order single risk loss approximations.

4. We will consider one of the potential problems (to the issuer) posed by CAT bonds, in that there is a debt which still has to be recognised on the issuer's balance sheet *ex post* the trigger (or after the occurrence of the catastrophic events). As a challenge to this potential problem, we will bring to the fore an alternative CAT bond instrument design, whereby the investor will recover his or her capital in the form of equity – this is called a contingent convertible catastrophe bond. We subsequently endeavour to investigate a framework to price, at the issue-date, such an instrument.

Before we present the contributions of the study, the following limitations must be noted.

1. We did not have access to a complete, sufficiently large enough, reliable and robust set of CAT bond pricing data[5]. Therefore, this unavailability of data limited us from using CAT bond pricing models which work solely under a risk-neutral pricing measure, and also inhibited us from finding market prices of risk. In consequence, this is a major reason for why our pricing models in this study operate under the real-world probability measure. We want to be able to test the models with real-life data, and it also makes comparisons of our results with those of previous studies more feasible. Moreover, given the lack of CAT bond pricing data, we assumed that random processes pertaining exclusively to catastrophe risk have the same distributional characteristics under the real-world and risk-neutral probability measures. By adhering to this assumption, we therefore did not attempt to identify the risk-preferences of catastrophe bond investors, and it should also be noted that attempting to identify these risk-preferences is a subjective and debatable exercise given the novelty of the catastrophe risk market.

2. It must be borne in mind that there is no single and widely accepted way to price CAT bonds, and also to estimate parameters of the various processes

[5] The Artemis Deal Directory online now records CAT bond spreads for many transactions with information in the public domain. However, much of the CAT bond data given on the website is opaque: spreads are listed, however, it is difficult to ascertain the precise structure of each and every bond. Given that we were specifically looking for a particular index-linked type of CAT bond data, it was not possible for us to find a complete set thereof. Also, we noticed that the deal directory only began providing more comprehensive data from 2016, and was unfortunately not very useful during the earlier stages of the work.

used in their valuation. This explains why there are a number of models in the literature to price them (see Chapter 2 for a review of most of these models), and hence the subjectivity behind the views of CAT-bond practitioners. But the existence of all these different types of models is also explained by the very fact that CAT bonds are not standardised (unlike options), and different models have been postulated for different CAT bond structures. Therefore, we just adopted a particular pricing approach and worked with it, providing more of a theoretical contribution to the literature as opposed to a more practical one. Finally, we point out that we did not have access to existing natural scientific models (such as those used by catastrophe-risk modelling agents such as AIR-Worldwide and Risk Management Solutions [RMS]) to perhaps model the natural catastrophes on a more sophisticated and detailed basis.

3. Since we only had insurance-loss data from a particular insurance-loss index (that is, the PCS index), we constrained ourselves to considering only index-linked insurance-linked securities.

4. The index loss data was only available from 1985 to 2011. We were unable to source further data, so as to more properly elicit trends from it. Also, the data was composite in nature - it included recorded insurance losses from a wide variety of natural catastrophes (e.g. windstorms, fire, earthquake etc. - see Chapter 2 for a further discussion on the data and its source). In an endeavour to model the data as precisely as possible, an argument can be put forward on the basis of deriving individual models on a more granular level, and then combining them all at a later stage. However, we did not have access to the data on a more granular level and moreover we did not have enough years' worth data to make adequate deductions. Also, in the interests of creating a simple, user-friendly and parsimonious model it is of interest not to model the data on a more granular level.

Given the above background and limitations of the study, we can nonetheless summarise the contributions of this study to the fields of (1) pricing of catastrophe-linked instruments, (2) modelling left-truncated heavy-tailed natural catastrophe data and (3) the design of novel natural catastrophe-linked instruments, as follows.

1. We present a robust and stepwise methodology to model heavy-tailed left-truncated catastrophe-related loss data as a non-homogeneous compound Poisson process. In doing so, we demonstrate

   (a) the need for considering a special case of the generalised extreme value distribution, especially when the support of the distribution is required to be non-negative. This aspect has been overlooked in many recent practically-orientated research papers. We also demonstrate the need for a "working parameter" when estimating the generalised extreme value distribution's parameters - this ensures the necessary positivity of certain parameters, in light of the support.

(b) a maximum product of spacings estimator for heavy-tailed severity distributions, accounting for left-truncation, that results from a function that is not unbounded (i.e. this function can be maximised, unlike the unbounded likelihood function issue which arises with maximum likelihood estimation in most cases which we consider).

(c) the rather surprising result that accounting for left-truncation gives rise to fitted heavy-tailed severity distributions that have a finite first moment for certain distributional forms. If the left-truncation feature is ignored, then most first moments are infinite for the considered fitted distributions.

(d) the application and refinement of Moran's log spacings goodness-of-fit test statistic to left-truncated data. Also, to the best of our knowledge, it is the first time that such a goodness-of-fit test statistic has been compared to the more traditional quadratic and supremum class goodness-of-fit test statistics. Encouragingly, our version of Moran's log spacings statistic gives them same reject/do-not-reject decisions in the hypothesis tests performed.

(e) an elementary importance-sampling technique to robustly simulate, via Monte Carlo simulation, realisations from our fitted heavy-tailed distributions (Burr, generalised Pareto and generalised extreme value distributions).

(f) evidence for higher loss-exceedance (i.e. tail) probabilities and lower simulated catastrophe bond prices when accounting for the left-truncation.

2. We derive simple, closed-form tail probability approximations for use in index-linked catastrophe bond prices, by using the weak convergence of compound renewal processes (and in consequence compound Poisson processes) to α-stable Lévy motion. To the best of our knowledge, this is the first time that catastrophe-related insurance loss processes, such as the Property Claims Services indices, have been modelled using an α-stable Lévy motion. We show that our approximations

   (a) are useful if only the first moment of the severity distribution exists. Therefore, when other approximation techniques based on higher moments or cumulants of the compound Poisson process cannot be used (for example, consider the mixed approximation technique of Ma and Ma (2013) and also that of Chaubey *et al.* (1998)), ours can.

(b) perform well compared to Monte Carlo simulation and first-order single loss approximations.

3. We present a new CAT bond pricing model, which accommodates for the fact that catastrophe risk processes may not have the same distributional characteristics under the real-world and risk-neutral probability measures. In consequence, our model can allow for higher than expected severities and frequencies, under the assumption that it follows a time-inhomogeneous compound Poisson process, than what the real-world data suggests. But another interesting feature is that our model encompasses the special case when catastrophe risk variables are posited to retain their distributional characteristics when moving from the risk-neutral probability measure to the real-world probability measure. A novel feature of our model is that it leads to neat pricing formulae under the real-world pricing measure, so all one needs to implement it is real-world insurance-loss data (and not pricing data). Our model still does, however, encompass the classical assumption that catastrophe risk and financial markets risk are independent.

4. We present the first formalisation of a new catastrophe instrument, namely the contingent convertible catastrophe bond. This instrument has its roots in the contingent convertible (Coco) bonds issued by banks.

   (a) We present a practical discussion around the design and features instrument, and consider how it will affect both the theoretical balance sheet and the true balance sheet (as dictated by International Financial Reporting Standards (IFRS) 7 now, and IFRS 9 from 2018)

   (b) We motivate the benefits of issuing the new contingent convertible catastrophe bond over a more traditional CAT bond and catastrophe-equity put. In contrast, we also discuss some of the difficulties with issuing such a new instrument.

   (c) Based on real-world probability measures, we present a valuation framework for this instrument, which is based on some novel mathematical tools. By using

      - the assumption of independence of catastrophe risk and financial markets risk,

- an exponential change of measure for catastrophe-related processes, and
- a Girsanov-like change of measure to remove correlation between standard Brownian motions,

we are able to present analytical pricing formulae for our new instrument.

5. We present simulations of prices for contingent convertible catastrophe bonds, and perform a comparison between Monte Carlo estimated prices and prices based on our analytical formulae. For the consequential quasi-exponentially-tilted random variables we are required to simulate from (as a result of the exponential change of measure), we present a modified simulation technique based on the Acceptance-Rejection algorithm. Such simulation exercises could be useful in pitch-books used by insurance-linked securities structuring agents (such as Aon Securities, Credit Suisse Insurance-linked securities, Nephila and Leadenhall Capital Partners) to pitch the idea of issuing such novel investments to prospective issuers.

In the final section of this chapter, a chapter outline is provided, indicating in which sections our various topics are addressed.

1.3 Organisation of the study

Chapter 2 provides an expansive, but not too detailed, literature review on a number of topics. It considers most of the tools and approaches used thusfar for CAT bond pricing, but in addition looks at related topics such as reinsurance pricing, loss distribution estimation techniques and incomplete market pricing. Finally, it ends off with a brief discussion around some other insurance-linked securities in issue today, or which have been issued in the past.

In Chapter 3, we set up a probabilistic framework and bring to the fore our generalised CAT instrument pricing model, which we use mainly for pricing at the issue-date of such instruments. Thereafter, we apply our developments to the sphere of CAT bond pricing. We also show that our model is but a generalisation of many of the CAT bond pricing models considered in the literature thusfar (for example, that of Ma and Ma (2013)).

In Chapter 4, we introduce the methodology for robustly modelling heavy-tailed left-truncated catastrophe insured loss data, based on the Property Claims services data. We then go on to illustrate how one implements a robust Monte Carlo simulation algorithm for distributions fitted via this approach, and apply this within the context of index-linked CAT bond pricing, using our model from Chapter 3. We end off by providing a comparison of CAT bond prices, when the data is modelled accounting for the left-truncation, and when the data is modelled ignoring it.

Chapter 5 introduces the idea of weakly approximating a compound renewal process by an α-stable Lévy motion. We then apply this to the compound Poisson process fitted to the data, and in consequence obtain simple approximations to CAT bond prices. Again, we perform a numerical illustration and compare our approximations to other approximations, such as Monte Carlo simulation and first-order single loss approximations.

It is in Chapter 6 where we embark on addressing the problem of underinsurance of catastrophe-related risks. Here, we bring to the fore our idea (based loosely on an instrument issued by Swiss Re in 2013) of a contingent convertible catastrophe bond. We present a motivation for it, and also discuss practical issues such as accounting implications and the benefits of it over, say, the already-existent catastrophe-equity puts and CAT bonds. We provide a valuation framework for the new instrument, and end off with a numerical illustration of prices and the effects varying parameters can have on these. A conclusion containing a summary of our work, as well as further research suggestions, is presented in Chapter 7.

Chapter 2

Literature review

In this chapter, we provide an overview of the literature on the theme of this thesis. Firstly, we offer a discussion on natural catastrophe risk and why it is suitable for methods of alternative risk transfer.

Secondly, we present a discussion on catastrophe (CAT) bonds, since this is the pivotal theme of our work. We then go on to discuss the loss distributions of natural catastrophe claims; an aspect which is crucial for the modelling of catastrophe bond prices. Finally, we provide a brief yet comprehensive discussion on the catastrophe bond (and also other catastrophe instrument) pricing to date.

2.1 Catastrophe risk and insurance-linked securities

For insurance companies, natural catastrophes present a peculiar case. Indeed, natural catastrophes "provide a principal justification for insurance" (Zeckhauser, 1995), however, it remains a necessary condition for insurers to control their exposures to such large and unpredictable losses in order to function economically. Consider, for example, an insurer selling more than five million auto-insurance policies. Here, no single car accident can be predicted with certainty, but by the law of large numbers the total number of accidents for the entire book of policies is fairly predictable. Whereas contrast to this scenario the case of a natural catastrophe — where either no policy is affected, or most (or even all) are affected. The law of large number appears to no longer apply in this case.

By considering this above simple example, it becomes evident why natural catastrophes are uninsurable risks from a textbook perspective. Natural catastrophes violate one of the basic tenets of insurance, that is the diversifiability and independence of risks (see Cummins and Geman (1995b)), and Nguyen and Lindenmeier (2014) in fact provide qualitative evidence in support of this claim). It is possible that a single catastrophe could lead to undesirable consequences such as bankruptcy (especially in the extreme situation) or reserve inadequacy. In contrast, Jaffee and Russell (1997) have argued that insuring natural catastrophes would not impede the operation of a successful private insurance market in this field. Despite the problems with the insurance of catastrophic losses, a large and successful property liability insurance and reinsurance industry exists (BNY Mellon, 2013), with a variety of key players from private insurance companies to global reinsurers and governments.

It has been reported that the insurance and reinsurance markets, in isolation, are unable to bear the yoke of all the risks policyholders insure against, in particular low-frequency high-severity tail events such as natural catastrophes. Indeed, such markets have become quite constrained (Swiss Re, 2012). Hurricane Andrew in 1992 and Hurricanes Katrina, Rita and Wilma in 2005 have cost insurers unprecedented losses of US $24 billion in 2007 prices (Cummins and Weiss, 2009) and US $114 billion (Swiss Re, 2008) respectively. Despite most of these insured losses being paid out, each of these hurricanes gave rise to various insurer insolvencies (Cummins *et al.*, 2002) and large fluctuations in insurance and reinsurance rates-on-line, mainly resulting from fluctuating reinsurance capacities. Hence, because of the "undercapitalisation of the reinsurance industry", there is default risk since both insurers and reinsurers are subject to the same risks. As a consequence of these hurricanes and other natural catastrophes (and also coming to the realisation of the aforementioned), there has – to a large extent – been a convergence of the capital and insurance markets (Cummins and Weiss, 2009; World Economic Forum, 2008). Because of this convergence, insurers assuming risk - say from a natural catastrophe - have been able to cede this risk not only to reinsurers, but to players in the capital markets as well.

From a perspective of contractual risk transfer, Cummins and Weiss (2009) earmarked five reasons as to why this convergence occurred. Firstly, the most important factor causing convergence is attributable to the increase in property values in geographical areas which are susceptible to catastrophic risk. The second most important factor driving convergence is the reinsurance underwriting cycle. The authors also warn about the fluctuating capacity constrains that reinsurers face, especially in times following catastrophic losses. A third factor is the improvement in communications and computer technologies. The last two factors driving convergence are the result of developments in firm-wide risk management strategies and the improvement in accounting, regulatory, tax and rating agency factors, respectively. Finally, another factor should be added to this list. This factor concerns the difficulty of putting in early-warning risk-mitigation strategies in place to lessen the impact of natural disasters.

Notwithstanding the reasons posited by Cummins and Weiss, there is reason for investors in capital markets to "invest" in the occurrence of natural catastrophes. A natural way for investors to gain access to such risk is via insurance-linked securities (ILS). By doing so, these investors are partial bearers of the insurance industry's risk. Doherty (1997) as well as a number of industry reports (see Swiss Re (2014) and BNY Mellon (2013)) argue theoretically that natural disaster risks are essentially uncorrelated with capital market variables such as aggregate consumption and interest rate levels (see also Froot *et al.* (1995)). Therefore, diversification opportunities are offered to investors. Moreover, the addition of collateralised ILSs to diversified portfolios may enhance risk/return opportunities for investors (Litzenberger *et al.*, 1996), and such high returns have been noted in the market review by Nguyen and Lindenmeier (2014). Finally, from the perspective of the insurer of catastrophic risks, securitisation offers the benefit of the diversification of ceded capital (Cox *et al.*, 2000) as well as the transfer of these large risks, in derivative or security form, into markedly larger pools of investment capital (Canabarro *et al.*, 2000).

Let us now magnify into a particular type of ILS. According to Cummins (2008b) as well as Barrieu and Albertini (2010), one of the most successful non-traditional risk-hedging instruments has been catastrophe bonds. Catastrophe bonds, commonly known in industry and practice as CAT bonds, are instruments "which seek to ensure capital requirements are met in the specific instance of a catastrophic event" (Baryshnikov *et al.*, 1998). Today, CAT bonds are a dependable source of capacity for insurers, reinsurers and governments. These bonds also provide sponsors with multi-year, multi-line coverage, and replace the year-to-year turbulences experienced by annual reinsurance premiums with a fixed cost over the duration of the bond. The market for CAT bonds has grown steadily from 2005 to 2016, average tranche size has increased (with the size of 2014 notes almost double of those issued in 2007) and issuances have reached a new record of a total issuance of $8.5 billion in total for the first three quarters of the year 2017 (Willis, 2017). Figure 2.1 illustrates the increasing trend, over time, in CAT bonds' and other ILSs' capital issued and capital outstanding. Note that the majority of the risk capital issued and outstanding stems from CAT bond issues.

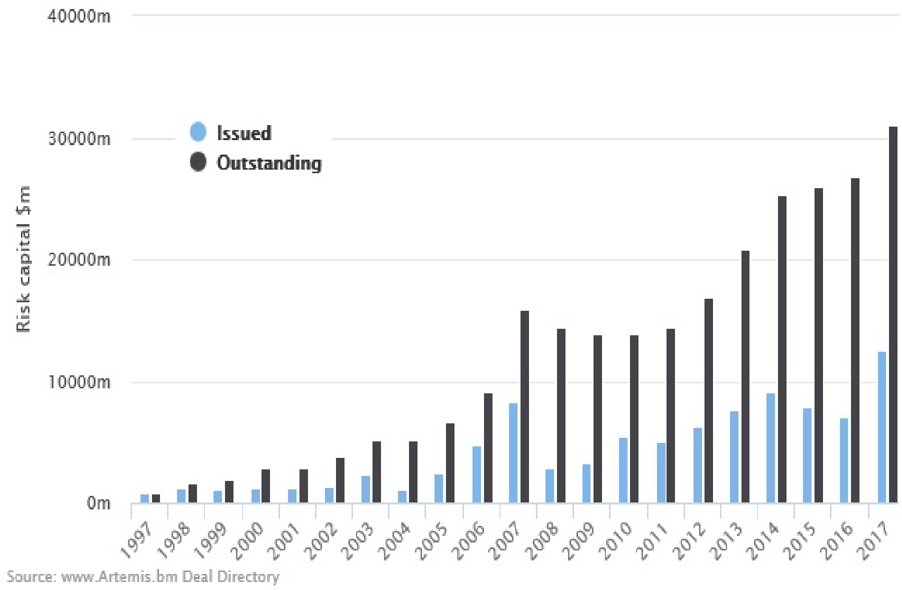

Fig. 2.1: CAT bond and other ILS capital issued and outstanding as at 1 January 2018. Source: `http://www.artemis.bm/deal_directory/cat_bonds_ils_issued_outstanding.html`.

Cummins and Weiss (2009) provide two justifications for the development and growth of CAT bond markets. Firstly, insured losses from large-scale natural catastrophes such as Hurricane Katrina and Hurricane Andrew are large relative to the capital resources of insurance and reinsurance markets, however, they are small relative to the magnitude of those from the capital markets. Secondly, because securities markets reduce information opacities and aid price setting (with prices being set by the equilibrium between supply and demand), capital markets are far more price efficient than insurance markets. On balance, it is reasonable to suggest that the CAT bond market will experience continued growth (see the industry reports of BNY Mellon (2013); Swiss Re (2014); Willis (2015)).

2.2 A primer on catastrophe (CAT) bonds

CAT bonds are a type of ILS. Their primary objective is to transfer reinsurance risk associated with natural catastrophes from companies, insurers and reinsurers as well as governments to the capital markets. Essentially, CAT bonds are interest-rate instruments which are structured similarly to corporate or government bonds; the difference is that the possible coupon payments and the repayment of the face value depend on the occurrence of certain or several natural (or perhaps man-made) catastrophe-related events.

Harrington and Niehaus (2003) provide justification for the issuance of CAT bonds: the reasons a financial entity may wish to issue such bonds are because of (1) the lower tax costs of equity financing and (2) the avoidance of the financial distress costs associated with subordinated debt issues. An additional benefit of CAT bonds is that they provide immediate liquidity *ex ante*, unlike pure insurance derivatives (such as, *inter alia* catastrophe options and swaps), which provide liquidity *ex post*. Both Lane (2004) and Cummins and Weiss (2009) summarise the structure of CAT bonds as follows (and see also Figure 2.2 below):

1. A sponsoring company (such as an insurer or reinsurer) sets up a special purpose vehicle (SPV) together with a structuring agent and legal counsel. This SPV, (a) issues the CAT bonds and (b) acts as a source of reinsurance protection. The SPV is typically established in an offshore tax haven (such as the Cayman Islands or Bermuda) where the issuer is not subject to any capital gains, income, profit, corporate and witholdings tax (Hammer and Singer, 2001). Note that according to IFRS 7, a catastrophe bond can be issued as a derivative or reinsurance contract. In order to be classified as a reinsurance contract, the SPV is required to be “appropriately licensed in the jurisdictions in which it conducts its business” (Barrieu and Albertini, 2010).

2. A contract between the SPV and issuer is established[1]. The issuer pays premiums to the SPV, while the SPV covers the risks from an identified pool of the insurer’s catastrophe exposure.

3. The SPV issues CAT bonds to investors. This may be done with the assistance of a structuring agent, who assists in making the issue more attractive. In this transaction, investors from the capital markets pay a principal to the SPV, and they typically receive a remunerate interest equal to LIBOR (for example) plus a risk premium. The risk premium is usually equivalent in magnitude to the premium paid by the insurer. After issue, a ratings agency provides a credit rating to the bond.

4. The SPV uses the principal received to purchase highly-rated short-term investments such as short-term US treasuries. These investments are held in a trust account, and these investments together with the trust account ensure that the entire CAT bond is fully collateralised.

In addition to the CAT bond structure described above, the trust account often swaps the fixed returns from the short-term investments for floating returns

[1] Note that if the issuer is an insurer or reinsurer, then the contract between the SPV and issuer is, technically speaking, reinsurance. Consequently, insurance premiums and payments are tax-free.

based on LIBOR[2] (as represented by LIBOR - X in Figure 2.2). The rationale for the swap is to protect the investors and the insurer from default risk and interest rate risk (Cummins and Weiss, 2009), and moreover to base the investors' investment income on a reference interest rate (Müller and Grandi, 2000). Note that this structure should not be viewed as a bond in the traditional sense: as can be seen from Figure 2.2, it is a complex interplay of a number of transactions.

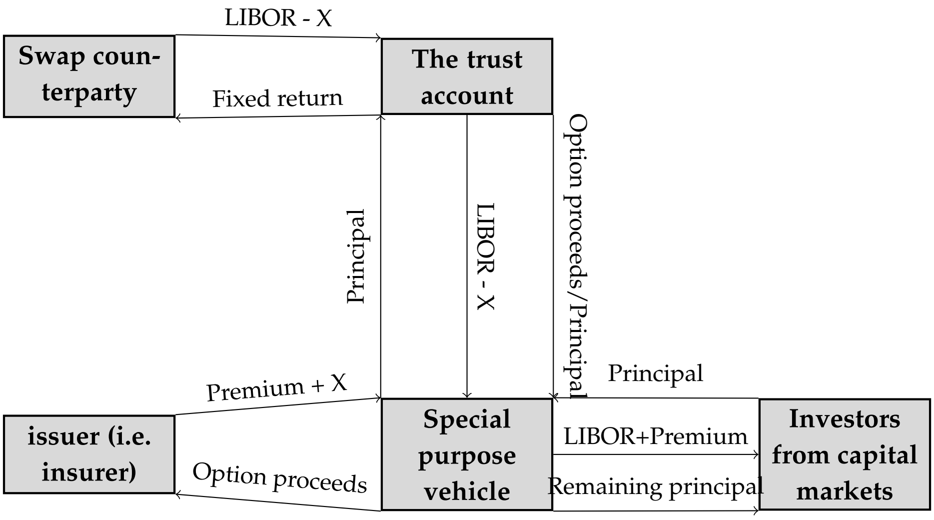

Fig. 2.2: Typical structure of a CAT bond.

Upon closer inspection, it is evident there is a call option embedded in the CAT bond structure: contingent upon the occurrence of a pre-specified event (the trigger event), capital is transferred from the SPV to the insurer in order to assist in paying out claims. The trigger event represents an equilibrium between the preferences of bond investors and the issuer. Notice that the difference between CAT bonds and corporate (or government) bonds – in the latter the payments depend on the creditworthiness of the issuer – lies in the fact that the trigger is outside the sphere of influence of the contract partners.

Depending on the type of CAT bond, if the trigger event occurs, then investors could lose their entire investment of the principal. In a **principal-at-risk** tranche, there is the risk that the full principal invested is not repaid to investors, while in a **principal-protected** tranche, the return of the principal is guaranteed. However, in this latter case, the occurrence of the trigger event affects spread and interest repayments, as well as the time at which the principal is repaid. Principal-protected tranches are not common, because they do not provide enough risk capital to issuing companies compared to principal-at-risk tranches (Canabarro *et al.*, 2000; Cum-

[2] London Interbank Offered Rate. Other money market fund rates may also be used, such as EURIBOR.

mins and Weiss, 2009).

Four broad types of CAT bond triggers exist, with the first three being more traditional, namely:

- *Indemnity triggers*, where the trigger is based on the actual losses of the issuer. One of the most significant risks associated with this trigger type is that of moral hazard[3]. The cover provided by this CAT bond is only effective if the issuer incurs a pre-determined level of losses (SwissRe, 2011); because of this construction, investors need to have detailed knowledge of the business of the issuer. In consequence, the issuer needs to provide certain information about its business. Note that most of the seminal CAT bonds in the 1990s were characterised by indemnity triggers.

- *Parametric triggers*, where the trigger is based on occurrences of pre-specified characteristics or criteria of a pre-specified natural disaster. For example. the trigger may be specified in relation to the Richter scale readings on the magnitude of an earthquake at a specific location. The risk assessment procedure for parametric CAT bonds is transparent for investors, but, the issuer could be exposed to basis risk (SwissRe, 2011). However, special parametric index triggers can be employed to decrease a issuer's basis risk by using a weighted number of locations (and not one single location). Note that CAT bonds with parametric triggers have significantly smaller issue volumes compared to CAT bonds with indemnity triggers.

- *Industry index triggers*, where the trigger is based on insurance-industry catastrophe loss indices, such as that provided by Property Claims Services (PCS) of the Insurance Services Office in the US or that provided by PERILS in Europe. We reiterate that CAT bonds based on such a trigger are commonly called **index-linked catastrophe (CAT) bonds**. Essentially, such triggers operate on the basis that the issuer recovers a percentage of total industry losses in excess of a predefined attachment point (or threshold level). These triggers suffer from the drawback of basis risk[4], because the performance of the loss index may not adequately reflect the performance of an insurer's pool of risks. However, using an industry index increases transparency for investors. Note that weighted industry indices can also be used: these can potentially

[3] In the current context, moral hazard refers to the act of the issuer deliberately over-estimating their losses. Harrington and Niehaus (2001) as well as Kunreuther (2006) point out that government assistance can encourage moral hazard to some extent, especially if governments provide aid to catastrophe-struck victims.

[4] Basis risk is the risk of a potential mismatch between the cashflows of the protection instrument and the losses it is supposed to be hedging.

decrease the issuer's basis risk by creating coverage which more closely aligns with the issuer's own portfolio of risks. A popular version of such a weighted industry index is a weighted PCS index.

- *Modelled loss triggers*, wherein a risk-modelling firm provides an evaluation of the catastrophe risk. Upon occurrence of a catastrophe, the risk-modelling firm's model evaluates the expected loss using the physical, observed parameters of the catastrophe, in order to provide an estimate of the expected losses of the issuer's portfolio (SwissRe, 2011). The CAT bond is then settled on the basis of this estimate, and not the issuer's actual losses. Basis risk to the issuer is low, however, the CAT bond trigger is not that transparent to investors.

Finally other less-important trigger types, on the basis of other types of catastrophic events (not necessarily natural catastrophes), do exist. For example, the *medical benefit ratio triggers* and also *pandemic-related triggers* are related to unexpected increases in medical benefit claims ratios in health insurance. Also, *mortality index triggers* are based upon the difference between expected and observed mortality rates. However, these triggers sit within the life-insurance related arm of ILS issuances. Finally, triggers can be combined into *multiple triggers*. Figure 2.3 shows the relative issue sizes of each of the types of CAT bonds (by trigger type). Notice that indemnity trigger bonds are most common, with industry index triggers being the second most common.

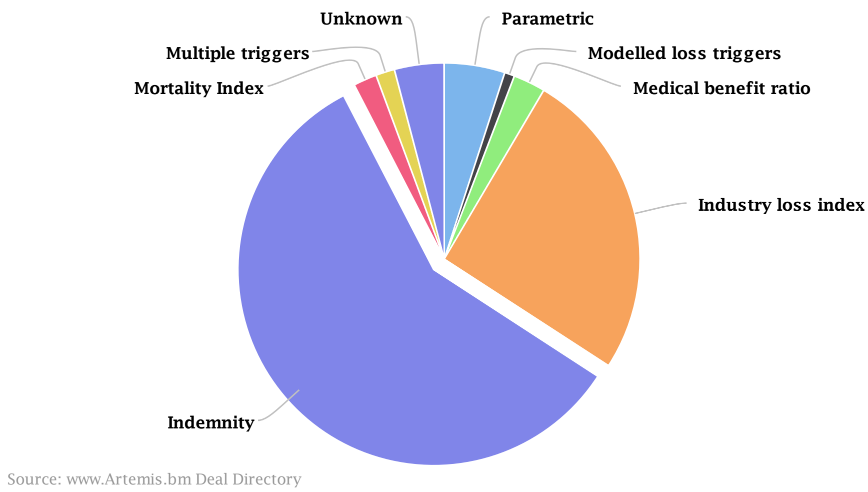

Fig. 2.3: CAT bond and ILS risk capital outstanding, by trigger type, as at 29 October 2017. Source: `http://www.artemis.bm/deal_directory/cat_bonds_ils_by_trigger.html`.

Another point, which is not commonly mentioned in the literature, is the concept of a trigger basis. Broadly speaking, this basis refers to the number of events to which the trigger is exposed to. Two types of basis exist: a *per-occurrence* trigger basis means that the trigger is based on a single event, while a *per-aggregate* trigger basis means the trigger is based on multiple events over a pre-defined term[5] or on an annual basis. The two different types of trigger bases would seem to justify different loss modelling techniques. This will be briefly discussed further in Section 2.3.

Typically, the structuring agent (which can also be a capital market division of a reinsurer) will be involved in the design as well as the placement of the bond. A modelling agent (such as RMS or AIR-Worldwide) is also called upon by the issuer to provide the necessary catastrophe modelling tools and to estimate the risk to which the issuer is exposed. The purpose of the risk analysis undertaken by the modelling agent is to estimate the various losses, which can be incurred, as well as the probabilities thereof. After the risk analysis has been performed, the issuer together with the structuring agent decides on the appropriate level of risk against which to protect.

From the discussion of the above structure of a typical CAT bond, it is evident that CAT bonds function in a similar manner to collateralised, multi-period reinsurance contracts (Froot, 2001, 2008). This is one of the reasons that has contributed to the convergence of capital and insurance markets, and has furthermore made reinsurance "interchangeable", in some sense, with CAT bonds.

According to an industry report by BNY Mellon (2013), 75% of the ILS sector is focused on US perils, those perils being mostly earthquakes and windstorms. The primary reason for this concentration is that granular US data of a high quality is available over long time periods, and this data is a requisite for the modelling of the natural catastrophes. Therefore, it is not surprising that industry index trigger CAT bonds comprise the second-highest capital outstanding (see Figure 2.3), with the majority of such bonds being written on the US-peril catastrophe loss index calculated by PCS (Swiss Re, 2012).

We close this section on the following note. Over time, it must be borne in mind that CAT bonds have become more standardised because of the common needs of the parties to the bond (Cummins, 2008a), but are still mostly traded over-the-counter and are not fully standardised (by a central counterparty, for example). However, according to industry practitioners, a small secondary market for CAT bonds does exist in Bermuda.

[5] This is the case for index-linked catastrophe bonds.

2.3 Loss distributions for catastrophe claims

Given the pervasiveness of the PCS index in index-linked catastrophe bonds, its modelling has received much attention in the literature. Let us begin with a short overview of this index.

PCS is a division of the Insurance Services Office (which is a subsidiary of Verisk Analytics) and is based in the USA. PCS is, in fact, recognised as the world's authority on catastrophe-related property and casualty insured losses (Burnecki *et al.*, 2000). PCS collects information on anticipated industry-wide insurance payouts covering different property lines, resulting from pre-specified natural catastrophes which occur in the USA (and its associated territories) and meet PCS-specified criteria. PCS publishes an "index" which quantifies these payouts – or insurance losses. The following distinct event types are represented in the data:

- Geophysical events, such as earthquakes, tsunamis and volcanic activities;
- Meteorological events, such as hurricanes and windstorms;
- Hydrological events, such as floods and landslides;
- Climatological events, such as temperature extremes, droughts and forest fires.

Now, a natural question would be to ask how this index is constructed. Since 1949, PCS has been responsible for the construction of a PCS index on catastrophic losses. PCS first assembles its estimates of insured property damage by utilising a variety of procedures including surveys of insurance companies, its National Insurance Risk Profile and its own surveys on the ground (Chicago Board of Trade, 1995), and then constructs the index using the method outlined briefly below.

One part of the index compilation involves a survey of agents, adjusters and companies. Based upon gross premium-written market shares, PCS surveys at least 70 per cent of the US insurance market. PCS then constructs an amalgam of claim estimates and individual losses reported on by these various sources and by using these figures as well as claim projections, PCS compiles a total industry estimate. Also, PCS indices track insured catastrophic loss estimates on each of a national, regional, and US-state basis (Burnecki *et al.*, 2000).

The PCS indices, available on both a national (US) basis and regional basis, represent the year-to-date aggregate amount of incurred insurance losses from natural catastrophes. Each index value provided by PCS represents \$100 million worth of damage. As an example, a value of 77.3 in 1968 means that $\$77.3 \times 100$ million equals \$7.73 billion in aggregate losses were incurred in 1968. Using this methodology, it is possible to construct the PCS index from data provided by PCS.

Up to now, various distributions, most being power laws and/or heavy-tailed laws (such as the Burr type XII, generalised Pareto and generalised extreme value distributions), have been fitted to data from this index (Burnecki *et al.*, 2000). Burnecki *et al.* (2000) also found that this index's log-return process is mean-reverting and that there is no significant autocorrelation of index returns, however, in light of global warming, one would expect a slight increasing trend in the PCS index. Note that the PCS index is not an index in the traditional sense, but rather a representative time series of aggregate losses arising from pre-defined natural catastrophes. So for example, all the aggregated insured losses resulting from hurricane Katrina in 2005 represented a single entry in the dataset. For further details on the construction of this index, we refer to the PCS report of Kerney (2013). However, we will also mention that to the best of our knowledge, no survivorship bias is represented in the data (that is, it is not necessarily the case that the same companies are represented over the entire data sample). Moreover, and for the purposes of completeness, note that the data includes losses incurred by insurers and reinsurers.

In the operation of the PCS index, one of the key defining features of a natural phenomenon classifying it as a catastrophe is that it is expected to cause losses exceeding a PCS-specified monetary threshold. In the 1980's, this threshold was set at \$5 million, but in 1997 it was increased to \$25 million, where it remains today. This threshold signifies **left-truncation**, in that both loss magnitudes and the number of such losses below the \$25 million threshold are not recorded. This phenomenon nestled within the operation of the PCS index was first studied by Chernobai *et al.* (2006). However, left-truncation features of datasets has also been studied in the context of the more general areas of operational risk by Mignola and Ugoccioni (2006), Luo *et al.* (2009) and Ergashev *et al.* (2016). Since natural catastrophes (and consequently the PCS index) give rise to such large insurance losses, one needs to be as precise as possible in the modelling process, and therefore the left-truncation present in the PCS index must be accounted for in attempts to model it.

CAT bonds, which (for example) are based on the PCS loss index, have the feature in that their underlying trigger event is dependent upon both the frequency and severity of natural disasters. It is this very feature which sets CAT bonds quite apart from conventional government or corporate bonds. Therefore, the loss process used in the CAT bond pricing should account for this dual dependency. In addition, for the sake of generality of application, the assumptions about the underlying distribution functions (for the severity) need to be kept at an absolute minimum.

It was mentioned in Section 2.2 that there were two bases to be considered in the design of CAT bond triggers, namely per-occurrence and per-aggregate bases. Due to their different natures, it appears fair to model the loss processes in each of these instances differently. The former should account for the occurrence of the first catastrophe with total insured losses in excess of some pre-defined threshold (the trigger), while the latter should account for the aggregate losses from multiple catastrophes, when these exceed some pre-defined threshold (the trigger). Our research focuses on the latter basis, as it most commonly appears in the literature and is also more commonly used in practical CAT bond settings (SwissRe, 2011).

Bearing this per-aggregate basis in mind, it seems suitable to model the loss process, in general, as a marked Poisson process (or even a compound renewal process). Embrechts and Meister (1997) state that there exists both statistical and theoretical support for Poisson-type assumptions on the loss processes and postulate three broad categories into which the distribution of the loss processes may fall:

1. A homogeneous or non-homogeneous compound Poisson process,
2. A mixed compound Poisson process, or
3. A doubly stochastic (or Cox) compound Poisson process.

Before proceeding with a discussion of these Poisson-type models in the current literature, it is necessary to indicate what is meant by the term "catastrophic loss". Embrechts and Meister (1997) define catastrophic losses as losses in excess of some threshold. Therefore, a motivation for a definition of a natural catastrophe can be developed. This research will define a natural catastrophe as an event caused by natural forces (such as floods, hurricanes or earthquakes) resulting in a large number of related individual losses, with the total losses exceeding some pre-specified threshold level[6].

The Poisson process has been extensively employed to model the occurrence of natural catastrophes (see, for example Cummins and Geman (1995b); Aase (1999); Lee and Yu (2002); Lee and Yu (2007); Burnecki and Kukla (2003); Cox *et al.* (2000); Jaimungal and Wang (2006a); 'Biagini *et al.* (2008b); Lin and Chang (2008); Härdle and Cabrera (2010); Braun (2011); Ma and Ma (2013) and Nowak and Romaniuk (2013)). However, fewer papers have applied their Poisson modelling techniques to the PCS loss data. Chernobai *et al.* (2006) modelled PCS recorded insurance losses from 1990 to 1999 using a non-homogeneous compound Poisson process with a

[6] So intuitively, a catastrophic loss can be used as a trigger for an index-linked CAT bond.

deterministic sinusoidal intensity rate function, while considering a range of distribution functions for the severity. An interesting feature of Chernobai *et al.* (2006)'s paper is that the left-truncation of the PCS insured losses data is accounted for — it must be borne in mind that the PCS indices only record losses in excess of a certain threshold. Lin and Chang (2008) posited that a Poisson process with constant intensity is inadequate for modelling PCS insured losses, and motivate the use of a Markov Modulated Poisson process instead. Ma and Ma (2013) model the PCS data from 1985 to 2010 in a similar fashion to Chernobai *et al.* (2006), however, they do not account for the left-truncation in the data. In fact, Ma and Ma (2013)'s calibration and estimation takes no cognisance of the changing threshold levels (or equivalently, left-truncation).

In summary, modelling the PCS loss index via some marked Poisson process (or even a compound renewal process) seems to have gained popularity in recent research articles. However, no robust Monte Carlo simulation methodology, in the context of simulating the PCS index (being modelled as a compound Poisson process, for example), has been brought to the fore. Moreover, it is apparent from Ma and Ma (2013)'s paper that there was no attempt to robustly simulate very large losses from the underlying heavy-tailed distributions (and actually ensure that they form part of a Monte Carlo sample). Note that, however, Monte Carlo simulation for the purposed of CAT bond pricing was considered by Vaugirard (2003b), but not in the context of the PCS data.

2.4 Pricing of catastrophe bonds

The pricing of CAT bonds (and other catastrophe instruments) is complex – this is mainly because these securities do not possess unique prices based upon arbitrage-free pricing techniques. The states of the world reflected in such securities are not mirrored by those states reflected by more traditional securities (i.e. more primitive assets, such as bonds and stocks, cannot span the payoffs from CAT bonds) (Cox *et al.*, 2000). Therefore, exact replication is impossible (Embrechts and Meister, 1997), and the admissability of no-arbitrage arguments and replication strategies is brought into question.

On consideration of the above, it seems clear as to why there have been discussions around a number of methods for pricing CAT bonds. To the best of our knoweldge, six differing theoretical frameworks have been postulated in the literature (Galeotti *et al.*, 2013; Braun, 2011). Firstly, standard actuarial pricing methodologies have been applied to CAT bonds by Lane (2000), Lane and Mahul (2008) and Bodoff and Gan (2009), all following from the work of Bühlmann (1984, 1985). Actuarial pricing refers to the usual pricing of insurance risks via real-world premium principles. The loss distribution is estimated via real-world probabilities, and a premium is often calculated on the basis of the moments of the distribution. Often, a margin for uncertainty is added to the estimated parameters of the distribution, and this appears to pervade in CAT bond pricing in practice (Braun, 2011). Secondly, an econometric CAT bond pricing model, based on observable catastrophe-related quantities, was proposed by Braun (2016) for pricing in the primary market.

Thirdly, utility-based approaches have been proposed by Embrechts and Meister (1997); Cox and Pedersen (2000); Reshetar (2008); Egami and Young (2008) and Dieckmann (2010) for numerous types of CAT bonds. Johnson *et al.* (2004), however, highlight general difficulties with the utility theory approach, in that it is demanding on the inputs required in asset pricing in general. Moreover, if such CAT bond pricing is nestled within a utility maximisation framework, Embrechts and Meister (1997) posit that a unique risk-neutral pricing measure can emerge in a natural way.

Fourthly, frameworks based on partial equilibrium and equilibrium pricing theory have been put forward. Aase (1999) employed a partial equilibrium framework with constant absolute risk aversion wherein CAT risk was regarded as systemic. Cox and Pedersen (2000) also selected equilibrium pricing theory and included a time separable utility function, while Christensen and Schmidli (2000a) introduced an exponential utility model within a similar framework. Aase (2001) extended his earlier work to a competitive equilibrium approach with constant relative risk aversion. The model of Cox and Pedersen (2000) was applied to pricing earthquake-related CAT bonds, in a discrete time framework, in the doctoral thesis of Shao (2015).

But quite rapidly, the pre-2001 theories coalesced around a fifth approach; that is a preference-free no-arbitrage framework. Much of the motivation for such an approach stemmed from the general incomplete market framework considered in Merton (1976). In these models, natural catastrophes are treated as idiosyncratic risks that can be (almost) fully diversified. Subsequent pricing methods have assumed that the risk neutral probability measure has been predetermined, usually by assuming that CAT risk is orthogonal to market risk, thus avoiding complicated changes of measure. This assumption is supported by the empirical studies of Hoyt and McCullough (1999) and Cummins and Weiss (2009), and disputed in Carayannopoulos and Perez (2015a) and Gürtler *et al.* (2016). Under an assumed risk-neutral probability measure, stochastic processes and distributions used to price CAT bonds retain the same characteristics as under the real world probability measure. Unfortunately, calibration of these models seems all but impossible due to a lack of pricing data. Many of the models simply propose a plausible index process, and then employ simulation in order to price and hedge CAT derivatives. However, those papers which have successfully managed to use real-world data (such as insurance losses, and not CAT bond prices) to estimate CAT bond prices (such as Ma and Ma (2013) and Nowak and Romaniuk (2013)) have reinterpreted Merton (1976)'s arguments as saying that catastrophe variables have the same distributional characteristics under the real-world and risk-neutral probability measures - this makes sense, especially in the context of Cox and Pedersen (2000)'s assumption in that aggregate consumption only depends upon financial markets risk processes. Let us now briefly review these no-arbitrage models.

Baryshnikov *et al.* (1998) undertook the first attempt to price catastrophe bonds within a no-arbitrage framework. The authors foresaw a growth in the CAT bond market, with near-continuous trading of these instruments in the future. The authors contended that because this was the case, CAT bond prices would reflect no-arbitrage or fair bond prices. They further argued that the risk-neutral and real-world probability measures would coincide[7] – this assumption significantly simplified their work by firstly allowing pricing to commence in the risk-neutral world directly and by secondly omitting the inclusion of investors' in CAT bonds risk preferences. With all of this in mind, a no-arbitrage pricing model for CAT bonds, employing a compound Poisson process to represent the accumulated losses resulting from catastrophes, was used. This pivotal work was able to accommodate for the findings of Barton and Nishenko (1997), in that certain catastrophic natural events possess power-law distributions associated with their loss (severity) distributions. This model was extended in Burnecki and Kukla (2003), and used by Dassios and Jang (2003b) to model stop-loss reinsurance contracts and CAT derivatives. Thereafter, various extensions of these initial works were applied to CAT futures, options, bonds, and swaps. See, for example, the developments in Geman and Yor (1997); Loubergé *et al.* (1999); Lee and Yu (2002); Bakshi and Madan (2002); Vaugirard (2003a,b, 2004); Muermann (2001, 2003, 2008); Zimbidis *et al.* (2007); Biagini *et al.* (2008b,a); Egami and Young (2008); Lin and Wang (2009a); Chang *et al.* (2010); Härdle and Cabrera (2010); Jarrow (2010); Braun (2011); Ma and Ma (2013); Nowak and Romaniuk (2013) and Jaimungal and Chong (2014). Upon reviewing these various articles, what is noticed is the difference in structure of all the catastrophe instruments (corroborating the fact that catastrophe instruments are mostly traded over-the-counter), thereby leading to slight differences in these arbitrage-free valuation frameworks.

[7] This assumption is similar to the arguments, in the CAT bond literature, based on the work of Merton (1976).

We now briefly discuss the sixth and final framework. The concerns surrounding incomplete markets and non-unique measures can be overcome by a judicious choice of change in measure, such as the Esscher transform (see the discussion in Gerber and Shiu (1996)). Also, in the context of insurance premium calculation on the basis of a non-homogeneous compound Poisson process, Delbaen and Haezendonck (1989) provided a martingale approach to premium calculation, in an arbitrage-free market and under a unique probability measure - in fact, their approach encompassed the Esscher transform. Another approach, similar to this, is the Wang transform[8] introduced by Wang (2000). Wang (2004) extended this previous work and applied it to CAT bond pricing. He demonstrated that the Wang transform extends the Sharpe ratio to credit-related risks characterised by skewed loss distributions. A universal pricing formula was devised, and it was consequently calibrated to a set of CAT bond prices from 1999. Valdez (2005) presented elliptical transforms with application in insurance premium calculation principles. Such elliptical transformations were shown to lead to the Wang and Wang Student-t distortions, as well as the Esscher premium principle. But Pelsser (2008) criticised the Wang transform, and showed that certain stochastic processes prices based on the Wang transform were not consistent with arbitrage-free prices, despite the measure being unique. He concluded that the Wang transform could not be a universal framework for the pricing of insurance and financial risks (but it may be an approximation). Finally, one could resort to extending the complete market option pricing framework to incomplete markets. Instead of eliminating the risk by a perfect hedging portfolio, these strategies invariably involve partial replication, with a tolerance for some residual risk. Such strategies usually generate a profit-and-loss process. An example of such a strategy has been the classical risk minimization approach of Föllmer and Schweizer (1986) (further developed in Föllmer and Schweizer (1989), Schweizer (1991, 1999), and Møller (2001)), which minimises fluctuations in the discounted profit-and-loss processes by using a quadratic criterion under an assumed minimal equivalent martingale measure.

Before we end, it is important to note that mixtures of the above procedures have also been used to price catastrophe-related insurance contracts or ILS, especially in what is called "engineering pricing". Tao *et al.* (2009) bring to the fore a model for earthquakes, based upon seismic risk assessment. They go on to describe the cashflows of earthquake CAT bond premiums, in complete and incomplete markets, by using geometric Brownian motion and jump diffusion processes. Also, Aslan *et al.* (2011) proposes a four-step engineering-based loss model for earth-

[8] Note that Labuschagne and Offwood (2010) found a connection between the Esscher transform and the Wang transform.

quakes. The study's results show a non-linear relationship between earthquake-related and financial parameters. It moreover integrates engineering-based, actuarial and financial stances into a unified pricing framework.

It is evident from our review of the CAT bond pricing literature that, in the context of most pricing rules (especially those based on no-arbitrage considerations), collective risk models such as the aggregate loss processes (ALP) – or loss distribution approach (LDA) models – (see Panjer and Willmot (1992) and Peters and Shevchenko (2015)) prevail as underlying mechanisms to capture the catastrophe risk. Indeed, this resonates what was said in Embrechts and Meister (1997), in that any insurance linked-security involving claim payments will possess either a compound Poisson process, a mixed compound Poisson process or a doubly-stochastic compound Poisson process as an important feature in the modelling exercise. This is especially true in the context of index-linked catastrophe bonds, as evinced, for example, in Ma and Ma (2013). Most work in CAT bond pricing assumes that the ALP is a compound Poisson process. However, we stress that even though such a model can add to intuition, it should not substitute actuarial judgement and insight.

2.5 Other insurance-linked securities linked to this study

The market for ILSs is large, as evinced in Section 2.1. Since many of the ILSs are designed to be particular to their issuers and investors, the market is also exceptionally varied, with many of these ILSs trading over-the-counter. It would be difficult to record each-and-every type of ILS here, so we rather describe those ILSs (different from CAT bonds) which will be relevant to this study.

We firstly consider catastrophe futures and options, of which most are linked to the PCS index (broadly called "PCS options"), were traded on the Chicago Board of Exchange from 1992 until 2000 (Braun, 2011). Such cash-settled instruments are derivatives whose payoffs depend on property and casualty insurance-related losses arising from natural catastrophes. (Re-)insurance companies would purchase them as a form of protection from large losses. Such derivatives have been studied by Cummins and Geman (1995b); Chang *et al.* (1996); Schradin (1996); Möller (1996); O'Brien (1997); Aase (1999, 2001) and Cox *et al.* (2004a), amongst others. According to Braun (2011), catastrophe derivatives focusing only on US hurricane risk have been re-launched by several exchanges, some of those exchanges being the Insurance Futures Exchange, European Exchange, Chicago Mercantile Exchange and the New York Mercantile Exchange.

In 1996, a new ILS innovation came to the fore: the first ever catastrophe-equity put (CAT-E-put) was issued. CAT-E-puts give their holder the right, but not the obligation, to issue convertible shares at a fixed price, should the accumulated losses of the holder exceed some pre-defined critical coverage limit (or threshold) during the lifetime of the option (Jaimungal and Wang, 2006b). Such a derivative structure provides its holder with the opportunity to raise additional funds to pay catastrophic losses, *ex-post* the catastrophes. However, like catastrophe futures and options, their market is also extinct (Wang, 2016). CAT-E-puts have attracted little scholarly attention, and have only been treated in the literature by Jaimungal and Wang (2006b), Chang *et al.* (2011) and Wang (2016).

Another interesting ILS which has attracted very little scholarly attention to date is the catastrophe swap. A catastrophe swap is an over-the-counter traded financial instrument wherein the protection buyer (fixed-rate payer) makes periodic payments to the protection seller (floating-rate payer) in exchange for a predefined binary loss-compensation payment contingent upon the occurrence of a predefined trigger event. According to a press-release by the World Bank, they recently organised a $206 million catastrophe swap for the government of the Philippines. Despite there being only mention of catastrophe swaps in early articles pertaining to ILSs, the sole paper to comprehensively analyse it is that by Braun (2011).

Chapter 3

Catastrophe instrument valuation model

In this chapter, we build up a probabilistic framework to form the base for CAT bond pricing. In doing so, we loosely follow the frameworks of Cox and Pedersen (2000), Ma and Ma (2013) and Shao (2015): we set up a risk-neutral probability measure for general CAT bond pricing, then apply this in the special case of index-linked CAT bonds. But in addition to this, we present a discussion on the various pricing measures at work within the framework. Let us begin with constructing the risk-neutral probability measure.

3.1 Constructing a combined product space from two individual spaces

We commence under the real-world probability measure. Under any probability measure the CAT bond's price depends on two emerging phenomena: financial market-related risk and catastrophe-related risk. Since the catastrophe-risk variables will give rise to jumps, we need to work in an incomplete markets setting and moreover note that complicated changes of measure could arise. To avoid this, we make the following assumption in line with much of the previous literature on pricing catastrophe-linked financial instruments. Evidence in support of this assumption has been found by Hoyt and McCullough (1999) and Cummins and Weiss (2009), but is disputed by Carayannopoulos and Perez (2015b) and Hagendorff *et al.* (2015).

Assumption 3.1. Catastrophe-risk variables and financial markets risk variables are independent in the real-world.[1]

This assumption is made in a myriad of research papers on catastrophe-linked instruments, including Baryshnikov *et al.* (1998), Cox and Pedersen (2000), Jarrow (2010), Braun (2011), Ma and Ma (2013) and Nowak and Romaniuk (2013). Assumption 3.1 affords us the possibility to treat catastrophe-risk variables independently from financial markets risk variables. Therefore, we can split up the CAT bond pricing into two separate problems under both the real-world and, later, under the risk-neutral probability measure.

Now, the following method is a convenient way to set up the required probability space for the model. We suppose that there exist two probability spaces : for the financial markets risk variables, the space is specified by $(\Omega_F, \hat{\mathscr{F}}_\infty, \mathbb{P}_F)$, where $\hat{\mathscr{F}}_\infty := \bigvee_{t\geq 0} \hat{\mathscr{F}}_t$ for the partial financial markets filtration $\mathbb{F} := (\hat{\mathscr{F}}_t)_{t\geq 0}$. Also, Ω_F is the respective sample space and $\mathbb{P}_F$ is the real-world probability measure for the financial markets risk variables. For the catastrophe-risk variables, the space is given by $(\Omega_C, \hat{\mathscr{C}}_\infty, \mathbb{P}_C)$, where $\hat{\mathscr{C}}_\infty := \bigvee_{t\geq 0} \hat{\mathscr{C}}_t$ for the partial catastrophe-risk filtration $\mathbb{C} := (\hat{\mathscr{C}}_t)_{t\geq 0}$. Also, Ω_C denotes the respective sample space and $\mathbb{P}_C$ denotes the real-world probability measure for the catastrophe markets risk variables. From these probability spaces, we can construct a product space $(\Omega, \mathscr{G}_\infty, \mathbb{P})$, where $\Omega := \Omega_F \times \Omega_C$, $\mathscr{G}_\infty := \hat{\mathscr{F}}_\infty \otimes \hat{\mathscr{C}}_\infty$, $\mathbb{P} := \mathbb{P}_F \otimes \mathbb{P}_C$ and also $\mathscr{G}_t := \hat{\mathscr{F}}_t \otimes \hat{\mathscr{C}}_t$. Notice how Assumption 1 is conveniently captured in the definition of $\mathbb{P}$.

We continue by introducing the following important assumption, key notation and a consequent note on terminology. We assume that both $\hat{\mathscr{C}}_0$ and $\hat{\mathscr{F}}_0$ are trivial. Now, two filtrations pertaining to financial markets risk variables and catastrophe risk variables respectively, are defined on the product space, Ω:

$$\mathscr{F}_t := \hat{\mathscr{F}}_t \otimes \hat{\mathscr{C}}_0, \qquad \mathscr{C}_t := \hat{\mathscr{F}}_0 \otimes \hat{\mathscr{C}}_t.$$

In consequence we say that a stochastic process Y, defined on $(\Omega, \mathscr{G}_\infty)$, depends on financial markets risk variables if Y is adapted to the filtration $\mathscr{F}_t$, on the product space. Moreover, we say that a stochastic process X, defined on $(\Omega, \mathscr{G}_\infty)$, depends on catastrophe-risk variables if X is adapted to the filtration $\mathscr{C}_t$, on the product space.

We now need to consider in depth the translation between the individual measurable spaces, $(\Omega_C, \hat{\mathscr{C}}_\infty)$ and $(\Omega_F, \hat{\mathscr{F}}_\infty)$, and the combined measurable product space $(\Omega, \mathscr{G}_\infty)$. We need to demonstrate that all random variables defined on this product space which are $\mathscr{F}_t$-measurable can be associated uniquely with a random variable, that is defined on $(\Omega_F, \hat{\mathscr{F}}_\infty)$, which is $\hat{\mathscr{F}}_t$-measurable. The same also needs to be demonstrated for those random variables which are $\mathscr{C}_t$-measurable. Now, observe that there exist two measurable projections, π_F and π_C:

$$(\Omega, \mathscr{G}_\infty) \xrightarrow{\pi_F} (\Omega_F, \hat{\mathscr{F}}_\infty) : (\omega_F, \omega_C) \mapsto \omega_F \tag{3.1}$$

$$(\Omega, \mathscr{G}_\infty) \xrightarrow{\pi_C} (\Omega_C, \hat{\mathscr{C}}_\infty) : (\omega_F, \omega_C) \mapsto \omega_C \tag{3.2}$$

respectively. A random variable $\hat{X} : (\Omega_F, \hat{\mathscr{F}}_\infty) \xrightarrow{\hat{X}} \mathbb{R}$ can be lifted to the product space via composition with the projection π_F. This can be achieved by defining the random variable, X, such that

$$(\Omega, \mathscr{G}_\infty) \xrightarrow{X} \mathbb{R} : (\omega_F, \omega_C) \mapsto \hat{X}(\omega_F) \qquad \text{i.e.} \quad X = \hat{X} \circ \pi_F.$$

In a similar fashion, a random variable $\hat{Y} : (\Omega_F, \hat{\mathscr{C}}_\infty) \xrightarrow{\hat{Y}} \mathbb{R}$ can be lifted to the product space via composition with the projection π_C, and this can be achieved by defining the random variable Y such that

$$(\Omega, \mathscr{G}_\infty) \xrightarrow{Y} \mathbb{R} : (\omega_F, \omega_C) \mapsto \hat{Y}(\omega_C) \qquad \text{i.e.} \quad Y = \hat{Y} \circ \pi_C.$$

Observe that $(\Omega, \mathscr{F}_t) \xrightarrow{\pi_F} (\Omega_F, \hat{\mathscr{F}}_t)$ is $\mathscr{F}_t$-measurable. Hence, X is $\mathscr{F}_t$-measurable if-and-only-if $\hat{X}$ is $\hat{\mathscr{F}}_t$-measurable. On the grounds of similar arguments, we can also say that Y is $\mathscr{C}_t$-measurable if-and-only-if $\hat{Y}$ is $\hat{\mathscr{C}}_t$-measurable.

In addition, suppose that X is $\mathscr{F}_t$-measurable (for $t \leq \infty$), that $\omega_C^1, \omega_C^2 \in \Omega_C$, and that $c_1 := X(\omega_F, \omega_C^1)$, $c_2 := X(\omega_F, \omega_C^2)$. Then, $(\omega_F, \omega_C^1) \in X^{-1}\{c_1\} \in \mathscr{F}_t$, so $X^{-1}\{c_1\} = \hat{F} \times \Omega_C$ for some $\hat{F} \in \hat{\mathscr{F}}_t$. But then $(\omega_F, \omega_C^2) \in X^{-1}\{c_1\}$ as well, which tells us that $c_1 = c_2$. Thus, we can define a $\hat{\mathscr{F}}_t$-measurable random variable $\hat{X}$ on Ω_F so that $X = \hat{X} \circ \pi_F$, i.e.

$$\hat{X}(\omega_F) := X(\omega_F, \omega_C) \qquad \text{for any } \omega_C \in \Omega_C.$$

Hence, there is a one-to-one correspondence between $\mathscr{F}_t$-measurable random variables on Ω and $\hat{\mathscr{F}}_t$-measurable random variables on Ω_F. On the basis of similar arguments, there also exists a one-to-one correspondence between $\mathscr{C}_t$-measurable random variables on Ω and $\hat{\mathscr{C}}_t$-measurable random variables on Ω_C.

We now show that distributional properties of random variables as well as independence between random variables on the individual spaces are transferred across to the product spaces.

Lemma 3.2. *It holds that* $\mathscr{F}_\infty \perp\!\!\!\perp_{\mathbb{P}} \mathscr{C}_\infty$.

Proof. Consider $B, C \in \mathscr{B}(\mathbb{R})$. Then,

$$\begin{aligned}\mathbb{P}(X \in B, Y \in C) &= \mathbb{P}_F \otimes \mathbb{P}_C \left(X^{-1}[B] \cap Y^{-1}[C]\right) \\ &= \mathbb{P}_F \otimes \mathbb{P}_C \left(\hat{X}^{-1}[B] \times \hat{Y}^{-1}[C]\right) \\ &= \mathbb{P}_F \left(\hat{X} \in B\right) \mathbb{P}_C \left(\hat{Y} \in C\right).\end{aligned}$$

If we take $C = \mathbb{R}$, we obtain that $\mathbb{P}(X \in B) = \mathbb{P}_F \left(\hat{X} \in B\right)$, so this shows that X has the same distribution under $\mathbb{P}$ as does $\hat{X}$ under $\mathbb{P}_F$. Similar considerations apply for the random variable Y. It then follows that

$$\mathbb{P}(X \in B, Y \in C) = \mathbb{P}(X \in B)\,\mathbb{P}(Y \in C),$$

so that X and Y are independent under $\mathbb{P}$. Thus, $\mathscr{F}_\infty \perp\!\!\!\perp_{\mathbb{P}} \mathscr{C}_\infty$. □

The following proposition provides an important result: it shows that Assumption 3.1 is transferred to the combined measurable product space. Define $\mathscr{F}_\infty := \bigvee_{t\geq 0} \mathscr{F}_t$ and $\mathscr{C}_\infty := \bigvee_{t \geq 0} \mathscr{C}_t$.

Proposition 3.3. *Suppose that X is $\mathscr{F}_\infty$-measurable and Y is $\mathscr{C}_\infty$-measurable. Then, it holds that*

$$\mathbb{E}_{\mathbb{P}}[XY|\mathscr{G}_t] = \mathbb{E}_{\mathbb{P}}[X|\mathscr{F}_t]\,\mathbb{E}_{\mathbb{P}}[Y|\mathscr{C}_t].$$

Proof. Observe that

$$\mathbb{E}_{\mathbb{P}}[XY|\mathscr{G}_t] = (\mathbb{E}_{\mathbb{P}_F}[\hat{X}|\hat{\mathscr{F}}_t] \circ \pi_F) \cdot (\mathbb{E}_{\mathbb{P}_C}[\hat{Y}|\hat{\mathscr{C}}_t] \circ \pi_C).$$

Indeed, $(\mathbb{E}_{\mathbb{P}_F}[\hat{X}|\hat{\mathscr{F}}_t] \circ \pi_F) \cdot (\mathbb{E}_{\mathbb{P}_C}[\hat{Y}|\hat{\mathscr{C}}_t] \circ \pi_C)$ is $\mathscr{G}_t$-measurable, so it suffices to show that

$$\int_G (\mathbb{E}_{\mathbb{P}_F}[\hat{X}|\hat{\mathscr{F}}_t] \circ \pi_F) \cdot (\mathbb{E}_{\mathbb{P}_C}[\hat{Y}|\hat{\mathscr{C}}_t] \circ \pi_C) \mathrm{d}\mathbb{P} = \int_G XY \mathrm{d}\mathbb{P}$$

for every $G \in \mathscr{G}_t$. By a monotone class argument, it suffices to show this for sets of the form $G = \hat{F} \times \hat{C}$, where $\hat{F} \in \hat{\mathscr{F}}_t$, $\hat{C} \in \hat{\mathscr{C}}_t$, since sets of this form generate $\mathscr{G}_t$. Now, by the Fubini-Tonelli Theorems (see Billingsley (1995, p. 233-234))

$$
\begin{aligned}
&\int_G (\mathbb{E}_{\mathbb{P}_F}[\hat{X}|\hat{\mathscr{F}}_t] \circ \pi_F) \cdot (\mathbb{E}_{\mathbb{P}_C}[\hat{Y}|\hat{\mathscr{C}}_t] \circ \pi_C) \mathrm{d}\mathbb{P} \\
&= \int_{\hat{F}} \int_{\hat{C}} \mathbb{E}_{\mathbb{P}_F}[\hat{X}|\hat{\mathscr{F}}_t](\omega_F) \cdot \mathbb{E}_{\mathbb{P}_C}[\hat{Y}|\hat{\mathscr{C}}_t](\omega_C) \mathbb{P}_F(\mathrm{d}\omega_F) \mathbb{P}_C(\mathrm{d}\omega_C) \\
&= \left(\int_{\hat{F}} \mathbb{E}_{\mathbb{P}_F}[\hat{X}|\hat{\mathscr{F}}_t](\omega_F) \mathbb{P}_F(\mathrm{d}\omega_F) \right) \left(\int_{\hat{C}} \mathbb{E}_{\mathbb{P}_C}[\hat{Y}|\hat{\mathscr{C}}_t](\omega_C) \mathbb{P}_C(\mathrm{d}\omega_C) \right) \\
&= \left(\int_{\hat{F}} \hat{X}(\omega_F) \mathbb{P}_F(\mathrm{d}\omega_F) \right) \left(\int_{\hat{C}} \hat{Y}(\omega_C) \mathbb{P}_C(\mathrm{d}\omega_C) \right) \\
&= \int_{\hat{F}} \int_{\hat{C}} \hat{X}(\omega_F) \hat{Y}(\omega_C) \mathbb{P}_F(\mathrm{d}\omega_F) \mathbb{P}_C(\mathrm{d}\omega_C) \\
&= \int_{\hat{F} \times \hat{C}} XY \, \mathrm{d}(\mathbb{P}_F \otimes \mathbb{P}_C).
\end{aligned} \tag{3.3}
$$

In particular, taking $Y = 1$ in Equation (3.3), we obtain

$$\mathbb{E}_{\mathbb{P}}[X|\mathscr{G}_t] = \mathbb{E}_{\mathbb{P}_F}[\hat{X}|\hat{\mathscr{F}}_t] \circ \pi_F. \tag{3.4}$$

Since the right-hand side of Equation (3.4) is $\mathscr{F}_t$-measurable, by conditioning both sides on $\mathscr{F}_t$ we see that

$$\mathbb{E}_{\mathbb{P}}[X|\mathscr{G}_t] = \mathbb{E}_{\mathbb{P}}[X|\mathscr{F}_t].$$

An analogous result holds if we consider $X = 1$ in Equation (3.3). Hence the result.

□

3.2 Constructing a combined risk-neutral product space

As in the case of many types of ILSs a CAT bond is not an insurance contract but rather a financial instrument, so it is to be priced using financial pricing techniques. As put forward in Cox *et al.* (2004a), if a liquid and large market for catastrophe-linked securities (for example, CAT bonds) exists, then standard derivatives pricing theory (see, for example, Harrison and Pliska (1981)) implies the existence of a risk-neutral measure, so ILSs such as CAT bonds can be priced. However, since CAT bonds rely on a process exhibiting jumps[2], the market is incomplete and, hence, no unique risk-neutral measure exists (Embrechts, 2000). So we proceed as follows.

In order to define a risk-neutral pricing measure for the financial markets risk variables on the space $(\Omega_F, \hat{\mathscr{F}}_\infty)$, we need to introduce the instantaneous short rate which defines the bank account and hence the notion of discounting. The instan-

[2] Index-linked securities are based on insurance loss indices such as the PCS index. Such indices are often modelled as jump processes, and we return to this point later on in the section.

taneous short rate process and bank account process are defined in Definitions 3.4 and 3.5. As a precursor to later developments within our pricing model, note that we shall adopt the bank account as the numéraire asset.

Definition 3.4. (Instantaneous short rate or risk-free interest rate). We call the progressively measurable[3] process, $\{r_t, t \geq 0\}$, the *instantaneous short rate* or *risk-free interest rate*.

Definition 3.5. (Bank account). Let A_t be the value of a riskless *bank account* at time $t \geq 0$. The bank account evolves according to the stochastic differential equation

$$\mathrm{d}A_t = r_t A_t \mathrm{d}t, \quad A_0 = 1,$$

which is solved by

$$A_t = \exp\left(\int_0^t r_u \mathrm{d}u\right).$$

We now assume the existence of a given risk-neutral pricing measure for the financial markets risk variables, $\mathbb{Q}_F$ on $(\Omega_F, \hat{\mathscr{F}}_\infty)$, which has been obtained from suitable calibrations of interest-rate and stock-price processes[4]. $\mathbb{Q}_F$ has the properties that (i) it is equivalent to $\mathbb{P}_F$ and (ii) discounted securities with payoffs based only on financial markets risk variables are $(\mathbb{Q}_F, \mathbb{F})-$martingales. Now, let

$$\hat{\xi}_\infty := \frac{\mathrm{d}\mathbb{Q}_F}{\mathrm{d}\mathbb{P}_F},$$

and let

$$\hat{\xi}_t := \mathbb{E}_{\mathbb{P}_F}[\hat{\xi}_\infty | \hat{\mathscr{F}}_t]$$

be the associated Radon-Nikodým derivative and likelihood process (for $t \geq 0$). Hence, using the projection, π_F, from Equation (3.1), we can define

$$\xi_\infty := \hat{\xi}_\infty \circ \pi_F$$

with its associated likelihood process (for $t \geq 0$)

$$\xi_t := \mathbb{E}_{\mathbb{P}}[\xi_\infty | \mathscr{F}_t].$$

[3] With a little abuse of notation, a stochastic process X is called progressively measurable or progressive if $(\omega, s) \mapsto X(\omega, s)$, $(\omega, s) \in \Omega \times [0, t]$ is $\mathscr{G}_t \otimes \mathscr{B}[0, t]-$measurable, where $\mathscr{G}_t$ is the filtration of the process up to time t, and $\mathscr{B}$ denotes the Borel algebra (Filipovic, 2009).

[4] We will specify (in further detail) the stock price process in Chapter 6.

Note that $\mathbb{E}_{\mathbb{P}}[\xi_\infty|\mathscr{F}_t] = \mathbb{E}_{\mathbb{P}_F}[\hat{\xi}_\infty|\hat{\mathscr{F}}_t] \circ \pi_F$. To see this, consider the following argument. Let $F \in \mathscr{F}_t$. Since $\hat{\xi}_t \circ \pi_F$ is $\mathscr{F}_t$-measurable, we must show that

$$\int_F \hat{\xi}_t \circ \pi_F \, \mathrm{d}\mathbb{P} = \int_F \hat{\xi}_\infty \circ \pi_F \, \mathrm{d}\mathbb{P}. \tag{3.5}$$

Since $F \in \mathscr{F}_t$, $F = \hat{F} \times \Omega_C$ for some $\hat{F} \in \hat{\mathscr{F}}_t$. Consider the left-hand side of Equation (3.5) and apply the Fubini-Tonelli theorems:

$$\begin{aligned}
\int_F \hat{\xi}_t \circ \pi_F \, \mathrm{d}\mathbb{P} &= \int_{\hat{F} \times \Omega_C} \hat{\xi}_t \circ \pi_F \, \mathrm{d}(\mathbb{P}_F \otimes \mathbb{P}_C) \\
&= \int_{\Omega_C} \int_{\hat{F}} \hat{\xi}_t(\pi_F(\omega_F, \omega_C)) \, \mathrm{d}(\mathbb{P}_F) \mathrm{d}(\mathbb{P}_C) \\
&= \int_{\hat{F}} \hat{\xi}_t(\omega_F) \, \mathrm{d}(\mathbb{P}_F).
\end{aligned}$$

The right-hand side of Equation (3.5) is identified as $\int_{\hat{F}} \hat{\xi}_\infty(\omega_F)\,\mathrm{d}(\mathbb{P}_F)$. By noting that $\hat{\xi}_t = \mathbb{E}_{\mathbb{P}_F}[\hat{\xi}_\infty|\hat{\mathscr{F}}_t]$ and using the partial averaging property of the conditional expectation (see Jacod and Protter (2003, Theorem 23.3)), we see that the left-hand side of Equation (3.5) equals the right-hand side. Hence, we have lifted ξ_t from the individual measurable space, $(\Omega_F, \hat{\mathscr{F}}_\infty)$, to the product measurable space, $(\Omega, \mathscr{G}_\infty)$. Now, the following definition is important.

Definition 3.6. (Zero coupon bond). The price of a *zero-coupon bond* with face value 1 and maturity time T is specified by $P(t,T)$ for $0 \le t \le T$. We moreover require that $P(T,T) = 1$.

Given the assumed existence of the risk-neutral pricing measure, $\mathbb{Q}_F$, we have that

$$P(t,T) = \mathbb{E}_{\mathbb{Q}_F}\left[\exp\left(-\int_t^T r_u \mathrm{d}u\right) | \hat{\mathscr{F}}_t\right].$$

So, the discounted bond price process $\left\{\frac{P(t,T)}{A_t}, 0 \le t \le T\right\}$ is a $(\mathbb{Q}_F, \mathbb{F})-$martingale, and A_t is the numéraire for $\mathbb{Q}_F$. More generally, by definition of $\mathbb{Q}_F$, the time t price of a financial claim, which depends only on financial markets risk variables, with time T payoff X_T^F, is given by

$$X_t^F = A_t \mathbb{E}_{\mathbb{Q}_F}\left[\frac{X_T^F}{A_T} | \hat{\mathscr{F}}_t\right]. \tag{3.6}$$

The next question pertains to the associated pricing measure for the catastrophe-risk variables. Since the catastrophe-risk variables will be assumed to follow a jump process, we will have several choices (Dassios and Jang, 2003a). We, however, attempt to remain general and suppose that the CAT bond investor prices using a probability measure, $\mathbb{Q}_C$, on $(\Omega_C, \hat{\mathscr{C}}_\infty)$ (equivalent to $\mathbb{P}_C$) defined by

$$\frac{d\mathbb{Q}_C}{d\mathbb{P}_C} := \hat{\beta}_\infty$$

where $\hat{\beta}_\infty$ is a strictly positive $\hat{\mathscr{C}}_\infty-$measurable random variable satisfying $\mathbb{E}_{\mathbb{P}_C}[\hat{\beta}_\infty] = 1$. Moreover, define the associated Radon-Nikodým derivative and likelihood process (for $t \geq 0$),

$$\hat{\beta}_t := \mathbb{E}_{\mathbb{P}_C}[\hat{\beta}_\infty | \hat{\mathscr{C}}_t]. \tag{3.7}$$

Note that $\hat{\beta}_\infty$ can, possibly, be chosen on the basis of some actuarial pricing framework, such as the Esscher transform (see, for example, Gerber and Shiu (1996)), the Wang transform (see Wang (2000) and Wang (2004)), or premium calculation principles based on the mean and variance of relevant insurance loss (see Goovaerts *et al.* (1984)). But also, we can consider $\hat{\beta}_\infty = 1$ - indeed, this is a case we consider further on in Section 3.3 below. Note that the purpose of $\hat{\beta}_\infty$ is to incorporate the CAT bond investor's risk preferences to price nature (or catastrophe) risk. In a similar manner as we did for $\hat{\xi}_\infty$ before, we can lift $\hat{\beta}_\infty$ to the combined product space, $(\Omega, \mathscr{G}_\infty)$, using Equation (3.2). Define

$$\beta_\infty = \hat{\beta}_\infty \circ \pi_C, \tag{3.8}$$

and its associated likelihood process (for $t \geq 0$) as

$$\beta_t := \mathbb{E}_{\mathbb{P}}[\beta_\infty | \mathscr{C}_t]. \tag{3.9}$$

Note that $\mathbb{E}_{\mathbb{P}}[\beta_\infty | \mathscr{C}_t] = \mathbb{E}_{\mathbb{P}_C}[\hat{\beta}_\infty | \hat{\mathscr{C}}_t] \circ \pi_C$. Moreover, notice that ξ_∞ and β_∞ are independent under $\mathbb{P}$.

On the basis of the aforementioned, it is evident that a risk adjustment has been included to obtain a pricing measure, $\mathbb{Q}_C$, for catastrophe risk. Here, $\mathbb{Q}_C$ is such that the time t price of a catastrophe claim, which depends on only catastrophe-risk variables, with payoff X_T^C is given by

$$X_t^C = P(t,T)\mathbb{E}_{\mathbb{Q}_C}[X_T^C | \hat{\mathscr{C}}_t]. \tag{3.10}$$

We are now in a position to define a product space coupled with its own associated pricing measure. Such a measure will, also, be in line with Assumption 3.1.

So, on the measurable space $(\Omega, \mathscr{G}_\infty)$, define the measure $\mathbb{Q}$, equivalent to $\mathbb{P}$, as

$$\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} := \xi_\infty \beta_\infty$$

where $\xi_t \beta_t := \mathbb{E}_{\mathbb{P}}[\xi_\infty \beta_\infty | \mathscr{G}_t]$. Now, if $\hat{F} \in \hat{\mathscr{F}}_\infty$ and $\hat{C} \in \hat{\mathscr{C}}_\infty$

$$\begin{aligned}
\mathbb{Q}(\hat{F} \times \hat{C}) &= \mathbb{E}_{\mathbb{P}}[\mathbb{I}_{\hat{F} \times \hat{C}} \xi_\infty \beta_\infty] \\
&= \mathbb{E}_{\mathbb{P}}\left[\left((\mathbb{I}_{\hat{F}} \hat{\xi}_\infty) \circ \pi_F\right) \cdot \left((\mathbb{I}_{\hat{C}} \hat{\beta}_\infty) \circ \pi_C\right)\right] \\
&= \mathbb{E}_{\mathbb{P}}\left[(\mathbb{I}_{\hat{F}} \hat{\xi}_\infty) \circ \pi_F\right] \mathbb{E}_{\mathbb{P}}\left[(\mathbb{I}_{\hat{C}} \hat{\beta}_\infty) \circ \pi_C\right] \\
&= \mathbb{E}_{\mathbb{P}_F}[\mathbb{I}_{\hat{F}} \hat{\xi}_\infty] \mathbb{E}_{\mathbb{P}_C}[\mathbb{I}_{\hat{C}} \hat{\beta}_\infty] \\
&= \mathbb{Q}_F(\hat{F}) \mathbb{Q}_C(\hat{C}).
\end{aligned}$$

Hence, we have that $\mathbb{Q} = \mathbb{Q}_F \otimes \mathbb{Q}_C$. We now go on to demonstrate that the independence properties assumed for the random variables defined on the real-world individual spaces, as well as conditional expectations, are preserved under this new combined product-measure, $\mathbb{Q}$. Notice that Lemma 3.2 and Proposition 3.3 can be viewed as applying to any arbitrary product measure. In consequence, we obtain the following Lemma and Proposition.

Lemma 3.7. *It holds that* $\mathscr{F}_\infty \perp\!\!\!\perp_{\mathbb{Q}} \mathscr{C}_\infty$.

Proposition 3.8. *Suppose that X is $\mathscr{F}_\infty$-measurable and Y is $\mathscr{C}_\infty$-measurable. Then, it holds that*

$$\mathbb{E}_{\mathbb{Q}}[XY|\mathscr{G}_t] = \mathbb{E}_{\mathbb{Q}}[X|\mathscr{F}_t]\,\mathbb{E}_{\mathbb{Q}}[Y|\mathscr{C}_t].$$

We also obtain an important, but helpful, corollary to Proposition 3.8.

Corollary 3.9. *With the same notation as in Proposition 3.8, the following hold:*

(i) $\mathbb{E}_{\mathbb{Q}}[X|\mathscr{G}_t] = \mathbb{E}_{\mathbb{Q}}[X|\mathscr{F}_t] \quad \forall t > 0;$ *and*

(ii) $\mathbb{E}_{\mathbb{Q}}[Y|\mathscr{G}_t] = \mathbb{E}_{\mathbb{Q}}[Y|\mathscr{C}_t] \quad \forall t > 0.$

Proof. Part (i) follows by setting $Y = 1$ in Proposition 3.8, while part (ii) follows by setting $X = 1$. □

We now go on to show how the pricing measure, $\mathbb{Q}$, can be considered to be a risk-neutral probability measure. Let $X_T^{F,C}$ denote a $\mathscr{G}_T-$measurable payoff at time T, i.e. a payoff which *may* depend on both financial and catastrophe variables.

Define the time t price of this claim by

$$X_t := A_t \mathbb{E}_{\mathbb{Q}} \left[\frac{X_T^{F,C}}{A_T} | \mathscr{G}_t \right].$$

By definition, the process $\left\{ \frac{X_T^{F,C}}{A_t}, 0 \le t \le T \right\}$ is a $(\mathbb{Q}, \mathbb{G})-$martingale, and hence $\mathbb{Q}$ is a risk-neutral pricing measure.

We now check if such a definition is consistent with Equations (3.6) and (3.10). If $X_T^{F,C}$ is purely financial (i.e. $X_T^{F,C}$ is $\mathscr{F}_T-$measurable), then

$$\begin{aligned} X_t^{F,C} &:= A_t \mathbb{E}_{\mathbb{Q}} \left[\frac{X_T^{F,C}}{A_T} | \mathscr{G}_t \right] \\ &= A_t \mathbb{E}_{\mathbb{Q}_F} \left[\frac{X_T^F}{A_T} | \hat{\mathscr{F}}_t \right] \circ \pi_F, \end{aligned}$$

where X_T^F is the corresponding $\hat{\mathscr{F}}_t$-measurable random variable such that $X_T^{F,C} = X_T^F \circ \pi_F$. Therefore, the price is consistent with the financial market price in Equation (3.6). Next, if $X_T^{F,C}$ is purely catastrophic (i.e. $X_T^{F,C}$ is $\mathscr{C}_T-$measurable), then

$$\begin{aligned} X_t^{F,C} &:= A_t \mathbb{E}_{\mathbb{Q}} \left[\frac{X_T^{F,C}}{A_T} | \mathscr{G}_t \right] \\ &= \mathbb{E}_{\mathbb{Q}} \left[\frac{A_t}{A_T} | \mathscr{C}_t \right] \mathbb{E}_{\mathbb{Q}} \left[X_T^{F,C} | \mathscr{C}_t \right] \\ &= P(t,T) \mathbb{E}_{\mathbb{Q}_C} \left[X_T^C | \hat{\mathscr{C}}_t \right] \circ \pi_C, \end{aligned}$$

where X_T^C is the corresponding $\hat{\mathscr{C}}_t$-measurable random variable such that $X_T^{F,C} = X_T^C \circ \pi_C$. Therefore, the price is consistent with the financial market price in Equation (3.10). Now that we have a risk-neutral probability measure on the product space, we are ready to proceed with valuing CAT bonds.

We end this section with a critical theorem, which gives us the opportunity to price a catastrophe-related contingent claim under real-world probability measures. This theorem will be very useful in pricing catastrophe instruments via the calibration of a compound Poisson process to real-world data. It will, however, be up to the user of the catastrophe instrument to select a particular form of the Radon-Nikodým derivative for the catastrophe-risk variables: such a selection could depend on, *inter alia*, how high the user would like the probabilities of catastrophe-risk variables (in relation to their probabilities under the real-world probability measure) to be. Let $\{V_t, 0 \le t \le T\}$ denote the catastrophe-related contingent-claim (i.e. catastrophe instrument) pricing process, with payoff at fixed time $T > 0$.

Theorem 3.10. *Consider the random variable $\tilde{C}_T : \Omega \mapsto \mathbb{R}$, which is $\mathscr{C}_T-$measurable, representing the payoff (dependent on only catastrophe-risk variables) of a catastrophe instrument at time T for $T > 0$. Then*

$$V_t = P(t,T)\beta_t^{-1}\mathbb{E}_{\mathbb{P}}[\tilde{C}_T\beta_T|\mathscr{C}_t] \tag{3.11}$$

is an arbitrage-free price for the catastrophe instrument, where $\beta_.$ is the Radon-Nikodym derivative process as specified in Equation (3.9).

Proof. Under the measure $\mathbb{Q}$,

$$V_t = \mathbb{E}_{\mathbb{Q}}\left[\exp\left(-\int_t^T r_s \mathrm{d}s\right)\tilde{C}_T|\mathscr{G}_t\right]$$

is an *arbitrage-free* price for the catastrophe instrument. Moreover,

$$\begin{aligned}
\mathbb{E}_{\mathbb{Q}}\left[\exp\left(-\int_t^T r_s \mathrm{d}s\right)\tilde{C}_T|\mathscr{G}_t\right] &= \xi_t^{-1}\beta_t^{-1}\mathbb{E}_{\mathbb{P}}\left[\exp\left(-\int_t^T r_s \mathrm{d}s\right)\tilde{C}_T\xi_T\beta_T|\mathscr{G}_t\right] \\
&= \xi_t^{-1}\beta_t^{-1}\mathbb{E}_{\mathbb{P}}\left[\exp\left(-\int_t^T r_s \mathrm{d}s\right)\xi_T|\mathscr{F}_t\right]\mathbb{E}_{\mathbb{P}}\left[\tilde{C}_T\beta_T|\mathscr{C}_t\right] \\
&= \mathbb{E}_{\mathbb{Q}}\left[\exp\left(-\int_t^T r_s \mathrm{d}s\right)|\mathscr{F}_t\right]\beta_t^{-1}\mathbb{E}_{\mathbb{P}}\left[\tilde{C}_T\beta_T|\mathscr{C}_t\right] \\
&= P(t,T)\beta_t^{-1}\mathbb{E}_{\mathbb{P}}\left[\tilde{C}_T\beta_T|\mathscr{C}_t\right],
\end{aligned}$$

where the second-last line follows by Proposition 3.3. □

3.3 Minding $\mathbb{P}$ and $\mathbb{Q}$

In this section, we provide a link between the real-world and risk-neutral probability measures. Given that we did not have access to CAT bond pricing data, our aim will be to justify why we use the real-world pricing measure for catastrophe-risk variables and set $\hat{\beta}_\infty = 1$ for the catastrophe-risk variables.

We begin by considering the incomplete market framework of Merton (1976). Such a framework has been used extensively in the literature when valuing derivatives with payoffs linked to natural catastrophes, see for example Bakshi and Madan (2002), Lee and Yu (2002), Vaugirard (2003a), Jaimungal and Wang (2006a), Lee and Yu (2007), Ma and Ma (2013), Nowak and Romaniuk (2013) and Chang and Chang (2017). On the grounds of its pervasiveness in the literature to date, the following assumption is made in our work, and we use it extensively.

Assumption 3.11. Investors are risk-neutral towards the risk posed by the natural catastrophe-risk variables.

More fully and in the context of pricing financial instruments, Assumption 3.11 states that in the overall economy natural catastrophes can be treated as idiosyncratic risks that can be (almost) fully diversified. The catastrophe risks will pose "non-systematic risk" and will, in consequence, carry a zero risk-premium. This can all be put in line with analysis offered in Section 3.2 by setting $\hat{\beta}_\infty = 1$. Therefore, the risk-neutral probability measure for the catastrophe-risk variables will coincide with the respective real-world probability measure[5], $\mathbb{P}_C$. Moreover, note that if $\hat{\beta}_\infty = 1$, then $\beta_\infty = 1$ on the basis of Equation (3.8).

However, it must be borne in mind that recent empirical catastrophe bond pricing literature has shown that catastrophe bonds do not have a zero risk premium (see, for example, Papachristou (2009), Braun (2016) and Gürtler *et al.* (2016)); this may carry over to other catastrophe-linked ILS instruments as well. Against this backdrop, it is possible to infer that pricing models based on the zero risk-premium assumption may give rise to values higher than those given by pricing models which assume a non-zero risk premium. In consequence, the usage of these pricing models may require additional margins added to the calculated value, or margins added to the parameters of the distributions associated with the jump process, all at the discretion of the issuer. Despite this, we remain true to Assumption 3.11 in our work for two reasons. Firstly, it is commensurate with many actuarial pricing techniques which according to Braun (2011) prevail in practice. But secondly, and most importantly, it can be adapted to the scope of underlying state variables, in the model, which are not investment assets and hence not tradeable. Hence, we can use real-world data to price, and this is useful given the scarcity of (and difficulty of obtaining) pricing data for many catastrophe-linked ILSs.

Therefore, we are in a position to state a corollary to Theorem 3.10.

Corollary 3.12. *With the same notation as in Theorem 3.10, if $\beta_\infty \equiv 1$ then $\tilde{C}_T$ carries a zero catastrophe-risk premium and*

$$V_t = P(t,T)\mathbb{E}_{\mathbb{P}}[\tilde{C}_T|\mathscr{C}_t] \tag{3.12}$$

is an arbitrage-free price for the catastrophe instrument.

[5] In fact, the formal declaration of a risk-neutral probability measure for catastrophe-risk variables is not necessary.

Proof. Since $\beta_\infty \equiv 1$, we have that $\beta_t = 1$ for all $t \geq 0$. The result then follows immediately from Theorem 3.10. □

Before we proceed with index-linked CAT bond (and the subsequent contingent-convertible catastrophe bond) valuation treated in Chapter 6), we present the following important remark on notation.

Remark. We now make the following identification, and in consequence sometimes abuse notation. We will identify every $\mathscr{F}_t$-measurable random variable $X : \Omega \mapsto \mathbb{R}$ with the corresponding $\hat{\mathscr{F}}_t$-measurable random variable $\hat{X} : \Omega_F \mapsto \mathbb{R}$. Similarly, we will identify every $\mathscr{C}_t$-measurable random variable $Y : \Omega \mapsto \mathbb{R}$ with the corresponding $\hat{\mathscr{C}}_t$-measurable random variable $\hat{Y} : \Omega_C \mapsto \mathbb{R}$.

3.4 Index-linked CAT bond valuation

We begin this section with a brief recapitulation on the structure of an index-linked (IL) CAT bond, since this is one of the key areas of focus in this thesis. IL CAT bonds are a type of CAT bond written on an industry loss index, such as the PCS loss index. In such a CAT bond, its payoff is triggered when estimated industry-wide losses from a catastrophic event, or sequence thereof, exceeds a predefined threshold level. IL CAT bonds can be of two types - zero-coupon, or coupon-paying. For the purposes of our research and the CAT bond types we endeavour to price, we suppose that the coupons are constant. Moreover, for the purposes of CAT bond valuation, we suppose that the short rate is constant. A modelled short rate process is, however, considered in Chapter 6. We now proceed to introduce some notation, which also will be commonly used throughout Chapters 4 and 5. Let:

- $T > 0$ denote the fixed term of the IL CAT bond.
- Z be the principal amount invested in the IL CAT bond, which the investor will receive back should the bond not trigger during its term. For the sake of simplicity, we suppose that $Z = 1$.
- V_0^{ZC} denote the price of the IL zero coupon CAT bond at issue date, $t_0 := 0$. V_0^{CP} denotes the value in the case of the IL coupon-paying bond case.
- $\{L_t, t \geq 0\}$ be an aggregate loss process capturing the behaviour of the index

upon which the IL CocoCat is based, where

$$L_t = \sum_{k=1}^{N_t} X_k. \tag{3.13}$$

The aggregate loss process is assumed to have a frequency component specified by a renewal process $N = \{N_t, t \geq 0\}$. We define a renewal process, $\{N_t, t > 0\}$, as

$$N_t = \max\left\{ n; \sum_{k=1}^{n} T_k \leq t \right\},$$

where the inter-arrival times $\{T_k, k \in \mathbb{N}\}$ are independent and identically-distributed (i.i.d.) positive random variables each with $\mathbb{E}[T_k] = 1/\gamma$, for some $\gamma > 0$, and with $\mathrm{Var}[T_k] = s_P^2$, where $s_P > 0$ (see Grandell (2012)). As a special case, one can suppose that N is a Poisson process[6] with deterministic time-dependent intensity specified by the real-valued function λ_t. Moreover, we have a sequence of i.i.d. severity-component random variables $\{X_k, k \in \mathbb{N}\}$ (independent of the frequency component), each with distribution function F_X and density f_X (which is assumed to exist).

- $\mathbb{T} := \{t_1, t_2, \ldots, t_{J-1}, t_J = T\}$, $J \in \mathbb{N}$, denote the set of J coupon-paying dates.
- $r > 0$ be the constant short rate.
- $\tau = \inf\{0 \leq t \leq T : L_t \geq D\}$, the first time the trigger level is met or exceeded. D is called the contractually-specified threshold level of the IL CAT bond in the zero-coupon case. For the coupon-paying case, we refer to the threshold level as D_{CP}.
- $c > 0$ be the constant coupon rate.
- $0 \leq \rho \leq 1$ be the constant recovery rate on the bond, should it be triggered.

We now consider the two payoff structures for CAT bonds. The first we consider is an index-linked zero-coupon (ZC) CAT bond, having maturity time $T > 0$. The structure of the payoff, C_T, of the IL ZC CAT bond is given by:

$$C_T = \begin{cases} 1 & \text{if } L_T < D \\ \rho & \text{if } L_T \geq D, \end{cases} \tag{3.14}$$

[6] This will then lead to the aggregate loss process being a compound Poisson process.

Theorem 3.13. *An arbitrage-free price at time* 0, V_0^{ZC}, *of the IL ZC CAT bond is given by*

$$V_0^{ZC} = e^{-rT}\left[\rho + (1-\rho)\mathbb{P}_C\left(L_T < D\right)\right] \tag{3.15}$$

assuming that L_t *follows a compound renewal process.*

Proof. By an application of Corollary 3.12,

$$\begin{aligned} V_0^{ZC} &= P(0,t)\mathbb{E}_{\mathbb{P}_\mathbb{C}}[\mathbb{I}_{\{L_T<D\}} + \rho\mathbb{I}_{\{L_T\geq D\}}], \\ &= e^{-rT}[\mathbb{P}_\mathbb{C}(L_T<D) + \rho\mathbb{P}_\mathbb{C}(L_T\geq D)] \\ &= e^{-rT}[\mathbb{P}_\mathbb{C}(L_T<D) + \rho(1-\mathbb{P}_\mathbb{C}(L_T<D))] \end{aligned}$$

and the result follows. □

Corollary 3.14. *If* L_t *is assumed to follow a compound Poisson process with a time-dependent intensity specified by* λ_t, *then* V_0^{ZC} *is given by Equation (3.15) with*

$$\mathbb{P}_C\left(L_T < D\right) = \sum_{n=0}^{\infty} e^{-\int_0^T \lambda_u \mathrm{d}u} \frac{\left(\int_0^T \lambda_u \mathrm{d}u\right)^n}{n!} F_X^{n*}(D)$$

where F_X^{n*} *denotes the* $n-$*fold convolution of* F_X *with itself*[7].

In the case of an IL coupon-paying (CP) CAT bond, for each coupon paying date $t \in \{t_1, t_2, ..., t_k = T\}$, the payoff, C_t, per unit nominal is

$$C_t = \begin{cases} c + \mathbb{1}_{\{t=T\}} & \text{if } L_t < D_{\mathrm{CP}} \\ \rho c + \rho\,\mathbb{1}_{\{t=T\}} & \text{if } L_t \geq D_{\mathrm{CP}}. \end{cases} \tag{3.16}$$

Observe that the coupon-paying bond has both its coupons and redemption amount written down by ρ should the threshold be exceeded. At time 0 the price, V_0^{CP}, of the index-linked CP CAT bond having maturity time $T > 0$ is given by Theorem 3.15 below.

Theorem 3.15. *An arbitrage-free price at time* 0, V_0^{CP}, *of the IL CP CAT bond is given by*

$$\begin{aligned} V_0^{CP} = &\sum_{i=1}^{J} c e^{-rt_i}\left[\rho + (1-\rho)\mathbb{P}_C\left(L_{t_i} < D_{CP}\right)\right] \\ &+ e^{-rT}\left[\rho + (1-\rho)\mathbb{P}_C\left(L_T < D_{CP}\right)\right]. \end{aligned} \tag{3.17}$$

assuming that L_t *follows a compound renewal process.*

[7] Note that the $n-$fold convolution is given by $\mathbb{P}_\mathbb{C}(X_1 + X_2 + ... + X_n < D)$ for each $n \in \mathbb{N}$ and for $n = 0$ is defined to be 1.

Proof. Since a coupon-paying bond can be stripped into individual zero-coupon bonds, the statement of the theorem follows by the repeated application of Theorem 3.10 to each coupon-paying date. □

Corollary 3.16. *If L_t is assumed to follow a compound Poisson process with a time-dependent intensity specified by λ_t, then V_0^{CP} is given by Equation (3.17) with*

$$\mathbb{P}_C\left(L_{t_i} < D_{CP}\right) = \sum_{n=0}^{\infty} e^{-\int_0^{t_i} \lambda_u \mathrm{d}u} \frac{\left(\int_0^{t_i} \lambda_u \mathrm{d}u\right)^n}{n!} F_X^{n*}(D_{CP})$$

where F_X^{n*} denotes the n–fold convolution of F_X with itself.

As can be observed from Equation (3.17) and its predecessors, the coupon-paying times are general. However, many index-linked CAT bonds pay quarterly coupons and we assume this for the purposes of the numerical illustrations in Chapters 4 and 5.

Chapter 4

Modelling of left-truncated catastrophe data with application to catastrophe bond pricing

4.1 Background

The modelling of extremal events – or, synonymously, the monetary losses in consequence of such events – is necessary in many different contexts. In the insurance markets, such an example is the attempt to model the threats posed by natural catastrophes, that is, catastrophe risk. Catastrophe risks can be viewed as high-severity, low-probability risks, and most resultant data will certainly be heavy-tailed in nature.

Today, insurers and reinsurers are seeking more ways to offer protection from catastrophe risk, specifically property risk arising from natural catastrophes. The development of the ILS market has offered insurers an avenue to do so (Burnecki *et al.*, 2000; Cox and Pedersen, 2000; Burnecki, Kukla and Taylor, 2005). In this chapter, we consider index-linked catastrophe bonds and how we can satisfactorily model the principle underlying process driving their monetary value. In particular, we magnify into a specific type of CAT bond, namely PCS Index-linked CAT bonds (Burnecki, Kukla and Taylor, 2005; Ma and Ma, 2013). Since such CAT bonds reference the left-truncated PCS loss index[1], we emphasise that there is a need to account for the left-truncation feature in the data.

In what follows, we examine the effect of firstly ignoring and secondly including this left-truncation when modelling PCS index-linked CAT bond prices, using the pricing rule outlined in Chapter 3 and under the assumption that catastrophe losses (which are one of the ultimate drivers of CAT bond prices) follow a time-inhomogeneous compound Poisson process. We primarily perform such an exer-

[1] Recall from Chapter 2 that the PCS index does not consider losses less than $25 million.

cise because explicitly accounting for left-truncation in the modelling can potentially lead to a better estimation of CAT bond prices. Better estimation techniques may, invariably, give rise to improved valuation techniques for all parties interested in pricing such index-linked CAT bonds.

In an effort to more precisely model the left-truncated data, we build on the methodology developed by Chernobai *et al.* (2006). Now, it must be noted that Chernobai *et al.* (2006) did not explicitly account for the heavy-tailed nature of the severity data to which they applied their fitting methodology. Primarily, the possibility of unbounded likelihoods was not addressed. That is, there may exist paths in the parameter space along which the likelihood shoots to infinity and maximum likelihood estimation (MLE), in consequence, breaks down. In a general context, the existence of such a phenomenon has been shown to occur in research by Juárez and Schucany (2004) and Ergashev *et al.* (2016). These authors also commented that it may be the case that maximum likelihood estimates can be obtained, but converge at a slower rate compared to classical MLE. But importantly, it was pointed out by Wong and Li (2006) that such a problem exists particularly in the case of heavy-tailed distributions such as the generalised Pareto (GP) and generalised extreme value distributions (GEV), which are often used to model heavy-tailed data.

Therefore, in this chapter we consider the problem of modifying Chernobai *et al.* (2006)'s methodology to make it more applicable to the modelling of heavy-tailed left-truncated data. That is why we consider the parameter estimation technique of the Maximum Product of Spacings (MPS), which does not suffer from the aforementioned problems[2], and apply it in the context of estimating the parameters of the the Burr type XII[3], GP and GEV distributions[4]. Note that recent evidence has shown that MPS performs better than MLE when it comes to considering the bias, mean-squared error and variance of estimated GEV parameters (see Soukissian and Tsalis (2015)).

The MPS estimation method of parameter estimation for continuous univariate distributions was developed independently by Cheng and Amin (1983) and Ranneby (1984), and it does not always give rise to the same parameter estimates found on the basis of MLE. It has been shown to be a robust parameter estimation technique, especially in the context of certain power laws, by Goldstein *et al.* (2004) and Wong and Li (2006). The intuition behind the MPS method of estimation lies in the probability integral transform: any sample of independent and identically-

[2] In fact, MPS does not even require the existence of a probability density or mass function.

[3] We will, loosely speaking, sometimes refer to the Burr type XII distribution as simply the Burr distribution

[4] Other distributions considered in this chapter are fitted via maximum likelihood estimation, since MPS is not necessary given bounded likelihood functions.

distributed realisations should be uniformly distributed with respect to the cumulative distribution function of the random variable from which they are sampled from. As a natural consequence of the MPS method of estimation, a goodness-of-fit test statistic was found by Cheng and Stephens (1989), which they called Moran's log spacing statistic. Unlike the more traditional goodness-of-fit tests which require Monte Carlo simulation to calculate their $p-$values (when parameters are estimated from the data), a key advantage of using this statistic is that no such simulation is necessary when the relevant parameters of the fitted distribution are estimated from the data.

In view of the above, this chapter is organised as follows. In Section 4.2 we present the methodology for modelling heavy-tailed left-truncated data in the compound time-inhomogeneous Poisson process framework. Specifically, we present the insurance risk process which will play a key role in the modelling of left-truncated data. We also outline and implement our parametric distribution-fitting methodology to fit the severity component of the compound Poisson process, based on a modified MPS estimation technique. In Section 4.3 we apply the methodology to the PCS index data. We also compare our approach to a straightforward, if naive, fitting approach. Additionally, we present a simple importance sampling technique which is useful in sampling from heavy-tailed distributions with only finite first moment, and apply it to the Burr and generalised extreme value distributions. In Section 4.4, we apply our fitting methodology within the practical context of valuing CAT bonds. We also go on to examine the effects on the CAT bond prices of ignoring the left truncation in the fitting of the insurance risk process. It must be noted that Sections 4.3 and 4.4 form the basis for many of our mathematical and numerical analyses in Chapters 5 and 6. Section 5.7 concludes and presents some practical findings, which may be of use for CAT bond market participants, and also useful for those researchers who have datasets with similar characteristics to ours.

4.2 Identification and validation of the compound Poisson process

4.2.1 Modified compound Poisson process

In the insurance literature, random summation models (for example, compound Poisson processes) have been used to model claims[5]. These models have two main components: one component characterises the frequency of an event giving rise to

[5] However, random summation models also have applications in other fields such as in operational loss and credit-default modelling in banking.

the loss, while the other the severity of the loss. Given that we ultimately wish to model the behaviour of the PCS loss index in this chapter and subsequent ones, we consider an extension to the classical aggregate loss process Panjer and Willmot (1992) to model left-truncated data:

$$L_t = \sum_{k=1}^{N_t} X_k \mathbb{I}_{\{X_k \geq H\}},\ t \geq 0 \quad (L_0 = 0) \tag{4.1}$$

where $X_1, X_2, \ldots$ is a sequence of i.i.d. random variables representing the loss severities, each having the distribution function $F(x) = \mathbb{P}(X \leq x)$, and H is the threshold level signifying the point of left-truncation. From here on, we refer to Equation (4.1) as the insured loss process Chernobai *et al.* (2006). The number of losses in the interval $(0, t]$ is driven by the point process $N = \{N_t, t \geq 0\}$, which is commonly called the loss arrival process and is independent of the loss severities sequence. We model the process N as a time-inhomogeneous Poisson process with a deterministic intensity function given by $\lambda(t)$. The distribution of N_t is given by

$$\mathbb{P}(N_t = n) = \frac{e^{-\int_0^t \lambda(s)\mathrm{d}s} \left(\int_0^t \lambda(s) ds\right)^n}{n!} \quad n = 0, 1, \ldots. \tag{4.2}$$

4.2.2 Parameter estimation in the context of heavy-tailed data

It must be borne in mind that we wish to (i) fit a severity distribution, defined on the non-negative real numbers, in order to fully characterise the process shown in Equation (4.1). In addition, we wish to (ii) model the arrival of all losses, and not just those in excess of the threshold. However, the data we have to fit such processes to is left-truncated. Therefore, we mention that there are three approaches to modelling left-truncated data which are possibly in line with our wishes, (i) and (ii) above. The first approach is based on Figure 4.1 (a) – the "naive" approach. Under the "naive" approach, the researcher treats the data set as complete, ignores the threshold level and fits a distribution to the data. Moreover, the observed frequency (estimated from the data) is treated as the true frequency. Clearly, this approach is unsatisfactory. The second possible approach to the distribution fitting (illustrated in Figure 4.1 (b) is the so-called "shifting approach", whereby the data are shifted to the left by the threshold amount, the shifted data are fitted to a non-truncated distribution and the fitted distribution is thereafter translated back to the right. It must be noted that this approach is problematic from a catastrophe (or even insured loss) modelling perspective since it ignores observations below the threshold level. Moreover, the act of ignoring small observations will not assist in adjusting

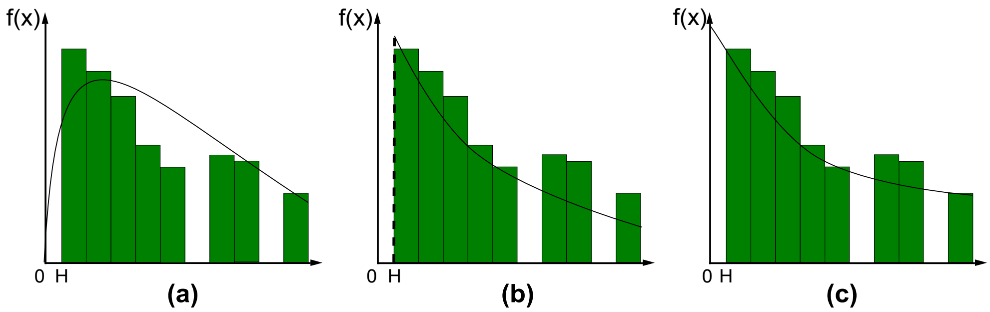

Fig. 4.1: Illustrations of (a) the density function for the losses obtained via the fitting approach ignoring H (i.e. the naive approach), (b) the truncated density function and (c) the unconditional complete-data density function under the conditional approach Chernobai *et al.* (2006).

the intensity estimates (for the arrivals of the losses) so that this adjustment accommodates for losses less than the threshold level. The third approach is based on Figure 4.1 (c) – the conditional complete-data approach (CCD approach) introduced by Chernobai *et al.* (2006). The CCD approach correctly specifies both the loss distribution and the frequency by observing that the data are only recorded above the threshold level H. It recovers full information about the distribution of the aggregate loss process. We use the CCD approach in this chapter.

We now briefly present the mechanics of the CCD approach, which, broadly speaking, directly fits the right tail of the distribution and scales up the frequency. Moreover, we discuss the potential refinements to this methodology, which make it more applicable to modelling heavy-tailed data.

In the context of (4.1), let $F_\gamma(x)$ be the (real-world) distribution function of each loss (arising from a natural catastrophe, for example), with its corresponding probability density function denoted by $f_\gamma(x)$.

We now specify the estimation of the underlying loss distribution. The observed sample in the time interval $[T_1, T_2]$ is of the form

$$\mathbf{x} = (x_1, x_2, \ldots, x_n) \in (H, \infty), \tag{4.3}$$

where n denotes the number of observations taking values in (H, ∞) and $x_1, x_2, \ldots, x_n$ denotes the values of these observations. It must be borne in mind that only the losses in excess of the threshold H are observed. Therefore, it becomes necessary to discern between the observed data set – that is, the observations in excess of H – and the complete data set – that is, all losses be they above or below H in magnitude. Parameters estimated using the observed losses will exhibit a superscript "O" in their notation, while those parameters to be used for modelling the unconditional complete-data loss process will exhibit the superscript "C".

There are two approaches to estimate the parameters of the severity distribution, MLE and MPS. Each approach is discussed below.

Maximum likelihood estimation (MLE)

Let the sample space for the total losses be denoted by $\mathscr{X}$. Therefore, the likelihood function of the observations $(\mathbf{x})$ on $\mathscr{X}$, with respect to the product of Lebesgue (for the severity) and counting measures (for the frequency), is given by

$$g_{\mathbf{x}}^{\text{MLE}}(\mathbf{x}, \lambda^{O}(t), \gamma^{C}) = \frac{1}{n!}\left(\int_{T_1}^{T_2} \lambda^{O}(s)\mathrm{d}s\right)^{n} e^{-\int_{T_1}^{T_2} \lambda^{O}(s)\mathrm{d}s} \prod_{k=1}^{n} \frac{f_{\gamma^C}(x_k)}{1 - F_{\gamma^C}(H)} \tag{4.4}$$

where F_{γ^C} denotes the unconditional complete-data (true) distribution function, f_{γ^C} its corresponding density function and γ^C its corresponding parameters. Equation (4.4) is written on the assumption that parameters can be consistently estimated by MLE.

In Equation (4.4), $F_{\gamma^C}(H)$ represents the probability of a random observation falling into the time interval $[0, H]$. Also, observe the usage of $\lambda^{O}(t)$ – this considers the time-inhomogeneous Poisson process, N_T^O, which counts only the losses with magnitudes in excess of H over the time interval $(0, T]$. This can be interpreted as a thinning of the complete Poisson process N_T^C and therefore N_T^O has intensity $\lambda^{O}(t) = \left[1 - F_{\gamma^C}(H)\right] \lambda^{C}(t)$. $\lambda^{C}(t)$ refers to the intensity of the complete Poisson process N_T^C.

Notice the independence between the frequency and the severity of the losses, conferred by the assumption that the insured loss process follows a compound Poisson process, be it time-homogeneous or time-inhomogeneous. Therefore, the estimation of the parameters γ^C of the distribution function can be done independently of the estimation of $\lambda^{O}(t)$ and so, this conveniently allows for the estimation of the parameters γ^C. If MLE estimation is required, we take natural logarithms and thereafter maximise the log-likelihood function of (4.4) with respect to the vector γ^C and obtain the maximum-likelihood estimate of γ^C as

$$\widehat{\gamma^C_{MLE}} = \underset{\gamma^C}{\operatorname{argmax}} \left[\log\left(\prod_{k=1}^{n} \frac{f_{\gamma^C}(x_k)}{1 - F_{\gamma^C}(H)}\right)\right]. \tag{4.5}$$

Equation (4.5) provides a method whereby estimates for the parameters of the distribution of X_i, $F_{\gamma^C}(x)$, but not for the distribution of N_t, can be found. The superscript "C" attached to $\widehat{\gamma_{MLE}}$ is used to specify that these are the parameters of the unconditional complete-data distribution F. Ergashev *et al.* (2016) warns against potentially unbounded likelihood functions when applying parameter estimation techniques such as MLE.

Maximum product of spacings estimation (MPS)

We now combine the CCD approach with the MPS estimation technique in order to find estimators for the parameters of the severity distribution, $\widehat{\gamma^C}$. With the same notation as in Section 4.2.2 the joint density on $\mathscr{H}$, for the data sample $x_1, x_2, \ldots x_n$ with corresponding order statistics $x_{(1)}, x_{(2)}, \ldots x_{(n)}$, is given by

$$g_{\mathbf{x}}^{\text{MPS}}(\mathbf{x}, \lambda^O(t), \gamma^C) = \frac{1}{n!}\left(\int_{T_1}^{T_2} \lambda^O(s)\mathrm{d}s\right)^n e^{-\int_{T_1}^{T_2}\lambda^O(s)\mathrm{d}s} \prod_{k=1}^{n}\left(\frac{F_{\gamma^C}(x_{(k)}) - F_{\gamma^C}(x_{(k-1)})}{1 - F_{\gamma^C}(H)}\right), \tag{4.6}$$

where $F_{\gamma^C}(x_{(n+1)}) := 1$ and $F_{\gamma^C}(x_{(0)}) := 0$. Now, the MPS parameter estimator for Equation (4.6), which we denote by $\widehat{\gamma^C_{MPS}}$, is the value of the vector γ^C which minimises the MPS function:

$$\mathscr{M}(\gamma^C) = -\sum_{k=1}^{n+1} \log\left(\frac{F_{\gamma^C}(x_{(k)}) - F_{\gamma^C}(x_{(k-1)})}{1 - F_{\gamma^C}(H)}\right), \tag{4.7}$$

That is, the estimator is

$$\widehat{\gamma^C_{MPS}} = \underset{\gamma^C}{\text{argmin}}\left[-\sum_{k=1}^{n+1} \log\left(\frac{F_{\gamma^C}(x_{(k)}) - F_{\gamma^C}(x_{(k-1)})}{1 - F_{\gamma^C}(H)}\right)\right]. \tag{4.8}$$

One potential problem which may arise with the MPS estimation technique is the treatment of existing tied observations in the sample of order statistics. Cheng and Amin (1983), however, show that the maximum product of spacings estimator can still be found in this instance. Therefore, we caution that it must be checked if there are any tied observations in the data set. See Cheng and Amin (1983) for further details.

Finally, the intensity for the observed Poisson process is obtained by fitting a (chosen) deterministic function to the aggregate number of events per unit interval over the time interval $[T_1, T_2]$. A periodogram of the frequency can be used to intuitively ascertain which deterministic function should be used for the intensity rate and the function can be fitted using non-linear least squares. The intensity for the complete Poisson process can thereafter be obtained by setting $\widehat{\lambda^C(t)} = \widehat{\lambda^O(t)} / \left[1 - \widehat{F}_{\gamma^C}(H)\right]$, where the unknown value $\widehat{F}_{\gamma^C}(H)$ is estimated from using the parameters found in (4.5) or (4.8).

4.2.3 Goodness of fit tests

Once a severity distribution has been fitted to the PCS loss data, it is necessary to determine whether or not it is actually appropriate to use. Goodness-of-fit tests are useful criteria, complimentary to a decision-making framework, for deciding whether or not a fitted distribution fits the data-set well. However, they need to be adapted to suit the left-truncated nature of the data. Goodness-of-fit tests for truncated and censored data have been studied by, amongst others, Dufour and Maag (1978), Gastaldi and Gastaldi (1992) and Guilbaud (1988). However, Moran's log spacing statistic has not applied in the context of left-truncated data.

Supremum and quadratic class goodness-of-fit tests

The first set of tests we use in this chapter are two supremum class test statistics – the Kolmogorov–Smirnov (KS) and Kuiper (V) tests – and two quadratic class test statistics – the Anderson-Darling (AD^2) and Cramér von Mises (W^2) tests. However, two problems arise. The first is that the parameters of the fitted distributions are estimated from the data, and so the limiting distributions for the various goodness-of-fit test statistics cannot be used (Ross, 2002). The second is that the data is left-truncated, and so the conventional goodness-of-fit test statistics for testing distributions fitted via the CCD approach need to be modified to reflect this feature of the data.

We now consider the first problem. We find ourselves in a situation where we would like to test the null hypothesis that the sample has a common distribution function $F_{\gamma^C}(x)$ with unknown γ^C. To make use of any of the goodness-of-fit tests based on the empirical distribution function, we need to estimate the parameters, γ^C, from the data. However, in such a situation (i.e. when parameters are estimated from the data), the critical values for these tests must be reduced (Burnecki, Misiorek and Weron, 2005). Ross (2002) supports the use of Monte Carlo simulations within this context, and furthermore Capasso *et al.* (2009) clearly outline the procedure to be used in such goodness-of-fit tests. Following Ross (2002), Capasso *et al.* (2009) and Burnecki, Misiorek and Weron (2005), we firstly calculate the test statistic, say d (be it a KS, V, AD^2 or W^2 test statistic), based on the estimated parameter vector $\widehat{\gamma^C}$ for the assumed distribution F. Secondly we generate, by Monte

Carlo simulation, a sufficiently large number (N) of independently-drawn samples from the distribution $F_{\widehat{\gamma^c}}(x)$ each of size M. For each of these M-sized samples, we re-estimate the parameter vector and using these re-estimated parameters calculate a single observation of the test statistic. Finally, we estimate the p-value as the proportion of times that the simulated test statistics exceed or are equal to d.

The second problem has already been considered by Chernobai *et al.* (2015). In their work, they derived an adjusted empirical distribution function (EDF) which, under the null hypothesis, depends on the fitted probability distribution function evaluated at the threshold level H. From this adjusted EDF, Chernobai *et al.* (2015) derived adjusted goodness-of-fit test statistics for the KS, V, AD^2 or W^2 tests. Table 4.1 illustrates the difference between the adjusted goodness-of-fit test statistics for left-truncated data samples versus the conventional goodness-of-fit test statistics for full or complete data samples (i.e. if no left-truncation is present).

Tab. 4.1: Comparison of EDF test statistics for full data samples versus left-truncated (LT) data samples. Note that $F_n(x)$ refers to the EDF of the data sample, $z_j = \widehat{F}_{\gamma^C}(x_{(j)})$ and $z_H = \widehat{F}_{\gamma^C}(H)$.

Test	Sample	Statistic	Computing formula
KS	*Full*	$\sqrt{n} \cdot \sup_x \lvert F_n(x) - \widehat{F}(x) \rvert$	$\sqrt{n} \cdot \max \left\{ \sup_j \left\{ \frac{j}{n} - z_j \right\}, \sup_j \left\{ z_j - \frac{j-1}{n} \right\} \right\}$
	LT	$\sqrt{n} \cdot \sup_x \lvert F_n(x) - \widehat{F^*}(x) \rvert$	$\frac{\sqrt{n}}{1-z_H} \cdot \max \left\{ \sup_j \left\{ z_H + \frac{j}{n}(1-z_H) - z_j \right\}, \sup_j \left\{ z_j - \left(\frac{j-1}{n}(1-z_H) \right) \right\} \right\}$
V	*Full*	$\sqrt{n} \left(\sup_x \left\{ F_n(x) - \widehat{F}(x) \right\} + \sup_x \left\{ \widehat{F}(x) - F_n(x) \right\} \right)$	$\sqrt{n} \left(\sup_j \left\{ \frac{j}{n} - z_j \right\} + \sup_j \left\{ z_j - \frac{j-1}{n} \right\} \right)$
	LT	$\sqrt{n} \left(\sup_x \left\{ F_n(x) - \widehat{F^*}(x) \right\} + \sup_x \left\{ \widehat{F^*}(x) - F_n(x) \right\} \right)$	$\frac{\sqrt{n}}{1-z_H} \left(\sup_j \left\{ z_H + \frac{j}{n}(1-z_H) - z_j \right\} + \sup_j \left\{ z_j - \left(\frac{j-1}{n}(1-z_H) \right) \right\} \right)$
AD^2	*Full*	$n \cdot \int_{-\infty}^{\infty} \frac{\left(F_n(x) - \widehat{F}(x)\right)^2}{\widehat{F}(x)\left(1-\widehat{F}(x)\right)} d\widehat{F}(x)$	$-n + \frac{1}{n} \sum_{j=1}^n (1-2j) \log(z_j) - \frac{1}{n} \sum_{j=1}^n (1 + 2(n-j)) \log(1-z_j)$
	LT	$n \cdot \int_{H}^{\infty} \frac{\left(F_n(x) - \widehat{F^*}(x)\right)^2}{\widehat{F^*}(x)\left(1-\widehat{F^*}(x)\right)} d\widehat{F^*}(x)$	$-n + 2n \log(1-z_H) - \frac{1}{n} \sum_{j=1}^n (1 + 2(n-j)) \log(1-z_j) + \frac{1}{n} \sum_{j=1}^n (1-2j) \log \lvert z_j - z_H \rvert$
W^2	*Full*	$n \int_{-\infty}^{\infty} \left(F_n(x) - \widehat{F}(x) \right)^2 d\widehat{F}(x)$	$\frac{n}{3} + \frac{1}{n} \sum_{j=1}^n (1-2j) \log(z_j) + \sum_{j=1}^n z_j^2$
	LT	$n \int_{-\infty}^{\infty} \left(F_n(x) - \widehat{F^*}(x) \right)^2 d\widehat{F^*}(x)$	$\frac{n}{3} + \frac{n z_h}{1-z_H} + \frac{1}{n(1-z_H)} \sum_{j=1}^n (1-2j) z_j + \frac{1}{(1-z_H)^2} \sum_{j=1}^n (z_j - z_H)^2$

Moran's log spacing statistic

The goodness-of-fit tests mentioned in Section 4.2.3 relied on Monte Carlo simulation. Although accurate in the limit, they suffer the drawback of sampling error as well as the need for excessive computing time, computing power and computer memory. Given this backdrop, and as a consequence of the MPS estimation technique, we also propose the use of Moran's log spacings statistic (see Cheng and Stephens (1989)). Such a test statistic does not require Monte Carlo simulation when parameters are estimated from the data, and is based on a chi-squared distribution approximation. We now consider the test and how it can be altered to apply to left-truncated data samples. Consider the fitted distribution function of the observed sample under the CCD approach:

$$\widehat{F^*}(x) = \begin{cases} \frac{\widehat{F}_{\gamma^C}(x) - \widehat{F}_{\gamma^C}(H)}{1-\widehat{F}_{\gamma^C}(H)} & \text{if } x \geq H, \\ 0 & \text{if } x < H, \end{cases} \tag{4.9}$$

where $\widehat{F}_{\gamma^C}(x)$ refers to the fitted distribution function on the entire support. Now suppose that γ^C has j components to be estimated. Then using $\widehat{F^*}$ as the distribution function, in a similar vein to Cheng and Stephens (1989) we can find Moran's log spacings statistic using the following stepsprocess.

(a) Calculate

$$\begin{aligned} \overline{\mathscr{M}}(\gamma^C) &= -\sum_{k=1}^{n+1} \log\left(\widehat{F^*}(x_{(k)}) - \widehat{F^*}(x_{(k-1)})\right) \\ &= -\sum_{k=1}^{n+1} \log\left(\frac{\widehat{F}_{\gamma^C}(x_{(k)}) - \widehat{F}_{\gamma^C}(x_{(k-1)})}{1-\widehat{F}_{\gamma^C}(H)}\right). \end{aligned}$$

(b) Find the test statistic, $\mathscr{T}$, specified by

$$\mathscr{T}(\gamma^C) = \frac{\overline{\mathscr{M}}(\gamma^C) + \frac{1}{2}j - C_1}{C_2}$$

where, if we set $m = n + 1$,

$$
\begin{aligned}
C_1 &= \gamma_m - \sqrt{\frac{1}{2} n}\,\sigma_m, \\
C_2 &= \frac{1}{\sqrt{2n}}\sigma_m, \\
\gamma_m &\approx m(\log m + \tilde{\gamma}) - \frac{1}{2} - \frac{1}{12m}, \\
\sigma_m &\approx m\left(\frac{\pi^2}{6} - 1\right) - \frac{1}{2} - \frac{1}{6m}, \qquad \text{and} \\
\tilde{\gamma} &\approx 0.57722 \qquad \text{(Euler's constant)}.
\end{aligned}
$$

(c) Reject the null hypothesis at a level of significance α (in that the sample comes from the distribution $\widehat{F^*}$) if $\mathscr{T}(\gamma^C) > \chi_n^2(\alpha)$.

4.3 Fitting the compound Poisson process to Property Claims Services catastrophe data

In this section, we compare results of the fitting procedure for the insured loss process using both the naive and CCD approaches. But first, we begin with a description of the data to which we apply our methodology, and check whether or not it is heavy-tailed.

4.3.1 Data description

We use the PCS data for the purposes of the application, and a detailed description of the index was provided in Section 2.3 of Chapter 2. Recall that this data represent a time series of the total insured losses (hereafter referred to as losses) from each PCS-identified natural catastrophes.

Moreover, it was mentioned in Sections 5.1 and 2.3 (from Chapter 2) that the PCS index exhibits left-truncation and involves a threhold level, H. Losses below that threshold level are not recorded in the index. Currently, that threshold level stands at \$25 million, and we have not seen any indication that it will change in the future. For the purposes of estimating the intensity for the PCS loss data, we consider the threshold to be \$25 million for all losses studied over the period of investigation. For the remainder of this Chapter and the next, we denote the \$25 million threshold by H (as mentioned in Section 4.2) – the minimum threshold above which losses are recorded. We consider the PCS index from January 1985 to July 2011 – and we stress that a large data set (as compared to that used by, *inter alia* Bur-

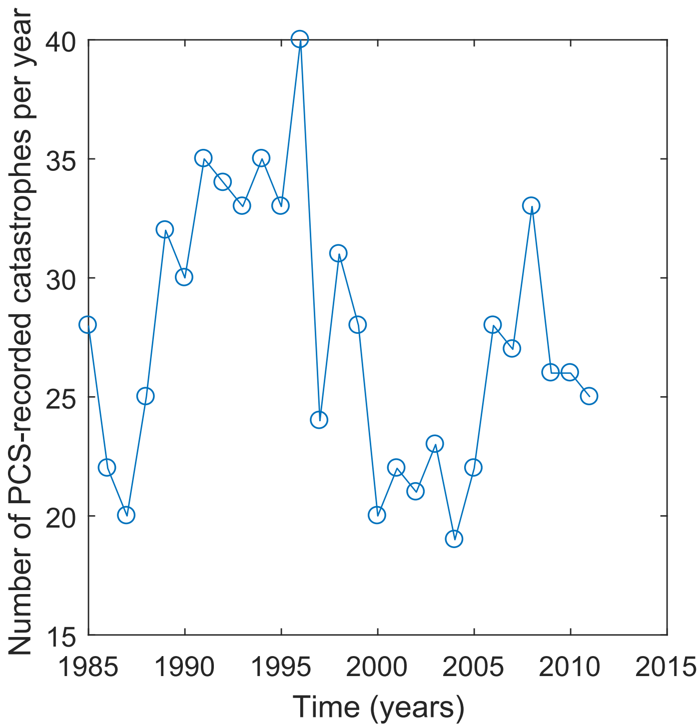

Fig. 4.2: Number of PCS-recorded catastrophes, per year, from January 1985 to July 2011.

necki *et al.* (2000) and Chernobai *et al.* (2006)) is necessary so as to have a sufficient number of "rare events" represented in the sample (Stanley *et al.*, 2007).

It is necessary to adjust the PCS index data for inflation. The data were adjusted to January 2012 using the US Department of Labor's Consumer Price Index (CPI). For the purposes of estimation and in-sample validation, claim amounts above the current PCS index threshold of $25 million are considered[6] - this choice of ours is based upon the fact that we wish to forecast the PCS losses, and we, furthermore, assume that this threshold will not be increased in the near future. Note that we only used data points that were in excess of the threshold in all of our analyses: as mentioned before some of the older data points, when adjusted for inflation, did not exceed the threshold level of $25 million. The adjusted PCS loss data exceeding the threshold level is depicted in both Figures 4.2 and 4.3.

[6] This only applied to losses prior to 1997 not reaching the $25 million threshold when adjusted to January 2012 terms, since the PCS index today only records losses in excess of $25 million.

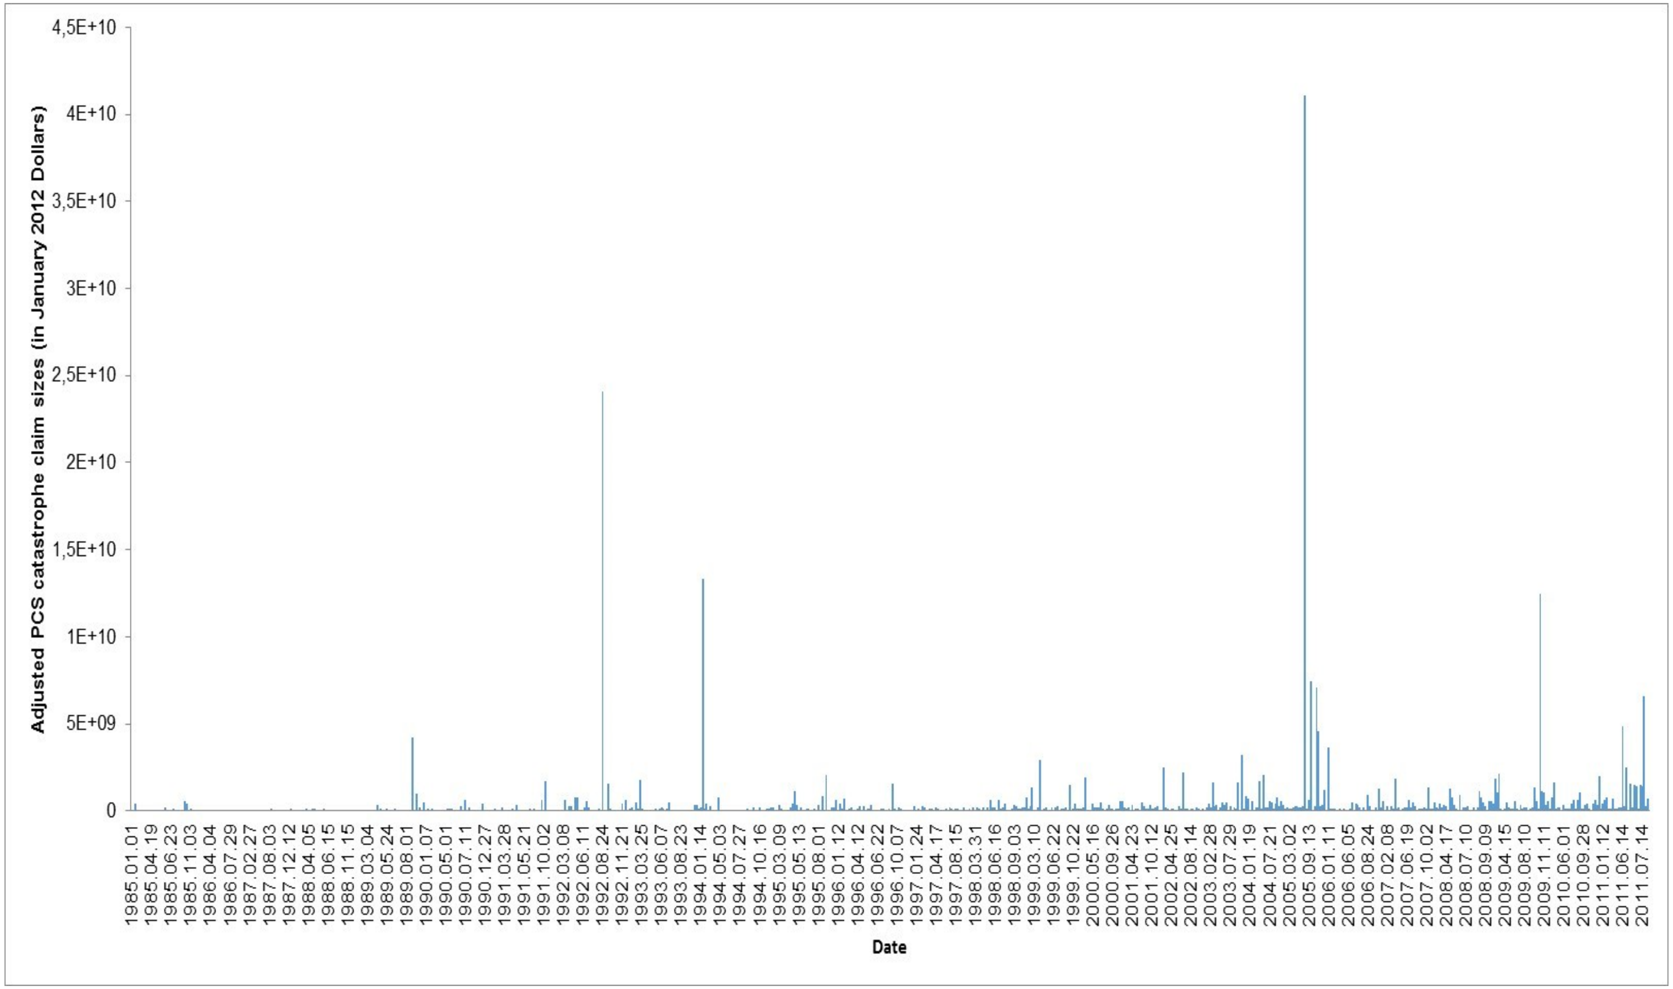

Fig. 4.3: PCS catastrophe claim severity data from January 1985 to July 2011.

We now briefly record some of the features of the data here. 742 catastrophes recorded in the index gave rise to losses in excess of the threshold level, with the largest total insured loss observed being that of Hurricane Katrina (on 25 August 2005), totalling approximately $46 billion (inflation-adjusted to January 2012). The second-largest loss was that from Hurricane Andrew on 24 August 1992, totalling approximately $25 billion. We furthermore note that (i) the maximum total loss recorded in our period of investigation is approximately 90 times the size of the empirical mean; and (ii) approximately 92% of the data are smaller than $1/50^{\text{th}}$ of the mean. This provides some evidence in support of the fact that the data set we work with is heavy-tailed.

4.3.2 Testing for heavy tails

We indeed have some cursory evidence that the PCS dataset we work with is heavy-tailed. But this remains to be investigated more formally. We check whether the underlying law of the data is a Gaussian or non-Gaussian stable distribution[7], by examining its rate of convergence using the discrimination algorithm of Burnecki *et al.* (2015, p.5-6). The idea is to examine the rate of convergence of the estimated index of stability of the proposed stable law, α, for sequential bootstrapped samples from the PCS data. If the estimated value of α is two, then the underlying distribution is possibly Gaussian, however, if the estimated value for α is less than two, then the underling distribution is most likely heavy-tailed and may indeed be a stable non-Gaussian distribution.

We bootstrap as follows. Firstly, we divide our dataset into consecutive but non-overlapping blocks each of length $K = 1, 2, \ldots, 10$. Secondly, we sum the values in each block and obtain aggregated data sets for each K. Finally, for each K we estimate α using the regression method of Koutrouvelis (1980). In Figure 4.4, we show the estimated values for α for the PCS data, for each K. For the first sample ($K = 1$; corresponding to the whole dataset), the value is lower, but for the other K's, the values increase and fluctuate around some α which is smaller than 2. This suggests that the PCS dataset exhibits non-Gaussian stable behaviour. Moreover, Figure 4.4 suggests the index of stability of the underlying stable law is less than 1: therefore, most heavy-tailed distributions fitted to this dataset will not possess a finite mean, which may limit their usage. We shall re-address this concern in Section 4.3.3.

To close this section, we remark on whether there are any significant autocorrelations present in the PCS index return series. Burnecki *et al.* (2000) found evidence for a mean-reversion structure in the index returns, for lag-periods in excess of one

[7] The stable distribution will be formally defined in Chapter 5.

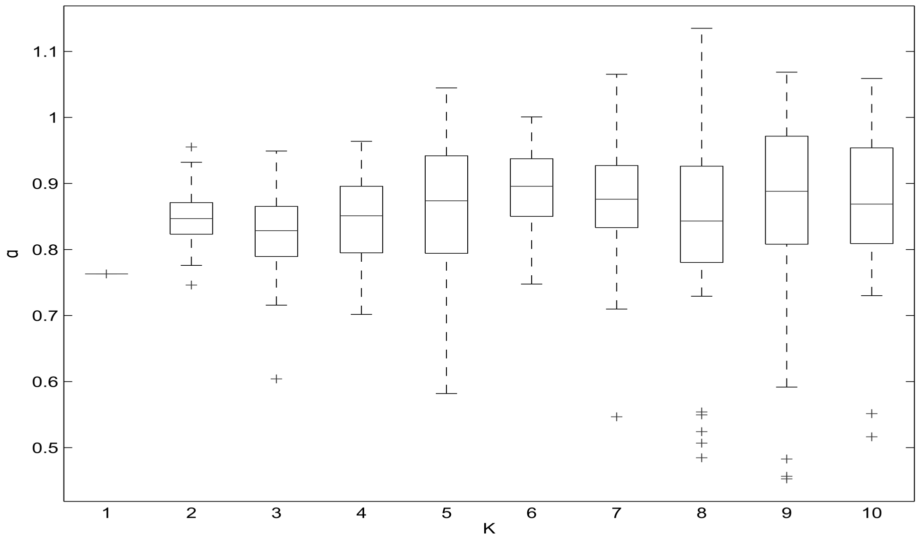

Fig. 4.4: Estimated values for the index of stability (α). The boxplots (for $K \in [1,10]$) were constructed from 100 bootstrapped samples each of length 1000.

year. We point out that this is dissimilar to most other financial data (see Mantegna and Stanley (1999); Plerou *et al.* (2000) and Burnecki *et al.* (2011)).

4.3.3 Loss distribution fitting

We propose to fit the following parametric severity distributions, having parameterisations as shown in Table 4.2. Notice the bias in our choice towards heavy-tailed and power law distributions (as is done, in other contexts, in Burnecki *et al.* (2000); Chernobai *et al.* (2006); Ma and Ma (2013); Mantegna and Stanley (1999); Bouchaud (2001) and Stanley *et al.* (2007). This is because of "rare events" skewing our data set to the right; indeed, such a judicious choice of parametric distribution is crucial when modelling heavy-tailed data. Power laws have been used extensively in the econophysics literature and they have been shown to arise in many high-frequency economic variables (Liu *et al.*, 1997; Coronel-Brizio and Hernandez-Montoya, 2005), such as return series, liquidity and volume. But also, there is some evidence for the fact that monetary losses arising from certain US natural catastrophes follow a (partial) power law (Barton and Nishenko, 1997).

Tab. 4.2: Parametric severity distributions considered to model the PCS losses, defined on the positive real numbers.

Distribution	Parameterisation	Constraints
Exponential	$\mathscr{E}(\mu) = \frac{1}{\mu} e^{\frac{1}{\mu}x}$	$\mu > 0$
Lognormal	$\mathrm{LN}(\mu,\sigma) = \frac{1}{\sqrt{2\pi\sigma^2 x}} e^{-(\log x-\mu)^2/2\sigma^2}$	$\mu, \sigma > 0$
Gamma	$\mathrm{G}(a,b) = \frac{1}{b^a \Gamma(a)} x^{a-1} e^{-x/b}$	$a, b > 0$
Weibull	$\mathrm{WEIB}(a,b) = \frac{b}{a}\left(\frac{x}{a}\right)^{b-1} e^{-(x/a)^b}$	$a, b > 0$
Burr type XII	$\mathrm{BURR}(\zeta,c,k) = \frac{\frac{kc}{\zeta}\left(\frac{x}{\zeta}\right)^{c-1}}{\left(1+\left(\frac{x}{\zeta}\right)^c\right)^{k+1}}$	$\zeta, c, k > 0$
Generalised Pareto	$\mathrm{GP}(k,\sigma) = \frac{1}{\sigma}\left(1+\frac{kx}{\sigma}\right)^{-\left(1+\frac{1}{k}\right)}$	$\sigma > 0, k \in \mathbb{R}$
Inverse Gaussian	$\mathrm{IG}(\mu,\lambda) = \sqrt{\frac{\lambda}{2\pi x^3}} \exp\left(-\frac{\lambda}{2\mu^2 x}(x-\mu)^2\right)$	$\mu, \lambda > 0$
Generalised extreme value	$\mathrm{GEV}(k,\mu,\sigma) = \frac{1}{\sigma}\exp\left(-\left(1+\frac{k(x-\mu)}{\sigma}\right)^{-\frac{1}{k}}\right) \times$	$1+\frac{k(x-\mu)}{\sigma} > 0$
	$\left(1+\frac{k(x-\mu)}{\sigma}\right)^{-1-\frac{1}{k}}$	$\sigma > 0, k \in \mathbb{R}\backslash\{0\}, \mu \in \mathbb{R}$

Some care needs to be taken when working with the GEV distribution. It must be noted that the support of the GEV distribution depends explicitly upon its parameters, in that $x > \mu - \frac{\sigma}{k}$ if $k > 0$ Embrechts *et al.* (2013). Hence, setting $\mu = \frac{\sigma}{k}$ defines a special case of the GEV distribution, which we call the modified GEV (MGEV) distribution, with its support being the positive real line, and we use this special case in our subsequent analysis. To ensure the condition that $k > 0$ remains, we introduce a "working parameter" κ and set $k = e^{\kappa}$ when optimising (4.7) for the MGEV. Secondly, it remains to be checked that

$$1 + \frac{k(x_i - \mu)}{\sigma} > 0$$

for each x_i in the observed sample $\mathbf{x}$. In words, this implies that no sample data point should fall outside the bounds of the distribution. Such a checking procedure is often overlooked (see Ma and Ma (2013) and Liu *et al.* (2014), for example).

But what is perhaps most important to consider is the problem with unbounded or infinite variance likelihood functions. In the case of the GEV (and in consequence the MGEV) distribution, as $x \downarrow \mu + \frac{\sigma}{k}$, the information matrix of the likelihood becomes infinite if $k > \frac{1}{2}$, and concerningly, if $k > 1$, likelihood functions may be unbounded. For the GP distribution, unbounded likelihoods could occur if $k > 1$. However, such a problem is not encountered if MPS is used. (Wong and Li, 2006)

Bearing such problems with the MGEV and GP distributions in mind, we fit the distributions shown in Table 4.2 using both the naive and the CCD fitting approaches. All the distributions except for the Burr[8], GP and MGEV distributions are fitted using MLE; the Burr, GP and MGEV distributions are fitted via MPS. Table 4.3 demonstrates the movements in the parameters and $F_{(\cdot)}(H)$ when the CCD fitting approach is used compared to the naive fitting approach.

Tab. 4.3: Estimated parameters (using MLE and MPS) and $F_{(\cdot)}(H)$ of the fitted distributions to the PCS data.

Distribution	$\widehat{\gamma}, F_{(\cdot)}(H)$	**Naive approach**	**CCD approach**
Exponential	μ $F_{(\cdot)}(H)$	5.63×10^8 4.35%	5.38×10^8 4.54%
Lognormal	μ, σ $F_{(\cdot)}(H)$	$18.59, 1.19$ 5.06%	$18.58, 1.49$ 14.86%
Gamma	a, b $F_{(\cdot)}(H)$	$0.54, 1.04 \times 10^9$ 14.84%	$6.79 \times 10^{-8}, 2.23 \times 10^9$ 100%
Weibull	a, b $F_{(\cdot)}(H)$	$3.37 \times 10^8, 0.66$ 16.5%	$5.56 \times 10^7, 0.38$ 52.32%
Burr type XII	ζ, c, k $F_{(\cdot)}(H)$	$7.21 \times 10^7, 1.86, 0.40$ 5.10%	$9.53 \times 10^7, 1.57, 0.70$ 7.78%
Generalised Pareto (GP)	k, σ $F_{(\cdot)}(H)$	$0.57, 1.96 \times 10^8$ 11.58%	$0.89, 1.26 \times 10^8$ 16.70%
Inverse Gaussian (IG)	μ, λ $F_{(\cdot)}(H)$	$5.63 \times 10^8, 1.30 \times 10^8$ 2.84%	$5.28 \times 10^8, 9.38 \times 10^8$ 6.28%
Modified GEV (MGEV)	k, σ $F_{(\cdot)}(H)$	$1.01, 9.99 \times 10^7$ $\approx 0\%$	$0.95, 9.99 \times 10^7$ 1.19%

From Table 4.3, it is clear that for the heavier-tailed distributions there is a change in parameter values when moving from one fitting approach to another. The mean of the fitted exponential distribution does not change by a large amount – this is expected since it is a light-tailed distribution.

The estimated fraction of missing data, $F_{(\cdot)}(H)$, is in all cases larger under the CCD fitting approach compared to the naive. Hence, the CCD fitting approach does appear to account for a degree of "true information loss", while it can be inferred that the naive fitting approach underestimates the proportion of missing data. But we do stress that the obtained estimates of $F_{(\cdot)}(H)$ are indeed dependent upon the choice of the distribution function and the threshold level of \$25 million. As expected, the exponential distribution, being a light-tailed distribution, adds back very little data. However the more heavy-tailed distributions, such as lognormal, Burr, Weibull and GP (power laws) as well as the MGEV, add back markedly higher portions of data.

We also note the following very important observation in the case of the fitted

[8] We fit the Burr distribution by MPS given that it is a three parameter distribution, and it may be possible that an unbounded likelihood may result (see Liu *et al.* (2015)).

Burr and MGEV distributions. The first and subsequent moments are infinite under the naive fitting approach: this is as expected, given our analysis on the index of stability in Section 4.3.2. However, under the CCD fitting approach, the first moment was finite for these two distributions. We point out that this non-existence of the first moment, in the case of the naive approach, restricts the usage of such distributions with such parameters if, for example, moments are required. Also, the usage of Monte Carlo simulation, which is based on the Strong Law of Large Numbers[9], is drawn into question if the naively fitted Burr and MGEV distributions are used.

Finally, we observe that $F_{\gamma^C}(H)$ is approximately 100% in the case of the gamma distribution fitted via the CCD approach. We provide evidence that the gamma distribution produces "true information loss" approximately 100% of the time. Also, the Weibull distribution appears to produce "true information loss" approximately 52% of the time, which is unrealistically high. We therefore exclude both the Weibull and the gamma distributions from subsequent analyses.

We now move on to assess the difference in parameter estimates (for the Burr, GP and MGEV distributions) when using MPS versus MLE under the CCD approach. Note that we do not use the MLE-estimated parameters for the Burr, GP and MGEV distributions in any subsequent analyses. We only present the assessment below to motivate why MPS-estimated parameters should be selected over MLE-estimated parameters. The MLE-estimated parameters, as well as $F_{\gamma^C}(H)$, are shown in Table 4.4.

Tab. 4.4: Estimated parameters (using MLE) and $F_{\gamma^C}(H)$ of the fitted distributions to the PCS data.

Distribution	$\widehat{\gamma^C_{MLE}}, \hat{F}_{\gamma^C}(H)$	**CCD approach (MLE)**
Burr type XII	ζ, c, k $\hat{F}_{\gamma^C}(H)$	$1.08 \times 10^8, 1.46, 0.73$ 7.82%
Generalised Pareto (GP)	k, σ $\hat{F}_{\gamma^C}(H)$	$0.75, 1.22 \times 10^8$ 17.36%
Modified GEV (MGEV)	k, σ $\hat{F}_{\gamma^C}(H)$	$0.85, 1.06 \times 10^8$ 0.19%

On comparison of Table 4.3 to Table 4.4, it is evident that there are small differences in parameter values. We now assess which of the MPS-estimated or MLE-estimated parameters may fit the data better through the use of quantile plots. See Figure 4.5 below.

For the fitted Burr distribution, it does not appear that there is much difference

[9] See Glasserman (2013, Section 1.1.1).

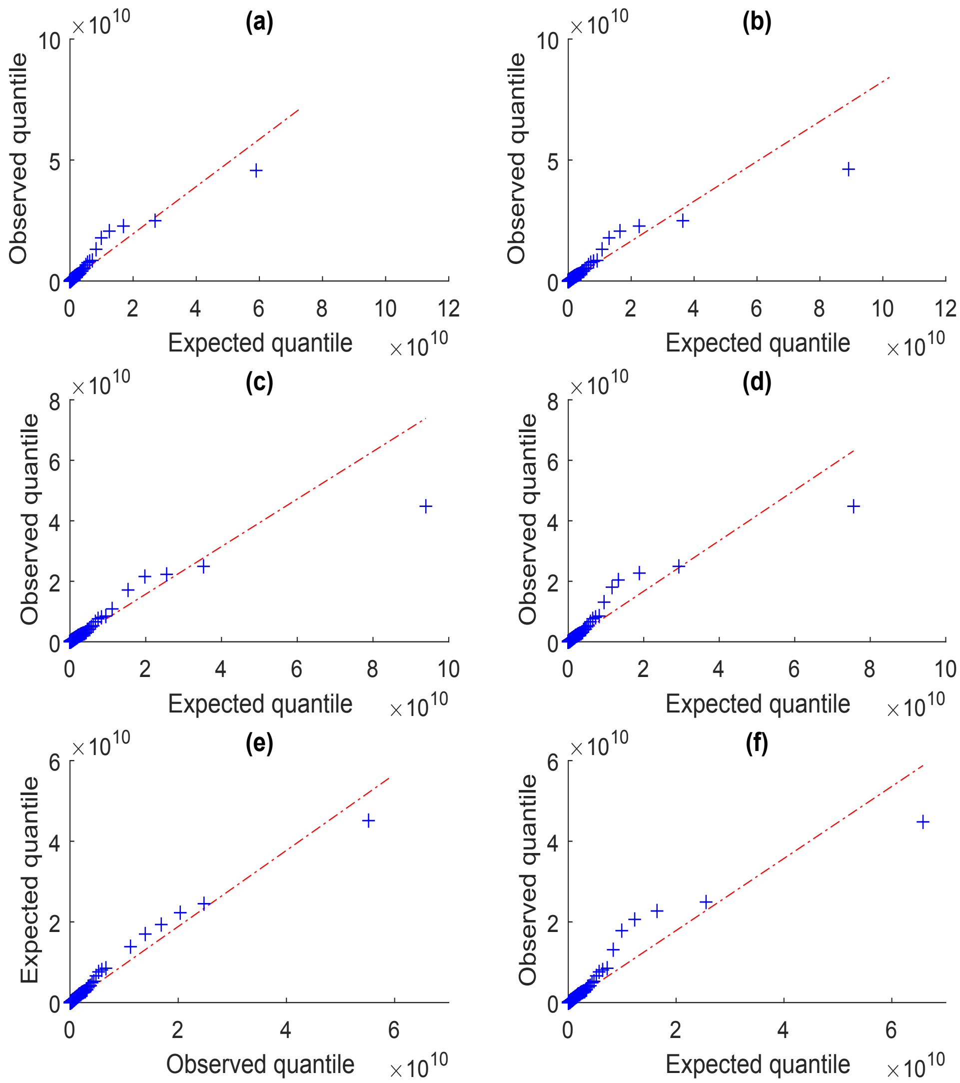

Fig. 4.5: Quantile plots for the PCS data with expected quantiles based on the (a) MPS-fitted Burr distribution; (b) MLE-fitted Burr distribution; (c) MPS-fitted GP distribution; (d) MLE-fitted GP distribution; (e) MPS-fitted MGEV distribution and (f) MLE-fitted MGEV distribution.

between using MPS or MLE to estimate parameters. However, for the GP and MGEV distribution, Figure 4.5 (c) and (e) show that most of the empirical upper quantiles (except for the most extreme observation, which pertains to the insured-losses caused by Hurricanes Katrina, Rita and Wilma) were closer to the straight line in the case of MPS as opposed to the case of MLE. Therefore, the quantile plots suggest than MPS could be a better method (when compared to MLE) to estimate parameters for the GP and MGEV distributions, when the data is heavy-tailed.

We end with the following remark. Given the outlying observation noted in Figures 4.5 (a) to (f), it could be suggested to trim the data and obtain a more robust fit to the lower quantiles of the distribution. However, we caution that the use of trimming must be determined by the purpose for which the fitted distributions will be used. For example, if it is for use in CAT bond pricing, then the largest observations do need to remain in the data set since they can significantly affect CAT bond prices. However, in the case of operational risk modelling for insurers, the data set could possibly be trimmed. We do not consider trimming here, but refer the reader to Section 5.5 of Chernobai *et al.* (2006) for further details.

4.3.4 Goodness-of-fit tests for fitted loss (severity) distributions

We now test the goodness-of-fit of the fitted severity distributions. Using the methodology outlined in Section 4.2.2, we tested the composite null hypothesis that the EDF, $F_n(X)$, belongs to a family of hypothesised fitted distributions. We begin with the distributions fitted by the MPS approach, and the results of these tests are presented in Table 4.5. For Moran's log spacing statistic, note that n was set equal to 741.

Tab. 4.5: Results of the in-sample goodness-of-fit tests for distributions fitted using MPS. p-values were obtained for the first four tests using 1,000 Monte Carlo simulations, and are provided below. Test statistics are denoted by D for KS, V for Kuiper, A^2 for AD, W^2 for Cramér von Mises and $\overline{\mathscr{M}}$ for Moran's log spacing statistic.

	Naive approach					**CCD approach**				
	D	V	A^2	W^2	$\overline{\mathscr{M}}$	D	V	A^2	W^2	$\overline{\mathscr{M}}$
Exponential	*9.23* (< 0.005)	*10.43* (< 0.005)	*157.06* (< 0.005)	*32.76* (< 0.005)	*925.88* (< 0.005)	*9.19* (< 0.005)	*9.74* (< 0.005)	*200.40* (< 0.005)	*38.56* (< 0.005)	*902.75* (< 0.005)
Lognormal	*2.19* (< 0.005)	*3.64* (< 0.005)	*7.76* (< 0.005)	*1.19* (< 0.005)	*877.77* (< 0.005)	*1.39* (0.07)	*2.15* (0.005)	*2.94* (0.22)	*0.49* (0.12)	*776.34* 0.17
BURR	*0.97* (0.01)	*1.66* (0.01)	*1.31* (< 0.005)	*0.18* (< 0.005)	*849.01* (< 0.005)	*0.68* (0.50)	*1.09* (0.36)	*0.33* (0.83)	*0.06* (0.72)	*724.51* (0.66)
Generalised Pareto	*3.25* (< 0.005)	*6.39* (< 0.005)	*15.56* (< 0.005)	*2.47* (< 0.005)	*962.72* (< 0.005)	*1.09* (0.03)	*2.23* (0.04)	*2.35* (0.16)	*0.49* (0.05)	*780.93* (0.15)
Inverse Gaussian	*3.29* (< 0.005)	*4.13* (< 0.005)	*14.40* (< 0.005)	*2.78* (< 0.005)	*891.53* (< 0.005)	*2.37* (< 0.005)	*3.04* (< 0.005)	*7.95* (< 0.005)	*1.42* (< 0.005)	*862.89* (< 0.005)
MGEV	*0.98* (0.03)	*1.78* (0.04)	*1.38* (0.02)	*0.51* (0.01)	*835.48* (0.01)	*0.88* (0.69)	*1.33* (0.52)	*0.46* (0.91)	*0.11* (0.88)	*697.84* (0.87)

It is clear from Table 4.5 that under the naive fitting approach, none of the posited distributions appear to fit the data well, except for the Burr and MGEV distributions to a small extent. Under the CCD fitting approach there does not appear to be enough evidence to reject the fit of the very heavy-tailed distributions, those being the Burr and MGEV – their p-values for all five statistical tests exceeded 0.05 significantly. The GP and lognormal distribution also appeared to fit the data to a rather limited extent, however, the heavier-tailed distributions appear to fit better. Finally, we point out that Moran's log spacings statistic performs well and in-line with the other four goodness-of-fit tests applied.

It is evident that the inverse Gaussian distribution appears not to fit the data well, under both the naive and CCD fitting approaches. This is in direct contrast to Ma and Ma (2013), who found that there was not enough evidence, based on the data, to reject the fit of their inverse Gaussian distribution.

On consideration of both the quantile plots as well as the goodness-of-fit tests, it appears that in the CCD case with MPS-estimated parameters, the Burr and MGEV distributions fit the data best and we use these distributions in Section 4.4 below.

As a short extension to our analyses above, we consider the goodness-of-fit of the MLE-fitted Burr, GP and MGEV distributions. Now in respect of the distributions fitted via the CCD approach on the basis of MLE, similar goodness-of-fit test statistics were noted to the CCD fitted distributions on the basis of MPS. Such statistics are reported in Table 4.6 below, and on balance it appears that in this case the Burr and MGEV distributions appeared to fit the data best. It was also interesting to note that Moran's log spacings statistic reported higher p-values for the MPS-estimated parameters compared to the MLE-estimated parameters; and furthermore that all the associated p-values for the MPS-estimated MGEV parameters were higher than those in the MLE-estimated case.

Tab. 4.6: Results of the in-sample goodness-of-fit tests for distributions fitted using MLE. p-values were obtained for the first four tests using $1,000$ Monte Carlo simulations, and are provided below. Test statistics are denoted by D for KS, V for Kuiper, A^2 for AD, W^2 for Cramér von Mises and $\overline{\mathscr{M}}$ for Moran's log spacing statistic.

	CCD approach (MLE)				
	D	V	A^2	W^2	$\overline{\mathscr{M}}$
BURR	*0.47* (0.67)	*1.69* (0.31)	*0.32* (0.83)	*0.05* (0.73)	*758.48* (0.32)
GP	*0.97* (0.02)	*1.42* (0.03)	*1.61* (0.09)	*0.22* (0.05)	*788.25* (0.11)
MGEV	*0.88* (0.58)	*1.26* (0.53)	*0.49* (0.83)	*0.07* (0.71)	*706.74* (0.81)

On consideration of both the quantile plots as well as the goodness-of-fit tests, it appears that in the CCD case, the Burr and MGEV distributions fit the data better on the basis of the MPS approach compared to MLE. Note that for the remainder of this Chapter, we primarily use the MPS-estimated parameters for the Burr and MGEV distributions, under the CCD approach.

We now turn to ending this section with some comments on the practical applicability of our findings. We used a dataset with data starting from 1985. It must be borne in mind that the intensity and severity distributions back in the 1980's and 1990's are different from what they are today, due to, *inter alia*, insurance-related regulations increasing the claim amounts, new and more developed reinsurance agreements as well as urbanisation and global warming (the latter's effects are, however, likely to be small). Therefore, one could posit that total catastrophic insured losses increase over time. We account for this, and allow for a linear trend in the intensity function of the fitted aggregate loss process (ALP - see Section 4.3.6). Alternatively, it may be possible to take an "actuarial approach" and add a margin to the fitted severity distribution's parameters, so as to account for increases in the severity over time. However, we do not aim to model the structural breaks, from time-to-time, in the PCS data explicitly: we simply intend for the PCS data to be a candidate dataset upon which we can illustrate our heavy-tailed modelling procedure.

4.3.5 Sampling from selected heavy-tailed distributions

We note here that a very important issue – that is, rare-event simulation – needs to be addressed when one is sampling from certain very heavy-tailed distributions (such as the BURR and GEV (and in consequence the MGEV) distributions for certain parameter ranges). Note that the problems associated with the simulation of rare events from such distributions have been studied extensively (see, for example Asmussen *et al.* (2000) and Asmussen and Kroese (2006)). However, such an issue has been overlooked by recent research on catastrophe bond prices (see Ma and Ma (2013) and Liu *et al.* (2014), for example). To illustrate this point, we use the Burr and MGEV distributions each fitted via the CCD approach above and consider $1,000$ samples of size $10,000$ each simulated from these distributions. For each sample, the sample mean is calculated and plotted in Figure 4.6 (a). In addition, the theoretical mean is calculated and included in the plot.

For both distributions, the sample mean is almost consistently lower than the theoretical mean. In the case of the Burr distribution, Burnecki, Misiorek and Weron (2005) point out that this can occur when the product of the Burr parameters, ck (i.e. the reciprocal of the extreme value index), is greater than (but very close to) one. In our case, ck is approximately 1.099.

An explanation for the aforementioned anomaly is the difficulty in sampling from an extremely heavy-tailed distribution. On the extremely rare occasion when an observation from the very far reaches of the right tail is indeed simulated, its magnitude dwarfs its associated probability and so this has the effect of markedly increasing the realised Monte Carlo estimate. However, because of the rareness of observing such large losses, most realised Monte Carlo estimates will not contain such high losses, hence the lower value of the Monte Carlo estimates. In fact, this phenomenon appears clearly in Figure 4.6 (a) and Figure 4.6 (c).

In an attempt to overcome such a difficulty in sampling Burr- and MGEV- distributed random numbers from the right tails, we employ importance sampling (see Asmussen *et al.* (2000)). As an importance sampling distribution Q_{IS} (with corresponding density function q_{IS}, assumed to exist, and $Q_{\text{IS}} \ll F$, F being the loss distribution we wish to sample from), we take a more heavy-tailed distribution for the losses - we consider a generalised Pareto distribution with $k = 1/0.73$ and $\sigma = 1.26 \times 10^8$. Notice that our selected Q_{IS} is a heavier-tailed version of the CCD-approach fitted GP distribution. We highlight that our chosen importance sampler may not be optimal, and so in repeating such a simulation other suitable importance samplers (heavier-tailed than the distribution we originally wish to sample from) may be considered.

Using the goodness-of-fit tests in Section 4.2.3, we tested to see whether the losses did indeed conform to the distribution Q_{IS}, and did not find any evidence to reject this claim. We thereafter proceeded as follows. We simulated each of the losses from the importance sampling distribution Q_{IS}, say $X_i^{Q_{\text{IS}}}$, and thereafter

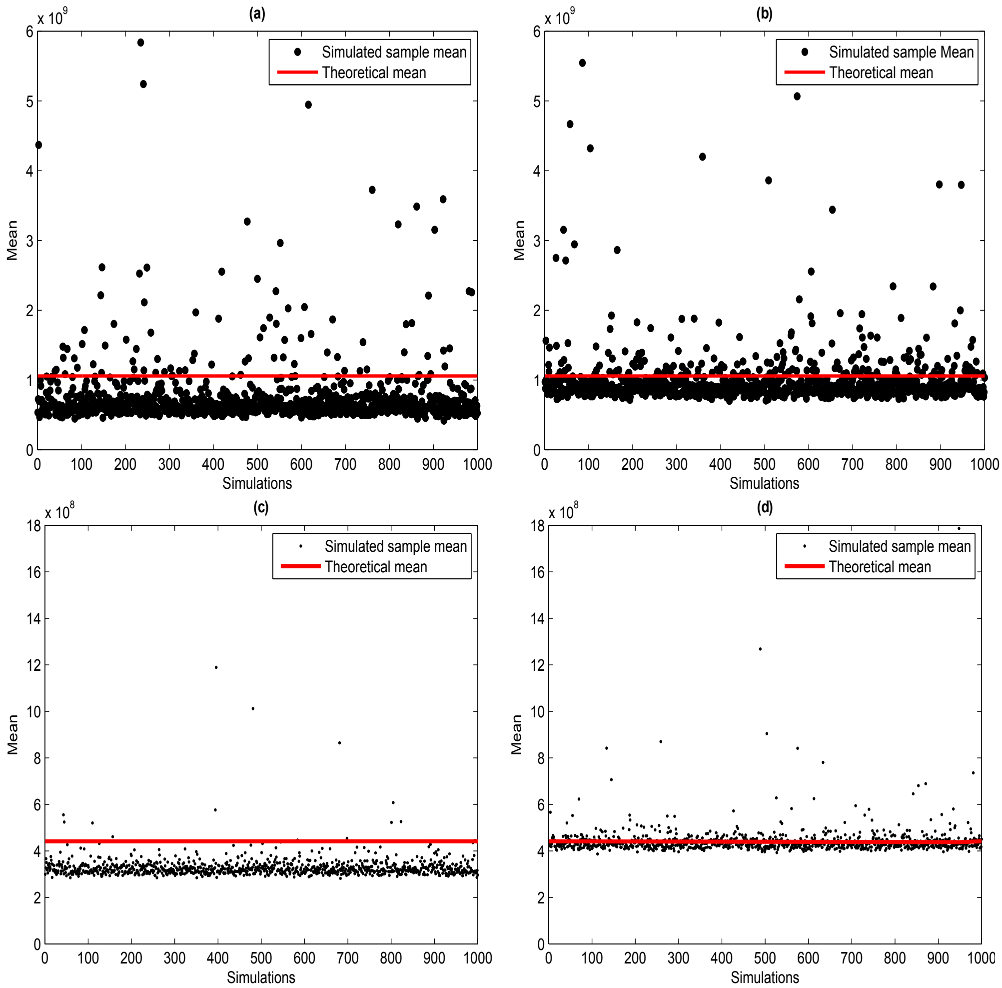

Fig. 4.6: Theoretical means and simulated sample means under: (a) standard simulation techniques (Burr); (b) importance sampling techniques (Burr); (c) standard simulation techniques (MGEV) and (d) importance sampling techniques (MGEV)

scaled each as demonstrated by Equation (4.10) to obtain the scaled losses we would employ as our X_i's in the Monte Carlo simulation algorithm presented in Section 4.4.2.

$$X_i := \frac{f\left(X_i^{Q_{\mathrm{IS}}}\right)}{q_{\mathrm{IS}}\left(X_i^{Q_{\mathrm{IS}}}\right)} X_i^{Q_{\mathrm{IS}}} \tag{4.10}$$

This importance sampling approach to simulating from the CCD-fitted Burr distribution produced better behaviour. The sample means of both the CCD-fitted Burr and MGEV distributions were closer to the theoretical mean (see Figure 4.6 (b) and Figure 4.6 (d)) and, equivalently, were not markedly below the theoretical mean as is the case without importance sampling.

4.3.6 Intensity estimation

We now turn to estimating the intensity embedded in the process demonstrated by Equation (4.1). The arrival times of the individual losses are modelled using a Poisson process. Poisson processes have been widely applied in modelling the arrivals of natural catastrophes for the purposes of pricing catastrophe-linked securities, see for example Christensen and Schmidli (2000b); Lee and Yu (2007); Burnecki and Kukla (2003); Burnecki, Kukla and Taylor (2005); Dassios and Jang (2003a); Cox *et al.* (2004b); Chernobai *et al.* (2006); Jaimungal and Wang (2006b); Biagini *et al.* (2008a); Härdle and Cabrera (2010); Braun (2011); Ma and Ma (2013); Nowak and Romaniuk (2013) and Têtu *et al.* (2015). In this chapter we use a time-inhomogeneous Poisson process, having a deterministic (time-dependent) intensity rate specified in such a way as to account for the seasonality (by year) observed in the arrivals of natural catastrophes. Figure 4.2 illustrates this annual seasonality – and it appears that a sinusoidal intensity may satisfactorily model the frequency.

To formalise, we suppose that the number of losses over the interval of study $(0,T]$ is driven by a time-inhomogeneous Poisson point process $N = \{N_t, t \geq 0\}$ with intensity given by $\lambda(t)$. We begin by considering the desirable aspects of such a function – such a function should ideally elicit the long-term mean of the process and should contain a sinusoidal term to account for the seasonality of the pervasive atmospheric perils inherent in the PCS-recorded natural catastrophes, as well as an exponential term to allow for extraneous events (such increases in the number of natural catastrophes, in particular periods of torrential rain and heightened cyclone activity, may be associated with severe La Niña). This exponential term should also be subject to some degree of seasonality. Finally, the function should elicit an upward linear trend (Lin and Wang, 2009a; Têtu *et al.*, 2015). We noticed that the intensity postulated by Ma and Ma (2013),

$$\lambda^1(t) = a + b\sin^2(t+c) + d\exp\left(\cos\left(\frac{2\pi t}{\omega}\right)\right), \tag{4.11}$$

for all $t \geq 0$ and with $a > 0$, $d > 0$ as well as $\omega > 0$, does indeed embrace the aforementioned characteristics. However, notice that the periodicity of the sine term is 2π, which is in excess of one year. So, the year-by-year seasonal behaviour of such catastrophes may not be well captured. Therefore, we propose the following alternative model which more satisfactorily captures the seasonality effects over a yearly time period:

$$\lambda^2(t) = a + bt + c\sin(2\pi(t+d)) + m\exp\left(\cos\left(\frac{2\pi t}{\omega}\right)\right), \tag{4.12}$$

for all $t \geq 0$ and with $a > 0$, $b > 0$, $d > 0$ as well as $\omega > 0$. We also note that Chernobai *et al.* (2006) analysed a special case of (4.12) with no exponential and linear trend. Chernobai *et al.* (2006) put forward the following model for the intensity:

$$\lambda^3(t) = a + b \cdot 2\pi\sin\left(2\pi(t-c)\right), \tag{4.13}$$

for all $t \geq 0$ and with $a > 0$. However, we noticed that if, for example, the time increases by one unit, there will be no change in the intensity predicted by this model (which is due to the presence of the constant 2π in the sine term) and this appears unsatisfactory if one is using $\lambda^3(t)$ to model the number of catastrophes per year. We therefore omit $\lambda^3(t)$ from further study. Finally, it may be more appropriate to simply use a constant annual intensity, which we calculate by

$$\lambda^4(t) = \frac{n}{T} = 27.92, \tag{4.14}$$

for all $t \geq 0$ and with n denoting the number of catastrophes over the period studied, and T denoting the length (in years) of this period. In light of Figure 4.7, notice that $\lambda^4(t)$ evidently underestimates the number of catastrophes per quarter, but such a constant intensity may, by chance, still lead to multiple catastrophes in a short period.

One potential way to estimate the intensity would be to fit $\lambda^i(t)$ (for $i \in \{1,2\}$) to the observed number of PCS-recorded catastrophes per year, and it appears that this was the way which was implemented by Ma and Ma (2013). However, a potential problem with such an estimation procedure is that it may not capture any seasonal trends in the data, since seasonality effects occur over a calendar year. A satisfactory alternative, which captures the seasonality in a better way, would be to fit the aggregate catastrophes sequentially using the mean-value function for the time-inhomogeneous Poisson process. We therefore proceed as follows. Firstly, we consider the vector recording the cumulative number of catastrophes over the period $(0,T]$,

$$V = [1,2,...,N]. \tag{4.15}$$

Secondly, define the function $g : \mathbb{R}^N \to \mathbb{R}^N$ to be

$$g([t_1,t_2,...,t_N]) = \left[\int_0^{t_1} \lambda^i(s)\mathrm{d}s, \int_0^{t_2} \lambda^i(s)\mathrm{d}s, ..., \int_0^{t_N} \lambda^i(s)\mathrm{d}s\right] \tag{4.16}$$

for $i \in \{1,2\}$ and where the vector $[t_1,t_2,...,t_N]$ represents the time, in years since time zero over the period $(0,T]$.

We then aim to minimise the distance between g and V (i.e. non-linear least squares). So in the third step we fit the parameters for $\lambda^i(s)$ by minimising (4.17).

$$\sum_{k=1}^{N}\left(\int_0^{t_k} \lambda^i(s)\mathrm{d}s - k\right)^2 \tag{4.17}$$

Finally, as we are implementing the conditional fitting approach of Chernobai *et al.* (2006), it remains to "scale up" the intensities by dividing the fitted intensity by $1-F_{\gamma^C}(H)$. The "scaling-up" is performed only in the pricing section of this chapter (see Section 4.4), but is omitted from the analysis in the remainder of this section.

Table 4.7 shows the estimated values of the parameters belonging to each of $\lambda^1(t)$, $\lambda^2(t)$ and $\lambda^4(t)$. An immediate first observation is that our parameters do not correspond to those estimated by Ma and Ma (2013). Moreover, from these results it appears that the mean number of catastrophes per quarter, as predicted by $\lambda^1(t)$ and $\lambda^2(t)$, is approximately 30. The change in parameters from $\lambda^1(t)$ to $\lambda^2(t)$ is also clear – it is sensible that only b and c change by a large amount.

Tab. 4.7: Parameter estimated for each of $\lambda^1(t)$, $\lambda^2(t)$ and $\lambda^4(t)$ (before scaling).

Intensity	Parameter	Estimate
$\lambda^1(t)$	a	27.75
	b	2.24
	c	-0.47
	d	1.301
	ω	4.76
$\lambda^2(t)$	a	24.93
	b	0.026
	c	5.6
	d	7.07
	m	10.3
	ω	4.76
$\lambda^4(t)$	-	27.92

Now, precise tests are required in order to assess how well the intensity functions fit the data. We follow in the spirit of Ma and Ma (2013) and use goodness-of-fit measures which allow us to compare the mean number of catastrophes (per year) based on the models to the observed number of catastrophes per year. The five goodness-of-fit measures used are the mean absolute error (MAE), the root mean-square error (RMSE), Theil's coefficient (U), the Nash-Sutcliffe coefficient of efficiency (E) and the index of agreement (D) as shown in (4.18) to (4.22). Their respective computational formulae are:

$$\text{MAE} = \frac{1}{N}\sum_{i=1}^{N}|O_i - P_i|, \tag{4.18}$$

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(O_i - P_i)^2}, \tag{4.19}$$

$$\text{U} = \frac{\sqrt{\frac{1}{N}\sum_{i=1}^{N}(O_i - P_i)^2}}{\sqrt{\frac{1}{N}\sum_{i=1}^{N}(O_i)^2} + \sqrt{\frac{1}{N}\sum_{i=1}^{N}(P_i)^2}}, \tag{4.20}$$

$$\text{E} = 1 - \frac{\sum_{i=1}^{N}(O_i - P_i)^2}{\sum_{i=1}^{N}\left(O_i - \bar{O}\right)^2}, \tag{4.21}$$

$$\mathrm{D} = 1 - \frac{\sum_{i=1}^{N} (O_i - P_i)^2}{\sum_{i=1}^{N} \left(|P_i - \bar{O}| + |O_i - \bar{O}|\right)^2}, \tag{4.22}$$

where O_i is the observed number of catastrophes in each year, P_i is the predicted number of catastrophes calculated via the expected value of the number of catastrophes in each year, N is the total number of years under study (27 years) and $\bar{O}$ is the mean number of observed catastrophes per year (Ma and Ma, 2013). Table 4.8 provides a summary of how well each of the fitted intensities adhere to the catastrophe arrivals. It is apparent that $\lambda^4(t)$ has the highest MAE, RMSE and U statistics, and also the lowest E and D statistics. Also, it appears that $\lambda^2(t)$ has a lower RMSE, MAE and U statistic (albeit slightly) as well as a higher E and D statistic than $\lambda^1(t)$. On balance, we postulate that $\lambda^2(t)$ has more in its favour, especially because it allows both positive and negative deviations from the parameter a and allows for an upward trend. Therefore, we select $\lambda^2(t)$ as the candidate to model the intensity of the time-inhomogeneous Poisson process for the PCS-recorded catastrophe losses. However, it is noticed that each of the postulated intensities do not give rise to a large number of catastrophes clustering within a certain quarter (such as those uncommonly-occurring quarters giving rise to 60 or more catastrophes, as is clear in Figure 4.7) - therefore, further research into stochastic intensity models may be necessary. Moreover, $\lambda^2(t)$ appears to have the effect of underestimating the count of catastrophes per quarter in those years when an excessively large number of catastrophes do occur, and overestimating the count in those years when very few catastrophes occur. These effects do, in fact, cancel out to some extent over time.

Tab. 4.8: Results from the goodness-of-fit measures for the proposed Poisson intensities.

Intensity	MAE	RMSE	U	E	D
$\lambda^1(t)$	13.8773	16.6507	0.0195	0.9948	0.9985
$\lambda^2(t)$	10.4469	12.7869	0.0149	0.9964	0.9991
$\lambda^4(t)$	21.2907	24.6184	0.0287	0.9923	0.9967

The integration of the fitted intensity function is depicted in Figure 4.7, together with the accumulated PCS losses. Visually, the intensity appears to fit the data well over all time periods considered. See Figure 4.7.

Finally, as we are implementing the conditional fitting approach, it remains to "scale up" the intensities by dividing the fitted intensity by $1 - F_{\gamma^C}(H)$. The "scaling-up" is performed in the pricing section of this chapter (see Section 4.4).

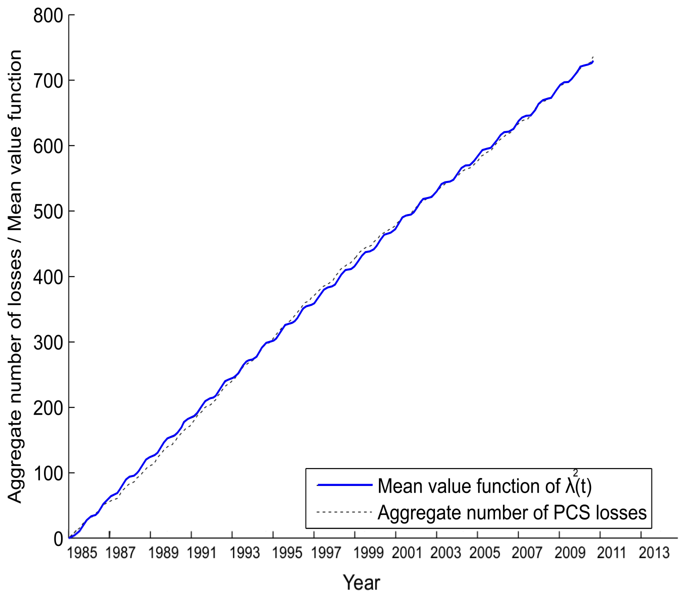

Fig. 4.7: Mean value function for fitted intensity ($\lambda^2(t)$) and aggregate number of PCS losses, from 1985 to 2011.

4.4 Pricing CAT bonds under the naive versus the conditional approach

4.4.1 Pricing framework

The aim of this section is to assess the difference in estimated CAT bond prices firstly using the naive fitting approach and secondly the CCD approach for parameter estimation. To do so, we adopt our CAT bond pricing framework outlined in Chapter 3 in order to assess this difference and come to a conclusion as to which fitting methodology has more in its favour. Note that two of the most pervasive classes of models used in the literature to price CAT bond are (i) the LFC model of Morton Lane (Lane, 2003) and (ii) approaches based on the compound Poisson process. We resort to the latter framework here (which is encompassed in Chapter 3), but emphasise that our fitting methodology is also equally applicable to the more practical pricing approach of Lane (2003).

As a final point to this preamble, in order to obtain a more tractable pricing formula it may be necessary to fit a process to the interest rates which impact CAT bond prices, and this has been analysed recently in Ma and Ma (2013) as well in Nowak and Romaniuk (2013). Various processes have been considered in the literature, such as the Cox-Ingersoll-Ross model of Cox *et al.* (1985) and the Vasicek model of Vasicek (1977). However, again in the interest of simplicity and restricting our attention to the effect of changing the fitting procedure on the insured loss process, both Chapters 4 and 5 assume that the interest rate has a constant value r, as mentioned in Chapter 3.

4.4.2 Monte Carlo simulation

In Section 3.4, we present Equations (3.15) and (3.17) in such a way so that it is straightforward to see how Monte Carlo simulation can be performed to numerically evaluate them. In this section, we describe how we simulate the prices of the CAT bonds specifically at time 0. The Monte Carlo framework presented has dual applicability, in that it can be applied in the context of both the naive and the CCD fitting approach.

As before, let T be the maturity date of either a zero coupon or coupon-paying CAT bond and let L_t be the value of the insured loss process at time t. For the zero-coupon bond, the only random variable of interest is the aggregate loss value at maturity L_T. For the coupon-paying bond, simulated values of the insured loss process are required on each coupon date, i.e. L_t for $t = 0.25, ..., T$. The method for simulating the insured loss process over the interval $(0, T]$ is described as follows. Firstly, simulate a single random number of events, N (with corresponding realisation n), in the specified time interval. Then, given $N = n$, simulate n losses from the specified distribution for X. This relies on the independence assumption of the loss process X and the arrival process N. Finally, summing the values of X_i which exceed \$25 million provides a realisation of the insured loss process L at the specified time.

In the simplest case where the insured loss process follows a time-homogeneous Poisson process with constant intensity λ, simulating N simply involves generating a Poisson random variable with intensity λt. Note that for the coupon-paying bond, it is necessary to simulate N_i for $i = 1, ..., m_T$, where m_T is the number of quarters in $(0, T]$ and $\sum_{i=1}^{m_T} N_i = N$. The values of L at each coupon date are given by $L_t = \sum_{i=1}^{M_t} X_i$ where $M_t = \sum_{i=1}^{m_t} N_i$, where m_t is the number of quarters in $(0, T]$. Where the insured loss process follows a time-inhomogeneous Poisson process with deterministic intensity $\lambda(t)$, the thinning algorithm shown in Burnecki and Weron (2005, p.324) is used to simulate L_t at each quarter. The key simulation steps of this algorithm are shown below (modified from Burnecki and Weron (2005)):

Key steps of algorithm to generate L_t in the case of the time-inhomogeneous Poisson process

1. Set $t^* = 0$ and $L^* = 0$

2. While $t^* < t$

(i) Generate a realisation of an exponential random variable, E, with parameter $\bar{\lambda}$. Note that $\bar{\lambda}$ should be chosen such that $\lambda(s) \leq \bar{\lambda}\ \forall\ s \leq T$.

(ii) Set $t^* = t^* + E$

(iii) Generate a uniformly-distributed random variable, U, on the interval $(0,1)$.

(iv) If $U > \lambda(t^*)/\bar{\lambda}$, then return to step (i) else generate a realisation, x, of the random variable X and set $L^* = L^* + x$.

3. Set $L_t = L^*$.

The same algorithm is used for both the zero coupon- and coupon-paying bonds. Once a sufficient number of realisations of L_t are simulated for the required time-to-maturity and/or coupon paying dates, the required probabilities in Equations (3.15) and (3.17) can be evaluated for any threshold level.

4.5 Numerical comparison of the naive versus the conditional complete-data fitting approaches

4.5.1 Zero coupon CAT bond simulation study

By ignoring the threshold level of the index, H, in the parametric distribution fitting procedure, it is expected that the upper quantiles of the fitted loss distribution will be underestimated. Therefore, given that the PCS index-linked CAT bond prices presented in this chapter are based on the upper quantiles of the insured loss process, it is expected that in general the prices generated using the naive approach will be greater than those generated using the CCD fitting approach. We aim to confirm this through the numerical computation of CAT bond prices.

In this section, we simulate using the following inputs: we let $D \in [3,740; 44,880]$ million US Dollars, and $T \in \{0.25, 0.5, ..., 2.5\}$ (years). These figures for D are taken from Ma and Ma (2013) on the basis of facilitating comparison of our results with theirs (in fact, the quarterly average loss is \$3,740 million, while three times the annual average loss is \$44,880 million). Moreover, we assume a constant $\rho = 50\%$, however, this can be altered depending on the term sheet of the CAT bond. For the purposes of the Monte Carlo simulation, we simulated $20,000$ paths for the insured loss process.

ZC CAT bond prices estimated by Monte Carlo simulation, with insured loss processes fitted by both the naive and the CCD approaches, are illustrated in Figure

4.8. The left-hand (resp. right-hand) column of Figure 4.8 shows the CAT bond pricing surface with respect to time-to-maturity and threshold level, using the fitted Burr distribution (resp. MGEV) for the underlying losses and the fitted NHPP. Both pricing surfaces simulated under the naive and the CCD fitting approaches appear to exhibit similar behaviour. As the threshold level decreases and the time-to-maturity increases, the price of the CAT bond decreases as expected. As one moves to a lower threshold level and to a greater time-to-maturity, a sharp drop in price is observed. The price decrease appears to be faster under the CCD fitting approach. Figure 4.8 (c) and Figure 4.8 (f) show the difference, in price, when using the naive versus the CCD fitting approaches. An immediate observation is that the prices under the naive fitting approach are, for the majority of the time, greater than those under the CCD fitting approach. But this effect is, in particular, more pronounced in the case of the MGEV distribution - a distribution which fitted our data better than all others considered. On balance, this all suggests that the tails of the fitted Burr and MGEV distributions are markedly heavier when the CCD fitting approach is used, compared to when the naive approach is used. Curiously, in the cases of this CCD-fitted Burr and MGEV distributions, if the importance sampling technique is not used and standard simulation is rather employed, Figure 4.8 (c) and Figure 4.8 (f) are inverted. This is consistent with what we would expect should very large losses not be sampled enough.

Overall the results seem to indicate the probability required in Equation (3.15), $\Pr(L_T > D)$, is larger if the CCD fitting approach is employed compared to the naive fitting approach, and is moreover pronounced in the case of very heavy-tailed distributions. Even though they did not fit as well as the Burr and MGEV distributions, simulation analyses similar to the above were performed for the lognormal and GP distributions[10], and analogous results were noted. This all echoes the findings of Chernobai *et al.* (2006): they also encountered smaller (or underestimated) probabilities in the naive fitting approach. From the point of view of an investor in such CAT bonds, especially if the investor is a prudent one such as a pension fund, a conservative approach to valuing assets may be key, and the CCD fitting approach appears to offer such an avenue.

4.5.2 Coupon-paying CAT bond simulation study

For the purposes of simulation, we let D be denoted by D_{CP} and let the coupon-paying times be from the vector $\{0.25, 0.5, ..., 2.5\}$ (years). ρ remained as for the ZC CAT bond case, and we let the coupon rate be a constant $c = 0.05$ per unit nominal.

[10] To simulate from the GP distribution, we did apply the procedure outlined in Section 4.3.5.

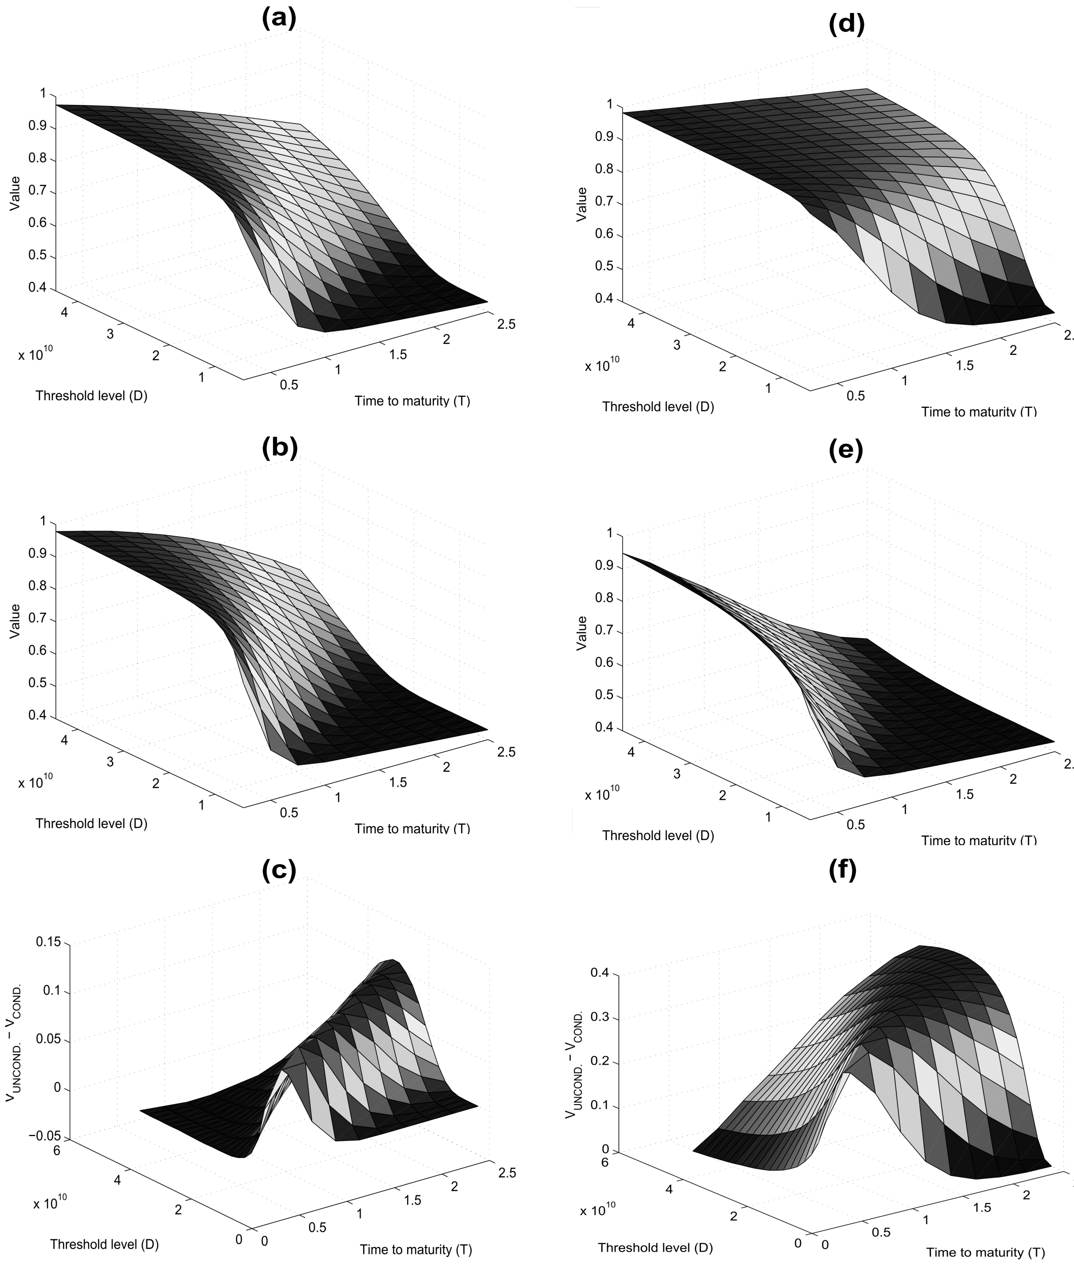

Fig. 4.8: Surfaces for (a) ZC CAT bond prices under Burr distribution, fitting via the naive approach; (b) ZC CAT bond prices under the Burr distribution, fitting via the CCD approach; (c) the difference between the ZC CAT bond prices under the naive approach (V_{UNCOND}) and the CCD approach (V_{COND}) under the Burr distribution; (d) ZC CAT bond prices under MGEV distribution, fitting via the naive approach; (e) ZC CAT bond prices under the MGEV distribution, fitting via the CCD approach; (f) the difference between the ZC CAT bond prices under the naive approach (V_{UNCOND}) and the CCD approach (V_{COND}) under the MGEV distribution.

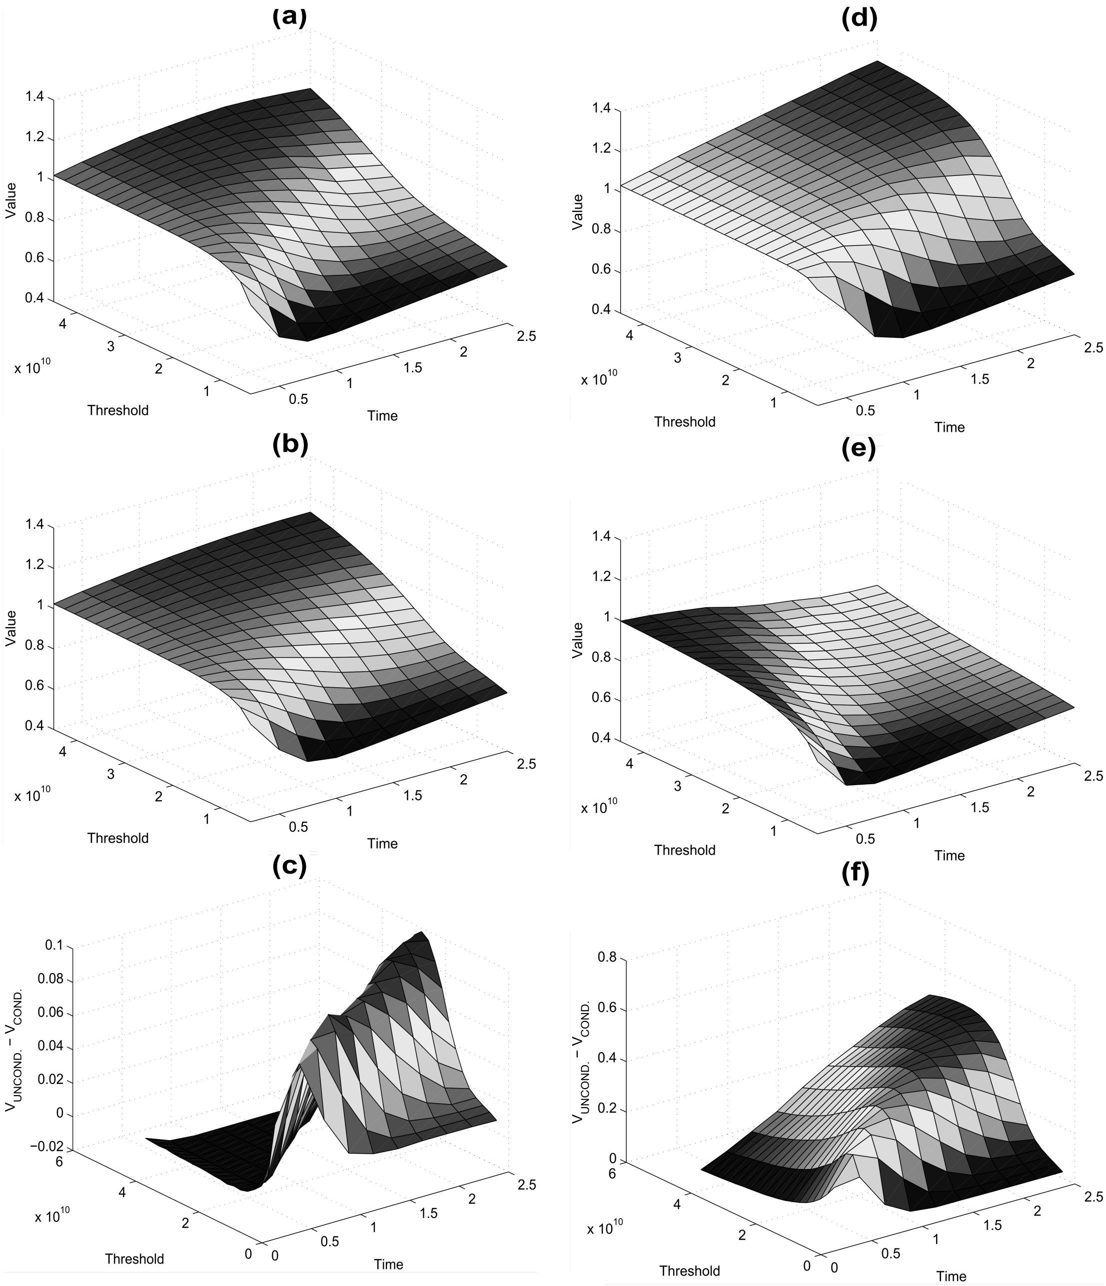

Fig. 4.9: Surfaces for (a) CP CAT bond prices under Burr distribution, fitting via the naive approach; (b) CP CAT bond prices under the Burr distribution, fitting via the CCD approach; (c) the difference between the CP CAT bond prices under the naive approach (V_{UNCOND}) and the CCD approach (V_{COND}) under the Burr distribution; (d) CP CAT bond prices under the MGEV distribution, fitting via the naive approach; (e) CP CAT bond prices under the MGEV distribution, fitting via the CCD approach; (f) the difference between the CP CAT bond prices under the naive approach (V_{UNCOND}) and the CCD approach (V_{COND}) under the MGEV distribution.

Figure 4.9 shows the pricing surfaces in the case of the CP CAT bonds, where simulation from the CCD-fitted Burr and MGEV distributions uses the importance-sampling simulation technique outlined in Section 4.3.5. Similar observations are noted to the ZC CAT bond case, most importantly in that the probability required in Equation (3.15), $\Pr(L_T > D)$, is larger on the basis of the CCD fitting approach compared to the naive fitting approach[11]. However, what is perhaps different to the ZC CAT bond case, specifically in the case of the MGEV distribution, is that there are a number of regions where the CAT bond price in the case of the CCD fitting approach is lower than the price calculated in the case of the naive fitting approach. However, we note that this difference occurs in (theoretical) CAT bond contracts with features that one would rarely find in practice – that is, a very low threshold level with a long term-to-maturity. But on balance, the benefits of the CCD fitting approach, for example from the perspective of a conservative investor valuing his or her assets for the purposes of financial reporting, are clear.

4.6 Concluding remarks

In this chapter, we built on an existing methodology for the fitting of compound time-inhomogeneous Poisson process to general left-truncated, heavy-tailed data. In particular, we studied a problem of identifying underlying distributions by means of the MPS-estimation procedure. A key feature of our methodology is that it avoids unbounded likelihoods and consequent problems with the numerical optimisation thereof. As a natural consequence of the MPS technique, we showed how a goodness-of-fit test statistic, called Moran's log spacings statistic, could be adapted to be used on left-truncated data. We, furthermore, presented a simple importance sampling technique which could be useful in sampling from certain heavy-tailed distributions, in particular the Burr and MGEV distributions, and showed that in implementing such importance sampling a heavier-tailed distribution needs to be used as the sampling distribution.

The proposed methodology was illustrated in the context of CAT bond pricing via the pricing methodology outlined in Chapter 3. We found evidence that naively fitting the compound Poisson process without accounting for the left truncation leads to higher CAT bond prices, compared to CAT bond prices when explicitly accounting for such truncation. Now, insurers and investors operating in the CAT bond environment could employ the CCD fitting approach should they (i) wish not to run the potential risk of overestimating CAT bond prices, and (ii) wish to

[11] The analysis for CP CAT bonds was also performed on the lognormal and GP distributions, and again similar results in respect of the exceedance (i.e. tail) probabilities were noted.

extrapolate insured losses below the PCS index's threshold without any further data to do so.

We stress that this research focuses on analysing the impacts of data mis- specification in light of the left truncation feature of the heavy-tailed PCS data. We point out that in both theory and practice, it is easy to ignore such truncations present in the data, and we furthermore point out that this is prevalent in recent literature. From our calculations and of foremost significance, it is noted that fitting distributions using the CCD approach considerably improved the fit and, moreover, gave rise to finite first moments for the severities, whereas, the naive approach did not. Hence, the fitted severities distributions' usage is not as limited in the CCD-fitting case as opposed to the naive-fitting case. Also, we found the Burr and MGEV severity distributions to fit the PCS data best. In our analysis of CAT bond prices and the impact of data misspecification, it was found that the CCD fitting approach gives rise to loss exceedance (i.e. tail) probabilities which are greater than those estimated by using the naive fitting approach.

We emphasise again that in our research we aimed to present a robust procedure for modelling heavy-tailed left-truncated data as a compound Poisson process, and the PCS dataset was the data to which we chose to apply our procedure. However in practice, our procedure can be applied to other such datasets as well. In fact, application of these computational techniques extends beyond the realm of insurance. These techniques could, for example, find a potential home in left-truncated datasets pertaining to biomedical applications (such as survival data, time-to-relapse or development of a disease (see Cain *et al.* (2011))), but also engineering-related survival modelling (such as modelling future lifetimes of pieces of equipment, given that such a piece has already been used for some time). Of interest may also be applications to operational risk modelling (for example, when data are reported only above some threshold (see Chavez-Demoulin *et al.* (2006); Shevchenko and Temnov (2009))), as well as modelling of insurance claims given the existence of a deductible (see Burnecki, Nowicka-Zagrajek and Wyłomańska (2005)).

4.7 Looking ahead

This chapter formed a foundation for the rest of the work in this thesis. We robustly fitted processes to the underlying drivers of catastrophe bond (and other catastrophe-related ILSs) prices, and we now go on to use these fitted processes in other related applications. In the next chapter, we shall treat the problem of complicated catastrophe bond valuation formulae (particularly those which do not exist in closed-form). In fact, we illustrate an avenue on how we can approximate such complicated valuation formulale, and furthermore explore some of its applications.

Chapter 5

An application of stable weak approximations to index-linked catastrophe bond pricing[1]

5.1 Background

Many catastrophe (CAT) bond pricing formulae rely on the cumulative distribution function of the aggregate loss process (ALP) underlying it. For example, if the ALP is a non-homogeneous compound Poisson process, having time-dependent intensity λ_t, then its cumulative distribution function is a consequence of Lemma 5.1 below[2] if we set $\mathbb{P}(N_t = n) = e^{-\int_0^t \lambda_u \mathrm{d}u} \frac{\left(\int_0^t \lambda_u \mathrm{d}u\right)^n}{n!}$. This cumulative distribution function is impossible to evaluate in closed-form for many severity distributions.

Lemma 5.1. *If $\{L_t, t \geq 0\}$ is a compound renewal process with frequency component $\{N_t, t \geq 0\}$ and i.i.d. severity components $X_1, X_2, \ldots, X_k$ (strictly-positive random variables) for some $k \in \mathbb{N}$, then its cumulative distribution function is specified by (for $s \in \mathbb{R}$)*

$$\mathbb{P}(L_t \leq s) = \begin{cases} \sum_{n=0}^{\infty} \mathbb{P}(N_t = n) F_X^{n*}(s) & \text{if } s > 0 \\ \mathbb{P}(N_t = 0) & \text{if } s = 0. \end{cases} \tag{5.1}$$

where $F_X^{n*}(s)$ represents the n-fold convolution of F_X with itself.

On the basis of the prevalence of ALPs – such as compound Poisson processes – in many CAT bond pricing frameworks, it is important to mention the difficulty in explicitly evaluating functionals of the ALP. However, it is possible to approximate the cumulative distribution function of the compound Poisson process using either Fourier transforms, numerical simulation techniques or approximation

[1] Much of this Chapter is published in Burnecki and Giuricich (2017).

[2] We state Lemma 5.1 without proof.

methods based on the moments and/or cumulants of the severity distribution assumed to underlie the ALP. So far, the literature on CAT bonds appears to mainly focus on numerical simulation techniques, such as Monte Carlo simulation. Monte Carlo simulation and Quasi-Monte Carlo simulation have been applied to numerically evaluating the pricing formulae (Burnecki, Misiorek and Weron, 2005; Nowak and Romaniuk, 2013). But Monte Carlo simulation methods have to be used with caution, since one is attempting to simulate heavy-tailed data (see Section 4.3.5 of Chapter 4), and therefore more advanced simulation techniques such as importance sampling need to be used.

Now, approximation methods for evaluating cumulative distribution functions of the compound Poisson processes have been used in the CAT bond literature to date (see Ma and Ma (2013)). Such approximations of the compound Poisson random variable include the normal approximation (see, for example, Bowers *et al.* (1997, Theorem 12.5.1)), the Edgeworth approximation (Pentikäinen, 1977; Beard, 2013), the gamma approximation (Seal, 1977; Sundt, 1982) , the Inverse-Gaussian approximation (Chaubey *et al.*, 1998) and the Esscher approximation (Esscher, 1932). For a comprehensive treatment of the comparison of their errors, see Seri and Choirat (2015). A problem noticed with the usage of these approximations is that their application is not always possible in cases when the underlying severity distributions have only finite first moments.

Given the aforementioned, it seems that despite attempts to numerically evaluate the ALP-based pricing formulae for CAT bonds, no simple yet approximate pricing formulae have been derived in the case of very heavy-tailed distributional assumptions. So, we endeavour to fill this gap and apply approximative formulae in the context of the IL CAT bond pricing model from Chapter 3.

In this chapter, we consider the class of α-stable distributions and their associated motions. To link the notion of α-stable Lévy motion to catastrophe modelling and CAT bond pricing we present an approximation of the ALP (assumed to underlie CAT bonds) to α-stable Lévy motion by invoking the concept of weak convergence of the former process to the latter. Indeed, this approximation technique arises from the discipline of ruin theory (see Furrer *et al.* (1997), Burnecki (2000) and Michna (2005)) and considers loss severity distributions belonging to the domain of attraction[3] of α-stable laws. We highlight that this approximation explicitly accounts for (and indeed accommodates for) the heavy-tailed nature of the losses (and, in consequence, the loss processes) giving rise to CAT bond price behaviour.

Having this goal in mind, Section 5.2 provides the necessary theoretical framework to construct the weak approximations to α-stable Lévy motion. In Section 5.3,

[3] See Furrer *et al.* (1997, Equation (3)).

we present our results and derivations on the convergence of the ALPs to α-stable Lévy motion. Section 5.4 applies the general results of Section 5.3 to compound Poisson processes. Based on specific candidates for the loss severity distributions for the ALP (taking the form of a compound Poisson process), we then present closed-form approximations for index-linked CAT bond prices. Section 5.5 briefly discusses the application of the techniques developed in Section 5.3 to the IL CAT bond pricing model developed in Chapter 3, and Section 5.6 shows our approximations at work in this practical setting. We then go on to compare our approximations to standard Monte-Carlo simulation exercises for simulating the ALP, and also to first-order single risk loss process approximations to tail probabilities of the ALP. Moreover, in this section we provide some guidance on how to test if our approximations are applicable to the situations considered. Section 5.7 provides a conclusion to this chapter.

5.2 Theoretical framework: weak convergence of the aggregate loss process to α-stable Lévy motion

We contextualise what we have in mind for the mathematical model for the ALP. We firstly consider a renewal process, as specified in Equation (3.4) from Chapter 3, with $\mathbb{E}[T_k] = {}^1\!/\gamma$ ($k \in \mathbb{N}$). In view of Equation (3.4), we simply assume that the renewal process governs the frequency component in the ALP and that successive losses $\{X_k, k \in \mathbb{N}\}$ form a sequence of i.i.d. random variables with $0 < \mathbb{E}[X_k] = \mu < \infty$. We, furthermore, allow for the case when the second and higher moments of the X_k's (and in particular, the variance) can be infinite. Therefore, and as is done in Chapter 3 (see Equation (3.13)), our ALP takes the form

$$L_t = \sum_{k=1}^{N_t} X_k \,, \tag{5.2}$$

which is, in essence, a compound renewal process. The remainder of this section will be devoted to discussing the definitions of α-stable distributions and processes, some assumptions necessary for this chapter and the idea of weak convergence.

5.2.1 α-stable distributions and α-stable Lévy motion

Much research has extensively studied the applicability of Gaussian distributions and processes in stochastic modelling. However, analyses of insurance as well as financial data often points to the existence of heavy tails (see, for example, Embrechts *et al.* (2013) and the recent work by Calderín-Ojeda *et al.* (2017)), calling the use of Gaussian distributions into question. The more general class of stable distributions can account for such a heavy-tailed feature in the data.

The theory regarding univariate stable distributions was developed in the early 1900s, and is covered in detail in the books of Janicki and Weron (1994), Samorodnitsky and Taqqu (1994) and the forthcoming book of Nolan (2015). Despite there being a number of ways to define it[4] we define a stable random variable in terms of its characteristic function to emphasise its parameters.

Definition 5.2. (Stable random variable - see also Samorodnitsky and Taqqu (1994) and Janicki and Weron (1994)). A random variable X is said to have a stable distribution, written $S_\alpha(\sigma_S, \beta_S, \nu_S)$, if there are parameters $0 < \alpha \leq 2$ (the index of stability), $-1 < \beta_S \leq 1$ (the skewness parameter), $\sigma_S > 0$ (the scale parameter) and $\nu_S \in \mathbb{R}$ (the shift) such that its characteristic function has the following form:

$$\mathbb{E}[\exp(i\theta X)] = \begin{cases} \exp\left\{-\sigma_S^\alpha |\theta|^\alpha \left(1 - i\beta_S(\text{sign}(\theta))\tan\left(\frac{\pi\alpha}{2}\right)\right) + i\nu_S\theta\right\} & \text{if } \alpha \neq 1 \\ \exp\left\{-\sigma_S|\theta|\left(1 + i\beta_S\frac{2}{\pi}(\text{sign}(\theta)\ln|\theta|)\right) + i\nu_S\theta\right\} & \text{if } \alpha = 1. \end{cases}$$

Also,

$$\text{sign}(\theta) = \begin{cases} 1 & \text{if } \theta > 0 \\ 0 & \text{if } \theta = 0 \\ -1 & \text{if } \theta < 0. \end{cases}$$

The parameters σ_S, β_S and ν_S are unique, and β_S is irrelevant when $\alpha = 2$.

The notation $X \sim S_\alpha(\sigma_S, \beta_S, \nu_S)$ means that the random variable X has a stable distribution with four parameters $\alpha, \sigma_S, \beta_S$ and ν_S. Furthermore, bear in mind that densities for stable random variables expressed in terms of elementary functions only exist in the cases when $\alpha = 2$ (the normal distribution), $\alpha = 1$ (the Cauchy distribution) and $\alpha = \frac{1}{2}$ (the Lévy distribution). Moreover, note that the rate of decay in the stable distribution depends mostly on the parameter α. We now define what is meant by an α-stable Lévy motion.

Definition 5.3. (α-stable Lévy motion - see Samorodnitsky and Taqqu (1994) and Janicki and Weron (1994)). A stochastic process $\{Z_\alpha(t) : t \geq 0\}$ is called an α-stable Lévy motion if

(i) $Z_\alpha(0) = 0$ a.s.;

[4] See Samorodnitsky and Taqqu (1994, Section 1.1).

(ii) Z_α has independent increments;

(iii) $Z_\alpha(t) - Z_\alpha(s) \sim S_\alpha((t-s)^{1/\alpha}, \beta_S, 0)$ for any $0 \le s < t < \infty$, for some $0 < \alpha \le 2$ and for $|\beta_S| \le 1$.

Observe that the process $Z_\alpha(t)$ introduced in Definition 5.3 has stationary increments, and it is Brownian motion when $\alpha = 2$. We now present some useful properties (see Samorodnitsky and Taqqu (1994) and Janicki and Weron (1994)) which will be relevant to our expositions in Sections 5.3, 5.4 and 5.6.

- *Property 1: existence of moments.* If $X \sim S_\alpha(\sigma_S, \beta_S, \nu_S)$, $\alpha \in (0,2)$ and $p \in (0,\alpha)$ then $\mathbb{E}[|X|^p] < \infty$, and if $p \in [\alpha, 2)$ then $\mathbb{E}[|X|^p] = \infty$.

- *Property 2: tail probability estimation.* If $X \sim S_\alpha(\sigma_S, \beta_S, \nu_S)$ and $\alpha \in (0,2)$, then

$$\begin{cases} \lim_{x\to\infty} x^\alpha \mathbb{P}[X > x] &= C_\alpha \frac{1+\beta_S}{2} \sigma_S^\alpha \\ \lim_{x\to\infty} x^\alpha \mathbb{P}[X < x] &= C_\alpha \frac{1-\beta_S}{2} \sigma_S^\alpha \end{cases}$$

  where

$$C_\alpha = \begin{cases} \frac{1-\alpha}{\Gamma(2-\alpha)\cos(\pi\alpha/2)} & \text{if } \alpha \neq 1 \\ \frac{2}{\pi} & \text{if } \alpha = 1. \end{cases}$$

- *Property 3: self-similarity.* α-stable Lévy motions are $1/\alpha$-self-similar. That is, for all $c > 0$ $\{X_\alpha(ct), t \ge 0\}$ and $\{c^{1/\alpha} X_\alpha(t), t \ge 0\}$ have the same finite- dimensional distributions.

We point out that Property 2 will feature prominently in our work - we will apply it in approximating tail probabilities for α-stable random variables. We do, however, highlight that there are other means to approximate such random variables and processes - see Kohatsu-Higa and Tankov (2010), for example.

Before we present the idea of weak convergence, we introduce an important assumption underlying our work. We assume that our sequence of i.i.d. loss random variables $X_1, X_2, \ldots$ underlying the ALP (see Equation (3.17)), each with finite mean μ, satisfies

$$\frac{1}{\varphi(n)} \sum_{k=1}^{n} (X_k - \mu) \quad \overset{\mathscr{D}}{\Rightarrow} \quad X_\alpha, \tag{5.3}$$

where $\overset{\mathscr{D}}{\Rightarrow}$ denotes convergence in distribution, $\varphi(n) = n^{1/\alpha} L(n)$ (where L is slowly varying at infinity, that is for all $t > 0$ $\lim_{x\to\infty} \frac{L(tx)}{L(x)} = 1$ (Embrechts *et al.*, 2013, Section 1.3.1)), and X_α is an α-stable random variable with $1 < \alpha < 2$. Based on this assumption, we say that our X_k's are in the domain of attraction of an α-stable random variable X_α, the former random variable only having a finite first moment.

5.2.2 Weak convergence of the aggregate loss process

Let $\mathbb{D} := \mathbb{D}[0,\infty)$ denote the space of all real-valued càdlàg functions on $[0,\infty)$ endowed with the Skorokhod topology[5]. Then $(\mathbb{D}, J_1)$ is a separable and complete metric space[6], where J_1 is the Skorokhod metric as defined in Skorokhod (1957).

Definition 5.4. (Weak convergence - see Billingsley (1995)) A sequence $\{X^{(n)} : n \in \mathbb{N}\}$ of stochastic processes is said to converge weakly in $(\mathbb{D}, J_1)$ to a stochastic process X if for every bounded functional f on $\mathbb{D}$ it follows that

$$\lim_{n\to\infty} \mathbb{E}\left[f(X^{(n)})\right] = \mathbb{E}[f(X)]. \tag{5.4}$$

When a sequence of stochastic processes satisfies Equation (5.4), we will write $X^{(n)} \overset{J_1}{\Rightarrow} X$.

We now turn to showing that the ALPs we consider in this chapter do converge to a stable Lévy motion, and moreover apply the results of Furrer *et al.* (1997) and Michna (2005) to ALPs. For our sequence of i.i.d. losses $\{X_k\}_{k=1}^{\infty}$ satisfying Equation (5.3) and $\{N(nt)\}_{n=1}^{\infty}$ a sequence of renewal processes, if

$$\frac{N(nt) - \lambda nt}{\varphi(n)} \quad \to \quad 0, \quad \text{as } n \to \infty \tag{5.5}$$

in probability in $(\mathbb{D}, J_1)$, for $1 < \alpha < 2$ and for a positive constant λ, then

$$\frac{1}{\varphi(n)} \sum_{k=1}^{N(nt)} (X_k - \mu) \quad \overset{J_1}{\Rightarrow} \quad \lambda^{1/\alpha} Z_\alpha(t) \quad \text{as } n \to \infty, \tag{5.6}$$

where $\{Z_\alpha(t), t > 0\}$ is an α-stable Lévy motion (Furrer *et al.*, 1997). Observe that if L is a constant, say d, then we may rewrite Equation (5.6) as

$$\frac{1}{n^{1/\alpha}} \sum_{k=1}^{N(nt)} (X_k - \mu) \quad \overset{J_1}{\Rightarrow} \quad d\lambda^{1/\alpha} Z_\alpha(t) \quad \text{as } n \to \infty. \tag{5.7}$$

Note that in light of the definition of the renewal process given by Equation (3.4), we can set $\lambda = \gamma$ in Equations (5.5) to (5.7) since N_t is a renewal process.

[5] See Lindvall (1973).
[6] See Lindvall (1973) for further details.

Now, the relationship specified in Equation (5.7) is crucial to our application of the weak convergence to an α-stable Lévy motion. However, Equation (5.7) provides no guidance on how to explicitly find the constant d and the explicit value of α. We now present a less general version of Nolan (2015)'s generalised central limit theorem, which shall assist in finding d. Suppose that $\{X_1, X_2, X_3, \ldots\}$ is an i.i.d. sequence of random variables each having common distribution function, F, firstly with $\mathbb{E}[X_k] = \mu$, secondly exhibiting tail probabilities that satisfy $\lim_{x\to\infty} x^{\alpha}(1-F(x)) = c^+ \geq 0$ and $\lim_{x\to\infty} x^{\alpha} F(-x) = c^- \geq 0$ for $1 < \alpha < 2$ and thirdly with $c^+ + c^- > 0$. Then

$$\frac{(X_1 + \ldots + X_n) - b_n}{a_n} \quad \stackrel{\mathscr{D}}{\Rightarrow} \quad S_\alpha(1, \beta_S, 0).$$

The sequences a_n and b_n, as well as the the constant β_S are given by

$$a_n = n^{1/\alpha}\left(\frac{\pi(c^+ + c^-)}{2\Gamma(\alpha)\sin\left(\frac{\alpha\pi}{2}\right)}\right)^{1/\alpha}, \; b_n = n\mu \;\text{ and }\; \beta_S = \frac{c^+ - c^-}{c^+ + c^-}.$$

Notice that the sequence $\{a_n\}_{n=1}^{\infty}$ is crucial in finding d. That is, $d = a_n / n^{1/\alpha}$ which depends on the distributional form of the X_k's. Moreover, it is clear that given a severity distribution (underlying the ALP), α will be equal to a function of its parameters.

We now present some comments on the applicability of Equation (5.7). In selecting ALPs to model the underlying catastrophe risk inherent in CAT bonds, heavy-tailed severity distributions are often selected. So, in the forthcoming sections of this chapter, we shall restrict ourselves to the application of the generalised central limit theorem to heavy-tailed distributions assumed to belong to the domain of attraction of an α-stable distribution with $1 < \alpha < 2$. Thus, the mean of these distributions is finite, however the higher-order moments need not be finite. We also point out that while classical Normal distribution (see Bowers *et al.* (1997, Theorem 12.5.1)) and also standard Brownian motion approximations to ALPs, (see, for example, Embrechts *et al.* (2013)) require the assumption of light-tailed underlying severity distributions, this assumption can be relaxed in the more general context of α-stable Lévy motion approximations.

5.3 Weak approximation to tail probabilities of compound renewal processes

Index-linked CAT bond pricing most often requires the evaluation of exceedance (or tail) probabilities such as $\mathbb{P}(L_T > D)$, where D is some positive constant and L_T is a kind of compound renewal process. An example of this is in the case of an index-linked CAT bond, considered in Section 3.4, that has some associated threshold level (i.e. D). Within the realm of their index-linked CAT bond pricing framework, Ma and Ma (2013) considered a mixed approximation to this tail probability. However their approximation, which relies on higher-order moments and cumulants of the loss process, was not always valid since many of their fitted severity distributions did not possess finite second-order (and higher) moments and cumulants. We now consider how to overcome this problem, and therefore approximate index-linked CAT bond prices (under heavy-tailed distributional assumptions) by invoking the convergence of the underlying compound renewal process to α-stable Lévy motion.

We emphasise that we consider the case when the severity distribution lies in the domain of attraction of an α-stable distribution with $1 < \alpha < 2$, which implies that only the first moment is finite. Therefore, we present a theorem which allows us to weakly approximate the probability $\mathbb{P}(L_T > D)$. The idea is to consider the compound renewal process over the time period $[0, T]$, and replace it by a suitable α-stable Lévy motion, and within this sphere apply known results on tail probability estimation to find a closed-form expression for the required probability. Our result is born from Furrer *et al.* (1997)'s Proposition 3 as well as from further work in Michna (2005), but with the premium rate in both their considered risk processes being set to zero. In consequence of the weak convergence, we only obtain one process in the limit (see Equations (5.5) and (5.6)).

Theorem 5.5. *Let $\{N_t : t \geq 0\}$ be a renewal process as specified in Equation 3.4, suppose that $1 < \alpha < 2$ for each of the X_k's and let the constant D be represented as $D = \gamma T \mu + M \gamma^{\frac{1}{\alpha}}$, where $M \in \mathbb{R}^+$. Then we have, for T a positive constant, that*

$$\lim_{\gamma \to \infty} \mathbb{P}\left(D - \sum_{k=1}^{N(T)} X_k < 0 \right) \quad \sim \quad C_\alpha \frac{1 + \beta_S}{2} T (d/M)^\alpha, \quad M \to \infty, \tag{5.8}$$

where $C_\alpha = (1 - \alpha)/\{\Gamma(2 - \alpha) \cos(\pi\alpha/2)\}$, β_S is the skewness parameter of the limiting α-stable distribution and where d is a constant that depends upon the distributional form of the X_k's.

Proof. For T a positive constant, define the sequence of random variables $(Q^{(\gamma)}(T), \gamma \in \mathbb{N})$ in L^1 as follows:

$$Q^{(\gamma)}(T) = \frac{\sum_{k=1}^{N(T)} X_k - D}{\gamma^{1/\alpha}}.$$

By a time transformation on N, we can specify that

$$\begin{aligned} Q^{(\gamma)}(T) &= \frac{\sum_{k=1}^{\tilde{N}(\gamma T)} X_k - D}{\gamma^{1/\alpha}} \\ &= \frac{1}{\gamma^{1/\alpha}} \sum_{k=1}^{\tilde{N}(\gamma T)} (X_k - \mu) + \mu \left(\frac{\tilde{N}(\gamma T) - \gamma T}{\gamma^{1/\alpha}} \right) + \frac{\gamma T \mu}{\gamma^{1/\alpha}} - \frac{D}{\gamma^{1/\alpha}}. \end{aligned}$$

where $\tilde{N}_t$ is a renewal process with unit intensity and $\mu = \mathbb{E}[X_k]$. By Equation (5.5), the second term converges to 0 in probability and by Equation (5.6), the first term converges weakly to $dZ_\alpha(T)$, both as $\gamma \to \infty$. The remaining terms converge to $-M$. As a consequence, $Q^{(\gamma)}(T)$ converges weakly to $dZ_\alpha(T) - M$. Since the random variable $dZ_\alpha(T) - M$ is continuous

$$\lim_{\gamma \to \infty} \mathbb{P}\left(Q^{(\gamma)}(T) > 0\right) = \mathbb{P}\left(dZ_\alpha(T) - M > 0\right).$$

By the self-similarity property of $Z_\alpha(t)$ (Property 3), it can be shown that

$$\mathbb{P}\left(dZ_\alpha(T) - M > 0\right) = \mathbb{P}\left(Z_\alpha(1) > \frac{M}{dT^{1/\alpha}}\right).$$

Finally, the statement follows by invoking the tail probability estimation of an α-stable random variable (Property 2). □

5.4 Application to compound Poisson processes and selected severity distributions

Notice that in Sections 5.2 and 5.3, we assumed that the ALP took the form of a compound renewal process. The ALPs that index-linked CAT bond pricing formulae commonly reference are compound Poisson processes (which are special cases of compound renewal processes), and in this section (and for the remaining parts of the chapter) we suppose that L_t is indeed a compound Poisson process. Moreover, in the context of index-linked CAT bond pricing note that underlying compound Poisson processes often assume a (heavy-tailed) parametric form for the underlying severity distribution component.

5.4.1 Selected parametric severity distributions

Three types of severity distributions are considered, and are all assumed to belong to the domain of attraction of an α-stable random variable (on the basis of their parameters). Their parameterisations are shown in Table 5.1 for convenience. Notice that the modified generalised extreme value distribution considered here is the classical generalised extreme value distribution with location parameter μ set equal to σ/k (see Section 4.3.3 of Chapter 4). This action is taken so that this distribution is defined on the positive real numbers, a support that is consistent with the modelling of insurance losses.

Tab. 5.1: Parametric severity distributions considered, defined on the strictly positive real numbers.

Distribution	Parameterisation	Constraints
Burr type XII	$\text{BURR}(\zeta,c,k)=\frac{\frac{kc}{\zeta}\left(\frac{x}{\zeta}\right)^{c-1}}{\left\{1+\left(\frac{x}{\zeta}\right)^{c}\right\}^{k+1}}$	$\zeta,c,k>0$
Generalised Pareto	$\text{GP}(k,\sigma)=\frac{1}{\sigma}\left(1+\frac{kx}{\sigma}\right)^{-\left(1+\frac{1}{k}\right)}$	$\sigma>0, k\in\mathbb{R}$
Modified generalised extreme value	$\text{MGEV}(k,\sigma)=$ $\frac{1}{\sigma}\exp\left\{-\left(\frac{k}{\sigma}x\right)^{-\frac{1}{k}}\right\}\left(\frac{k}{\sigma}x\right)^{-1-\frac{1}{k}}$	$\sigma>0, k\in\mathbb{R}$, $k\neq 0$

5.4.2 Compound Poisson processes

We now apply Theorem 5.5 to both time-homogeneous and time-inhomogeneous Poisson processes, each which can underlie the considered ALP, L_t. The following two corollaries specify the results for calculating the required tail probabilities, $\mathbb{P}(L_T>D)$.

Corollary 5.6. *(to Theorem 5.5). Let $\{N_t^{HP}: t\geq 0\}$ be a time-homogeneous Poisson process with constant intensity λ, suppose that $1<\alpha<2$ for each of the X_k's and let D be represented as $D=\lambda T\mu+M\lambda^{\frac{1}{\alpha}}$, where $M\in\mathbb{R}^+$. Then, for a positive constant T*

$$\lim_{\lambda\to\infty}\mathbb{P}\left(D-\sum_{k=1}^{N^{HP}(T)}X_k<0\right)\quad\sim\quad C_\alpha\frac{1+\beta_S}{2}T(d/M)^\alpha,\quad M\to\infty. \tag{5.9}$$

Proof. By setting $\gamma=\lambda$, the statement is an immediate consequence of Theorem 5.5. □

Corollary 5.7. *(to Corollary 5.6). Let $\{N^{INHP}(t) : t \geq 0\}$ be a time-inhomogeneous Poisson process with a time-dependent intensity $\lambda(t) > 0$, suppose that $1 < \alpha < 2$ for each of the X_k's and let D be represented as $D = \mu T\Lambda + M\Lambda^{\frac{1}{\alpha}}$, where $M \in \mathbb{R}^+$ and $\Lambda = \frac{1}{T}\int_0^T \lambda(s)ds$. Then, for a positive constant T,*

$$\lim_{\Lambda\to\infty} \mathbb{P}\left(D - \sum_{k=1}^{N^{INHP}(T)} X_k < 0\right) \sim C_\alpha \frac{1+\beta_S}{2} T(d/M)^\alpha, \quad M \to \infty. \tag{5.10}$$

Proof. Let us define a time-homogeneous Poisson process $\{N_t^{HP} : t \geq 0\}$ with constant intensity $\Lambda = \frac{1}{T}\int_0^T \lambda(s)\mathrm{d}s$. This process at time T has the same distribution as the process N^{INHP}. Hence, by Corollary 5.6 we obtain the thesis. □

Notice that Corollaries 5.6 and 5.7 are applicable when the X_k follow a Burr, generalised Pareto (GP) or modified generalised extreme value (MGEV) distribution each belonging to the domain of attraction of an α-stable random variable – that is, each of the distributions satisfy Equation (5.3). Indeed, it was noted from Section 4.3.4 in Chapter 4 that the Burr and MGEV distributions, as well as the GP distribution to a lesser extent, fitted[7] the severity data. In order to apply the corollaries to the Burr, GP and MGEV distributions, it is necessary to find the associated index of stability α as well as the associated skewness parameter β_S for each distribution, and then find d. Under the assumption that $1 < \alpha < 2$, this was all done by applying the generalised central limit theorem from Section 5.2.2, and Table 5.2 provides the relevant values found. Notice that d is the same in the case of both the GP and MGEV distributions.

Tab. 5.2: Values of α, β_S and d for the selected severity distributions.

Distribution	α	β_S	d
BURR(ζ, c, k)	ck	1	$\zeta \left\{ \frac{\pi}{2\Gamma(ck)\sin\left(\frac{ck\pi}{2}\right)} \right\}^{1/ck}$
GP(k, σ)	$1/k$	1	$\frac{\sigma}{k} \left\{ \frac{\pi}{2\Gamma(\frac{1}{k})\sin\left(\frac{\pi}{2k}\right)} \right\}^{k}$
MGEV(k, σ)	$1/k$	1	$\frac{\sigma}{k} \left\{ \frac{\pi}{2\Gamma(\frac{1}{k})\sin\left(\frac{\pi}{2k}\right)} \right\}^{k}$

[7] The lognormal distribution also fitted the data to a lesser extent, but in addition to a finite mean possessed a finite variance as well. When using the lognormal distribution as a candidate for the severity distribution for the ALP, approximating tail probabilities of the ALP becomes a trivial exercise since arguments based on the classical central limit theorem can be used (see, for example, the approximation given in Bowers *et al.* (1997, Theorem 12.5.1). We do not consider such approximations in this Chapter.

5.5 Application to index-linked catastrophe bond pricing

In this section, we briefly speak about the application of the previous sections to the index-linked CAT bond pricing model outlined in Chapter 3. We point out that the approximation formulae derived in Section 5.3 – namely Equations (5.9) and (5.10) – are directly applicable in the context of this CAT bond pricing model (see Equations (3.15) and (3.17) in Chapter 3).

Before proceeding, take note of the less general model setting we are now operating in. We assume that the ALP, L_t, is a compound time-inhomogeneous Poisson process with frequency component $\{N_t, t > 0\}$ having intensity function $\lambda_t > 0$ and the severity distribution component assumed to be as in Section 5.4. However, we do point out that the usage of our approximation formulae (namely Equation (5.8)) are not limited to the case where L_t is a compound Poisson process. They can also be used in the more general case where L_t is a compound renewal process, and the index-linked CAT bond pricing framework, from Chapter 3, we use can accommodate for this generalisation.

5.6 Numerical illustration and comparison exercise

5.6.1 Preliminary remarks

In order to better understand the behaviour of our weak approximation in the context of our index-linked CAT bond pricing model from Chapter 3, we give a numerical illustration. Moreover, we assess the behaviour of our approximation on the basis of heavy-tailed distributions commonly used in the literature on index-linked CAT bond pricing.

We point out that the weak approximations for the compound Poisson process found in Corollaries 5.6 and 5.7, can be inserted into the CAT bond pricing formulae specified by Equations (3.15) and (3.17), assuming that the constants D and D_{CP} are indeed the threshold levels of the index-linked CAT bonds. We do so, and numerically compare our CAT bond prices obtained via these simple approximations to firstly Monte-Carlo (MC) estimated prices (see Chapter 4 Section 4.4.2 for an outline of the simulation procedure) and secondly to a simple first-order single risk loss process (FSRLP) approximations. The idea behind the FSRLP is that if we assume the i.i.d. severities, X_i's, are sub-exponential[8] and that the number of losses N_T is Poisson distributed over the time interval of interest (with intensity λ_t), then

[8] See Embrechts *et al.* (2013, Definition 1.3.3): a distribution function F, with support $(0, \infty)$ is sub-exponential if for all $n \geq 2$, $\lim_{x \to \infty} \frac{1-F^{n*}(x)}{1-F(x)} = n$, where F^{n*} denotes the $n-$fold convolution of F with itself.

(see Daley *et al.* (2007); Stam (1973) and Peters and Shevchenko (2015)),

$$\mathbb{P}(L_T > D) \approx \mathbb{E}[N_T]\,\mathbb{P}(X_1 > D).$$

We re-highlight that the aim of our work in this section is two-fold. Firstly, it is to compare the performance of our weak approximation against Monte Carlo simulation (which, in the limit, will give the exact price under the pricing model used). Secondly, it is also a first foray into the comparison of the performance of our weak approximation to that of the FSRLP approximation, each relative to Monte Carlo simulation. We do not, however, wish to conclude, in general, on the relative superiority or inferiority of our weak approximation to the FSRLP approximation.

For the purposes of this illustration, we consider the ALP fitted in Section 4.3 of Chapter 4, i.e. a compound time-inhomogeneous Poisson process with a time-dependent intensity. That is, the fitted intensity function is

$$\lambda(t) = 24.93 + 0.026t + 5.61\sin\{2\pi(t+7.07)\} + 10.30\exp\left\{\cos\left(\frac{2\pi t}{4.76}\right)\right\}. \tag{5.11}$$

We use Equation (5.11) in conjunction with each of the considered severity distributions from Table 5.1. For the severity distributions considered in Table 5.1, we adopt the parameters fitted by the conditional complete-data (CCD) approach in Section 4.3 of Chapter 4. We stress that these parameters are based on Property Claims Services (PCS) insured loss data – data that are indeed left-truncated – so we select the CCD-fitted parameters in order to account for the left-truncation feature. These (CCD-approach estimated) parameters were shown in in Table 4.3 in Chapter 4.

For the CCD-approach fitted GP, Burr and MGEV distributions, we point out that their associated α's are all, *a fortiori*, within the range $(1,2)$, and in consequence all these distributions belong to the domain of attraction of an α-stable distribution with $1 < \alpha < 2$. Therefore, all distributions have only finite first moments. However, notice that for the MGEV distribution the power law exponent is very close to 1 that it will produce numerical errors in simulation[9]. Because of this, we omit the MGEV distribution from further analysis.

5.6.2 Is λ (or γ) large enough for Theorem 5.5 to apply?

In any attempt to apply Theorem 5.5 to the PCS data, an immediate question arises. Is the fitted intensity λ, be it constant or (as in Equation (5.11)) time- dependent,

[9] This was indeed checked and such numerical errors did arise. For a further discussion on the prevalence of these numerical errors, see Burnecki, Misiorek and Weron (2005).

satisfactorily large enough? Note that for the PCS data set we study, the constant intensity is approximately 27.91 events per year, while the time-dependent fitted intensity function begins at 24.93. We checked this issue by analysing if for the fitted λ, $\mathbb{P}(Q^{(\gamma)}(T) > 0)$ is sufficiently close to $\mathbb{P}(dZ_\alpha(T) - M > 0)$ for each of the Burr and GP severity distributions for fixed threshold levels. To estimate the latter probability, we used the algorithm of Nolan (1997), while we estimated the former probability via $100,000$ Monte Carlo simulations.

We begin our analysis in this section by bearing in mind the assumption that the PCS data follows a compound Poisson process. Furthermore, we suppose that the term of the CAT bond, T, is fixed at 2.5 years. If, for the fitted λ, $\mathbb{P}(Q^{(\gamma)}(T) > 0)$ is sufficiently close to $\mathbb{P}(dZ_\alpha(T) - M > 0)$ for each of the Burr and GP severity distributions for fixed threshold levels D, then we could argue that Theorem 5.5 is applicable. Consider Figure 5.1 below, which plots these probabilities for each of the severity distributions. Note, again, that $\mathbb{P}(Q^{(\gamma)}(T) > 0)$ is estimated via $100,000$ Monte Carlo simulations, while $\mathbb{P}(dZ_\alpha(T) - M > 0)$ is calculated, to a tolerance of 1×10^{-8}, using the algorithm of Nolan (1997).

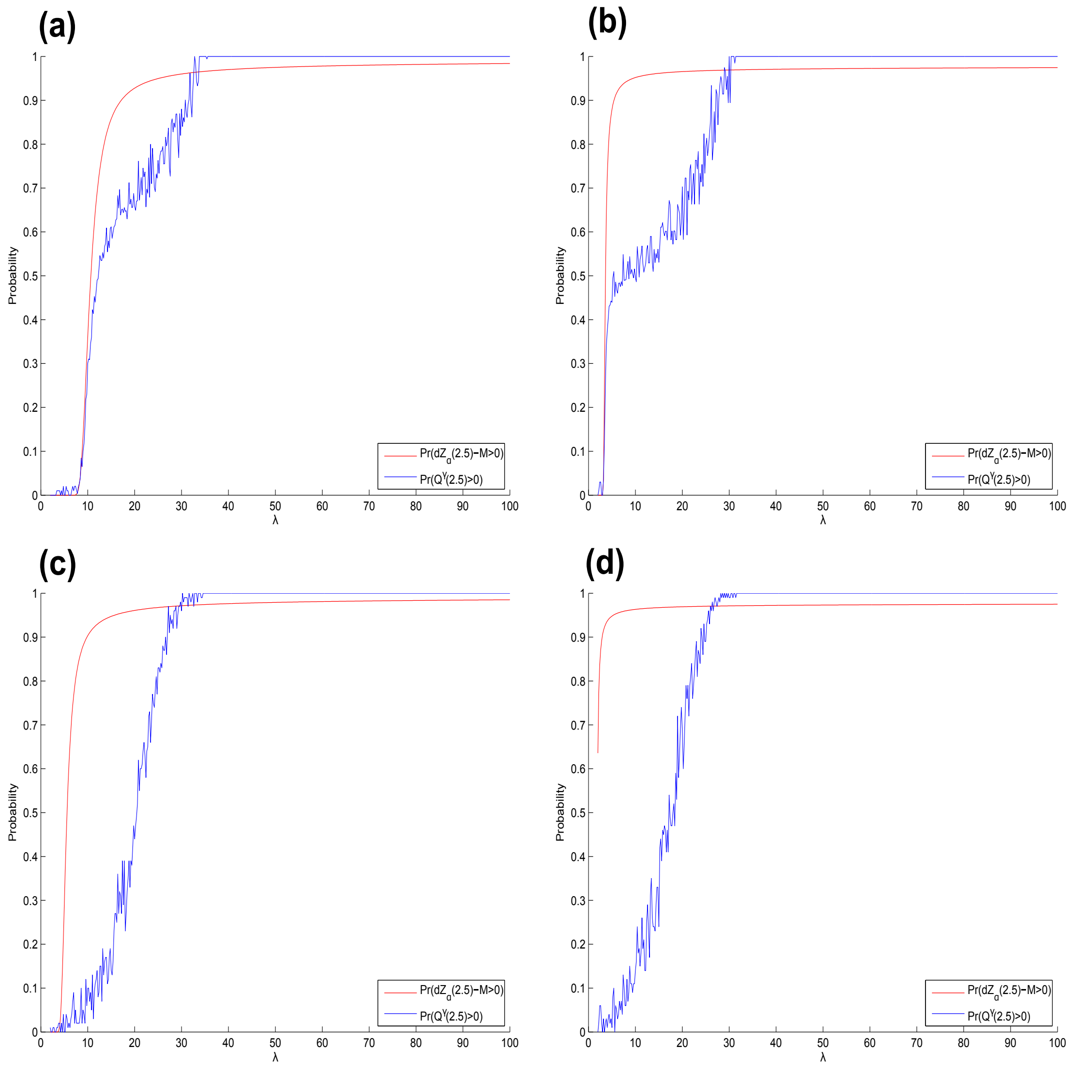

Fig. 5.1: Convergence of $\mathbb{P}(Q^{(\gamma)}(2.5) > 0)$ to $\mathbb{P}(dZ_\alpha(2.5) - M > 0)$ for a range of values of λ under different assumptions for the ALP and D. For $D = \$7.8 \times 10^{10}$, (a) illustrates the case when the severity distribution is GP and (b) when it is Burr. For $D = \$1.45 \times 10^{11}$, (c) illustrates the case when the severity distribution is GP and (d) when it is Burr.

For our data (and as is evinced in Figure 5.1), as λ tended to infinity, $\mathbb{P}(Q^{(\gamma)}(2.5) > 0)$ converged to $\mathbb{P}(dZ_\alpha(2.5) - M > 0)$ as expected. At $\lambda = 24.93$, the two probabilities appeared to be sufficiently close to one another for both severity distributions: $\mathbb{P}(dZ_\alpha(2.5) - M > 0)$ lay within a three standard deviation error bound of $\mathbb{P}(Q^{(\gamma)}(2.5) > 0)$ for each distribution. But also, the approximation appeared to be fairly accurate for all values of λ in excess of 30, and also for low values of λ less than 10.

On balance, and under the assumption that the PCS dataset follows a compound Poisson process, it would appear that the use of Theorem 5.5 is suitable. However, we do caution that its usage does not lead to an exact result, but rather a suitable approximation. We end by emphasising that if Theorem 5.5 is applied to other data sets, it is important to check that γ is large enough in order to apply it.

5.6.3 Numerical simulation results and comparisons

Tables 5.3 and 5.4 show some numerical values, for illustrative purposes, for the case of pricing index-linked ZC and CP CAT bonds (at time 0). The CAT bond threshold level is assumed to be in line with that used by Ma and Ma (2013) (that is $D, D_{\text{CP}} \in [3,740; 44,880]$ million), and we set $T \in (0,2]$ years, $\rho = 0.5$ and $r = 0.06$ in Equations (3.15) and (3.17). After approximating prices via our weak approximation and the FSRLP approximation, and estimating prices by Monte Carlo simulation, we calculate relative errors. The relative errors of the weak approximation prices for each of the ZC and CP IL CAT bonds ($V_0^{\text{ZC, CP (WEAK)}}$) to Monte Carlo simulation prices (i.e. $\varepsilon^{\text{WEAK}}$) and secondly of the FSRLP approximated prices ($V_0^{\text{ZC, CP (FSRLP)}}$) to Monte Carlo simulation prices (i.e. $\varepsilon^{\text{FSRLP}}$)are calculated by $\left(V_0^{\text{ZC, CP (WEAK)}} - V_0^{MC}\right)/V_0^{MC}$ and $\left(V_0^{\text{ZC, CP (FSRLP)}} - V_0^{MC}\right)/V_0^{MC}$ respectively. Note that we used $N = 100,000$ Monte Carlo simulations.

Before we present the results, note that the pricing surfaces (based on our weak approximations), and relative error comparisons (to Monte Carlo simulation), are shown in Figures 5.2 and 5.3. More precise relative error comparisons, for $T = 1$ year, for each severity distribution (Burr and GP) are illustrated in Figure 5.4.

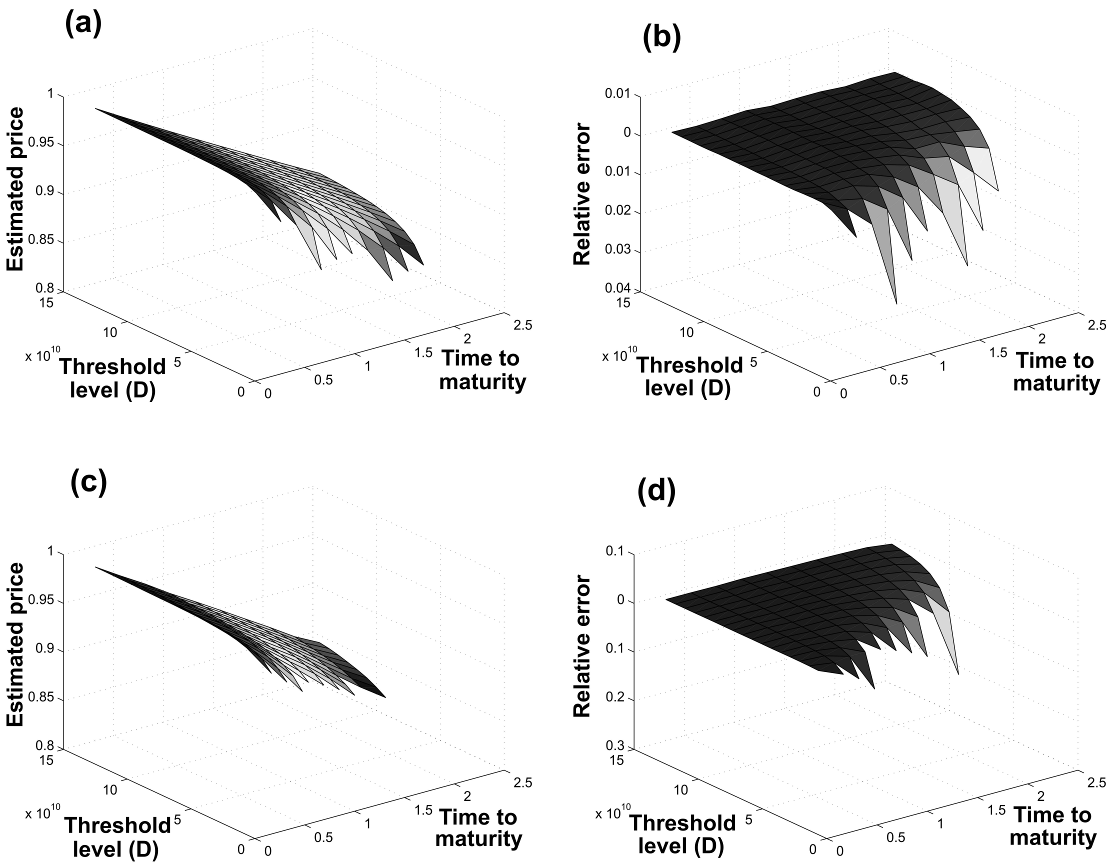

Fig. 5.2: Pricing surfaces, obtained by using the stable weak approximations, for index-linked ZC CAT bonds under the assumption of (a) GP distributed losses (and the relative error to Monte Carlo estimation in Panel (b)); and (c) Burr distributed losses (and the relative error to Monte Carlo estimation in Panel (d)). Note the omission of points where the weak approximations were non-applicable – hence affording the figures their jagged look.

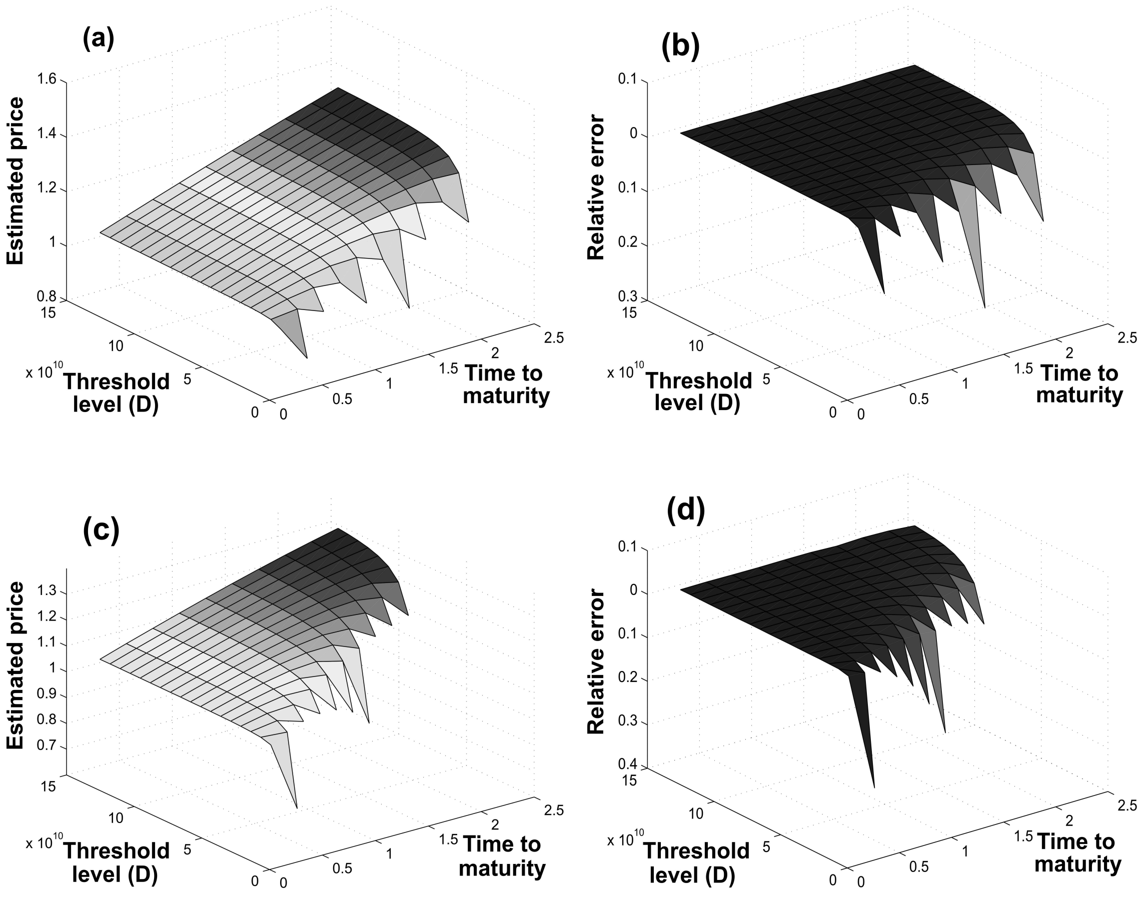

Fig. 5.3: Pricing surfaces, obtained by using the stable weak approximations, for index-linked CP CAT bonds under the assumption of (a) GP distributed losses (and the relative error to Monte Carlo estimation in Panel (b)); and (c) Burr distributed losses (and the relative error to Monte Carlo estimation in Panel (d)). Note the omission of points where the weak approximations were non-applicable – hence affording the figures their jagged look.

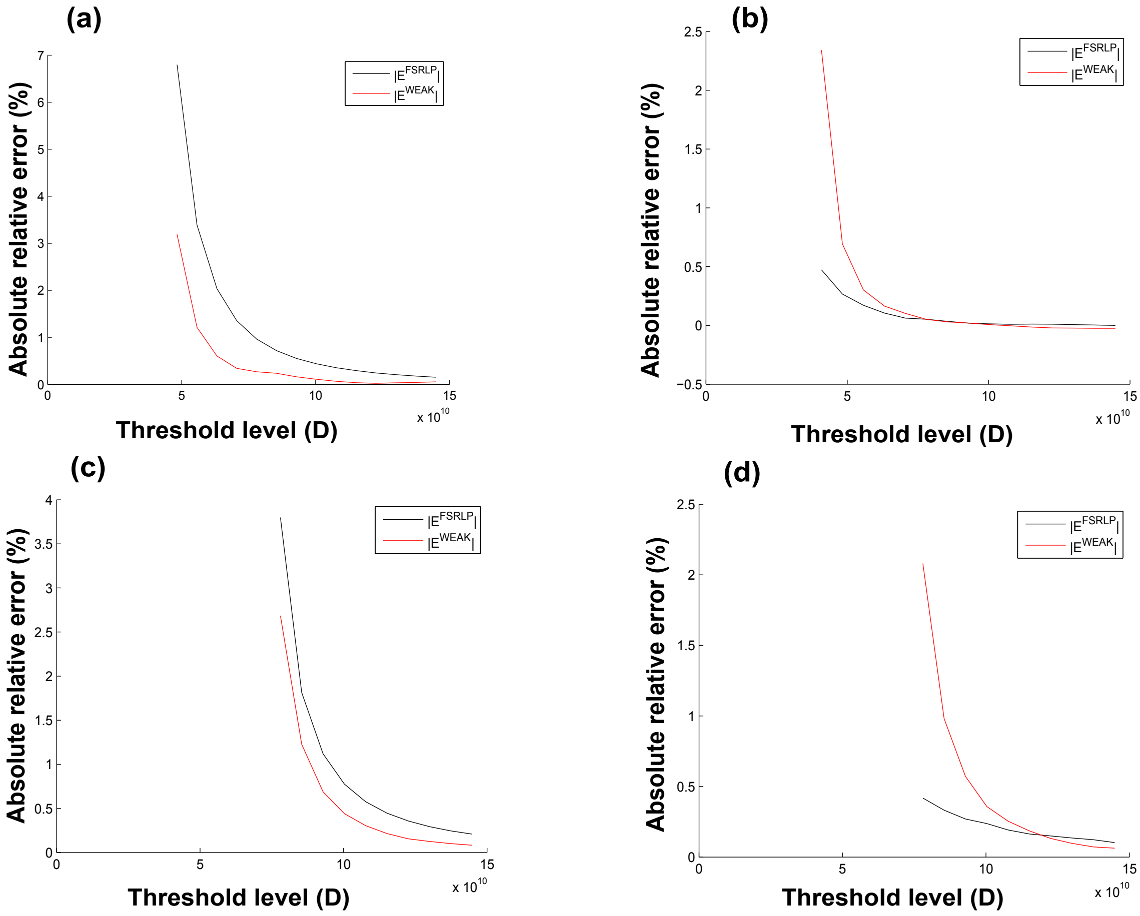

Fig. 5.4: Absolute relative error comparisons for each severity distribution, for (a) GP distributed losses in the ZC CAT bond case (and the CP case in (b)); and (c) Burr distributed losses in the index-linked ZC CAT bond case (and the index-linked CP CAT bond case in (d)), for $T = 1$ year. Note the different threshold levels, for each severity distribution, from which the weak approximation became applicable.

Tab. 5.3: Comparison of index-linked ZC CAT bond prices, under the weak approximation, Monte Carlo estimation and the FSRLP approximation, for different threshold levels D and different underlying loss severity distributions. $\hat{\sigma}_{MC}$ is the standard error of the Monte Carlo estimate.

D	**Severity distribution**	**MC price** $\pm 3\frac{\hat{\sigma}_{MC}}{\sqrt{N}}$	$\lvert\varepsilon^{WEAK}\rvert$ **(%)**	$\lvert\varepsilon^{FSRLP}\rvert$ **(%)**
7.8×10^{10}	GP	$0.8697 \pm 1.82 \cdot 10^{-3}$	0.3	1.0
1.45×10^{11}	GP	$0.8829 \pm 8.93 \cdot 10^{-4}$	0.06	0.3
8.61×10^{12}	GP	$0.8828 \pm 2.21 \cdot 10^{-4}$	≈ 0	0.01
7.8×10^{10}	Burr	$0.8345 \pm 3.04 \cdot 10^{-3}$	2.8	4.1
1.45×10^{11}	Burr	$0.8722 \pm 1.69 \cdot 10^{-3}$	0.4	0.8
8.61×10^{12}	Burr	$0.8722 \pm 4.57 \cdot 10^{-3}$	≈ 0	≈ 0

Tab. 5.4: Comparison of index-linked CP CAT bond prices, under the weak approximation, Monte Carlo estimation and the FSRLP approximation, for different threshold levels D_{CP} and different underlying loss severity distributions. $\hat{\sigma}_{MC}$ is the standard error of the Monte Carlo estimate.

D_{CP}	**Severity distribution**	**MC price** $\pm 3\frac{\hat{\sigma}_{MC}}{\sqrt{N}}$	$\lvert\varepsilon^{WEAK}\rvert$ **(%)**	$\lvert\varepsilon^{FSRLP}\rvert$ **(%)**
7.8×10^{10}	GP	$1.2467 \pm 1.61 \cdot 10^{-3}$	0.1	0.1
1.45×10^{11}	GP	$1.2571 \pm 8.29 \cdot 10^{-4}$	≈ 0	≈ 0
8.61×10^{12}	GP	$1.2570 \pm 9.77 \cdot 10^{-4}$	≈ 0	≈ 0
7.8×10^{10}	Burr	$1.2035 \pm 3.74 \cdot 10^{-3}$	0.3	0.5
1.45×10^{11}	Burr	$1.2432 \pm 2.17 \cdot 10^{-3}$	0.21	0.25
8.61×10^{12}	Burr	$1.2431 \pm 6.84 \cdot 10^{-3}$	≈ 0	≈ 0

The behaviour of our approximation appears to vary according to the underlying severity distributional assumption. It appears, that within the context of our numerical calculations, the closer to unity the power law exponent of the severity distribution, the less well-behaved the approximation (compared to Monte Carlo estimation and the FSRLP approximation) for low threshold levels and times to maturity. This was indeed the case with the MGEV distribution, hence the reason for its omission

We now consider the cases of the fitted GP and Burr distributions, which both have associated α's further away from unity (1.099 and 1.370 respectively). Looking at Tables 5.3 and 5.4 as well as Figures 5.2 and 5.3, it seems clear that for high threshold levels and long terms to maturity, the weakly approximated and the Monte Carlo-estimated CAT bond prices are similar. This phenomenon is particularly evinced as the threshold levels D and D_{CP} tend to infinity. We also note that the weakly approximated prices, in both the context of the IL ZC and the CP CAT bonds, are similar to not only the Monte Carlo-estimated prices but also the prices based on the FSRLP approximation, an approximation that is known to work quite well in the context of compound Poisson process tail probability estimation (see Peters and Shevchenko (2015, Chapter 7).). In fact, the weakly approximated prices are, mostly, closer to the Monte Carlo estimated prices on the basis of absolute relative error (as the threshold level approaches infinity) than the FSRLP approximated prices (particularly in the case of the ZC CAT bonds). This may be so because our approximation is indeed useful in the case of heavy-tailed data. However, the panels in Figure 5.4 demonstrate that this behaviour is not always the case for lower threshold levels (for all severity distributions): for index-linked ZC CAT bonds the weak approximation is better in relative error, while for index-linked CP CAT bonds the FSRLP is better in relative error. Given that the approximation is used more often in the index-linked CP CAT bond pricing formula (see Equation (3.17)) and also at different durations in time, we put forward that the weak approximation fares more poorly at shorter time durations than longer durations, for lower threshold levels. This appears plausible from the point-of-view that threshold exceedance (i.e. tail) probabilities are vanishingly small for shorter and shorter durations. However, the FSRLP performed better at shorter durations for lower threshold levels (i.e. lower than 7.8×10^{10}) - further numerical analysis into this indeed did reveal that this was the case.

We now remark on the following two challenges as regards our weak approximation usage in CAT bond pricing. Firstly, we noticed in our numerical analysis that the resultant approximations of Theorem 5.5 led to the calculation of probabilities in excess of one. We omit these price calculations from the pricing surfaces in Figures 5.2 and 5.3, hence their jagged look. Since we are invoking a weak approximation to probabilities (see Theorem 5.5 as well as Corollaries 5.6 and 5.7) that are not restricted to the interval $[0, 1]$, we can expect there to be probabilities in excess of one, so this limits the usage of the weak approximation especially for index-linked CAT bonds with low threshold levels and/or short terms to maturity. The second drawback concerns the following: the approximation does not lead to real numerical values for small threshold levels, D – in fact, for small threshold levels the approximation is not valid since we cannot find a positive value for M (since in Theorem 5.5, M depends on D, for fixed D). But this is as expected – for both distributional assumptions – since the approximation result in Equation (5.8) is for when

$M \to \infty$ and indeed such small threshold levels are not evident in index-linked CAT bond pricing practice.

5.7 Concluding remarks

In this chapter, we examined index-linked ZC and CP CAT bonds under the pricing framework outlined in Chapter 3. As the pricing formulae obtained under this framework were not in closed-form, we invoked a weak approximation of the ALP – assumed to be a compound renewal process – to α-stable Lévy motion. Thereafter, we presented Theorem 5.5 which allowed one to weakly approximate tail probabilities pertaining to this type of ALP, and then specialised the approximation to the case of a time-inhomogeneous compound Poisson process. The contributing feature of our approximation is that it can be used in the case of heavy-tailed distributions (Theorem 5.5 - in particular for when only the first moment is finite), and as a useful check for existing approximations such as the FSRLP approximation. Since such tail probabilities are essential in computing CAT bond prices, we emphasise the applicability of our weak approximation in obtaining a way to indeed approximate such prices.

Additionally, we highlight that our weak approximation has other beneficial applications in the insurance sector. Our fast, simple and relatively accurate (compared to Monte Carlo simulation and FSRLP approximations) weak approximation may be applied in the context of other CAT bond pricing models such as in computing attachment probabilities for use in the LFC model of Lane (2003), especially when the losses follow heavy-tailed distributional assumptions. Moreover, our approximation can be applied in approximating loss-exceedance (i.e. tail) probabilities for reinsurance portfolios, and also in estimating premiums for catastrophe excess-of-loss reinsurance contracts.

We end this section by briefly re-mentioning some of the limitations and general advantages of our weak approximation to the ALP, noticed in the context of Burr and GP-distributed severity components. In this research, we pointed out three limitations. Firstly, we demonstrated that the application of our approximation needs to be assesses on a case-by-case basis. It remains to be checked (by the user of our approximation) that the fitted intensity (or its integration) is large enough for our approximation to be applied. Secondly, our weak approximation did not perform well relative to Monte Carlo estimation and the FSRLP approximations for shorter time horizons, and also for the lognormally-distributed severity assumption. Thirdly, in certain extreme cases our weak approximation did lead to probabilities in excess of one; therefore this needs to be checked for. However, we posit that our weak approximation does still have merit. Our weak approximation was found to be simple and computationally inexpensive to implement. Moreover, our weak approximation performed in line with the more traditional approaches of Monte Carlo estimation and the FSRLP approximation in most situations, but for longer time horizons our approximation fared better, in the context of all our numerical illustrations. But most importantly, we emphasise our weak approximation's applicability in situations where data follows a heavy-tailed law – this is, ultimately, because our weak approximations only require the finiteness of the first moment of the loss severity distribution.

5.8 Looking ahead

We have now studied, to a detailed extent, the problem of pricing theoretically-structured catastrophe bonds, and also numerically simulating their values. But lingering questions which remain pertain to the potential improvements on the structure of the catastrophe bond. What if the recovery amount was not returned in the form of cash to the investor in the catastrophe bond, but rather in the form of equity of the issuer? Would this adjustment to the structure of catastrophe bonds appeal to a wider pool of investors, and be able to sell off extremely rare tail risk at a suitable price and return? We consider an answer to this question in the next chapter, and in consequence present the first formalisation of and pricing-framework for the so-called contingent convertible catastrophe bond.

Chapter 6

Contingent convertible catastrophe bonds - a case for equity conversion

6.1 Background

Given the pervasiveness of urbanisation in natural catastrophe-prone areas and also the untoward impacts of global warming, insurers, reinsurers and governments have been suffering from substantial natural catastrophe-related losses. Insurers typically deal with this ever-increasing risk by either re-insuring this risk in the reinsurance market or securitising it in the capital markets. Since the capital markets have access to larger, more diversified and more liquid pools of capital as opposed to the equity of reinsurers, such capital markets possess a notable advantage over reinsurance markets when it comes to the financing of catastrophe risk (Durbin, 2001). Therefore, the search for ways of accessing the alternative, rich and robust sources of capital – from the capital markets – for the financing of contagion-risk and catastrophe-risk exposed entities has ignited a wave of innovative financial products. In this context, ILSs have been at the fore, with the most prominent type of such products being the CAT bond which, as discussed fully in Chapter 3, is a fully-collateralised debt security which pays off on the occurrence of a pre-defined catastrophic event (Cummins, 2008a). Other examples have been catastrophe options, catastrophe futures and CAT-E puts. However, the market places for each of these instruments are now extinct, given low trading volumes (Braun, 2011; Wang, 2016).

On the basis of, firstly, the demise of such instruments' marketplaces, secondly, the recent expansion in academic literature on these instruments, thirdly, the increase in globally-occurring natural catastrophe risk and fourthly, the growth of the ILS catastrophe-bond market (SwissRe, 2009), it seems plausible to suggest the following. There may be a need for novel and alternative sources of funding for catastrophe-prone entities, varying away from the more traditional CAT bonds.

Recently in the banking industry, "contingent capital" instruments, such as contingent convertible (Coco) bonds, have gained the support of various academics, practitioners, economists, regulators and banks as a potential avenue to reduce the need for bailouts of institutions that are classified as "too-big-to-fail" (Rüdlinger, 2015; Flannery, 2016). Contingent capital instruments are a type of debt instrument with a loss-absorbing mechanism: that is, they are automatically converted into common equity or written down when a pre-specified trigger event occurs (Flannery, 2016). It is in the very specification of these contingent capital instruments where we see their application to insurance and reinsurance. According to our interviews with industry practitioners and a number of press releases online, many global insurers and also reinsurers such as the Munich Reinsurance Company, SwissReinsurance Company, Hannover Reinsurance Company and the Reinsurance Group of America, are often referred to as being "too-big-to-fail". Given the success of contingent capital instruments, such as Coco bonds, we now specify a special type of Coco bond for insurers and reinsurers (both of which, in this research, will collectively be referred to as "issuers"). We believe that the issuance of such a special Coco bond, which we shall call a CocoCat, will help stabilise their issuers' balance sheets in times of distress; particularly in times of extreme natural catastrophes potentially spurring on large non-independent insurance-related losses.

In view of the above, this Chapter is organised as follows. Section 6.2 provides a brief discussion of what we believe is the first CocoCat to be issued, and thereafter goes on to formalise the mechanics, structure and features of such a CocoCat, being cognisant of the literature on traditional Coco bonds. Thereafter, we attempt to give a valuation framework for a specific type of CocoCat - namely one linked to an insurance loss index such as the Property Claims Services (PCS) index. We call this type of CocoCat an Index-linked (IL) CocoCat.

Section 6.3 describes the necessary joint asset, loss and interest-rate processes under the real-world (or physical) probability measure, needed to price the IL CocoCat in the context of our model. Also, an important assumption is introduced: we assume that natural catastrophe risk and financial markets risk are independent - and such an assumption allows for convenient pricing formulae.

Section 6.4 uses the dynamics of the various processes driving the price of the IL CocoCat under the real-world measure to price it under a risk-neutral measure, by employing a specific exponential measure change. Thereafter, analytical IL CocoCat pricing formulae are derived. In order to simplify pricing formulae and to avoid numerical simulation of the financial markets variables (namely interest-rates and share prices), we employ an exponential change of measure for the loss process and also a Girsanov-like measure change to remove the assumed correlation between interest-rates and share prices.

Section 6.5 uses the pricing formulae derived in Section 6.4 in order to empirically study the behaviour of the IL CocoCat prices with changes in the model's various parameters. Such an analysis is useful from a contract design perspective for the issuer, for the inclusion of pitch books of ILS structurers, but is also important to the investor in the IL CocoCat (and other types of CocoCats as well). Finally, in Section 6.6, we state our conclusions and recommendations for further research into this new and interesting topic.

6.2 CocoCats - a case for equity convertibles

6.2.1 Background and instrument design

In October 2013, a new reinsurance-hybrid security was placed in the capital markets. The Swiss Reinsurance Company (SwissRe) pioneered the creation of a CHF 175 million contingent convertible bond, primarily to sell off hurricane tail risk to a wider pool of investors. The hybrid security has a term of 32 years, pays an annual coupon of 7.5% and redeems at par, unless triggered. Such returns are reasonable in the ILS markets, wherein for example EU and USA-based markets investors typically demand returns of 5 to 7%. The trigger event is unusual in the classical context of Coco (and CAT) bonds, in that it is a dual trigger: the bond triggers if either a 1-in-200 year Atlantic hurricane[1] occurs during the term (which is unusual in traditional capital-raising exercises), or SwissRe's solvency ratio (as determined by the Swiss Solvency Test, and reported to the Swiss Financial Market Supervisory Authority at the statutory reporting date) falls below 135%. Should either trigger event occur, investors lose their entire principal (see SwissRe (2013)).

Such novel hybrid securities are, indeed, appealing to both issuers and investors. In a low interest-rate environment and a rising equity market, high-yielding coupon-rates on such novel debt issuances cannot be overlooked by capital markets investors. Moreover, such issuances are attractive to these investors from the perspective of diversification, in that firstly catastrophe risks are remote from financial market risks; and secondly such novel securities differ from the more traditional insurance-linked securities. Nevertheless, it must be borne in mind that the diver-

[1] According to press releases by Reuters, Bloomberg and SwissRe, the probability of the occurrence of such an event is low compared to catastrophic events upon which catastrophe bonds are more commonly based (such as 1-in-30 or 1-in-50 year events).

sification benefits in the case of SwissRe's issue are not as pronounced as in the case of a traditional catastrophe (CAT) bond. This is due to the existence of the dual trigger, part of which is based on financial market events. But ultimately, such novel securities – if structured differently to that of SwissRe's – can potentially offer rare opportunities to recoup the full amount of principal invested over time should equity prices rise, *ex post* the catastrophe. Finally from the issuer's perspective, such securities could help satisfy regulatory solvency requirements and could reduce probabilities of default *ex post* under Solvency II frameworks in the EU and (in certain cases) the capital-requirement frameworks set out by the National Association of Insurance Commissioners in the USA. Moreover, the coupon payments on such bonds can provide a degree of tax relief for banks (Rüdlinger, 2015) and potentially insurers and reinsurers. But most importantly, such instruments can help issuers shift catastrophe-related tail risk off their balance sheets via a novel way and can provide certainty on the capital to be received *ex post* the trigger.

We now nestle the special Coco bond issue by SwissRe into a more formal setting, and attempt to formalise what is meant by a CocoCat. Since the market for such securities is still in its infancy and given the little scholarly attention to date, it appears clear to posit that, to the extent of our knowledge, no formal definition of a contingent convertible catastrophe (CocoCat) bond exists in the academic literature. In the corporate liabilities sphere, a Coco bond is defined to be a debt instrument[2] (for accounting purposes, categorised as an ordinary liability) that, upon the occurrence of a pre-defined trigger event, converts into common equity via some pre-specified conversion mechanism, or suffers a full write down. Spiegeleer and Schoutens (2012) maintain that, in the context of an issuing bank, the trigger event for the Coco is most often a state of possible financial non-viability. Therefore, the purpose of Coco instruments is to stabilise the balance sheet in times of financial distress or contagion effects, and also to allow for a decrease in the systemic risk faced by large financial institutions (Rüdlinger, 2015). However, it must be borne in mind that the conversion to common equity exposes the Coco bond investors to future potential losses and furthermore exposes the existing shareholders to the risk of dilution upon conversion. So, both Coco bond investors and existing shareholders have a greater vested interest in monitoring the risk budget of the financial institution, leading to better risk monitoring by both these parties (De Martino *et al.*, 2010). Furthermore, given the impending risk of a dilution, existing shareholders may demand a higher required return on their equity stakes. From the issuer's perspective this higher required return will, to a large extent, be counterbalanced by

[2] The debt instrument can be zero-coupon, a fixed-coupon or (more commonly) a floating-coupon (with a fixed spread) bond.

the reduction in risk from the conversion feature. Therefore, the specification of the conversion feature in a Coco bond's structure – and also in the context of a CocoCat – will be important from the perspective of counterbalancing the additional return required by existing shareholders.

We view a CocoCat as a special type of Coco bond. Coco bonds are characterised by two important features - the *conversion trigger* and the *conversion mechanism* - and we attempt to apply them within the sphere of CocoCats. As mentioned by Rüdlinger (2015), the conversion trigger sets out one or several events that trigger the conversion mechanism of the Coco bond, while the conversion mechanism explicitly defines, in the bond covenant, what happens to the Coco bond directly after the trigger event. Therefore, we consider a CocoCat to be a Coco bond that has a trigger[3] linked to the occurrence of a single or sequence of predefined natural catastrophes, and a conversion mechanism whereby the bond either (i) converts into common equity of the issuer (therefore increasing the size of common equity in issue), at a predefined conversion rate as specified in the bond covenant, or (ii) is written down (both principal and coupons) by a fixed percentage which is specified in the covenant. The latter conversion mechanism is reminiscent of the typical structures of various CAT bonds in issue today (for example, see Cummins and Weiss (2009)).

In view of our proposed definition for CocoCats, we offer the following remarks in light of its practical applicability in the insurance and reinsurance settings. Firstly, SwissRe's 2013 placement is loosely an example of such a CocoCat, with the trigger being indemnity-based. Secondly, we note that in the context of bank-issued Cocos with bank-related triggers, in order to protect the issuing bank's reputation there is usually a tendency not to write down the principal amount invested or defer coupon payment (Bishop *et al.*, 2009). As mentioned by some industry practitioners we interviewed, this behaviour is also evident in the CAT bond landscape. When triggered, many CAT bonds are not fully written down, but instead begin to pay smaller coupons over a longer period of time (compared to the original term), with the principal repayment potentially being delayed to a time point after maturity. Therefore, by undertaking such actions it is clear that both banks and issuers of CAT bonds do not wish to send negative signals through to current and future investors. This highlights one of the many potential benefits that CocoCats have to offer to issuers, since such actions (as in the case of CAT bonds) are not necessary.

We now present the benefits we foresee CocoCats to offer. Firstly, there can be more certainty in timing (i.e. debt is converted at the time of trigger) of the principal

[3] A CocoCat's trigger is allowed to be of the same form as the traditional CAT bond triggers, namely a parametric, index-linked, modelled or indemnity trigger.

recoveries from the CocoCats since there will be no unexpected delays in principal repayments. Secondly, the structure of the CocoCat can allow for the amount of principal recuperation for the issuer, as well as the total amount to be injected into equity (belonging to the investor) to be fixed in advance. Finally and most importantly, CocoCats can accommodate for a needed reduction in ordinary liabilities and also a provision of immediate *ex ante*[4] liquidity to immediately pay insurance claims. Notice that the equity conversion feature of the CocoCat can allow for a possible boost in the *ex post* solvency margin of the issuer under the Solvency II regime, upon financial stress caused by the impact of natural catastrophe-related and clustered claims on its balance sheet.

We point out, however, that trigger conversion mechanisms such as write downs can be penal from both the CocoCat issuer's and CocoCat investor's perspective. From the investor's perspective, a high coupon rate prior to the equity conversion will be necessary in light of the write-down risk but must, nonetheless, be commensurate with the risk of trigger of the CocoCat. While from the issuer's perspective, the attractiveness of such a CocoCat will limit the size of the capital market that can be tapped into, as certain investors may not be able or willing to tolerate the risk of a full write down of the principal. On consideration of all of the aforementioned, we propose it may be suitable for CocoCats to be issued on the basis of an equity conversion trigger - in that a certain proportion of the bond's principal is recovered in common equity of the issuer, hence potentially increasing the capital reflected in their balance sheet or in their risk assessment exercises. We subsequently continue with this impetus: from now on, we assume that the CocoCat converts into equity upon trigger.

As final evidence in respect of our case for equity conversion-based CocoCats, we motivate the benefits of issuing CocoCats over traditional CAT bonds. Much of this evidence was pointed out in Georgiopoulos (2016), which we use as a basis for the simple design of our proposed CocoCat. Firstly, CocoCats afford issuers the opportunity to transfer insurance risk without the need to deal or trade their investments in offshore jurisdictions, such as Bermuda and the Cayman Islands, where ILSs are mostly traded. The CocoCat can be directly issued by the issuer via an underwriter or with the help of a structuring agent, in its own local (or judiciously selected foreign) debt market, in a similar fashion to the way a Coco bond is issued. Secondly, investors can more easily trade in CocoCats compared to CAT bonds, and may not need a qualified investor clause to trade in them (Georgiopoulos, 2016). Another attractive feature for the issuer of the CocoCat is that the setup of a special

[4] This *ex ante* capital provision is also a pleasing advantage of a CocoCat over a CAT-E put, since the *ex post* capital provision from the latter may potentially not materialise (i.e. credit risk).

purpose vehicle (and also the total return swap required for a CAT bond setup) to issue the debt and act as a type of collateral for the debt is not needed, therefore reducing the instrument's setup expenses. Thirdly because of the non-existence of a special purpose vehicle for their issuance, unlike traditional ILSs, CocoCats do not require a special reinsurance intermediary for the promotion, issue and sale of the debt - an investment bank can underwrite the issue (Georgiopoulos, 2016). Finally, the trigger mechanisms of CocoCats can easily be based on third-party catastrophe risk models, and specialised in-house model development will, therefore, not be required. We point out that this is also an advantage of issuing an index-linked CAT bond over an indemnity or parametric one, but that basis risk[5] can result for both a CocoCat with an index-linked trigger, and an index-linked CAT bond. However, we emphasise that in the case of an insurer or a reinsurer acting as the issuer, the point of a CocoCat is not complete protection against catastrophe-related insured loss payouts, but rather partial financial protection stemming from the capital markets, a markedly larger market compared to the reinsurance market alone.

Therefore, on the basis of the above information, we propose the formal structure for the CocoCat to be as illustrated in Figure 6.1. In reference to Figure 6.1, the operation of a CocoCat is discussed below.

2.1.1 The investors transfer the bond's principal to the issuer, where this is then reflected immediately as an ordinary liability in the balance sheet. The proceeds will then be invested by the issuer in liquid treasuries (possibly at a haircut), with discounted mean terms approximately on par with that of the CocoCat in order to avoid credit risk.

2.1.2 The issuer may organise to swap the fixed return for a floating return, especially if the CocoCat's coupons are floating in nature. A floating return is often used in order to base investors' expected returns on a reference interest-rate (Müller and Grandi, 2000).

2.1.3 Prior to the trigger of the CocoCat, the issuer will use the floating return as well as insurance profits to pay the coupon on the CocoCat, on a pre-specified tenor, to the investors. The coupon will be based on a reference interest-rate (such as LIBOR), and will also include a fixed spread to allow for the catastrophe risk.

2.1.4 If the CocoCat has not been triggered, at maturity the investors will receive their full principal back in cash, together with the final coupon.

[5] Basis risk is the risk of a potential mismatch between the cashflows of the protection instrument and the losses it is supposed to be hedging. We point out that taking on basis risk may be too costly, *ex post*, for the issuer.

2.1.5 If triggered, the CocoCat will terminate and the liability will be written off the issuer's balance sheet. The issuer will redeem the principal from the treasuries and it will use a predefined proportion of this principal to cover earmarked excess (catastrophic) claims. The remaining proportion of the principal from the treasuries will be converted into (new) common equity, and will belong to the investors in the CocoCat, thereby increasing the issuer's total equity in issue. So, that is, the investors recover a proportion of their principal in equity of the issuer. Figure 6.2 illustrates the impact of a CocoCat's trigger on the equity and liabilities of the issuer. Notice that after the CocoCat has been triggered, there is a reduction in liabilities (arising from their repayment as a result of the catastrophe), a wipe-out of a proportion of the CocoCat's debt-based value, and a conversion of the remainder of the CocoCat's value into new common equity.

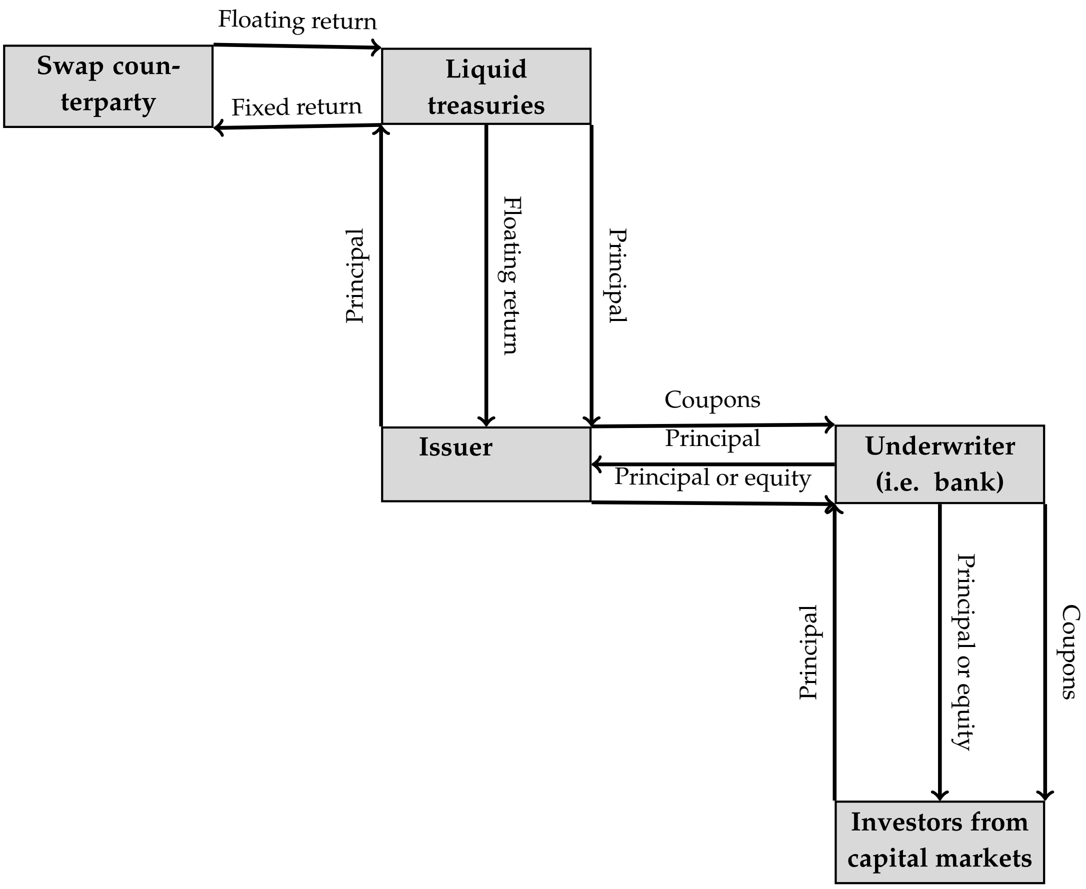

Fig. 6.1: Proposed structure for the CocoCat, in the case of it being underwritten by a bank.

With a simplified design structure for the CocoCat in mind we now endeavour

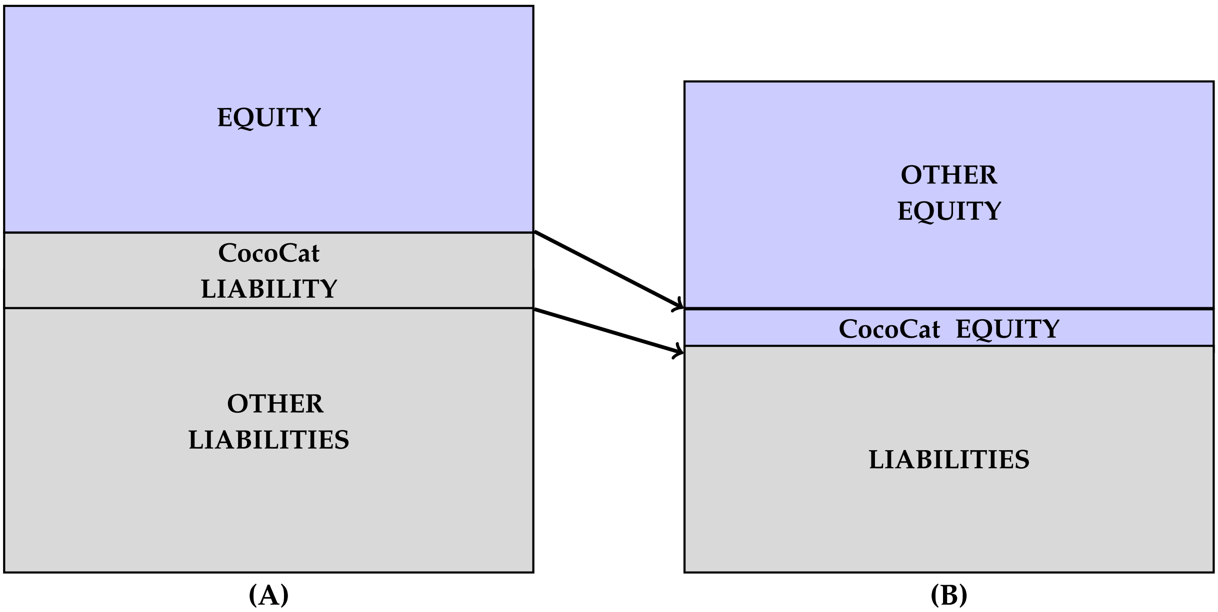

Fig. 6.2: Projected effect of the CocoCat's trigger on the equity and liabilities of the issuer: (A) provides a simplified overview of the balance sheet structure prior to the trigger of the CocoCat, while (B) provides the balance sheet overview after the trigger of the CocoCat. Notice the decrease in liabilities, as a result of the write-down.

to analyse and refine the structure more fully, with reference to the "anatomy" of Coco bonds specified by Spiegeleer and Schoutens (2012). We firstly consider the *conversion trigger*. The natural catastrophe-related trigger event used in the specification of the CocoCat need not be one of the four types of more general Coco bond triggers - accounting, market, regulatory or multi-variate - as put forward by Spiegeleer and Schoutens (2012). This is because of the different purpose and nature of the CocoCat, and also that the CocoCat is a security specific to insurers and reinsurers, and not exclusively banks. The catastrophe-based trigger itself is inherently different from explicit indicators of the financial health of the issuer (such as the solvency margin), but these types of triggers can be indicators of the overall financial health of the insurance and reinsurance industry as a whole. Thus, we posit that the catastrophe-based trigger is a kind of systemic trigger (which is a trigger linked to the overall financial health of the industry within which the instrument operates, as introduced by Pazarbasioglu *et al.* (2011)).

Before proceeding, it is important to note that the trigger should follow four criteria (Rüdlinger, 2015). We now discuss whether, in general, a natural catastrophe-based trigger is compliant with these four criteria.

2.2.1 **Clarity of the trigger event.** The trigger should carry the same message in whatever jurisdiction the issuer operates. It is not possible for all of the universally-accepted CAT bond triggers (parametric, indemnity, index-linked or modelled) to be CocoCat triggers. Indeed, indemnity-based triggers would meet this criterion of clarity with difficulty, given that insurance loss reporting differs from one jurisdiction to the next. Index-linked triggers (if based on a particular insurance loss index) and parametric-based triggers would meet this criterion. The trigger could also encompass an accounting-related trigger, such as in the case of SwissRe's CocoCat, but due care and attention would need to be taken since different accounting regulations apply in different jurisdictions.

2.2.2 **Objectivity of the trigger event.** The trigger should be well-documented in the CocoCat's prospectus, and known at the date of issue. Categorically, there should be no scope to alter the definition of the trigger or change the way that the bond is converted into equity during the term of the CocoCat.

2.2.3 **Transparency of the trigger event.** The catastrophe-related trigger should be simple to understand, and observable for both the investor and the issuer at regular intervals of time. This may be a problem for all four of the CAT bond triggers forementioned, since much of the information is proprietary to the company (especially in the case of an indemnity-based trigger) or proprietary to another catastrophe-modelling company (especially in the case of an [industry] index-linked trigger). This criterion is satisfied in the case of an industry-index trigger. However, we argue that the issuer should endeavour, under strict confidentiality clauses, to provide the information on the evolution of the trigger process to the investors.

2.2.4 **Functionality of the trigger event.** Trigger categories for Coco bonds are defined by the functionality condition: the trigger should be an appropriate measure for the state of financial distress of the issuer, or the financial market within which the issuer operates. We argue that the functionality of natural catastrophe-based triggers does give an indication of the financial distress of a particular issuer, since the expected future loss claims and potential claim contagion specific to the issuer, in respect of these catastrophes, will be linked to the occurrence of these catastrophes. Furthermore, the occurrence of catastrophes is not under the control of the issuers, providing further justification for the choice of such objective trigger mechanisms.

We now turn to considering the *conversion mechanism* which, as mentioned before, specifies the procedure to follow at conversion, and the potential loss (of principal) to the investor in the CocoCat at conversion. We follow in the spirit of Rüdlinger (2015), and for our proposed CocoCat structure consider the conversion fraction, price and rate.

The conversion fraction, ζ, sets out the proportion of the contingent bond's face value which is converted into common equity, at a contractually-specified conversion price. Thus, if $\zeta = 1$, the full face value of the bond is converted into equity. However, for the purpose of our proposed CocoCat, we suppose that $0 < \zeta < 1$: thus, the CocoCat investor loses[6] a proportion of $1 - \zeta$ of his or her invested principal, and has the remaining proportion of ζ converted into new common equity. Since the main purpose of a CocoCat is to provide immediate, liquid funding in the event of a trigger (which is directly or indirectly linked to the claims experience of the issuer), we propose that the CocoCat debt is written down in the balance sheet, and this written down amount is earmarked to immediately cover any worse-than-expected catastrophe-related claims (directly or indirectly linked to the trigger). This has the effect of de-levering the issuer's balance sheet. We do caution, however, that the choice of the conversion fraction is a subjective but critical one. Recall that the purpose of ILSs is not to provide full protection in the case of adverse experience, but rather be complementary to a comprehensive catastrophic risk-management framework. So, the issuer needs to carefully decide on ζ, which ultimately defines how much capital the issuer receives, per unit nominal, *ex post* the catastrophe. Factors which will impact on the magnitude of ζ are firstly the projected future catastrophe-related claims experience (which is difficult to assess with accuracy) as well as risk budget, and secondly the investor base to which the security will be marketed and issued. Finally, the impact of the consequent equity dilution needs to be accounted for.

The conversion price, K_P, is the key indicator for the potential loss the CocoCat investor can incur on conversion. In fact, K_P can be interpreted as the share price of the CocoCat's issuer at which the fraction, ζ, of the bond's face value is converted into common equity upon trigger. Rüdlinger (2015) presents three possibilities on how the conversion price can be set, and we apply this to the context of CocoCats: (a) fixed in the CocoCat prospectus, (b) the market share price upon conversion or (c) a function of, *inter alia*, the known share price at time of trigger. We argue that from one perspective, option (a) may be unsuitable for investors. Evidence for moderate decreases in insurance firms' share prices after the occurrence of mega natural catastrophes has been found recently by Hagendorff *et al.* (2015), which is intuitively expected. Therefore, there exists the risk of setting the conversion price

[6] In fact, the investor may not lose the entire amount of $1 - \zeta$: depending on the conversion price, he or she may even gain upon trigger.

higher than the market price, creating an adverse effect for the investors of the CocoCat, and consequently investors will receive less share. This could reduce the marketability and relative attractiveness of the CocoCat. Moreover, this conversion price could allow for moral hazard from the side of the issuer and also from current shareholders with substantial stakes in the company. That is, setting the conversion price at an unacceptably high level will materially affect the equity stake CocoCat investors will recover upon trigger. However, from another perspective there exists the consequential risk if the conversion price is set too low compared to the share price at trigger. Suppose that in the case of $1 nominal, $\zeta = 0.4$, $K_P = 5$ and the share price at conversion is 25: in this case, the investor will receive an equity stake of $3, which is three-times higher than the nominal amount he or she invested. This is very attractive to the investor, however it is also penal to the issuer since they would need to find excess funds to back this equity in times of crisis. This penal effect can, to a large extent, be controlled through a judicious selection of K_P and ζ.

On balance, a key benefit of option (a) is as follows: depending on the size of the CocoCat issue this option may be preferred by current shareholders in the issuing firm because it could potentially restrain dilution of their holdings if the current share price is severely depressed as a result of the impact of natural catastrophes but the strike price is comparatively high (compare to Spiegeleer and Schoutens (2012) in the case of conventional Coco bonds). Upon consideration of option (b) for the conversion price, it would appear that its usage would be optimal for CocoCat investors since few (if any) losses on share price differences will result for the investors. Moreover, (b) is also beneficial because CocoCat investors have the potential to gain material stakes in the issuing firm if the size of the CocoCat issue is large relative to the total equity in issue. Also, the scope for the moral hazard identified in option (a) is not explicitly present. Current shareholders, however, will simply have to accept a dilution of their shareholdings (Spiegeleer and Schoutens, 2012), which may be undesirable from their perspective. Option (c) is also a possibility in CocoCat design and can allow for more flexiblity, which could improve the attractiveness of such an issue to both investor and issuer. However, it can still suffer from the drawbacks mentioned for option (a). There are some further arguments in favour of such function-based conversion prices concerning the reduction in manipulation of share prices by the CocoCat investors - see Section 3.3 of Rüdlinger (2015) for similar arguments based on conventional Coco bonds. However, we caution that this avenue may complicate pricing considerations and frameworks. In this research, we accommodate for options (a), (b) and (c), in light of the choice for K_P.

We end this section with some comments on the practical use of CocoCats within the context of an insurance or reinsurance company. CocoCats are primarily intended to be of assistance in the management of economic and also solvency capital for an issuer: as postulated by Besson *et al.* (2009), capital's critical function is to absorb risks undertaken by the company, be they worse or more contagious than expected. We also believe that our proposed CocoCat lends itself to a situation where existing shareholders may not be called upon frequently to provide additional capital in situations of worse-than-expected risk. Requiring additional funds from existing shareholders is unfavourable (Besson *et al.*, 2009).

We also reiterate that the proposed CocoCat instrument is not intended as an ingenious financial instrument to achieve full indemnity against catastrophic losses. Rather, it is to be a complement to and also an integral part of a comprehensive and consistent catastrophe risk-management framework. In consequence, it should adhere to the framework elements put forward by Pazarbasioglu *et al.* (2011), those being enhanced supervision, a robust economic capital base, transparent disclosure which better informs markets, and a clear resolution regime. Most CocoCats will, for the time-being, be unstandardised, over-the-counter traded and tailor-made (to the issuer) instruments, since their market is new. So careful scrutiny is necessary in developing and managing such issues on behalf of the issuer, the investors and the insurance market regulators. Although it is specific to each issue and is a difficult task, a careful balance between the potential benefits of CocoCats to the issuer, and the rewards reaped by investors, needs to be achieved without the introduction of additional moral hazard and information asymmetries.

6.2.2 Possible accounting and regulation

By design, CocoCats are financial instruments and not insurance or reinsurance contracts. As a result, their accounting and regulatory treatment will fall under the auspices of the International Financial Reporting Standards (IFRS) 7, and also US Generally Accepted Accounting Principles (GAAP). However, the exact accounting methodology affecting their recognition in the financial statements is difficult for us to pinpoint.

CocoCats have two components: a contractual obligation to deliver cash, as well as an equity conversion feature. Firstly, the cash obligation comprises the payment of the coupons and the principal amount, which are all mandatory payments up until the time of trigger. International Accounting Standard (IAS) 39 applies to financial contracts that take the form of insurance, which principally involve the transfer of financial risks and derivatives embedded in insurance contracts. Therefore, the expected coupons and principal repayments will be a financial liability (as defined by IAS 32), recognised at cost. However, the equity conversion feature will be accounted for differently. After the trigger, investors recover a portion in equity, which must still be recognised in the financial statements. This recovery can be viewed as an embedded derivative, and IAS 39 requires that an embedded derivative is separated from its host contract and accounted for as a derivative instead. Therefore, the equity conversion feature will be recognised as a separate financial liability at fair value through profit-and-loss (FVTPL). This separated recognition could lead to increased volatility in the income statement since the issuer's technical liabilities are presently not marked-to-market. Also, note that prior to conversion no part of the CocoCat will be recognised as equity. Only after conversion will the CocoCat lead to an equity component in the financial statements.

When IFRS 9 comes into effect in 2018, any financial liability can be measured at FVTPL provided that it contains at least one embedded (but not closely related) derivative which sufficiently modifies its cashflows. Therefore, under the new regulatory regime the entire CocoCat will be recognised at FVTPL.

In the US, the treatment will differ slightly. The National Association of Insurance Comissioners recognises any properly-structured ILS as a reinsurance contract (Braun, 2017), so CocoCats will also fit under this banner for now.

Under Solvency I in the EU and the National Association of Insurance Commissioners regulation in the USA, only indemnity-linked (and not index-linked) financial contracts could be recognised as regulatory capital (Braun, 2011). However, under Solvency II, properly structured ILSs (such as CocoCats) can be incorporated into the issuer's Solvency Ratio calculation, consequently decreasing the Solvency Capital Requirement (see SwissRe (2009) and Braun (2017)).

6.2.3 Comparison to other catastrophe-linked ILSs

CocoCats are an ILS that form a unique class of their own. They are similar to CAT-E puts, in that the issuer will sell some of their share to the investor should the trigger occur. However, CAT-E puts suffered from the drawback of credit risk. CocoCats do not, since they provide capital (*ex ante* the trigger) at the outset of the contract. Moreover, CocoCats can be much longer in term than CAT-E puts - the term will depend on the trigger type.

Very much like index-linked catastrophe bonds, catastrophe swaps and CAT-E puts, index-linked CocoCats can expose the issuer to basis risk (which is not the case with industry-loss warranties and reinsurance), especially if the instrument is targeted at hedging a particular portfolio of the issuer's liabilities. But, a pleasing advantage of index-linked ILSs is that they may remove the pervasiveness of moral hazard, which is, unfortunately, an issue when it comes to reinsurance. Finally, it is possible to recover the full principal invested in a CocoCat, even if it is converted (should the equity perform well in the future), but such full recovery is not always possible for CAT bonds. Table 6.1 summarises the key points of this comparison.

Tab. 6.1: Brief comparison of CAT bonds, CAT-E puts and CocoCats (the superscript, *, indicates that the specification can vary beyond what is mentioned in the table).

	CAT bond	**CAT-E put**	**CocoCat**
Term	3-5 years	1-5 years*	Depends on trigger
Capital provision	*Ex-post*	*Ex-post*	*Ex-ante*
Possibility of full principal recovery *ex-post* catastrophe	No	N/A	Yes
Trigger of payment	Index-linked Indemnity Pure parametric Parametric index Modelled loss Multiple triggers	Strike vs. share	Index-linked Indemnity Pure parametric Parametric index Modelled loss Multiple triggers
Moral hazard	Little if index-linked	None	None if index-linked
Basis risk	Little if pure parametric	Large - smaller if variance-linked	Little if pure parametric
Existence of market	OTC and exchange	Extinct	Very small
Counterparty default risk	Low (collateralised)	High	Low
Accounting treatment	Depends on trigger	Financial instrument	Financial instrument

6.3 Index-linked CocoCat: model

6.3.1 Model setup, assumptions and properties

We now turn to focusing on a particular type of CocoCat, and introduce the workings, notations and basic definitions necessary for its analytical pricing. Before reading on, note that we do refer to Chapter 3 extensively. We suppose that the trigger is index-linked, and is furthermore in line with many of the index-linked triggers that CAT bonds are based upon (see, for example, the review by Cummins (2008a), as well as Haslip and Kaishev (2010); Ma and Ma (2013); Nowak and Romaniuk (2013) and Gatzert *et al.* (2014)). As an example, consider one of the most commonly-issued index-linked CAT bonds: the type that is dependent upon the Property Claims Services (PCS) industry index. The trigger, for most of these bonds, is defined to be the point in time when the accumulated losses from the PCS index exceeed some contractually-specifed threshold level. From here on, we suppose that the CocoCat is based on the PCS loss index, however, note that there is no loss of generality in terms of the type of index which can be used in the model. Other indices (such as that of PERILS in the EU) may also be used in our framework.

Before presenting the processes capturing the behaviour of the financial markets and catastrophe-risk variables, we introduce some notation specifically for the CocoCat. Let:

- $T > 0$ denote the term of the IL CocoCat. For IL CocoCats, we suppose that the term will be in line with that of commonly-issued index-linked CAT bonds. However, for parametric CocoCats based on very rare tail risks (such as the SwissRe CocoCat mentioned in Section 6.2.1) it makes sense for the term to be much longer.

- Z be the principal amount invested in the CocoCat, which the investor will receive back should the CocoCat not trigger during its term.

- V_0^{COCO} denote the price of the IL CocoCat at issue date, $t_0 := 0$. We are only interested in the issue-date price.

- $\{S_t, t \geq 0\}$ be the share price process of the issuing firm.

- $\{L_t, t \geq 0\}$ be an aggregate loss process capturing the behaviour of the index upon which the IL CocoCat is based. The aggregate loss process is assumed to have a frequency component specified by a (possibly non-homogenous) Poisson process $N = \{N_t, t \geq 0\}$ with deterministic intensity specified by the real-valued function λ_t, and a sequence of i.i.d. severity-component continuous random variables $\{X_k, k \in \mathbb{N}\}$ (independent of the frequency component), each with distribution function F_X and density f_X. We assume that $\int_0^T \lambda_s \, \mathrm{d}s < +\infty$.

- $\mathbb{T} := \{t_1, t_2, ..., t_{N-1}, t_N = T\}$ denote the set of N coupon-paying dates.
- Δ denote the constant yearly time period between coupon payment dates t_{i-1} and t_i for $i \in \{1, 2, ..., n\}$. According to Jarrow (2010), in the context of CAT bonds this should be either one (1/12 year), three (3/12 year) or six months (6/12 year).
- $R(t, t_{i-1}, t_{i-1} + \Delta)$ be the Δ-year simple forward LIBOR rate per annum, at time $t \geq 0$.
- $\{r_t, t \geq 0\}$ be the riskless spot rate process per annum, continuously compounded.
- $c \geq 0$ be the constant spread for the IL CocoCat (i.e. the catastrophe risk premium).
- ζ be the contractually-specified conversion fraction for the IL CocoCat, as introduced before.
- $D > 0$ be the threshold level for the trigger, specified in the IL CocoCat's prospectus.
- $\tau = \inf\{0 \leq t \leq T : L_t \geq D\}$, the first time the trigger level is met or exceeded. D is called the contractually-specified threshold level of the IL CocoCat.
- K_P be the pre-specified conversion price as introduced in Section 6.2.1. As mentioned before, K_P can be a pre-specified constant K (we consider this case later on), but it can also be equal to S_τ or $f(S_\tau)$ for some real-valued function f.

Based on all of the above notations, our modelling assumptions are encompassed by the following system of stochastic differential equations (SDEs) and identities under the real-world measure $\mathbb{P}$:

$$
\begin{aligned}
\mathrm{d}r_t &= (\bar{\vartheta}_t - a_r r_t)\mathrm{d}t + \sigma_r \mathrm{d}W_t^1, && (6.1)\\
S_t &= S_t^{\mathscr{C}} S_t^{\mathscr{F}}, && (6.2)\\
S_t^{\mathscr{C}} &= \exp\left(-\alpha L_t + \alpha\kappa \int_0^t \lambda_u \mathrm{d}u\right), && (6.3)\\
S_t^{\mathscr{F}} &= S_0 Y_t, && (6.4)\\
dY_t &= \mu_S Y_t \mathrm{d}t + \sigma_S Y_t \mathrm{d}W_t^2, Y_0 = 1 && (6.5)\\
\mathrm{d}\langle W_t^1, W_t^2\rangle &= \rho \mathrm{d}t, && (6.6)\\
L_t &= \sum_{k=1}^{N_t} X_k. && (6.7)
\end{aligned}
$$

6.3.2 Remarks on the model

Interest-rate

We select a special case of the Extended Vasicek Model of Hull and White (1990). This model is one of the most historically important interest rate models, and is still frequently used, especially by actuaries, in risk-management practice (Brigo and Mercurio, 2007). In the case of a single parameter depending upon time, such a model is furthermore fairly straightforward to calibrate to the initial yield-curve (and will be consistent with it), and has analytically tractable zero-coupon bond pricing formulae. However, what was also an important deciding factor for us was that the Extended Vasicek interest rate process remains, in form, as an Extended Vasicek process under any Girsanov transformation with a constant kernel (see Theorem 6.1). In fact, the entire analysis which follows could also be accomplished for any other interest-rate model as long as the latter property is satisfied. For example, the Vasicek single-factor model of Vasicek (1977) and the Longstaff single-factor model of Longstaff (1989) could be considered. The Cox-Ingersoll Ross model can only be considered in the case when the conversion price is set at a constant, since the effect of interest rates disappears in the valuation procedure.

The special case of the Extended Vasicek model, which we consider, takes on the form as shown in Equation (6.1). It is a three-parameter model, which does admit negative interest rates (which have been observed in many economies, e.g. Switzerland), wherein the yield is linear in r_t. Consider Equation (6.1): under $\mathbb{P}$, $\bar{\vartheta}_t := \vartheta_t - \rho\sigma_r\sigma_S$, and a_r are model parameters, σ_r is the instantaneous volatility and W_t^1 is a standard Brownian motion.

We select $\bar{\vartheta}_t$ in such a way that allows us to fit the the model to the current yield-curve. This gives our analysis a more practical angle. If we denote by $f^M(0,T)$ the market instantaneous forward rate at time 0, that is

$$f^M(0,T) = -\frac{\partial \log P^M(0,T)}{\partial T},$$

where $P^M(0,T)$ is the market discount factor for maturity time T, then as highlighted in Brigo and Mercurio (2007, Chapter 3),

$$\vartheta_t = \frac{\partial \log f^M(0,t)}{\partial t} + a_r f^M(0,t) + \frac{\sigma_r^2}{2a_r}\left(1 - e^{-2a_r t}\right), \tag{6.8}$$

with its integration

$$\int_t^T \vartheta_u \mathrm{d}u = f^M(0,T) - f^M(0,t) - a_r \log\left(\frac{P^M(0,T)}{P^M(0,t)}\right) + \frac{\sigma_r^2}{2a_r}\left(T - t - \frac{e^{-2a_r t}}{2a_r}\left(1 - e^{-2a_r(T-t)}\right)\right),$$

both under the risk-neutral probability measure $\mathbb{Q}_F$. Also, the reason behind the specification of $\bar{\vartheta}_t$ as $\vartheta_t - \rho\sigma_r\sigma_S$ will be made clear in Theorems 6.1 and 6.2.

Theorem 6.1. *Consider the (general) dynamics of the Extended Vasicek model under any probability measure,* $\bar{\mathbb{P}}$*:*

$$\mathrm{d}r_t = (\hat{\vartheta}_t - \hat{a}_r r_t)\mathrm{d}t + \hat{\sigma}_r \mathrm{d}\bar{W}_t \tag{6.9}$$

where $\bar{W}_t$ *is a standard Brownian motion under* $\bar{\mathbb{P}}$*,* $\hat{\vartheta}_t > 0$ *is a deterministic function of time,* $\hat{\sigma}_r$ *and* $\hat{a}_r$ *are positive constants. Then if* $\bar{\bar{\mathbb{P}}}$ *is defined by a Girsanov transformation with constant kernel,* γ*, i.e. for all* $t > 0$*,*

$$\left.\frac{\mathrm{d}\bar{\bar{\mathbb{P}}}}{\mathrm{d}\bar{\mathbb{P}}}\right|_{\hat{\mathscr{F}}_t} = \hat{\eta}(t),$$

for

$$\hat{\eta}(t) := e^{\gamma \bar{W}_t - \frac{1}{2}\gamma^2 t},$$

then the dynamics of the interest-rate process under $\bar{\bar{\mathbb{P}}}$ *still follow the Extended Vasicek model, that is*

$$\mathrm{d}r_t = (\tilde{\vartheta}_t - \hat{a}_r r_t)\mathrm{d}t + \hat{\sigma}_r \mathrm{d}\bar{\bar{W}}_t$$

where $\bar{\bar{W}}_t := \bar{W}_t - \gamma t$ *is a standard Brownian motion under* $\bar{\bar{\mathbb{P}}}$ *and*

$$\tilde{\vartheta}_t = \hat{\vartheta}_t + \gamma\hat{\sigma}_r,$$

Proof. Under $\bar{\bar{\mathbb{P}}}$, the dynamics of the Extended Vasicek model can be expressed as

$$\mathrm{d}r_t = (\hat{\vartheta}_t - \hat{a}_r r_t)\mathrm{d}t + \hat{\sigma}_r(\mathrm{d}\bar{\bar{W}}_t + \gamma \mathrm{d}t). \tag{6.10}$$

After a little algebra, Equation (6.10) can again be expressed in the form of the Extended Vasicek interest-rate model:

$$\mathrm{d}r_t = (\hat{\vartheta}_t + \gamma\hat{\theta}_r - \hat{a}_r r_t)\mathrm{d}t + \hat{\sigma}_r \mathrm{d}\bar{\bar{W}}_t.$$

□

It must also be noted that the Extended Vasicek model as shown in Equation (6.9) admits an analytical simplification to the expression,

$$\mathbb{E}^{\bar{\mathbb{P}}}\left[\exp\left(-\int_t^T r_u \mathrm{d}u\right)|\tilde{\mathscr{F}}_t\right],$$

where $\tilde{\mathscr{F}}_t$ is the filtration generated by the process shown in Equation (6.9) up until time t. Let us define

$$P(r,t,T,\hat{\vartheta},\hat{a}_r,\hat{\sigma}_r) := \mathbb{E}^{\bar{\mathbb{P}}}\left[\exp\left(-\int_t^T r_m \mathrm{d}m\right)|\tilde{\mathscr{F}}_t\right], \tag{6.11}$$

where $\hat{\vartheta}$, a deterministic function of time, has had its dependence on time suppressed for notational and readability purposes only. In order to find Equation (6.11), we need to find the distribution of $\int_t^T r_m \mathrm{d}m$ conditional on $\tilde{\mathscr{F}}_t$. Since, for a fixed $s < t$,

$$r_t = r_s e^{-\hat{a}_r(t-s)} + \int_s^t e^{-\hat{a}_r(t-u)}\hat{\vartheta}_u \mathrm{d}u + \hat{\sigma}_r \int_s^t e^{-\hat{a}_r(t-u)}\mathrm{d}\bar{W}_u,$$

is conditionally normal with mean and variance specified by

$$\mathbb{E}^{\bar{\mathbb{P}}}\left[r_t|\tilde{\mathscr{F}}_s\right] = r_s e^{-\hat{a}_r(t-s)} + \int_s^t e^{-\hat{a}_r(t-u)}\hat{\vartheta}_u \mathrm{d}u, \quad \text{and} \tag{6.12}$$

$$\mathrm{Var}^{\bar{\mathbb{P}}}\left[r_t|\tilde{\mathscr{F}}_s\right] = \frac{\hat{\sigma}_r^2}{2\hat{a}_r}\left[1 - e^{-\hat{a}_r(t-s)}\right] \tag{6.13}$$

respectively, we have that

$$\begin{aligned}\int_t^T r_m \mathrm{d}m &= r_t \int_t^T e^{-\hat{a}_r(m-t)}\mathrm{d}m + \int_t^T\int_t^m e^{-\hat{a}_r(m-u)}\hat{\vartheta}_u \mathrm{d}u\,\mathrm{d}m + \hat{\sigma}_r \int_t^T\int_t^m e^{-\hat{a}_r(m-u)}\mathrm{d}\bar{W}_u\,\mathrm{d}m \\ &= \frac{r_t}{\hat{a}_r}\left[1 - e^{-\hat{a}_r(T-t)}\right] + \frac{1}{\hat{a}_r}\int_t^T \hat{\vartheta}_u \mathrm{d}u - \int_t^T e^{-\hat{a}_r(T-u)}\hat{\vartheta}_u \mathrm{d}u + \frac{\hat{\sigma}_r}{\hat{a}_r}\int_t^T \left(1 - e^{-\hat{a}_r(T-u)}\right)\mathrm{d}\bar{W}_u\end{aligned}$$

is conditionally normally distributed with mean and variance specified by

$$\mathbb{E}^{\bar{\mathbb{P}}}\left[\int_t^T r_m \mathrm{d}m|\tilde{\mathscr{F}}_t\right] = \frac{r_t}{\hat{a}_r}\left[1-e^{-\hat{a}_r(T-t)}\right] + \frac{1}{\hat{a}_r}\int_t^T \hat{\vartheta}_u \mathrm{d}u - \int_t^T e^{-\hat{a}_r(T-u)}\hat{\vartheta}_u \mathrm{d}u \quad \text{and}$$

$$\mathrm{Var}^{\bar{\mathbb{P}}}\left[\int_t^T r_m \mathrm{d}m|\tilde{\mathscr{F}}_t\right] = \frac{\hat{\sigma}_r^2}{\hat{a}_r^2}\left[T-t+\frac{2}{\hat{a}_r}e^{-\hat{a}_r(T-t)} - \frac{1}{2\hat{a}_r}e^{-2\hat{a}_r(T-t)} - \frac{3}{2\hat{a}_r}\right]$$

respectively. Using standard arguments based on the moment-generating function of the normal distribution, we obtain that

$$P(r,t,T,\hat{\vartheta},\hat{a}_r,\hat{\sigma}_r) = \exp\left(A(t,T) - B(t,T)r_t\right) \tag{6.14}$$

where

$$B(t,T) = \frac{1}{\hat{a}_r}\left[1-e^{-\hat{a}_r(T-t)}\right], \quad \text{and} \tag{6.15}$$

$$A(t,T) = -\frac{1}{\hat{a}_r}\int_t^T \hat{\vartheta}_u \mathrm{d}u + \int_t^T e^{-\hat{a}_r(T-u)}\hat{\vartheta}_u \mathrm{d}u + \frac{1}{2}\mathrm{Var}^{\bar{\mathbb{P}}}\left[\int_t^T r_m \mathrm{d}m|\tilde{\mathscr{F}}_t\right]. \tag{6.16}$$

Now, if $\hat{\vartheta}$ is a strictly positive constant, we are within the sphere of the classical Vasicek model and in consequence the analytical solution to Equation (6.11) retains the same form as in (6.14) but with $B(t,T)$[7] and $A(t,T)$ specified below (adapted from Brigo and Mercurio (2007, Section 3.3)):

$$B(t,T) = \frac{1}{\hat{a}_r}\left[1-e^{-\hat{a}_r(T-t)}\right], \quad \text{and}$$

$$\begin{aligned} A(t,T) &= -\frac{\hat{\vartheta}}{\hat{a}_r}\left(T-t-\frac{1-e^{-\hat{a}_r(T-t)}}{\hat{a}_r}\right) + \frac{\hat{\sigma}_r^2}{2\hat{a}_r^2}\int_t^T\left(1-e^{\hat{a}_r(u-(T-t))}\right)^2 \mathrm{d}u \\ &= (B(t,T)-T-t)\left(\frac{\hat{\vartheta}}{\hat{a}_r} - \frac{\hat{\sigma}_r^2}{2\hat{a}_r^2}\right) - \frac{\hat{\sigma}_r^2}{4\hat{a}_r}B(t,T)^2. \end{aligned}$$

It must be noted that if the measure $\bar{\mathbb{P}}$ is the risk-neutral probability measure for the particular model being considered, then Equation (6.14) denotes the zero-coupon bond price under the assumption of interest rates following a (Extended) Vasicek model. In the case of the Extended Vasicek model (and again assuming that $\bar{\mathbb{P}}$ is the risk-neutral probability measure), as a natural consequence of Equations (6.14), (6.15), (6.16) and (6.8) the zero-coupon bond pricing formula can be shown to be given by

$$P(r,t,T,\hat{\vartheta},\hat{a}_r,\hat{\sigma}_r) = A(t,T)\exp\left(-B(t,T)r_t\right),$$

[7] This term remains the same as in the case of the Extended Vasicek model.

where

$$B(t,T) = \frac{1}{\hat{a}_r}\left[1 - e^{-\hat{a}_r(T-t)}\right], \quad \text{and}$$
$$A(t,T) = \frac{P^M(0,T)}{P^M(0,t)} \exp\left(B(t,T)f^M(0,t) - \frac{\hat{\sigma}_r^2}{4\hat{a}_r^2}\left(1 - e^{-2\hat{a}_r t}\right)B(t,T)^2\right).$$

Share price
We choose the share price process in a similar fashion to, amongst others, Cox *et al.* (2004a), Jaimungal and Wang (2006a), Lin and Wang (2009b) and Wang (2016). Notice that our share price process comprises two components: $S_t^{\mathscr{C}}$ and $S_t^{\mathscr{F}}$, the former being the component driven by catastrophe-risk variables and the latter the component driven by financial markets risk variables. More specifically, when catastrophic events affecting an issuer occur, share prices can be expected to decrease since these large claims must be paid. This is accounted for in the share price process, hence the negative dependency of the process on L_t. Under $\mathbb{P}$, S_0 is the initial share price, μ_S the long-run mean of Y_t and σ_S the instantaneous volatility of Y_t. The constant $\alpha > 0$ represents the effect of the catastrophic losses on the logarithm of the share price. The greater the value of α, the more serious the effect of the catastrophe losses through the term αL_t. Moreover, as in Wang (2016), the mean number of claims $\int_0^t \lambda_u \mathrm{d}u$ is included to compensate positively (to some extent) for the presence of downward jumps in the share price. The constant $\kappa > 0$ governs the manifestation of this effect, and is selected on the basis of Lemma 1 below. Note that W_t^2 is standard Brownian motion under $\mathbb{P}$.

Aggregate loss
We follow a classical approach to modelling aggregate loss process, in that we employ a time-inhomogeneous compound Poisson process to govern the behaviour of the IL CocoCat's underlying index. Embrechts and Meister (1997) state that such a process is a suitable candidate to model catastrophic losses for catastrophe-related derivatives. We also choose such a process since it can capture over-dispersion in the catastrophe arrivals.

Correlations
We assume that the interest-rate and share price processes are dependent: this is captured by a correlation coefficient of ρ.

6.3.3 Instrument operation

From the time of issue, the IL CocoCat holder will receive (at contractually-specified constant intervals Δ) floating coupon payments based on the Δ-year LIBOR rate, $R(t_{i-1}, t_{i-1}, t_{i-1} + \Delta)$, plus a spread c. Hence, the floating payment received at each coupon-date is $R(t_{i-1}, t_{i-1}, t_{i-1} + \Delta) + c$ per unit nominal. At time T, the principal of the bond, $Z > 0$, is received back by the investor, unless the IL CocoCat is triggered earlier.

During the term of the IL CocoCat, the issuer will monitor the performance of the underlying index, and will record the cumulative losses giving rise to the process L_t. If the trigger event is to occur, that is $L_\tau > D$ for some $\tau \in [0, T]$ (if it exists), then the IL CAT bond-type leg of the IL CocoCat terminates, and the investor receives a share in the common equity of the firm, at a conversion price equal to K_P. That is, the investor will recover $\zeta Z / K_P$ units of shares of the issuer, at a value of $(\zeta Z / K_P) S_\tau$ for Z nominal invested in the IL CocoCat.

6.3.4 Selection of conversion price K_P

As we mentioned before, there are three general cases one can consider for the conversion price, and we consider each of them. For the case when K_P is assumed to be a real-valued function of the share price, we suppose that it takes on the form $K_P := S_\tau^\nu$, for $\nu \in (0, 1]$. Taking such a functional of S_τ allows for flexibility in the design of the IL CocoCat, and also analytical expressions for the price of the IL CocoCat. Firstly, it allows both the investor and the issuer to take into account their views on the impact of catastrophe-related losses on the issuer's share price. For example, if the investor believes that the market does not satisfactorily capture the impact of large catastrophic losses on its assessment of the share's value, then ν should be set equal to a value less than one so that the investor purchases the share at a value cheaper than market value. Secondly, it can allow both the investor and the issuer to account for their views on the future market performance of the issuer's share. For example, if the issuer believes the share price will rise (independent of the catastrophic losses), then ν could be set equal to a value less than 1, so that the investor does not gain a large portion of share ownership. Finally, note that if $\nu = 1$, then the conversion price is set at the share price at time-of-conversion.

6.4 Analytical risk-neutral pricing of index-linked CocoCat

We now price the index-linked CocoCat within the context of our model, commencing with a generic conversion price K_P, at the issue date $t_0 := 0$ under our risk-neutral probability measure $\mathbb{Q}$. All notation is assumed to follow that in Section 6.3.1, and the assumptions as well as results therein are also assumed to hold. As specified in Section 6.3.1,

$$R(t, t_{i-1}, t_{i-1} + \Delta) \tag{6.17}$$

is the forward LIBOR at time t for the interval $[t_{i-1}, t_{i-1} + \Delta]$. In particular, $R(t, t, t + \Delta)$ is the LIBOR at time t: since Δ is assumed to be constant, we denote the LIBOR process by $\{R_t, t \geq 0\}$.

$$B(0,t) := \exp\left(-\int_0^t r_u \mathrm{d}u\right) \tag{6.18}$$

is the discounted riskless bank account associated with the progressively measurable process $\{r_t, t \geq 0\}$, and

$$P(0,t) := \mathbb{E}^{\mathbb{Q}}\left[B(0,T)|\mathscr{F}_t\right], \qquad \forall t \in [0,T]. \tag{6.19}$$

Now, under $\mathbb{Q}$, it is possible to find an expression for the price of the CocoCat at issue under expectation. This is of crucial importance and we formalise it as a main Fact.

Fact. The issue-date, (or time-zero) risk-neutral price, V_0^{COCO}, of an IL CocoCat is

$$V_0^{\mathrm{COCO}} = \mathbb{E}^{\mathbb{Q}}[I_1 + I_2 + I_3], \tag{6.20}$$

where

$$I_1 := \sum_{i=1}^{N} (R_{t_{i-1}} + c)\Delta Z \mathbb{I}_{\{\tau > t_i\}} B(0, t_i);$$
$$I_2 := \frac{\zeta Z}{K_P} S_\tau \mathbb{I}_{\{\tau \leq T\}} B(0, \tau);$$
$$I_3 := Z \mathbb{I}_{\{\tau > T\}} B(0, T).$$

Notice that the expectation in the Fact comprises three terms. I_1 represents the coupon payments (linked to LIBOR) inclusive of the spread, while I_3 represents the capital repaid at maturity should no default occur prior. I_2 represents the recovery upon conversion-to-equity. Notice that if K_P is set to be equal to S_τ, we consequently obtain the pricing formula for a specific type of CAT bond which pays out ζZ immediately upon the time of trigger (i.e. time τ). We emphasise this feature of our model, which lends it to a broader suite of applications. Some CAT bond issues, in practice, are of this nature so in consequence our valuation framework may find potential applicability in this instance. However, in our framework below we shall consider a specific function of S_τ, which leads to three cases for the conversion price.

Now we show how our model changes under the martingale measure $\mathbb{Q}$. Under $\mathbb{Q}$, the discounted share price process $\left\{\exp\left(-\int_0^t r_u \mathrm{d}u\right) S_t, t \geq 0\right\}$ must, by definition of a risk-neutral measure, be a martingale with respect to the filtration $\mathscr{G}_t$. This is gives rise to the following theorem.

Theorem 6.2. *Let*

$$\kappa = \frac{1}{\alpha}\left(1 - (\mathscr{L} f_X)(\alpha)\right), \tag{6.21}$$

where $(\mathscr{L} f_X)(\alpha) := \int_0^\infty e^{-\alpha y} f_X(y) \mathrm{d}y$ is the Laplace transform of f_X, the density function of each severity component X over the positive support of X evaluated at α. Then there exists the risk-neutral measure $\mathbb{Q} = \mathbb{Q}_F \otimes \mathbb{P}_C$ and the catastrophe-risk and financial markets risk variables under this measure are captured by the following system of equations:

$$\mathrm{d}r_t = (\vartheta_t - a_r r_t)\mathrm{d}t + \sigma_r \mathrm{d}W_t^1, \tag{6.22}$$

$$S_t = S_t^{\mathscr{C}} S_t^{\mathscr{F}}, \tag{6.23}$$

$$S_t^{\mathscr{C}} = \exp\left(-\alpha L_t + \alpha\kappa \int_0^t \lambda_u \mathrm{d}u\right), \tag{6.24}$$

$$S_t^{\mathscr{F}} = S_0 Y_t, \tag{6.25}$$

$$\mathrm{d}Y_t = r_t Y_t \mathrm{d}t + \sigma_S Y_t \mathrm{d}\tilde{W}_t^2, Y_0 = 1 \tag{6.26}$$

$$\mathrm{d}\langle \tilde{W}_t^1, \tilde{W}_t^2 \rangle = \rho \mathrm{d}t, \tag{6.27}$$

$$L_t = \sum_{k=1}^{N_t} X_k, \tag{6.28}$$

where ϑ_t is the risk-neutral parameter for the interest-rate process given explicitly in Equation (6.34) and $\tilde{W}_t^1$ and $\tilde{W}_t^2$ are two Brownian motions under the measure $\mathbb{Q}_F$.

Proof. Equations (6.24) and (6.28) are both an immediate consequence of Assumption 2, in that the processes retain their distributional forms as well as parameters when moving from $\mathbb{P}$ to $\mathbb{Q}$. Now, we consider how to find Equations (6.22), (6.25), (6.26) and (6.27) from Equations (6.1), (6.4), (6.5) and (6.6) respectively.

In the first step we prove that the discounted stock price process,

$$\left\{\exp\left(-\int_0^t r_u \mathrm{d}u\right) S_t^{\mathscr{F}}, t \geq 0\right\},$$

is an $\hat{\mathscr{F}}_t$-martingale under the appropriate chosen market martingale-measure, $\mathbb{Q}_F$.

Classical arguments based on Itô's formula show that this requirement is equivalent to Equation (6.26). We show now how to choose the measure $\mathbb{Q}_F$ to obtain this equation out of Equation (6.5). Define

$$\begin{aligned} B_t^1 &= W_t^1 \quad \text{and} \\ B_t^2 &= \frac{1}{\sqrt{1-\rho^2}} W_t^2 - \frac{\rho}{\sqrt{1-\rho^2}} W_t^1. \end{aligned}$$

By Lévy's Theorem[8], B_t^1 and B_t^2 are two Brownian motions respectively. Moreover, the covariance between B_t^1 and B_t^2 is zero, so the two Brownian motions are uncorrelated. Let

$$\gamma^1 := \rho\sigma_S \quad \text{and} \tag{6.29}$$

$$\gamma_u^2 := \sigma_S\sqrt{1-\rho^2} + \beta_u. \tag{6.30}$$

for some β_u which will be specified later.

We define the risk neutral measure $\mathbb{Q}_F$ on the financial market using inverse of Girsanov martingale, which is also a martingale:

$$\left.\frac{\mathrm{d}\mathbb{Q}_F}{\mathrm{d}\mathbb{P}_F}\right|_{\hat{\mathscr{F}}_t} = \eta(t) := \exp\left(-\frac{1}{2}\int_0^t [(\gamma^1)^2 + (\gamma_u^2)^2]\mathrm{d}u - \int_0^t \gamma^1 \mathrm{d}B_u^1 - \int_0^t \gamma_u^2 \mathrm{d}B_u^2\right).$$

Now by the multidimensional Girsanov Theorem[9] the processes

$$\begin{aligned} \tilde{B}_t^1 &:= B_t^1 + \int_0^t \gamma^1 \mathrm{d}u, \quad \text{and} \\ \tilde{B}_t^2 &:= B_t^2 + \int_0^t \gamma_u^2 \mathrm{d}u, \end{aligned}$$

are standard (uncorrelated) Brownian motions under the measure $\mathbb{Q}_F$. Then, by a further application of Lévy's Theorem, we can define two new correlated Brownian motions, under the measure $\mathbb{Q}_F$, such that

$$\tilde{W}_t^1 := \tilde{B}_t^1 = W_t^1 + \int_0^t \gamma^1 \mathrm{d}u, \tag{6.31}$$

$$\tilde{W}_t^2 := \rho\tilde{B}_t^1 + \sqrt{1-\rho^2}\tilde{B}_t^2 = W_t^2 + \int_0^t \left(\rho\gamma^1 + \sqrt{1-\rho^2}\gamma_u^2\right)\mathrm{d}u, \tag{6.32}$$

$$\mathrm{d}\langle \tilde{W}_t^1, \tilde{W}_t^2\rangle = \rho\,\mathrm{d}t.$$

[8] See Karatzas and Shreve (2012, Chapter 3).
[9] See Karatzas and Shreve (2012, Chapter 3).

Now by choosing

$$\beta_u = \frac{\mu_S - r_u - \sigma_S^2}{\sigma_S\sqrt{1-\rho^2}} \tag{6.33}$$

we obtain $\mu_s - \sigma_S\left(\rho\gamma^1 + \sqrt{1-\rho^2}\gamma_u^2\right) = r_u$. So, inserting this into Equation (6.5) and using (6.32), we obtain Equation (6.26).

From Equation (6.31) and Theorem 6.2 it follows that the Extended Vasicek interest-rate model is preserved, that is, Equation (6.22) holds true. In this case, the new parameter is specified by

$$\vartheta_t := \bar{\vartheta}_r + \sigma_r\gamma^1, \quad \text{and} \tag{6.34}$$

In the second step, we prove that the chosen κ satisfies $\mathbb{E}^{\mathbb{P}_C}\left[S_t^{\mathscr{C}}|\hat{\mathscr{C}}_s\right] = S_s^{\mathscr{C}}$ for $s < t$. In the context of Assumption 3.11, we hence require that, for $s < t$,

$$\mathbb{E}^{\mathbb{P}_C}\left[\exp\left(-\alpha L_t + \alpha\kappa\int_0^t \lambda_u \mathrm{d}u\right)|\hat{\mathscr{C}}_s\right] = \exp\left(-\alpha L_s + \alpha\kappa\int_0^s \lambda_u \mathrm{d}u\right), \tag{6.35}$$

which can be rewritten as

$$\mathbb{E}^{\mathbb{P}_C}\left[\exp\left(-\alpha(L_t - L_s) + \alpha\kappa\int_s^t \lambda_u \mathrm{d}u\right)|\hat{\mathscr{C}}_s\right] = \mathbb{E}^{\mathbb{P}_C}\left[\exp\left(-\alpha(L_t - L_s) + \alpha\kappa\int_s^t \lambda_u \mathrm{d}u\right)\right] = 1. \tag{6.36}$$

Now consider $\mathbb{E}^{\mathbb{P}_C}\left[\exp\left(-\alpha(L_t - L_s)\right)\right]$. This can be simplified as follows.

$$\begin{aligned}
\mathbb{E}^{\mathbb{P}_C}\left[\exp\left(-\alpha(L_t - L_s)\right)\right] &= \mathbb{E}^{\mathbb{P}_C}\left[\mathbb{E}^{\mathbb{P}_C}\left[\exp\left(-\alpha(L_t - L_s)\right)|N_t - N_s\right]\right] \\
&= \mathbb{E}^{\mathbb{P}_C}\left[\exp\left(-\alpha\sum_{k=1}^{N_t - N_s} X_k\right)|N_t - N_s\right] \\
&= \mathbb{E}^{\mathbb{P}_C}\left[\{(\mathscr{L}f_X)(\alpha)\}^{N_t - N_s}\right] \\
&= G_{N_t - N_s}\left((\mathscr{L}f_X)(\alpha)\right) \\
&= \exp\left\{\left[(\mathscr{L}f_X)(\alpha) - 1\right]\int_s^t \lambda_u \mathrm{d}u\right\}, \qquad (6.37)
\end{aligned}$$

$$\tag{6.38}$$

where $G_{N_t - N_s}$ is the probability generating function of the Poisson random variable with mean $\int_s^t \lambda_u \mathrm{d}u$. From Equation (6.36) we have thus $\kappa = \frac{1}{\alpha}\left(1 - (\mathscr{L}f_X)(\alpha)\right)$.

In the last step of the proof we consider the discounted share price process, $\left\{\exp\left(-\int_0^t r_u \mathrm{d}u\right) S_t, t \geq 0\right\}$. Note that by design, its expectation for every $t > 0$ is finite. Using Corollary 3.9, we have, for $s < t$,

$$
\begin{aligned}
\mathbb{E}^{\mathbb{Q}}\left[S_t^{\mathscr{C}} S_t^{\mathscr{F}} \exp\left(-\int_0^t r_u \mathrm{d}u\right) |\mathscr{G}_s\right] &= \mathbb{E}^{\mathbb{Q}}\left[\mathbb{E}^{\mathbb{Q}}\left[S_t^{\mathscr{C}} S_t^{\mathscr{F}} \exp\left(-\int_0^t r_u \mathrm{d}u\right) |\mathscr{G}_s \vee \mathscr{F}_t\right] |\mathscr{G}_s\right] \\
&= \mathbb{E}^{\mathbb{Q}}\left[S_t^{\mathscr{F}} \exp\left(-\int_0^t r_u \mathrm{d}u\right) \mathbb{E}^{\mathbb{Q}}\left[S_t^{\mathscr{C}} |\mathscr{G}_s\right] |\mathscr{G}_s\right] \\
&= S_s^{\mathscr{C}} \mathbb{E}^{\mathbb{Q}}\left[S_t^{\mathscr{F}} \exp\left(-\int_0^t r_u \mathrm{d}u\right) |\mathscr{G}_s\right] \\
&= S_s^{\mathscr{C}} S_s^{\mathscr{F}} \exp\left(-\int_0^s r_u \mathrm{d}u\right),
\end{aligned}
$$

where the last line follows from the first step of the proof. □

We now evaluate the three terms in Equation (6.20) separately.

6.4.1 Coupon payments

We consider $\mathbb{E}^{\mathbb{Q}}[I_1]$, which will be split into the first LIBOR-referencing coupon, and the remaining subsequent ones. One avenue to use in simplifying this term would be to make an approximation to the forward LIBOR rates, and suppose that they are always a constant spread above the riskless rate (see Jarrow (2010), for instance). During the periods of financial stability surrounding the 2007/2008 financial crisis, this assumption held to a considerable extent in practice (taking the OIS rate to be the riskless rate). However, in times of crises, empirical work showed that this spread scales up significantly and the assumption becomes questionable (Hull and White, 2013). We do not use this approximate approach in our analyses but point out this drawback in light of the fact that it has been used in theoretical CAT bond pricing (with the view of obtaining closed-form solutions).

(i) The first LIBOR-referencing coupon (which is known at the outset) can be found by risk-neutral valuation formula:

$$
\mathbb{E}^{\mathbb{Q}}\left[\left(R_0 + c\right)\Delta Z \mathbb{I}_{\{\tau > t_1\}} B(0, t_1)\right],
$$

which is, under Assumption 2, equal to

$$
\left(R_0 + c\right) P(r_0, 0, t_1, \vartheta, a_r, \sigma_r) \mathbb{P}\left(L_{t_1} < D\right) \Delta Z, \tag{6.39}
$$

where R_t and $P(0,t)$ are defined in Equations (6.17) and (6.19), respectively, r_0

is the initial instantaneous interest rate and

$$\mathbb{P}\left(L_{t_i} < D\right) = \exp\left(-\int_0^{t_i} \lambda_u \mathrm{d}u\right) \sum_{n=0}^{\infty} \frac{\left(\int_0^{t_i} \lambda_u \mathrm{d}u\right)^n}{n!} F_X^{n*}(D). \quad (6.40)$$

where $F_X^{n*}(D)$ denotes the $n-$fold convolution of single loss distribution function F_X with itself, evaluated at the positive argument D.

(ii) For the second to n^{th} LIBOR-referencing coupon, we change measure to the respective forward measure. For each $i \in \{2, ...n\}$, we use the t_i forward measure $\mathbb{Q}^{t_i}$, defined to be the forward measure for the numéraire process $P(0,t_i)$ (see Björk (2009, Definition 26.6)). Thus the value of the i^{th} coupon payment at the issue-date is,

$$\begin{aligned}\mathbb{E}^{\mathbb{Q}}\left[\left(R_{t_{i-1}} + c\right)\Delta Z \mathbb{I}_{\{\tau > t_i\}} B(0,t_i)\right] = Zc\mathbb{P}(L_{t_i} < D) P(r_0, 0, t_i, \vartheta, a_r, \sigma_r)\Delta \\ + Z\mathbb{P}(L_{t_i} < D)\mathbb{E}^{\mathbb{Q}}\left[R_{t_{i-1}} B(0,t_i)\right]\Delta,\end{aligned} \quad (6.41)$$

with the respective probabilities given by Equation (6.40). Now consider $\mathbb{E}^{\mathbb{Q}}\left[R_{t_{i-1}} B(0,t_i)\right]\Delta$. By changing to the t_i forward measure[10]), recalling that the forward LIBOR is a $\mathbb{Q}^{t_i}$ martingale[11] and noting the definition of forward LIBOR[12]:

$$\begin{aligned}\mathbb{E}^{\mathbb{Q}}\left[R_{t_{i-1}} B(0,t_i)\right]\Delta &= P(0,t_i)\mathbb{E}^{\mathbb{Q}^{t_i}}\left[R_{t_{i-1}}\right]\Delta \\ &= P(0,t_i) R(0, t_{i-1}, t_{i-1} + \Delta)\Delta \quad (6.42) \\ &= P(0,t_i)\frac{1}{\Delta}\left(\frac{P(0,t_{i-1})}{P(0,t_i)} - 1\right)\Delta \\ &= P(0,t_{i-1}) - P(0,t_i). \quad (6.43)\end{aligned}$$

Inserting Equation (6.43) into Equation (6.42), and adopting our notation for the zero-coupon bond price under the Extended Vasicek model gives the required value at the issue date of the remaining coupon payments, i.e.:

$$\begin{aligned}\mathbb{E}^{\mathbb{Q}}\left[\sum_{i=2}^{N}\left(R_{t_{i-1}} + c\right)\Delta Z \mathbb{I}_{\{\tau > t_i\}} B(0,t_i)\right] = Z\sum_{i=2}^{N}\mathbb{P}(L_{t_i} < D)\left[P(r_0, 0, t_{i-1}, \vartheta, a_r, \sigma_r)\right. \\ \left. + (1 - c\Delta) P(r_0, 0, t_i, \vartheta, a_r, \sigma_r)\right].\end{aligned} \quad (6.44)$$

[10] See Björk (2009, Definition 26.6).
[11] See Björk (2009, Proposition 26.7).
[12] See Brigo and Mercurio (2007, p.29).

6.4.2 Redemption amount

We now consider $\mathbb{E}^{\mathbb{Q}}[I_3]$. By analogous reasoning to Section 6.4.1, we obtain

$$\mathbb{E}^{\mathbb{Q}}\left[Z\mathbb{I}_{\{\tau>T\}}B(0,T)\right] = ZP(r_0,0,T,\vartheta,a_r,\sigma_r)\mathbb{P}(L_T \geq D). \tag{6.45}$$

6.4.3 Conversion feature

Finally, we consider $\mathbb{E}^{\mathbb{Q}}[I_3]$ for a particular type of conversion price. Recall that K_P can either be set as a constant, as the share price at time-of-conversion or be specified as the function $K_P := S_\tau^\nu$, for $\nu \in (0,1]$ as specified in Section 6.3.4. The restriction of ν to this interval is necessary for the Laplace transform of f_X not to be infinite. Notice that if $\nu = 1$, then the conversion price is set equal to the share price. Our analysis below is general and we simply consider either case of

4.3.1 $K_P := S_\tau^\nu$, or

4.3.2 $K_P := K$.

We analyse each in turn below.

Case 1. In this case, we are interested in simplifying

$$\mathbb{E}^{\mathbb{Q}}[S_\tau^{1-\nu}\mathbb{I}_{\{\tau\leq T\}}B(0,\tau)]$$

which can be expanded as

$$\mathbb{E}^{\mathbb{Q}}\Bigg[\exp\left(-\alpha(1-\nu)L_\tau + \int_0^\tau (1-\nu)\left(\alpha\kappa\lambda_u - \frac{\sigma_S^2}{2}\right)\mathrm{d}u + \sigma_S(1-\nu)\tilde{W}_\tau^2 - \nu\int_0^\tau r_u \mathrm{d}u\right) \times S_0\mathbb{I}_{\{\tau\leq T\}}\Bigg]. \tag{6.46}$$

To evaluate Equation (6.46). we will use a specific change of measure. Since N_t in the process L_t is a time-inhomogeneous Poisson process with the intensity λ_t, then the process $X_t^\# := (t, L_t)$ is a particular case of the so-called piecewise deterministic Markov process (see Palmowski and Rolski (2002)), which is defined in Definition 6.3 below. Note that we only present a sketch of the auxiliary facts necessary for the change of measure - the main results upon which our change of measure is based are Theorem 6.5 and the exponential martingale specified by Equation (6.49). For a more detailed presentation of the sketches of the auxiliary facts we present below which beyond the realm of our application, see Palmowski and Rolski (2002) and also Davis (1993, p.52-62).

Definition 6.3. (Palmowski and Rolski, 2002, p.776) Let $\mathscr{I}$ be a countable set and d_f be a function mapping $\mathscr{I}$ to the natural numbers, $\mathbb{N}$. Suppose that E is a state space[13] consisting of the pairs $x = (v, z)$, where z belongs to an open subset $\mathscr{O}_v$ of $\mathbb{R}^{d_f(v)}$ and v assumes a finite number of values from $\mathscr{I}$. A piecewise deterministic Markov process, X_t, is characterised by the following:

(i) $\chi = \sum_{i=1}^{d_f(v)} g^{v,i}(x)\partial/\partial z_i$ for each $v \in \mathscr{I}$, a vector field[14] which determines the flow[15], $\phi_v(t,z)$, of χ and where $g^{v,i} : E \mapsto \mathbb{R}$ are locally Lipschitz continuous[16] functions;

(ii) $\lambda(\cdot)$, the force of transition[17] of the process; and

(iii) $Q(\cdot,\cdot)$, the transition kernel[18] of the process.

Before proceeding, we furthermore denote the boundary of $\mathscr{O}_v$, $\bar{\mathscr{O}}_v \backslash \mathscr{O}_v$ where $\bar{\mathscr{O}}_v$ is the closure of $\mathscr{O}_v$, by $\partial\mathscr{O}_v$ and the boundary of E, $\bar{E}\backslash E$ where $\bar{E}$ is the closure of E, by ∂E. Moreover, let

$$\begin{aligned} \partial^*\mathscr{O}_v &= \left\{ z \in \partial\mathscr{O}_v : z = \phi_v(t,z') \text{ for some } (t,z') \in \mathbb{R}_+ \times \mathscr{O}_v \right\}, \\ \Gamma &= \{(v,z) \in \partial E : v \in \mathscr{I}, z \in \partial^*\mathscr{O}_v\} \text{ (the active boundary)}, \quad \text{and} \\ t^*(v,z) &= \sup\{t > 0; \phi_v(t,z) \text{ exists and } \phi_v(t,z) \in \mathscr{O}_v\}. \end{aligned}$$

Now let T_n denote the consecutive jumps of the process X_t from Definition 6.3. Then, as mentioned in Palmowski and Rolski (2002, p.776), we assume that

$$\lim_{n\to\infty} T_n = \infty \tag{6.47}$$

almost surely under the measure which we consider the piecewise deterministic Markov process.

We now specialise the above to our particular case which pertains to a time-inhomogeneous compound Poisson process. Therefore, in this case one has to consider the empty active boundary, Γ, take the external state index of this process to be equal to one, and take the state space, E, to be $\mathbb{R}^2$ - that is, $x^\# := (t,y) \in \mathbb{R}^2$ is a value of the process $X_t^\#$. Moreover, the differential operator of $X_t^\#$ is taken to be $\chi f(t,x^\#) := \frac{\partial}{\partial t} f(t,x^\#)$, the transition kernel of the process satisfies $Q(x^\#, \mathrm{d}y) = F_X^{\mathbb{Q}}(\mathrm{d}y)$

[13] Note that E must be a Borel space

[14] χ is, in fact, a first-order differential operator - see Davis (1993, p.53).

[15] See Davis (1993, p.53).

[16] See Davis (1993, p.53).

[17] See Davis (1993, p.58).

[18] See Davis (1993, p.58).

(where $F_X^{\mathbb{Q}}$ is the severity distribution of the process L_t under $\mathbb{Q}$) and its jump intensity, $\lambda(t,y)$, equals λ_t. We now proceed to present an informal statement of Theorem 5.3 (without proof) in Palmowski and Rolski (2002), which we use to define a new probability measure. But first, an informal definition is necessary.

Definition 6.4. (Palmowski and Rolski, 2002, Equation 1.1, p.768) Consider a Markov process, X_t^M defined on a filtered probability space and having an extended generator[19], $\mathscr{A}$. If, for some strictly positive function h, the process

$$E^h(t) = \frac{h(X_t^M)}{h(X_0^M)} \exp\left(-\int_0^t \frac{(\mathscr{A}h)(X_s^M)}{h(X_s^M)} \mathrm{d}s\right) \tag{6.48}$$

is a martingale, then it is furthermore said to be an exponential martingale and we therefore say that h is a good function.

Theorem 6.5. *(Palmowski and Rolski, 2002, Theorem 5.3, p. 777) Assume that h is a good function satisfying $H(x) < \infty$ for all $x \in E$, where $H(x) = \int_E h(y)Q(x,\mathrm{d}y)$. Suppose furthermore that Equation (6.47) holds. Then on a new probability space, having the sample space specified by $[0,\infty)$, the process X_t is a piecewise deterministic Markov process with unchanged differential operator χ and the following jump intensity and transition kernel:*

$$\tilde{\lambda}(x) = \frac{\lambda(x)H(x)}{h(x)}, \qquad \tilde{Q}(x,\mathrm{d}y) = \frac{h(y)}{H(x)}Q(x,\mathrm{d}y).$$

Now, in the context of the present research, taking $h(y) = e^{-\alpha(1-\nu)y}$ in Equation (6.48) as a good function (trivially) and applying it to the extended generator for the piecewise linear Markov Process (as given in Palmowski and Rolski (2002, p.776)),

$$\mathscr{A}f(t,x) = \frac{\partial}{\partial t}f(t,x) + \lambda_t\left(\int_0^\infty f(t,x+y)\,F_X(\mathrm{d}y) - f(t,x)\right),$$

of the process L_t produces the following exponential martingale:

$$\bar{\eta}^{(\nu)}(t) := \exp\left(-\alpha(1-\nu)L_t + \varphi(\alpha(1-\nu),t)\right), \tag{6.49}$$

where

$$\varphi(\alpha(1-\nu),t) := \left[1 - (\mathscr{L}f_X)(\alpha(1-\nu))\right]\int_0^t \lambda_u\,\mathrm{d}u. \tag{6.50}$$

and $(\mathscr{L}f_X)(\alpha(1-\nu))$ is the Laplace transform of f_X, for the non-negative support of the random variable X, evaluated at the argument $\alpha(1-\nu)$. Following Theorem

[19] The general form of the extended generator for a Markov process is given in Section 2.2 of Palmowski and Rolski (2002).

6.5, we can now define a new probability measure $\mathbb{P}^{(\nu)}$ via the Radon-Nikodym derivative:

$$\left.\frac{\mathrm{d}\mathbb{P}^{(\nu)}}{\mathrm{d}\mathbb{P}_C}\right|_{\hat{\mathscr{C}}_t} = \bar{\eta}^{(\nu)}(t), \tag{6.51}$$

on which our process L_t remains a compound Poisson process but with altered severity distribution and intensity:

$$\lambda_t^{(\nu)} = (\mathscr{L} f_X)(\alpha(1-\nu))\lambda_t, \quad \text{and} \tag{6.52}$$

$$F_X^{(\nu)}(\mathrm{d}x) = \frac{e^{-\alpha(1-\nu)x} F_X^{\mathbb{Q}}(\mathrm{d}x)}{(\mathscr{L} f_X)(\alpha(1-\nu))}. \tag{6.53}$$

Note that when $\lambda_t = \lambda$ (a constant), then many things simplify. For example,

$$\varphi(\alpha(1-\nu),t) = \lambda \varphi(\alpha(1-\nu))t, \tag{6.54}$$

where

$$\varphi(\alpha(1-\nu)) := \lambda[1-(\mathscr{L} f_X)(\alpha(1-\nu))]. \tag{6.55}$$

By applying the measure change specified by Equation (6.51) and applying Assumption 3.11, we can rewrite Equation (6.46) as

$$\begin{aligned}
&\mathbb{E}^{\mathbb{Q}_F \otimes \mathbb{P}^{(\nu)}}\left[\exp\left(\int_0^\tau (1-\nu)\left(\alpha\kappa\lambda_u - \frac{\sigma_S^2}{2}\right)\mathrm{d}u - \varphi(\alpha(1-\nu),\tau) + \sigma_S(1-\nu)\tilde{W}_\tau^2 - \nu\int_0^\tau r_u \mathrm{d}u\right)\right. \\
&\left.\times S_0 \mathbb{I}_{\{\tau \le T\}}\right] \\
&= S_0 \int_0^T \mathbb{P}_\tau^{(\nu)}(\tau \in \mathrm{d}s)\left\{\exp\left(\int_0^s (1-\nu)\left(\alpha\kappa\lambda_u - \frac{\sigma_S^2}{2}\right)\mathrm{d}u - \varphi(\alpha(1-\nu),s)\right)\right. \\
&\quad\left.\times \mathbb{E}^{\mathbb{Q}_F}\left[\exp\left(\sigma_S(1-\nu)\tilde{W}_s^2 - \nu\int_0^s r_u \mathrm{d}u\right)\right]\right\},
\end{aligned} \tag{6.56}$$

where the last line follows by the independence of $\tilde{W}_t^1$ and $\tilde{W}_t^2$ from the measure change, and $\mathbb{P}_\tau^{(\nu)}$ is the density function of τ under $\mathbb{P}^{(\nu)}$. Notice that Equation (6.56) is a product of three terms, the first the distribution of τ under $\mathbb{P}^{(\nu)}$, the second an exponential term and the third an expectation, under $\mathbb{Q}_F$, of two correlated $\mathbb{Q}$-Brownian motions. We now present the following helpful lemma. Before reading further, recall that in our model we had two correlated standard Brownian motions under $\mathbb{Q}$, $\tilde{W}_t^1$ and $\tilde{W}_t^2$, with correlation coefficient ρ. The purpose of this lemma is to remove $\tilde{W}_s^2$ from the expectation in Equation (6.56).

Lemma 6.6. *Consider a new probability measure, $\bar{\mathbb{Q}}_F$, given by the Radon-Nikodym derivative specified by*

$$\left.\frac{\mathrm{d}\bar{\mathbb{Q}}_F}{\mathrm{d}\mathbb{Q}_F}\right|_{\hat{\mathscr{F}}_t} := \eta^*(t) := \exp\left((1-\nu)\sigma_S\tilde{W}_t^2 - \frac{1}{2}(1-\nu)^2\sigma_S^2 t\right)$$

Then, under the probability measure $\bar{\mathbb{Q}}$ the process $\tilde{\tilde{W}}_t^1 := \tilde{W}_t^1 - \rho\sigma_S(1-\nu)t$ is a standard Brownian motion.

Proof. Note that by a change of measure from $\bar{\mathbb{Q}}$ to $\mathbb{Q}$, for $\zeta > 0$,

$$\begin{aligned}\mathbb{E}^{\bar{\mathbb{Q}}_F}\left[e^{\zeta\tilde{W}_t^1}\right] &= \mathbb{E}^{\mathbb{Q}_F}\left[\exp\left(\zeta\tilde{W}_t^1 + (1-\nu)\sigma_S\tilde{W}_t^2 - \frac{1}{2}(1-\nu)^2\sigma_S^2 t\right)\right] \\ &= \exp\left(-\frac{1}{2}(1-\nu)^2\sigma_S^2 t\right)\mathbb{E}^{\mathbb{Q}_F}\left[(\zeta,(1-\nu)\sigma_S)\cdot(\tilde{W}_t^1,\tilde{W}_t^2)\right].\end{aligned} \tag{6.57}$$

Since the distribution of $(\tilde{W}_t^1, \tilde{W}_t^2)$ is bivariate normal, we can show that Equation (6.57) is equal to

$$\begin{aligned}&\exp\left(-\frac{1}{2}(1-\nu)^2\sigma_S^2 t + \frac{1}{2}\left\{\zeta^2 t + 2\rho\sigma_S(1-\nu)\zeta t + (1-\nu)^2\sigma_S^2 t\right\}\right) \\ &= \exp\left(\frac{1}{2}\zeta^2 t + \rho\sigma_S(1-\nu)\zeta t\right).\end{aligned}$$

□

In consequence, we are able to proceed with simplifying

$$\mathbb{E}^{\mathbb{Q}_F}\left[\exp\left(\sigma_S(1-\nu)\tilde{W}_s^2 - \nu\int_0^s r_u \mathrm{d}u\right)\right] \tag{6.58}$$

from Equation (6.56). Now, change measure from $\mathbb{Q}_F$ to $\bar{\mathbb{Q}}_F$ to obtain that Equation (6.58) is equal to

$$\exp\left(\frac{1}{2}(1-\nu)^2\sigma_S^2 s\right)\mathbb{E}^{\bar{\mathbb{Q}}_F}\left[\exp\left(-\nu\int_0^s r_u \mathrm{d}u\right)\right] \tag{6.59}$$

where $\{r_t, t \geq 0\}$ is the interest-rate process now under the measure $\bar{\mathbb{Q}}_F$. Theorem 6.1 reminds us that the Extended Vasicek interest-rate model specified in (6.22), under $\mathbb{Q}_F$, remains an Extended Vasicek interest-rate model under $\bar{\mathbb{Q}}_F$ with the parameter

$$\vartheta_t^* = \vartheta_t + \sigma_r\rho\sigma_S(1-\nu).$$

The remaining parameters, a_r and σ_r, of the Extended Vasicek interest-rate model remain the same. It can also be shown under the measure $\bar{\mathbb{Q}}$, by a simple application of Itô's Lemma, that the process $\{\tilde{r}_t, t \geq 0\}$, where $\tilde{r}(t) := \nu r(t)$ is an Extended Vasicek interest-rate process with $\bar{\mathbb{Q}}$-dynamics given by

$$\mathrm{d}r_t = (\vartheta_t^\circ - a_r^\circ r_t)\mathrm{d}t + \sigma_r^\circ \mathrm{d}\tilde{\bar{W}}_t^1,$$

where

$$\vartheta_t^\circ = \nu\vartheta_r^* \ ; \ a_r^\circ = \nu a_r \ ; \ \sigma_r^\circ = \nu\sigma_r,$$

Thus, we can rewrite Equation (6.59) as

$$\exp\left(\frac{1}{2}(1-\nu)^2\sigma_S^2 s\right)\mathbb{E}^{\bar{\mathbb{Q}}}\left[\exp\left(-\int_0^s \tilde{r}_u \mathrm{d}u\right)\right] \tag{6.60}$$

Using Equation (6.14), we can simplify Equation (6.60) and obtain that it equals to

$$\exp\left(\frac{1}{2}(1-\nu)^2\sigma_S^2 s\right)P(r_0, 0, s, \vartheta^\circ, a_r^\circ, \sigma_r^\circ). \tag{6.61}$$

Hence, finally, Equation (6.56) can be expressed as an integral involving an infinite sum, that is

$$S_0\int_0^T \mathbb{P}_\tau^{(\nu)}(\tau \in \mathrm{d}s)\exp\left(\int_0^s (1-\nu)\left(\alpha\kappa\lambda_u - \frac{\sigma_S^2}{2}\right)\mathrm{d}u - \varphi(\alpha(1-\nu), s) + \frac{1}{2}(1-\nu)^2\sigma_S^2 s\right)$$
$$\times P(r_0, 0, s, \vartheta^\circ, a_r^\circ, \sigma_r^\circ) \tag{6.62}$$
$$= S_0\int_0^T A(s)\,P(r_0, 0, s, \vartheta^\circ, a_r^\circ, \sigma_r^\circ)\,B(s)\,\mathrm{d}s, \tag{6.63}$$

where

$$A(s) = \exp\left(-(\sigma_S^2/2)\nu(1-\nu)s + (1-\nu)[\varphi(\alpha, s) - \varphi(\alpha(1-\nu), s)] - \int_0^s \lambda_u^{(\nu)}\mathrm{d}u\right)$$
$$B(s) = \lambda_s^{(\nu)}\sum_{n=0}^{\infty}\left[\frac{\left(\int_0^s \lambda_u^{(\nu)}\mathrm{d}u\right)^{n-1}}{(n-1)!} - \frac{\left(\int_0^s \lambda_u^{(\nu)}\mathrm{d}u\right)^{n}}{n!}\right]F_X^{(\nu)n*}(D)$$

and where $F_X^{(\nu)n*}(D)$ denotes the $n-$fold convolution of $F_X^{(\nu)}$ with itself, evaluated at the argument D, and $\varphi(\alpha, s) := \varphi(\alpha(1-\nu), s)|_{\nu=0}$. Recall that $F_X^{(\nu)}$ and $\lambda_t^{(\nu)}$ are given in (6.53) and (6.52), respectively and P in (6.14).

We are now in a position to state the time-zero risk-neutral price of an IL CocoCat in analytical form. We formalise this all in Theorem 6.7.

Theorem 6.7. *In an analytical form, the time-zero risk-neutral price, V_0^{COCO}, of an IL CocoCat with conversion price equal to S_τ^ν for $\nu \in (0,1]$, and assuming the dynamics given in Equations (6.22) to (6.28), is given by*

$$V_0^{COCO} = Z\left(I_1^E + I_2^E + I_3^E\right), \tag{6.64}$$

where

$$\begin{aligned}
I_1^E &= \big(R_0 + c\big) P(r_0, 0, t_1, \vartheta^\circ, a_r^\circ, \sigma_r^\circ) \mathbb{P}(L_{t_1} < D) \Delta \\
&\quad + \sum_{i=2}^{N} \mathbb{P}(L_{t_i} < D) \left[P(r_0, 0, t_{i-1}, \vartheta^\circ, a_r^\circ, \sigma_r^\circ) + (1 - c\Delta) P(r_0, 0, t_i, \vartheta^\circ, a_r^\circ, \sigma_r^\circ) \right]; \\
I_2^E &= S_0 \zeta \int_0^T A(s)\, P(r_0, 0, s, \vartheta^\circ, a_r^\circ, \sigma_r^\circ)\, B(s)\, \mathrm{d}s; \\
I_3^E &= P(r_0, 0, T, \vartheta^\circ, a_r^\circ, \sigma_r^\circ) \mathbb{P}(L_T \geq D)
\end{aligned}$$

and I_1^E is identified in Equations (6.39) and (6.44), I_2^E in Equation (6.63) and I_3^E in Equation (6.45).

Case 2. In the final part of this section, we consider the interesting case pertaining to when K_P is a constant[20] of K. Therefore, to analyze $\mathbb{E}^{\mathbb{Q}}[I_3]$ we will only concern ourselves with simplifying

$$\mathbb{E}^{\mathbb{Q}}[S_\tau \mathbb{I}_{\{\tau \leq T\}} B(0, \tau)]$$

which can be expanded as

$$\begin{aligned}
&\mathbb{E}^{\mathbb{Q}}\left[S_0 \exp\left(-\alpha L_\tau + \alpha\kappa \int_0^\tau \lambda_u \mathrm{d}u + \int_0^\tau r_u \mathrm{d}u - \frac{\sigma_S^2}{2}\tau + \sigma_S \tilde{W}_\tau^2 \right) \mathbb{I}_{\{\tau \leq T\}} \exp\left(- \int_0^\tau r_r \mathrm{d}u \right) \right] \\
&= \mathbb{E}^{\mathbb{Q}}\left[S_0 \mathbb{I}_{\{\tau \leq T\}} \exp\left(-\alpha L_\tau + \int_0^\tau \left(\alpha\kappa\lambda_u - \frac{\sigma_S^2}{2} \right) \mathrm{d}u + \sigma_S \tilde{W}_\tau^2 \right) \right].
\end{aligned} \tag{6.65}$$

Since L_t is a time-inhomogeneous compound Poisson process, then the measure change considered in Section 6.4.3 part (i) can be applied with $\nu = 0$. We denote $\eta(t) = \eta^{(0)}(t)$ for the density process (6.49). Note that $\varphi(\alpha, t) = \varphi(\alpha(1-\nu), t)|_{\nu=0}$.

[20] It is not possible to use Theorem 6.7 to deduce the result, since it is impossible to analytically evaluate P in Equation (6.63).

Upon consideration of this measure change, notice that Equation (6.65) can be rewritten as

$$
\begin{aligned}
&\mathbb{E}^{\mathbb{Q}}\left[S_0\mathbb{I}_{\{\tau\leq T\}}\exp\left(-\alpha L_\tau+\varphi(\alpha,\tau)-\varphi(\alpha,\tau)+\int_0^\tau\left(\alpha\kappa\lambda_u-\frac{\sigma_S^2}{2}\right)\mathrm{d}u+\sigma_S\tilde{W}_\tau^2\right)\right]\\
&=\mathbb{E}^{\mathbb{Q}_F\otimes\mathbb{P}^{(0)}}\left[S_0\mathbb{I}_{\{\tau\leq T\}}\exp\left(\int_0^\tau\left(\alpha\kappa\lambda_u-\frac{\sigma_S^2}{2}\right)\mathrm{d}u-\varphi(\alpha,\tau)+\sigma_S\tilde{W}_\tau^2\right)\right]\\
&=S_0\int_0^T\mathbb{P}_\tau^{(0)}(\tau\in\mathrm{d}s)\exp\left(\int_0^s\left(\alpha\kappa\lambda_u-\frac{\sigma_S^2}{2}\right)\mathrm{d}u-\varphi(\alpha,s)\right)\mathbb{E}^{\mathbb{Q}_F}[e^{\sigma_S\tilde{W}_s^2}],
\end{aligned}
\tag{6.66}
$$

where the last line follows because $\tilde{W}_s^2$ is independent of the measure change for all $s\in[0,T]$, and $\mathbb{P}_\tau^{(0)}$ is the density function of τ under $\mathbb{P}^{(0)}$. By considering the moment generating function of $\tilde{W}_t^2$, we obtain that Equation (6.66) is equal to

$$
\begin{aligned}
&S_0\int_0^T\mathbb{P}_\tau^{(0)}(\tau\in\mathrm{d}s)\exp\left(\alpha\kappa\int_0^s\lambda_u\mathrm{d}u-\varphi(\alpha,s)\right)\\
&=S_0\int_0^T\mathbb{P}_\tau^{(0)}(\tau\in\mathrm{d}s)\\
&=S_0\mathbb{P}_\tau^{(0)}(\tau\leq T) \qquad (6.67)\\
&=S_0\left[1-\exp\left(-\int_0^T\lambda_u^{(0)}\mathrm{d}u\right)\sum_{n=0}^{\infty}\frac{\left(\int_0^T\lambda_u^{(0)}\mathrm{d}u\right)^n}{n!}F_X^{(0)n*}(D)\right], \qquad (6.68)
\end{aligned}
$$

where $F_X^{(0)n*}(D)$ denotes the $n-$fold convolution of $F_X^{(0)}$ with itself, evaluated at the argument D. Notice that the simplification of the expression, $\mathbb{E}^{\mathbb{Q}}[S_\tau\mathbb{I}_{\{\tau\leq T\}}B(0,\tau)]$, as given by Equation (6.67) is highly intuitive. The discounted value of the share price at the time of trigger is simply the existing share price, S_0, multiplied by the probability of trigger under a probability measure that explicitly adjusts for the aggregate loss process.

We can now give the time-zero risk-neutral price of an IL CocoCat with constant conversion price in analytical form. We present this in Theorem 6.8.

Theorem 6.8. *In an analytical form, the time-zero risk-neutral price, $V_0^{COCO,\,K}$, of an IL CocoCat with constant conversion price equal to K, and assuming the dynamics given in Equations (6.22) to (6.28), is given by*

$$
V_0^{COCO,\,K}=Z\left(I_1^E+I_2^E+I_3^E\right). \tag{6.69}
$$

where I_1^E is identified in Equations (6.39) and (6.44), I_2^E in Equation (6.68) (multiplied by ζ/k) and I_3^E in Equation (6.45).

6.4.4 Remarks on Theorems 6.7 and 6.8

We now provide a brief motivation behind the use of Theorems 6.7 and 6.8 for calculating issue-date (i.e. time-zero) IL CocoCat prices. Compared to the use of Monte Carlo simulation for approximating Equation (6.20) directly, the use of our Theorems seems more accurate. Firstly, if the convolutions of the losses under the measure $\mathbb{P}^{(\nu)}$ and $\mathbb{P}^{(0)}$ are known in closed form (for example, if the severities are assumed to follow a gamma distribution), we can approximate the price of the IL CocoCat (and in particular the value of the conversion feature) directly without the need for Monte Carlo simulation. Hence, without any simulation we will be able to compute Equation (6.63) by truncating the summation at a computationally-suitable upper limit, and by discretising the integral. Moreover, without any simulation we will be able to compute Equation (6.68) by truncating the summation at a computationally-suitable upper limit). Such a case can indeed arise if we assume that, under the probability measure $\mathbb{P}$, the losses follow a gamma distribution[21].

Secondly, for more heavy-tailed distributional assumptions of the loss random variables, it is necessary to use Monte Carlo simulation to numerically evaluate the integrals in Equations (6.62) and (6.67). That is, one has to simulate the empirical distribution of the stopping times τ's. This is still more accurate, and in fact faster, than evaluating Equation (6.20) directly by Monte Carlo since in the former case only one process, L_t, has to be simulated compared to the latter, where it is necessary to simulate three processes, namely L_t, r_t and S_t. Such an evaluation can also be done using data from the initial yield curve, especially if the Extended Vasicek model is used (as opposed to the classical Vasicek model, which is not necessarily consistent with the initial yield curve to which it is calibrated).

6.5 Empirical illustration

In this section we present numerical experiments as a first foray into the price behaviour of the IL CocoCat. Gaining an understanding of the IL CocoCat price behaviour, for varying parameters, is crucial in the design stage of such an instrument. Moreover, it could be instrumental in preparing illustrations for pitch books which ILS structurers could use to market new ILS issue types to potential issuers.

In our simulations we pay particular attention to the conversion feature as this will be the item of most interest when issuing the instrument. Also, since exact

[21] The use of Equations (6.62) and (6.67) requires first an exponential tilting of the loss random variable, and thereafter an n-fold convolution. Now, an exponentially-tilted gamma distributed random variable is again gamma distributed, and so too is its convolution. Moreover, the Laplace transform of a gamma random variable exists in closed-form.

closed-form solutions are not available for the underlying loss severity distributions we choose, we use Monte Carlo simulation based on importance sampling, for the loss process. 100,000 simulations are used in each respective instance after applying the importance sampling. Even though it is possible within the context of our model, on the grounds of simplicity we endeavour not to consider a calibration to the initial yield curve. Hence, in this section we consider, simply, the classical Vasicek model for interest rates (which is indeed a special case of the Extended Vasicek model with the parameter ν taken to be a constant). The classical Vasicek model was also considered by Jaimungal and Wang (2006a) for the interest-rate process in valuing 5-year catastrophe-equity put options.

In order to obtain numerical values for the IL CocoCat prices, we need to specify a base set of parameters. Such parameters are specified in Table 6.2.

Tab. 6.2: Selected parameter values for the IL CocoCat price numerical illustration.

Compound Poisson loss process parameters		**Values**
λ_t	Intensity of catastrophe loss process	Equation (5.11)
c_b	Shape parameter of Burr severity distribution	1.57
k_b	Shape parameter of Burr severity distribution	0.7
ζ_b	Scale parameter of Burr severity distribution	9.53×10^7
Interest-rate process parameters under $\mathbb{Q}$		
r_0	Initial instantaneous interest rate	0.02
ϑ	Model parameter	0.04
σ_r	Instantaneous volatility	0.1
a_r	Model parameter	1.125×10^{-3}
Issuer's share price process parameters under $\mathbb{P}$		
S_0	Initial share price	10
ρ	Correlation coefficient of share and interest-rate processes	-0.5
α	Effect of losses on log share price	5.81×10^{-11}
IL CocoCat parameters		
K_P	Constant conversion price	Varies
ν	Power parameter for conversion price	Varies
ζ	Conversion fraction	0.2
Δ	Tenor: time between coupon payments	0.25
c	Constant spread (CAT risk premium)	0.1
Z	Nominal amount	1
T	Term	Varies
D	Threshold level	Varies

Firstly, it is necessary to select parameters for the compound Poisson process underlying the IL CocoCat. We do this by using the fitted intensity, shown in Equation (4.12) in Chapter 4 and furthermore in Equation (5.11) in Chapter 5.

Moreover, we want our process to be based on a heavy-tailed underlying severity distribution, so that low-frequency and high-severity disasters are indeed accounted for. Evidence for heavy-tailed underlying distributional properties in catastrophe-related economic and insured losses has been found by Levi and Partrat (1991); Burnecki *et al.* (2000); Milidonis and Grace (2008); Braun (2011) and Ma and Ma (2013). In all endeavours to assess which heavy-tailed distribution fits such data, it appears from a review of the literature on catastrophe-related ILS that the Burr distribution consistently comes out best. In this section, we indeed use the Burr type XII distribution with probability density function, $f(x, \zeta_b, c_b, k_b)$, specified in Table 4.2 from Chapter 4.

For illustrative purposes, we consider the Burr distribution fitted in Chapter 4 (see Table 4.3): note that $c_b = 1.57$, $k_b = 0.7$ and $\zeta_b = 0.53 \times 10^7$. In passing, note that it would be possible to add risk margins to our various parameters of the loss process derived from real-world data. This is in line with some of the traditional actuarial approaches (see for example the recent work by Chang and Chang (2017) for a brief discussion on this point, and also Braun (2011) for a brief mention), but as always, such a choice is subjective and specific to the modeller. Additionally, it must be noted that severity distributions (and also intensity functions) fitted to both real-world and additional synthetically-simulated data from sophisticated natural catastrophe models can also be employed in our context. In discussion with industry practitioners it would appear that this is the approach used to price natural catastrophe-related instruments in practice.

Secondly, we consider the interest-rate parameters. We reiterate that simulation of the interest-rate process is not necessary given the simplification of the IL CocoCat pricing formulae in Theorems 6.7 and 6.8. We let the risk-neutral parameters, ϑ, a_r and σ_r, equal 0.02, 0.04, 0.1 and 1.125×10^{-3} respectively. Moreover, we specify the correlation coefficient, ρ, to be -0.5. All of these parameters are in line with Jaimungal and Wang (2006a) (who used the Vasicek interest-rate model), Lo *et al.* (2013) and Wang (2016) (both of whom used the Cox-Ingersoll-Ross interest-rate model), as well as previous literature concerning the modelling of interest-rates in the context of insurance (see, for example, Duan and Simonato (2002) and Chang *et al.* (2011)).

Thirdly, we consider the share price process' parameters. We set the initial share price, S_0, to be 10. The effect of the catastrophic losses on the logarithm of the share price, α, is found in a similar fashion to Jaimungal and Wang (2006a). We let α represent the percentage drop in the log share price per unit of expected loss, that is

$$\alpha = \frac{\delta}{E^{\mathbb{P}}[X_k]},$$

and we consider the case, as in Jaimungal and Wang (2006a), where $\delta = 0.02$. In consequence, $\alpha = 5.81 \times 10^{-11}$. Despite α being very small in size, its effect is still prevalent since it is multiplied by L_t in Equation (6.24), which is relatively much larger than α.

Finally, we give thought to the various parameters for the IL CocoCat itself. We set the contractually-specified conversion fraction, ζ, equal to 20%, which is in line with previous literature on CocoCats (see Georgiopoulos (2016)), and we let the tenor be 3 months (in line with Jarrow (2010)). For illustrative purposes we let the IL CocoCat spread be 10%, and set the nominal value of the bond, Z, equal to 1. Also, we will let certain parameters vary, namely the term (T), the threshold level (D) and ν, in order to assess how the IL CocoCat time-zero price varies with changes in the aforementioned parameters.

We now comment on how one can estimate numerical values for Equations (6.62) and (6.67) via Monte Carlo simulation under the measures $\mathbb{P}^{(\nu)}$ and $\mathbb{P}^{(0)}$ respectively. In order to evaluate these integrals, it is necessary to develop empirical distributions for the stopping time τ under the respective measures by simulating the paths for the process L_t by considering the value of L_t for each $t < T$. To simulate paths for L_t under the measures $\mathbb{P}^{(\nu)}$ and $\mathbb{P}^{(0)}$, it is necessary to know the intensity and severity distributions under these measures and the necessary links are provided by Equations (6.52) and (6.53) respectively. Equation (6.52) is easy to find given that we numerically know the Laplace transform of the severity random variable. However, coping with Equation (6.53) is not immediately obvious for Burr-distributed severity random variables. This is primarily because the exponentially-tilted Burr distribution does not have a density that is known and computable. Therefore, to simulate losses from Equation (6.53), we employed the acceptance-rejection algorithm (Von Neumann, 1951). We point out that such a simulation technique can be used in any setting where the measure change is specified by Equation (6.53). Note that such a simulation technique[22] is common in the sphere of simulating random variables, under other probability measures, from those random variables under a given, complex measure change (Asmussen and Glynn, 2007). We proceed as follows: if $f(x, \zeta_b, c_b, k_b)$ is the pdf of the Burr-distributed losses under $\mathbb{Q}$ (which is invertible and known), then we can generate from the density $f^{(\nu)}(x, \zeta_b, c_b, k_b)$ by generating first from $f(x, \zeta_b, c_b, k_b)$ directly[23], and then applying

[22] Note that exponential tilting is a case in point.
[23] To simulate efficiently from the heavy-tailed density $f(x, \zeta_b, c_b, k_b)$, we use importance sampling

the acceptance-rejection algorithm with the constant $c_R := [(\mathscr{L} f_X)(\alpha(1-\nu))]^{-1}$. Under the measure $\mathbb{P}^{(\nu)}$, the simulation algorithm is summarised in the following key steps. The algorithm for $\mathbb{P}^{(0)}$ is analogous.

Key steps of algorithm to generate L_t

---

1. For the interval t_{i-1} to t_i, generate a Poisson realisation, n_i, with intensity $\int_{t_{i-1}}^{t_i} \lambda_u \mathrm{d}u$.

2. Generate n_i realisations from $f^{(\nu)}(x, \zeta_b, c_b, k_b)$ by the acceptance-rejection algorithm:

  (i) Generate U from $\mathscr{U}(0,1)$ and independent X with density $f(x, \zeta_b, c_b, k_b)$.

  (ii) If $U < \frac{f^{(\nu)}(X,\zeta_b,c_b,k_b)}{c_R f(X,\zeta_b,c_b,k_b)}$ then accept X else return to step (i).

  (iii) Continue until n_i realisations from $f^{(\nu)}(x, \zeta_b, c_b, k_b)$ are drawn.

3. Continue until time t_N, consecutively adding the accepted realisations.

Before presenting the numerical results, we now give a brief overview of the remainder of Section 6.5. The aspect of the IL CocoCat which will matter most will be the conversion feature (i.e. I_2 from Theorem 6.7), and consequently we will focus on the analysis of different types thereof. Two broadly different conversion prices are considered for the conversion feature and are each evaluated by Monte Carlo simulation. Probabilities relating to the loss process in Equations (6.39), (6.44) and (6.45) (i.e. $\mathbb{P}(L_{t_i} < D) = \mathbb{P}(\tau >= t_i)$) are also evaluated by simulating the distribution of the stopping times via Monte Carlo. However, the remaining terms relating to the interest-rate process in Equations (6.39), (6.44) and (6.45) as well as in Equation (6.62) are evaluated via the closed-form solutions available. In light of this, Section 6.5.1 considers the case when the IL CocoCat has a constant conversion price of K_P, while Section 6.5.2 looks at the case when the conversion price is a function of the share price, that is $K_P = S_\tau^\nu$ for $\nu \in (0,1]$. The latter includes the special case when $K_P = S_\tau$, which is of interest since the conversion price is set to the share price at time-of-conversion. Finally, Section 6.5.3 compares the price behaviour of the IL CocoCat across three cases: when K_P is a constant, $K_P = S_\tau$ and $K_P = S_\tau^{0.5}$.

6.5.1 Case 1: K_P is a constant

We begin by studying the behaviour of the IL CocoCat price in the context of a changing constant conversion price (K_P) and threshold level (D). This is illustrated

(see Section 4.3.5 of Chapter 4).

in Figure 6.3 panel (a). Note that, by the design of Equations (6.22) to (6.26), the interest-rate process does not affect the value of this IL CocoCat's conversion feature, $(\zeta/K_P)\mathbb{E}^{\mathbb{Q}}[S_\tau Z\mathbb{I}_{\{\tau\leq T\}}B(0,\tau)]$, at all.

From Figure 6.3 panel (a), it is clear that for all constant strike prices, the IL CocoCat prices level out around a value of approximately 1.5 per unit nominal as the threshold level increases. Intuitively this makes sense as chances of trigger at such high threshold levels become smaller and smaller, making the conversion feature less valuable and the coupon and redemption payments more valuable. In fact, for these high threshold levels the IL CocoCat appears to behave like a corporate bond (ignoring credit risk). However, for lower threshold levels the conversion feature which comprises a partial write down of the principal invested, is more valuable compared to the redemption (or principal) amount and subsequently IL CocoCat prices are lower.

It is also interesting to study the IL CocoCat price behaviour over different terms, the numerical results of which are shown in Figure 6.3 panel (b). It shows that the longer the term, the lower the price of the IL CocoCat. This is because for longer terms, there is a greater probability of trigger and in consequence there is a higher risk of an investor losing $1-\zeta$ per unit nominal. But as expected, this decline in IL CocoCat price is slower for those with higher threshold levels.

6.5.2 Case 2: $K_P = S_\tau^\nu$

We now study the case when the conversion price is a function of the share price at time-of-conversion, in the context of a changing value of ν and threshold level D. Figure 6.4 panel (a) shows this. Also, note that changing interest-rate model parameters does have an impact on the price of an IL CocoCat structured in this way, which prompted further endeavours to study this effect. The results of this investigation are shown in Figure 6.4 panel (b), whereby we analyse the effect of altering the two interest-rate process parameters, σ_r and θ_r, on the IL CocoCat price for different terms.

Consider Figure 6.4 panel (a). As a function of ν and threshold level, similar behaviour to the case when the conversion feature is set at a constant level (i.e. Case I) is seen. Given that the conversion fraction ζ is relatively small at 0.2, the results seen make intuitive sense as the conversion fraction has a small effect on the overall price of the IL CocoCat. The time-zero value of the coupons and principal amount have a markedly greater impact on the price.

If ζ is increased in size, of which it can be, then the IL CocoCat price behaviour is observed to change and the time-zero value of the conversion price has a visible effect on increasing the price of the IL CocoCat. In fact, from our numerical simu-

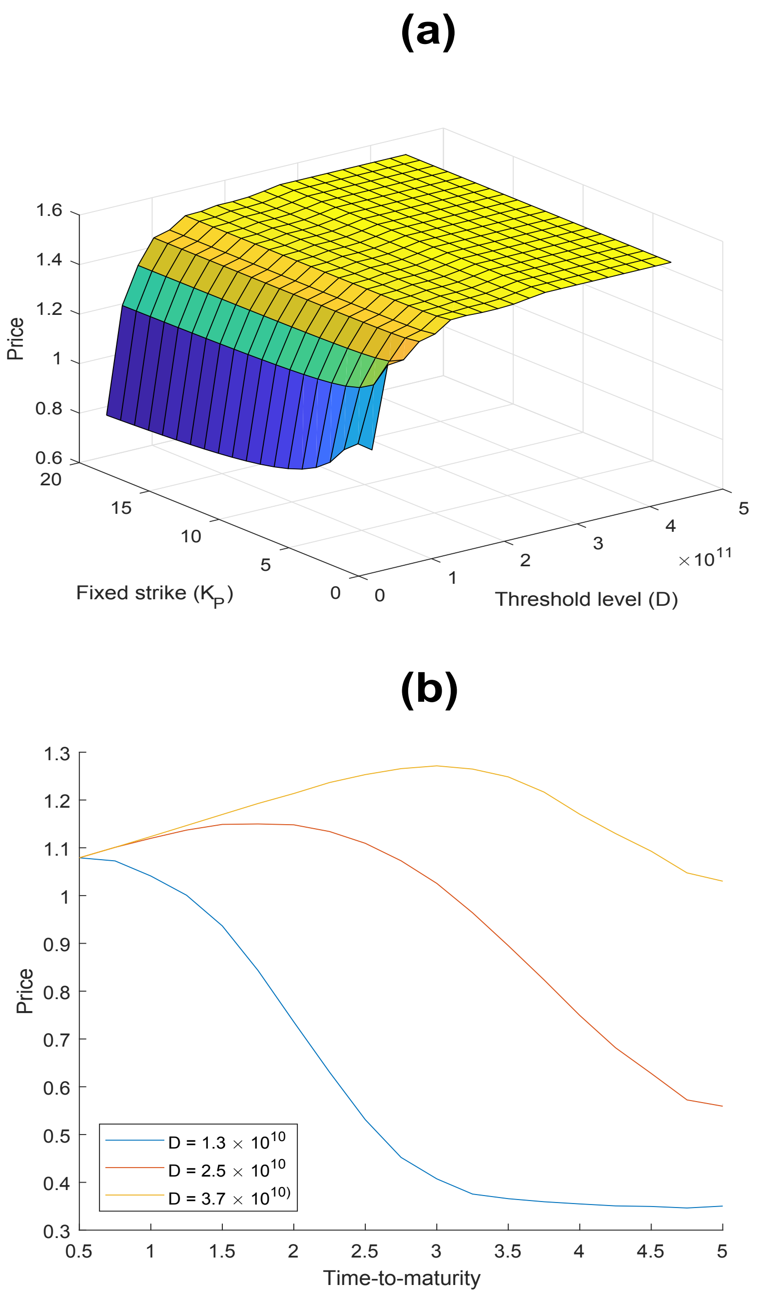

Fig. 6.3: (a) Price of IL CocoCat as a function of conversion price (K_P) and threshold level (D), calculated using 100,000 Monte Carlo simulations of the loss process, L_t. (b) Price of IL CocoCat as a function of term (T) for three different threshold levels, calculated using 100,000 Monte Carlo simulations of the loss process, L_t, taking $K_P = 8$.

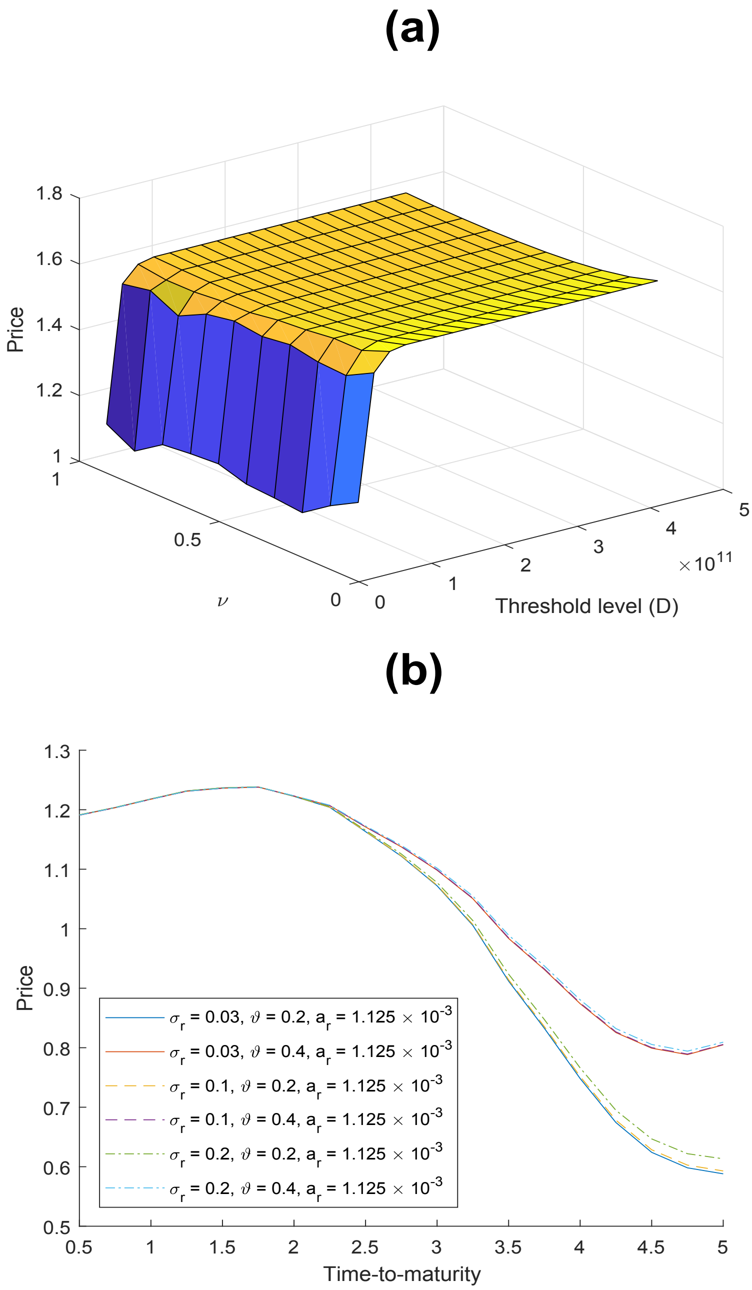

Fig. 6.4: (a) Price of IL CocoCat as a function of ν and threshold level (D), calculated using 100,000 Monte Carlo simulations of the loss process, L_t. (b) Price of IL CocoCat as a function of term (T) for three different interest-rate volatilities and two different values for the parameter ϑ, calculated using 100,000 Monte Carlo simulations of the loss process, L_t, taking $\nu = 0.5$.

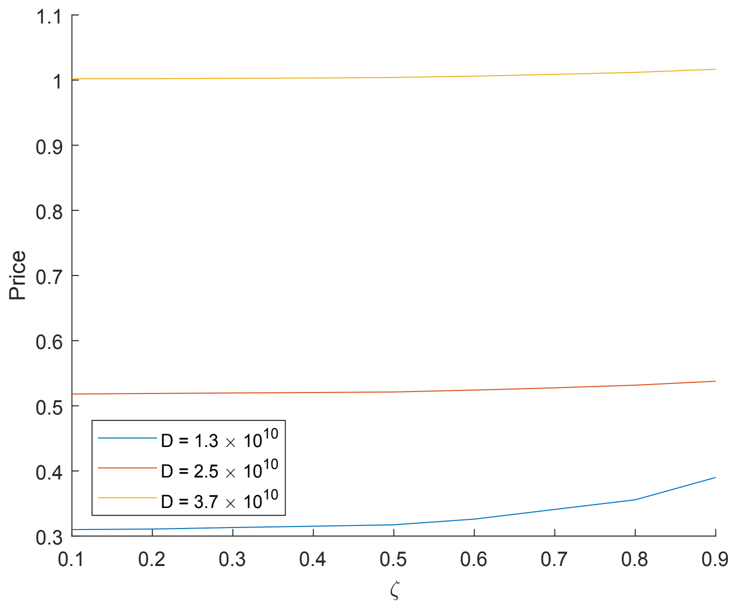

Fig. 6.5: (a) Price of IL CocoCat as a function of ζ and three different threshold levels (D), calculated using 100,000 Monte Carlo simulations of the loss process, L_t, and taking $v = 0.5$.

lations illustrated in Figure 6.5, it is clear that from values for ζ of 0.8 and higher, higher time-zero prices for the IL CocoCat can arise. However, we highlight that such a high value for ζ is potentially unrealistic in practice, since not much of the IL CocoCat will be available for use by the issuer as capital relief and moreover there is a large risk of dilution. Moreover, on the basis of other numerical analyses which we performed, it was interesting to note that the lower the conversion price (i.e. the smaller v is), the greater the effect of increasing ζ on the time-zero IL CocoCat price.

Furthermore, we studied the effect of changing interest-rates on the price of the IL CocoCat, by analysing the effects of changing instantaneous volatility and the parameter ϑ in Figure 6.4 panel (b). Only for medium to long terms, did changing interest-rates have a visible effect, with the parameter ϑ being the greatest influence leading to this change. Note that at longer terms, the different interest-rates had a large impact on the time-zero values of the coupons and principal amount, hence leading to the different values observed for the IL CocoCat time-zero price for these terms.

We close this section with the following remark. From our numerical simulations, it seems evident that the IL CocoCat price is most sensitive to the threshold level, D. This is particularly the case for lower threshold levels, but for higher threshold levels (i.e. those in excess of 10^{10}), it does not make much of a difference. The latter case is true since the probabilities of trigger for these higher threshold levels are extremely small and do not differ to a noticeable extent for small changes in these very high threshold levels.

6.5.3 Comparison of three conversion prices

Finally, we compare 5-year IL CocoCat time-zero prices for three[24] different conversion prices: when K_P is a constant at a value of 8 (denoted by $V_0^{\mathrm{COCO},1}$), $K_P = S_\tau$ (denoted by $V_0^{\mathrm{COCO},2}$) and $K_P = S_\tau^{0.5}$ (denoted by $V_0^{\mathrm{COCO},3}$). Table 6.3 lists the time-zero prices of IL CocoCat for the different conversion prices, $V_0^{\mathrm{COCO},1}$, $V_0^{\mathrm{COCO},2}$ and $V_0^{\mathrm{COCO},3}$, for different threshold levels, D. We focus our comparison mainly on lower threshold levels since the differences in price, per threshold level, were visible to at least three significant figures. As we remarked in Section 6.5.2, for higher threshold levels price differences across different threshold levels were exceptionally small.

Tab. 6.3: Time-zero IL CocoCat prices for different threshold levels, across the three conversion price structures, $V_0^{\mathrm{COCO},1}$, $V_0^{\mathrm{COCO},2}$ and $V_0^{\mathrm{COCO},3}$.

D	$V_0^{\mathbf{COCO},1}$	$V_0^{\mathbf{COCO},2}$	$V_0^{\mathbf{COCO},3}$
1.3×10^{10}	0.318	0.291	0.299
1.8×10^{10}	0.402	0.277	0.301
2.3×10^{10}	0.531	0.441	0.488
2.9×10^{10}	0.692	0.622	0.671
3.4×10^{10}	0.909	0.914	0.944
4.0×10^{10}	1.262	1.277	1.296
9.5×10^{10}	1.490	1.499	1.519
2.5×10^{11}	1.528	1.528	1.528
3.5×10^{11}	1.528	1.528	1.528

For each of $V_0^{\mathrm{COCO},1}$, $V_0^{\mathrm{COCO},2}$ and $V_0^{\mathrm{COCO},3}$, it is still the case that the price levels out for very high threshold levels. This effect was also detected in Section 6.5.1 and 6.5.2: the conversion price (and hence feature) has little effect on the IL CocoCat time-zero price. However, Table 6.3 shows that – within the constraints of our simulation exercise – $V_0^{\mathrm{COCO},2}$ is always less than or equal to $V_0^{\mathrm{COCO},3}$, which is expected since the conversion option is more valuable due to the conversion price being lower than the share price at time-of-conversion. However, it is interesting to analyse the relationship between $V_0^{\mathrm{COCO},1}$ and $V_0^{\mathrm{COCO},2}$. At lower threshold

[24] For ease of notation, we depart slightly from the notation we used in Theorems 6.7 and 6.8.

levels, it is clear from Table 6.3 that $V_0^{\text{COCO},1}$ exceeds $V_0^{\text{COCO},2}$. For these low threshold levels, it is likely that conversion will occur more quickly (than conversion for higher threshold levels), and moreover such small losses and the effect of changing interest-rates over a shorter term will not greatly impact the initial share price of 10. Hence, at conversion time an investor in the IL CocoCat with a constant conversion price of 8 will, in all likelihood, purchase the share for a value less than its current value at conversion time, leading to a higher value for the conversion feature. However, for higher threshold levels (i.e. 3.4×10^{10} and higher), $V_0^{\text{COCO},2}$ is greater than or equal to $V_0^{\text{COCO},1}$. At such high threshold levels, the large insured losses now have a more marked impact on the share price, leading to a potentially depressed share price at conversion time, which may well be lower than the fixed conversion price of 8. Additionally, it takes more time for the loss process to reach the higher threshold level which allows for further uncertainty in the interest-rate movements. Overall, this interest-rate uncertainty coupled with the possibility of depressed share prices will, in all likelihood, lead to a lower value for the conversion feature.

We end by emphasising that it is up to the issuer to select an appropriate threshold level for the contractually-specified conversion price. The threshold level should not be set too high (which is the option which will be preferred by the investor) so that conversion never happens. In addition, the conversion price should not be set too low, so that the probability of trigger is high (despite the lower IL CocoCat price). This is a situation which may be unfavourable to the investor.

6.6 Concluding remarks

In this chapter, we formalised the design of a CocoCat, which is a type of ILS similar to the traditional Coco bonds issued by banks. We linked it to the already existing Coco bond literature, and moreover included a brief review of an existing CocoCat issued by SwissRe in 2013. Moreover, we motivated as to why there is a potential need for insurers and reinsurers to issue CocoCats. Amongst other reasons, CocoCats allow their issuers to tap into a broader pool of financing offered by the capital markets.

Subsequently, we went on to price a special type of Cococat, namely an IL CocoCat linked to the PCS loss index. The Extended Vasicek interest-rate model was chosen for the interest-rate dynamics. We were able to find intuitive, analytical expressions for the price via, firstly, our assumption of independence between catastrophe-risk and financial markets risk variables. Secondly, an exponential change of measure allowed us to separately deal with financial markets as well as catastrophe-risk variables, and a Girsanov-like transformation allowed us to synthetically remove a Brownian motion from the expectation containing two correlated Brownian motions. Finally, we arrived at an analytical expression for the conversion feature (and hence the price) which only required simulation of the loss process in order to empirically estimate the distribution of the time-of-trigger of the equity conversion feature. We did note that Monte Carlo simulation could be used to estimate the value of the conversion feature of the IL CocoCat directly. However, our simplification to an analytical formula had more in its favour, since only one process had to be simulated, namely the insured loss index, while the interest-rate and stock price processes did not have to be simulated.

We finished off our work in this Chapter by presenting a numerical analysis into the prices of the IL CocoCat, and we believe such an analysis is crucial in the design stage of this instrument and for use in pitch books. The prices we obtained in our analyses conformed to intuition: the higher the threshold level of the IL CocoCat, the greater the price. But for exceptionally high threshold levels, IL CocoCat prices did not vary much which may be a limitation of the instrument. We also found evidence suggesting that IL CocoCat price behaviour was quite sensitive to three design aspects: interest-rates, the threshold level and the conversion fraction, with the former impacting the price to a considerable extent in light of the fact that the IL CocoCat is a fairly long-term investment. We found evidence for the IL CocoCat price being most sensitive to the threshold level, especially for lower threshold levels and not, considerably, for higher ones. But also, since the IL CocoCat is linked to interest-rates, the price behaviour with changing interest-rates was studied and for terms of three to five years the instrument's price was fairly sensitive to changes in the mean-reversion rate of the assumed interest-rate process. Finally, since the conversion fraction impacts the value of the conversion feature, price sensitivity towards the conversion fraction was also found.

Chapter 7

Conclusions and suggestions for further research

7.1 Conclusions

This thesis set out to contribute to the ever-growing twin fields of index-linked catastrophe instrument design and valuation. We recapitulate on our contributions, given the limitations highlighted in Section 1.2.

In Chapter 3, a catastrophe instrument pricing model (for instruments based on a loss process), which is especially useful for valuing index-linked catastrophe instruments such as zero-coupon and coupon-paying index-linked catastrophe bonds, was proposed. By assuming that catastrophe-risk and financial markets risk variables are independent, it affords us the ability to calibrate to the catastrophe risk processes and financial markets risk variables independently. In Chapter 6, our proposed model was extended to the valuation of contingent convertible catastrophe bonds.

Chapter 4 was devoted to building on Chernobai *et al.* (2006)'s procedure for modelling left-truncated data via a compound non-homogeneous Poisson process. The contribution we made included the act of modifying the fitting process introduced by Chernobai *et al.* (2006) so that it becomes systematically applicable in the context of data that is not only left-truncated, but heavy-tailed as well. We also proposed a Monte Carlo importance sampling algorithm which ensured that large losses are satisfactorily simulated. Our fitting procedure was applied to relatively recent data from the left-truncated Property Claims Services index, and we concluded that the insurance losses from such an index are suitably modelled by a Burr and generalised extreme value distribution.

Chapter 5 looked into finding simple closed-form approximations to zero-coupon and coupon-paying index-linked catastrophe bond prices. The catastrophe-risk variables were assumed to follow a compound non-homogeneous Poisson process, and had heavy-tailed underlying severity distributions. A unique feature of our approximation is that it only demands finiteness of the first moment of the aggregate loss processes, a feature of heavy-tailed distributions. Our approximation compared favourably with Monte Carlo simulation, as well as the first-order single risk loss process approximation.

In Chapter 6, we comprehensively specified, compared and analysed a relatively new insurance-linked security, namely the contingent convertible catastrophe bond. In an effort to make our analysis as practical as possible, we also looked into the potential accounting treatment (in the UK and US) for such an instrument, and concluded that the instrument should be valued at fair value through profit-and-loss under IFRS 9. We went on to find an analytical expression for its time-zero value, and emphasised that an advantage of using our analytical expression over for example, a direct Monte Carlo simulation, is that only simulation of the underlying loss process (assumed to be a compound non-homogeneous Poisson process) was required. The chapter ended by presenting a first numerical foray into the price behaviour of this instrument, and concluded that the instrument's price behaviour was quite sensitive to three aspects - its threshold level, interest-rates and the conversion fraction. Such an analysis is crucial for ILS structurers when designing contingent convertible catastrophe bonds.

So overall, this thesis has made four contributions to the valuation of index-linked catastrophe instruments. However, we would have appreciated working more closely with those in industry actively working with insurance-linked securities for these three main reasons:

- We believe that a collaboration[1] with industry practitioners would have assisted us in obtaining robust index-linked catastrophe-bond pricing data.
- We also believe that this would have offered us the opportunity to be exposed to more sophisticated natural catastrophe models (such as those developed by Risk Management Solutions, AIR Worldwide and PERILS).
- We finally are of the opinion that a closer relationship with practitioners would have allowed us to align our index-linked catastrophe instruments' setup, design and specification with what they would prefer.

In general what our research further highlights is that there is a greater need for collaboration between academics and practitioners in firstly, designing suitable alternative-risk transfer instruments and secondly, robustly modelling the random processes underlying such instruments. We hope that this call is answered in the near future.

[1] Indeed, this came up in discussions between the author and Professor Paul Embrechts.

7.2 Suggestions for further research

Some ideas for future research reside in the following suggestions:

7.2.1 Modelling catastrophe data:

- Our distribution fitting problem does seem to fit squarely into the ambit of extreme value modelling. Therefore, it may also be suggested to look into the peaks-over-threshold modelling methodology and fitting the excess of each loss in the Property Claims Services index above a selected threshold to a generalised Pareto distribution. To cater for the effects of left-truncation, recent work by Beirlant *et al.* (2017) may be considered, as well as the classical reference of Beirlant *et al.* (2006). However, we do caution that the choice of a suitable threshold level will be difficult, since it may be higher than the loss amounts included in the Property Claims Services index. One solution here would be to split the distribution fitting into two separate exercises: fitting a candidate distribution to the data below the selected threshold, and fitting the generalised Pareto distribution to the excessses above the threshold. It must be noted, however, that the parameters of the fitted generalised Pareto distribution are sensitive to the choice of threshold level. And importantly, existing methods to find such threshold levels are rather *ad hoc*: one of the most frequently used procedures relies on a rough inspection of the estimated mean-excess plot and selecting the threshold at the point where the mean-excess plot begins to follow a straight line.

- The goodness-of-fit tests, modified to suit left-truncated data, of Chernobai *et al.* (2015) have not been assessed in terms of their (empirical) size control and power. To check this, one could consider simulation experiments based on data generated from the distributions we fitted to the Property Claims Services data. Within this context, different threshold levels could also be considered. The same analysis could be performed for Moran's log spacings statistic.

- In Chernobai *et al.* (2015) adjusted goodness-of-fit test statistics were derived, but were not proven to satisfy the limiting distributional behaviours which we know hold for most of the conventional goodness-of-fit test statistics (see Darling (1957); Anderson and Darling (1954) and also Stephens (1974)). It would be interesting to perform simulation analyses to check the quality of the adjusted computing formulae on the basis of these limiting distributional behaviours, on the basis of heavy-tailed data.

- It would be instructive to look into the natural catastrophe scientific simulation models[2] used by catastrophe-modelling vendor firms such as AIR Worldwide and Risk Management Solutions. An interesting exercise would be to compare the results of our simulations of the Property Claims Services index to those predicted by these vendor models.
- The catastrophe risk securitised in natural-catastrophe insurance-linked securities often related to events with recurrence periods of 50 years or more. In the (smaller) datasets we use, one may argue that there is simply not enough (extreme) data on these such historical observations. Therefore, it may be suggested to augment our current dataset with synthetic observations (based on myriads of scenarios created by natural catastrophe scientific simulation models, such as those created by AIR Worldwide and Risk Management Solutions) in order to create new probability distributions which assign more probability to extreme events.
- If we had access to primary and secondary catastrophe bond market data, we could perhaps expand our contribution by attempting to find real-world evidence that risk modellers and investors do underestimate the expected losses arising from the Property Claims Services loss index (due to the left-truncation feature).
- Finally, one could consider the application of our modelling procedure for heavy-tailed left-truncated data to other areas, such as operational risk modelling in banks.

7.2.2 Pricing of catastrophe bonds:

- One could consider different payoff structures for catastrophe bonds: an interesting payoff structure would be to include an adjustment to the recovery amount based upon the amount by which the insured losses exceed the threshold level.
- Our catastrophe bond pricing framework could be empirically considered in the context of interest-rates modelled by a single- or multi-factor interest rate model, such as the Cox-Ingersoll-Ross single factor model or the G2++ or even the Longstaff-Schwartz two-factor models. It would be more beneficial, and from a practical perspective more useful, if these theoretical yield curves (on the basis of these aforementioned models) were calibrated to the market yields for the markets within which catastrophe bonds are commonly issued.

[2] Since these are proprietary models, it would be difficult to obtain access to them.

- A further interesting study would be the selection of an optimal interest-rate model to model the interest-rates upon which catastrophe bonds are commonly based. If one had access to a large database of catastrophe bond prices, one could partly perform such an exercise by calibrating the interest rate process parameters to the interest rates which the catastrophe bonds reference and consequently ascertain how closely the true catastrophe bond prices are "replicated" by the model's prices.

- If suitable index-linked catastrophe bond pricing data (either from the primary or secondary markets) can be obtained, it would be a compelling exercise to calibrate the parameters of the insurance-loss process (under our catastrophe bond pricing model) to these prices. This would then do away with our assumption that catastrophe bond investors are risk-neutral to catastrophe risk (see Assumption 3.11 in Chapter 3). Also, the exercise could allow for market prices of risk to be extracted.

- The distributional-fitting results of Chapter 4 could be applied within the context of the LFC model (see Lane (2003)). Also, one could perform a comparison exercise, effectively comparing catastrophe bond prices under the LFC model to those of our catastrophe bond pricing model from Chapter 3.

- Certain catastrophe bonds in issue today have multiple (most commonly dual) triggers, so it could be possible to consider a multidimensional approach to modelling the underlying aggregate renewal processes. In doing so, we could price such catastrophe bonds using a multidimensional version of the pricing framework from Chapter 3. As an extension to the analysis, we could possibly investigate a weak multidimensional convergence of the aggregate renewal processes to multidimensional $\alpha-$stable Lévy processes in view of attempting to obtain closed-form approximations to the catastrophe bond prices.

- Since there is no canonical model for the pricing of catastrophe bonds (and indeed other long-dated insurance contracts), alternative pricing models for stylized catastrophe bonds could be considered. In line with the Benchmark Pricing approach of Platen and Heath (2006), a loading pricing idea combining the theoretically possible minimal price of a catastrophe bond with its formally-obtained risk-neutral price could be investigated. Such a pricing technique could allow for these long-dated catastrophe bonds to be valued less expensively and with a higher investment return compared to the pricing approach proposed in Chapter 3 of this thesis.

7.2.3 Pricing of contingent-convertible catastrophe bonds:

- Since the contingent-convertible catastrophe bond has attracted little scholarly attention to date, further research into it is clearly called for. It would be interesting to look into other types of contingent-convertible catastrophe bonds, such as ones based on parametric triggers, and to study their price behaviour. It would also be interesting to consider other design features such as converting the recovery upon trigger into the equity of an entity other than the issuer.

- We also point out the following interesting, potential application of the work by Braun (2011) in the context of valuing contingent-convertible catastrophe bonds. Suppose that we only modelled an index consisting of losses from earthquakes and/or hurricanes[3]. If we did not assume that for the catastrophe-risk variables the real-world and risk-neutral probability measures coincided[4], and rather used a risk-neutral measure for the catastrophe-risk variables, then we could use Braun's implied intensities backed out from catastrophe swap transactions. In doing this, we could possibly ensure that the contingent- convertible catastrophe bond is consistently priced with other instruments in the catastrophe risk markets. Moreover, note that these backed-out intensities could also be used for catastrophe bond pricing in the context of our model in Chapter 3.

[3] With a more detailed breakdown of PCS data, it is indeed possible to extract this information.
[4] That is, if we did not use Assumption 3.11 but made the slightly weaker assumption that a random variable retains its distributional characteristics when moving from the real-world to the risk-neutral probability measure.

Glossary of acronyms

ALP	Aggregate loss process
CAT	Catastrophe
CAT-E-put	Catastrophe-equity put
CCD approach	Conditional complete-data approach
CP	Coupon-paying
CPI	Consumer Price Index
Coco	Contingent convertible (bond)
CocoCat	Contingent convertible catastrophe bond
EDF	Empirical distribution function
FSRLP	First-order single risk loss approximation
GAAP	Generally Accepted Accounting Principles
GEV	Generalised Extreme Value (distribution)
GP	Generalised Pareto (distribution)
HP	(Time-)homogeneous Poisson (process)
IAS	International Accounting Standards
i.i.d.	Independent and identically-distributed
IFRS	International Financial Reporting Standards
IHP	(Time-)inhomogeneous Poisson (process)
IL	Index-linked
ILS	Insurance-linked security
LDA	Loss distribution approach
LIBOR	London Interbank Offered Rate
MAE	Mean absolute error
MC	Monte Carlo
MGEV	Modified Generalised Extreme Value (distribution)
MLE	Maximum likelihood estimation
MPS	Maximum product of spacings
RMSE	Root mean-square error
OTC	Over-the-counter
PCS	Property Claims Services
SPV	Special purpose vehicle
ZC	Zero-coupon

Bibliography

Aase, K. (1999). An equilibrium model of catastrophe insurance futures and spreads, *The Geneva Papers on Risk and Insurance Theory* **24**(1): 69–96.

Aase, K. K. (2001). A Markov model for the pricing of catastrophe insurance futures and spreads, *Journal of Risk and Insurance* **68**(1): 25–49.

Anderson, T. W. and Darling, D. A. (1954). A test of goodness of fit, *Journal of the American Statistical Association* **49**(268): 765–769.

Aslan, Z., Damnjanovic, I. and Mander, J. (2011). Engineering the finance of earthquake risk, *Proceedings of the Ninth Pacific Conference on Earthquake Engineering*.

Asmussen, S., Binswanger, K. and Højgaard, B. (2000). Rare events simulation for heavy-tailed distributions, *Bernoulli* **6**(2): 303–322.

Asmussen, S. and Glynn, P. W. (2007). *Stochastic simulation: algorithms and analysis*, Vol. 57, Springer, Berlin.

Asmussen, S. and Kroese, D. P. (2006). Improved algorithms for rare event simulation with heavy tails, *Advances in Applied Probability* **38**(2): 545–558.

Bakshi, G. and Madan, D. (2002). Average rate claims with emphasis on catastrophe loss options, *Journal of Financial and Quantitative Analysis* **37**(1): 93–115.

Barrieu, P. and Albertini, L. (2010). *The handbook of insurance-linked securities*, Vol. 525, John Wiley & Sons.

Barton, C. and Nishenko, S. (1997). Natural disaster: Forecasting economic and life losses, USGS Fact Sheet-US Geological Survey, *Marine and Coastal Geology Program, St. Petersburg, Florida* .

Baryshnikov, Y., Mayo, A. and Taylor, D. R. (1998). Pricing of CAT bonds. Working paper.
URL: *http://citeseerx.ist.psu.edu*

Beard, R. (2013). *Risk theory: the stochastic basis of insurance*, Vol. 20, Springer, Berlin.

Beirlant, J., Alves, I. F. and Reynkens, T. (2017). Fitting tails affected by truncation, *Electronic Journal of Statistics* **11**(1): 2026–2065.

Beirlant, J., Goegebeur, Y., Segers, J. and Teugels, J. L. (2006). *Statistics of extremes: theory and applications*, John Wiley and Sons, New York.

Besson, J.-L., Dacorogna, M. M., De Martin, P., Kastenholz, M. and Moller, M. (2009). How much capital does a reinsurance need?, *The Geneva Papers on Risk and Insurance. Issues and Practice* **34**(2): 159–174.

Biagini, F., Bregman, Y. and Meyer-Brandis, T. (2008a). Pricing of catastrophe insurance options under immediate loss reestimation, *Journal of Applied Probability* **45**(3): 831–845.

Biagini, F., Bregman, Y. and Meyer-Brandis, T. (2008b). Pricing of catastrophe insurance options written on a loss index with reestimation, *Insurance: Mathematics and Economics* **43**(2): 214–222.

Billingsley, P. (1995). *Convergence of probability measures*, John Wiley & Sons.

Bishop, D., Liu, E. Z., Murray, P. and Solomonia, T. (2009). Contingent Convertible Bonds: Comments on the December 2009 Consultative Document-Strengthening the Resilience of the Banking Sector, *Archived Advisory Document – Comments to the Bank of International Settlements* .

Björk, T. (2009). *Arbitrage theory in continuous time*, Oxford University Press, Oxford.

BNY Mellon (2013). The Disaster Gap. The Bank of New York Mellon Corporation.
URL: *www.artemis.bm/downloads*

Bodoff, N. M. and Gan, Y. (2009). An Analysis of the Market Price of CAT Bonds. Casualty Actuarial Society E-Forum.
URL: *https://www.casact.org/pubs/forum/09spforum/02Bodoff.pdf*

Bouchaud, J. P. (2001). Power laws in economics and finance: some ideas from physics, *Quantitative Finance* **1**: 105–112.

Bowers, N. H., Gerber, H., Hickman, J., Jones, D. and Nesbitt, C. (1997). *Actuarial mathematics*, Schaumberg, IL: Society of Actuaries.

Braun, A. (2011). Pricing catastrophe swaps: A contingent claims approach, *Insurance: Mathematics and Economics* **49**(3): 520–536.

Braun, A. (2016). Pricing in the primary market for CAT bonds: new empirical evidence, *Journal of Risk and Insurance* **83**(4): 811–847.

Braun, A. (2017). Evolution or Revolution How Solvency II will Change the Balance between Reinsurance and ILS, *Journal of Insurance Regulation [Forthcoming]* .

Brigo, D. and Mercurio, F. (2007). *Interest rate models-theory and practice: with smile, inflation and credit*, Springer, Berlin.

Bühlmann, H. (1984). The general economic premium principle, *ASTIN Bulletin* **14**(1): 13–21.

Bühlmann, H. (1985). Premium calculation from top down, *ASTIN Bulletin* **15**(2): 89–101.

Burnecki, K. (2000). Self-similar processes as weak limits of a risk reserve process, *Probability and Mathematical Statistics* **20**(2): 261–272.

Burnecki, K., Gajda, J. and Sikora, G. (2011). Stability and lack of memory of the returns of the Hang Seng index, *Physica A* **390**(18): 3136–3146.

Burnecki, K. and Giuricich, M. N. (2017). Stable weak approximation at work in index-linked catastrophe bond pricing, *Risks* **5**(4): 64.

Burnecki, K. and Kukla, G. (2003). Pricing of zero-coupon and coupon CAT bonds, *ACTA Applied Mathematics (Warsaw)* **30**(3): 315–324.

Burnecki, K., Kukla, G. and Taylor, D. (2005). Pricing of catastrophe bonds, *Čižek, P., Härdle, W and Weron, R (Eds): Statistical tools for finance and insurance*, Springer, Berlin, pp. 93–114.

Burnecki, K., Kukla, G. and Weron, R. (2000). Property insurance loss distributions, *Physica A* **287**(1): 269–278.

Burnecki, K., Misiorek, A. and Weron, R. (2005). Loss distributions, *Čižek, P., Härdle, W and Weron, R (Eds): Statistical tools for finance and insurance*, Springer, Berlin, pp. 289–317.

Burnecki, K., Nowicka-Zagrajek, J. and Wyłomańska, A. (2005). Pure risk premiums under deductibles., *Čižek, P., Härdle, W and Weron, R (Eds): Statistical tools for finance and insurance*, Springer, Berlin, pp. 427–452.

Burnecki, K. and Weron, R. (2005). Modeling of the risk process., *Čižek, P., Härdle, W and Weron, R (Eds): Statistical tools for finance and insurance*, Springer, Berlin, pp. 319–337.

Burnecki, K., Wylomanska, A. and Chechkin, A. (2015). Discriminating between light-and heavy-tailed distributions with limit theorem, *PloS one* **10**(12): e0145604.

Cain, K. C., Harlow, S. D., Little, R. J., Nan, B., Yosef, M., Taffe, J. R. and Elliott, M. R. (2011). Bias due to left truncation and left censoring in longitudinal studies of developmental and disease processes, *American Journal of Epidemiology* **173**(9): 1078–1084.

Calderín-Ojeda, E., Fergusson, K. and Wu, X. (2017). An EM Algorithm for Double-Pareto-Lognormal Generalized Linear Model Applied to Heavy-Tailed Insurance Claims, *Risks* **5**(4): 60.

Canabarro, E., Finkemeier, M., Anderson, R. R. and Bendimerad, F. (2000). Analyzing insurance-linked securities, *The Journal of Risk Finance* **1**(2): 49–75.

Capasso, M., Alessi, L., Barigozzi, M. and Fagiolo, G. (2009). On approximating the distributions of goodness-of-fit test statistics based on the empirical distribution function: The case of unknown parameters, *Advances in Complex Systems* **12**(02): 157–167.

Carayannopoulos, P. and Perez, M. F. (2015a). Diversification through catastrophe bonds: lessons from the subprime financial crisis, *The Geneva Papers on Risk and Insurance-Issues and Practice* **40**(1): 1–28.

Carayannopoulos, P. and Perez, M. F. (2015b). Diversification through catastrophe bonds: lessons from the subprime financial crisis, *The Geneva Papers on Risk and Insurance-Issues and Practice* **40**(1): 1–28.

Chang, C.-C., Lin, S.-K. and Yu, M.-T. (2011). Valuation of Catastrophe Equity Puts With Markov-Modulated Poisson Processes, *Journal of Risk and Insurance* **78**(2): 447–473.

Chang, C. W. and Chang, J. S. (2017). An Integrated Approach to Pricing Catastrophe Reinsurance, *Risks* **5**(3): 51.

Chang, C. W., Chang, J. S. and Lu, W. (2010). Pricing catastrophe options with stochastic claim arrival intensity in claim time, *Journal of Banking & Finance* **34**(1): 24–32.

Chang, C. W., Chang, J. S. and Yu, M.-T. (1996). Pricing catastrophe insurance futures call spreads: a randomized operational time approach, *Journal of Risk and Insurance* pp. 599–617.

Chaubey, Y. P., Garrido, J. and Trudeau, S. (1998). On the computation of aggregate claims distributions: some new approximations, *Insurance: Mathematics and Economics* **23**(3): 215–230.

Chavez-Demoulin, V., Embrechts, P. and Nešlehová, J. (2006). Quantitative models for operational risk: extremes, dependence and aggregation, *Journal of Banking and Finance* **30**(10): 2635–2658.

Cheng, R. C. H. and Amin, N. A. K. (1983). Estimating parameters in continuous univariate distributions with a shifted origin, *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 394–403.

Cheng, R. and Stephens, M. (1989). A goodness-of-fit test using Moran's statistic with estimated parameters, *Biometrika* **76**(2): 385–392.

Chernobai, A., Burnecki, K., Rachev, S., Trück, S. and Weron, R. (2006). Modelling catastrophe claims with left-truncated severity distributions, *Computational Statistics* **21**(3-4): 537–555.

Chernobai, A., Rachev, S. T. and Fabozzi, F. J. (2015). Composite goodness-of-fit tests for left-truncated loss samples, *Handbook of financial econometrics and statistics*, Springer, pp. 575–596.

Chicago Board of Trade (1995). *PCS catastrophe insurance options: A user's guide*, CBOT, Chicago.

Christensen, C. V. and Schmidli, H. (2000a). Pricing catastrophe insurance products based on actually reported claims, *Insurance: Mathematics and Economics* **27**(2): 189–200.

Christensen, C. V. and Schmidli, H. (2000b). Pricing catastrophe insurance products based on actually reported claims, *Insurance: Mathematics and Economics* **27**(2): 189–200.

Coronel-Brizio, H. F. and Hernandez-Montoya, A. R. (2005). On fitting the Pareto–Levy distribution to stock market index data: Selecting a suitable cutoff value, *Physica A: Statistical Mechanics and its Applications* **354**: 437–449.

Cox, J. C., Ingersoll Jr, J. E. and Ross, S. A. (1985). A theory of the term structure of interest rates, *Econometrica* pp. 385–407.

Cox, S. H., Fairchild, J. R. and Pedersen, H. W. (2000). Actuarial and economic aspects of securitization of risk, *ASTIN Bulletin* **20**(1): 157–193.

Cox, S. H., Fairchild, J. R. and Pedersen, H. W. (2004a). Valuation of structured risk management products, *Insurance: Mathematics and Economics* **34**(2): 259–272.

Cox, S. H., Fairchild, J. R. and Pedersen, H. W. (2004b). Valuation of structured risk management products, *Insurance: Mathematics and Economics* **34**(2): 259–272.

Cox, S. H. and Pedersen, H. W. (2000). Catastrophe risk bonds, *North American Actuarial Journal* **4**(4): 56–82.

Cummins, J. D. (2008a). CAT Bonds and Other Risk-Linked Securities: State of the Market and Recent Developments, *Risk Management and Insurance Review* **11**(1): 23–47.

Cummins, J. D. (2008b). The Bermuda insurance market: an economic analysis, *The Bermuda Insurance Association, Hamilton, Bermuda.* .

Cummins, J. D., Doherty, N. and Lo, A. (2002). Can insurers pay for the "big one"? Measuring the capacity of the insurance market to respond to catastrophic losses, *Journal of Banking & Finance* **26**(2): 557–583.

Cummins, J. D. and Geman, H. (1995a). Pricing catastrophe insurance futures and call spreads: An arbitrage approach, *Journal of Fixed Income* **4**(4): 46–57.

Cummins, J. D. and Geman, H. (1995b). Pricing catastrophe insurance futures and call spreads: An arbitrage approach, *The Journal of Fixed Income* **4**(4): 46–57.

Cummins, J. D. and Weiss, M. A. (2009). Convergence of insurance and financial markets: Hybrid and securitized risk-transfer solutions, *Journal of Risk and Insurance.* **76**(3): 493–545.

Daley, D., Omey, E. and Vesilo, R. (2007). The tail behaviour of a random sum of subexponential random variables and vectors, *Extremes* **10**(1-2): 21–39.

Darling, D. A. (1957). The Kolmogorov-Smirnov, Cramer-von Mises tests, *The Annals of Mathematical Statistics* **28**(4): 823–838.

Dassios, A. and Jang, J. (2003a). Pricing of catastrophe reinsurance and derivatives using the Cox process with shot noise intensity, *Finance and Stochastics* **7**(1): 73–95.

Dassios, A. and Jang, J.-W. (2003b). Pricing of catastrophe reinsurance and derivatives using the Cox process with shot noise intensity, *Finance and Stochastics* **7**(1): 73–95.

Davis, M. H. A. (1993). *Markov models and optimization*, Chapman & Hall, London.

De Martino, G., Libertucci, M., Marangoni, M. and Quagliariello, M. (2010). Countercyclical contingent capital (ccc): possible use and ideal design, *Working Paper, Bank of Italy: Banking and Finance Supervision Area.* .

Delbaen, F. and Haezendonck, J. (1989). A martingale approach to premium calculation principles in an arbitrage free market, *Insurance: Mathematics and Economics* **8**(4): 269–277.

Dieckmann, S. (2010). By force of nature: explaining the yield spread on catastrophe bonds, *Available at SSRN 1082879* .

Doherty, N. A. (1997). Innovations in managing catastrophe risk, *Journal of Risk and Insurance* pp. 713–718.

Duan, J.-C. and Simonato, J.-G. (2002). Maximum likelihood estimation of deposit insurance value with interest rate risk, *Journal of Empirical Finance* **9**(1): 109–132.

Dufour, R. and Maag, U. R. (1978). Distribution results for modified Kolmogorov-Smirnov statistics for truncated or censored, *Technometrics* **20**(1): 29–32.

Durbin, D. (2001). Managing natural catastrophe risks: The structure and dynamics of reinsurance, *The Geneva Papers on Risk and Insurance. Issues and Practice* **26**(2): 297–309.

Egami, M. and Young, V. R. (2008). Indifference prices of structured catastrophe (CAT) bonds, *Insurance: Mathematics and Economics* **42**(2): 771–778.

Embrechts, P. (2000). Actuarial versus financial pricing of insurance, *The Journal of Risk Finance* **1**(4): 17–26.

Embrechts, P., Klüppelberg, C. and Mikosch, T. (2013). *Modelling extremal events: for insurance and finance*, Vol. 33, Springer, Berlin.

Embrechts, P. and Meister, S. (1997). Pricing insurance derivatives, the case of CAT-futures, *Proceedings of the 1995 Bowles Symposium on Securitization of Risk, Georgia State University Atlanta, Society of Actuaries, Monograph M-FI97-1*, pp. 15–26.

Ergashev, B., Pavlikov, K., Uryasev, S. and Sekeris, E. (2016). Estimation of truncated data samples in operational risk modeling, *Journal of Risk and Insurance.* **83**(3): 613–640.

Esscher, F. (1932). On the probability function in the collective theory of risk, *Scandinavian Actuarial Journal* **1932**(3).

Filipovic, D. (2009). *Term-Structure Models. A Graduate Course.*, Springer, Berlin.

Flannery, M. J. (2016). Stabilizing large financial institutions with contingent capital certificates, *Quarterly Journal of Finance* **6**(02): 1650006.

Föllmer, H. and Schweizer, M. (1986). Hedging of contingent claims, *M. H. A. Davis and R. J. Elliott (Eds), "Applied Stochastic Analysis", Stochastics Monographs* **5**: 389.

Föllmer, H. and Schweizer, M. (1989). Hedging by sequential regression: An introduction to the mathematics of option trading, *ASTIN Bulletin* **19**(S1): 29–42.

Froot, K. A. (2001). The market for catastrophe risk: a clinical examination, *Journal of Financial Economics* **60**(2): 529–571.

Froot, K. A. (2008). The intermediation of financial risks: Evolution in the catastrophe reinsurance market, *Risk Management and Insurance Review* **11**(2): 281–294.

Froot, K. A., Kim, M. and Rogoff, K. (1995). The law of one price over 700 years, *Technical report*, National Bureau of Economic Research.

Furrer, H., Michna, Z. and Weron, A. (1997). Stable Lévy motion approximation in collective risk theory, *Insurance: Mathematics and Economics* **20**(2): 97–114.

Galeotti, M., Gürtler, M. and Winkelvos, C. (2013). Accuracy of premium calculation models for CAT bonds – an empirical analysis, *Journal of Risk and Insurance* **80**(2): 401–421.

Gastaldi, T. and Gastaldi, T. (1992). A Kolmogorov-Smirnov test procedure involving a possibly censored or truncated sample, *Communications in Statistical Theory* **22**(1): 31–39.

Gatzert, N., Pokutta, S. and Vogl, N. (2014). Convergence of Capital and Insurance Markets: Consistent Pricing of Index-Linked Catastrophic Loss Instruments, *Technical report*, Friendrich-Alexander University Erlangen-Nurnberg (FAU).

Geman, H. and Yor, M. (1997). Stochastic time changes in catastrophe option pricing, *Insurance: Mathematics and economics* **21**(3): 185–193.

Georgiopoulos, N. (2016). The valuation of contingent convertible catastrophe debt under simple solvency and liquidity covenants, *Journal of Risk (Forthcoming)* **18**(6).

Gerber, H. U. and Shiu, E. S. (1996). Actuarial bridges to dynamic hedging and option pricing, *Insurance: Mathematics and Economics* **18**(3): 183–218.

Glasserman, P. (2013). *Monte Carlo methods in financial engineering*, Vol. 53, Springer, New York.

Goldstein, M. L., Morris, S. A. and Yen, G. G. (2004). Problems with fitting to the Power-Law distribution, *The European Physical Journal B - Condensed Matter and Complex Systems* **41**(2): 255–258.

Goovaerts, M., De Vylder, F. and Haezendonck, J. (1984). Insurance premiums, *North-Holland, Amsterdam* **1**.

Grandell, J. (2012). *Aspects of risk theory*, Springer, Berlin.

Guilbaud, O. (1988). Exact Kolmogorov-type tests for left-truncated and/or right-censored data, *Journal of the Americal Statistical Association* **83**(401): 213–221.

Gürtler, M., Hibbeln, M. and Winkelvos, C. (2016). The impact of the financial crisis and natural catastrophes on CAT bonds, *Journal of Risk and Insurance* **83**(3): 579–612.

Hagendorff, B., Hagendorff, J. and Keasey, K. (2015). The Impact of Mega-Catastrophes on Insurers: An Exposure-Based Analysis of the US Homeowners Insurance Market, *Risk Analysis* **35**(1): 157–173.

Hammer, V. and Singer, A. (2001). Insurance derivatives: a tax angle, *Taxation of Financial Products* **2**: 29.

Härdle, W. K. and Cabrera, B. L. (2010). Calibrating CAT bonds for Mexican earthquakes, *Journal of Risk and Insurance.* **77**(3): 625–650.

Harrington, S. E. and Niehaus, G. (2001). Government insurance, tax policy, and the affordability and availability of catastrophe insurance, *Journal of Insurance Regulation* **19**(4): 591–612.

Harrington, S. E. and Niehaus, G. (2003). Capital, corporate income taxes, and catastrophe insurance, *Journal of Financial Intermediation* **12**(4): 365–389.

Harrison, J. M. and Pliska, S. R. (1981). Martingales and stochastic integrals in the theory of continuous trading, *Stochastic processes and their applications* **11**(3): 215–260.

Haslip, G. G. and Kaishev, V. K. (2010). Pricing of reinsurance contracts in the presence of catastrophe bonds, *Astin Bulletin* **40**(01): 307–329.

Hoyt, R. E. and McCullough, K. A. (1999). Catastrophe insurance options: Are they zero-beta assets?, *Journal of Insurance Issues* pp. 147–163.

Hull, J. C. and White, A. (2013). LIBOR vs. OIS: The derivatives discounting dilemma, *Working paper. Available at SSRN 2211800* .

Hull, J. and White, A. (1990). Pricing interest-rate-derivative securities, *The Review of Financial Studies* **3**(4): 573–592.

Jacod, J. and Protter, P. E. (2003). *Probability essentials*, Springer, New York.

Jaffee, D. M. and Russell, T. (1997). Catastrophe insurance, capital markets, and uninsurable risks, *Journal of Risk and Insurance* pp. 205–230.

Jaimungal, S. and Chong, Y. (2014). Valuing clustering in catastrophe derivatives, *Quantitative Finance* **14**(2): 259–270.

Jaimungal, S. and Wang, T. (2006a). Catastrophe options with stochastic interest rates and compound Poisson losses, *Insurance: Mathematics and Economics* **38**(3): 469–483.

Jaimungal, S. and Wang, T. (2006b). Catastrophe options with stochastic interest rates and compound Poisson losses, *Insurance: Mathematics and Economics* **38**(3): 469–483.

Janicki, A. and Weron, A. (1994). *Simulation and Chaotic Behaviour of Stable Processes*, Marcel Dekker, New York.

Jarrow, R. A. (2010). A simple robust model for CAT bond valuation, *Finance Research Letters* **7**(2): 72–79.

Johnson, C., Omar, Y. and Ouwehand, P. (2004). Valuing risky income streams in incomplete markets, *Applied Mathematical Finance* **11**(3): 227–258.

Juárez, S. F. and Schucany, W. R. (2004). Robust and efficient estimation for the generalized Pareto distribution, *Extremes* **7**(3): 237–251.

Karatzas, I. and Shreve, S. (2012). *Brownian motion and stochastic calculus*, Vol. 113, Springer, Berlin.

Kerney, G. (2013). Everything You Need to Know about the PCS Catastrophe Loss Index.
URL: *http://www. iso. com/. access*

Kohatsu-Higa, A. and Tankov, P. (2010). Jump-adapted discretization schemes for Lévy-driven SDEs, *Stochastic Processes and their Applications* **120**(11): 2258–2285.

Koutrouvelis, I. A. (1980). Regression-type estimation of the parameters of stable laws, *Journal of the American statistical association* **75**(372): 918–928.

Kunreuther, H. (2006). Disaster mitigation and insurance: Learning from Katrina, *The Annals of the American Academy of Political and Social Science* **604**(1): 208–227.

Labuschagne, C. C. and Offwood, T. M. (2010). A note on the connection between the esscher–girsanov transform and the wang transform, *Insurance: Mathematics and Economics* **47**(3): 385–390.

Lane, M. (2004). The viability and likely pricing of CAT bonds for developing countries, *Catastrophe Risk and Reinsurance: A Country Risk Management Perspective. Risk Books, London* pp. 239–268.

Lane, M. N. (2000). Pricing risk transfer transactions1, *ASTIN Bulletin* **30**(2): 259–293.

Lane, M. N. (2003). Rationale and results with the LFC CAT bond pricing model, *Lane Financial Ltd., Chicago* .

Lane, M. N. and Mahul, O. (2008). Catastrophe risk pricing: an empirical analysis. Working paper., *The World Bank* .

Lee, J. P. and Yu, M. T. (2002). Pricing default-risky CAT bonds with moral hazard and basis risk, *Journal of Risk and Insurance.* pp. 25–44.

Lee, J. P. and Yu, M. T. (2007). Valuation of catastrophe reinsurance with catastrophe bonds, *Insurance: Mathematics and Economics* **41**(2): 264–278.

Levi, C. and Partrat, C. (1991). Statistical analysis of natural events in the united states, *ASTIN Bulletin* **21**(2): 253–276.

Lin, S.-K. and Chang, C.-C. (2008). Pricing catastrophe insurance products in Markov jump diffusion models, *Journal of Financial Studies* **2**(16): 1–33.

Lin, X. S. and Wang, T. (2009a). Pricing perpetual american catastrophe put options: A penalty function approach, *Insurance: Mathematics and Economics* **44**(2): 287–295.

Lin, X. S. and Wang, T. (2009b). Pricing perpetual American catastrophe put options: A penalty function approach, *Insurance: Mathematics and Economics* **44**(2): 287–295.

Lindvall, T. (1973). Weak convergence of probability measures and random functions in the function space d [0,8), *Journal of Applied Probability* pp. 109–121.

Litzenberger, R. H., Beaglehole, D. R. and Reynolds, C. E. (1996). Assessing catastrophe reinsurance-linked securities as a new asset class, *Journal of Portfolio Management* **23**: 76–86.

Liu, J., Xiao, J., Yan, L. and Wen, F. (2014). Valuing catastrophe bonds involving credit risks, *Mathematical Problems in Engineering* **2014**.

Liu, S., Wu, H. and William, Q. (2015). Understanding and addressing the unbounded "likelihood" problem, *American Statistician* **69**(3): 191–200.

Liu, Y., Cizeau, P., Meyer, M., Peng, C. K. and Stanley, H. E. (1997). Correlations in economic time series, *Physica A* **245**(3-4): 437–440.

Lo, C.-L., Lee, J.-P. and Yu, M.-T. (2013). Valuation of insurers contingent capital with counterparty risk and price endogeneity, *Journal of Banking & Finance* **37**(12): 5025–5035.

Longstaff, F. A. (1989). A nonlinear general equilibrium model of the term structure of interest rates, *Journal of Financial Economics* **23**(2): 195–224.

Loubergé, H., Kellezi, E. and Gilli, M. (1999). Using catastrophe-linked securities to diversify insurance risk: A financial analysis of CAT bonds, *Journal of Insurance Issues* pp. 125–146.

Luo, X., Shevchenko, P. V. and Donnelly, J. B. (2009). Addressing the impact of data truncation and parameter uncertainty on operational risk estimates, *Journal of Operational Risk* **4**(2): 3–26.

Ma, Z. G. and Ma, C. Q. (2013). Pricing catastrophe risk bonds: A mixed approximation method, *Insurance: Mathematics and Economics* **52**(2): 243–254.

Mantegna, R. N. and Stanley, H. E. (1999). *Introduction to econophysics: correlations and complexity in finance*, Cambridge University Press.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* **3**(1): 125–144.

Michna, Z. (2005). On approximations of risk process with renewal arrivals in alpha-stable domain, *Probability and Mathematical Statistics* **25**(1): 173.

Mignola, G. and Ugoccioni, R. (2006). Sources of uncertainty in modeling operational risk losses, *Journal of Operational Risk* **1**(2): 33–50.

Milidonis, A. and Grace, M. F. (2008). Tax-deductible pre-event catastrophe loss reserves: The case of florida, *ASTIN Bulletin* **38**(1): 13–51.

Möller, M. (1996). Pricing PCS-options with the use of Esscher-transforms, *Aktuarielle Ansätze für Finanzrisiken, AFIR* pp. 1299–1310.

Møller, T. (2001). Risk-minimizing hedging strategies for insurance payment processes, *Finance and Stochastics* **5**(4): 419–446.

Muermann, A. (2001). Pricing catastrophe insurance derivatives, *Financial Markets Group Discussion Paper* (No. 400).

Muermann, A. (2003). Actuarially consistent valuation of catastrophe derivatives, *The Wharton Financial Institution Center Working Paper Series* pp. 03–18.

Muermann, A. (2008). Market price of insurance risk implied by catastrophe derivatives, *North American Actuarial Journal* **12**(3): 221–227.

Müller, A. and Grandi, M. (2000). Weather derivatives: a risk management tool for weather-sensitive industries, *Geneva Papers on Risk and Insurance. Issues and Practice* pp. 273–287.

Nguyen, T. and Lindenmeier, J. (2014). Catastrophe risks, CAT bonds and innovation resistance, *Qualitative Research in Financial Markets* **6**(1): 75–92.

Nolan, J. P. (1997). Numerical calculation of stable densities and distribution functions, *Communications in statistics. Stochastic models* **13**(4): 759–774.

Nolan, J. P. (2015). *Stable Distributions - Models for Heavy Tailed Data*, Birkhauser, Boston. In progress, Chapter 1 online at *www.academic2.american.edu/∼jpnolan*.

Nowak, P. and Romaniuk, M. (2013). Pricing and simulations of catastrophe bonds, *Insurance: Mathematics and Economics* **52**(1): 18–28.

O'Brien, T. (1997). Hedging strategies using catastrophe insurance options, *Insurance: Mathematics and Economics* **21**(2): 153–162.

Palmowski, Z. and Rolski, T. (2002). A technique for exponential change of measure for Markov processes, *Bernoulli* **8**(6): 767–785.

Panjer, H. and Willmot, G. (1992). *Insurance risk models*, Society of Actuaries, Schaumburg.

Papachristou, D. (2009). Statistical Analysis of the Spreads of Catastrophe Bonds at the Time of Issue, *Working Paper, presented at the 39th ASTIN Colloquium.* .

Pazarbasioglu, C., Zhou, J., Leslé, V. L. and Moor, M. (2011). Contingent capital: Economic rationale and design features. International Monetary Fund: Monetary and Capital Markets Department, Staff Discussion Note.

Pelsser, A. (2008). On the applicability of the wang transform for pricing financial risks, *ASTIN Bulletin* **38**(1): 171–181.

Pentikäinen, T. (1977). On the approximation of the total amount of claims, *ASTIN Bulletin* **9**(03): 281–289.

Peters, G. W. and Shevchenko, P. V. (2015). *Advances in Heavy Tailed Risk Modeling: A Handbook of Operational Risk*, John Wiley & Sons, New York.

Platen, E. and Heath, D. (2006). *A benchmark approach to quantitative finance*, Springer, Berlin.

Plerou, V., Gopikrishnan, P., Rosenow, B., Amaral, L. A. N. and Stanley, H. E. (2000). Econophysics: financial time series from a statistical physics point of view, *Physica A* **279**(1): 443–456.

Ranneby, B. (1984). The maximum spacing method. An estimation method related to the maximum likelihood method, *Scandinavian Journal of Statistics* pp. 93–112.

Reshetar, G. (2008). Pricing of multiple-event coupon paying CAT bond, *Available at SSRN 1059021* .

Ross, S. (2002). *Simulation*, Academic Press. San Diego, USA.

Rüdlinger, M. (2015). *Contingent Convertible Bonds: An Empirical Analysis of Drivers and Announcement Effect*, PhD thesis, University of St. Gallen.

Samorodnitsky, G. and Taqqu, M. S. (1994). *Stable Non-Gaussian Random Processes*, Chapman & Hall, London.

Schradin, H. R. (1996). PCS catastrophe insurance options - a new instrument for managing catastrophe risk, *Contribution to the 6th AFIR International Colloquium*.

Schweizer, M. (1991). Option hedging for semimartingales, *Stochastic processes and their Applications* **37**(2): 339–363.

Schweizer, M. (1999). A guided tour through quadratic hedging approaches, *Technical report*, Discussion Papers, Interdisciplinary Research Project 373: Quantification and Simulation of Economic Processes.

Seal, H. L. (1977). Approximations to risk theory's $F(x,t)$ by means of the gamma distribution, *ASTIN Bulletin* **9**(1-2): 213–218.

Seri, R. and Choirat, C. (2015). Comparison of approximations for compound poisson processes, *ASTIN Bulletin* **45**(03): 601–637.

Shao, J. (2015). *Modelling catastrophe risk bonds*, PhD thesis, University of Liverpool.

Shevchenko, P. V. and Temnov, G. (2009). Modeling operational risk data reported above a time-varying threshold, *arXiv preprint, arXiv:0904.4075* .

Skorokhod, A. V. (1957). Limit theorems for stochastic processes with independent increments, *Theory of Probability & Its Applications* **2**(2): 138–171.

Soukissian, T. and Tsalis, C. (2015). The effect of the generalized extreme value distribution parameter estimation methods in extreme wind speed prediction, *Natural Hazards* **78**(3): 1777–1809.

Spiegeleer, J. D. and Schoutens, W. (2012). Pricing contingent convertibles: A derivatives approach, *The Journal of Derivatives* **20**(2): 27–36.

Stam, A. (1973). Regular variation of the tail of a subordinated probability distribution, *Advances in Applied Probability* pp. 308–327.

Stanley, H. E., Gabaix, X., Gopikrishnan, P. and Plerou, V. (2007). Economic fluctuations and statistical physics: Quantifying extremely rare and less rare events in finance, *Physica A* **382**(1): 286–301.

Stephens, M. A. (1974). EDF statistics for goodness of fit and some comparisons, *Journal of the American statistical Association* **69**(347): 730–737.

Sundt, B. (1982). Asymptotic behaviour of compound distributions and stop-loss premiums, *ASTIN Bulletin* **13**(02): 89–98.

Swiss Re (2008). Natural Catastrophes and Man-Made Disasters in 2007: High Losses in Europe. Sigma No. 1/2008.
URL: *http://www.swissre.com/library*

Swiss Re (2012). What are Insurance Linked Securities (ILS), and why should they be considered?
URL: *http://www.casact.org/community/affiliates/CANE/0912/Cat-Bond.pdf*

Swiss Re (2014). Insurance linked securities market update. Swiss Re Capital Markets.
URL: *http://www.swissre.com/reinsurance/insurers/ils*

SwissRe (2009). The role of indices in transferring insurance risk to the capital markets. Sigma No. 4/2009.
URL: *http://www.swissre.com/library*

SwissRe (2011). The Fundamentals of Insurance-linked Securities - Transforming Insurance Risk into transparent and tradeable Capital Market Products.
URL: *http://www.swissre.com/library*

SwissRe (2013). Annual Report: Contingent capital instruments.
URL: *http://reports.swissre.com/2013/financial-report/financial-statements/notes-to-the-group-financial-statements/6-debt-and-contingent-capital-instruments.html*

SwissRe (2016). Natural catastrophes and man-made disasters in 2015: Asia suffers substantial losses. Sigma 01/2016.
URL: *http://www.swissre.com/library*

Tao, Z., Tao, X. and Li, P. (2009). Pricing Model for Earthquake CAT Bonds, *Business Intelligence and Financial Engineering, 2009. BIFE'09. International Conference on*, IEEE, pp. 740–744.

Têtu, A., Lai, V. S., Soumaré, I. and Gendron, M. (2015). Hedging Flood Losses Using CAT Bonds, *Asia Pacific Journal of Risk and Insurance* **9**(2): 149–184.

Valdez, E. A. (2005). Probability transforms with elliptical generators, *The 36th International ASTIN Colloquium, Zurich, Switzerland*.

Vasicek, O. (1977). An equilibrium characterization of the term structure, *Journal of Financial Economics* **5**(2): 177–188.

Vaugirard, V. E. (2003a). Pricing catastrophe bonds by an arbitrage approach, *The Quarterly Review of Economics and Finance* **43**(1): 119–132.

Vaugirard, V. E. (2003b). Valuing catastrophe bonds by Monte Carlo simulations, *Applied Mathematical Finance* **10**(1): 75–90.

Vaugirard, V. E. (2004). A canonical first passage time model to pricing nature-linked bonds, *Economics Bulletin* **7**(2): 1–7.

Von Neumann, J. (1951). Various techniques used in connection with random digits, *Applied Mathematical Series. National Bureau of Standards, Washington DC.* **12**: 36–38.

Wang, S. S. (2000). A class of distortion operators for pricing financial and insurance risks, *Journal of Risk and Insurance* pp. 15–36.

Wang, S. S. (2004). CAT bond pricing using probability transforms, *Geneva Papers* **278**: 19–29.

Wang, X. (2016). Catastrophe equity put options with target variance, *Insurance: Mathematics and Economics* **71**: 79–86.

Willis (2015). ILS Market Update Q4 2014. Willis Capital Markets & Advisory.
URL: *http://www.willis.com/Documents/Publications/Services/WCMA*

Willis (2017). ILS Market Update: Wake-up call. New York, USA.
URL: *http://www.willis.com/Documents/Publications/Services/WCMA*

Wong, T. S. T. and Li, W. K. (2006). A note on the estimation of extreme value distributions using maximum product of spacings, *Time Series and Related Topics*, Institute of Mathematical Statistics, pp. 272–283.

World Economic Forum (2008). Convergence of insurance and capital markets.
URL: *http://www3.weforum.org/docs*

Zeckhauser, R. (1995). Insurance and catastrophes, *The Geneva Papers on Risk and Insurance Theory* **20**(2): 157–175.

Zimbidis, A. A., Frangos, N. E. and Pantelous, A. A. (2007). Modeling earthquake risk via extreme value theory and pricing the respective catastrophe bonds, *ASTIN Bulletin* **37**(1): 163–183.