Essays in Tail Risk and Asset Pricing in Credit Markets

Reinhard Fellmann

Department of Statistics
London School of Economics and Political Science

This dissertation is submitted for the degree of
Doctor of Philosophy

October 2021

I would like to dedicate this thesis to my parents,
especially to the wonderful memories of my beloved mother.
I miss her every day!

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it). The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without my prior written consent. I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party. I declare that my thesis consists of less than 100,000 words.

Reinhard Fellmann
October 2021

Acknowledgements

I want to thank my supervisors, Erik Baurdoux and Andrea Vedolin, for their guidance and support. I am extremely grateful to Erik Baurdoux for his invitation to the Centro de Investigación en Matemáticas (CIMAT) in Guanajuato, Mexico. I am also extremely thankful to Andrea Vedolin for her invitation to the Department of Economics at Yale University. I want to thank both supervisors for their support to attend the long programme "Broad Perspectives and New Directions in Financial Mathematics" at the University of California, Los Angeles (UCLA). This research was financially supported by the Engineering and Physical Sciences Research Council (EPSRC) and the Doctoral Training Centre in Financial Computing and Data Science. I would also like to thank Tomaso Aste (internal examiner) and Rosario Nunzio Mantegna (external examiner) for their valuable feedback. I would also like to Doyne Farmer, Bryan Kelly, Xavier Gabaix and Campbell Harvey, who inspired this thesis. I am incredibly grateful to my friend Alberto Ciacci for the excellent feedback on computational work. I would also like to thank Julius Bonart, Tiziana Di Matteo, Angelos Dassios, Constantinos Kardaras, Pauline Barrieu, Milt Mavrakakis, Philip Treleaven, Rama Cont, René Carmona, Stephen Bourchier and Thomas Rauter for helpful comments, discussions and encouragement. I would also like to thank Ann Law, the LSE Systemic Risk Centre and IHS Markit Ltd. for providing credit default swap data. I would also like to thank Penny Montague, Imelda Noble Andolfo and Ian Marshall for their administrative support. I would also like to thank Patrick Haller, Benjamin Fritsche, Yantcho Kondov, Julia Langer and Conor Campbell. Most importantly, I sincerely thank my wife Karen for her unconditional love, support and care throughout this exciting journey.

Personal Note

All content and results presented in this thesis were produced between November 2018 and December 2020. Revision and correction period: June to September 2021.

Abstract

A measure of tail risk in credit markets is essential to understand the behaviour of credit default swap prices. This thesis presents three tail-risk measures based on dynamic power-law models with multiple time-varying tail parameters. The models use univariate and cross-sectional returns of sovereign and corporate credit default swaps to estimate the tail risk at each point of time. The power-law is considered a plausible statistical hypothesis for the tail distribution of returns and measure of tail risk in credit markets. The dynamic power-law exponent is time-varying and persistent. The tail exponent series for 35 European sovereign credit default swaps vary around the mean of 3.0, consistent with the inverse cubic law in other asset classes. Tests show that past exposure to extreme event risk significantly impacts future credit default swap prices and returns. A one-standard-deviation increase in tail risk forecasts an average increase in sovereign credit default swap spreads of 7.6 bps in US credit markets, which is highly significant. These results are robust out-of-sample. The forecasting power of tail risk is also robust to controlling for 25 alternative predictors. Furthermore, tail risk has substantial explanatory power for the cross-section of expected returns in US corporate credit default swaps. Cross-sectionally, firms with high positive loadings on past tail risk earn average expected annual returns 8.1% higher than credit default swaps with low tail risk covariation. Protection sellers increase spreads for credit default swaps bearing high sensitivity to tail risk. This tail risk premium is different from the premiums on market risk, idiosyncratic volatility and coskewness, and robust to considering alternative risk factors. These findings are consistent with asset pricing models that relate tail risk to expected returns and risk premiums.

Keywords: Tail Risk, Risk Premiums, Power-Laws, Index Estimation, Credit Default Swaps, Empirical Asset Pricing

Table of contents

List of figures

List of tables

Overview

Various economic and financial time series are known to exhibit tail distributions of returns with a power-law structure, pioneered by Mandelbrot (1963) and Fama (1963). In finance, the power-law behaviour is associated with the tail risk, which can have important effects on asset prices (Kelly and Jiang, 2014). While previous literature has mostly focused on equity markets (Jansen and De Vries, 1991; Mantegna and Stanley, 1996; Lux, 1996; Gopikrishnan *et al.*, 1999; Plerou *et al.*, 1999), there has not yet been any extensive study on power-laws in sovereign and corporate credit default swaps and their application asset pricing. This gap in the literature motivated this thesis to study time-varying extreme event risk in credit markets.

The main goal of this thesis is to investigate the effects of time-varying tail risk on prices and returns of sovereign and corporate credit default swaps. The major obstacle to this investigation is finding a sound measure of tail risk over time. Ideally, one would directly construct a tail risk measure from the underlying credit default swap time series as discussed in Chapter 1. However, estimating time-varying tail risk in univariate time series is challenging because of the infrequent nature of extreme events. A sufficiently long time series is required to provide enough information for an accurate estimation of the tail risk. However, long time series may include tail events from the distant past with little information about the current tail risk. To overcome this problem, we conceive two different return aggregation methods. First, in Chapter 3, we estimate the economy-wide tail risk by aggregating returns of sovereign credit default swaps across the term structure of default. Aggregated return data significantly reduce the lookback window from several to one year. The second method to measure the tail risk is built on the intuition that the tail risk of individual firms is closely related to the economy-wide tail risk. In Chapter 5, we estimate the market-wide tail risk by aggregating returns of corporate credit default swaps. If a common underlying process drives tail risk dynamics of firms, the cross section of extreme returns can be used to accurately measure prevailing tail risk in the economy. The main advantage of using the cross section of extreme returns is shortening the lookback window to one month of daily returns, which only captures the most recent tail risk.

Our framework fuses econophysics with asset pricing. Our empirical framework is based on a reduced form description for the distribution of tail returns. The set of tail returns is defined as the observations exceeding some high threshold value, denoted x_{min}. We assume that the tail of asset return i behaves according to the tail probability distribution

$$P(x > x_{min}) = \left(\frac{x}{x_{min}}\right)^{-\alpha}, \tag{1}$$

where $0 < x_{min} < x$. Equation 1 states that returns greater than some minimum x_{min} obey a power law. There are two key parameters that determine the shape of tail. The first parameter is α known as the "scaling parameter" in physics or the "tail exponent" in finance. Low values of α correspond to fat tails and a high probability of extreme events. Throughout this thesis, we estimate α for observations above x_{min} using the Hill (1975) estimator.

The second key parameter is x_{min} known as the lower bound on the power-law behaviour or tail threshold. Choosing the tail threshold is a notoriously difficult task and has important effects on the tail exponent estimate. Throughout this thesis, we discuss and apply three methods to estimate the tail threshold. First, Chapters 1 and 2 use the Kolmogorov-Smirnov method to identify the tail threshold (here denoted by m_t). The Kolmogorov-Smirnov technique is a quantitative traceable method, which removes the non-power-law portion of the distribution from the estimation of the scaling parameter. Then the fit to the tail distribution has a simple functional form that allows us to test the level of agreement between the tail returns and the best-fit power-law model. However, the main drawback is that the Kolmogorov-Smirnov method can lead to abrupt changes of the tail threshold caused by statistical origins and not by fundamental changes of the risk behaviour. Consequently, this causes significant variations in the tail exponent, tail length and percentile. Second, to overcome this issue, we propose a new selection method based on a smoothing technique for the tail threshold in Chapter 3. This smoothing technique estimates the new optimal tail threshold (here denoted by u_t) from a moving average of the tail percentile. The smoothened tail threshold eliminates unwanted fluctuations due to abrupt changes in the tail length. The disadvantage of this approach is that it is computationally expensive, especially on a large cross section of assets. Therefore, we decrease the computational complexity by reducing the number of parameters required for the tail exponent estimation. Third, we replace the Kolmogorov-Smirnov method with a fixed sample size approach. A fixed percentage of the total sample is used to estimate the tail threshold (here denoted by a_t). This heuristic method is quite common in the quantitative finance literature. Gabaix *et al.* (2006) advocate a simple rule fixing the a-threshold at 5% (95%) quantile for power-law estimation of the lower (upper) tail. We follow these authors by applying fixed percentiles in the remaining Chapters 4, 5, and 6.

We implement the dynamic power-law estimator using daily returns of sovereign and corporate credit default swaps. Chapter 1 analyses trading prices of sovereign credit default swaps provided by a derivatives dealer from the G16 industry group. The remaining chapters use composite level information provided by IHS Markit through the LSE Systemic Risk Centre. While we study the tail risk of sovereign credit markets in different regions (Europe, United States, and global), we also move from single time series to panel data, for example cross-maturity data in Chapter 3, and cross-country data in Chapter 4. The motivation for shifting to panel data is twofold. First, a larger cross section allows a reduced lookback window, and therefore avoids using outdated information. Second, a larger cross section captures more extreme events, and therefore more accurately estimates the tail exponent. Because the cross section of sovereign credit default swaps is rather small, Chapter 5 and 6 derive the tail risk measures from a large cross section of US corporate credit default swaps.

Chapters 1 and 2 analyse the time-varying tail risk in European sovereign credit default swaps from January 2005 to March 2017. This study is performed on 35 univariate time series of seven countries and five maturities. Chapter 3 investigates the time-varying tail risk in US credit markets from January 2008 to March 2017. This analysis is based on a panel approach, aggregating daily returns from ten maturities of US sovereign credit default swaps (cross-maturity data). Chapter 4 explores the time-varying tail risk for specific credit default swap maturities from the cross section of 46 sovereign credit default swaps from January 2009 to March 2017 (cross-country data). Chapters 5 and 6 extend the time-varying tail risk analysis to US corporate credit default swaps from January 2009 to March 2017. In Chapter 5, we estimate the tail risk of an industry from the cross section of firms. Furthermore, we measure the market-wide tail risk from the cross section of extreme returns of all 675 firms. Throughout this thesis, we only consider price information of credit default swaps for senior unsecured debt traded under the Cum Restructuring (XR and XR14) clause.

This thesis is structured as follows. Chapter 1 explores whether the power-law is a plausible hypothesis for sovereign credit default swap returns. Our statistical framework for discerning and quantifying power-law behaviour in empirical data builds on the seminal paper of Clauset *et al.* (2009). This approach combines maximum-likelihood fitting methods with goodness-of-fit tests based on the Kolmogorov-Smirnov statistic in a static (non time-varying) environment. We extend this procedure to estimate the time-varying tail statistics by using a rolling window approach. Furthermore, we explore whether stylised facts of equity markets, which are analogous to power or scaling laws in statistical physics, can be found in credit markets, e.g. the cubic law of large returns (Gopikrishnan *et al.*, 1998), the scaling behaviour of return distributions over shorter and longer horizons (Gopikrishnan *et al.*, 1999), and the asymmetries in the distribution of returns (Plerou *et al.*, 1999).

Chapter 2 analyses the factors affecting the change of the tail exponent over time. The time-varying tail risk measure proposed in the previous chapter, accounts for the variation in extreme returns and the variation in tail length (variation in tail threshold). A challenge of this estimation approach is to separately quantify changes in the tail exponent from day t to the consecutive day-ahead $t+1$ due to fluctuations in tail returns and variations in tail length, especially when both effects coincide. To solve this problem, we developed a novel tail exponent decomposition method. The idea of the decomposition method is to quantify the tail exponent changes due to tail returns and tail length variations separately. We perform this study on different rolling windows with different lengths and compare the evidence across asset classes and aggregation methods.

Chapter 3 links our research on tail exponent estimation with asset pricing. We analyse whether the time-varying tail exponent impacts future prices of US sovereign credit default swaps. To test this hypothesis, we build the tail exponent estimator for the dynamic power-law structure by aggregating returns for different maturities of the same underlying asset (US sovereign credit default swaps with maturities from 1 to 30 years). This cross-maturity approach captures more tail returns than univariate time series, which avoids accumulating years of tail observations with no causality and information about the current market situation and tail risk. If the tail risk is persistent, tail risk should positively forecast price increases of credit default swaps because a positive tail risk shock increases the price required by the protection seller to sell a credit default insurance. We investigate this hypothesis with a series of predictive regressions, where the dependent variable is the future credit default swap price for different maturities. We explore the robustness of the forecasting power of the tail risk measure to controlling for a large set of 25 alternative predictors.

The following two chapters relate to aggregation properties of power-laws. Power-laws have excellent aggregation properties that hold under various transformation rules (see e.g., Jessen and Mikosch (2006) or Gabaix *et al.* (2006) for a summary). For instance, a finite sum of independent power-law distributed variables with tail exponent α is also power-law distributed with the same tail exponent α. The previous chapter aggregates power-law distributed returns of US sovereign credit default swaps with maturities from 1 to 30 years. When we combine power-law variables from credit default swaps with different maturities, the general rule is that the smallest tail exponent from the univariate time series dominates the power-law exponent for the aggregated time series. To identify whether a specific maturity tends to dominate the power-law, Chapter 4 studies the relationship between tail risk and maturities in the cross section of returns of global sovereign credit default swaps. The characterisation of extreme event risk as a function of maturity is referred to as the "term structure of tail risk". The term structure of tail risk helps to describe tail risk and maturity patterns such as increasing, decreasing, hump-shaped or flat term structure patterns. If power-law exponents are not significantly different for different maturities, the tail risk's term structure is considered "flat", implying that short- and long-dated contracts bear a similar level of extreme event risk.

Chapter 5 studies the degree of commonality in time-varying tail exponents across firms that share a common factor. If tail risk is time-varying and the evolution of tail risk for different factors is highly correlated, then the cross section of extreme events can be used to identify the common, time-varying element of tail risk. Therefore, we investigate correlation effects of the time-varying tail risk for two factors. First, corporate credit default swaps are assigned to ten industries based on the IHS Markit industry classification scheme. We estimate ten industry tail exponents, month-by-month, in the cross section of daily returns in the industry portfolios. Second, we conduct the same study based on firm idiosyncratic default risk. We split the sample of corporate credit default swaps into five non-overlapping subsets and estimate the cross-sectional tail risk for each risk category. This study is directed to provide empirical evidence that the cross section of extreme events from individual firms can be used for modelling the common firm-level tail dynamics.

Chapter 6 examines the pricing of aggregated tail risk in the cross section of expected returns in US corporate credit default swaps. We investigate whether firms with high sensitivity to tail risk carry a risk premium over short and long horizons (out-of-sample, for one to twelve-month holding periods). We estimate the aggregated tail risk from the cross section of daily returns of corporate credit default swaps. Our approach applies the Hill (1975) tail risk estimator to the cross section of extreme returns at the end of each month t. In each month $t+1$, we estimate each firm's tail risk betas (sensitivity) by regressing monthly returns over 60 months on the end-of-month tail risk exponent. We then sort firms into quintile portfolios based on their estimated sensitivity to tail risk. The difference in expected returns between quintile 5 and quintile 1 is the risk premium. We study whether the tail risk premium is significant for one-, three-, six- and twelve-month out-of-sample data. Furthermore, we test the hypothesis that alternative explanatory characteristics explain differences in expected returns, such as market risk beta (Fama and MacBeth, 1973), idiosyncratic volatility (Ang *et al.*, 2006b), coskewness risk (Harvey and Siddique, 2000), and idiosyncratic default risk. Finally, we check whether the risk premium for exposure to tail risk is robust to controlling for these alternative factors.

Chapter 1

Tail Risk in Credit Markets: A Dynamic Power-Law Model

1.1 Introduction

Various economic and financial time series are known to exhibit distributions with a power-law decay (Gabaix, 2009). In finance, the power-law behaviour is associated with the tail risk of asset prices. A series of power-laws are reported for many asset classes including foreign exchange rates (Guillaume *et al.*, 1997), stocks in developed markets (Plerou *et al.*, 1999), financial market indices (Gopikrishnan *et al.*, 1999), equity trading volumes (Gabaix *et al.*, 2003), and cross-sectional returns of US equities (Kelly and Jiang, 2014). While previous literature has mostly focused on equity markets, there has not yet been any extensive research in credit markets. We help to close this gap in the literature by studying the tail risk in sovereign credit default swaps.

The goal of this research is to investigate the statistical behaviour of extreme values in credit markets over time. The chief obstacle to this investigation is to define a simple and accurate measure of tail risk, which allows for time-varying tail thresholds and tail percentiles. Conventional techniques in quantitative finance limit the tail analysis to arbitrary, fixed percentiles of the total sample size, for example, 5% of the upper order statistics such as in Kelly and Jiang (2014). However, fixed percentiles do not account for variations of the tail length over time, which may lead to misspecifications of the tail risk. Furthermore, fixed percentiles have a weak theoretical foundation and might therefore not be robust (Danielsson *et al.*, 2016).

To overcome this problem, we devise an objective and quantitative traceable method to estimate the time-varying tail statistics. The non time-varying methodology is inspired by Clauset *et al.* (2009). We extend this procedure to estimate the time-varying tail percentiles and tail risk with a rolling window approach. We utilise the Kolmogorov-Smirnov method to determine the smallest tail return, denoted m_t in period t. The threshold value of $m_t > 0$, also known as the lower bound on the power-law behaviour, then specifies the length of the tail (k_t) by separating the tail from the body of the distribution. The tail length (k_t) in period t is defined as the number of returns greater than m_t. The Kolmogorov-Smirnov method defines the lower bound $m_t > 0$ on the power-law behaviour by minimising the maximum deviation between theoretical cumulative distribution function for a power law and the empirical cumulative distribution function for a power-law model that best fits the data above this threshold value.

Our empirical framework centres on a power-law description for the tail distribution of returns. The set of tail returns is defined as the tail observations falling above the threshold value ($0 < m_t < x_t$) at time t. We assume that the tail returns of sovereign credit default swaps behave according to a power-law distribution, such that

$$p(x_t > m_t) = \frac{\alpha_t - 1}{m_t}\left(\frac{x_t}{m_t}\right)^{-\alpha_t}, \tag{1.1}$$

where α_t is the scaling parameter, x_t the tail returns and m_t the threshold value of the tail at time t. The key parameter α_t is also referred to as credit tail risk exponent. In Equation 1.1, primary factors of the model are the tail risk exponent α_t and the tail threshold m_t, which determine the shape of the tail. The power-law distribution for tail returns is particularly interesting for empirical and theoretical research, because of their heavy-tail characteristics, which means rare and infrequent extreme events are far more likely than they would be in, for example, a Gaussian distribution. Low values of the tail risk exponent α_t correspond to fatter tails and a high probability of extreme events (high tail risk), and vice versa. In finance, the scaling behaviour of α_t is associated with the tail risk of asset prices and has an essential function for pricing of credit risk. The heavy tails can take so extreme values such that the mean (for $\alpha_t < 2$) and standard deviation of the distribution (for $\alpha_t < 3$) can be undefined.

In contrast to previous dynamic power-law models, Equation 1.1 is a tail risk model with both time-varying tail threshold and tail percentiles. Compared to models with fixed percentile, which only have time-varying tail thresholds, our model allows for simultaneous changes of the threshold value m_t and tail length k_t. This is essential for measuring the tail risk exponent, because the number of tail events impacts the tail exponent and standard error. Furthermore, we find empirical evidence that the tail risk exponent α_t varies with the number of tail events and the default risk of a country. During the European sovereign credit crisis, the positive tail percentile and default risk are highly positively correlated (93%). After the peak of the financial crisis, the positive tail percentile decreases in length, and the negative tail percentile strongly correlates with default risk (-81%). Taking these fluctuations of the tail percentiles into consideration is essential for accurate estimations of the tail risk.

We build a tail risk measure from the dynamic power-law model in Equation 1.1. We identify the tail threshold using the Kolmogorov-Smirnov distance. An accurate estimation of the minimum tail return is crucial for precisely measuring the tail exponent. Then we estimate the tail exponent for observations above the threshold value using the Hill (1975) estimator. The method of choice for fitting the hypothesised power-law distribution to observed data is the method of maximum likelihood, which provably gives accurate parameter estimates in the limit of large sample size (see, for example, Barndorff-Nielsen and Cox (1994) and Wasserman (2013)). Various other estimation methods for measuring the tail behaviour have been proposed, i.e. de Haan and Resnick (1980), Hall (1982), Mason (1982), Davis and Resnick (1984), Csorgo *et al.* (1985), and Hall *et al.* (1985).[1] Recent research in statistics of extreme values shows that the Hill (1975) estimator performs well even in the presence of dependent and heterogeneous data (Kelly, 2014). However, the tail exponent can almost always be retrieved from a time series regardless of whether the data genuinely fit a power-law distribution or not. Most previous studies in finance do not quantitatively assess whether the power-law is a plausible model. For this reason, we use a goodness-of-fit test, which generates a p-value that quantifies the plausibility of the power-law hypothesis. To support the power-law assumption in credit markets, we fit our data to competing distributions, estimate those p-values, and compare it to the p-values of the power-law model. This approach is inspired by the research by Clauset *et al.* (2009). Finally, we conduct a comparison between fixed and time-varying tail percentiles. We show that small misspecification of the tail length can lead to significant deviations of the tail risk exponent and false rejection of the power-law hypothesis.

[1] Other estimation methods for measuring the tail behaviour include the asymptotic estimate constructed for the index of a stable distribution with convergence at a logarithmic rate by de Haan and Resnick (1980); the estimates of an exponent of regular variation with convergence at an algebraic rather than a logarithmic rate by Hall (1982); the seminal paper on the tail estimation of distributions with exponential-like upper tails by Mason (1982); the estimation approaches from the classical extreme value theory Pickands *et al.* (1975) and Davis and Resnick (1984); the kernel estimator approach by Csorgo *et al.* (1985); the adoptive estimator by Hall *et al.* (1985), and others.

Our research on power-laws in credit markets draws on several strands of literature in physics, economics and finance. Firstly, our power-law model is build on a widely accepted statistical framework for discerning and quantifying power-law behaviour in empirical data (Clauset *et al.*, 2009). We extend this methodology by introducing a rolling window approach to estimate the power-law dynamics over time. Gabaix (2009) surveys well-documented empirical power-laws, such as the size of cities (Krugman, 1996), size of firms (Axtell, 2001), increase in CEO compensation proportion to the average size of firms (Gabaix and Landier, 2008), income and wealth (Atkinson and Piketty, 2007), international trade (Hinloopen and van Marrewijk, 2006), among other power-laws in economic and finance. Bouchaud (2001) discusses ideas from physical statistics to shed some light on the origin of power-law distributions and power-law correlations in financial time series. While the origin of power-laws in finance remains a controversial topic (Gabaix *et al.*, 2003; Farmer *et al.*, 2004; Plerou *et al.*, 2004), there is evidence that many financial time series obey a power-law decay in their tails (Mantegna and Stanley, 1996; Guillaume *et al.*, 1997; Plerou *et al.*, 1999; Gopikrishnan *et al.*, 1999; Gabaix *et al.*, 2003, 2006; Kelly and Jiang, 2014; Kyle and Obizhaeva, 2016; Gabaix, 2016).

Secondly, the tail distribution of returns has been analysed in a series of studies that price fluctuations are distributed according to a power-law with exponent $\alpha \simeq 3$. Gopikrishnan *et al.* (1998) investigate the power-law behaviour for the New York Stock Exchange (NYSE), the American Stock Exchange (AMEX), and the National Association of Securities Dealers Automated Quotation (NASDAQ). Gopikrishnan *et al.* (1998) find an asymptotic power-law behaviour for the cumulative distribution with an exponent $\alpha \simeq 3$ for these three major US stock markets (40 million data points). Gopikrishnan *et al.* (1998) call the Equation 1.1 with an exponent of $\alpha \simeq 3$ the "cubic law" or "inverse cubic law" of returns. Furthermore, Gopikrishnan *et al.* (1999) study the distribution of price fluctuations of international stock market indices over different time scales, denoted δt, for returns from 1 minute up to more than 1 month. Gopikrishnan *et al.* (1999) find the distribution of S&P 500 index returns to be consistent with a non-stable power-law functional form ($\alpha \simeq 3$) until four days ($\delta t \leq 4$ days), after which an onset of convergence to Gaussian behaviour is exhibited. Gopikrishnan *et al.* (1999) confirm the robustness of the asymptotic power-law behaviour for other international stock market indices (i.e. Nikkei and Hang-Seng index). The particular scaling exponent $\alpha \simeq 3$ is consistent with a finite variance, but moments higher than 3 are unbounded. Plerou *et al.* (1999) examine the statistical properties of stock price fluctuations for 1,000 publicly-traded US companies with the largest market capaitalisation. Plerou *et al.* (1999) find that the cumulative distribution of individual-firm returns is consistent with an asymptotic power-law behaviour with a tail exponent $\alpha \simeq 3$. These return distributions appear to retain the same functional form until approximately 16 days ($\delta t \leq 16$ days). For longer time scales, Plerou *et al.* (1999) also report a convergence to Gaussian behaviour. Furthermore, Plerou *et al.* (1999) examine firms of different market capitalisation, where small firms tend to have a higher volatility than large firms. Adjusted for different levels of volatility, the cumulative distributions of the normalized returns have similar functional forms with exponent $\alpha \simeq 3$ for small and large firms. An empirical study of Gabaix *et al.* (2003) report that the "inverse cubic law" of Equation 1.1 is rather "universal", holding over as many as 80 standard deviations for some stock markets, with δt ranging from one minute to one month, across different sizes of stocks, different time periods, and also for different stock market indices. Gabaix *et al.* (2003) test the universality of Equations 1.1 by analysing the 35 million transactions of the 30 largest stocks on the Paris Bourse from 1994 to 1999. The analysis shows that power-laws obtained for US stocks also hold for a distinctly different stock market, consistent with the possibility that power-law behaviour might be universal across the large set of stocks and indices. Furthermore, Gabaix *et al.* (2003) analyse the cumulative distribution of the absolute values of the normalised 15 minute returns of the 1,000 largest firms in the Trades And Quotes (TAQ) database from 1994 to 1995 (12 million observations). Gabaix *et al.* (2003) find that normalised tail returns are distributed with a power law exponent of $\alpha = 3.1 \pm 0.1$ between 2 and 80 standard deviations of returns. It also appears that the inverse cubic law holds internationally (Lux, 1996; Gopikrishnan *et al.*, 1999; Gu *et al.*, 2008). For example, Makowiec and Gnacinski (2000) study the Warsaw Stock Exchange Index (WIG) for five years (1995-2000) and found a scaling exponent of $\alpha = 3.06$ for negative tails. Various other studies support the cubic power-law across asset classes Guillaume *et al.* (1997) and Plerou *et al.* (2005). For example, Dacorogna and Pictet (1998) studies intraday price movements in foreign exchange markets and report a power-law with $\alpha \simeq 3$ for short time scales. Similarly, GARCH models generate power-laws, but need to be fine-tuned to replicate the exponent of 3 (Gabaix *et al.*, 2006).

A tail exponent $\alpha \simeq 3$ contradicts the “stable Paretian hypothesis” of Mandelbrot (1963), which proposes that financial returns follow a Lévy stable distribution. A Lévy distribution has an exponent $\alpha \leq 2$, which is inconsistent with the empirical evidence, e.g. Fama (1963), McCulloch (1996) and Rachev and Mittnik (2000). However, there are also some dissenting views to the cubic law in literature. For example, Yan *et al.* (2005) investigate daily returns of 104 stocks from the Shanghai and Shenzhen Stock Exchanges from 1994 to 2001 and argue that the tail exponent is $\alpha = 2.44$ for positive and $\alpha = 4.29$ for negative tail returns. Zhang *et al.* (2007) remove the opening and close returns of high-frequency data for the Shanghai Stock Exchange Composite Index and show that the tail exponents are much closer to inverse cubic. Other studies in emerging markets also report divergences from the cubic law, e.g. Wang and Hui (2001), Lee and Lee (2004), Pan and Sinha (2007) and Mu *et al.* (2010).

Thirdly, a closely related topic to the cubic law concerns the time scale δt which defines the returns. Generally speaking, the tail distribution evolves from power-law at small time scales to Gaussian at large scales (Ghashghaie *et al.*, 1996). This behaviour is often called "aggregational Gaussianity" (Gopikrishnan *et al.*, 1999; Cont, 2001; Zhao, 2010). Longer-horizon return distributions are shaped by two opposite forces (Gabaix, 2009). The first force is that a finite sum of independent power-law distributed variables with exponent α is also power-law distributed, with the same exponent α. (Gabaix, 2009) state that returns remain power-law distributed if the time-series dependence between returns is not too large. The second force is the central limit theorem. This implies that if the distribution of returns has a finite second moment ($\alpha \geq 2$), one would expect convergence to a Gaussian. However, several studies suggest the distributions of returns retain their power-law functional form for time scales up to days, weeks or even month (Plerou *et al.*, 1999; Gopikrishnan *et al.*, 2000). (Gabaix, 2009) conclude that under return aggregation, the central part of the distribution becomes more Gaussian, while the tail return distribution remains a power-law with exponent α but have an ever smaller probability (larger tail exponent), so that they may not even be detectable in practice. This is emprically found in several studies. (Plerou *et al.*, 1999) analyse about 16,000 companies from the CRSP database from 1962 to 1996. The authors find that the cumulative distribution for normalised positive and negative returns of individual firms retain their power-law functional form for time scales $\delta t = 1, 2, 4, 8$ and 16 days. The estimate of the positive tail exponent increases monotonically, whereas for the negative tail the tail exponent estimates are approximately constant. The scaling behaviour of the distribution of normalised positive returns appears to break down for time scale beyond 16 days and indicates a slow convergence to Gaussian behaviour. Interestingly, (Plerou *et al.*, 1999) find that the scaling behaviour of the distribution of normalised negative returns does not converge to Gaussian and fat-tails, even at yearly horizons preserve the cubic law. In another study, Gopikrishnan *et al.* (1999) study distribution of price fluctuations of financial market indices for different time scales. For the cumulative distribution of returns for the S&P 500 index, Gopikrishnan *et al.* (1999) reports a stable functional form for different time scales up to approximately 4 days for positive and negative tails. For larger time scales the results are consistent with a slow convergence to Gaussian behaviour (Lux, 1996; Gopikrishnan *et al.*, 1999). It is also found that return distributions retain their power-law form for different time scales in other asset classes, such as for foreign exchange rates (Dacorogna *et al.*, 1995). Gabaix (2009) conclude that the existing literature shows that while high frequencies offer the best statistical resolution to investigate the tails, power laws still appear relevant for the tails of returns at longer horizons, such as a month or even a year.

Fourthly, asymmetric returns are well-documented in financial literature. Research in behavioural finance offers evidence of loss aversion and asymmetric perception of risk in equity markets (Benartzi and Thaler, 1995; Shleifer, 2000; Barberis and Thaler, 2003). Merton *et al.* (1985) states that it is difficult to see a clear theoretical explanation for extreme events being symmetric. For the cumulative distribution of returns for the S&P 500 index, Gopikrishnan *et al.* (1999) report slightly heavier negative tails on 12 out of 17 different time scales. For individual stocks, (Plerou *et al.*, 1999) provide evidence for heavier negative tails on all 18 time scales ranging from $\delta t = 5$ minutes to $\delta t = 1024$ days. Furthermore, LeBaron and Samanta (2005) report evidence for asymmetric perception of tail risk in equity markets. Stoyanov *et al.* (2017) study tail asymmetry of different types of markets, before and after the financial crisis of 2008.

Finally, since at least Mandelbrot (1963) and Fama (1963), economist have argued that unconditional distributions of financial returns are characterised by volatility clustering, heavy tails, aggregational gaussianity and power-law scaling. To accommodate the extremal properties of financial returns, several models have been proposed in the econometric literature. Engle (1982) and Bollerslev (1986) introduce the generalised autoregressive conditional heteroskedasticity (GARCH) model to model conditional return distributions with heavy-tailed i.i.d. innovations. In order to further capture extreme returns, Bollerslev (1987) incorporates Student t shocks in the GARCH model. Subsequent models, combine Extreme Value Theory (EVT) to model tail behaviour of the asset returns with traditional generalised autoregressive conditional heteroskedasticity (GARCH) models. For example, McNeil and Frey (2000) combines the Peak over Threshold (POT) approach based on the generalised Pareto distribution (GPD) and the traditional GARCH model for estimating conditional Value at Risk (VaR). The EVT-GARCH approach of McNeil and Frey (2000) performs especially well in risk management applications such as Value at Risk (VaR). Kuester *et al.* (2006) finds that a hybrid method, combining a heavy-tailed generalized autoregressive conditionally heteroskedastic (GARCH) filter with an extreme value theory-based approach, performs best overall, for predicting VaR in a univariate context. Our dynamic power-law approach differs in explicitly allowing for time-variation in the tail exponent (α_t), above and beyond any volatility dynamics in returns.

We contribute to the literature on power-laws in finance in several ways. We extend the static power-law model to a dynamic power-law model for measuring tail risk (exponent) over time by using a rolling window. Unlike recent research on power-law models in finance (Kelly, 2014), our model allows for time-varying tail thresholds and tail percentiles depending on the risk of the underlying asset. We provide a quantitative approach to assess whether the power-law is a plausible model to estimate the tail risk (power-law hypothesis test). This is the first study that examines the dynamic tail risk in credit default swap markets.

Our main contribution is an empirical analysis of the time-varying tail risk in European credit markets. We implement the dynamic tail risk model in Equation 1.1 using credit default swaps returns of seven countries and five maturities from 2005 to 2016. We find that the power-law model is a feasible model for measuring the dynamic tail risk in credit markets. We estimate the tail exponent for rolling time windows with a length of two and four trading years. For both time windows, the average tail length is sufficiently large to estimate the tail statistics and perform the power-law hypothesis test.[2] The power-law hypothesis test assesses whether the tail observations genuinely fit a power-law distribution. We find that the power-law is a plausible hypothesis for positive, negative and absolute logarithmic returns of sovereign credit default swaps, independently of the time window.[3] We confirm these findings for normalised log-returns.[4] Furthermore, we compare the power-law model to competing tail distributions. We find that other models, such as log-normal or exponential, are not superior over the power-law model. We conclude that the power-law model is a feasible model for measuring tail risk in credit default swaps.

We investigate whether the inverse cubic law of equity markets also holds in credit markets. Therefore, we estimate the tail exponent for over 750,000 data sets of credit default swap returns. We observe that the average tail exponent across all countries and maturities follows the inverse cubic law. France and Portugal perfectly follow the inverse cubic law in the tail distribution of credit default swap price fluctuations. Other countries lie within the range of one standard error. The inverse cubic law holds for different regions and maturities

[2] The average tail length is 51 (74) using a rolling time window with a two-year (four-year) lookback period.

[3] For a time window of two (four) years, the average daily p-value from 2009 to 2016 for positive tail returns is 40% (39%), for negative tail returns is 46% (47%), and for absolute returns is 41% (30%).

[4] For a time window of two (four) years, the average daily p-value from 2009 to 2016 for normalised positive tail returns is 39% (39%), for normalised negative tail returns is 46% (46%), and for normalised absolute returns is 42% (39%).

of sovereign credit default swaps, which supports the robustness of the inverse cubic law. The inverse cubic law has important implications for asset and credit risk pricing. Power-law exponents outside the Lévy-stable region $0 < \alpha < 2$ supports the existence of a finite second moment for return distributions. The existence of the first two moments (mean and volatility) provides a fundamental basis for the usage of financial theories and co-variance-based techniques in credit risk management.

We next analyse the cumulative distribution at extreme values for different time scales ranging from daily (denoted $\delta t = 1$) to monthly returns (denoted $\delta t = 22$). Sovereign credit default swap returns follow a power-law decay in their tails for different time intervals. It is noteworthy that tail returns of sovereign credit default swaps follow the inverse cubic law on short time scales up to one week ($\delta = [1;5]$). Furthermore, there is evidence of two remarkable patterns. Firstly, we observe that the credit tail risk decreases with increasing time scales for all credit default swap maturities. Secondly, the term structure of credit tail risk evinces that credit default swaps with shorter-dated maturities exhibit heavier tails than longer-dated maturities, meaning that extreme price fluctuations are considerably more frequent in shorter-dated contracts. These patterns are persistent among different time scales, time windows and hold for the vast majority of countries. We infer that sellers of short-dated credit default swaps bear the highest credit tail risk on short time scales.

Finally, we investigate the asymmetry of tail risk exponents in credit default swap markets. Our hypothesis is that the upper tail exponent exhibits a high probability of extreme events compared to the lower tail exponent. While we associate positive tail returns with large price increases usually occurring in periods of financial distress, negative tail returns are related to large price decreases usually occurring in periods of financial stabilisation. We find that individual time series of credit default swaps have a higher probability of positive returns above the tail threshold, whereas the positive threshold value is higher than the negative one. The average extreme return in times of financial distress is more significant than the average extreme return in times of financial recovery. Furthermore, cross-maturity time series contain a higher likelihood of extreme price increases in times of financial turbulence. Surprisingly, we find that the core region has a higher tail risk despite a lower probability of default, which can be explained by the significant impact of volatility on the tail distribution. We explain the implications of tail risk asymmetry for risk management and portfolio allocation.

The rest of this chapter is organised as follows. Section 1.2 defines the notion of credit tail risk, describes the characteristics of a credit default swap, data and selection criteria.

Section 1.3 extensively discusses the empirical methodology of the power-law model for estimating tail risk in sovereign credit markets. Section 1.4 reports the results of the power-law hypothesis test, the cubic power-law, the tail exponent on different time scales, and finally, the asymmetric behaviour of tail shocks. Section 1.5 concludes.

1.2 Definitions and Data

In this section, we briefly explain the main characteristics of sovereign credit default swap contracts, qualitatively define the notion of "credit tail risk", and elaborate on our data set of sovereign credit default swaps.

1.2.1 Sovereign Credit Default Swap

Sovereign credit derivatives are contingent liability claims with payoffs that are linked to the creditworthiness of a country. Credit derivatives on sovereign debt allow market participants to hedge, trade and manage risk associated with certain debt-related events, e.g. changes of credit rating, releases of economic fundamentals or fiscal data. It can be regarded as a form of insurance against the credit risk of default of the underlying government debt. Basic credit derivatives include spread options, total return swaps, and credit default swaps, where the latter is the most liquid one (Blanco *et al.*, 2005).

This study uses the information in credit default swaps to provide a direct measure of the tail and default risk in sovereign credit markets. Unlike government bonds, sovereign credit default swaps have constant maturities (also called tenor), and the underlying instrument is always valued at par value (Schönbucher, 2003). Credit default swaps are one of the most common types of credit derivative and concentrate liquidity in one instrument. In a sovereign credit default swap contract, the protection seller is obligated to compensate the protection buyer in the event of default of a country before the maturity of the contract. In case there is no default event before maturity, the protection seller pays nothing. The protection buyer pays a constant quarterly fee each period until either default occurs or the contract matures, whichever is first. If a default event occurs, the accrued premium is also paid in this period. In this study, we refer to the annualised premium as the credit default swap price. The International Swaps and Derivatives Association (ISDA) documentation states that credit default events typically include bankruptcy, failure to pay, obligation default or acceleration, a repudiation or moratorium for sovereign entities, or a restructuring. If a credit event occurs before the maturity of the contract, there are two forms of settlement: physical and cash settlement. In the case of physical settlement, which is most commonly in credit default swap contracts, the protection seller pays the par value to the protection buyer in return for the delivery of the reference government bond (or a set of bonds). However, if the reference obligation is not specified, any senior unsecured obligation of the reference country may be delivered. In the case of a cash settlement, the protection seller pays the notional amount minus the post-default market value of the reference issue, or a predetermined fraction of the notional amount.

1.2.2 The Definition of Credit Tail Risk Exponent and Credit Tail Risk

Credit tail risk is the risk associated with infrequent events of extreme magnitude that have a significant impact on the default behaviour and credit pricing of the underlying asset. Empirical studies in equity markets provide evidence that tail risk has a significant impact on asset prices and varies over time, i.e. Kelly and Jiang (2014). We devise a measure of time-varying tail risk in credit markets that is motivated by asset pricing with extreme value theory and is directly estimated from sovereign credit default swap returns. Our dynamic tail risk methodology is an extension of the static power-law model by Clauset *et al.* (2009). We introduce the Kolmogorov-Smirnov method to define the minimum magnitude of the smallest extreme event (threshold value m_t) in period t. The lower bound m_t separates the body from the tail of the distribution, and consequently determines the length of the tail k_t and the size of the time-varying tail percentile. For estimating the credit tail risk, we choose only strictly positive price changes of sovereign credit default swaps, which are associated with an increase or jump in default risk. We assume that only extreme positive fluctuations hold information about the tail risk of unexpected default behaviour. After defining the number of tail events k_t, we apply the Hill (1975) estimator to those tail returns to measure the exponent or scaling parameter α_t in period t. This scaling exponent α_t is referred to as credit tail risk exponent or sovereign credit tail risk exponent. A key assumption of our model is that the tail returns ($x_t > m_t$) behaves accordingly to a power-law over time (Equation 1.1).

The credit tail risk exponent is inversely related to the concept of credit tail risk. Low values of the credit tail risk exponent correspond to fatter tails and a high probability of extreme events. Vice versa, high values of the credit tail risk exponent correspond to thinner tails and a lower probability of rare events. However, measuring fluctuations of tail risk in univariate time series remains a major challenge due to the rare nature of extreme events.

1.2.3 Data

We estimate the dynamic tail risk exponent using daily sovereign credit default swap prices from January 2005 to May 2016. The length of the time series is 2,951 days. The main data source for this investigation is a derivatives dealer from the G16 industry group.[5] This credit default swap dealer provides trading prices of credit default swaps on a daily basis. The concept of trading prices and quotes are fundamentally different. The trading price is the last price at which the credit default swap is traded, whereas a quote is a (non-binding) price indication at which price level a dealer would enter a transaction. The trading price is more informative, as it reflects the actual price and the implied probability of default at which dealers buy or sell credit derivatives.

Sovereign credit default swaps need to meet a range of criteria for the dynamic tail risk analysis. Firstly, the underlying asset of the credit default swap must be a country of the European Monetary Union. The European Monetary Union consists of 19 countries with the common currency, the Euro. The countries in the Eurozone as of 2016 are Austria, Belgium, Cyprus, Estonia, Finland, France, Germany, Greece, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, Netherlands, Portugal, Slovakia, Slovenia and Spain. As we perform this study on countries in the Eurozone, the first criterion implies that the Euro has been adopted by 2005. This criterion excludes a range of countries such as Slovenia (2007), Cyprus (2008), Malta (2008), Slovakai (2009), Estonia (2011), Latvia (2014) and Lithuania (2015). Secondly, we exclude non-EU member states, which adopted the Euro, but are relatively unimportant in terms of economic size. These countries are Andorra, Monaco, San Marino and the Vatican City State. Furthermore, we exclude countries with considerable trading interruptions (trading holds), such as in the case of Greece. The second criterion is that a credit default swap time series has sufficient price information for stable and reliable estimations of the tail risk exponent. Starting in January 2009, we apply a daily rolling backwards-looking time window of two and four trading years. To be considered in the sample, we require price information on more than 75% of all trading days (2,213 days) between January 2005 and May 2016. We exclude Finland, Luxembourg and the Netherlands due to insufficient price information. The first backwards-looking time window of four trading years includes price information from January 2005 to January 2009, and the first backwards-looking time window of two trading years includes price information

[5]The G16 is an industry group comprising the largest derivatives dealers: Bank of America-Merrill Lynch, Barclays Capital, BNP Paribas, Citi, Crédit Agricole, Credit Suisse, Deutsche Bank, Goldman Sachs, HSBC, JP Morgan, Morgan Stanley, Nomura, RBS, Societe Generale, UBS and Wells Fargo Bank.

from January 2007 to January 2009. Finally, we only use price information under the most common conventions regarding the structuring clause and transaction currency. Bai and Wei (2012) state that sovereign credit default swaps usually trade under the Cum Restructuring clause on senior unsecured debt.[6,7] After filtering our data set for these criteria, we are left with seven countries, namely, Belgium, France, Germany, Ireland, Italy, Portugal and Spain. The first three countries are usually referred to as so-called "core countries", whereas the last four countries are considered as "peripheral countries". For each country we have consistent trading data (univariate time series) for five maturities, also called "tenor", namely 1, 3, 5, 10 and 30 years. In some cases, there are small gaps of price information. We fill the missing information by assuming that the credit default swaps price remains unchanged from the previous day.

1.3 Empirical Methodology

In this section, we establish a simple and accurate methodology for fitting the power-law model to financial time series. Many of the statistical measures we describe have been discussed in previous research on power-laws. Our objective is to bring them together to establish a complete framework for analysing power-laws in finance. Furthermore, we want to demonstrate that heuristic methods lead to erroneous estimations of the tail statistics.

The methodology to estimate and validate the power-law model consists of four essential parts. Firstly, the Kolmogorov-Smirnov (KS) distance method defines the tail threshold (m). The threshold value is the lowest tail return of the distribution. Consequently, it determines the length of the tail (k), which is also known as the scaling range in statistical physics. The tail percentile (ρ) is defined as the number of tail returns (k) over the length of the time series (n). An accurate measure for the scaling range is essential to estimate the tail exponent (α). Secondly, we estimate the tail exponent from the tail observations using the Hill (1975) estimator. Note that the length of the tail and tail exponent are mutually dependent on each other $\alpha(k)$. Therefore, misspecification of the tail threshold introduces estimation errors

[6]There are four common restructuring types, namely the Cum Restructuring (CR), Modified Restructuring (MR), Modified Modified Restructuring (MM), and Ex-Restructuring (XR). Credit default swaps with restructuring clause (CR) usually have a higher price compared to credit defaults without restructuring (XR). If credit default swaps are priced correctly, the following inequality holds true: CR $\geq$ MM $\geq$ MR $\geq$ XR.

[7]Credit default swaps can be issued for four different seniority levels of the debt within the capital structure: senior, subordinated, junior and preferred. Sovereign credit default swaps are commonly traded on senior unsecured debt (Bai and Wei, 2012).

of the tail exponent. Thirdly, we introduce a quantitative approach to assess whether the power-law assumption is a plausible model, by performing a hypothesis test based on the resampling method proposed by Preis *et al.* (2011). Lastly, we want to rule out the possibility that other distributions might provide a better fit to the data. For this purpose, we compare the power-law to competing models.

1.3.1 Estimation of the Tail Threshold

The first step is to define the correct scaling range for fitting the empirical return distribution to the power-law model. Therefore, we need to specify whether the entire distribution or only a fraction of the data should be considered for fitting. For most financial time series, the power-law only holds for a fraction above a certain threshold value. The lower bound of the scaling range is the smallest tail return, also known as the tail threshold, while the upper bound is the largest tail return. Before calculating the estimate of the tail exponent, observations below the tail threshold, which may follow some other distribution have to be eliminated.

There are several techniques to choose the optimal threshold value. A heuristic approach to select the tail threshold relies on visual analysis of the distribution. A visual method to determine the tail threshold is to locate the value beyond which the complementary cumulative distribution function (CCDF) of the distribution becomes approximately straight on a double logarithmic scale, i.e. Resnick and Stărică (1997). Another technique is to express the scaling parameter as a function of the scaling threshold and to visually identify the point beyond which the scaling exponent appears stable, i.e. Dimitropoulos *et al.* (2007).

A common technique used in quantitative finance is to limit the tail analysis to arbitrary percentiles. For example, Plerou *et al.* (1999) and Gopikrishnan *et al.* (1999) only use returns larger than two, three or five standard deviations or within a range of standard deviations. Doyne Farmer *et al.* (2004) limit the analysis to the most significant returns only, such as the largest $\sqrt{n}$ or $^1/_{10}n$, where n is the number of returns. These measures are subjective and underestimated model uncertainty, such as sensitivity to noise or fluctuations in the tail of the distribution (Stoev *et al.*, 2006). Kelly and Jiang (2014) selects a fixed percentage of 5% of the total sample size for estimating the power-law tail exponent in equity markets. The equity tail risk negatively predicts real economic activities and significantly correlates with tail risk measures extracted from index options. We show that in credit markets, the number of tail data vary over time and correlate with the default risk of the underlying asset. The upper tail percentile increases during periods of financial distress. Vice versa, during periods of financial stabilisation, the lengths of the upper tail of returns decreases, whereas the size of the lower tail increases. We find that during the European debt crisis, the correlation between the implied probability of default and the positive tail length is +0.93 from 2009 to 2012. After the peak of the financial crisis and during times of financial stabilisation (2012-2014), the correlation between the implied probability of default and the size of lower tail percentile is -0.81.[8] Our findings reject the idea to limit the tail analysis to (fixed) arbitrary percentiles, and favour a more objective method to estimate a time-varying tail length in financial time series.

The shortcomings of the heuristic methods motivated our new approach. Compared to fixed percentiles, which only allow for time-varying tail thresholds, our approach allows for simultaneous changes of the threshold value (m_t) and tail length (k_t) over time. In this research, we utilise the Kolmogorov-Smirnov statistics to estimate the optimal threshold. The value of m_t subsequently determines the tail length by separating the tail from the body of the distribution. The tail length (k_t) in period t is defined as the number of returns greater than the scaling threshold m_t. This approach is partially inspired by the static threshold model by Clauset *et al.* (2007). We move from a static to a dynamic framework to model the optimal tail length over time. The fundamental objective of this approach is simple. We select the tail threshold $\hat{m}_t$ such that the probability distributions of the empirical data and the best-fit power-law model become as similar as possible above $\hat{m}_t$ at time t.[9] If we select $\hat{m}_t$ higher than the true threshold m_t, then we effectively reduce the number of tail data, which increases sensitivity to fluctuations in the tail of the distribution and decreases the goodness-of-fit test. Conversely, if we select $\hat{m}_t$ below the true threshold value m_t, the distribution will differ because of the fundamental difference between the empirical data and model by which we are describing it. The optimal threshold value has the smallest absolute distance (D_t) between empirical data and the best-fit power-law model at time t. For non-normal data, the Kolmogorov-Smirnov statistic is one of the most popular choice for quantifying the distance between two probability distributions (Press *et al.*, 1992). The Kolmogorov-Smirnov test functions as the distance metric, defining maximum distance between the complementary distribution functions (CDFs) of the empirical data and the fitted model, such that

$$D_t = \sup_{\{x_t > m_t\}} |S(x_t) - P(x_t)|, \tag{1.2}$$

[8]The term "peak" refers to the highest average credit default swap price across all countries and maturities.

[9]Here and elsewhere the ˆ symbol is used to denote estimates derived from data such as $\hat{m}_t$ or $\hat{\alpha}_t$; hatless symbols denote the true values, which are often unknown in practice.

where $\sup_{\{x_t > m_t\}}$ is the supremum of the set of distances. Here $S(x_t)$ is the empirical CDF for the tail returns $x_{i,t}$ above the tail threshold m_t, and $P(x_t)$ is the CDF for a power-law model that best fit the data over the tail region $x_t > m_t$ in period t. The estimate $\hat{m}_t$ is then the threshold value of m_t that minimises the distance D_t.

Misspecifications of the tail threshold may cause estimation bias and statistical errors. In case the minimum tail return is underestimated, the tail risk exponent is fitted to a part of the distribution which does not obey a power-law, resulting in a biased estimate of the α_t parameter. Conversely, if the minimum tail return is overestimated, relevant tail returns will be ignored, which increases both the statistical error of the tail exponent and the bias due to finite sample size effects. In light of these facts, heuristic rules such as picking a fixed percentage of the total sample size (Doyne Farmer *et al.*, 2004; Kelly and Jiang, 2014) may lead to erroneous conclusions about the tail length and risk exponent. To demonstrate the misspecification more clearly, we estimate the tail threshold and length using fixed percentages of 5% and 10% of the upper order statistics. The sample consists of more than 50,000 time series of sovereign credit default swaps with a uniform length of four trading years. Then, we examine the disparities of the tail length and threshold between the arbitrary method and the estimates by the Kolmogorov-Smirnov method. Firstly, we find that the arbitrary 5% percentiles on average overestimate the tail threshold by 25.65%, which results in an underestimation of the true tail length by 30.68%. Secondly, we observe that fixed 10% percentiles underestimate of the tail threshold on average by 33.81%, which overestimates the true tail length by 36.59%. We confirm our hypothesis that fixed tail percentiles lead to misspecifications of tail thresholds and lengths in sovereign credit default swaps. Figure 1 shows an example for Belgium, which points out the differences between the two approaches. It also shows that the tail lengths vary each year. Finally, we want to note that the misspecifications of the tail threshold have a significant impact on the estimation of the credit tail risk exponent, which is discussed in the following section.

1.3.2 Estimation of the Tail Risk Exponent

After defining the threshold value for the tail returns, the second step is the correct fitting of the power-law scaling parameter, also known as the tail risk exponent. Recall from Equation 1.1 that the time-varying power-law distribution for continuous data is defined as

$$p(x_t) = \frac{\alpha_t - 1}{m_t}\left(\frac{x_t}{m_t}\right)^{-\alpha_t},$$

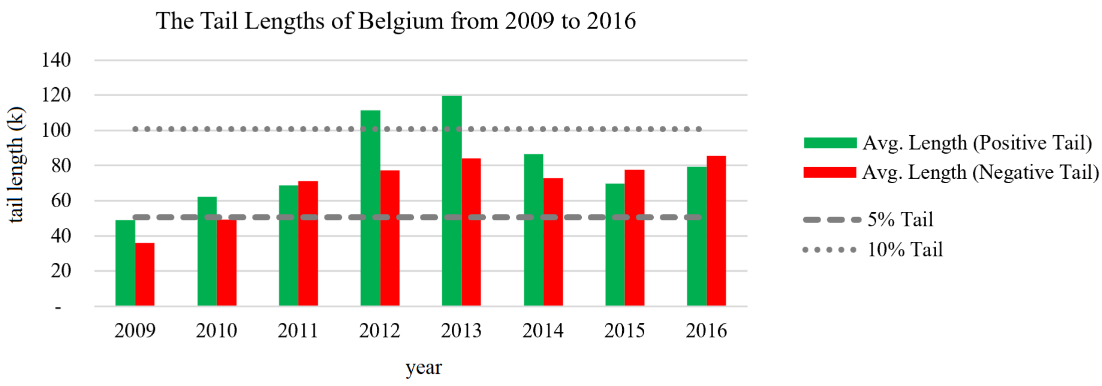

Fig. 1.1 shows the average of the optimal lengths for Belgium in each year from 2009 to 2016. The green bars show the length of the positive tail and the red bars the length of the negative tail. We observe that fixed percentiles for 5% and 10% (*grey lines*) systemically under- and overestimate the true tail lengths, while the term "true" refers to the tail length estimated by the Kolmogorov-Smirnov method. This result holds for different credit default swap maturities and countries.

where α_t is the tail exponent and m_t is the value of the tail threshold at time t. The method of choice for fitting the hypothesised power-law distribution to observed data is the method of maximum likelihood, which provably gives accurate parameter estimates in the limit of large sample size (Barndorff-Nielsen and Cox, 1994; Wasserman, 2013).[10] In the case of continuous data, the maximum likelihood estimator for the tail exponent is equivalent to the well-known Hill (1975) estimator. Given a set of tail returns $x_{i,t} > m_t$, we want to find the tail exponent of a_t for the power-law model that is most likely to have generated these tail returns. The probability that the tail returns are drawn from the power-law model with tail exponent α_t is proportional to

$$p(x_t|\alpha_t) = \prod_{i=1}^{k_t} \frac{\alpha_t - 1}{m_t} \left(\frac{x_{i,t}}{m_t} \right)^{-\alpha_t}. \tag{1.3}$$

This probability is called the likelihood of the tail returns given the model. The set of tail returns are most likely to have been drawn by the model with tail exponent α_t that maximises this function. In fact, it is convenient to maximise not the probability itself, but its logarithm

[10] Alternative methods for estimating the scaling parameter are based on linear regression, such as the fitting of a log-transformed histogram with constant bins, logarithmic bins, or rank-frequency plot (without bins). The maximum likelihood estimate holds the best results, while most of the regression methods yield significantly biased value (Clauset *et al.*, 2009).

($\mathscr{L}_t$), which has its maximum in the same place, such that

$$\mathscr{L}_t = \ln p(x_t|\alpha_t) = \ln \prod_{i=1}^{k_t} \frac{\alpha_t - 1}{m_t} \left(\frac{x_{i,t}}{m_t} \right)^{-\alpha_t}, \tag{1.4}$$

which is equivalent to

$$= k_t \ln(\alpha_t - 1) - k_t \ln m_t - \alpha_t \sum_{i=1}^{k_t} \ln \frac{x_{i,t}}{m_t}. \tag{1.5}$$

Setting $\delta \mathscr{L}_t / \delta \alpha_t = 0$, we obtain the maximum likelihood estimate for the tail exponent α_t, such that

$$\hat{\alpha}_t = 1 + k_t \left[\sum_{i=1}^{k_t} \ln \frac{x_{i,t}}{m_t} \right]^{-1}, \tag{1.6}$$

where $x_{i,t}$, $i = 1, ..., k_t$ are observed tail returns of x_t such that $x_{i,t} \geq m_t$ in period t. The method of maximum likelihood accurately estimates the tail exponent under mild regularity conditions, if the tail returns are independent, identically-distributed drawn from a power-law distribution with parameter α_t, then in the limit of large sample size $k_t \to \infty$, $\hat{\alpha}_t$ converges to α_t almost surely.[11] The Hill (1975) estimator is shown to be asymptotically normal (Hall, 1982) and consistent (Mason, 1982).

Equation 1.6 shows that the estimation of the tail exponent depends on the threshold value of the tail. To demonstrate the impact of the threshold on the tail exponent, we compare the Kolmogorov-Smirnov method to fixed percentiles. We calculate estimates of the tail exponent using the tail threshold of the Kolmogorov-Smirnov method and the tail threshold for commonly used arbitrary tail percentiles, i.e. 5% and 10%. We find that misspecifications of the tail thresholds through arbitrage tail percentiles lead to significant divergences of the tail exponents. Compared to the tail exponent estimates using the Kolmogorov-Smirnov method, the 5% percentile overestimates the tail risk exponent by +0.39 (or +12.3%) across all countries and maturities. Limiting the tail exponent analysis to the largest 10% of the sample, induces average misspecification of the tail risk exponent by -0.75 (or -26.3%). These results clearly demonstrate that fixed percentiles cause misspecifications of the tail risk exponent.

[11] For the proof of this theorem, see, for example Pitman (2018).

The standard error and corresponding confidence interval of the tail exponent $\hat{\alpha}_t$ is derived from the width of the likelihood maximum, such that

$$\sigma_t = \frac{\hat{\alpha}_t - 1}{\sqrt{k_t}} + O\left(\frac{1}{k_t}\right), \tag{1.7}$$

where the higher-order correction is positive.[12] Equation 1.7 states that the maximum likelihood estimate of the continuous power-law is asymptotically Gaussian, with variance $(\alpha_t - 1)^2/k_t$ (Muniruzzaman, 1957). The bias for a finite data set decays as $O(1/k_t)$ for any choice of m_t. The corrections $O(1/k_t)$ can be derived from the sampling distribution of the parameter $\hat{\alpha}_t$, or in other words, the distribution of deviations from the true estimate α_t due to finite-sample fluctuations. Clauset *et al.* (2009) report that the error becomes smaller than 2% of the value of α_t when $\alpha_t > 1$ and $k_t > 50$. For tail exponents fluctuating around the tail exponent value of 3, as often observed in credit default swaps, the error is smaller than 1% given $k_t > 40$. Hence, in some of the future analysis, we only consider samples with a sufficiently large number of tail observations. For finite data, the bias is present, but usually small compared to the statistical error of the tail exponent estimator, which decays as $O(k_t^{-1/2})$.

As the number of tail observations becomes large, the tail exponent estimation becomes exact. This may comprise some challenges for quantifying the dynamic tail risk in financial time series. Firstly, there might be difficulties in finding a sufficiently large number of tail events in a short time series, because extreme events are relatively rare and infrequent. Intuitively, to overcome this problem, one might suggest extending the time window until a sufficiently large number of tail events is found for stable tail risk estimations. However, this induce another challenge, because longer time series likely hold tail events from the distant past, unrelated to a current economic or financial situation of the underlying asset. It can even include data from a previous crisis without causality to the most recent events. Therefore, a long time series should be treated with caution, as it may become problematic to draw conclusions from the evolution of the dynamic tail exponent to recent changes of economic variables or to use it for forecasting purposes (see, section 3). Again, to overcome

[12]The standard error is given by the curvature of the likelihood function at the location of the maximum, which is related to the Fisher information of the function Barndorff-Nielsen and Cox (1994). We assume that $\alpha_t > 1$, since smaller values of $\alpha_t \leq 1$ are not normalisable and hence cannot occur in nature. The first proof of the asymptotic Gaussian distribution of the maximum likelihood estimate, and its relation to the Fisher information, may be found in Fisher (1922).

this problem, one could argue in favour of intraday data (change of data granularity).[13] The granularity of data and period can be chosen in a way that it obtains a sufficiently large number of tail observations. For example, if a daily return time series does not hold sufficient information, intraday returns might possess enough tail events. In the case of credit default swaps, derivative dealers provide the earliest records for tick (intraday) data for sovereign credit default swaps from 2008, and consistently for all countries of the eurozone after 2010. However, we do not have access to intraday credit default swap data. Therefore, we work with daily returns, trying to keep the time series reasonably short but large enough for accurate estimations of the tail risk exponent.

1.3.3 Testing the Power-Law Hypothesis

In the previous two sections, we introduced estimation methods for the tail threshold and tail risk exponent. However, these parameters can always be retrieved from credit default swap data regardless of whether the data genuinely fit a power-law distribution or not. For this reason, it is essential to use a quantitative approach to assess whether the power-law hypothesis is a plausible model given some empirical time series. The power-law hypothesis test requires the simulation of synthetic time series, the calculation of the p-value, the assessment of accuracy, and finally, the definition of the decision criteria. Thereby, our approach follows the initial proposal of Clauset *et al.* (2009).

A sensible approach for assessing the model plausibility, is to test whether the detected deviations between empirical data and a hypothesised power-law model can be explained as mere random effects. Therefore, we generate a set of synthetic time series with the properties of the empirical time series, namely, the tail lengths (k_t) and tail risk exponent ($\hat{\alpha}_t$) at each point of time t. The length of the time series (n_t) and the tail lengths (k_t) of the synthetic time series is equivalent to the empirical one. To generate the synthetic time series, we use a semiparametric approach. We generate a new time series of length (n_t), where the random tail returns above ($x_{i,t} > m_t$) are drawn from a hypothesised power-law distribution with tail exponent ($\hat{\alpha}_t$) and probability of (k_t/n_t). The body of the distribution (n_t-k_t) follows the empirical non-power-law distribution. Returns of the centre part of the distributions are selected uniformly at random from the empirical time series ($x_{i,t} < m_t$) with the probability ($1 - k_t/n_t$). Repeating the process above for a large number of iterations (h), we generate a complete set of synthetic time series that follows a power-law with (α_t) above (m_t), which is used to access the power-law hypothesis. Therefore, we fit each synthetic time series

[13] Granularity refers to the data (return) frequency, i.e. every minute, hour or day.

individually to its own power-law model. We calculate the Kolmogorov-Smirnov statistics for the synthetic data sets relative to the best fit power-law for the simulated data set, and not relative to the original power-law distribution from which the data was drawn. Then, the p_t-value is defined to be the fraction of the synthetic distances ($D_{t,i}^{(s)}$) that are larger than the empirical distance ($D^{(e)}$), such that

$$p_t = \frac{\text{number of } (D_{t,i}^{(s)} > D^{(e)})}{\text{number of simulations } (h)}. \tag{1.8}$$

The accuracy of the p_t-value in Equation 1.8 depends on the number of iterations (h). The error (ε_t) of the p_t-value decreases with an increasing number of simulations. Clauset *et al.* (2009) state that a good rule of thumb is to simulate at least $1/4\ \varepsilon^{-2}$ sets. Thus, if we require the p_t-value to be accurate to about two decimal digits ($\varepsilon_t = 0.01$), it involves the simulation of at least 2,500 synthetic time series. A remark on the computational complexity: for the dynamic tail exponent analysis, the computational time is approximately one day per credit default swap series from 2009 to 2016. The estimation encompasses approximately 2,000 tail exponent per country and maturity and 2,500 iterations for accurate p-values, resulting in 5 million computations. The computational time increases exponential with the length of the time series (n) and numbers of iterations (h). To reduce the computational time, we accept a slightly higher error ($\varepsilon = 0.015$) for the p-value. This decreases the number of iterations to 1,000 without changing the results of the power-law hypothesis.

Finally, we need to make a decision about the power-law hypothesis. A large p-value (close to 1) means that the difference between the empirical time series and hypothesised model could be attributed to statistical fluctuations alone. However, even when the data are drawn from a power-law distribution, it is extremely unlikely that the observed distribution exactly follows the power-law form, because of the random nature of the sampling process. Clauset *et al.* (2009) state that $p \leq 0.1$ is a relatively conservative choice to rule out a power-law distribution. We adopt this conservative decision criterion to make a decision whether tail returns of credit default swaps follow a power-law tail distribution.

It is important to understand that a large p-value does not unquestionably mean that the power-law is the correct distribution for the tail returns. It is much more likely that a small number of tail returns follow a power-law distribution, and consequently, the p-value is large, even when the power-law is the incorrect model for the empirical returns. This is not a shortcoming of the methodology, but is attributable to the fact that it is harder to rule out

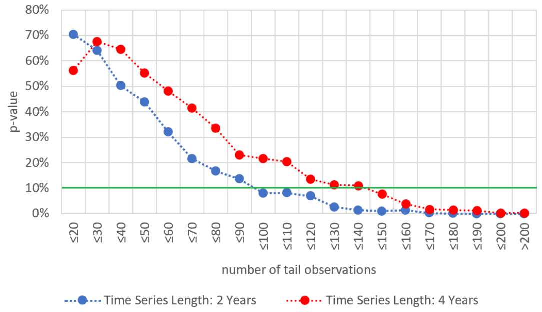

(a) The p-value as a function of tail observations for dynamic time windows of 2 and 4 years.

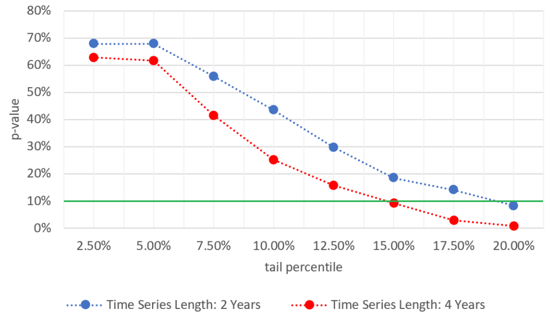

(b) The p-value as a function of the tail percentile for dynamic time windows of 2 and 4 years.

Fig. 1.2 shows the p-values for different bins of tail observations and tail percentiles. Shorter tails exhibit higher p-values compared to longer tails. Given a certain bin size, we observe that longer time series have higher p-values compared to shorter time series. Given a certain tail percentile, we illustrate that shorter time series have higher p-values compared to longer time series.

the power-law hypothesis (or any other model) in the presence of short tails.

Next, we illustrate the relationship between the p-values and the number of tail observations and tail percentiles. Figure 1.2 shows the average p-values as a function of the number of tail returns and size of tail percentiles for dynamic time windows of two (*blue*) and four (*red*) trading years of sovereign credit default swap returns from 2005 to 2016. We observe that a small number of tail observations and lower tail percentiles generally have higher p-values.

Figure 1.2a reports the average p-values for 20 non-overlapping, consecutive bins of size 10 for the number of estimated tail observations using the Kolmogorov-Smirnov method. The first label on the horizontal axis in Figure 1.2a, here ≤ 20, indicates the upper bound of the first bin, ranging from 11 to 20. The second bin, ranging from 21 to 30, has the label "≤ 30", and so on. Given a particular bin, we observe that the dynamic two-year time window (shorter time series) has mostly lower p-values compared to the four-year time window. This may be to due to the fact that the more extended time series have a higher probability of finding k_t-many tail observations, which fit the power-law model well for a given tail bin. There is evidence that for a given number of tail observations, the p-value increases with the increasing length of the time series.

Figure 1.2b reports the average p-values for non-overlapping, consecutive tail percentiles estimated by the Kolmogorov-Smirnov method. The first label on the horizontal axis in Figure 1.2b is "$\leq 2.5\%$", which indicates the upper bound of the first tail percentile ranging from 0% to 2.5%. The second tail percentile, ranging from 2.5% to 5.0%, carries the label "$\leq 5.0\%$", and so on. Given a certain tail percentile, we illustrate that shorter time series have higher *p*-values compared to longer time series. We observe this pattern because, for a specific tail percentile, the tail with the lower number of tail observations generally has a higher probability of being drawn from a power-law distribution.

It is important to note, that, since we fit the power-law form to only the part of the distribution above the tail threshold (m_t), the value of (m_t) effectively controls how many returns will be included in the tail estimation. As the tail percentile becomes large, above 20%, the *p*-values for the two dynamic time windows fall below the 10% decision criteria indicated by the green horizontal line in Figure 1.2b. This appears intuitively plausible, because when the tail threshold becomes smaller, a larger fraction of the total observations is considered for the tail analysis. However, a large number of observations cannot be considered as tail events anymore. Observations from the body of the distributions might behave much more as ordinary fluctuations without power-law decay, which consequently result in lower *p*-values.

The power-law hypothesis test fundamentally depends on the correct estimation of the tail length and percentile. Over- or underestimations of the correct tail length (percentile) may lead to false conclusions about the plausibility of the power-law model. To demonstrate the effect on the results of the power-law hypothesis test, we compare the *p*-values of arbitrary with true tail percentiles. For simplicity, we use a static time window with a length of 2175 observations from 2008 to 2016. We estimate the tail exponents (positive tail exponents) and the corresponding *p*-values (2,500 iterations) for seven countries using daily log-returns of sovereign credit default swaps with a one-year tenor. Table 1.1 contains information about the relative length of the tail $k_t(\%)$, the number of tail observations are stated in (), and the *p*-values comparing the three percentiles: tails determined by Kolmogorov-Smirnov method, and fixed percentiles of 2.5% and 5.0%. We denote the 2.5% and 5.0% arbitrary percentile as $\rho_{2.5\%}$ and $\rho_{5.0\%}$, and the tail percentile determined by the Kolmogorov-Smirnov method as ρ_{KS}. The decision criteria for the power-law hypothesis is $p > 0.10$.

Evaluating the findings of Table 1.1, we observe some interesting facts about the tail percentiles and the resulting *p*-values. Firstly, the power-law is a plausible fit for positive tail returns determined by the Kolmogorov–Smirnov distance method for all countries from 2008 to 2016 (Table 1.1). The tail lengths are sufficiently large to estimate the tail statistics. Germany, Italy and Portugal exhibit *p*-values above 90%, which means that the differences between the empirical data and the hypothesised model can be merely attributed to random effects caused by the sampling process.

Secondly, the power-law hypothesis is usually rejected when the arbitrary percentile is significantly larger than the tail percentile estimated by the Kolmogorov-Smirnov method (overestimation). An arbitrary percentile larger than true percentile encompasses returns from the centre of the distribution which usually exhibit a different behaviour to a power-law. Including returns from the body of distribution into the tail analysis, leads to misspecifications of the tail exponents, the synthetic distances ($D_{t,i}^{(s)}$) and consequently the p_t-values. These misspecifications are more profound for shorter tails. For example, an arbitrary percentile of $\rho_{5.0\%}$ significantly overestimates the correct percentile for Germany, Ireland, Portugal and Spain. As a result, the p-values are close to zero for those countries. Table 1.1 provides also evidence that even small overestimations (<5%) result in flawed p-values. This effect is particularly noticeable for France, Ireland and Italy. In the case of France, an arbitrary percentile of 5% misspecifies the true tail by only five tail observations (overestimation of 4.81%), which decreases the p-value from 89.44% to 0.52%. In the case of Ireland, using an arbitrary percentile of 2.5% includes only two returns from the body of the distribution, which causes a sharp decline of the p-value from 73.12% to 0.24%. Note that due to the relatively short tail length, a small deviation of only 3.77% from the true percentile already leads to a profound estimation error of the tail exponent and thus explain the low p-value. On the other hand, in the case of Italy, the overestimation of tail percentile is less significant due to a longer tail. The deviation from the true tail percentile is only one observation (less than 1%). Consequently, the p-value of the arbitrary percentile is lower compared to the Kolmogorov-Smirnov method. As the misspecification is almost neglectable for longer tails,

Country	KS-Method		2.5% Percentile		5.0% Percentile	
	tail length (k)	p-value	tail length (k)	p-value	tail length (k)	p-value
Belgium	5.15% (112)	0.7184	2.50% (55)	0.0008	5.00% (109)	0.6904
France	4.78% (104)	0.8944	2.50% (55)	0.0048	5.00% (109)	0.0052
Germany	3.82% (83)	0.9404	2.50% (55)	0.5920	5.00% (109)	0.0008
Italy	4.97% (108)	0.9176	2.50% (55)	0.0280	5.00% (109)	0.6500
Ireland	2.44% (53)	0.7312	2.50% (55)	0.0024	5.00% (109)	0.0000
Spain	3.59% (78)	0.7648	2.50% (55)	0.0020	5.00% (109)	0.0000
Portugal	2.71% (59)	0.9720	2.50% (55)	0.7580	5.00% (109)	0.0012

Table 1.1 shows the tail percentiles in %, number of tail observations in () and the corresponding p-values for three tests performed on sovereign credit default swap time series with maturity of one year from 2008 to 2016. Using arbitary percentiles to determine the tail exponent, not just measures the credit tail risk inaccurately, it also results in imprecise p-values, which may lead to misjudging the power-law hypothesis.

the power-law assumption is not ruled out for both methods.

Thirdly, we find that the power-law hypothesis is interpreted correctly in cases where the tail percentiles estimated by the Kolmogorov-Smirnov method are similar or only slightly larger than the arbitrary tail percentile (underestimation), and both methods hold similar tail exponents. Let us assume the case, where a lower percentile is a subset of the larger percentile. For example, the true percentile is 2.75% and the arbitrary percentile is 2.5%. Both percentiles have similar tail exponents and standard errors. As a result, both percentiles lead to the correct interpretation of the power-law hypothesis. We have a similar case for Belgium and Portugal. The true percentiles are slightly larger than the arbitrary percentiles. The tail exponents and standard errors are similar under both methods. As a result, the lower arbitrary percentiles hold the correct interpretation of the power-law hypothesis, however with lower p-values compared to the true estimates.

In this section, we introduced a quantitative framework to assess the power-law model. We retrieve accurate p-values ($\pm\varepsilon = 0.015$) for a large number of simulations (above 1,000 iterations). For credit default swap returns, we defined conservative decision criteria. The power-law hypothesis is not ruled out for p-values larger than 0.10. Finally, we elaborate on the importance of the correct specification of the tail length. Over- and underestimations of the actual tail percentile most likely results in flawed p-values. The p-values are highly sensitive to misspecification errors, especially in the presence of a small number of tail observations. The power-law hypothesis only holds in exceptional cases when the tail percentile is misspecified.

1.3.4 Model Comparision

The methodology outlined in the previous section provides a reliable framework to test whether credit default swap returns are plausibly drawn from a power-law distribution. However, even if the tail returns fit well to a power-law distribution, it does not rule out the possibility that another distribution might provide a better fit. Therefore, we compare the power-law distribution with other heavy-tailed distributions, such as the exponential and log-normal distribution. Well-known methods for model comparisons are the cross-validation criterion (Stone, 1974), the Bayes factors (Kass and Raftery, 1995), the minimum description length approach (Grünwald, 2007) and the likelihood ratio test (Vuong, 1989). We choose the likelihood ratio test because it can tell which distribution is the better fit if neither of the two candidate distributions is ruled out under the Kolmogorov-Smirnov test as a potential fit to the data.

The fundamental idea of the likelihood ratio test is to calculate the likelihood of the tail returns under competing distributions. Consider two different types of distributions with probability density functions $p_{1,t}(x_t)$ and $p_{2,t}(x_t)$ at time t. The likelihoods of a given data set is

$$L_{1,t} = \prod_{i=1}^{k_t} p_{1,t}(x_{i,t}), \quad L_{2,t} = \prod_{i=1}^{k_t} p_{2,t}(x_{i,t}), \tag{1.9}$$

where $L_{1,t}$ is the likelihood of the power-law distribution and $L_{2,t}$ is the likelihood of the competing distribution. The distribution with the higher likelihood value is consequently the better fit to the observations. By preference, we calculate the ratio of the two likelihoods $R_t = L_{1,t}/L_{2,t}$ and take the logs, such as

$$\mathscr{R}_t = \sum_{i=1}^{k_t} [\ln p_{1,t}(x_{i,t}) - \ln p_{2,t}(x_{i,t})] = \sum_{i=1}^{k_t} [\ell_{i,t}^{(1)} - \ell_{i,t}^{(2)}] \tag{1.10}$$

where $\mathscr{R}_t$ denotes the log-likelihood ratio and $\ell_{i,t}^{(j)} = \ln p_{j,t}(x_{i,t})$ can be seen as the log-likelihood for a single measurement $x_{i,t}$ within distribution j at time t. There are three possible outcomes for the log-likelihood ratio. In the case that the value is positive and sufficiently far from zero, the power-law distribution is the better fit to the tail returns. Vice versa, if the log-likelihood ratio is negative and statistically significant from zero, the competing distribution is the better fit. If the log-likelihood ratio is close to zero, the observed sign of $\mathscr{R}_t$ does not serve as a reliable indicator of which model is favoured.

The log-likelihood ratio, and so the observed sign of $\mathscr{R}_t$, does not definitively indicate which model is the better fit. To make a better quantitative judgment about whether the observed values of $\mathscr{R}_t$ is statistically significant from zero, we need to calculate the p_t-value of the log-likelihood ratio. To this end, we approximate the variance σ_t^2 on $\mathscr{R}_t$, which is defined as

$$\sigma_t^2 = \frac{1}{k_t} \sum_{i=1}^{k_t} \left[(\ell_{i,t}^{(1)} - \ell_{i,t}^{(2)}) - (\bar{\ell}_{i,t}^{(1)} - \bar{\ell}_{i,t}^{(2)}) \right]^2. \tag{1.11}$$

The variance is used to calculate the probability p_t. The probability p_t that the observed log likelihood ratio has a magnitude as large as or larger than the measured value of $|\mathscr{R}|$ is given by

$$p_t = \frac{1}{\sqrt{2\pi k_t \sigma^2}} \left[\int_{-\infty}^{-|\mathscr{R}|} e^{-t^2/2k_t\sigma^2} dt + \int_{|\mathscr{R}|}^{\infty} e^{-t^2/2k_t\sigma^2} dt \right]. \tag{1.12}$$

The p_t-value provides an estimate of the probability that we exhibit a given value of $\mathscr{R}_t$ when the true value of $\mathscr{R}_t$ is in fact close to zero. If the p_t-value is small (i.e. $p < 0.10$), then it is unlikely that log-likelihood ratio is a chance result, and so the observed sign is a trustworthy indicator of which model is the better fit. If p_t is large, then the likelihood ratio test is not statistically significant from zero and neither of the two distributions is superior.

As stated in Table 1.2, we find that the p >0.10, which suggests that the alternative distributions do not provide a better fit to the tail events of sovereign credit default swaps. Given that the p-values for the power-law hypothesis test are reasonably high, we conclude that the power-law distribution is a plausible model for our data.[14]

	Belgium	France	Germany	Italy	Ireland	Portugal	Spain
p|PL, 504	0.49	0.55	0.53	0.56	0.53	0.57	0.51
$p\|\mathscr{R}$(exp)	$0.42^{(+)}$	$0.48^{(+)}$	$0.48^{(+)}$	$0.53^{(+)}$	$0.46^{(+)}$	$0.46^{(-)}$	$0.52^{(+)}$
$p\|\mathscr{R}$(log)	$0.40^{(-)}$	$0.49^{(-)}$	$0.40^{(-)}$	$0.45^{(-)}$	$0.36^{(-)}$	$0.45^{(-)}$	$0.39^{(-)}$
p|PL, 1008	0.50	0.54	0.52	0.53	0.56	0.58	0.56
$p\|\mathscr{R}$ (exp)	$0.49^{(-)}$	$0.41^{(+)}$	$0.57^{(+)}$	$0.49^{(+)}$	$0.43^{(+)}$	$0.39^{(+)}$	$0.47^{(+)}$
$p\|\mathscr{R}$(log)	$0.38^{(-)}$	$0.49^{(-)}$	$0.38^{(-)}$	$0.37^{(-)}$	$0.41^{(-)}$	$0.50^{(-)}$	$0.45^{(-)}$

Table 1.2 reports the average p-values for the fit to the power-law model for each country and two lookback windows in rows 1 and 4. We also report the average p-values for the log-likelihood ratios $\mathscr{R}$ exponential (exp) and log-normal (log). These p-values for the log-likelihood ratios are not statistically significant. Positive values of the log-likelihood ratios, expressed by the sign (+), indicate that the power-law model is favoured over the alternative. Negative values of the log-likelihood ratios, expressed by the sign (-), indicate that the alternative is favoured over the power-law model. However, if the alternative model is favoured, the values of log-likelihood ratios are close to zero. This means that the observed result is purely the product of fluctuations, and it does not imply that the alternative distribution is truly a better fit.

Table 1.2 shows the p-values for the dynamic tail exponent assuming a power-law distribution. The power-law hypthesis is not ruled out for all countries. However, for some countries, the alternative distribution might be a better fit, which is marked by a change of sign (-) of the log-likelihood ratio. However, given that the differences of the log-likelihoods are close to zero, it is rather unlikely that the alternative distribution is truely a better fit to the data.

[14]Note, in general, small data sets should be treated with caution. Namely, it is difficult to rule out alternative fits to such data, even when they are truly power-law distributed, and conversely, the power-law form may appear to be a good fit even when the data are drawn from a non-power-law distribution.

1.4 Results

We consider sovereign credit default swap returns over various consecutive time windows and from different countries, maturities and regions in the European Monetary Union, to evaluate whether the power-law is a statistically plausible model, and ascertain whether the tail exponent supports the presence of a finite second moment. Secondly, we investigate the possible existence of universal scaling properties in credit tail risk for different time scales ranging from daily to monthly log returns. Thirdly, we explore the uniformity of tail behaviour in sovereign credit default swap markets in terms of symmetry between upper and lower tails across different countries, maturities and regions. Finally, we elaborate on the implications of these findings on risk management.

1.4.1 Power-Law Tails in Sovereign Credit Markets

In this section, we investigate whether the power-law is a statistically plausible model in sovereign credit markets. Therefore, we consider daily log-returns of sovereign credit default swaps of seven countries and five maturities from 2005 to 2016. To estimate the tail exponents over time, we use daily rolling time windows of two and four trading years. Shorter time windows may not hold enough tail observations for reliable estimations of the tail exponents whereas longer time windows may carry tail events from the distant past.[15] To make results comparable across different lengths of rolling windows, we use the same starting and end date and, consequently, the same number of time windows (1,916) between January 2009 and May 2016. Starting in January 2009, we calculate the tail statistics for both backward-looking time windows on each trading day. Firstly, we estimate the tail percentile by minimising the maximum distance between the power-law model and the observed quantile and estimate the maximum likelihood estimator for the dynamic tail exponent. Then, we conduct the power-law hypothesis test for the fitted model. During the period from 2009 to 2016, we estimate 22,992 tail exponents per country and maturity for positive, negative, absolute and normalised returns.[16] To assess the power-law hypothesis, we perform over 804 million

[15]We calculated the number of tail observations for a time window of one year using the Kolmogorov-Smirnov method on 133,200 samples. The average tail length for negative and positive tails is 32 and 33 observations. This means that the negative and positive tails are too short for extracting reliable parameter estimates for a rolling one-year time window.

[16]The number of tail exponents is calculated as follows. We estimate 1,916 tail exponents between January 2009 and May 2016, for positive, negative and absolute tail returns, two rolling windows, and normalised and unnormalised returns. The product is 22,992 tail exponents for each country and maturity.

a) Power-law statistics for a two-year rolling time window:
average p-value and tail length for positive tails from 2009 to 2016

	Belgium		France		Germany		Italy		Ireland		Spain		Portugal	
Maturity	*p*	*k*	*p*	*k*	*p*	*k*	*p*	*k*	*p*	*k*	*p*	*k*	*p*	*k*
01Y	0.30	52	0.53	47	0.43	48	0.54	52	0.42	52	0.41	64	0.40	47
03Y	0.33	56	0.30	52	0.54	33	0.41	48	0.30	61	0.37	53	0.44	43
05Y	0.34	52	0.30	52	0.37	42	0.50	47	0.35	50	0.33	57	0.47	47
10Y	0.20	58	0.37	37	0.31	59	0.43	46	0.51	50	0.37	50	0.64	38
30Y	0.31	52	0.51	39	0.23	75	0.53	44	0.42	57	0.32	57	0.28	66
Average	0.30	54	0.40	46	0.38	51	0.48	47	0.40	54	0.36	56	0.45	48

b) Power-law statistics for a four-year rolling time window:
average p-value and tail length for positive tails from 2009 to 2016

	Belgium		France		Germany		Italy		Ireland		Spain		Portugal	
Maturity	*p*	*k*	*p*	*k*	*p*	*k*	*p*	*k*	*p*	*k*	*p*	*k*	*p*	*k*
01Y	0.23	88	0.50	78	0.39	75	0.55	72	0.32	99	0.33	108	0.36	61
03Y	0.43	71	0.36	62	0.49	66	0.47	63	0.30	88	0.45	66	0.55	48
05Y	0.38	74	0.40	57	0.25	73	0.35	67	0.47	69	0.25	83	0.36	97
10Y	0.25	83	0.34	44	0.19	84	0.36	64	0.61	88	0.50	52	0.66	60
30Y	0.13	104	0.50	54	0.22	102	0.46	72	0.44	79	0.36	74	0.43	70
Average	0.28	84	0.42	59	0.31	80	0.44	68	0.43	85	0.38	77	0.47	67

Table 1.3 shows the power-law statistics for (a) a two-year and (b) a four-year rolling time window. We find that the average p-value for the power-law is reasonably large ($p > 0.10$) for all countries and maturities for both backward looking time windows from 2009 until 2016. Hence, the power-law hypothesis is not ruled out. For this time span (1,916 days), the reported p-values are the average of the daily p_t values. The tail length varies over time and for each country and maturity. The reported tail length k is the average of the daily tail length (k_t). The average tail length is between 12% - 15% of the sample size for the two-year moving time window and between 7% - 12% for the four-year moving time window. The key result of this investigation is that the tail returns of credit default swaps are well fit by a power law, where the average tails are large enough for reliable estimations of the tail exponent.

simulations to calculate the p-values.[17] After these extensive computations, we conclude that the power-law distribution is a plausible model for tail returns of sovereign credit default swaps.

Table 1.3(a) and (b) state the average length of the positive tails and average p-values for each country and maturity, for a two- and four-year rolling time window from 2009 to 2016. Both time windows hold on average, a sufficiently large number of tail observations for reliable estimations of the tail exponents. The average p-values suggest that the power-law

[17] The number of simulations is calculated as follows. We consider 804,720 tail exponents for seven countries and five maturities. Given a 1,000 iterations per combination, we calculate 804,720,000 p-values. This is computationally expensive, taking months of parallel computing on six high-performance computers.

hypothesis is not ruled out for sovereign credit default swaps, independently of the time window.[18] The average p-values for positive tails are similar for both time windows (41% for a two-year and 40% for a four-year rolling window). Furthermore, the negative tail returns are plausibly drawn from power-law distributions as negative tails possess average p-values of 45% and 47% for both time windows. For the two-year rolling time window, the lowest and highest average p-values of the tail exponents have Belgium with 10-year maturity (20%) and Portugal with 10-year maturity (64%). The average p-values across the five maturities for a country ranges from 30% for Belgium to 48% for Italy. For the four-year rolling time window, the smallest and highest average p-value for the tail exponent has Belgium with 30-year maturity (13%) and Portugal with 10-year maturity (67%). The average p-values across the five maturities for a country ranges from 28% for Belgium to 47% for Portugal. Finally, we compute the log-likelihood ratio and corresponding p-values for competing heavy-tailed models. We find that alternative heavy-tailed distributions, such as the log-normal and exponential distribution, are unlikely a better fit. We conclude that the power-law distribution is truly a plausible model for positive and negative tail returns. These results also hold for absolute returns and normalised returns for both rolling windows over the same period.

1.4.2 The Inverse Power-Law and Lévy-Stable Distributions

In this section, we discuss the dynamics of the tail exponent in sovereign credit default swap markets across countries, maturities and regions. Motivated by research in equity and foreign exchange markets, we investigate whether tail exponents in credit markets share some common statistical characteristic. A series of studies in equity markets report that price fluctuations are distributed accordingly to a power-law with tail exponent $\alpha \simeq 3$ (Gopikrishnan *et al.*, 1998, 1999; Plerou *et al.*, 1999; Gopikrishnan *et al.*, 2000; Cont, 2001; Gabaix *et al.*, 2007; Gabaix, 2009).[19] To examine whether the inverse cubic law holds in credit markets, we consider daily log-returns of sovereign credit default swaps of seven countries and five maturities. Similar to the previous section, we conduct this study on two rolling time windows of two- and four-years from January 2009 to May 2016. In this period, we estimate

[18] Some researchers might be familiar with the use of p-values to confirm rather than rule out hypotheses for experimental data. In the latter case, one quotes a p-value for a null model, a model other than the model the experiment is attempting to verify. Usually, one then considers low values of p to be "good", since they indicate that the null hypothesis is unlikely to be correct. In our study, by contrast, we use the p-value as a measure of the hypothesis we are trying to verify, and hence high values are considered as "good" (Clauset *et al.*, 2009).

[19] The respective probability density functions tails thus decay with $\alpha + 1 \simeq 4$.

1,916 (9,580) tail exponent for each individual time series (country). However, we limit our analysis to distributions with at least 40 tail observations (minimum tail length $k_t = 40$) in each time window. Shorter tails observe higher standard errors, which may distort the results. We calculate the so called "average dynamic credit tail risk exponent" by averagering the exponent for all estimates with sufficiently long tails. Table 1.4 presents the main results of the average dynamic credit tail risk exponent discussed in the following paragraphs.

Firstly, we find that the average of the dynamic credit tail risk exponent across all countries and maturities obeys the inverse cubic law for the two-year time window, such that

$$p(x_{i,t} > m_t) = \left(\frac{x_{i,t}}{m_t}\right)^{-\alpha_t} \text{ with } \alpha = 2.99. \tag{1.13}$$

We find that the following distribution of the tail exponent: $\alpha < 1.25$=0 observations, $1.25 \leq \alpha < 1.75$=140 observations (0.4%), $1.75 \leq \alpha < 2.25$=1,386 observations (4.0%), $2.25 \leq \alpha < 2.75$=10,402 observations (29.8%), $2.75 \leq \alpha < 3.25$=12,636 observations (36.2%), $3.25 \leq \alpha < 3.75$=7,368 observations (21.1%), $3.75 \leq \alpha < 4.25$=2,494 observations (7.1%), $4.25 \leq \alpha < 4.75$=478 observations (1.4%), and $\alpha > 4.75$=5 observations. The most probable value for the credit tail risk exponent is $\alpha_{\text{MP}} = 3$ for a rolling two-year window. The inverse cubic law is also found for using a four-year rolling time window. The average dynamic credit tail risk exponent across all time series is $\alpha = 3.27$ for a rolling four-year window. While the credit tail risk exponent only considers the right tail of the distribution, we also confirm the inverse cubic law for the left tails for both time windows ($\alpha = 3.20$ for a two-year and $\alpha = 3.48$ for a four-year rolling window).

Furthermore, we calculate the average of dynamic credit tail risk exponent for a specific country. The column "Average" in Table 1.4(a) reports the average of dynamic credit tail risk exponent across the five maturities using a two-year rolling window. We find that France and Portugal almost perfectly obey the inverse cubic law with average dynamic credit tail risk exponents of 2.99 and 3.01. All other countries closely fluctuate around the exponent $\alpha \simeq 3$, whereas Belgium marks the lower limit and Italy the upper limit of the range $\alpha = [2.85; 3.24]$. The average credit tail exponents for all countries lie within the range of one standard error, which supports the robustness of the inverse cubic law. The column "Average" in Table 1.4(b) reports similar results for a four-year rolling window. The average of the dynamic credit tail exponents for most countries lie within the range of one standard error, which supports the robustness of the inverse cubic law independently of the rolling window.

Secondly, we observe that inverse cubic law holds for different maturities of sovereign credit default swaps. For each maturity, we calculate the averages of the dynamic credit tail risk exponents across different countries. We find that average credit tail risk exponents for different maturities closely fluctuate around the tail exponent of $\alpha \simeq 3$. The results reveal that the average tail risk exponents increase in maturity from 1- to 10-year credit default swap contracts. Table 1.4(a) shows that average dynamic credit tail risk exponents is 2.81, 2.95, 3.10 and 3.16 for 1-, 3-, 5- and 10-year tenors for a two-year window. Credit insurances with the shortest time to maturity of the contract have the thickest tails in the probability distributions. Credit default swap contracts with a 1-year tenor have on average the highest credit tail risk or highest probability of extreme events ($\alpha = 2.81$), whereas the 10-year contracts possess the lowest credit tail risk ($\alpha = 3.16$) for a two-year rolling window. Interestingly, the average dynamic credit tail risk exponent ($\alpha = 2.92$) is slightly lower for contracts with 30-years to maturity. This hump-shaped pattern is consistent for different lengths of rolling time windows.

Thirdly, we investigate the dynamic credit tail risk exponent across different regions. The core region includes Belgium, France and Germany, which are characterised by lower credit default swap prices.[20] The countries within the peripheral region are Italy, Ireland, Portugal and Spain. Credit default swaps prices for peripheral countries are usually more expensive and exhibit a high volatility. We normalise returns by the average volatility for the period from 2005 to 2016, so that the normalised distributions all have a standard deviation of 1. We find that the average tail exponents for the core region is 2.87 and for the peripheral region is 3.08. Interestingly, core countries with a lower probability of default have a higher credit tail risk compared to peripheral countries. The credit tail risk exponents for both regions have exponents close to the inverse cubic law.

Finally, we investigate the corresponding equity tail risk for those seven countries. To estimate the tail exponent, we use negative tail returns of national stock market indices from 2009 to 2016.[21] The last column of Table 1.4(a) reports that across seven indices, the average dynamic equity tail risk exponents is $\alpha = 3.31$ for a two-year rolling time window. Our

[20] We state the average cross-maturity price from 2005 to 2016 in increasing order: Germany is 29 bps (24), France is 55 bps (47), Belgium is 73 bps (70), Italy is 131 bps (125), Spain is 140 bps (128), Ireland is 201 bps (235), Portugal is 270 bps (331).

[21] We estimate the equity tail risk (left tail exponent) for the seven European countries using the following equity indices: Euronext Brussels (BEL 20) for Belgium, Cotation Assistée en Continu (CAC 40) for France, Deutscher Aktien Index (DAX 30) for Germany, Milano Indice di Borsa (FTSE MIB 40) for Italy, Irish Stock Exchange Quotient (ISEQ All) for Ireland, Índice Bursátil Español (IBEX 35) for Spain, and Portuguese Stock Index (PSI 20) for Portugal. The power-law hypothesis is not ruled out for all countries, and the tails are

a) Credit and equity tail risk for a two-year rolling time window

	Credit Default Swaps						Equity Index
	01Y	03Y	05Y	10Y	30Y	Average	
Belgium	2.69	2.81	2.96	2.90	2.88	2.85	3.29
France	2.69	2.87	2.96	3.51	2.95	2.99	3.18
Germany	2.64	2.84	3.16	2.64	2.50	2.76	3.36
Italy	2.82	3.25	3.26	3.35	3.49	3.24	3.09
Ireland	2.66	2.78	3.07	3.20	2.76	2.89	3.48
Spain	3.10	3.03	3.16	3.41	3.23	3.19	3.47
Portugal	3.05	3.04	3.12	3.15	2.67	3.01	3.29
Average	2.81	2.95	3.10	3.16	2.92	2.99	3.31

b) Credit and equity tail risk for a four-year rolling time window

	Credit Default Swaps						Equity Index
	01Y	03Y	05Y	10Y	30Y	Average	
Belgium	2.69	3.63	3.34	3.59	2.96	3.24	3.80
France	2.88	3.60	3.49	3.73	3.30	3.40	3.57
Germany	2.66	3.26	3.12	2.69	2.84	2.91	3.61
Italy	3.16	3.70	3.65	3.79	3.90	3.64	3.97
Ireland	2.48	2.86	3.33	3.34	3.16	3.04	3.68
Spain	3.00	3.23	3.20	3.84	3.84	3.42	3.74
Portugal	3.09	3.38	3.14	3.37	3.24	3.24	3.74
Average	2.85	3.38	3.32	3.48	3.32	3.27	3.73

Table 1.4 reports the average credit tail risk and equity tail risk from 2009 to 2016. Extreme events in sovereign credit default swaps are consistent with the so-called inverse cubic law, where the tail exponent $\alpha \simeq 3$. This statistical characteristic holds for both rolling time windows in credit markets. Furthermore, we state the corresponding equity tail risk for national stock market indices. We confirm previous findings that tail returns in equity markets obey the inverse cubic law (for the two year rolling time window). Note that the credit tail risk implied in sovereign credit default swaps is higher compared to the equity tail risk independently of the time window.

results are in line with previous findings that tail returns of international stock market indices follow a power-law with tail exponent $\alpha \simeq 3$ (Gopikrishnan *et al.*, 1999). Furthermore, we find that on average sovereign credit default swaps tend to have a higher probability of extreme events compared to the corresponding equity index. In other words, the average credit tail risk implied in credit default swaps is almost always higher compared to the corresponding equity tail risk independently of the time window.

These findings have important implications for risk management and derivative pricing. One implication of the inverse cubic law is that there are many more tail events than would occur if the underlying distribution were Gaussian. The standard deviation of daily log-returns

sufficiently large for stable tail exponent estimations.

of sovereign credit default swaps is approximately 2% from 2005 to 2016. A 10-standard deviations event is a day in which the credit default swap price moves by at least 20%. From our data set, the reader can see that those price flucuations are not rare. Between 2005 and 2016, essentially every month, a 10-standard deviations event occurs for one of the 35 time series in the sovereign credit market. The cubic law quantifies that notion. Under the inverse cubic law, the chances of a 10 standard deviation event is, respectively, $5^3 = 125$ times less likely than a two standard deviation event, whereas if the distribution of returns was Gaussian, the chances of a 10 standard deviation event would be 10^{22} times less likely than a two standard deviation event.[22] A second implication is that the inverse cubic law implies the existence of a finite second moment, because the estimated power-law exponents are outside the Lévy-stable region $0 < \alpha < 2$. As the particular value $\alpha \simeq 3$ is consistent with a finite variance, and it means that credit default swaps returns are not Lévy distributed. A Lévy distribution is either Gaussian, or has infinite variance, α<2. The existence of the variance is a key foundation for risk estimation and portfolio optimisation where covariance–based methods are primarily used.

1.4.3 Credit Tail Risk Across Different Time Scales

In this section, we investigate whether credit tail risk exponents exhibit similar dynamics on different time scales. The time scale measures the distance between two observations and is denoted as δt. We calculate log-returns of sovereign credit default swaps for time scales ranging from daily ($\delta t = 1$ day), over weekly ($\delta t = 5$ days), to monthly ($\delta t = 22$ days). In total, we consider seven time scales $\delta t = [1, 2, 3, 4, 5, 10, 22]$ days. Section 1.4.2 reports that the distributions of daily tail returns decay as a power-law with an tail exponent $\alpha \simeq 3$ for different countries, maturities and regions. Motivated by previous findings in equity markets (Lux, 1996; Gopikrishnan *et al.*, 1999; Plerou *et al.*, 1999), we explore whether the distribution retains its power-law functional form for longer time scales in credit markets. Similar to our previous studies, we estimate the credit tail risk exponent for 35 individual time series of seven countries and five maturities using two rolling windows. Secondly, we investigate the relationship between time scales and maturities of credit default swaps. To account for different levels of volatility, we also perform this study for normalised log-returns. To ensure comparability of the asymptotic behaviour on different time scales, we exclude time series with less than 40 tail observations. Figure 1.3 shows the average

[22] The formula is: $(\text{standard deviation}/2)^{\alpha}$.

credit tail risk exponents and standard errors based on 67,060 estimates for each time scale.[23] Table 1.5 reports the credit tail risk exponents for different time scales and maturities. In the following, we discuss these results in more detail and its implications on risk management.

Firstly, we find that inverse cubic law partially holds true for different time scales. The average of the dynamic credit tail risk exponent across all countries and maturities obeys the inverse cubic law for short time scales $\delta t = [1;5]$ for two-year time windows. Figure 1.3(a) shows that the average credit tail risk exponents across countries and maturites lie within a close range of $3 < \alpha < 3.3$ for short time scales. Figure 1.3(b) shows that these findings are consistent by accounting for different levels of volatility. In both cases, the distributions retain its power-law functional form (the inverse power law) for positive tail returns within +/- 1 standard error. For longer time scales, such as bi-monthly or monthly returns, the distributions slowly converge to Gaussian behaviour. Our results in credit markets are in line with findings in equity markets. Gopikrishnan *et al.* (1999) find the distribution of S&P 500 index returns to be consistent with a non-stable power-law functional form ($\alpha \simeq 3$) until four days ($\delta t \leq 4$ days), after which an onset of convergence to Gaussian behaviour is exhibited. Figure 1.3(c) and (d) shows the empirical cumulative distribution of log-returns only decays with an average tail exponent of $\alpha \simeq 3$ for $\delta t = 1$ days. Interestingly, the inverse cubic law does not hold beyond $\delta t = 1$ for four year rolling windows . For all other time scales $\delta t = [2;22]$, the average dynamic credit tail risk exponents lie outside of +/- 1 standard error range. However, the distributions of returns retain a similar power-law functional form ($\alpha = [3.50;4.00]$) from $\delta t = 2$ days until approximately $\delta t = 10$ days, after which an onset of convergence to Gaussian behaviour is found. Overall, we conclude that credit tail risk (exponent) decreases (increases) with increasing time scales. The inverse cubic law only holds for short time scales and time windows. For longer time scales, such as $\delta t = 10$ and $\delta t = 22$, the inverse cubic law does not hold independently of the time window, and show convergence to Gaussian behaviour.

Secondly, we find that that credit default swap sellers, also known as protection sellers, bear the highest credit tail risk of short-dated credit default swaps on shorter time scales. For example, on a daily time scale ($\delta t = 1$), Table 1.5 reports that credit default swaps with a one-year tenor have an average credit tail risk exponent across all countries of $\alpha = 2.81$. On

[23]The average credit tail risk exponent for a time scale is calculated as follows: seven countries (7x) multiplied by five maturities (5x) multiplied by 1,916 time-varying tail exponents from 2009 to 2016. In total, we calculate 1,877,680 credit tail risk exponents for (un-) normalised log-returns (2x), two rolling time windows (2x) and seven time scales (7x). To compute these results, we use six high-performance computers, parallel computing, and require approximately two months of runtime.

Average dynamic credit tail exponent of credit default swap returns for time scales from one day to one month from 2009 to 2016

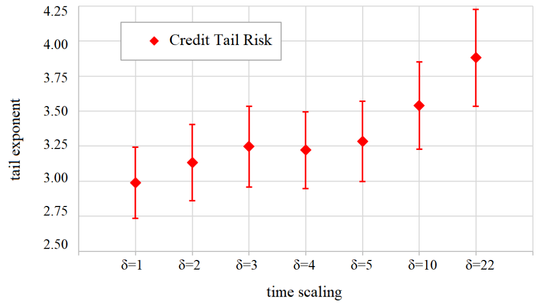

(a) Unnormalised tail returns (two years)

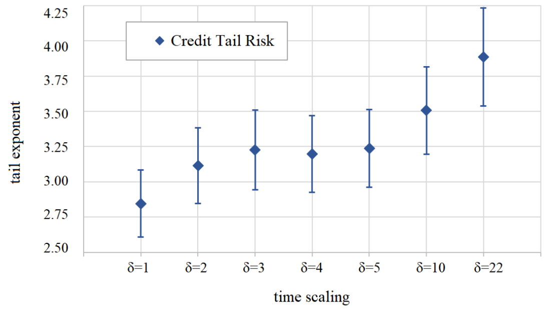

(b) Normalised tail returns (two years)

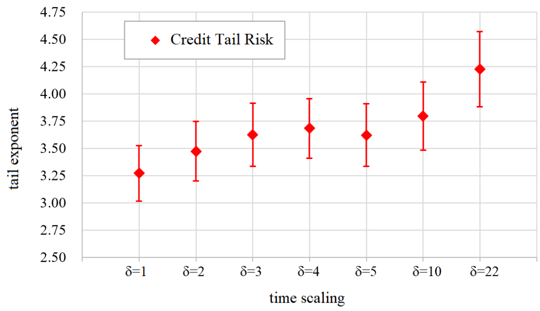

(c) Unnormalised tail returns (four years)

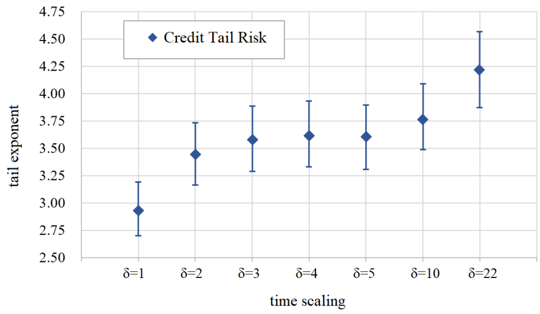

(d) Normalised tail returns (four years)

Fig. 1.3 shows the average tail risk exponents for positive tails of sovereign credit default swap returns over seven time scales for log-returns and normalised log returns. The bars indicate +/- 1 standard error of the tail risk exponent. The average tail risk exponent and standard error for a specific time scale are calculated based on 67,060 tail exponents of seven countries and five maturities from January 2016 until May 2016. The top row states the results for a rolling backwards-looking time window of two years. We find that the distribution of returns for various choices of δ, ranging from 1 day ($\delta = 1$) to 5 days ($\delta = 5$) have a similar functional form. The average tail risk exponents for positive tail returns suggest that the inverse cubic law holds for time scales up to a week ($\delta = 5$). For larger time scales ($\delta = 10$ and $\delta = 22$) our results are consistent with the break-down of the scaling behaviour, i.e. convergence to Gaussian. The bottom row shows the tail risk exponents for a four-year time window. The inverse cubic law only holds for daily log-returns ($\delta = 1$). For larger time scales ($\delta = 10$ and $\delta = 22$), the inverse cubic law does not hold independently of the time window. Overall, we conclude that tail exponent decreases with increasing time scales. This implies that the size of extreme returns decreases as a function of time, and a larger tail index might result from smaller price fluctuations on longer time horizons. Vice versa, the smaller tail indexes on shorter time scales indicate a larger probability of extreme credit default swap returns.

The average dynamic credit tail risk exponent on different time scales for five maturities

	Two-Year Time Window						Four-Year Time Window					
	Credit Default Swap Maturity						Credit Default Swap Maturity					
	01Y	03Y	05Y	10Y	30Y	Avg.	01Y	03Y	05Y	10Y	30Y	Avg.
δt=1	2.81	2.95	3.10	3.16	2.93	2.99	2.85	3.38	3.32	3.48	3.32	3.27
δt=5	3.07	3.25	3.33	3.39	3.39	3.28	3.10	3.42	3.73	3.91	3.95	3.62
δt=10	3.38	3.46	3.48	3.64	3.73	3.54	3.54	3.56	3.57	4.01	4.29	3.80
δt=22	3.65	3.86	3.82	3.85	4.22	3.88	4.02	4.63	4.12	4.10	4.26	4.23
Global	3.23	3.38	3.43	3.51	3.57	3.42	3.38	3.75	3.69	3.87	3.95	3.73

Table 1.5 shows the average sovereign credit tail risk on different time scales and maturities for two- and four-year time windows from 2009 to 2016. Reading the results horizontally, credit tail risk decreases with increasing credit default swap maturity. This implies that long-dated credit default swaps usually have a lower credit tail risk. Reading the results vertically states that the credit tail risk decreases with increasing time scale. Both results together suggest that sellers of short-dated credit default swaps bear the highest credit tail risk on shorter time scales (2.81), while longer-dated credit default swaps bear the lowest credit tail risk on longer time scales (4.22) for the two-year rolling time window. These results hold independently of the rolling time window.

the other hand, we observe that protection sellers of longer-dated credit default swaps bear a lower credit tail risk (higher average credit tail risk exponent) on longer time scales. For example, Table 1.5 reports that on a monthly time scale ($\delta t = 22$ days), credit default swaps with a 30-year tenor have an average credit tail risk exponent across all countries of $\alpha = 4.22$. For the rolling two-year time window, we also find that the credit tail risk decreases with increasing time scale δt. Table 1.5 reports a monotonically increasing pattern between average credit tail risk exponents and time scales. This pattern is consistent among all maturities. Our findings might appear to be a troublesome anomaly for rational expectations. Accordingly to conventional wisdom, the chances of a rare event happening over a long time period seem to be much higher, which should be reflected in the tail risk exponent of long-dated insurance contracts. However, the credit tail risk exponents do not reflect these rational expectations. Therefore, we propose some possible non-quantitative explanations for this anomaly.

Investors may anticipate the impact of potential extreme events on asset prices in the near future, but may experience difficulties in quantifying rare (or unknown) events in the distant future. In case of sovereign distress, investors can consider extreme events in risk modelling using past experiences. For example, over 20 countries experienced financial distress, or defaulted on domestic or external debt between 2000 and 2016.[24] However,

[24]List of sovereign debt crises between 2000 and 2009: Ecuador (2000), Zimbabwe (2000), Morocco (2000), Kenya (2000), Côte d'Ivoire (2000), Surinam (2001–02), Argentina (2001), Nigeria (2001), Madagascar (2002), Myanmar (2002), Paraguay (2003), Dominica (2003–05), Venezuela (2004), Grenada (2004–05), Nigeria (2004), Cameroon (2004), Dominican Republic (2005), Argentina (2005–16), Zimbabwe (2006), Ecuador (2008), Côte d'Ivoire (2011), Greece (2012), Argentina (2014), and Greece (2015).

market participants may experience difficulties in predicting extreme events in the distant future. For example, the risk of a global health crisis and its negative repercussions for the global economy is nearly impossible to predict and quantify in advance for the far future. The inability to quantify (unidentified) risks in the distant future might be an explanation of why tail risk is not higher for long-dated credit default swaps. However, this argument does not justify the higher credit tail risk implied in short-dated credit default swaps as a near-term default simultaneously triggers the insurance mechanism of long-dated credit default swaps.

A possible explanation for higher tail risk (lower exponent) for short-dated credit default swaps is that financial shocks or distress can be seen as a temporary problem which requires near-term solutions. Consequently, when a short-term solution is likely to be found, it might not impact the long-term credit tail risk of a country. Furthermore, it can be considered as unlikely that a government remains a "close to default situation" for decades. Either a country defaults on its obligations, restructures its debt or manages a turn-around. This may provide a logical argument for a lower tail risk implied in long-dated credit insurance contracts. Comparing shorter- with longer-dated credit default swaps, we observe much fewer tail events for the longer-dated ones. Surprisingly, the number of tail observations are substantially less even with lower tail threshold values. For example, contracts with one-year maturity have on average 104 tail observations and a tail threshold value of 8.38% considering a four-year rolling time window from 2009 to 2016. During the same period, 30-year contracts have on average 91 tail observations and a tail threshold value of only 4.16%. These findings support our hypothesis that financial distress has a far larger impact on short-dated credit risk compared to long-dated one.

The dependence of the credit tail risk on the time scale factor has implication on risk and portfolio management. Assume that smaller values of δt, i.e. daily to weekly returns, represent shorter-term investors, whereas higher values of δt represent investors with longer-term investment horizons. Portfolio managers with shorter investment horizons are interested in daily to weekly changes in credit default and tail risk, whereas monthly returns or changes of risk may be more relevant to long-term investors. Our empirical results suggest that for shorter time horizons, it becomes more important to incorporate the additional downside risk from fat tails into the risk-return trade-off. Short-dated credit default swaps capture more extreme events (even with a higher tail threshold) compared to longer-dated insurance contracts. This also holds for volatility adjusted returns. The observed patterns hold for the vast majority of countries and are persistent among different time scales independently of the time window.

1.4.4 Asymmetric Perception of Credit Tail Risk

In this section, we explore the uniformity of tail behaviour in credit default swap markets. Motivated by research in equity markets (LeBaron and Samanta, 2005; Stoyanov *et al.*, 2017), we investigate the differences in tail behaviour across countries and among regions. We associate positive tail returns with large price increases of the credit default swaps, which usually occur during a financial crisis or distress. Vice versa, we associate negative tail returns with large price decreases, which indicate an improvement of the financial situation, reduction in risk, deleveraging or bailouts. Our hypothesis is that the tail risk exponents of sovereign credit default swaps exhibit asymmetric tail behaviour during financial distress and recovery. Merton *et al.* (1985) already argued that it is difficult to see a clear theoretical explanation for extreme events being symmetric. We want to show that credit default swaps imply a higher probability of positive tail returns, whereas the average positive tail returns are more significant than average negative tail returns. We infer that financial crises occur more frequently, while financial recoveries potentially take longer. In equity and foreign exchange markets, this phenomenon is colloquially known as "up the stairs, down the elevator". We expect to observe the reversed pattern in credit default swap markets.

To evaluate the asymmetry of tail risk, we separately estimate positive and negative tail statistics using daily log-returns ($\delta t = 1$) of sovereign credit default swaps. Firstly, we investigate the asymmetric perception of tail risk using non-aggregated data of seven countries and four maturities (28 individual time series). Applying a rolling four-year time window, we estimate a total of 107,296 positive and negative tail exponents from 2009 to 2016.[25] Secondly, we extend our previous studies to cross-sectional (aggregated) data. The cross-maturity (CM) method aggregates the tail returns of four maturities (01Y - 10Y) in one time series for each country. The cross-maturity tail exponent is calculated using a rolling backwards-looking time window of only one trading year (252 observations). Hence, the length of the aggregated time series is equivalent to the four trading years of the non-aggregated method. For the same period from 2009 to 2016, we estimate 26,824 exponents for both tails. Thirdly, to assess the asymmetry of tail risk for different regions, we construct a cross-maturity and cross-country (CMCC) time series. For the core region, the CMCC time series aggregates the tail returns for four maturities of Belgium, France, and Germany. For the peripheral region, the CMCC time series aggregates the tail returns for four maturities of Ireland, Italy, Spain and Portugal.

[25] We reduce the computational complexity by limiting this analysis to four instead of five maturities.

The regional tail exponent is calculated using a rolling time window of only half a year.[26] We estimate 7,664 exponents for both tails and regions between 2009 and 2016. Similar to the previous sections, we impose the minimum tail length of 40 observations for positive and negative tail returns. This condition is applied for aggregated and non-aggregated time series. Time series with shorter tails are ignored. To assess if the difference between the average positive and negative tail exponent is statistically significant, we compute the pooled standard deviation, standard error and construct confidence intervals at a 99% confidence level.

Asymmetry of Upper and Lower Exponents

	01Y		03Y		05Y		10Y		Cross-Maturity	
	Pos.	Neg.	Pos.	Neg.	Pos.	Neg.	Pos.	Neg.	Pos.	Neg.
Belgium	2.69	2.55	3.63	3.40	3.34	3.59	3.59	3.70	2.53	2.95
France	2.88	2.97	3.60	3.27	3.49	3.70	3.73	3.76	3.29	3.13
Germany	2.66	3.03	3.26	3.52	3.12	3.66	2.69	3.41	2.73	3.03
Ireland	2.48	2.70	2.86	2.99	3.33	3.34	3.34	3.47	2.89	3.10
Italy	3.16	3.24	3.70	4.02	3.65	3.89	3.79	4.08	3.51	3.58
Portugal	3.09	3.26	3.38	3.57	3.14	3.69	3.37	3.78	3.18	3.31
Spain	3.00	3.39	3.23	3.77	3.20	3.81	3.84	3.97	3.44	3.68

Table 1.6 shows the positive and negative tail exponent for four maturities and seven countries from 2009 to 2016. The positive tail exponent is lower in 25 cases, which indicates a higher probability of returns above the tail threshold in times of financial turbulences compared to tail returns during a financial recovery. The difference between the tail exponents is statistically significant for 23 individual time series and is not statistically significant for France (10Y) and Ireland (05Y). Only Belgium (01Y and 03Y) and France (03Y) tend to have thicker negative tails than positive tails. In these three cases, negative price swings (decreases in insurance prices) above the tails threshold are more likely to happen than extreme positive returns (increases in insurance prices). The last columns show the positive tail exponents for aggregated returns of cross-maturity (CM) time series. We observe that tail returns of most credit default swaps imply a higher probability of positive tail events. Those extreme returns are more profound for positive tails.

Firstly, we compare the positive (right) tail exponent with the negative (left) tail exponent. We find that the positive tail exponent is lower in 25 out of 28 individual time series from 2009 to 2016. The difference between the tail exponents is statistically significant for 23 individual time series. We observe the biggest differences for Germany (05Y and 10Y), Portugal (05Y and 10Y) and Spain (03Y and 05Y). When the positive tail exponent is lower than the negative tail exponent, it implies that the return distribution has thicker positive than negative tails. However, this result must be treated with some caution, as the difference in tail exponent might be caused by differences in the tail thresholds. Therefore,

[26] The length of the CMCC time series is calculated by multiplying the number of countries per region, the number of maturities and the lookback window. The length of the time series is 1512 trading days for the core region and 2016 trading days for the peripheral region.

we adjust for these differences by fixing the tail thresholds to a constant value of 5%. Then we estimate the tail exponents for the returns exceeding the fixed tail threshold. Given the same cut-off point for both sides of the distribution, we find that the positive tail exponent is lower in 24 out of 28 individual time series. These results also confirm our previous finding with time-varying percentile and asymmetric tail threshold. We infer that credit default swaps with a lower positive tail exponent have a higher probability of returns above the threshold in times of financial distress compared to tail returns during a financial recovery.

Moreover, we observe that the positive tail threshold is higher compared to the negative tail threshold. The difference in threshold values is statistically significant for 21 out of 28 time series. A higher positive tail threshold indicates that the (minimum) positive tail return has a more profound impact on sovereign credit default prices during periods of financial distress than the (minimum) negative tail return during a financial recovery. However, this result must be treated with caution, as the difference in tail threshold might be caused by differences in tail length. Therefore, we control for these differences by fixing the tail length to arbitrary percentiles. We calculate the positive and negative tail threshold for the 5% and 95% percentiles. Given the same tail length for both sides of the distribution, we find that the positive tail threshold is higher for all 28 individual time series. Furthermore, the average tail returns of the right side of the distribution are significantly larger than the average tail returns of the left side for all 28 individual time series.[27] These results also confirm our previous finding with time-varying percentile and asymmetric tail threshold. In combination, these results conclude that credit default swaps with a lower positive tail exponent have a higher probability of extreme returns whereas those large price increases are more significant in times of financial distress than large price decreases in periods of a financial recovery.

Secondly, using cross-maturity time series, we observe that positive tail exponents are smaller than negative tail exponents for all countries except France. The difference between the two tails is the largest for Belgium (0.30) and Germany (0.30). The peripheral countries, such as Italy (0.06), Ireland (0.22), Spain (0.13) and Portugal (0.24) have a smaller dispersion asymmetry. We infer from these results that sovereign credit default swaps imply a higher probability of extreme price fluctuations during periods of financial distress compared to periods of financial stabilisation, i.e. deleveraging of governments' debt. Similar to our previous study on individual time series, we adjust for differences in tail thresholds by fixing

[27] We confirm this result also for the same cut-off point for both tails. Given a tail threshold value of 5%, we find that the average tail return is larger for the positive tail in 27 out of 28 individual time series and statistically significant in 25 cases.

it to a constant value of 5%. Considering the same threshold value for both tails, we observe that the positive tail exponents are lower for all countries using cross-maturity tail returns. Furthermore, we consider the differences in tail length as a possible reason for the asymmetry of tail risk. Therefore, we compute the tail statistics for both sides of the distribution using the upper and lower 5% percentile. Both alternative methods confirm our previous results with time-varying percentiles and asymmetric tail threshold. We conclude that aggregated returns of cross-maturity time series have a higher probability of extreme returns in times of financial turbulence than in times of financial recovery. Peripheral countries such as Italy, Ireland and Portugal, show that those extreme returns are more profound for the upper tail (compared to the lower tail) of the distribution. Our findings are supported given the evidence of severe and increasing sovereign debt, and more frequent crashes of credit markets in Europe over the last decade, particularly among peripheral countries.[28] Surprisingly, the average positive tail return of core countries tends to exceed peripheral countries.[29]

This raises another question: is there an asymmetric perception of tail risk within the core and peripheral region? Countries of the core region, such as Belgium, France and Germany, are characterised by lower debt ratios, lower credit default swap prices, and hence, smaller implied probabilities of default. Vice versa, countries of the peripheral region have higher debt ratios and implied default probabilities. The peripheral region encompasses countries such as Italy, Ireland, Portugal and Spain. We estimate the regional credit tail risk exponent using cross-maturity and cross-country (CMCC) data. We find strong evidence of asymmetric tail risk within different regions. For the core region, the positive tail exponent (3.06) is lower than the negative tail exponent (3.21). While the return distribution has thicker positive

[28]The peripheral countries experienced the following financial crisis: **Ireland** required support from the European Union's European Financial Stability Facility (EFSF) and the International Monetary Fund (IMF) in November 2010. The 2008-2014 financial crisis in **Spain**, also known as the Great Spanish Depression, started in 2008 during the world financial crisis of 2007–08. In 2012, Spain became a late participant in the European sovereign debt crisis when the government was unable to bail out its financial sector, and consequently had to apply for a €100.0 billion rescue programme provided by the European Stability Mechanism (ESM). The 2010–2014 **Portuguese** financial crisis was part of the broader downturn of the Portuguese economy that began in 2001 and possibly ended in 2016–17. The period from 2010 to 2014 was probably the most problematic part of the entire economic crisis. In 2011, the Portuguese government applied for €78.0 billion bail-out package from the International Monetary Fund (IMF), the European Financial Stabilisation Mechanism (EFSM), and the European Financial Stability Facility (EFSF) to prevent an insolvency situation.

[29]Using time-varying percentiles, the average positive tail returns of Belgium, France and Germany are 10.70%, 11.69%, and 12.94%, whereas the average tail returns of Italy, Ireland, Portugal and Spain are 10.75%, 7.21%, 9.72% and 9.14%. Given the same tail percentile for all time series, the average positive tail return is even stronger for core countries: Belgium (14.22%), France (15.05%), Germany (17.34%), Italy (12.73%), Ireland (10.21%), Portugal (12.11%) and Spain (12.44%).

than negative tails, the frequency of positive and negative shocks is relatively equal in the core region.[30] We find a much larger difference between the tail exponents in the peripheral region. The positive tail exponent is 3.39 compared to 3.80 for the opposite tail from the empirical distribution. The differences are statistically significant in both regions. We also confirm the tail asymmetry within the core and peripheral regions by adjusting for differences in the tail thresholds and tail lengths. We conclude that tail exponents for aggregated tail returns of regional time series exhibit a higher probability for positive tail returns in both regions, whereas the average tail returns are higher in periods of financial distress.

Much more remarkable is that the credit tail risk exponent of the core region exhibits higher tail risk despite a lower probability of default than the peripheral region. Inspired by the volatility smile for options in equity markets, we propose the following possible explanation in credit default swaps markets.[31] Assume that a credit default swap of a core country has a large distance to the implied default barrier. This is similar to a far out-of-the money (OTM) option in equity markets. The implied default barrier can be seen an equivalent to the strike price of the option, which triggers an event. The distance to the implied default barrier is shorter for credit default swaps of peripheral countries. In other words, peripheral countries with a higher probability of default are closer to the implied default barrier (strike price). Credit default swaps of core countries become more expensive in times of high volatility than those of low volatility, whereas credit default swaps of peripheral countries have a lower level of volatility, that might have a weaker impact on tails of return distributions (see Figure 1.4). This result suggests that a higher level of volatility makes the right tail thicker, while the left tail thinner, resulting in more right-skewed return distributions (when the underlying country of a credit default swap is further away from the implied default barrier). To support our explanation, we eliminate the impact of volatility on tail distribution by normalising the

[30]Additional tail statistics support our argument that positive and negative shocks are relatively symmetrical in the core region. The average positive and negative number of tail events (126 and 129 observations) and tail thresholds are similar (6.76% and 6.39%).

[31]A volatility smile refers to a graph shape (a U-shape pattern), which is the result from plotting the different strike prices and different volatilities (so-called implied volatilities) of a group of European options (calls and puts) with the same underlying asset and expiration date. The U-shaped pattern is often referred to as the volatility skew or smile and exists in all major stock index markets today. Typically, the implied volatility is higher when the underlying asset price is deeply in- or out-of-the-money compared to at-the-money options. The existence of the volatility smile is often attributed to fear of large downward market movements, sometimes known as "crash-o-phobia" (Andersen and Andreasen, 2000). The steepness of the smile decreases with increasing option maturities. We observe that the credit tail risk decrease with increasing credit default swap maturity. Using cross-country time series and time-varying tail percentiles, the credit tail is 3.14, 3.19, 3.42, and 3.51 for 01Y, 03Y, 05Y and 10Y maturities. We confirm this pattern using contant tail percentiles and tail threshold values. We refer to this pattern as the "term structure of credit tail risk".

returns of the regional time series. After the normalisation, we find that there is no significant difference between the credit tail risk exponent for the core and peripheral region. Therefore, we conclude that the core region with a large distance to the implied default barrier exhibit higher credit tail risk due to the significant impact of volatility on the tail distribution.

The asymmetry between positive and negative tails has important implications for financial risk management and stress testing. Positive and negative returns should not be combined (in the form of absolute returns) in tail risk models because extreme returns might have different heavy-tailed distributions. This may lead to over-or under-hedging of default risk depending on the portfolio allocation. Finally, the combination of power-law distributions with different risk exponents for positive and negative tail returns ignores the directional risk of extreme market movements.

Rolling Cross-Sectional Standard Deviation of Core and Peripheral Region

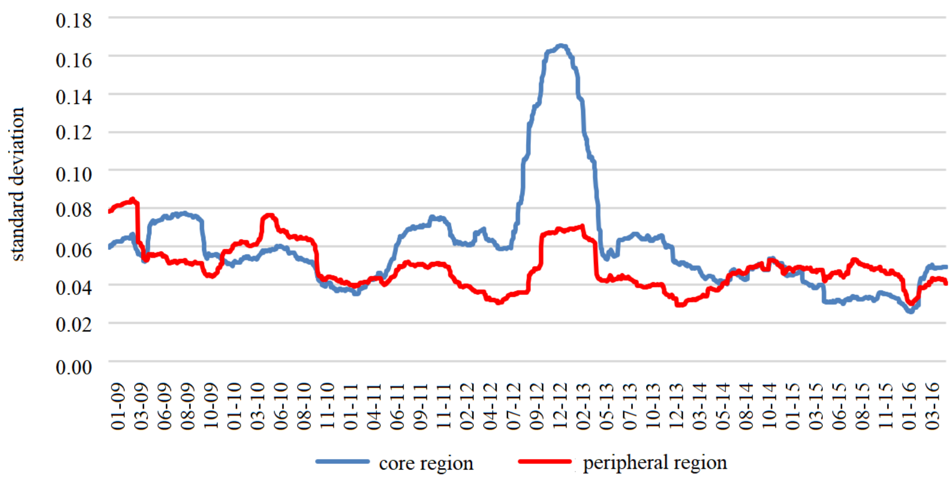

Fig. 1.4 shows the rolling standard deviation of the cross-sectional time series for the core (blue) and peripheral (red) region from January 2009 until May 2016. The standard deviation is calculated based on daily cross-sectional returns for each region using a daily rolling time window of 126 days, which is equivalent to the lookback window used for the tail exponent estimation. We observe that standard deviation of the core region, and consequently the volatility, exceeds the standard deviation of the peripheral region in times of financial distress, especially during the peak of European debt crisis in 2012. Credit default swaps of core countries become more expensive in times of high volatility than those of low volatility, whereas credit default swaps of peripheral countries have a lower level of volatility, that might have a weaker impact on tails of return distributions. Consequently, the higher level of volatility makes the right tail thicker, while the left tail thinner, resulting in more right-skewed return distributions. This effect might be more profound when the underlying country of a credit default swap is further away from the implied default barrier. The significant impact of volatility on the tail distributions of core countries, might be a plausible explanation for the higher credit tail risk in the core region.

1.5 Conclusion

A measure of tail risk in credit markets is essential to understand the behaviour of asset prices. Extreme value techniques based on arbitrary percentiles systemically under- or overestimate the actual tail length when the tail percentile fluctuates with the default risk. Consequently, as risk changes through time, these techniques lead to imprecise tail risk estimates. We present a dynamic tail risk model based on quantitative traceable method that overcomes this difficulty. We propose a credit tail risk measure based on a dynamic power-law model with multiple time-varying tail parameters. We estimate tail risk in the credit market by exploiting extreme returns of credit default swaps on sovereign debt. Our sample consists of sovereign credit default swaps with five tenors of seven countries within the European Monetary Union from 2005 to 2016.

We find that our dynamic power-law model is a feasible measure for credit tail risk. While the power-law hypothesis test does not rule out models based on time-varying tail percentiles, it rejects conventional methods based on arbitrary percentiles and threshold values. The power-law hypothesis is performed on negative, positive, absolute and normalised tail returns of sovereign credit default swaps. Furthermore, alternative distributions do not better fit the tail returns.

We find a cubic power-law distribution of large credit default swap fluctuations for over 750,000 data sets. The average tail risk exponent is $\alpha \simeq 3$ within a range of one standard error, which supports the cubic law's universality. This result is robust for different countries, maturities and regions, and holds for daily and weekly returns. Our finding suggests that credit tail risk exponents are outside the Lévy-stable region $0 < \alpha < 2$. This implies the existence of a finite second moment for return distributions, which has important implications for risk and portfolio management.

The empirical evidence suggest that credit tail risk decreases with increasing time scales. Furthermore, credit default swaps with shorter-dated maturities exhibit a higher probability of large price fluctuations, persistent among different time scales. These results suggest credit default swaps with short-maturities exhibit the highest credit tail risk on short time scales, which holds for unnormalised and normalised returns independently of the time window. We provide different possible explanations for these patterns.

Finally, we find significant tail risk asymmetries, which explain differences in extreme returns during periods of financial distress and financial recovery. Firstly, we find that individual credit default swaps exhibit a higher probability of extreme increases in prices during periods of financial distress than the probability of extreme price decreases during financial recovery periods. Thereby, the minimum magnitude of the smallest positive tail returns is significantly higher than the negative one. The average (positive) tail return during periods of financial distress is significantly higher than the average (negative) tail return during periods of financial stabilisation. These findings also hold for aggregated returns of cross-maturity time series. Furthermore, we find the tail asymmetry among different regions. Despite having a lower probability of default, the core region exhibits a higher credit tail risk than the peripheral region. This puzzling result can be explained by the significant impact of volatility on the tail distribution. Credit default swaps of core countries (region) with cheap insurance prices pay off in high tail risk states and thus are valuable hedges against extreme events.

Chapter 2

Decomposition of the Tail Exponent

2.1 Introduction

Various economic and financial time series are known to exhibit distributions with a power-law decay (Gabaix *et al.*, 2006). Power-laws are reported for different asset classes, such as foreign exchange rates (Guillaume *et al.*, 1997), individual stocks (Plerou *et al.*, 1999), financial market indices (Gopikrishnan *et al.*, 1999), trading volume (Gabaix *et al.*, 2003), cryptocurrencies (Begušić *et al.*, 2018) and credit default swap markets (Chapter 1). For most financial time series, the power-law only holds for a fraction of the time series, above a certain threshold value (m). The threshold value determines the number of tail observations, which is referred to as k. In the statistical literature, there is an ongoing debate on the number of tail observations (k) required to accurately estimate the tail risk exponent (α). The literature to date considers three groups of estimation methods on which the criterion of k are based: (i) heuristic rules such as the Eye-Balling method or arbitrary (fixed) percentiles, (ii) minimising the mean squared error of the tail exponent estimator in the probability dimension, and (iii) minimising the maximum deviation in the quantile dimension. Chapter 1 reports that heuristic rules, such as fixed percentiles, systemically under- or overestimate the time-varying tail length (k_t), and consequently the dynamic tail risk exponent (α_t). Instead, the number of tail returns vary over time and correlate with the risk of the underlying asset. Therefore, we presented a quantitative traceable dynamic tail risk model with time-varying tail threshold (m_t) and length (k_t) in Chapter 1. However, a challenge of this approach is to quantify changes in the tail risk exponent ($\Delta\alpha_{t,t+1}$) from day t to the consecutive day-ahead $t+1$ due to fluctuations in tail returns and variations in tail length, especially when both effects coincide. To overcome this challenge, Chapter 2 presents a new decomposition method. The decomposition of the tail risk exponent allows quantifying daily changes in tail risk due to changes in tail returns and tail length separately. Based on these new insights, Chapter 3 proposes a novel methodology to select the optimal time-varying tail length based on minimising the average distance between the empirical distribution and the theoretical power-law distribution over a moving time window.

The literature offers different methods to determine the optimal tail length. Broadly generalised, the methods can be divided into three groups. The first group consists of heuristic-based approaches. Heuristic models are frequently used in applications and focus on finding the stable region of the optimal tail length, where the tail length increases, the variance is subsiding, and the bias of the tail estimator has not yet become dominant. The techniques of finding the tail length by observing the stable regions in the Hill plot are known as the Eye-Balling method or the automated form of the Eye-Balling method (Resnick and Stărică, 1997). Another heuristic method is to limit the tail analysis to arbitrary percentiles. For example, Plerou *et al.* (1999) and Gopikrishnan *et al.* (1999) define the number of tail returns (k) by only using returns larger than two, three or five standard deviations or within a range of standard deviations. Doyne Farmer *et al.* (2004) limit the analysis to the most significant observed returns only, such as the largest $\sqrt{n}$ or $\frac{1}{10}n$, where n is the length of the sample. More recently, Kelly and Jiang (2014) utilise a fixed tail percentile at a 5% level for the estimation of tail risk using cross-sectional stock returns. The advantage of fixed tail percentiles is that fluctuations in tail risk are only attributed to changes in returns above the tail threshold and not due to variations in tail length. However, the main disadvantage is that fixed tail percentiles ignore fluctuations in the tail length over time. Furthermore, these methods have a weak theoretical foundation and underestimated model uncertainty (Stoev *et al.*, 2006).

The second branch of literature derives from the theoretical statistical literature. These methods are based on the minimisation of the mean squared error of the tail exponent estimator in the probability dimension, and balance the asymptotic variance and bias components. Ergun (2016) defines the probability dimension such that

$$p = \frac{k}{n}\left(\frac{x}{m_k}\right)^{-\alpha_{(k)}}$$

where k is the number of tail events above the lower bound on the power-law behaviour. The return that produces the smallest maximum difference along all the tail observations defines the lower bound (or cut-off point) of m_k, where k is chosen as the optimal number of tail returns to estimate the thickness of the tail. When the size of the tail (k) is small, the variance of the tail exponent estimate is large, while the use of a large number of tail observations (k) introduces a large bias in the estimation. Several procedures have been introduced for choosing the optimal tail length k in the sense of asymptotic minimal mean squared errors (see, for example, Dekkers and de Haan (1993), and Beirlant *et al.* (1996)). Hall (1990) and Danielsson *et al.* (2001) employ bootstrap procedures to minimise the asymptotic mean squared errors. Drees and Kaufmann (1998) exploit the same bias and variance trade-off, but use the maximum random fluctuation of the estimator to locate the point where the trade-off is optimal. These methods, based on minimisation of the mean squared error in the probability dimension, are asymptotically consistent, but have unsatisfactory finite sample properties (Danielsson *et al.*, 2016).

The third group of methods derives from minimising the maximum deviation in the quantile dimension. There are a variety of measures for quantifying the distance between two probability distributions. A commonly used measure for non-normal data is the Kolmogorov-Smirnov test.[1] Generally speaking, the Kolmogorov-Smirnov test evaluates whether empirical data seem a plausible random sample from a given probability distribution, by comparing the maximum difference (denoted D) between the empirical and a parametric distribution. The fundamental idea behind the Kolmogorov-Smirnov distance metric is introduced in Chapter 1 and has been extended to a dynamic estimate of the time-varying tail length (k_t). Danielsson *et al.* (2016) compare the results of simulated time series between the Kolmogorov-Smirnov distance to other methods mentioned above. He concludes that the Kolmogorov-Smirnov distance metric is the preferred approach to determine the tail length (k).

This chapter addresses the challenge of using the Kolmogorov-Smirnov distance metric for dynamic tail risk exponent estimations. The Kolmogorov-Smirnov distance metric estimates the dynamic tail threshold value (m_t), which defines the optimal tail length (k_t) for each rolling time window at t. Then the Hill (1975) estimator determines the daily rolling tail risk exponent (α_t) for k_t-many tail observations. We calculate the total difference of the tail exponent by taking the difference of two consecutive tail exponents, i.e. α_t and α_{t+1}. In the time-varying estimation framework, variations of the tail risk exponent ($\Delta\alpha$) from day t to the following day $t+1$ can be caused by changes of tail returns, changes of tail length, or both effects simultaneously. Across asset classes and time windows, we observe simultaneous changes of tail returns and tail length in 50% to 70% of all cases. If both effects coincide, it becomes difficult for economists to quantify each factor independently. This problem is particularly striking in the presence of jumps in the tail length from t to $t+1$. We observe

[1]The Kolmogorov-Smirnov test statistic goes back to publications by Kolmogorov (1933) and Smirnov (1948).

variations of the daily tail risk exponent ($\Delta\alpha$) from t to $t+1$ due to jumps in the tail length in 20% to 35% of all cases across asset classes and time windows. Longer lookback windows usually tend to have fewer jump events than short time windows. Abrupt changes of tail length usually result in larger movements of the tail exponent ($\Delta\alpha$), which do not necessarily reflect the actual change in tail risk. These jumps in tail length are certainly correct from a statistical standpoint (minimum of the maximal deviation), but they make consistent tail risk estimations more difficult for economists. Currently, there is no method to separately quantify changes in the time-varying tail risk exponent due to these two effects. We close this gap in the literature by introducing a decomposition method for the dynamic tail exponent.

The main goal of the decomposition analysis is to quantify the factors independently, which cause daily variations of the dynamic tail risk exponent. Firstly, we measure the variation of dynamic tail risk exponent caused by changes in tail returns from t to $t+1$. Therefore, we estimate the two tail exponents using the same tail threshold m_t, but different subsets of tail returns, $x_{i,t} > m_t$ on day t, and $x_{i,t+1} > m_t$ on day $t+1$. We control for fluctuations due to varying tail threshold by keeping m_t constant. The difference between these two tail exponents, measures the variation of the tail risk only due to the changes of the tail returns. Secondly, we assess the variation of dynamic tail risk exponent due to changes in the threshold values. Remember that the threshold value m_t defines the tail length k_t. We estimate two tail exponents using the same set of tail returns $x_{i,t+1} > m$ on day $t+1$, but different tail thresholds m_t and m_{t+1}. This time we control for fluctuations due to varying tail returns by keeping the set of tail returns $x_{i,t+1}$ unchanged. The difference between these two tail exponents measures the change in tail risk due to changes in the tail length. The sum of both differences is equal to the total variation of the tail exponent from t to $t+1$, which confirms that the decomposition is done correctly.

We estimate the dynamic tail risk exponent and perform the decomposition analysis for sovereign credit default swaps and the corresponding national stock market indices. We select a sample of four core (Austria, Belgium, France, Germany) and four peripheral countries (Ireland, Italy, Portugal, Spain) within the Eurozone. IHS Markit Ltd. provides daily composite quotes of sovereign credit default swap prices from January 2005 to March 2017 and daily prices of national stock market indices from January 1999 to March 2017.

The rest of this chapter is organised as follows. Section 2.2 introduces the decomposition method, which quantifies the changes in the time-varying tail risk exponent due to different factors. Section 2.3 describes the data and selection criteria for sovereign credit default swaps and equity indices. Section 2.4 reports the results of the decomposition analysis for different asset classes, followed by concluding remarks.

2.2 Methodology

2.2.1 Estimation of the Dynamic Tail Risk Exponent

Different estimation methods to determine the tail exponent have been proposed in the literature (Hill, 1975; Pickands *et al.*, 1975; de Haan and Resnick, 1980; Hall, 1982; Mason, 1982; Davis and Resnick, 1984; Csorgo *et al.*, 1985; Hall *et al.*, 1985). The most popular method for estimating the tail exponent of heavy-tailed distributions is the Hill (1975) estimator. Recent research in statistics of extreme values shows that the Hill (1975) estimator performs well even in the presence of dependent and heterogeneous data (Kelly, 2014). In Chapter 1, we derive the time-varying Hill (1975) estimator in Equations 1.3 to 1.6, which is applied in this chapter. Equation 1.6 defines the time-varying Hill (1975) estimator as

$$\hat{\alpha}_t = 1 + k_t \left[\sum_{i=1}^{k_t} \ln \frac{x_{i,t}}{m_t} \right]^{-1},$$

where $x_{i,t}$, $i = 1,...,k$ are observed returns in the tail of x_t such that $x_{i,t} \geq m_t$. The dynamic tail length is denoted by k_t in period t, which we discuss in the following section.

2.2.2 Estimation of the Dynamic Tail Length

The estimation of the optimal tail length derives from the classic Kolmogorov–Smirnov distance method. The Kolmogorov-Smirnov distance matches the empirical and theoretical distribution to find the optimal tail length for the heavy-tailed distributions. The distance is measured in the quantile rather than the probability dimension. Using the quantile domain is justified by the fact that a probabilistic error in the tail region translates into a significant distortion in the quantile dimension. Also, the quantile dimension is the domain that economists care about. Consequently, we focus on the quantile dimension rather than the probability dimension for modelling the optimal tail length.

Given that we measure over the quantile domain, we need a penalty function for deviations from the empirical distribution. Well-known penalty functions are the mean squared error (MSE), the root mean squared error (RMSE), the mean absolute error (MAE), in addition to various other penalty functions that weigh deviations differently. While the mean squared error approach penalises large deviations disproportionately more than small deviations, the mean absolute error punishes errors in a direct and linear proportion. As we focus on tail events only, we do not require additional penalisation, like squared differences. In this research, we use the penalty function of the Kolmogorov-Smirnov statistic for fitting the tail of the power-law distribution. Similar to Chapter 1, we calculate the time-varying Kolmogorov-Smirnov distance (D_t), which defines the lower threshold value $\hat{m}_t$. The estimate $\hat{m}_t$ is the value of m_t that minimises D_t, defines the cutoff point between the body and the tail of the distribution, and subsequently determines the optimal length k_t at time t.

2.2.3 The Decomposition of Tail Exponent

In the time-varying tail risk estimation framework, differences of the tail risk exponent ($\Delta\alpha$) from t to $t+1$ can be caused by changes of tail returns (Δx_i), changes of the tail threshold (Δm), or both effects simultaneously. A change in the tail threshold (Δm), most probably implies a change in the optimal tail length (Δk). If both effects coincide, it is not evident how much each factor contributes to the change in tail risk exponent. Therefore, we introduce a decomposition method, which enables us to quantify changes in the time-varying tail risk exponent caused by these two factors separately.

Firstly, we calculate the total difference of the tail exponents from day t to the consecutive day-ahead $t+1$. We measure the total difference of the tail exponent due to potential changes in the tail returns and tail threshold such that

$$\Delta\alpha_{t,t+1} = \alpha_{t+1}(x_{i,t+1} > m_{t+1}) - \alpha_t(x_{i,t} > m_t) \tag{2.1}$$

where the tail risk exponent α_t is estimated on the tail returns $x_{i,t}$ exceeding the tail threshold m_t. For estimating α_{t+1} the rolling window with constant length drops the oldest return as a new return is available on day $t+1$, which implies that α_t and α_{t+1} have overlapping data. Then α_{t+1} is estimated on the set of tail returns $x_{i,t+1}$ above the tail threshold m_{t+1} on the following trading day $t+1$.

Secondly, we measure the daily difference of the tail exponent due to the change in the tail threshold m from t to $t+1$. Therefore, we compute the tail risk exponents using the set of daily tail returns at $t+1$ and the tail thresholds m_t and m_{t+1}, such that

$$(\Delta\alpha_{t,t+1}|\Delta m_{t,t+1}) = \alpha_{t+1}(x_{i,t+1} > m_{t+1}) - \alpha_{t+1}(x_{i,t+1} > m_t), \tag{2.2}$$

where the tail thresholds determine the tail length k_t and k_{t+1} given the same tail returns $x_{i,t+1}$ at $t+1$. Note that larger fluctuations in the tail threshold lead to significant variations in the length, and consequently, to larger fluctuations in the tail risk exponent.

Thirdly, we measure the daily difference of the tail risk exponent due to changes in the tail returns $x_i > m$ from t to $t+1$. We subtract the daily variation of the tail exponent due to the change in the tail threshold m in Equation 2.2 from the total differnce in Equation 2.1, such that

$$\begin{aligned}(\Delta\alpha_{t,t+1}|\Delta x(i)_{t,t+1}) &= [\alpha_{t+1}(x_{i,t+1} > m_{t+1}) - \alpha_t(x_{i,t} > m_t)] \\ &\quad - [\alpha_{t+1}(x_{i,t+1} > m_{t+1}) - \alpha_{t+1}(x_{i,t+1} > m_t)] \\ &= \alpha_{t+1}(x_{i,t+1} > m_t) - \alpha_t(x_{i,t} > m_t).\end{aligned} \tag{2.3}$$

We compute the tail exponents α_t and α_{t+1} using the same tail threshold m_t. By keeping the tail threshold constant, we measure the variations in the tail exponent which are only contributed by changes of the tail returns from t to $t+1$. However, a constant tail threshold does not necessarily mean a constant number of tail events k. There are three possible options for changes in the tail length $k = [-1, 0, +1]$, given $\Delta\alpha_{t,t+1} \neq 0$ and $\Delta m_{t,t+1} = 0$. When the rolling time window moves from t to $t+1$, the oldest observation is dropped, and the latest return on date $t+1$ is added to the new subsample. If the dropped return is a non-tail event, and the new return is a tail event, then the tail length k changes by +1 observation. Vice versa, if the oldest observation is a tail event and the newest one is a non-tail event, then the tail length k changes by -1 observation. If a tail event gets replaced by a new tail event, then the tail length does not change $k = 0$. In all three cases, the change in $\alpha_{t,t+1}$ is exclusively caused by changes in tail returns. However, in most cases $\Delta\alpha_{t,t+1} \neq 0$ coincide with $\Delta m_{t,t+1} \neq 0$ and $\Delta x(i)_{t,t+1} \neq 0$. Therefore, we develop the tail decomposition algorithm to quantify changes in the dynamic tail risk exponent due to the two effects separately. We apply the decomposition method to sovereign credit default swap and equity indices.

2.3 Data and Selection Criteria

2.3.1 Credit Default Swap Data

We estimate the dynamic tail risk exponent using daily sovereign credit default swap prices from January 2005 to March 2017. The main data source for this investigation is IHS Markit Ltd., to which a group of leading market participants (including the G16 banks) contribute credit default swap quotes on a daily basis. Based on these quotes, IHS Markit calculates the daily composite quotes.[2] In order to form a composite, IHS Markit requires at least three distinct contributors submitting price information.

All credit default swap data is subject to a range of tests, which ensure data quality. These assessments encompass so-called 'logical' and 'relative' tests. The first logical test is the curve buildability test. It checks for valid survival and default probabilities using the bootstrapping method in the ISDA CDS Standard Model. Credit default swaps with unreasonable probabilities are rejected. The second test is the backwardation test, which examines the relationship across restructuring types. There are four common restructuring types, namely the Cum Restructuring (CR), Modified Restructuring (MR), Modified Modified Restructuring (MM), and Ex-Restructuring (XR). The backwardation test validates the following inequality for the contributed spreads for each maturity: $CR \geq MM \geq MR \geq XR$. This implies that credit default swaps with restructuring clause (CR) have a higher price compared to credit defaults without restructuring (XR). The first relative test is the stable data test, which measures the frequency of price updates and liquidity. If a credit default swap quote is updated infrequently or does not change, the data is excluded. We observe such stale data for our sample of sovereign credit default swaps before 2005, hence we cannot use a longer lookback window. In 2005, sovereign credit default swap prices received on average only one price update per week. The prices are refreshed three-time weekly by 2007. From 2008 onwards, we notice daily updates of sovereign credit default swap prices. Further, if a credit default swap curve fails the liquidity test, it is not used in the composite calculation. An entity is defined to be illiquid if it receives price updates from 13 or fewer banks or market maker. Secondly, the outlier test removes price information if a market maker genuinely has a different opinion of the value of the credit default swap or has not updated the price if the market of the underlying asset has moved. If a credit default price fails any of four tests,

[2] The G16 is an industry group comprising the largest derivatives dealers: Bank of America-Merrill Lynch, Barclays Capital, BNP Paribas, Citi, Crédit Agricole, Credit Suisse, Deutsche Bank, Goldman Sachs, HSBC, JP Morgan, Morgan Stanley, Nomura, RBS, Societe Generale, UBS and Wells Fargo Bank.

Markit rejects the received data.

To be included in the dynamic tail risk analysis, credit default swap data is subject to a range of criteria. The European Monetary Union consists of 19 countries with the common currency, the Euro. The countries in the Eurozone as of 2019 are Austria, Belgium, Cyprus, Estonia, Finland, France, Germany, Greece, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, Netherlands, Portugal, Slovakia, Slovenia and Spain. Starting in January 2009, we apply a daily rolling backwards-looking time window of four trading years. Hence, to estimate the tail exponent, we rely on the price information from 2005 to 2008. The first criterion is that a time series has sufficient price information to estimate the tail length and exponent. To be considered in the sample, we require >50% or more than two years of return information within the rolling time window of four years. As we perform this study on countries in the Eurozone, the second criterion is that the Euro has been adopted by 2005. This excludes a range of countries: Slovenia (2007), Cyprus (2008), Malta (2008), Slovakai (2009), Estonia (2011), Latvia (2014) and Lithuania (2015). Thirdly, we exclude non-EU member states, which adopted the Euro but are rather unimportant in terms of economic size. These countries are Andorra, Monaco, San Marino and the Vatican City State. Finally, we exclude countries with considerable trading interruptions (trading holds), such as Greece. We focus the tail exponent decomposition on a data set, which includes eight countries, namely, Austria, Belgium, France, Germany, Ireland, Italy, Portugal and Spain. The first four countries are considered as core countries, whereas the last four countries are considered as peripheral countries. For every given country, we consider 10 credit default swap maturities, also called "tenor", with length from 1 to 30 years.[3] In some cases, there are small gaps of price information. We fill the missing information by assuming that the credit default swaps price remains unchanged from the previous day. Bai and Wei (2012) discusses the common convention regarding the structuring clauses and transaction currencies. They state that sovereign credit default swaps usually trade under the Cum Restructuring clause. The Markit dataset consists of four different seniority levels of the debt within the capital structure: senior, subordinated, junior and preferred. We take the price information for senior unsecured debt.

2.3.2 Stock Market Data

We estimate the dynamic equity tail risk exponent using daily returns of national stock market indices from January 1999 to March 2017. The main data source is the WRDS database

[3] Credit default swap tenors: 1, 2, 3, 4, 5, 7, 10, 15, 20 and 30 years.

of IHS Global Insights. The national stock market indices encompass the Austrian Traded Index (ATX Index), NYSE Euronext Brussels (BEL 20 Index), the French stock market index Cotation Assistée en Continu (CAC 40 Index), the German blue-chip stock market index (DAX 30 Index), the primary stock market index for the Borsa Italiano (FTSE MIB 40 Index), the Irish stock market index (ISEQ All Share Index), Spain's principal stock exchange (IBEX 35 Index), and the main stock exchange of Portugal (PSI-20 Index). From this data, we compute daily log-returns using end-of-day price indices. Weekends and non-trading days were removed from the dataset.

The motivation to use equity indices in addition to credit default swap data is twofold. Firstly, we want to validate the results from the decomposition of the credit tail risk exponents with another asset class, using the same backwards-looking time window of four years. We expect some minor differences due to the different nature of the securities, but generally similar findings. Secondly, credit default swaps are a relatively new asset class which comes with the limitation of shorter lookback windows for the dynamic tail risk analysis. We observe stale sovereign credit defaults with infrequent price updates and little liquidity before 2005. If we want to decompose the tail risk exponent using longer lookback window, we must use a different asset class such as stock market indices. The motivation for using longer lookback windows is that we want to analyse the impact of the window length on the variation of the tail threshold, tail length and tail exponent. We expect that longer lookback windows reduce the variation of these variables. A lower variation of the tail threshold and length means a stable estimation of the tail risk with less abrupt changes. We also expect to gain insights between stable measures (over longer time windows) and adaptable measures (over shorter time windows).

2.4 Results

2.4.1 Decomposition of Tail Exponents in Credit Markets

We estimate the dynamic tail risk exponent using daily returns of sovereign credit default swaps for eight countries and ten maturities from January 2005 to March 2017. Starting in January 2009, we apply a daily rolling backwards-looking time window with a length of four years. There are 2135 successive time windows for each country and maturity between January 2009 and March 2017. Given a set of returns for each daily rolling backwards-looking time window, the time-varying Kolmogorov-Smirnov distance metric estimates the tail threshold value (m_t), which defines the optimal tail length (k_t). Then, the Hill (1975) estimator uses returns above the tail threshold value (m_t) to estimate the daily tail risk exponent (α_t). The tail threshold value (m_t), the tail returns ($x_{i,t} > m_t$), the optimal tail length (k_t), and the daily tail risk exponent (α_t) are all interdependent in a non-linear relationship. The total variation in tail risk is the difference between two consecutive tail exponents (α_{t+1}-α_t). In our sample of 170,800 tail risk exponents, there are 26,695 (15.63%) non-zero differences in credit tail risk exponents.[4] In the following paragraphs, we analyse the causes, decompose the factors, and quantify the changes of the tail risk exponent due to changes in tail returns, changes in tail length and jumps in tail length.

There are three causes of tail risk fluctuations in this research: i) fluctuations of the tail risk exponent caused by changes in tail returns, ii) fluctuations of the tail risk exponent caused by changes in the tail threshold, and iii) fluctuations of the tail risk exponent caused by both effects simultaneously. We observe that in 56.24% (15,014) of all cases, fluctuations of the dynamic credit tail risk are only caused by changes in tail returns and not due to changes in the threshold value. In other words, 100% of the fluctuations of the dynamic credit tail risk are exclusively attributed to changes in extreme events. Only 3.45% (921) of all events belong to the second category. However, in 40.31% (10,760) of all occurrences, fluctuations of the dynamic credit tail risk coincide with changes in the tail threshold and tail returns, and it becomes difficult to quantify these factors separately. This is particularly problematic in the presence of abrupt changes of the tail threshold and length, which may result in larger fluctuations of the tail risk exponent. These abrupt changes of the tail threshold may be caused by statistical origins inherent to Kolmogorov-Smirnov distance metric estimates, and not by fundamental changes of the risk behaviour. Due to the interdependency of factors (x_t, m_t, k_t), it is difficult to quantify the impact of each parameter on α_t separately using a distance-based approach in dynamic risk management. To solve this problem, we developed a new decomposition method, which quantifies changes in the tail risk exponent due to changes in tail returns and tail length separately.

We focus our analysis on the 40.31% (10,760) of all cases, where fluctuations of the tail risk exponent coincide with changes in the tail threshold and tail returns. To better understand the fluctuations of the tail risk exponent, we deconstruct changes in tail risk exponent into its single components. Firstly, we calculate the daily variation of the credit tail risk due to changes in tail returns, by varying the set of returns from t to $t+1$ but keeping the threshold value (m_t) constant for both consecutive dates. We find that 11.98% of the changes

[4]We analyse eight (8) countries, ten (10) maturities on 2,135 consecutive days. The product of these three factors is 170,800.

in dynamic credit tail risk exponent are attributed to changes in tail returns. The average absolute change of the tail exponents is 0.2549, but only the minority of 0.0307 is due to absolute changes of tail returns. Secondly, we calculate the daily variations of the credit tail risk due to changes in the tail thresholds, by varying the threshold values m_t to m_{t+1} but using the same set of returns at $t+1$. We observe that 88.02% (0.2242) of the changes in tail risk exponents are explained by changes in the tail threshold, which is determined by the Kolmogorov-Smirnov distance metric. When tail threshold and tail returns vary at the same time, the average of absolute changes of the tail risk exponent is 6.79%. The average absolute change of the tail risk exponent is significantly larger compared to such cases where the tail risk exponent solely changes due to changes in tail returns (0.66%). However, after removing the effect due to the varying tail threshold from the total variation of the tail risk, the percentage of variation due to changing tail returns are similar (0.81%). This indicates that the tail exponent decomposition works correctly. Note that these results are averages of the decomposition of 10,760 tail exponents of eight countries and ten maturities and with some regional differences.[5]

Next, we analyse the fluctuations of credit tail risk due to abrupt and non-abrupt changes of the optimal tail length. When the rolling time window moves from t to $t+1$, the oldest observation is dropped, and the latest return on date $t+1$ is added to the new subsample. If the dropped return is a non-tail event, and the new one is a tail event, the tail length k changes by +1 observation. Vice versa, if the oldest observation is a tail event and the newest one is a non-tail event, the tail length k changes by -1 observation. A change in the credit tail risk exponent without a change in tail length ($\Delta k = 0$) implies that a tail return got replaced by another tail return. We define non-abrupt changes of the tail length as $\Delta k = [-1, 0, +1]$. We define smaller or larger variations of the tail length as abrupt changes or jumps in tail length. The abrupt changes account for 33.87% (9,043) and the non-abrupt changes for 66.13% (17,655) of the variation in the tail length, given there is a change in the tail exponent. In the presence of non-abrupt changes in the tail length, 80.94% of the tail risk variations are assigned to changes in the tail returns. However, suppose the time-varying Kolmogorov-Smirnov distance metric estimation leads to an abrupt change in the tail length. In that case, it is difficult to differentiate whether changes of tail returns or the jump in tail length caused the change in tail risk. We observe that in the presence of abrupt changes in tail length, 77.92% of the variation in tail risk is attributed to jumps and only 22.08% due to changes in

[5]The change in tail risk due to tail returns only: Germany 8.45%, Spain 10.01%, Belgium 11.79%, Portugal 12.80%, Austria 14.92%, Ireland 15.71%, Italy 16.83% and France 21.19%. The equally-weighted average across all countries and maturities is 11.98% as reported above.

tail returns. Therefore, even though the Kolmogorov-Smirnov distance metric is one of the preferred approaches to estimate the tail threshold (Danielsson *et al.*, 2016), it must be treated with caution in the time-varying tail risk analysis. Interestingly, the average jump size is close to zero for most maturities and countries. This might indicate that the effect of jumps are only temporary, and there is some form of mean reversion to the long-term average tail length.

In the following section, we analyse changes in tail risk for national stock market indices, to validate our results across asset classes, and to analyse the effects using longer time windows.

2.4.2 Decomposition of Tail Exponents in Equity Markets

In this section, we analyse the fluctuations of the tail risk exponents in the national stock market indices. Equity tail risk is associated with rare events of extreme magnitude in the lower tail of the return distribution. The goal of this research is to validate the results from the decomposition analysis of the credit tail risk exponent using return information from a different asset class. Given that we begin with the dynamic credit tail risk analysis before the European debt crisis in 2009, our analysis is limited to a lookback window of four years, because, before 2005, the credit default swap prices exhibit stale and illiquid data. Compared to credit defaults swaps, which is a relatively new asset class, stock market indices have much longer historical time series. This property allows us to perform the equity tail risk analysis on longer lookback windows. Chapter 1 reports that longer time windows usually result in more tail observations, reduces sampling noise and estimations errors. Therefore, we also investigate the impact of longer time windows on the decomposition of tail exponents.

We estimate the dynamic equity tail risk exponent using daily log-returns of national stock market indices from January 2005 to March 2017.[6] We keep the methodology unchanged for estimating the tail statistics in equity markets. The Kolmogorov-Smirnov distance metric estimates the tail threshold, and the Hill (1975) estimator the equity tail risk exponent. Firstly, to compare the results of credit tail exponent decomposition, we use the same input parameters for the equity tail risk analysis. Starting in January 2009, we estimate the tail statistics using a four-year lookback window on 2135 successive days. We observe a similar

[6]The average dynamic equity tail risk exponent (and standard deviation) for eight countries (indices) using a rolling four-year time window from 2009 to 2017, in increasing order: Belgium (BEL 20) α=3.7895 (0.4180), Germany (DAX 30) α=3.8989 (1.1577), France (CAC 40) α=4.0506 (1.0790), Spain (IBEX 35) α=4.1444 (0.9542), Portugal (PSI 20) α=4.1530 (1.0264), Ireland (ISQE All) α=4.1827 (0.9062), Italy (MIB 40) α=4.2891 (1.5343), and Austria (ATX) α=4.3264 (0.6447).

average tail length for credit default swaps (85) and stock market indices (82) for a four-year time window. Secondly, we analyse the impact of longer time windows on the decomposition analysis. We decompose the equity tail risk exponent using three extended time windows of six, eight and ten years. The average tail length is 110, 138 and 170 tail observations for six-, eight- and ten-year time windows, respectively. Compared to the four-year time window, the standard error reduces by 22.31%, 31.14% and 35.58% for these longer time windows.[7] In order to make the decomposition results comparable among different time windows, we keep the starting date unchanged and estimate 2135 successive tail statistics for each of the three extended time windows.

Firstly, we compare the credit and equity tail exponent decomposition for a rolling time window of four years. We compute a total of 17,080 equity tail statistics for eight national stock market indices and observe 2,427 (14.10%) changes in the daily equity tail risk exponent. The relative number of changes in equity markets is similar to credit markets (15.63%) over the same period from 2009 to 2017, indicating that equity markets are similarly exposed to changes in tail properties (tail length and tail returns). We find that in 62.38% (1,514) of all cases, fluctuations of the dynamic equity tail risk are caused by changes in tail returns only and not due to changes in the threshold value. In only 3.09% (75) of all events, the total variation of the dynamic equity tail risk is caused by adjustments of the tail threshold. In 34.53% (838) of all cases, variations of tail exponent occur with simultaneous changes of the tail threshold and tail returns. The average absolute change of the tail exponent is 0.3549, but only 0.0350 is due to changes in tail returns of equity indices. This implies that changes in the tail threshold explain the majority of 90.36% (0.2718) of the changes in the dynamic equity tail risk. These findings are in line with the previous results in credit markets (Section 2.4.1).

We analyse the fluctuations of equity tail risk due to abrupt and non-abrupt changes of the tail length given a four-year time window. Recall that we define non-abrupt changes of the tail length as $\Delta k = [-1, 0, +1]$, and smaller or larger variations in Δk as abrupt changes (or jumps). The abrupt changes account for 29.30% (711) and non-abrupt changes for 70.70% (1,716) of the variation in the tail length given there is a change in the equity tail exponent. In the presence of non-abrupt changes of the tail length, 78.98% of the variation in equity tail risk is attributed to changes in tail returns. However, in cases of abrupt changes in tail length, the majority of variation of tail risk of 81.84% is caused by jumps in k, and only

[7] The average standard error of the tail risk exponent is 0.4087 for four-year, 0.3175 for six-year, 0.2814 for eight-year, and 0.2633 for a ten-year rolling time window.

the remaining 18.16% by changes in tail returns. Surprisingly, these results are consistent among countries, lying within a narrow range. Similar to sovereign credit default swaps, the average jump size of k is also close to zero for all national stock indices, which indicates some form of mean reversion to a long-term tail percentile. We observe coherent results of the decompositions of national stock market indices and sovereign credit default swaps for the same lookback window.

Secondly, we compare the equity tail decompositions for longer lookback windows. We compute the tail statistics for eight national stock market indices using rolling time windows of six-[8], eight-[9] and ten-years[10]. The stepwise increase of the time windows by two years, increases the tail length by 34.24%, 25.25% and 23.20% for each step, and simutanously decreases by 22.31%, 31.14% and 35.58% the standard error of the equity tail risk exponent. Also, the standard deviation of the standard error decreases by 14.43%, 27.67% and 36.10%. For each of the three time windows, we compute 17,080 equity tail risk exponents, tail thresholds and lengths. We observe that equity tail risk exponent becomes less adaptable to changes with increasing time windows. This means that the frequency of changes in the equity tail risk exponent decreases from 2,427 to 1,920 (-20.89%) for an increase in the rolling time window from four to ten years. We find that the number of variations in equity tail risk caused by changes in tail returns and tail threshold, account for approximately 36% independently of the length of the time window. The number of variations in equity tail risk caused by changes in tail returns amount for approximately 61% independently of the length of the time window. Notably, these percentages only vary within a small range for different time windows, countries and asset classes, which suggest some form of universality.

Next, we perform the decomposition analysis for different time windows to quantify the magnitude of changes in tail risk due to variations in tail returns and tail threshold,

[8] We report the average equity tail risk exponent for a lookback window of six years and eight national stock market indices from 2003 to 2017 in alphabetical order: ATX α=4.1690, BEL 20 α=3.9261, CAC 40 α=3.6618, DAX 30 α=3.5762, IBEX 35 α=4.1858, ISQE All α=4.2386, MIB 40 α=3.6108, and PSI 20 α=3.8717.

[9] We report the average equity tail risk exponent for a lookback window of eight years and eight national stock market indices from 2001 to 2017 in alphabetical order: ATX α=3.7774, BEL 20 α=4.1228, CAC 40 α=3.6144, DAX 30 α=3.3565, IBEX 35 α=4.0118, ISQE All α=4.3654, MIB 40 α=3.4333, and PSI 20 α=3.9472.

[10] We report the average equity tail risk exponent for a lookback window of ten years and eight national stock market indices from 1999 to 2017 in alphabetical order: ATX α=3.4437, BEL 20 α=4.3744, CAC 40 α=3.6753, DAX 30 α=3.3344, IBEX 35 α=4.0116, ISQE All α=4.4537, MIB 40 α=3.5340, and PSI 20 α=3.8487. Note that the computational time increases exponentially with increasing time windows.

separately. We find that variations in the tail threshold explain between 90-93% of changes in the tail exponent, while variations in tail return only account for 7-10%. These results hold independently of the length of the time window. While the magnitude of different factors remains constant, the average absolute change in equity tail risk decreases by 63.20% from 0.3549 for the four-year time window to 0.1306 for the ten-year time window. This is also reflected by the decrease of the standard deviation of the equity tail risk exponent by 61.50%. The decomposition analysis indicates that longer time windows lead to a smoother transitioning of the equity tail risk over time.

Finally, we investigate the impact of abrupt changes in the tail threshold on the equity tail risk exponent for longer time windows. We observe that the number of abrupt changes of the tail length decreases with increasing lookback windows. We count 711 abrupt changes of the tail length for four-year, 659 for six-year, 659 for eight-year, and 570 for a ten-year time window. It is noteworthy that the relative number of abrupt changes in the tail length remains steady between 29% and 32% because also the number of changes in tail risk decreases with increasing time windows. An increase of the time window from four to ten years decreases the magnitude of the average absolute change in the tail exponent by 64.13%. The average absolute change in the tail exponent reduces from 0.4840 to 0.1736 in the presence of a jump. These results suggest a smoother transitioning of the tail length and tail exponent during different market situations for longer time windows. However, even when the average absolute change of the tail exponent in the presence of jump decreases with increasing time windows, the variation of the tail threshold still dominates changes in equity tail risk. The jumps account for 81.84%, 78.76%, 79.88% and 76.50% of the total variation in equity risk exponent for four-, six-, eight- and ten-years. On the other hand, in the presence of non-abrupt changes, the changes in tail returns dominate the total variation in equity risk. The changes in tail returns account for 78.98%, 83.09%, 80,24% and 78.51% of the total variation in equity risk for four-, six-, eight- and ten-years.

The benefits of longer lookback windows motivate us to explore a new sampling method for credit default swaps. This new sampling approach is built on aggregated returns of cross-maturity time series, which is briefly mentioned in Section 1.4.4. In the following section, we analyse changes in credit tail risk of cross-maturity time series.

2.4.3 Decomposition of Tail Exponents in Cross-Sectional Data

We have learned from the equity tail risk analysis that longer time windows increase the number of tail observations, reduce the sampling noise and lower estimation errors. The equity tail risk decomposition shows that longer lookback windows substantially decrease the number of jumps in tail length and reduce the magnitude of changes in the tail exponent, especially in the presence of jumps. A lower standard deviation of the tail risk exponent indicates a smoother transition of the tail risk over time. Unfortunately, we cannot directly utilise these benefits of longer time windows for individual time series of credit default swaps, because, before 2005, we exhibit stale prices and illiquid data. Therefore, to overcome the problem of shorter time series in credit markets, we propose a new sampling approach. The new sampling approach relies on cross-maturity data of individual credit default swaps on the same underlying asset. We construct a longer time series by aggregating returns of different maturities into a new time series. Then the new time series is used to perform the credit tail risk analysis and decomposition. We find that some benefits of longer time windows can be recovered by using cross-maturity sampling in credit markets.

Firstly, we construct the cross-maturity time series by aggregating returns of ten different maturities ranging from 1 to 30 years. For this study we select a lookback window of 252 trading days (one year) for each individual time series. Consequently, the aggregated time series has a length of 2,520 observations, which is equivalent to a time series with a length of ten years. We find that the average tail length (right tail) of aggregated time series is 188 observations (7.47%), which is an increase of at least twofold compared to the individual (original) time series in Section 2.4.1.[11] The average tail length of the cross-maturity time series is similar to the average tail length in equity markets (170 observations) with a look-back window of ten years. This implies that a main benefit of longer time series can be recovered using cross-maturity time series. Further, the standard error of the credit tail risk exponent for cross-maturity data decreases by 33.01% compared to the standard error of individual time series. This result is similar to the decrease of the standard error by 35.58% for longer time windows in equity markets. Two main advantages of longer time series could be recovered using cross-maturity time series. In addition to more tail information and lower standard errors, another significant advantage is that cross-maturity data capture more recent tail events as the lookback window is only one year. Hence, the tail risk estimator does not rely on tail events from the distant past with no causality to the current market situation.

[11]The number of upper tail even and percentages in increasing order: Germany has 140 tail observations (5.54%), Portugal has 152 tail observations (6.04%), Belgium has 183 tail observations (7.24%), Ireland has 200 tail observations (7.92%), Italy has 201 tail observations (7.98%), Austria has 203 tail observations (8.06%), France has 210 tail observations (8.34%), and Spain has 219 tail observations (8.67%). The average tail length across countries is 188 observations. The average positive tail lengths for individual credit default swaps is 85 observations (Section 2.4.1).

Secondly, we estimate the sovereign credit tail risk exponents for cross-maturity time series. The estimation approach is similar to the previous approach used for individual time series in credit and equity indices. The Kolmogorov-Smirnov distance metric estimates the optimal tail length (k_t) for the cross-maturity time series of 2,520 observations. Importantly, note that we do not estimate and aggregate the optimal tail length for each individual time series separately. We use the tail returns of the cross-maturity tail to measure the sovereign credit tail risk exponent utilising the Hill (1975) estimator. From January 2009 until March 2017, we estimate 2,135 successive daily tail statistics. Using a daily rolling time window, we replace the oldest ten returns every day. In our sample of 17,080 sovereign credit tail exponents, the tail risk exponent changes in 7,222 cases. Compared to the individual time series (15.63%), the cross-maturity tail risk measure (42.30%) indicates a higher adaptability to changes in tail risk. Even with a higher level of adaptability, the average standard deviation of the tail risk exponent is approximately equal for the cross-maturity and individual time series.

Thirdly, we decompose tail exponent of the cross-maturity time series and draw comparisons to our previous results. We find that in 45.43% (3,281) of all cases, fluctuations of the dynamic sovereign credit tail risk are only caused by changes in tail returns and not due to changes in the threshold value. In only 1.56% (113) of all events, the entire fluctuations of the dynamic sovereign credit tail risk are caused by variations of the tail threshold. These variations occur because of returns close to, but below the threshold value change, which then trigger small movements of the tail threshold itself. However, in 53.01% (3,828) of all occurrences, the fluctuations of the dynamic sovereign credit tail risk coincide with changes of the tail threshold and tail returns, and it becomes difficult to differentiate between these two effects. Unlike our previous findings for individual credit default swaps and equity indices, the majority fluctuations of the dynamic sovereign credit tail risk are due simutanous changes of the tail threshold and returns.

We focus our analysis on fluctuations of the dynamic sovereign credit tail risk where changes in tail threshold coincide with changes in tail returns. The average absolute change of the tail exponent is 0.1952 across all eight countries with some regional differences.[12] To understand the variations in the sovereign credit tail risk exponent due to both factors separately, we decompose the changes of the tail risk exponent into its single components. Firstly, we calculate the daily variation of the sovereign credit tail risk due to changes in tail

[12]The average absolute change of dynamic sovereign credit tail risk where changes in tail threshold coincide with changes in tail returns in increasing order: Italy 0.1082, France 0.1286, Austria 0.1590, Germany 0.1700, Belgium 0.1829, Ireland 0.2308, Spain 0.2473, and Portugal 0.3348.

returns, by varying the set of returns from t to $t+1$ but keeping the threshold value constant for both consecutive dates. Secondly, we calculate the daily variations of the sovereign credit tail risk due to changes in the tail thresholds, by varying the threshold values m_t to m_{t+1} but using the same set of cross-maturity returns at $t+1$. Our tail exponent decomposition finds that only 18.48% (0.0361) of changes in dynamic credit tail risk exponent are attributed to changes in tail returns, which is similar to our previous results. Fluctuations in tail threshold account for 81.52% (0.1591) of the changes in dynamic credit tail risk exponent. When changes in tail threshold coincide with changes in tail returns, then the average absolute change in the tail risk exponent is 5.49%. This is significantly larger compared to such cases where the tail risk exponent only changes due to change in tail returns (0.52%). Note that these results are averages of the decomposition analysis of 3,828 tail exponent of eight countries using cross-maturity sampling.

Finally, we analyse the fluctuations of sovereign credit tail risk due to abrupt and non-abrupt changes of the tail length. The analysis is similar to the previous ones, but with a small twist. When the rolling time window moves from t to $t+1$, the oldest ten observations are dropped (one return per maturity), and the latest returns on date $t+1$ are added to the new subsample. If the dropped returns are non-tail events, and all new returns for each maturity are tail events, then the tail length Δk changes by +10 observations. The other extreme is that the ten oldest observations are tail events and the newest ones are non-tail events. In that case, the tail length Δk of the cross-maturity time series changes by -10 observations. In the daily time-varying approach, variations in the tail between -10 to +10 observations are considered as non-abrupt changes in tail length. Variations smaller or larger than this range, are considered as abrupt changes in tail length or jumps. The abrupt changes account for 20.22% (1,460) and non-abrupt changes for 79.78% (5,762) of the variation in the tail length, given that there is a change in the tail exponent. The percentage of non-abrupt changes increases by 9.47% (from 70.31%) and by 13.65% (from 66.13%) compared to equity (ten-year time window) and individual time series (four-year time window). The decrease in jumps suggests a smoother transitioning of the tail length over time.

In the presence of non-abrupt changes of the tail length, 65.59% of the tail risk variations are explained by changes in the tail returns, and 34.41% by changes in tail length. Compared to previous results, the fluctuations in tail risk due to changes in tail length is higher, because, in the cross-maturity analysis, non-abrupt events have a wider range of $\Delta k = [-10, +10]$. As discussed previously, the Kolmogorov-Smirnov distance metric estimation causes abrupt changes in the time-varying tail length. We face the same challenge for cross-maturity time series. The decomposition analysis reveals that in the presence of abrupt changes in tail length, 75.73% of the variation in tail risk is attributed to jumps and only 24.27% due to changes in tail returns, which is almost identical to the results of the decomposition analysis of individual credit default swaps. Interestingly, the average jump size is close to zero for all countries. This indicates that the effect of jumps might be only temporary and there is also mean reversion to the long-term tail percentile for cross-maturity data.

2.5 Conclusion

We provide evidence that risk changes occur on average three times per month in credit and equity markets. At least one these changes is caused by simultaneous changes of the tail returns and length. This challenges economists to differentiate between these two effects. To overcome this challenge, we presented a new decomposition method of the dynamic tail risk that quantifies these factors independently. We find that approximately 10% of tail risk changes can be explained through fluctuations in tail returns. This implies that a majority of tail risk changes might be caused by statistical origines inherent to the dynamic Kolmogorov-Smirnov method.

The tail length does not always transition smoothly. We observe that in approximately 30% of all cases, the tail length changes abruptly. While approximately 20% of the variation in tail risk can be explained through fluctuations in tail returns, the majority is caused by jumps in tail length. These findings are consistent among credit and equity markets.

We perform the tail risk decomposition for different time windows in equity markets. We provide evidence that longer time windows leads to more stable estimation of the dynamic tail risk exponent and decreases the number of abrupt changes in the tail length. While longer time windows capture more tail events, the frequency of tail risk changes decreases, which indicates a lower potential to adopt to new market situations.

The benefits of longer time windows motivated us to explore a new sampling method based on aggregating returns for different maturities of the same underlying asset. The cross-maturity approach captures more tail returns than individual time series, which reduces estimation errors. Estimations consider more recent returns which results in a higher adaptability without increasing standard deviation of the tail risk exponent. The decomposition analysis reveals that cross-maturity time series substantially decrease the number of jumps in tail length by approximately 40%.

Chapter 3

Asset Pricing and Tail Risk in US Sovereign Credit Default Swaps

3.1 Introduction

The objective of this chapter is to explore the effects of time-varying credit tail risk in asset markets. The major obstacle to this investigation is a sound measure of credit tail risk over time. Ideally, one would directly construct a tail risk measure from the underlying credit default swap time series. However, modelling time-varying tail risk in univariate time series is challenging because of the infrequent nature of extreme events. A sufficiently large number of tail events might be found in long time series. However, long time series may include tail events from the distant past (previous crises) with no causality to the current market situation. It may become problematic to draw conclusions from the evolution of the tail exponent to recent changes of financial or economic variables or to use it for forecasting purposes.

To overcome this problem, we conceive a cross-maturity estimation approach that captures short- and long-term risk dynamics across the term structure of default. Power-laws have very favourable aggregation properties, such as taking the sum of two (independent) power-law distributions gives another power-law distribution (Gabaix, 2016).[1] Hence, if tail distributions for single maturities possess similar tail behaviour, then the cross-maturity

[1] Gabaix (2009) provides an overview of aggregation properties for variables with power-law tails. The property of being distributed according to a power-law is conserved under addition, multiplication, polynomial transformation, min, and max. A general rule is that, when two power-law variables are combined, the fattest power-law (highest risk, smallest tail exponent) dominates. Jessen and Mikosch (2006) provides further details and derivations of power-law properties.

returns can be used to identify the common factor of their tail risk at each point of time.[2]

Our framework is built on a reduced-form description for the upper tail distribution of returns. Let $x_t = (x_{1,t}, ..., x_{m,t})$ denote the cross section of returns for m maturities in period t. Given time t, the upper tail distribution is defined as a set of returns falling above the threshold value u_t. We assume that the upper tail of credit default swap returns follows a continuous power-law distribution for the tail observations ($0 < u_t < x_t$). It is described by a probability density $p(x)$, such that

$$p(x_t) = \frac{\alpha_t - 1}{u_t} \left(\frac{x_t}{u_t} \right)^{-\alpha_t}, \tag{3.1}$$

where α_t is the tail exponent and u_t is the threshold value of the tail at time t. The tail exponent and threshold value are the key parameters of the credit tail risk model, which determine the shape of the tail. The credit tail risk exponent fundamentally depends on the upper tail threshold of u_t, which defines the number of extreme events, and hence, the size of the tail percentile. Low values of the tail risk exponent α correspond to "fatter" tails and a higher probability of extreme events.

In contrast to previous power-law research, Equation 3.1 is a dynamic model of cross-maturity tail returns with three time-varying parameters.[3] The tail risk exponent α_t, threshold value u_t and tail percentile ρ_t vary with the set of return information at t. Although different maturities can have different levels of tail risk, dynamics are the same for all credit default swaps, because they are driven by the common process α_t. Thus, we refer to α_t as "credit tail risk" exponent at time t, and we refer to the tail structure in Equation 3.1 as the "dynamic power-law" model.

We build the sovereign credit tail risk measure from the dynamic power-law model in Equation 3.1. We derive the threshold value of u_t in two steps. First, we estimate the mo-

[2]We utilise the excellent aggregation properties of power-laws to construct the term structure of credit tail risk in Chapter 4. We find the same level of credit tail risk independently of credit default swap maturities.

[3]Power-law research in finance and economics often uses heuristic rules to determine the tail length. A commonly used technique in quantitative finance is to take a fixed percentage of the total sample. Doyne Farmer *et al.* (2004), for instance, use the largest $\sqrt{n}$, where n is the size of the sample. Kelly and Jiang (2014), for instance, consider the 5% sample fraction to estimate the tail risk using cross-sectional returns of US stocks. Fix tail percentiles have the advantage that tail risk fluctuations are only attributed to changes in tail returns, not due to variations in tail length. However, fixed percentages are somewhat arbitrary. Different distributions have different optimal tail lengths, which vary over time and with the risk of the underlying asset (see, Chapter 1). Our approach allows for time-varying tail thresholds, lengths and percentiles.

mentary threshold value, denoted m_t, using the time-varying Kolomogrov-Smirnov distance metric. The momentary threshold value is determined such that the probability distribution of the empirical data and the best-fit power-law model become as similar as possible above initial threshold value m_t. From the fitted momentary threshold value m_t we can subsequently determine the tail length k_t. The ratio k_t/n_t, where n_t is the length of the emprical data set, defines the tail percentile ρ_t. To balance fluctuations of momentary threshold value caused by statistical origins inherent to time-varying Kolmogorov-Smirnov distance metric, we calculate an equally weighted mean of ρ_t from $t-h$ to t, where h determines the length of h-day average. Second, the threshold value u_t can be found using information of the average tail percentile ρ_t^* by backward induction. The average tail percentile ρ_t^* defines the average time-varying tail length k_t^*. From k_t^*, we derive the optimal time-varying threshold value u_t at time t. Then the Hill (1975) estimator determines the sovereign credit tail risk exponent α_t given u_t.

We find that the sovereign credit tail risk is persistent over time. We estimate α_t for each month and observe a monthly AR(1) coefficient of 0.967. The result indicates that the credit tail risk exponent might have significant predictive power for spreads of default insurances. Thus, the estimated persistence of tail shocks is offering the first hint that α_t is a potentially important factor of equilibrium prices. We find that the lower and upper tail dynamics have a correlation of 71% . Both tails together explain on average 42% of the future US CDS spreads.

Our first contribution is an empirical analysis of the credit tail risk on credit default swap spreads. We test the hypothesis that sovereign credit tail risk forecasts credit default swap spreads. Predictive regressions show that a one-standard-deviation increase in the sovereign credit tail risk predicts an increase in future US credit default swaps spreads of 3.4, 4.4, 5.6, 6.8, 8.2, 8.4, 8.8, 10.3, 9.9, and 9.9 basis points (bps) based on 1-, 2-, 3-, 4-, 5-, 7-, 10-, 15-, 20-, and 30-years to maturity. The corresponding t-statistics are 3.1, 4.3, 5.4, 6.4, 7.3, 8.3, 8.2, 10.4, 8.5 and 8.4. On average, a one-standard-deviation increase in sovereign credit tail risk predicts an increase in future US credit default swap spreads of 7.6 bps, which is highly significant, with a t-value of 7.0. The credit tail risk achieves impressive levels of predictability, reaching an average R^2 value of 34%. These results are robust out-of-sample. Estimated credit tail risk coefficients and their statistical significance are robust to controlling for a large set of alternative predictors. We find that on average, the sovereign credit tail risk measure dominates alternative predictors in bivariate predictive regressions of the credit default swap spreads. Noteworthy, when the credit tail risk is combined with the Smoothed US Recession Probability (Chauvet, 1998) it attains the highest levels of predictability for short-dated CDS by reaching R^2 values above 30%. On the contrary, for longer-dated contracts, the credit tail risk and the US Economic Policy Uncertainty Index achieve impressive levels of predictability, reaching R^2 levels above 60%.

This paper draws on several strands of literature. Since at least Mandelbrot (1963) and Fama (1963), economists have discussed that unconditional return distributions are heavy-tailed and well described by a power-law probability distribution. Gabaix (2009) seminal paper documents that various economic and financial time series exhibit distributions with a power-law decay (i.e. income and wealth, the size of cities and firms, stock market returns, trading volume, international trade, and executive pay). Numerous studies report power-laws across asset classes, such as foreign exchange rates (Guillaume *et al.*, 1997), individual stocks (Plerou *et al.*, 1999), financial market indices (Gopikrishnan *et al.*, 1999), trading volume (Gabaix *et al.*, 2003), cross-sectional equity returns Kelly and Jiang (2014), and cryptocurrencies (Begušić *et al.*, 2018). More recent empirical work suggests that tail returns of credit default swap markets follow a power-law with time-varying tail exponent (see, Chapter 1).[4] We show that empirical studies of time-varying tail behaviour and asset prices are closely linked.

The second strand of literature addresses is an ongoing debate about the number of tail observations that have to be used in the estimation of the power-law tail exponent. In finance, the power-law only holds true for a small fraction of the time series above a certain threshold value. The literature to date considers three groups of estimation methods for the tail threshold: (i) heuristic rules such as the Eye-Balling method or arbitrary (constant) percentiles[5], (ii) minimising the mean squared error of the tail exponent estimator in the prob-

[4]Previous research documents time variation in the tail behaviour in equities and futures markets, i.e. Quintos *et al.* (2001); Galbraith and Zernov (2004); Werner and Upper (2004); Wagner (2003); Kelly and Jiang (2014).

[5]The techniques of finding the tail length by observing the stable regions in the Hill plot are known as the Eye-Balling method or the automated form of the Eye-Balling method (Resnick and Stărică, 1997). Another heuristic method is to limit the tail analysis to arbitrary percentiles. For example, Plerou *et al.* (1999) and Gopikrishnan *et al.* (1999) only use returns larger than two, three or five standard deviations or within a range of standard deviations. Doyne Farmer *et al.* (2004) limit the analysis to the most significant observed returns only, such as the largest $\sqrt{n}$ or $1/10n$. More recently, Kelly and Jiang (2014) utilise a fixed tail percentile at a 5% level for the estimation of tail risk using cross-sectional stock returns. The advantage of fix tail percentiles is that fluctuations in tail risk are only attributed to changes in tail returns, and not due to variations in tail length. However, Chapter 1 shows that the tail length fluctuates with the risk of the underlying asset. Furthermore, heuristic methods have a weak theoretical foundation and underestimate model uncertainty (Stoev *et al.*, 2006).

ability dimension[6], and (iii) minimising the maximum deviation in the quantile dimension. In a simulation study Danielsson *et al.* (2016) demonstrate that minimising the Kolomogrov-Smirnov distance between the fitted Pareto type tail and the observed quantile, performs best. However, if the tail length (tail percentiles) changes over time, extreme value techniques based the Kolomogrov-Smirnov distance may result in jumps in the tail risk exponent due to abrupted changes in tail length and not due to significant changes in tail returns (see, Chapter 2.4). In this paper, we present a simple model extension to overcome this issue.

The most popular estimator for the tail exponent of heavy-tailed distributions is the Hill (1975) estimator. Various other estimation methods have been proposed to measure the tail behaviour, such as the asymptotic estimate constructed for the index of a stable distribution with convergence at a logarithmic rate convergence by de Haan and Resnick (1980); the estimates of an exponent of regular variation with convergence at an algebraic rather than a logarithmic rate by Hall (1982); the seminal paper on the tail estimation of distributions with exponential-like upper tails by Mason (1982); the estimation approaches from the classical extreme value theory Pickands *et al.* (1975) and Davis and Resnick (1984); the kernel estimator approach by Csorgo *et al.* (1985); and the adoptive estimator by Hall *et al.* (1985). Recent research in statistics of extreme values shows that the Hill (1975) estimator is consistent even in the presence of dependent and heterogeneous data (Kelly, 2014) and is, therefore, our preferred method of choice.

The rest of the chapter is organized as follows. In Section 3.2, we introduce the methodology for the dynamic power-law model to estimate sovereign credit tail risk from cross-maturity data. Section 3.3 describes the data and selection criteria. Section 3.4 reports the empirical results of the tail risk measure for sovereign credit default swaps in US credit markets. Section 3.5 concludes.

[6]Several procedures have been introduced for choosing the optimal tail length k in the sense of asymptotic minimal mean squared errors such as Dekkers and de Haan (1993) and Beirlant *et al.* (1996). Hall (1990) and Danielsson *et al.* (2001)) employ bootstrap procedures to minimises the asymptotic mean squared error. Drees and Kaufmann (1998) exploit the same bias and variance trade-off, but use the maximum random fluctuation of the estimator to locate the point where the trade-off is optimal. These methods, based on minimisation of the mean squared error in the probability dimension are asymptotically consistent, but have unsatisfactory finite sample properties.

3.2 Empirical Methodology

The Sovereign Credit Tail Risk Model

We begin our methodology with a brief definition of the basic quantities involved in the credit tail risk estimation. We aim to fit a power-law distribution over the upper tail of the cross-maturity log returns of US sovereign credit default swaps. Let x_t represent the quantity whose distribution we want to analyse in period t. For most financial time series, the power-law only holds for a fraction above a certain threshold value u_t. The threshold value is the smallest tail return of the cross-maturity time series. Recall from Equation 3.1 that the dynamic power-law model for cross-maturity returns is defined as

$$p(x_t) = \frac{\alpha_t - 1}{u_t} \left(\frac{x_t}{u_t} \right)^{-\alpha_t},$$

where $\alpha_t > 1$ is the tail exponent for the returns x_t above u_t at time t. Firstly, we estimate the credit tail risk exponent α_t. Estimating α_t correctly requires the threshold value u_t, which also defines the length of the tail (k_t) and the upper tail percentile (ρ_t). For the moment, we assume that the value of u_t is known. The method of choice for fitting the dynamic power-law model to the observed tail returns is the method of maximum likelihood, which provably provides accurate parameter estimates in the limit of large numbers (Barndorff-Nielsen and Cox, 1994; Wasserman, 2013). Assuming that the tail returns are drawn from a distribution that follows a power-law for $x_t > u_t$, we can derive the maximum likelihood estimator of the sovereign credit tail risk exponent, such that

$$\frac{1}{\hat{\alpha}_t} = \frac{1}{k_t} \sum_{i=1}^{k_t} \ln \frac{x_{i,t}}{u_t} \tag{3.2}$$

where $x_{i,t}$, $i = 1, ..., k_t$ are the numbers of tail returns in period t. In the non-time-dependent version, this equation version is known as the Hill (1975) estimator. Therefore, we refer to this equation as the time-varying Hill estimator. Research in theoretical statistics reports that the Hill (1975) estimator is known to be asymptotically normal (Hall, 1982) and consistent (Mason, 1982) (i.e., $\hat{\alpha}_t \to \alpha_t$ in the limit of large k_t).

Secondly, the calculation of the tail threshold (u_t) consists of two steps. First, the Kolmogorov-Smirnov distance method defines the time-varying tail threshold (m_t), which is called momentary tail threshold in this chapter.[7] The fundamental idea behind the time-varying threshold value (m_t) is simple. Given a subsample of the cross-maturity time series, we select the threshold value ($\hat{m}_t$) that makes the probability distributions of the empirical data and the best-fit power-law model as similar as possible above the tail threshold $\hat{m}_t$ at time t. There are a range of methods for quantifying the distance between two probability distributions, but for non-normal data the commonest is the Kolmogorov-Smirnov statistic (Press *et al.*, 1992). The Kolmogorov-Smirnov test functions as the distance metric, defining maximum distance between the complementary distribution functions (CDFs) of the empirical data and the fitted model, such that

$$D_t = \sup_{\{x_t > m_t\}} |F(x_t) - P(x_t)|, \tag{3.3}$$

where $\sup_{\{x_t > m_t\}}$ is the supremum of the set of distances. Here $F(x_t)$ is the empirical CDF for cross-maturity returns larger than momentary threshold m_t, and $P(x_t)$ is the CDF for a power-law model that best fit the data over the tail region $x_t > m_t$ in period t. The estimate $\hat{m}_t$ is then the value of m_t that minimises the distance D_t at time t. From the fitted tail above $\hat{m}_t$ we can subsequently determine the optimal length k_t. Then, the tail percentile is defined as the number of tail returns over the length of the time series, such that $\rho_t = k_t/n_t$. Note that the number of cross-maturity returns for sovereign credit default swaps is constant $n_t = n$ for the rolling time window in each period t. The length of the rolling time window is 2,520 days for cross-maturity data (each of the ten maturities has a lookback window of 252 days).

To balance jumps of m_t caused by statistical origins inherent to time-varying Kolmogorov-Smirnov distance metric, we calculate a new tail threshold u_t based on a simple moving average of the tail percentile ρ_t.[8] This second step of the calculation is straightforward. We calculate the equally weighted mean, denoted ρ_t^*, from $t-h$ to t

$$\rho_t^* = \frac{1}{h}\sum_{i=1}^{h} \rho_{t+1-h} \tag{3.4}$$

where h determines the length of the h-day average. For $h = 1$, ρ_t is the momentary tail percentile on day t. Then, the new threshold value u_t can be derived using the ρ_t^* by backward induction. The value of ρ_t^* implies the average time-varying tail length k_t^*. Then, the value

[7] The estimation of the tail threshold (m) is based on Kolmogorov-Smirnov distance metric (Clauset *et al.*, 2009). Chapter 1 extends this approach to a time-varying measure of the tail threshold (m_t).

[8] Using the decomposition method for cross-maturity time series from Section 2.4.3, we observe 202 abrupt changes in tail length for US credit default swap data.

of k_t^* determines the smallest tail return $x_{k^*,t}$, which is the cutoff point between the body and the tail of the distribution. This cutoff point is the new tail threshold value u_t. We choose h=126 trading days (half a year of trading). This approach reduces unwanted fluctuations in the tail exponent caused by abrupt changes in tail length.

The sovereign credit tail risk model has some significant advantages over other commonly known power-law models. Firstly, the sovereign credit tail risk model takes advantage of favourable aggregation properties of power-laws, which allows for the aggregation of returns across the term structure of default. Aggregated return data can significantly reduce the lookback window and dependency on data from the distant past.[9] Consequently, shorter time windows lead to higher adaptability to changes in credit tail risk.[10] Using cross-maturity instead of univariate time series, significantly increase the number of tail observations and reduce estimation errors in tail risk exponents.[11] Furthermore, using cross-maturity data lead to a smoother transitioning of tail risk over time, because tail risk changes occur more frequently and less often abrupt compared to univariate data.[12] Secondly, our sovereign credit tail risk model incorporates a simple smoothing technique for the time-varying tail threshold. The main advantage is that the smoothing method successfully balances abrupt changes in tail thresholds, which eliminates unwanted fluctuations in credit tail risk caused by abrupt changes in tail length.[13] Eliminating tail risk fluctuations caused by abrupt changes in tail

[9]Section 2.4.3 provides evidence that the cross-maturity sampling allows reducing the lookback window from ten years in univariate time series to one year in cross-maturity time series without loss of tail information. The average tail length of the cross-maturity time series is 237 observations of US sovereign credit default swaps using a time window of one year for each maturity. Using the same amount of data points (2,520 returns) in univariate time series, we observe on average 170 tail observations in equity markets (see, Section 2.4.2).

[10]A higher level of adaptability is measured by the average number of tail risk changes. Cross-maturity time series using m_t with a time window of one year record 990 (46.37%) changes in US credit tail risk from 2009 to 2017. Whereas univariate time series using m_t with a time window of ten years count on average 240 (11.15%) changes in equity tail risk from 2009 to 2017 (see, Section 2.4.2).

[11]The average positive tail length is 237 observations (9.40%) for cross-maturity time series with a one-year time window. The average positive tail length is 85 observations (8.43%) for individual credit default swaps using univariate time series with a four-year time window (see, Section 2.4.1).

[12]A smoother transitioning of tail risk is measured by a decrease in the average absolute change in tail risk exponent from 0.1311 for the univariate time series to 0.1185 for the cross-maturity time series. Given a change in tail length ($\Delta k \neq 0$), abrupt changes in tail length occur less frequently in cross-maturity time series (20.40%) compared to univariate time series with a four-year (33.86%) and ten-year (29.69%) time window.

[13]The smoothing technique reduces the average absolute change in tail length from 19.34 to 2.30 observations and narrows the range of daily changes in tail length to $\Delta k = [-7; +7]$. Note that only changes outside the range of $\Delta k = [-10; +10]$ are considered as abrupt changes in tail length (see, Section 2.4.3). Furthermore, the smoothing technique reduces the standard deviation of the tail length by 19.85% and the standard deviation of

length has an additional smoothing effect on the tail risk exponent without compromising on higher tail risk adaptability initially gained through tail return aggregation.[14]

3.3 Corporate and Sovereign Credit Default Swap Data

We estimate the dynamic tail exponent using daily returns of corporate and sovereign credit default swap from January 2008 to March 2017. IHS Markit Ltd. provides the data for this investigation. A group of leading market participants (including the G16 banks) contribute corporate and sovereign credit default swap data on a daily basis.[15] Based on this information IHS Markit computes the daily composite quotes. IHS Markit requires at least three distinct contributors submitting price information in order to form a composite.

To ensure data quality, corporate and sovereign credit default swap data is subject to a series of tests. These tests comprise so-called 'logical' and 'relative' assessment of the data quality. The first logical test is the curve buildability test which has great importance for our term structure of credit tail risk. The curve buildability test checks for valid survival and default probabilities using the bootstrapping method in the ISDA CDS Standard Model. Credit default swaps with unreasonable changes of default probabilities are rejected. Having an accurate term structure of the default probabilities is essential to build the term structure of tail risk in Chapter 4.

Credit default swaps trade under four common restructuring types, namely the Cum Restructuring (CR), Modified Restructuring (MR), Modified Modified Restructuring (MM), and Ex-Restructuring (XR). The second test is the backwardation test, which evaluates the relationship between restructuring types. The backwardation test validates the following inequality for the contributed spreads for each maturity: $CR \geq MM \geq MR \geq XR$. This implies that credit default swaps with restructuring clause (CR) have a higher price compared

the tail threshold by 45.57%.

[14] Using u_t instead of m_t eliminates tail risk fluctuations caused by abrupt changes in tail length, which has an additional smoothing effect on the tail risk exponent. The smoother transitioning of the tail risk is measured by a decrease in the average absolute change in US credit tail risk from 0.0750 for m_t to 0.0144 for u_t. This also reduces the standard deviation of the tail risk exponent by 30.63%. Another advantage of the smoothing technique is that the average estimation error of the credit tail risk exponent decreases by 20.51%.

[15] The G16 is an industry group comprising the largest derivatives dealers: Bank of America-Merrill Lynch, Barclays Capital, BNP Paribas, Citi, Crédit Agricole, Credit Suisse, Deutsche Bank, Goldman Sachs, HSBC, JP Morgan, Morgan Stanley, Nomura, RBS, Societe Generale, UBS and Wells Fargo Bank.

to credit defaults without restructuring (XR). Bai and Wei (2012) discusses the standard convention regarding the structuring clauses and transaction currencies. They state that sovereign credit default swaps usually trade under the Cum Restructuring clause. We use credit default swaps quoted under the Cum Restructuring clause, for both, corporate and sovereign credit default swap data.

The relative tests deal with price liquidity and outliers. The stable data test measures the frequency of price updates and liquidity. An entity is defined to be liquid if it receives contributions from 13 or more contributing banks and market maker; otherwise it is defined to be illiquid. If traders tend to update an entity infrequently or not at all, the data become stale. Illiquid and stale price information are excluded from this data set.

Furthermore, the outlier test removes data which are genuinely different to the value of the credit default swap. The outlier test computes a provisional median spread based on all submitted quotations for each tenor that passed the curve buildability, backwardation, and stale data tests. The outlier test then calculates a weighted sum of squared deviations across maturities from the provisional median. Finally, it then ranks the submitted quotations and rejects the ones with highest deviations.

Credit default swap data that pass the qualitative and relative tests are admitted in this study. To be included in the dynamic tail risk analysis, sovereign and corporate credit default swaps are subject to a range of criteria. We only consider credit default swaps for senior unsecured debt traded under the Cum Restructuring clause.[16] We impose the requirement that an entity has price information on at least 75% of all trading days between January 2009 and March 2017.

[16]The Markit dataset consists of four different seniority levels of the debt within the capital structure: senior, subordinated, junior and preferred. They state that sovereign credit default swaps usually trade under the Cum Restructuring clause. We consider price information for senior unsecured debt for all sovereign and corporate credit default swaps.

3.4 Empirical Results I: Sovereign Credit Tail Risk

3.4.1 US Credit Tail Risk Estimates

We estimate the dynamic credit tail risk exponent using short- and long-term risk dynamics across the term structure of default of US sovereign credit default swaps from January 2008 to March 2017. In order to construct the aggregated time series, we select a lookback window of 252 trading days (one year) for each credit default swap maturity, which is unchanged over time. Consequently, aggregated time series have 2,520 observations, which record sufficient tail information to estimate the tail statistics at each point of time ($k_t > 50$). We estimate 2,135 successive daily tail statistics from January 2009 until March 2017. We find that the dynamic power-law model is a feasible model for cross-maturity data and accept the power-law hypothesis for positive and negative tail returns. On average the smallest positive tail returns is 6.14% and the average tail percentage is 9.41%. The average sovereign credit tail risk exponent is 2.41 (average standard error is 0.12), which is at disagreement with the inverse cubic law ($\alpha \simeq 3$) found for univariate and cross-maturity time series of European sovereign credit default swaps (see, Chapters 1 and 2).[17] This means US credit default swaps exhibit a heavier positive tail than European credit default swaps. Important for risk modelling, the US tail risk exponent is outside the Lévy-stable region $0 < \alpha < 2$, which implies the existence of a finite second moment (i.e. variance).

The difference between the average positive (2.41) and negative (2.46) tail exponent is not significant at a 99% confidence level. US sovereign credit default swaps have similar probabilities of extreme returns above the positive and negative tail threshold value. However, similar probabilities of tail events do not provide information about the magnitude of positive and negative tail returns. While we associate positive tail returns with large price increases usually occurring in times of financial turbulences, negative tail returns are related to large price decreases usually occurring in periods of financial stabilisation. Firstly, we observe that the positive tail exhibits a higher average tail threshold value at a 99% confidence level. Secondly, the average extreme return of the positive tail (13.94%) is more significant than the negative one (11.34%). We confirm that extreme returns are more profound in the upper tail of the distribution even for same threshold values for both tails ($u = [0.05, 0.10]$).

[17]The sovereign credit tail risk measure using cross-maturity data of European sovereign credit default swaps from January 2009 to March 2017 in increasing order: France 3.10, Austria 3.16, Italy 3.37, Belgium 3.45, Germany 3.46, Ireland 3.58, Spain 3.74, Portugal 4.07. Note that these results are based on estimations using the momentary threshold value m_t and a rolling time window of 252 trading days for each maturity.

3.4.2 Predicting US Sovereign Credit Default Swap Spreads

We test the hypothesis that credit tail risk forecasts US sovereign credit default swap spreads with a series of predictive regressions. All regressions are conducted at a monthly frequency. The dependent variables are the US sovereign credit default swap spreads (1-30 years) at a monthly frequency. To illustrate economic magnitudes, all reported predictive coefficients are scaled to be interpreted as the effect of a one-standard-deviation increase in the regressor on future credit default swap spreads. Tables 3.1 and 3.2 show that the sovereign credit tail risk measure has large and significant forecasting power over all maturities. A one-standard-deviation increase in the upper tail risk predicts an increase in future credit default swap spreads of 3.4, 4.4, 5.6, 6.8, 8.2, 8.4, 8.8, 10.3, 9.9, and 9.9 basis points (bps) based on 1-, 2-, 3-, 4-, 5-, 7-, 10-, 15-, 20-, and 30-years to maturity. The corresponding t-statistics are 3.1, 4.3, 5.4, 6.4, 7.3, 8.3, 8.2, 10.4, 8.5 and 8.4. On average, a one-standard-deviation increase in tail risk predicts an increase in future US credit default swap spreads of 7.6 bps, which is highly significant with a t-value of 7.0. The credit tail risk measure achieves impressive levels of predictability, especially for longer-dated maturities. Across all maturities, the tail risk measure reaches an average R^2 value of 34%. The credit tail risk measure has significant forecasting power over longer time horizons. A one-standard-deviation increase in the tail risk predicts an increase in future average credit default swap spreads of 5.9, 4.4 and 2.8 bps based on data for three-months, six-months and one-year, respectively. The corresponding t-statistics are 5.7, 4.4 and 2.8.

Tables 3.1 and 3.2 compare the forecasting power of credit tail risk measure with a large set of alternative economic variables from the monthly publications from the Federal Reserve Bank of St. Louis, monthly historical US housing market data from Shiller (2015), and forecasting variables studied in a survey by Welch and Goyal (2008).[18] On an average level, the US Economic Policy Uncertainty Index, developed by (Baker *et al.*, 2016), is the only out of 25 predictors with a better average performance compared to the proposed credit tail risk measure, reaching a R^2 level of almost 40%. However, the US Economic Policy Uncertainty Index does not outperform the credit tail risk measure for some longer maturities. For short-dated credit default swaps (1-5 years) some economic variables generally perform well, such as Commercial Bank Credit, Total Consume Credit, and Industrial Production Index. A one-standard-deviation increase in measures of uncertainty (i.e. S&P 500 Volatility, VIX Index, US Economic Policy Uncertainty Index, Smoothed US Recession Probabilities) strongly increases future credit default swap prices for short-dated contracts (1-5 years). For

[18] Amit Goyal provides updated data through 2019.

longer-dated contracts, only the Real Housing Price Index and the US Economic Policy Uncertainty Index possess similar predictive power as the credit tail risk estimate. Interestingly, a one-standard-deviation increase in the Real Housing Price Index decreases future credit default swap prices, which is counterintuitive. We conclude that the credit tail risk yields strong predictability, especially for long-dated maturities.

We next carry out bivariate regressions using our tail risk measure alongside each predictor to assess the robustness of the tail risk's forecasts after controlling for alternative economic variables. Table 3.3 shows the results for credit tail risk measure for the most liquid maturities (five- and ten-year) after considering each predictor, as well as the average values across all ten maturities.[19] On average, the credit tail risk has significant predictive power across all credit default swap maturities with a t-statistic above 3.4. The average and cross-maturity coefficient remains above 6.7 compared to 7.6 in the univariate case. The overall level of predictability is lower for one- and two-year maturities (average predictability of only 14% and 22%). For these two maturities, the credit tail risk, when combined with the Smoothed US Recession Probability by Chauvet (1998) attains the highest levels of predictability, reaching R^2 values of 32% and 37%.[20] Credit tail risk, when combined with with the US Economic Policy Uncertainty Index, achieves impressive levels of predictability, reaching an average R^2 value of 50%, which is the highest among all economic variables. The index spikes near tight presidential elections, Gulf Wars I and II, the 9/11 attacks, the 2011 debt-ceiling dispute and other major battles over fiscal policy; events, which potentially have a long-term impact on the default risk of the United States.

The positive tail exponent has a correlation of 71% with the negative tail exponent. However, the upper tail risk dominates the lower tail risk in a bivariate predictive regression of the sovereign credit default swap spread. The tail dynamics of both tails together explain on average 42% of the future US credit default swap spreads across all maturities and performs particularly well for maturities longer than four-years.

[19]Five- and ten-year contracts are the most liquid maturities in terms of lower bid-ask spreads and trading volume as argued in Pan and Singleton (2008) and Berndt and Obreja (2010).

[20]The Federal Reserve Bank of St Louis publishes an economic indicator developed by Chauvet (1998) on a monthly basis called "Smoothed US Recession Probabilities." The probabilities are obtained from a dynamic-factor Markovswitching (DFMS) model applied to four monthly coincident variables: non-farm payroll employment, the index of industrial production, real personal income excluding transfer payments, and real manufacturing and trade sales. Note, it is important to specify predictive relations based on what market participants could have actually known at the time. For additional details and a related study using "real-time" dataset of coincident monthly variables, see (Chauvet and Piger, 2008).

US Credit Default Swap Spread Predictability: Univariate Predictor Performance (1/2)

		US - 01Y Maturity			US - 02Y Maturity			US - 03Y Maturity			US - 05Y Maturity		
		Coeff.	t-stat.	R^2	Coeff.	t-stat.	R^2	Coeff.	t-stat.	R^2	Coeff.	t-stat.	R^2
(1)	US Credit Tail Risk	3.42	3.06	8.9	4.43	4.32	16.3	5.55	5.40	23.3	8.20	7.26	35.5
(2)	Inflation	0.47	0.85	0.7	0.60	1.12	1.3	0.80	1.44	2.1	1.05	1.57	2.5
(3)	Default Return Spread	-2.22	-1.93	3.7	-2.15	-1.96	3.8	-2.01	-1.74	3.1	-2.36	-1.70	2.9
(4)	LT Government Yield	3.94	3.58	11.8	3.80	3.61	11.9	4.38	4.04	14.5	5.66	4.41	16.8
(5)	Risk Free Rate	1.21	0.91	0.9	0.75	0.59	0.4	0.40	0.30	0.1	-0.54	-0.34	0.1
(6)	Corporate Bond Yield (AAA)	4.53	4.21	15.6	4.32	4.19	15.5	4.85	4.55	17.7	5.85	4.59	18.0
(7)	Default Yield Spread (DFY)	-4.38	-4.02	14.4	-4.86	-4.81	19.4	-5.44	-5.23	22.2	-6.76	-5.49	23.9
(8)	Corporate Bond Returns (LT)	-1.23	-1.05	1.1	-1.20	-1.08	1.2	-0.88	-0.75	0.6	-0.88	-0.63	0.4
(9)	Term Spread	3.84	3.46	11.1	3.74	3.53	11.5	4.34	3.98	14.1	5.69	4.41	16.9
(10)	Rate of Return (Long Term)	0.04	0.03	0.0	0.27	2.49	0.1	0.61	0.52	0.3	0.25	0.25	0.1
(11)	Treasury Bills	0.08	0.06	0.0	-0.25	-0.21	0.0	-0.48	-0.38	0.2	-1.22	-0.81	0.7
(12)	Volatility S&P 500	5.05	4.80	19.3	5.76	6.02	27.4	6.25	6.34	29.5	6.81	5.56	24.3
(13)	US Recession Probabilities	6.15	6.22	28.7	5.90	6.25	28.9	5.96	5.94	26.9	6.13	4.87	19.8
(14)	Gold Future	-2.05	-1.78	3.2	-0.59	-0.52	0.3	0.39	0.34	0.1	2.17	1.56	2.5
(15)	VIX	4.51	4.16	15.3	5.48	5.62	24.8	6.45	6.60	31.2	7.96	6.90	33.2
(16)	Real Home Price Index	-1.40	-1.19	1.4	-3.03	-2.77	7.4	-4.89	-4.52	17.5	-8.66	-7.73	38.3
(17)	Commerical & Industrial Loans	-2.48	-2.14	4.5	-3.86	-3.62	12.0	-5.71	-5.50	23.9	-9.45	-8.99	45.7
(18)	Industrial Production Index	-3.67	-3.29	10.2	-4.47	-4.36	16.5	-5.80	-5.71	25.3	-8.41	-7.52	37.0
(19)	Commercial Bank Credit	-3.34	-2.91	8.1	-4.47	-4.25	15.8	-6.11	-5.96	27.0	-9.56	-9.05	46.1
(20)	Inflation Expectation (UoM)	-1.35	-1.16	1.4	-0.47	-0.42	0.2	0.33	0.28	0.1	1.92	1.37	1.9
(21)	Federal Surplus or Deficit	-1.59	-1.37	1.9	-1.99	-1.79	3.2	-2.70	-2.35	5.4	-4.03	-2.97	8.4
(22)	Economic Policy Uncertainty	3.55	3.19	9.6	5.22	5.30	22.6	6.90	7.34	35.9	9.94	10.21	52.1
(23)	Consumer Credit Cards	0.67	0.57	0.3	0.14	0.13	0.0	-1.44	-1.23	1.6	-3.94	-2.92	8.2
(24)	Total Consumer Credit	-3.32	-2.91	8.1	-4.61	-4.43	17.0	-6.40	-6.40	29.9	-10.09	-10.15	51.8
(25)	Employment-Population Ratio	0.79	0.66	0.5	-0.15	-0.13	0.0	-1.54	-1.30	1.7	-4.39	-3.22	9.7
(26)	CPI (Growth)	0.92	0.79	0.6	2.03	1.84	3.4	3.07	2.71	7.1	4.36	3.26	10.0

Table 3.1 reports results from monthly predictive regressions of US sovereign credit default swaps with maturities from 1 to 5 years. For comparison, all reported predictive coefficients are scaled to be interpreted as the change in basis points (bps) resulting from a one-standard-deviation increase in each predictor variable. The R^2 values are percentages.

US Credit Default Swap Spread Predictability: Univariate Predictor Performance (2/2)

	US - 07Y Maturity			US - 10Y Maturity			US - 15Y Maturity			US - 30Y Maturity			US - Average		
	Coeff.	t-stat.	R^2	Coeff.	t-stat.	R^2	Coeff.	t-stat.	R^2	Coeff.	t-stat.	R^2	Coeff.	t-stat.	R^2
(1)	8.37	8.27	41.6	8.83	8.22	41.3	10.25	10.40	53.0	9.88	8.44	42.6	7.56	7.03	33.5
(2)	0.98	1.57	2.5	0.99	1.49	2.3	1.32	1.96	3.9	1.42	1.97	3.9	1.02	1.58	2.7
(3)	-2.39	-1.84	3.4	-2.13	-1.54	2.4	-0.62	-0.43	0.2	-2.62	-1.72	3.0	-2.09	-1.61	2.8
(4)	3.92	3.10	9.1	2.40	1.74	3.0	3.52	2.53	6.2	3.29	2.18	4.7	3.90	3.16	9.9
(5)	-0.80	-0.53	0.3	-1.27	-0.80	0.7	-0.75	-0.46	0.2	-0.86	-0.49	0.2	-0.27	-0.13	0.3
(6)	3.76	2.96	8.4	2.15	1.55	2.4	3.54	2.54	6.3	3.14	2.08	4.3	4.05	3.34	11.1
(7)	-5.25	-4.31	16.2	-4.72	-3.57	11.7	-6.05	-4.64	18.3	-6.11	-4.30	16.2	-5.54	-4.58	18.0
(8)	-1.41	-1.07	1.2	-1.57	-1.13	1.3	0.35	0.24	0.1	-1.84	-1.20	1.5	-1.10	-0.83	0.9
(9)	4.01	3.17	9.5	2.57	1.86	3.5	3.54	2.53	6.3	3.40	2.25	5.0	3.93	3.16	9.9
(10)	0.01	0.01	0.0	0.26	0.18	0.0	-0.46	-0.32	0.1	-0.42	-0.27	0.1	0.10	0.33	0.1
(11)	-1.45	-1.03	1.1	-1.74	-1.17	1.4	-0.76	-0.50	0.3	-1.45	-0.88	0.8	-0.94	-0.63	0.5
(12)	5.38	4.46	17.1	5.06	3.88	13.5	4.83	3.57	11.7	5.94	4.18	15.4	5.71	4.86	19.9
(13)	4.83	3.92	13.8	4.13	3.09	9.1	5.05	3.76	12.9	5.34	3.70	12.5	5.47	4.94	20.4
(14)	3.99	3.17	9.5	5.92	4.69	18.6	5.30	3.98	14.2	5.94	4.18	15.4	2.85	2.10	8.2
(15)	6.00	5.09	21.2	5.12	3.92	13.8	5.79	4.40	16.8	6.42	4.56	17.8	6.09	5.24	22.3
(16)	- 8.80	-8.79	44.6	- 9.72	-9.53	48.6	-10.05	-9.69	49.4	-10.53	-9.21	46.9	-7.45	-6.92	33.2
(17)	-8.35	-8.03	40.2	-8.02	-6.89	33.1	-8.67	-7.48	36.8	-8.85	-6.91	33.2	-7.18	-6.39	29.9
(18)	-6.14	-5.25	22.3	-4.71	-3.57	11.7	-6.80	-5.38	23.2	-6.10	-4.30	16.2	-5.90	-5.03	21.0
(19)	-8.27	-7.81	38.8	-7.79	-6.53	30.8	-8.53	-7.20	35.1	-8.73	-6.69	31.8	-7.32	-6.46	30.4
(20)	2.96	2.28	5.2	4.12	3.06	8.9	5.33	3.98	14.2	4.50	3.04	8.8	2.33	1.66	5.2
(21)	-3.35	-2.60	6.6	-3.16	-2.30	5.2	-3.55	-2.54	6.3	-4.13	-1.37	7.4	-3.16	-2.25	5.8
(22)	9.18	9.79	50.0	9.32	9.05	46.0	9.37	8.73	44.2	10.16	8.87	45.0	8.22	8.03	39.6
(23)	-2.36	-1.82	3.3	-1.35	-0.97	1.0	-2.97	-2.12	4.5	-1.67	-1.09	1.2	-1.77	-1.30	2.7
(24)	-8.51	-8.25	41.5	-7.79	-6.58	31.1	-8.80	-7.62	37.7	-8.87	-6.90	33.1	-7.53	-6.84	32.5
(25)	-4.60	-3.63	12.1	-5.12	-3.84	13.3	-5.05	-3.67	12.3	-4.97	-3.32	10.3	-3.33	-2.44	7.7
(26)	5.28	4.35	16.5	6.23	4.97	20.5	4.48	3.28	10.1	5.94	4.16	15.3	4.25	3.30	11.0

Table 3.2 reports results from monthly predictive regressions of US sovereign credit default swaps with maturities from 7 to 30 years, as well as the average across all ten maturities. For comparison, all reported predictive coefficients are scaled to be interpreted as the change in basis points (bps) resulting from a one-standard-deviation increase in each predictor variable. The R^2 values are percentages.

US Credit Default Swap Spread Predictability: Bivariate Predictor Performance

	US - 05Y Maturity					US - 10Y Maturity					US - Average				
	Tail Coeff.	Tail t-stat.	Coeff.	t-stat.	R^2	Tail Coeff.	Tail t-stat.	Coeff.	t-stat.	R^2	Tail Coeff.	Tail t-stat.	Coeff.	t-stat.	R^2
(2)	8.07	7.01	0.36	0.65	35.8	8.74	7.97	0.24	0.47	41.5	7.42	6.78	0.38	0.72	33.9
(3)	8.16	7.32	-2.17	-1.95	38.0	8.78	8.28	-1.94	-1.83	43.3	7.52	7.07	-1.92	-1.80	35.9
(4)	7.22	6.49	3.76	3.38	42.4	8.80	7.87	0.08	0.08	41.3	7.02	6.39	2.06	1.93	37.3
(5)	8.21	7.21	0.03	0.03	35.5	8.79	8.14	-0.65	-0.53	41.5	7.57	7.00	0.26	0.21	33.8
(6)	7.27	6.75	4.26	3.95	44.6	8.78	7.94	0.23	0.21	41.3	7.01	6.52	2.52	2.43	39.1
(7)	6.91	6.50	-4.94	-4.63	47.4	8.15	7.50	-2.57	-2.36	44.6	6.56	6.24	-3.81	-3.64	42.0
(8)	8.27	7.33	-1.31	-1.16	36.4	8.93	8.42	-2.03	-1.91	43.5	7.64	7.13	-1.49	-1.38	35.0
(9)	7.23	6.52	3.82	3.43	42.6	8.75	7.84	0.30	0.27	41.4	7.02	6.40	2.12	1.97	37.3
(10)	8.21	7.24	0.45	0.39	35.6	8.83	8.19	0.35	0.33	41.4	7.56	7.00	0.20	0.18	33.6
(11)	8.19	7.25	-1.11	0.92	36.0	8.81	8.25	-1.62	-1.41	42.5	7.55	7.01	-0.84	-0.72	34.0
(12)	7.02	6.76	5.20	5.00	48.9	8.09	7.64	3.20	3.02	46.5	6.60	6.42	4.19	4.16	44.6
(13)	7.10	6.54	4.33	3.98	44.7	8.31	7.58	2.02	1.84	43.3	6.60	6.25	3.79	3.71	43.4
(14)	8.61	7.05	-1.07	-0.88	36.0	7.69	6.85	3.03	2.70	45.5	7.56	6.62	0.00	-0.02	36.6
(15)	6.40	6.28	6.03	5.90	52.8	8.02	7.31	2.69	2.44	44.8	6.30	6.00	4.19	4.07	44.1
(16)	4.44	3.04	-5.56	-3.76	43.8	4.26	3.23	-6.75	-5.04	53.7	4.77	3.44	-4.12	-2.97	40.3
(17)	4.30	3.63	-6.97	-5.79	52.3	6.41	5.18	-4.32	-3.44	47.8	5.23	4.25	-4.16	-3.38	40.9
(18)	5.69	5.23	-6.02	-5.52	51.1	8.32	7.04	-1.22	-1.03	42.0	6.18	5.43	-3.31	-2.99	40.1
(19)	4.12	3.39	-7.07	-5.69	51.9	6.65	5.18	-3.77	-2.87	46.0	5.12	4.05	-4.22	-3.31	40.6
(20)	8.76	7.18	-1.45	-1.19	36.4	8.51	7.30	0.84	0.72	41.6	7.82	6.75	-0.68	-0.60	35.4
(21)	7.74	6.91	-2.67	-2.37	39.1	8.54	7.89	-1.66	-1.52	42.7	7.23	6.70	-1.89	-1.73	35.6
(22)	4.49	4.44	7.81	7.71	60.3	5.68	5.56	6.62	6.47	59.3	4.72	4.56	5.97	5.77	50.0
(23)	7.70	6.61	-1.83	-1.57	37.1	9.14	8.19	1.15	1.03	42.0	7.65	6.83	0.32	0.28	34.3
(24)	3.64	3.15	-7.91	-6.74	56.4	6.61	5.18	-3.84	-2.96	46.3	4.91	3.94	-4.60	-3.72	42.2
(25)	7.70	6.28	-1.31	-1.05	36.2	8.11	7.00	-1.88	-1.59	42.8	7.42	6.37	-0.36	-0.29	34.7
(26)	5.17	5.14	-7.34	-7.21	58.3	7.66	6.64	-2.82	-2.42	44.7	5.76	5.24	-4.35	-4.02)	43.6

Table 3.3 reports results from monthly predictive regressions of US sovereign credit default swaps with maturities from 5 to 10 years, as well as the average values across all ten maturities. The table repeats the analysis of Tables 3.1 and 3.2 but instead reports bivariate regressions using the credit tail risk alongside each alternative economic variables. For each maturity and the average values, the first two columns are the coefficient estimate and t-statistic for the credit tail risk measure, whereas the third and fourth columns are the coefficient and t-statistic for the alternative predictor.

We also investigate the out-of-sample predictability of the tail risk measure on US sovereign credit default swap prices. We perform a univariate predictive regression of the credit default swap spread on credit tail risk using data only through month t (beginning with five years of monthly data from 2009 to 2013, $t = 60$, to allow for a sufficiently large initial sample size). This parameter (regression coefficient) is used to forecast the credit default swap spread at $t+1$. Each month, the estimation window is extended by one month obtaining an updated predictive coefficient, and a new out-of-sample forecast of the succeeding month's credit default swap spread is performed. This estimation approach is repeated until the full sample has been processed. This approach reflects the information set investors would use in real-time because coefficients only rely on data through t. The forecasting errors from this procedure are employed to calculate the out-of-sample R^2 such that

$$R^2 = 1 - \frac{\sum_t (s_{m,t+1} - \hat{s}_{m,t+1|t})^2}{\sum_t (s_{m,t+1} - \bar{s}_{m,t})^2} \tag{3.5}$$

where $\hat{s}_{m,t+1|t}$ is the out-of-sample forecast of the credit default swap spread at $t+1$ based on data through t, and $\bar{s}_{m,t}$ is the historical average credit default swap spread through t. A negative R^2 means that the predictor performance is worse compared to setting forecasts equal to the historical mean value $\bar{s}_{m,t}$. The same recursive out-of-sample method is also applied to each of the alternative predictors. The results from the out-of-sample analysis are reported in Table 3.4. Credit tail risk forecasts successfully predict the out-of-sample credit default swap spreads. The credit tail risk measure has the highest average R^2 compared to other predictors. The R^2 is 40%, 59%, 64%, and 65% for credit default swap spreads with 1-, 5-, 10-, and 30-years to maturity. It is also notable that the total outstanding consumer credit (owned and securitized), commercial and industrial loans, commercial bank credit perform exceptionally well for the five-year tenor. We assess the predictive power based on the average of the monthly consecutive p-values and R^2 from January 2014 until March 2017. The credit tail risk measure demonstrates statistically significant out-of-sample performance for 5-, 10-, and 30-years to maturity (at the 1% significance level or better). Other measures of uncertainty, such as the volatility of the S&P 500, VIX and Smoothed US Recession Probabilities also demonstrates statistically significant out-of-sample performance for contracts with a maturity of 1-year (however with a lower R^2-values). In summary, the predictive regressions suggest that our credit tail risk exponent is positively and significantly related to credit default swaps spreads.

US Credit Default Swap Predictability: out-of-sample R^2 (%)

		01Y	05Y	10Y	30Y	Avg.
(1)	Credit Tail Risk	0.55	0.70	0.78	0.75	**69.4**
(2)	Inflation (Goyal and Welch)	0.11	0.04	0.04	0.04	5.6
(3)	Default Return Spread	0.03	0.02	0.02	0.02	2.4
(4)	LT Government Yield	0.40	0.28	-0.11	-0.02	13.8
(5)	Risk Free Rate	-0.35	-0.14	-0.02	-0.04	-13.8
(6)	Corporate Bond Yield on AAA	0.42	0.33	-0.03	0.06	19.5
(7)	Default Yield Spread	0.04	0.27	0.22	0.24	19.3
(8)	Long Term Corporate Bond Returns	-0.02	-0.01	-0.02	-0.02	-1.7
(9)	Term Spread	0.35	0.27	-0.11	-0.03	12.0
(10)	Long Term Rate of Return	-0.02	-0.01	0.00	-0.01	-1.0
(11)	Treasury Bills	-0.17	-0.11	-0.03	-0.05	-8.7
(12)	Stock Volatility S&P500	0.15	0.26	0.20	0.21	20.6
(13)	US Recession Probabilities	0.26	0.22	0.17	0.18	21.0
(14)	Gold Future	-0.19	0.03	0.17	0.15	4.1
(15)	VIX	0.32	0.45	0.31	0.33	35.1
(16)	Real Home Price Index	-0.23	0.67	0.77	0.73	48.2
(17)	Commercial and Industrial Loans	-0.07	0.77	0.63	0.61	48.4
(18)	Industrial Production Index	0.44	0.65	0.23	0.34	41.4
(19)	Commerical Bank Credit	0.35	0.73	0.62	0.62	58.2
(20)	Inflation Expectation (UoM)	-0.27	0.03	0.19	0.19	3.7
(21)	Federal Surplus or Deficit	0.09	0.10	0.05	0.04	6.7
(22)	Economic Policy Uncertainty Index	0.28	0.63	0.62	0.64	54.0
(23)	Consumer Loans, Credit Cards	-0.16	0.19	-0.18	-0.12	-6.8
(24)	Total Consumer Credit	0.39	0.85	0.61	0.64	62.2
(25)	Employment-Population Ratio	-0.89	0.03	0.19	0.11	-13.8
(26)	CPI (Growth Rate)	-0.05	0.09	0.18	0.15	9.1

Table 3.4 shows the out-of-sample forecasting R^2 in percent from predictive regressions of the US credit default swap spreads for one-, five-, ten-, and thirty-years to maturity. We perform a univariate predictive regression of the credit default swap spreads on estimated credit tail risk and alternative predictors using data only through month t (beginning with five years of data from 2009 to 2013, t= 60, to allow for a sufficiently large initial sample size). Predictive coefficient estimates only use data through t, which are then used to forecast US credit default swap spreads at $t+1$. The out-of-sample R^2 is calculated as $1-\sum_t(s_{m,t+1}-\hat{s}_{m,t+1|t})^2/\sum_t(s_{m,t+1}-\bar{s}_{m,t})^2$, where $\hat{s}_{m,t+1|t}$ is the out-of-sample forecast of the credit default swap spreads at $t+1$ based on data through t, and $\bar{s}_{m,t}$ is the historical average credit default swap spread through t. A negative R^2 means that the predictor performance is worse compared to setting forecasts equal to the historial mean value $\bar{s}_{m,t}$.

3.5 Conclusion

A measure of tail risk is essential for understanding the behaviour of asset prices. Chapter 3 presents the second tail risk measure, which addresses two main challenges of the dynamic tail risk model in Chapter 1. Firstly, the Kolmogorov-Smirnov method causes abrupt changes of the tail threshold and tail length. This results in changes in tail exponent which could not be attributed to changes in extreme returns. To overcome this problem, the new dynamic power-law model incorporates a smoothing technique for the time-varying tail threshold that eliminates tail exponent fluctuations due to abrupt changes in the tail length. Secondly, estimating time-varying tail risk in univariate time series is challenging because of the infrequent nature of extreme events. This problem is particularly acute in univariate time-series of relatively new assets with short historical data or the absense of high-frequency data. Measuring tail risk in univariate data often relies on long time series, which may include tail events from the distant past (previous crises) with no causality to the current market situation. The new dynamic tail risk model overcomes this difficulty by aggregation of returns across the term structure of default. Aggregated return data significantly reduce the lookback window from several to one year and increase the adaptability of the tail risk measure.

We implement the dynamic tail risk estimator using daily returns from ten maturies of US sovereign credit default swaps with 1-year to 30-year tenor. The evidence suggests that credit tail risk has significant predictive power for corporate and sovereign credit default swaps returns. We find that a one-standard-deviation increase in tail risk forecasts an average increase in US sovereign credit default swap spreads of 7.6 bps, which is highly significant. We explore the robustness of the forecasting power of the credit tail risk measure to controlling for 25 alternative predictors. We conclude that increases in tail risk significantly predict increases in credit default swap prices.

Chapter 4

Term Structure of Tail Risk in Global Sovereign Credit Default Swaps

4.1 Introduction

The objective of this chapter is to explore the tail risk among different maturities of sovereign credit default swaps (CDS). We present a methodology for estimating tail risk for a specific maturity using extreme CDS returns from a wide range of countries (cross-sectional or cross-country data). Similar to the term structure of default probabilities (Delianedis and Geske, 2003), we develop the term structure of credit tail risk. Essentially, the term structure of credit tail risk is the relationship between tail risk and different maturities of insurance contracts.

We define three primary shapes of the term structure of tail risk. A downward-sloping term structure implies that the short term tail risk is lower than the long term tail risk. This shape is considered normal because the likelihood of rare events happening over a long time window is higher, which should be considered in the tail risk exponent. An upward-sloping term structure means that the short term tail risk is higher than the long term tail risk. Such a shape can be expected in periods of financial distress and when spreads invert. Chapter 1 exhibits an upward-sloping term structure for European credit defaults during the European sovereign debt crisis. Thereby, shorter-dated CDS maturities demonstrate a higher probability of extreme price fluctuations. Finally, a flat term structure means that tail risk exponents are approximately equal for all maturities. A flat term structure implies that extreme events impact the tail dynamics similarly and are independent of CDS maturity.

We estimate the term structure of credit tail risk exponent using daily cross-country data of sovereign credit default swaps from January 2009 to March 2017. It is challenging to measure fluctuations in the term structure of credit tail risk using univariate time series because a long time series is required to capture enough tail events. To reduce the dependency on data from the distant past, we introduce a panel-based estimator that captures common variation in credit tail risk of sovereign credit default swaps for each maturity. Our estimate for the term structure of tail risk relies on daily returns from 46 countries in seven regions. By pooling returns for a specific CDS maturity from different countries (cross-country returns), we can reduce the lookback window to monthly and non-overlapping time intervals. We estimate the credit tail risk, month-by-month, for each maturity applying the power-law estimator to a set of daily returns of sovereign credit default swaps.

We find evidence that the tail exponents are approximately equal for all CDS maturities. The ANOVA results show no significant difference in tail risk at a level of 0.001 among those maturities. This suggests that the cross-country portfolios with different maturities have no significant differences in the level of credit tail risk.

4.2 Data

Sovereign credit default swap data must meet certain criteria to be considered in the cross section of the sample. Markit classifies 92 countries into ten geographical regions, including Africa (14), Asia (11), Caribbean (4), Europe (18), East Europe (17), Latin America (14), Middle East (10), North America (2), Oceania (1), and Pacific (1). The number of countries in each region is stated in brackets. To be consistent with the US sovereign credit default swaps in Chapter 3, we only include price information for senior unsecured debt traded under the Cum Restructuring (XR) clause. For the analysis of the term structure of credit tail risk, a country is required to have price information on at least 75% of all trading days across four maturities between January 2009 and March 2017. This term structure analysis is performed on credit default swaps with 1-, 5-, 10- and 30-year tenor. We extend the US data by 45 countries from seven regions, which meet these criteria. The European data set includes 14 countries, namely, Austria, Finland, Ireland, Italy, Portugal, Spain, Sweden, Belgium, Germany, France, Norway, United Kingdom, Denmark, and the Netherlands. The second-largest data set includes nine counties from Latin America, including Chile, Colombia, El Salvador, Guatemala, Panama, Peru, Mexico, Uruguay, and Venezuela. The eight East European countries include Kazakhstan, Poland, Slovakia, Latvia, Russia, Czech Republic, Bulgaria, and Croatia. Qatar, Bahrain, Egypt, Emirate of Abu Dhabi, Israel, Lebanon, and Turkey belong to the qualified data set of Middle Eastern countries. In Asia, China, Indonesia, Japan and Thailand fulfil the minimum requirements. The data set of African countries is small, only Morocco and South Africa are meeting the criteria. From the Caribbean region, only the Dominican Republic is included in our analysis. Sovereign credit default swaps without adequate price information are disregarded. We exclude 46 countries (50%) due to a lack of sufficient price information.

4.3 Methodology

We introduce this model extension to cross-country credit default swap returns to estimate the term structure of credit tail risk. Chapter 1 discusses the challenges of modelling time-varying tail risk for a single maturity using univariate time series. As tail events rarely occur by definition, one requires a long history of returns to estimate the tail risk from univariate time series. The main disadvantage is that those time series might hold tail events from the distant past unrelated to the present situation. Furthermore, many sovereign credit default swaps have a relatively short history. To overcome these issues, we exploit information about the credit tail risk for a specific maturity using cross-country returns.

We define the tail as the set of returns ($x_{m,t}$) exceeding some high threshold ($a_{m,t}$) for a specific CDS maturity m at time t. Similar to Equation 1.1 in Chapter 1, we assume that the cross section of tail returns of sovereign credit default swaps for a specific CDS maturity m behave according to

$$P(x_{m,t} > a_{m,t}) = \left(\frac{x_{m,t}}{a_{m,t}}\right)^{-\alpha_{m,t}}. \tag{4.1}$$

Equation 4.1 states that tail returns across a specific CDS maturitity m obey the dynamic power-law with tail exponent $\alpha_{m,t}$. The threshold value $a_{m,t}$ defines where the centre of the distribution finishes and the tail begins for maturity m. We define threshold $a_{m,t}$ as the 95th percentile of the cross-country return set for a specific CDS maturity. The selection of the threshold can have a significant impact on tail exponent $\alpha_{m,t}$. An inappropriately low threshold might contaminate tail exponent estimates using data from the centre of the return distribution, whose behaviour might vary significantly from tail returns. On the other hand, a high threshold can result in noisy estimates resulting from too few data points. While we use the Kolmogorov-Smirnov method to define the minimum magnitude of the smallest extreme event in Chapter 1 to 3, this method is computationally too expensive to conduct for larger cross-sectional time-series. Gabaix *et al.* (2006) advocate a simple rule fixing the a-exceedence probability at 5% for power-law estimation. We follow these authors by applying a similar simple rule in this chapter.

Next, we estimate the credit tail risk exponent, month-by-month, for each maturity applying the power-law estimator to the set of daily tail returns for all sovereign credit default swaps in month t. Applied to the aggregated cross-country tail returns each month, the credit tail risk exponent estimate takes the form

$$\hat{\alpha}_{m,t} = 1 + k_{m,t} \left[\sum_{i=1}^{k_{m,t}} \ln \frac{x_{i,m,t}}{a_{m,t}} \right]^{-1}, \tag{4.2}$$

where $x_{m,i,t}$ is the ith daily return that falls above the extreme value threshold $a_{m,t}$ for maturity m during month t. The number of tail observations which exceeds the threshold value for maturity m is denoted as $k_{m,t}$. Observations that do not exceed the threshold value of $a_{m,t}$ are discarded.

We estimate the term structure of credit tail risk based on cross-country data. This significantly reduced the lookback window. Each cross-maturity data set has 1,012 returns per month. Assuming that we define $a_{m,t}$ as the 95th percentile of the cross-country data each month, we have a sufficiently large data set to estimate the tail statistics. We estimate the tail risk exponent for each maturity based on approximately 50 daily tail returns within a month.

4.4 Results

We find that the term structure of credit tail risk is flat between January 2009 and March 2017. The credit tail risk is 3.45 (0.24) for 1-year, 3.49 (0.22) for 5-years, 3.46 (0.23) for 10-years, and 3.37 (0.18) for 30-years. The variance is stated in brackets. We use the Levene test to ascertain whether the variances in credit tail risk are equal across different maturities (null hypothesis). The high p-value of 0.4218 indicates that the null hypothesis is not rejected at the predetermined alpha level of 5%. We confirm this result using the Brown–Forsythe test (p=0.5261), which uses the median instead of the mean. Furthermore, we perform the Levene's test on pooled ranked scores, which is also known as the nonparametric Levene test. In the presence of skewness, the nonparametric Levene test provides robust results (p=0.7097). The skewness is slightly positive for credit tail risk exponents ranging from

0.2973 to 0.4825.

The ANOVA results indicate no significant difference between the tail risk exponents among maturities at different acceptance levels of 0.001, 0.01 and 0.05. This suggests that all cross-country portfolios show statistically the same level of credit tail risk independently of the maturity. It also means that credit default swap seller bears the same credit tail risk for short- and long-term contacts for a large portfolio of diversified idiosyncratic country risks. While the implied probability of default increases with maturity, the impact of tail risk is perceived equally among maturities.

Chapter 5

Tail Risk Dynamics in US Corporate Credit Default Swaps

5.1 Introduction

Returns of securities are known to be correlated because they relate to market movements. This chapter will examine the assumption that tail risks of different industries are correlated. A major motivation for this study is the assumption that tail risks of different assets possess similar dynamics. If tail distributions possess similar dynamics, then the cross section of crash events for individual companies can be used to identify the common process of tail risk across firms at each point in time (Kelly and Jiang, 2014).

Our primary goal is to investigate the correlation effects of the time-varying tail risk. Ideally, one would directly build an estimator of aggregate tail risk dynamics from the underlying time series of individual firm returns, in analogy to dynamic volatility estimated from a GARCH model. However, as previously discussed in Chapter 2 and 3, dynamic tail risk estimates are difficult to model based on short and low-frequency (univariate) time series model due to the infrequent nature of extreme returns.

To overcome this problem, we introduce a panel estimation approach that captures common fluctuations in the tail risks of individual firms within the same sector. We assume that tail distributions of companies within the same industry possess similar behaviour. Hence, extreme events for individual firms can be utilised to measure the common element of tail risk. A sufficiently large number of firms within a sector is essential to provide accurate tail exponent estimates because only a small fraction of data is informative about the tail distribution.

Our approach to estimating sector-specific tail risk relies on commonality in the tail risk of individual firms, which in turn utilises information about the tail risk in the cross section of returns. We estimate the dynamic tail risk exponent using daily returns of US corporate credit default swaps from January 2009 to March 2017. At time t, the upper (positive) tail distribution is defined as the set of firm returns falling above some extreme positive threshold $a_{s,m,t}$ for some sector s and credit default swap maturity m. We assume that the upper tail returns of firm i in sector s behaves according to a power-law, such that

$$P(x_{i,m,t} > a_{s,m,t}) = \left(\frac{x_{i,m,t}}{a_{s,m,t}}\right)^{-\alpha_{s,m,t}}, \tag{5.1}$$

where $0 < a_{s,m,t} < x_{i,m,t}$. The key parameters of Equation 5.1, $a_{s,m,t}$ and $\alpha_{s,m,t}$ determine the tail risk dynamics of each sector s in period t. In contrast to previous approaches in Chapter 1 to 3, Equation 5.1 is a model of sector tail dynamics, where a large number of firms is available. Because in a sufficiently large cross section, enough firms will experience individual tail events, we can use a monthly lookback window and non-overlapping sets of returns.

Because individual credit default swap returns contain information about the likelihood of market-wide extremes, the cross section of all firms can be used to approximate the credit tail risk in the US economy. We also utilise Equation 5.1 to identify the common component of tail risk by merging all individual firms (independently of the sector) into one set. Then, our panel estimation approach uses the cross section of extreme events in all US corporate credit default swaps to estimate the tail exponent. We refer to $\alpha_{m,t}$ for all firms as the credit tail exponent for the US economy or the "US credit tail risk". The US credit tail risk estimate derived from all corporate credit default swaps is also used for the portfolio analysis in the final chapter (Chapter 6).

We find that the average correlation between the tail exponent of the US economy and the ten sectors is 71% and 69% for 5- and 10-year maturities of corporate credit default swaps. For both maturities, the consumer services and financial industry's tail risk exponent has the highest correlation with the aggregated tail risk measure, reaching levels above 80%. We find that the tail dynamics among sectors have an average correlation of 48% for credit default swaps with a 5-year maturity, where consumer goods and services reach the highest correlation of 75%. The average tail risk sector correlation is slightly lower (44%) for 10-year insurance contracts. Furthermore, the difference between average tail risk exponents for the US economy and most sectors (7 out of 10) is not significant at a 5% level for corporate credit default swaps with a 5-year maturity. Finally, we provide evidence that tail risk dynamics might be governed by a common process as firms with low and high credit default risk bear a similar level of tail risk. There is no significant correlation between the tail risk exponent and the level of credit default risk at a 5% level.[1]

5.2 Data

We estimate the dynamic tail risk exponent for each sector by using daily logarithmic returns of US corporate credit default swaps from January 2009 to March 2017. The data are provided by IHS Markit Ltd. through the LSE Systemic Risk Centre. The total number of US firms is 1,951 in this period. IHS Markit classifies the US companies into ten sectors, including basic material (107), consumer goods (247), consumer services (330), energy (178), financials (375), healthcare (116), industrials (210), technology (110), telecommunication services (87) and utilities (191). The number of companies in each sector is stated in brackets. To be consistent with sovereign credit default swap data, we only consider price information for senior unsecured debt traded under the Cum Restructuring (XR and XR14) clause. Trading on the XR14 definitions begins on September 22, 2014. To be included in the tail risk analysis, corporate credit default swap data are subject to the same range of logical and relative tests as sovereign credit default swaps in Chapter 2 to 4 (curve buildability test, backwardation test, stable data and liquidity test, and the outlier test).

In Section 5.4.2 and Chapter 6, we empose an additional selection criteria. To be considered in the cross section of the sample, a company is required to have price information on at least 75% of all trading days (at least 1,605 daily price data) between January 2009 and March 2017. Through this additional criteria, the total number is reduced from 1,951 to 675 US companies.

[1] A static tail risk analysis finds that the average tail exponent is 2.77 for 675 US firms from January 2009 to March 2017 (univariate time series, non-aggregated data, no rolling window). Over the same period, the tail exponent is 2.76 for aggregated returns of these 675 US firms. Furthermore, we find that the correlation between the tail exponent and the implied probability of default is 4.88%. There is no significant correlation between the tail risk exponent and the level of credit default risk at a 5% level.

5.3 Methodology

The estimation approach for the tail exponent is similar to the previous chapter, but with a twist. Instead of estimation across countries, we estimate the tail exponent across corporate credit default swaps based on common factors, denoted s, for example industry classification (Section 5.4.1) or risk group (Section 5.4.2).

We estimate the tail exponent month-by-month for each s by applying the power-law estimator of Hill (1975) to the set of daily CDS returns for all firms in s in month t. Applied to the aggregated cross section each month, the tail exponent estimate takes the form

$$\hat{\alpha}_{s,m,t} = 1 + k_{s,m,t} \left[\sum_{i=1}^{k_{s,m,t}} \ln \frac{x_{i,s,m,t}}{a_{s,m,t}} \right]^{-1}, \tag{5.2}$$

where $x_{i,s,m,t}$ is the i^{th} daily return in sector s that falls above the threshold $a_{s,m,t}$ of a particular sector during month t. The variable $k_{s,m,t}$ is the total number of such exceedences and determines the tail length within month t, which can be different for each s and t. The Hill (1975) estimator for the tail exponent uses only those extreme values that exceed the tail threshold $a_{s,m,t}$ and discards returns below the cut-off point. The tail exponent is estimated on a non-overlapping set of returns within each month t.

In different periods, different firms will experience extreme events, which affect the month-by-month estimation of the tail exponent. Although the identities of extreme returns are unknown, the number of tail events $k_{s,m,t}$ is known because the tail is defined by a fixed percentile of the pooled returns (the largest 5% of price increases that month). Because the tail percentile is a fixed quantile of the cross section of returns, the tail threshold $a_{s,m,t}$ varies over time. The tail threshold $a_{s,m,t}$ expands and contracts with variations in volatility. Hence, common time-variation in volatility is largely factored into the tail estimates' construction, helping to offset the effect of volatility dynamics on the tail exponent.

We derive two simple extensions from Equation 5.2. Firstly, because individual CDS returns contain information about the likelihood of market-wide extremes, the cross section of all firms can be used to approximate the credit tail risk in the US economy. This extension is also used in the following Chapter 6. Secondly, to analyse the relationship between credit default risk and tail risk, we pool returns across firms in five risk categories (from low to high) based on the average credit default probability from January 2009 to March 2017. (see, Section 5.4.2).

5.4 Results

5.4.1 Credit Tail Risk Dynamics within Sectors

The power-law aggregation method concerning firm tails to the sector and sovereign tail builds is based on our specification that all credit default swaps share a common factor as in Equation 5.2. To provide empirical evidence for this specification, we group the sample of US corporate credit default swaps into non-overlapping subsets according to the Markit industry code classification. A precise estimate of the credit tail exponent requires a sufficiently large number of tail observations. Within each industry, we calculate the cross-sectional credit tail risk by pooling daily observations within a month. The smallest industry (87 firms in the telecommunication services industry) has a moderately large subset of approximately 1,900 returns per month. The largest industry (375 financial companies) has a subset of approximately 8,250 observations per month. To estimate the credit tail risk of the US economy overall, we group all corporate credit default swap returns across these ten sectors. We then show that credit tail risk estimates are highly correlated across industries and with the credit tail risk of the US economy.

We find a high correlation among different industries, whereas the financial sector has one of the highest correlations with the overall credit tail risk of the US economy. Panel A of Table 5.1 shows that the sector credit tail risks are highly correlated with the overall credit tail of the US economy, ranging between 41% and 89%. The average correlation between the tail exponent of the US economy and the ten sectors is 71% for corporate credit default swaps with a 5-year maturity. The average correlation of the credit tail risk among sectors is 48%, whereas the correlation between consumer goods and consumer service is the highest with 75%. Panel B of Table 5.1 shows that these findings hold for both, five and ten-year credit default swaps tenors. The credit tail risk of the financial sector has the highest correlation of 88% with the overall credit tail risk, as the financial industry generally has higher exposure to risks of other sectors of the US economy, e.g. due to spill-over effects. All correlation estimates in Table 5.1 are highly statistically significant at a 1% level.

Furthermore, we provide evidence that credit tail risk is similar between the US economy and different industries. Panel C of Table 5.1 reports the average tail risk exponents for each sector for 5- and 10-year maturities from January 2009 to March 2017. The difference between average tail risk exponents for the US economy and 7 out of 10 sectors is not significant at a 5% level for corporate credit default swaps with a 5-year maturity, namely: Basic Material, Energy, Financials, Healthcare, Industrials, Technology and Telecommunication.

Correlation of dynamic credit tail risk exponents among sectors

		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
US Credit Tail Risk		0.77	0.80	0.89	0.64	0.81	0.69	0.71	0.70	0.41	0.70
Panel A: Correlation among sectors, 5-year tenor											
Basic Material	(1)	1.00									
Consumer Goods	(2)	0.63	1.00								
Consumer Services	(3)	0.68	0.75	1.00							
Energy	(4)	0.52	0.45	0.52	1.00						
Financials	(5)	0.62	0.57	0.65	0.42	1.00					
Healthcare	(6)	0.49	0.48	0.61	0.40	0.46	1.00				
Industrials	(7)	0.54	0.57	0.62	0.44	0.46	0.61	1.00			
Technology	(8)	0.61	0.52	0.65	0.41	0.54	0.40	0.49	1.00		
Telecommunication	(9)	0.12	0.34	0.31	0.35	0.19	0.30	0.31	0.20	1.00	
Utilities	(10)	0.54	0.45	0.53	0.57	0.60	0.46	0.47	0.41	0.18	1.00

		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
US Credit Tail Risk		0.75	0.74	0.81	0.69	0.88	0.61	0.68	0.57	0.43	0.72
Panel B: Correlation among sectors, 10-year tenor											
Basic Material	(1)	1.00									
Consumer Goods	(2)	0.57	1.00								
Consumer Services	(3)	0.56	0.59	1.00							
Energy	(4)	0.52	0.37	0.53	1.00						
Financials	(5)	0.62	0.67	0.69	0.50	1.00					
Healthcare	(6)	0.46	0.34	0.48	0.37	0.57	1.00				
Industrials	(7)	0.52	0.47	0.45	0.47	0.58	0.40	1.00			
Technology	(8)	0.40	0.35	0.42	0.44	0.47	0.36	0.26	1.00		
Telecommunication	(9)	0.28	0.30	0.27	0.37	0.23	0.11	0.29	0.15	1.00	
Utilities	(10)	0.51	0.38	0.45	0.65	0.59	0.47	0.47	0.41	0.36	1.00

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel C: Sector Credit Tail Risk										
CDS Tenor: 5Y	2.80	2.98	2.85	2.80	2.77	2.78	2.73	2.80	2.66	2.50
CDS Tenor: 10Y	2.78	2.93	2.91	2.77	2.76	2.69	2.69	2.74	2.64	2.53

Table 5.1 reports time series correlation between monthly credit tail risk estimated from the cross section of credit default swaps for ten industries with CDS tenor of five years (Panel A) and ten years (Panel B). The correlation between monthly sector credit tail risk and the US credit tail risk is stated in the top row above Panels A and B. The US credit tail risk is estimated from the cross section of all corporate credit default swaps. The US credit tail risk is 2.72 and 2.71 for five- and ten-year tenor. Panel C states the average credit tail risk exponent of each sector from 2009 to 2017.

5.4.2 Credit Tail Risk Dynamics within Risk Categories

We find evidence that credit tail risk fluctuations are governed by a single process independently of the idiosyncratic default risk. To provide empirical evidence for this claim, credit default swaps are sorted into five (non-overlapping) portfolios (risk categories) based on their average implied default probability from January 2009 and March 2017. Firms are sorted from low to high default risk and grouped into five equal-sized portfolios (risk categories). To be considered in the cross section of the sample, a firm is required to have price information on at least 1,605 days (75%) between January 2009 and March 2017. This criteria reduces the number of companies to 675 for credit default swap contracts with a maturity of five-year. Thus, each risk category contains an equal amount of 135 companies. Risk category 1 covers the firms with the lowest average default probability from January 2009 to March 2017. The average credit default swap price in risk category 1 is 44.48 bps. Risk category 5 contains the companies with the highest average default probability over the same period. The average credit default swap price in risk category 5 is 592.40 bps. We find that average credit default risk is significantly different between risk categories 1 and 5. The portfolios (risk categories) are only constructed once and sorted from low to high default risk. There is no (monthly or annual) rebalancing or transition of companies between risk categories.

Next, we analyse whether credit default swaps with low and high cross sectional default risk have a similar level of tail risk. If firms with a high probability of default are more prone to extreme events than firms with a low default probability, we would expect fatter tails and a higher probability of extreme returns. Each month, we estimate the tail exponent for each risk category using the cross section of extreme events each day. We compute 99 tail exponent for each risk category from January 2009 to March 2017. Based on the month-by-month estimates, we calculate the average (equal-weighted) credit tail risk exponent.

We find that the average tail exponent, tail threshold and tail returns are similar for firms with low and high credit default risk. In risk category 1, the average credit tail risk exponent is 2.92 with a tail threshold value of 3.2%. The average daily return of the right tail in risk category 1 is 6.3%. In risk category 5, the average credit tail risk exponent is 2.95 with a tail threshold value of 3.3%. The average extreme return in risk category 5 is 6.5%. The differences in the tail exponent, threshold value and average tail return are not significant at 5% level. Among all five risk categories, the average tail exponent is 2.86, with an average tail threshold of 3.1% and an average extreme return of 6.3%. In summary, portfolios of firms with low and high idiosyncratic default risk do not significantly differ in their tail risk statistics.

Because industry-level tail risks are highly correlated over time, we conduct the same correlation analysis on quintile portfolios with different credit default risks. We find that tail risks are highly correlated over time, ranging between 52% and 79%. The average correlation between all risk portfolios is 66%. The correlation between portfolio 1 and the other portfolios monotonically decreases with increasing default risk. Thus, the low- and high-risk portfolios show the lowest correlation (52%). Next, we analyse the correlation of tail risk between quintile portfolios and the US credit tail risk. Time series correlation of quintile tails with the US credit tail risk range between 76% and 90% (85% on average). The results suggest a high degree of co-movement in tail risk of portfolios with different idiosyncratic default risks.

5.5 Conclusion

A measure of extreme risk is essential to understand the behaviour of insurance prices. If the tail risk changes through time, extreme value techniques based on univariate time series (Chapter 1) or aggregated data from a small sample (Chapter 3) are suboptimal of providing reliable tail measures. However, if the tail risk of all firms are driven by a common process, the cross section of extreme returns can be used to accurately measure prevailing tail risk in the economy. We presented two simple studies, providing empirical support to use crash events of individual firms for modelling common tail risk variation even when firms possess distinct characteristics and different idiosyncratic default risks.

Evidence suggests that tail risk dynamics between sectors and the economy are highly correlated. But also tail risk dynamics of low and high default risk companies are highly correlated with the economy with no significant difference in the correlation. A high correlation between tail risk exponents in different industries and risk categories suggest that individual extreme returns likely contain information about the likelihood of market-wide extremes (aggregated market tail risk). Consequently, the cross section of extreme events for individual firms can be used for modelling the tail risk of the economy.

While the dynamics in tail risk are highly correlated, there is no significant correlation between credit default risk and tail risk. Most industries and the economy bear a similar level of tail risk. Firms with a high probability of default are similarly exposed to extreme events as firms with a lower default probability with no significant difference in the average tail return. We conclude that credit default swaps with a relatively low price provide a good and cheap hedge against tail risk.

These empirical results can be understood from the perspective of aggregation properties of variables with power-law tails: i) if the tail distributions possess similar levels of tail risk, then the entire cross section of firms can be used for modelling tail risk as there is no significant difference between low and high default risk companies; and ii) because the tail risk is time-varying and the evolution of tail risk is highly correlated, the cross section of extreme events can be used to identify the common, time-varying element of tail risk. In summary, dynamic tails estimated from distinct subsets of data (grouped by industry classification and firm-specific credit default risk) display a high degree of co-movement, providing empirical support for the specification in Equation 5.2 to estimate the credit tail risk of the US economy from firm-specific tail risk.

Chapter 6

Tail Risk and Risk Premiums in US Corporate Credit Default Swaps

6.1 Introduction

In this chapter, we test the hypothesis that tail risk helps to understand differences in cross-sectional expected returns of US corporate credit default swaps. If a firm is sensitive to tail risk in credit default swap markets, a protection seller require a higher future insurance premium to enter an insurance contract. Otherwise, it is unattractive for a protection seller to enter a credit default swap because the insurance premium remains unchanged for the duration of the contract, and there are no (future) upside gains precisely when the risk of rare events increases. If protection sellers are averse to tail risk, corporate credit default swaps with high predictive loadings to extreme events increase sharply in price and thus have higher future expected returns. On the other hand, firms with low or negative credit tail risk loadings have comparatively lower future expected returns. This study shows that the cross section of expected CDS returns reflects a premium for bearing tail risk.

Our main contribution is an empirical analysis of the tail risk on corporate credit default swaps. We propose a time-varying credit tail risk model that is directly estimable from the cross section of corporate credit default swap returns. We find that the cross section of corporate CDS returns reflects a premium for tail risk sensitivity. Cross-sectionally, firms that covary highly with tail risk earn average expected annual returns 8.1% higher than credit default swaps with low tail risk covariation. Sellers of protection demand additional compensation for default insurance contracts with high sensitivities to extreme events. Past credit tail risk sensitivity is highly persistent and provides enough divergences in future credit tail risk sensitivity to justify those large return spreads. We show that the credit tail risk premium is different from the premiums on market risk, idiosyncratic volatility (Ang *et al.*, 2006b), and coskewness (Harvey and Siddique, 2000), and robust to controlling for these alternative risk factors.

6.2 Research Design

We assume that cross section of tail returns for individual firms (credit default swaps) obey the dynamic power law structure in Equation 5.1. Because a sum of idiosyncratic power law shocks inherits the tail behaviour of individual shocks, we use the cross section of firm-level shocks to estimate the aggregated tail risk in credit default swap markets. Thus, we estimate the common component of tail exponent α_t by applying the power-law estimator of Hill (1975) as in Equation 5.2. The tail exponent α_t is computed, month-by-month, using using the set of daily CDS returns for 675 firms in month t.

Next, we estimate tail risk sensitivities of individual credit default swaps with predictive regressions of the form $E_t[x_{i,t+1}] = \mu_i + \beta_i \alpha_t$, where β_i measures sensitivity of firm i on the cross-sectional credit tail risk α_t. Thus, we refer to the β coefficient as the "credit tail risk sensitivity". Corporate credit default swaps with high values of β_i are most sensitive to credit tail risk, and thus have high expected returns. Conversely, corporate credit default swaps with low or negative values of β_i are good hedges, because their insurance prices tend to be lower, and their expected future returns fall.[1]

Each month, we assess the credit tail risk sensitivities for each credit default swap in regressions, using the most recent 60 months of data. Credit default swaps are then sorted into quintile portfolios based on their estimated credit tail risk sensitivities. Each quintile portfolio contains 135 firms. Portfolios are reconstructed each month. We calculate the average one-, three-, six- and twelve-month portfolio returns for equal- and risk-weighted assets in a one-year post-formation window. Table 6.1 reports out-of-sample portfolio returns as there is no overlap between sample used for estimating tail risk sensitivities and post-formation returns used for evaluating portfolio performances.

[1] At the time of portfolio formation, credit default swaps with low tail risk sensitivities tend to be 9.4% cheaper on average compared to credit default swaps with high tail risk sensitivities.

6.3 Results

6.3.1 Tail Risk and the Cross Section of Expected Returns

Panel A of Table 6.1 reports the twelve-month average returns of credit tail risk quintiles. Credit default swaps in the high tail risk quintile earn equal-weighted annual returns 8.1% higher than credit default swaps in the low tail risk portfolio, with a t-statistic of 4.22. Next, we calculate portfolio returns such that the aggregated investment is zero. This is achieved by balancing proceeds from selling credit default swaps in the low quintile portfolio with expenses from buying credit default swaps in the high quintile portfolio. Credit default swaps in the high tail risk portfolio earn annual returns 8.3% higher than credit default swaps in the low tail risk portfolio, with zero net investment. The difference has a robust t-statistic of 4.14.

Next, we optimise each quintile portfolio by underweighting credit default swaps with higher default risk and overweighting lower default risk firms. Firm's default risk is assessed based on the average credit default swap price in the month prior to portfolio formation. Panel A of Table 6.1 reports that portfolios with high tail risk quintile earn optimised annual returns 18.5% higher than low tail risk portfolios, with a t-statistic of 3.40. The high minus low zero net investment portfolio yields similar return statistics. Furthermore, Panel A of Table 6.1 shows a monotonically increasing pattern between credit tail risk sensitivities and future returns. Our results are consistent with studies in equity markets. Kelly and Jiang (2014) find that stocks with high loadings on past equity tail risk earn an equal-weighted annual return 4.0% higher than stocks with low equity tail risk loadings.

Panel B, C and D of Table 6.1 report average out-of-sample returns for post-formation periods of six-, three-, and one-month. Portfolio returns over shorter horizons have similar qualitative behaviour as those over longer investment horizons. The optimised high minus low portfolios earn 21.1% annualised (t=2.84), 26.6% annualised (t=2.88) and 12.0% annualised (t=2.37). Portfolio returns with shorter investment horizons retain the same monotonic pattern that is observed for the twelve-month post-formation window.

Finally, we examine if credit default swaps simply sorted on past credit tail risk sensitivity provide enough divergences in future credit tail risk sensitivity to justify those large return spreads between quintile portfolios. We sort credit default swaps into quintile portfolios at time t based on tail risk sensitivity, and then examine future credit tail risk sensitivities over the following one-, three-, six- and twelve-months. At the beginning of each month t, we re-

Credit tail risk beta-sorted portfolio returns

	Low	2	3	4	High	High-Low	*t-stat.*
Panel A: Twelve-month returns							
Equal-weighted portfolios	0.28%	1.71%	4.09%	5.65%	8.37%	8.08%	*4.22*
Zero net investment	0.45%	2.29%	4.26%	6.37%	8.78%	8.33%	*4.14*
Risk-optimised portfolios	3.21%	6.39%	13.30%	17.20%	21.68%	18.46%	*3.40*
Zero net investment	4.06%	8.45%	13.98%	19.59%	22.67%	18.61%	*3.22*
Panel B: Six-month returns							
Equal-weighted portfolios	-0.06%	0.99%	2.08%	3.50%	4.60%	4.66%	*4.00*
Zero net investment	0.08%	1.49%	2.17%	4.14%	4.76%	4.68%	*3.86*
Risk-optimised portfolios	0.99%	2.87%	6.11%	9.48%	11.53%	10.54%	*2.84*
Zero net investment	1.56%	4.10%	6.50%	11.03%	11.97%	10.41%	*2.72*
Panel C: Three-month returns							
Equal-weighted portfolios	-0.11%	0.53%	0.83%	1.68%	2.31%	2.42%	*3.93*
Zero net investment	-0.04%	0.78%	0.91%	2.04%	2.31%	2.36%	*3.85*
Risk-optimised portfolios	-0.09%	0.97%	1.80%	3.79%	5.31%	5.40%	*2.88*
Zero net investment	0.05%	1.43%	1.98%	4.55%	5.37%	5.31%	*2.82*
Panel D: One-month returns							
Equal-weighted portfolios	-0.04%	0.16%	0.25%	0.49%	0.62%	0.66%	*3.63*
Zero net investment	-0.01%	0.24%	0.27%	0.58%	0.59%	0.60%	*3.54*
Risk-optimised portfolios	-0.41%	-0.23%	-0.03%	0.18%	0.55%	0.97%	*2.37*
Zero net investment	-0.43%	-0.23%	-0.04%	0.25%	0.49%	0.93%	*2.31*

Table 6.1 reports return statistics for quintile portfolios constructed on the basis of corporate credit tail risk. Each month, we reconstitute portfolios based on predictive tail risk sensitivity using monthly data over the previous five years. Portfolios are based on US corporate credit default swaps on senior unsecured debt, traded under the Cum Restructuring clause. Panel A reports average out-of-sample one-year holding period portfolio returns. The equal-weighted portfolios consists of one credit default swap contract for each firm in the corresponding beta category. The optimised portfolios underweight credit default swaps with higher default risk and overweight lower default risk firms in each quintile. Below the equal-weighted and risk-optimised portfolios are the returns for zero net investment portfolios. Panel B, C and D report average out-of-sample returns for six-, three-, and one-month holding periods. For all holding periods, average portfolio returns demonstrate the same monotonic pattern, where returns increase with increasing tail risk sensitivity. The right-most columns report results for high minus low beta portfolios and corresponding t-statistics.

construct the quintile portfolios based on their updated credit tail risk sensitivity. We find a monotonically increasing pattern in future tail risk sensitivity of credit default swaps sorted on realised tail risk sensitivity. That is, credit default swaps with low (high) tail risk sensitivity continue to have low (high) tail risk sensitivity going forward. The differences in future tail risk sensitivities between high and low portfolios, sorted on realised tail risk sensitivity, is almost unchanged to differences in past tail risk sensitivities. Spreads in future tail risk sensitivities are statistically significant at the 1% level. Hence, past tail risk sensitivity seems to predict future tail risk sensitivity, with significant future tail risk spreads between high and low portfolios. We conclude that there is a strong relation between past credit tail risk sensitivity, future returns, and future tail risk sensitivity. Past credit tail risk is rewarded in cross section of future returns.

However, this relationship does not consider other firm characteristics in the cross section of credit default swaps returns, which we discuss in the following section.

6.3.2 Alternative Risk Factors and Expected Returns

In this section, we test the hypothesis that alternative explanatory characteristics explain differences in expected returns across credit default swaps on corporate senior secured debt. We calculate quintile portfolio returns with respect to various sources of risk, such as market risk beta, credit default risk, idiosyncratic volatility (Ang *et al.*, 2006b), and coskewness risk (Harvey and Siddique, 2000). If investors are averse to these sources of risk, firms with high risk loadings should have higher expected returns.

The first characteristic we examine is market risk beta, which measures the sensitivity of an individual credit default swap compared to the market risk of all corporate credit default swaps. We estimate market beta of individual credit default swaps with linear regressions of returns of firm i on market returns, using the most recent 60 months of data. Next, we investigate whether high credit default risk in the month prior to portfolio formation is associated with high future returns. The third characteristic is idiosyncratic volatility. Ang *et al.* (2006b) show that assets with high idiosyncratic volatility have low average returns. Idiosyncratic volatility is approximated as the standard deviation of returns prior to portfolio formation. There are two reasons why we estimate idiosyncratic volatility using returns and not return residuals from the Fama-French three-factor model (Fama and French, 1993). Firstly, there is little evidence that market capitalisation (small caps over big caps) and book-to-market ratio (value stocks over growth stocks) explain returns in credit default swap markets. Secondly, Kelly and Jiang (2014) argue that firm volatility measured as the standard deviation of residuals from the Fama-French three-factor model (Fama and French, 1993) and volatility directly computed from raw returns, hold similar results. Therefore, we do not estimate idiosyncratic volatility from the Fama-French model, because considering these "stock-related" factors in measuring return volatility of credit default swaps may add noise or distort results. Finally, we examine how past coskewness exposure relates to future returns (Kraus and Litzenberger, 1976). Coskewness is estimated from regressions of firm returns on squared market returns, using the most recent 60 months of data. Harvey and Siddique (2000) find that assets with low coskewness tend to have high average returns.

Each month, we estimate the market risk sensitivity (beta), default risk on corporate debt, idiosyncratic volatility, and coskewness for each credit default swap. Then, we sort corporate credit default swaps into equally-weighted quintile portfolios based on their realised factor loadings. Portfolios are reconstituted each month, based on their updated factor loadings. We calculate the average one-, three-, six- and twelve-month returns in a one-year post-formation window. Table 6.2 reports the average out-of-sample returns in each equally-weighted quintile portfolio for different investment horizons. There is no overlap between data used for estimating each factor and post-formation portfolio returns.

Table 6.2 shows a monotonically increasing pattern between future expected returns and realised market risk sensitivity (beta). Panel A reports that credit default swaps with high market risk sensitivity earn an average return of 8.2% per annum, whereas the low market risk portfolio only earns 0.1% per annum. Compared to other firm characteristics, quintile portfolios with high market tail risk loadings have the highest average post-formation return (except for credit tail risk). The spread in average twelve-month returns between high and low market risk portfolios is 8.1% (t=4.58) per annum, which is the highest spread among all five risk factors. Panel B, C and D of Table 6.2 report that monotonic return patterns and spreads between high and low market risk portfolios are robust for shorter time windows. Our results are consistent with the earliest research in equity markets, for example, Jensen *et al.* (1972), who find higher returns for holding stocks with high market risk.

Table 6.2 exhibits that high exposure to idiosyncratic volatility risk tends to be positively related to future expected returns. While some economic theories suggest that investors demand higher returns for holding assets with high volatility (for example, Merton (1987), Barberis and Huang (2001), Malkiel and Xu (2002) and Ewens *et al.* (2013)), other findings are directly opposite to these theories. Ang *et al.* (2006b) show that stocks with high idiosyncratic volatility risk have low average returns. We find that the volatility effect of

Ang *et al.* (2006b) works in the opposite way in credit default swap markets. While stocks with high idiosyncratic volatility have low returns, credit default swaps with high realised volatility earn high average returns. Table 6.2 reports monotonically increasing returns from low to high quintiles for one-, three-, six- and twelve-month windows. There is a strongly significant difference of 4.6% per annum (t=3.09) between the average twelve-month returns of the quintile portfolio with the highest idiosyncratic volatility and the quintile portfolio with the lowest idiosyncratic volatility. We find that the positive relation between realised volatility and average returns holds for different window lengths to measure firms' volatility.[2] Our findings suggest that firms with larger realised volatility require higher expected returns in credit default swap markets.

However, we find an interesting anomaly, that credit default swaps with very high volatility tend to have lower expected returns. While credit default swaps in the lowest volatility quintile have an annualised volatility of 7.5%, credit default swaps in the highest volatility quintile have a considerably higher average volatility of 23.6% per annum. The highest 10 credit default swaps have an average volatility of 36.5% per annum. These credit default swaps tend to underperform the highest volatility portfolio by, on average, 0.34% (1.47%) over the next month (twelve-months). This is similar to the anomaly of Ang *et al.* (2006b), who find that stocks with very high idiosyncratic volatility have remarkably low returns.

Furthermore, we consider coskewness risk to help understand the cross-sectional return variations in credit default swaps. Harvey and Siddique (2000) show that stocks with low or negative coskewness have high expected returns. We find that the high minus low coskewness portfolio yields a twelve-month return of 6.6% (t=4.04). While average returns increase with

[2]For each month, we also compute the idiosyncratic volatility using daily continuously compounded returns over the previous one-, three-, six- and twelve-months (four window lengths). After forming the quintile portfolios, sorted on historical volatility, we calculate the average expected returns of each quintile for four investment horizons (one-, three-, six- and twelve-month). We find that expected return differences between high and low volatility quintiles (volatility risk premium) for all 4$\times$4 portfolios are strongly significant. The expected volatility risk premium is 4.5% per annum ($t = 3.27$) for credit default swaps ranked according to past twelve month volatility. The annualised volatility risk premium is 5.9% for six-month returns (t=3.27), 6.2% for three-month returns (t=2.69), and 5.6% for one-month returns (t=2.95). There is no overlap between the historical returns used for estimating volatilities and expected future returns for evaluating the volatility risk premium. The volatility risk premium is higher among quintile portfolios sorted on one-month volatility. We find the average cross-sectional premium over the next month is 9.6% per annum (t=3.88). In addition, we compute the long-term volatility using monthly frequency returns over 36-month and 60-month horizons. We observe the same qualitative patterns that are statistically significant as using a 12-month horizon. Portfolios with higher volatility have higher average returns and quintile spreads between high and low are statistically significant. The volatility risk premium is 4.3% for twelve-month returns (t=2.91), 6.2% for six-month returns (t=2.91), 6.8% for three-month returns (t=3.58), and 6.7% for one-month returns (t=3.82) using a 36-month horizon. The volatilities for monthly returns over a 60-month horizon are reported in Table 6.2.

coskewness for credit default swaps, Harvey and Siddique (2000) find the opposite effect for stocks, where lower coskewness is associated with higher expected returns. Interestingly, both factors, idiosyncratic volatility and coskewness, tend to work in the opposite direction in credit markets. The following is our intuition for the opposing coskewness effect in credit markets. Everything else being equal, assume that price changes in corporate bonds, due to changes in default risk, move in the opposite direction to credit default swap prices. Thus, in ex-post returns, bonds with low (negative) coskewness should have corresponding credit default swaps with high (positive) coskewness. Everything else being equal, risk-averse investors should prefer risky bonds that are right-skewed to risky bonds that are left-skewed. Hence, corporate bonds that decrease in skewness (i.e., that make the return distribution more left-skewed), but have otherwise identical risk-characteristics, are less desirable and should demand higher expected returns. This should also be reflected in credit default swaps. Credit default swaps on bonds with negative coskewness must have higher expected returns than credit default swaps on bonds with identical risk-characteristics but zero- or positive coskewness. Therefore, coskewness may be an important factor for expected returns of credit default swap returns because of induced asymmetries in ex-post returns.

Finally, we identify a puzzling anomaly that firms with high credit default risk have low average returns over a twelve-month horizon. Panel A of Table 6.2 reports that portfolios with low default risk firms earn an average return of 6.0% per annum, whereas portfolios with high default risk firms only earn 3.1% per annum. This pattern is robust for different window lengths to calculate the historical probability of default from 1- to 36-months. However, this anomaly breaks down for the one-month returns. Panel D exhibits a monotonically increasing pattern between one-month returns and pre-formation one-month average default risk. The spread in average one-month returns between high and low credit default risk portfolios is 14.5% annualised (t=3.03%), which is the highest return among all factors.

Our results provide first evidence that firm characteristics explain differences in future expected returns of credit default swaps. There are multiple alternative explanatory characteristics that might have explanatory power in the cross section of returns, for example, leverage, debt ratios, the momentum effect (Jegadeesh and Titman, 1993; Carhart, 1997), cokurtosis risk (Scott and Horvath, 1980; Dittmar, 2002), liquidity risk (Pástor and Stambaugh, 2003), and downside beta risk (Ang *et al.*, 2006a).[3] We demonstrate next that high minus low tail risk return spreads are robust to controlling for these alternative factors.

[3] It seems appropriate to add a note of caution concerning the cokurtosis risk. Our tail risk exponent estimates suggest that the fourth and higher moments are not defined in cross-sectional returns of corporate credit default swaps.

Returns of equal-weighted credit default swap portfolios sorted by realised factor loadings

	Low	2	3	4	High	High-Low	*t-stat.*
Panel A: Twelve-month returns							
Market Risk Beta	-0.10%	1.90%	4.17%	6.08%	8.04%	8.14%	*4.70*
Firm's Credit Default Risk	5.95%	4.94%	3.94%	2.40%	3.09%	-2.86%	*1.95*
Idiosyncratic Volatility	1.26%	3.14%	4.11%	5.74%	5.86%	4.60%	*3.09*
Coskewness	0.05%	2.74%	5.69%	4.99%	6.66%	6.61%	*4.04*
Credit Tail Risk Beta	0.28%	1.71%	4.09%	5.65%	8.37%	8.08%	*4.22*
Panel B: six-month returns							
Market Risk Beta	-0.19%	0.80%	1.99%	3.21%	5.29%	5.48%	*4.51*
Firm's Credit Default Risk	2.46%	2.39%	1.68%	1.74%	2.98%	0.52%	*2.01*
Idiosyncratic Volatility	0.41%	1.67%	2.01%	3.14%	3.88%	3.46%	*2.80*
Coskewness	-0.08%	1.69%	2.67%	2.85%	3.99%	4.07%	*3.58*
Credit Tail Risk Beta	-0.06%	0.99%	2.08%	3.50%	4.60%	4.66%	*4.00*
Panel C: Three-month returns							
Market Risk Beta	-0.15%	0.34%	0.85%	1.36%	2.83%	2.98%	*4.95*
Firm's Credit Default Risk	0.72%	0.75%	0.76%	1.00%	2.06%	1.35%	*2.58*
Idiosyncratic Volatility	0.08%	0.73%	1.00%	1.42%	2.02%	1.94%	*3.30*
Coskewness	-0.11%	0.72%	1.28%	1.57%	1.78%	1.88%	*3.22*
Credit Tail Risk Beta	-0.11%	0.53%	0.83%	1.68%	2.31%	2.42%	*3.93*
Panel D: One-month returns							
Market Risk Beta	-0.04%	0.11%	0.19%	0.38%	0.85%	0.90%	*4.70*
Firm's Credit Default Risk	-0.22%	0.08%	0.16%	0.47%	1.00%	1.21%	*3.03*
Idiosyncratic Volatility	0.01%	0.16%	0.33%	0.34%	0.64%	0.63%	*3.35*
Coskewness	-0.03%	0.21%	0.37%	0.53%	0.40%	0.42%	*2.72*
Credit Tail Risk Beta	-0.04%	0.16%	0.25%	0.49%	0.62%	0.66%	*3.63*

Table 6.2 reports the equal-weighted average returns of US corporate credit default swaps for quintile portfolios formed on the basis of market risk (beta), firm's credit default risk, idiosyncratic volatility and coskewness. For each month, we reconstitute portfolios based on updated realised factor loadings. Portfolios consists of one credit default swap contract for each firm in the corresponding quintiles (equal-weighted). Portfolios are based on credit default swaps on senior unsecured debt, traded under the Cum Restructuring clause. For comparison purposes, we add average returns for quintile portfolios created on the basis of credit tail risk sensitivities. Panel A reports average out-of-sample one-year holding period portfolio returns. Portfolios with high credit tail and market risk earn the highest average returns. Panel B, C and D report average out-of-sample returns for holding periods of six-, three-, and one-month. Portfolios with high exposure to corporate default risk yield the highest average returns for a one-month investment horizon. For all holding periods, average portfolio returns demonstrate the same monotonic pattern, where returns increase with increasing market risk sensitivity, idiosyncratic volatility, cosknewness and credit tail risk sensitivity. The right-most columns report returns for high minus low portfolios and corresponding t-statistics.

6.3.3 Credit Tail Risk, Alternative Risk Factors and Expected Returns

In this section, we empirically assess whether the credit tail risk premium is robust to controlling for alternative risk factors. These risk factors are market risk beta, idiosyncratic volatility, and coskewness.

Market and Credit Tail Risk

Table 6.2 shows that both credit tail risk sensitivity and market beta have very robust, predictive power for the cross section of credit default swap returns. While market risk equally captures the effect of upside and downside risk, credit tail risk is an asymmetric measure that explicitly emphasises large positive price fluctuations associated with a significant increase in implied default risk. We now measure the magnitude of the reward for exposure to credit tail risk sensitivity while explicitly controlling for market risk beta.

To control for the effect of market risk, we first construct quintile portfolios ranked on market risk. Then, within each market risk quintile, we sort credit default swaps into three equally-weighted portfolios based on their credit tail risk sensitivity. Both market risk and credit tail risk regressions are estimated over the same 60-month horizon. We calculate the expected future returns over the next one-, three-, six- and twelve months. After forming the 5×3 market and credit tail risk portfolios, we compute the differences in expected returns between the highest and the lowest market risk portfolio for each credit tail risk tercile. This control method creates a set of portfolios with near-identical spreads of market risk. The spreads in market risk between fifth and first quintiles are 1.51, 1.55 and 1.55.[4] Hence, the spread of credit tail risk tercile portfolios controls for differences in market risk.

Panel A of Table 6.3 reports average twelve-month returns of the 12 market beta and credit tail risk portfolios. Each portfolio has credit default swaps of 45 firms. The column labelled "High-Low" reports the expected twelve-month premiums of the credit tail risk terciles controlling for market risk. The difference between the highest and the lowest market risk quintile for low tail risk portfolios is 5.6% per annum, with a t-statistic of 2.11. The high minus low market risk quintile for mid-tail risk portfolios is 7.8% per annum, with a t-statistic of 2.74. For high tail risk portfolios, the difference between the highest and the

[4]Before the characteristic control procedure, the average market risk for each tercile portfolio is 0.48, 0.81 and 1.20. After the characteristic control procedure, each tercile portfolio's average market risk is similar (0.77, 0.83, and 0.89), but still increasing with tail risk. However, the differences in market risk between fifth and first quintiles are near-identical: 1.51, 1.55 and 1.55.

lowest market risk quintile is much larger in magnitude at 11.0% per annum and statistically significant (t=3.43). While spreads in market risk between first and fifth quintile are almost identical for all levels of tail risk, the expected risk premium (difference in spreads) increases with credit tail risk.

The difference in expected risk premium is 5.5% per annum between high and low tail risk portfolios for the same market risk level. Indeed, the second and third credit tail risk terciles have 2.0 and 3.1 times higher factor loadings than the first tercile. Thus, market risk cannot account for expected returns from bearing credit tail risk.

Panel A of Table 6.3 reports average twelve-month returns of 12 market risk and credit tail risk portfolios. The rows labelled "low", "mid" and "high" report the average returns for each market risk quintile given the level of credit tail risk. In Panel A, the patterns within each credit tail risk tercile are interesting. As market risk increases from the lowest to the highest market risk quintile, the expected returns monotonically increase. This pattern holds for all three tail risk terciles. Across the low to high market beta quintiles, the average beta coefficient increases from 0.08 to 1.61.[5] Furthermore, as the credit tail risk sensitivity increases (low, mid, high), the high minus low quintile returns increase (5.6%, 7.8%, 11.0%). This effect is quite pronounced for high levels of credit tail risk. In the third credit tail risk tercile, the average return for low market risk is -0.8% per annum and high market risk is 10.2% per annum. The monotonic return patterns and spreads between high-low market risk quintiles are statistically significant for one-, three- and six-month out-of-sample returns.

Our reasons for these patterns are as follows. Credit default swaps with low market beta tend to have low expected returns even when credit tail risk is high (reading down the "Low" column). The market beta equally captures the effect of upside and downside movements when the overall credit default swap market increases or decreases. Low market beta firms have lower expected returns when tail risk is high, because (i) firms may recover less despite a market recovery, and (ii) the tail risk sensitivity level is significantly less (approximately 65%) compared to high tail risk credit default swaps in the other quintiles. Furthermore, low levels of tail risk sensitivity tend to be persistent in future. While higher credit tail risk is usually rewarded in the cross section of future returns, this may not apply to low market beta firms. In contrast, firms with high exposure to market risk tend to have high expected returns, which increase in credit tail risk (reading down the "High" column). The tail risk loadings monotonically increase from low to high market beta. Thus, portfolios with high market risk

[5]The average beta coefficient (average return per annum) is 0.08 (-0.1%) for the first quintile, 0.50 (1.9%) for the second, 0.83 (4.2%) for the third, 1.14 (6.1%) for the fourth, and 1.61 (8.0%) for the fifth market beta quintile. The average returns for each market beta quintile are consistent with returns reported in Table 6.2.

Double-sorted portfolio returns

	Low	2	3	4	High	High-Low	*t-stat.*	Average
Panel A: Market Risk Quintiles and Credit Tail Risk								
	Future Annual Return (t+12)							
Low Credit Tail Risk	0.56%	0.92%	3.17%	4.83%	6.11%	5.55%	*2.11*	3.12%
Mid Credit Tail Risk	0.02%	1.31%	4.86%	5.04%	7.79%	7.77%	*2.74*	3.80%
High Credit Tail Risk	-0.84%	3.49%	4.49%	8.41%	10.20%	11.03%	*3.43*	5.15%
	Future Monthly Return (t+1) - Annualised							
Low Credit Tail Risk	0.04%	0.68%	0.24%	5.25%	12.68%	12.63%	*2.77*	3.78%
Mid Credit Tail Risk	-0.88%	2.05%	2.37%	3.50%	6.48%	7.36%	*3.00*	2.70%
High Credit Tail Risk	-0.66%	1.04%	4.19%	4.89%	11.69%	12.34%	*2.98*	4.23%
	Low	2	3	4	High	High-Low	*t-stat.*	Average
Panel B: Volatility Quintiles and Credit Tail Risk								
	Future Annual Return (t+12)							
Low Credit Tail Risk	0.47%	-0.12%	1.23%	3.79%	2.68%	2.21%	*0.92*	1.61%
Mid Credit Tail Risk	1.08%	3.95%	2.43%	5.07%	7.68%	6.60%	*1.96*	4.04%
High Credit Tail Risk	2.23%	5.56%	8.65%	8.34%	7.23%	5.00%	*2.16*	6.40%
	Future Monthly Return (t+1) - Annualised							
Low Credit Tail Risk	-0.75%	1.01%	1.17%	2.96%	3.98%	4.73%	*1.56*	1.67%
Mid Credit Tail Risk	0.95%	1.80%	3.48%	4.34%	10.02%	9.06%	*2.52*	4.12%
High Credit Tail Risk	0.16%	2.94%	7.33%	4.98%	9.20%	9.04%	*2.25*	4.92%
	Low	2	3	4	High	High-Low	*t-stat.*	Average
Panel C: Coskewness Quintiles and Credit Tail Risk								
	Future Annual Return (t+12)							
Low Credit Tail Risk	-0.42%	0.86%	1.24%	1.73%	3.93%	4.35%	*1.79*	1.47%
Mid Credit Tail Risk	-0.27%	3.59%	4.54%	4.87%	6.68%	6.94%	*2.47*	3.88%
High Credit Tail Risk	0.84%	3.74%	11.26%	8.39%	9.37%	8.53%	*2.96*	6.72%
	Future Monthly Return (t+1) - Annualised							
Low Credit Tail Risk	-0.52%	0.70%	1.32%	2.01%	2.81%	3.33%	*1.38*	1.26%
Mid Credit Tail Risk	-0.60%	3.81%	2.68%	5.60%	5.34%	5.94%	*1.78*	3.37%
High Credit Tail Risk	0.14%	2.92%	9.39%	11.41%	6.11%	5.97%	*2.01*	5.99%

Table 6.3 reports expected future returns of credit default swaps for double-sorted portfolios that are formed on the basis of market beta (Panel A), idiosyncratic volatility (Panel B), coskewness risk (Panel C) and tail risk sensitivity. For each month, we estimate each market beta, volatility, coskewness and tail risk sensitivity over a 60-month horizon using monthly frequency returns. At the beginning of each month, firms are sorted independently into quintiles based on the market beta, volatility and coskewness risk (columns). Then, within each quintile, firms are ranked into terciles based on (low, mid, high) credit tail risk sensitivity (rows). We then report the average out-of-sample returns for each equal-weighted 5×3 portfolio over the next one- and twelve-months. Portfolios are based on credit default swaps on senior unsecured debt, traded under the Cum Restructuring clause. The column labelled "High-Low" reports the return spread (premium) for the high minus low quintile portfolio. The next column reports the t-statistics for the "High-Low" difference. The column labelled "Average" reports the average expected return across quintiles for low-, mid- and high credit tail risk.

contain credit defaults with high tail risk sensitivity levels, which is usually rewarded in the cross section of future returns. Hence, firms with high market and tail risk have the highest expected returns, because (i) the cross section of returns reflect a premium for bearing market risk, and (ii) a cross-sectional premium for holding credit default swaps with high tail risk. This explains the spreads in average returns across the market risk quintiles.

Volatility and Credit Tail Risk

Table 6.2 shows that high exposure to past volatility is associated with higher future returns for the cross section of credit default swaps. While credit tail risk only considers extreme upside risk, volatility treats extreme upside and extreme downside movements symmetrically. We now estimate the premium for exposure to credit tail risk while explicitly controlling for short- and long-term volatility effects.

Next, to control the volatility effect, we first form equally-weighted quintile portfolios with respect to idiosyncratic volatility in Panel B of Table 6.3. We perform the first ranking of credit default swaps according to past volatility risk. Then, within each volatility quintile, we sort credit default swaps into three groups (low, mid, high) based on their credit tail risk sensitivity. Both volatility and credit tail risk sensitivity are estimated over the same 60-month horizon. We then construct the 5×3 portfolios within each doubly-sorted group and report the expected future returns over the next one-, three-, six- and twelve months. After forming the volatility and credit tail risk portfolios, we average each volatility quintile's expected returns over the three credit tail risk terciles. This control approach creates a set of tail risk portfolios with almost identical levels of volatility. The low-, mid- and high credit tail risk portfolios have, on average, an annualised volatility of 15.1%, 15.0% and 16.0%. Thus, these credit tail risk tercile portfolios control for differences in volatility.[6]

Panel B of Table 6.3 reports average twelve-month returns of the 5×3 volatility and credit tail risk portfolios. Each portfolio has credit default swaps of 45 firms. Panel B shows a positive relation between past credit tail risk and future expected returns. The future expected returns increase in all quintiles, except for the highest volatility quintile. The column labelled "Average" reports average expected twelve-month returns of the credit tail risk terciles by averaging across the quintiles. This analysis examines the tail risk premium controlling for

[6] We also compute individual credit default swap volatility using daily continuously compounded returns over the previous 12 months. Then we repeat the same double-sorting procedure and portfolio formations. Panel B of Table 6.5 reports almost identical levels of historical volatility of 36.8%, 36.1% and 35.5% for low-, mid- and high credit tail risk portfolios.

volatility. The average tail risk premium of 4.8% per annum is the difference in average returns between the third and first tercile portfolio that controls for volatility risk. We find that the tail risk premium is statistically significant, with a t-statistic of 3.16. Credit tail risk premiums have the same qualitative behaviour for short-horizon portfolio returns. The annualised credit tail risk premium is 4.8%, 4.5% and 3.2% for six-, three-, and one-month returns. Hence, controlling for volatility in returns, there is a strong predictive relation between past tail risk and expected future returns in the cross section of credit default swaps.[7]

Furthermore, we assess the high minus low volatility portfolio for low-, mid- and high-tail risk sensitivity. Table 6.2 reports that volatility risk is positively related to future expected returns. The volatility premium is persistent for short- and long-term volatility measures and holds for investment horizons from one- to twelve months. Panel B of Table 6.3 reports that credit default swaps with high volatility exhibit high expected returns over the next twelve months. However, the high minus low volatility portfolio only exhibits significant return differences for portfolios with mid- and high credit tail risk. Low tail risk portfolios fail to show this volatility premium. Interestingly, these different predictive patterns of volatility premiums for low-, mid- and high-tail risk portfolios hold for one-, three-, six- and twelve months out of sample returns.

To help interpret the differences in volatility premiums, we investigate the relationship between historical and future volatility of portfolios. Firstly, we examine whether credit default swaps ranked on historical volatility provide enough variation in future volatility to rationalise the reward for bearing high volatility credit default swaps. Secondly, we examine further the link between historical and future volatility exposure for different levels of tail risk.

For our first exercise, we sort credit default swaps into quintile portfolios at time t based on pre-formation volatility, and then examine future volatilities within each quintile over the next one-, three-, six- and twelve months. At the beginning of each month t, we reconstruct the portfolios based on updated historical volatility. We also consider different effects of short- and long-term volatility. We compute historical volatilities on daily returns over the previous 12 months, and monthly returns over a 60-month horizon. We find a strictly increasing pattern for future volatility of credit default swaps sorted on historical volatility. This means

[7] We repeat this exercise for individual credit default swap volatility using daily continuously compounded returns over the previous 12 months. Expected returns increase with tail risk in all volatility quintiles over the next 12 months. The credit tail risk premium is 5.7% per annum, with a t-statistic of 3.76. The annualised credit tail risk premium is 6.5% for six-month returns (t=3.85), 6.3% for three-month returns (t=3.45), and 4.8% for one-month returns (t=2.49), while explicitly controlling for the volatility effect.

that credit default swaps with low (high) volatility continue to have low (high) volatility going forward. However, there is an interesting trend in $t+1$, $t+3$, $t+6$ and $t+12$ within each quintile for past twelve month volatility. Table 6.5 clearly shows that credit default swaps with lower historical volatility slightly higher volatility going forward. On the other hand, credit default swaps in the highest volatility quintile tend to decrease in future volatility. These trends are persistent from $t+1$ to $t+12$. While the difference in future volatility between the first and fifth quintile decreases, it remains highly positive and statistically significant at the 1% level, which helps to explain the future return premium for volatility. Note that the historical volatility over the previous 12 months has no overlapping returns with future volatility over the next 12 months. Nevertheless, the difference in annualised volatility between the first and fifth quintile remains high at 15.9% (t=6.11) with a corresponding return premium of 4.5% per annum (t=3.27). Hence, historical volatility seems to predict future volatility, with significant spreads in both historical and future volatility. Thus, spreads in historical volatility are reflected in future return spreads, which are consistent with future volatility spreads. We conclude that past realised volatility is rewarded in the cross section of future returns.

While the relation between volatility and risk premium holds for single-sorted portfolios, we observe that the volatility premium is insignificant for low tail risk portfolios of doubly-sorted portfolios. For low tail risk portfolios, the difference in historical long-term (short-term) volatility between high minus low portfolios is 18.0% (24.6%) and statistically significant at the 1% level. However, these low tail risk portfolios fail to show predictive patterns of volatility premiums for one-, three-, six- and twelve-month out of sample returns. Therefore, in our second exercise, we investigate future volatilities on double-sorted portfolios. We sort credit default swaps into quintile portfolios at time t based on pre-formation volatility, and then perform the second ranking on pre-formation tail risk sensitivity. We then form 5×3 portfolios and report future volatilities over the next one-, three-, six- and twelve-month periods. We find that the strictly increasing pattern for future volatility sorted on historical volatility holds for low-, mid- and high tail risk terciles. This pattern holds for (short- and long-term) future volatilities over the next one-, three-, six- and twelve-month periods. Low tail risk portfolios have future volatility spreads with similar magnitude as the historical volatility spread, both highly statistically significant. This means that high-low spreads for historical and future volatility do not provide a possible explanation for the insignificant volatility premium for low tail risk portfolios. Hence, it remains to be explored why portfolios with low credit tail risk do not hold a statistically significant premium for credit default swaps with high levels of volatility.

Historical and Future Volatility (Long-Term)

	Low	2	3	4	High	High-Low	*t-stat.*	Average
Panel A: Future Volatility Sorted on Historical Volatility								
Historical Volatility	7.5%	12.4%	15.3%	17.9%	23.6%	16.1%	35.29	15.3%
Future Volatility (t+1)	7.5%	12.3%	15.2%	17.7%	23.3%	15.9%	34.64	15.2%
Future Volatility (t+3)	7.3%	12.0%	14.9%	17.4%	22.8%	15.5%	33.31	14.9%
Future Volatility (t+6)	7.1%	11.6%	14.7%	17.2%	22.2%	15.1%	31.76	14.6%
Future Volatility (t+12)	6.8%	11.1%	14.4%	16.8%	21.1%	14.3%	28.15	14.0%

	Low	2	3	4	High	High-Low	*t-stat.*	Average
Panel B: Future Volatility Double-Sorted on Historical Volatility and Credit Tail Risk								
	Future Volatility (t+12)							
Low Credit Tail Risk	5.5%	11.1%	14.0%	16.1%	20.9%	15.4%	14.57	13.5%
Mid Credit Tail Risk	6.6%	11.1%	14.5%	16.6%	20.6%	14.0%	19.65	13.9%
High Credit Tail Risk	8.3%	11.1%	14.7%	17.7%	21.8%	13.5%	17.97	14.7%
	Future Volatility (t+6)							
Low Credit Tail Risk	5.7%	11.5%	14.5%	16.7%	22.3%	16.6%	17.57	14.1%
Mid Credit Tail Risk	6.8%	11.5%	14.7%	17.1%	21.2%	14.4%	23.39	14.3%
High Credit Tail Risk	8.9%	11.9%	15.0%	17.8%	23.1%	14.2%	18.62	15.3%
	Future Volatility (t+3)							
Low Credit Tail Risk	5.9%	11.9%	14.9%	17.2%	23.0%	17.1%	19.28	14.6%
Mid Credit Tail Risk	6.9%	11.9%	14.9%	17.3%	21.6%	14.7%	25.88	14.5%
High Credit Tail Risk	9.1%	12.2%	15.1%	17.8%	23.7%	14.6%	18.34	15.6%
	Future Volatility (t+1)							
Low Credit Tail Risk	6.0%	12.2%	15.1%	17.6%	23.7%	17.7%	20.73	14.9%
Mid Credit Tail Risk	7.1%	12.1%	15.1%	17.6%	22.0%	14.9%	27.79	14.8%
High Credit Tail Risk	9.3%	12.4%	15.3%	17.9%	24.2%	14.9%	18.53	15.8%
	Historical Volatility (t+0)							
Low Credit Tail Risk	6.0%	12.4%	15.3%	17.7%	24.1%	18.0%	21.47	15.1%
Mid Credit Tail Risk	7.2%	12.3%	15.2%	17.8%	22.3%	15.1%	28.54	14.9%
High Credit Tail Risk	9.4%	12.6%	15.4%	18.1%	24.5%	15.1%	18.68	16.0%

Table 6.4 reports the relationship between past realised volatility and future volatility. For each month, we compute the historical volatility of individual credit default swaps on realised returns over a 60-month horizon using monthly frequency returns. Panel A examines whether credit default swaps ranked on historical volatility provide enough variation in future volatility to rationalise the future reward for bearing high volatility firms. We rank firms into quintiles at the beginning of each month based on historical volatility calculated over the previous 60 months. We then examine the future volatility within each quintile over the next one-, three-, six- and twelve months. Panel B examines the link between historical and future volatility for different levels of credit tail risk. We sort credit default swaps into historical volatility quintiles for each month and then perform the second sorting based on pre-formation tail risk. We then form 5×3 portfolios and report future volatilities over the next one- to twelve-month periods. The column labelled “High-Low” reports the difference between future volatilities of quintile portfolios "High" and portfolio "Low". The next column reports the t-statistics for the “High-Low” difference. The column labelled “Average” reports the average of all quintile portfolios.

Historical and Future Volatility (Short-Term)

	Low	2	3	4	High	High-Low	*t-stat.*	Average
Panel A: Future Volatility Sorted on Historical Volatility								
Historical Volatility	2.6%	7.2%	11.1%	14.6%	25.3%	22.7%	20.95	12.2%
Future Volatility (t+1)	2.8%	7.3%	11.3%	14.6%	25.0%	22.2%	20.04	12.2%
Future Volatility (t+3)	3.3%	7.7%	11.7%	14.9%	24.1%	20.8%	17.95	12.3%
Future Volatility (t+6)	4.1%	8.6%	12.5%	15.6%	23.0%	18.9%	15.07	12.8%
Future Volatility (t+12)	6.0%	10.5%	14.4%	17.3%	20.9%	14.9%	10.55	13.8%

	Low	2	3	4	High	High-Low	*t-stat.*	Average
Panel B: Future Volatility Double-Sorted on Historical Volatility and Credit Tail Risk								
	Future Volatility (t+12)							
Low Credit Tail Risk	5.5%	10.5%	14.2%	16.4%	18.0%	12.5%	5.59	12.9%
Mid Credit Tail Risk	5.8%	10.4%	13.5%	17.0%	21.3%	15.5%	7.23	13.6%
High Credit Tail Risk	6.7%	10.8%	15.3%	18.5%	23.5%	16.8%	6.56	15.0%
	Future Volatility (t+6)							
Low Credit Tail Risk	3.5%	8.4%	12.7%	15.0%	22.3%	18.8%	8.14	12.4%
Mid Credit Tail Risk	4.1%	8.4%	11.9%	15.7%	22.5%	18.4%	10.80	12.5%
High Credit Tail Risk	4.8%	8.9%	12.9%	16.3%	24.2%	19.4%	8.83	13.4%
	Future Volatility (t+3)							
Low Credit Tail Risk	2.7%	7.7%	11.9%	14.7%	24.6%	21.9%	10.19	12.3%
Mid Credit Tail Risk	3.2%	7.5%	11.4%	15.0%	23.3%	20.1%	12.75	12.1%
High Credit Tail Risk	3.9%	7.9%	11.8%	15.1%	24.4%	20.5%	10.22	12.6%
	Future Volatility (t+1)							
Low Credit Tail Risk	2.1%	7.3%	11.3%	14.6%	26.0%	23.8%	12.16	12.3%
Mid Credit Tail Risk	2.7%	7.2%	11.2%	14.7%	24.1%	21.4%	13.49	12.0%
High Credit Tail Risk	3.5%	7.4%	11.3%	14.6%	24.9%	21.4%	11.09	12.3%
	Historical Volatility (t+0)							
Low Credit Tail Risk	1.9%	7.2%	11.1%	14.7%	26.4%	24.6%	13.22	12.3%
Mid Credit Tail Risk	2.6%	7.1%	11.1%	14.6%	24.5%	21.9%	13.60	12.0%
High Credit Tail Risk	3.4%	7.3%	11.2%	14.5%	25.1%	21.7%	11.32	12.3%

Table 6.5 examines the relationship between past realised volatility and future volatility. For each month, we compute the historical volatility of individual credit default swaps using continuously compounded returns over the past 12 months. Panel A examines whether credit default swaps ranked on historical volatility provide enough variation in future volatility to rationalise the future reward for bearing high volatility firms. We rank firms into quintiles at the beginning of each month based on historical volatility calculated over the previous 12 months. We then examine the future volatility within each quintile over the next one-, three-, six- and twelve months. Panel B examines the link between historical and future volatility for different levels of credit tail risk. We sort credit default swaps into historical volatility quintiles for each month and then perform the second sorting based on pre-formation tail risk. We then form 5×3 portfolios and report future volatilities over the next one- to twelve-months. The column labelled “High-Low” reports the difference between future volatilities of quintile portfolios "High" and portfolio "Low". The next column reports the t-statistics for the “High-Low” difference. The column labelled “Average” reports the average of all quintile portfolios.

Coskewness and Credit Tail Risk

In Section 6.3.2, we discuss that credit default swaps with positive (negative) exposure to coskewness risk have high (low) expected future returns. While Harvey and Siddique (2000) find the opposite effect in equity markets, we provide a simple explanation for credit default swap markets. Credit default swaps with high coskewness exhibit a higher probability that a firm's default risk increases more in relation to increases in default risk of the market (average default risk of all firms). Hence, sellers of credit default insurances increase prices more, which induces asymmetric ex-post returns of credit default swaps with positive coskewness. Since both tail risk and coskewness risk capture the effect of asymmetric higher moments and higher default risk, we now evaluate the expected premium for exposure to credit tail risk, while explicitly controlling for coskewness.

We control for the effect of coskewness by forming quintile portfolios sorted on past coskewness. Then, within each quintile, we sort credit default swaps into three groups (low, mid, high) based on their sensitivity to credit tail risk. Both credit tail risk sensitivity and coskewness are estimated over the same 60-month horizon. After constructing the 5×3 coskewness and credit tail risk portfolios, we average the expected future returns of each tail risk tercile over the five coskewness portfolios. This control measure creates a set of tercile portfolios with almost identical levels of coskewness risk of 0.52, 0.54, 0.55 for low-, mid- and high credit tail risk. Hence, these tail risk portfolios control for differences in coskewness.

Panel C of Table 6.3 reports average returns of the 12 coskewness and tail risk portfolios for out-of-sample returns of one- and twelve-months. The column labelled "Average" reports the average expected returns of the three tail risk portfolios controlling for coskewness risk. The low credit tail risk portfolio has an average expected 12-month return of 1.5%. The high credit tail risk portfolio has an average expected 12-month return of 6.7%. Controlling for coskewness, the expected premium for credit tail risk is 5.2% per annum. The premium for credit tail risk has a robust t-statistic of 3.46. Note that the credit tail risk premium is 4.7% per annum for a monthly holding period, with t-statistics of 2.60. The credit tail risk in the second and third terciles have 3.4 and 5.9 times higher factor loadings than the first tercile. Hence, coskewness risk cannot account for the premium for exposure to higher credit tail risk.

Panel C of Table 6.3 reports consistently low returns for low credit tail risk portfolios within each coskewness quintile. The difference between high and low credit tail risk portfolio is statistically significant for the third, fourth, and fifth quintile for 12-month out-of-sample returns. Furthermore, the first and second coskewness quintiles have low average returns (average within quintile) of 0.1% and 2.7% per annum. In contrast, the fourth and fifth coskewness quintile have average returns of 5.0% and 6.7% per annum. This effect is also quite pronounced for one-month returns. The (low-low) portfolio with low coskewness and low tail risk has a negative return of -0.4% per annum, whereas the portfolio with high coskewness and high tail risk earns 9.4%. The difference of 9.8% between these two portfolios is statistically significant, with a t-statistic of 3.62.

The reasons for these patterns are as follows. Coskewness is effectively the covariance of a credit default swap's return with the volatility of the market. A credit default swap with negative (or low) coskewness tends to have low returns when market volatility is high. These are usually, but not always, periods of financial distress, where prices of credit default swaps tend to increase rapidly. As mentioned previously, the volatility of the market treats upside and downside risk symmetrically, so both extreme upside and extreme downside returns of the market have the same volatility.

The first and second coskewness quintiles have low average returns of 0.1% and 2.7% because low coskewness has limited capacity to account for credit default swaps with tail risk. Credit default swaps with negative (or low) coskewness exhibit a lower probability that a firm's default risk disproportionally increases when market default risk increases. The prices of credit default swaps with (large) negative coskewness tend to decrease when markets increase, but the prices of these credit default swaps may also decrease when the market recovers quickly. On the other hand, prices of credit default swaps with (large) positive coskewness tend to increase in periods of financial distress (high volatility of the market). However, these credit default swaps' prices may also increase during a rapid market recovery (or recover less in relation to the market). In contrast, credit tail risk concentrates only on the former effect by explicitly considering significant increases in default risk, and consequently, large increases in insurance prices. Therefore, high coskewness credit default swaps capture high sensitivity to credit tail risk, because extreme price increases tend to coincide with periods of high market volatility. On the contrary, low coskewness credit default swaps only capture significant price increases (tail returns) of idiosyncratic tail risk shocks in periods of low market volatility. This implies that these low coskewness credit default swaps do not catch high sensitivity to credit tail risk in periods of high market volatility (systemic market shocks). The limited capacity to account for high sensitivity to tail risk, also explains the low and statistically insignificant tail risk premium for the first and second coskewness quintiles. Furthermore, the low coskewness quintile contains firms with sensitivity to credit tail risk, ranging from -0.7 to 3.9 (average 1.5). In comparison, the high coskewness quintile comprises firms with sensitivity to credit tail risk, ranging from 2.4 to 6.8 (average 4.4). On average, the range of credit tail risk sensitivity is 3.2 times higher for the high coskewness quintile compared to the low quintile. Hence, the latent exposure to high credit tail risk explains the large average returns for credit default swaps with high coskewness.

Across the first coskewness quintile, the coskewness ranges from -0.09 to 0.03. The first coskewness quintile has very little coskewness. The small (negative) return of -0.4% per annum for the low-low portfolio is due to the low asymmetry of credit default swaps. The distribution of coskewness across credit default swaps is negatively skewed (-2.34) and is negative each month, on average (-0.09). This means that credit default swaps do not change their behaviour across periods of high and low market volatility. High market volatility also tends to coincide with periods of financial distress when default correlations between various firms tend to increase sharply. Thus, these credit default swaps provide a hedge against changes in market volatility, because increasing volatility usually presents a deterioration in diversification opportunities. Therefore, these credit default swaps are attractive and have lower expected returns. Finally, credit default swaps that decrease in price when market volatility increases, also tend to have negatively skewed returns in subsequent periods.[8]

6.4 Conclusion

The cross section of credit default swap returns reflects a premium for tail risk. Firms sensitive to tail risk in credit markets have high expected returns. The risk-return relation is consistent with an economy where protection sellers increase insurance prices for assets which strongly covary with market shocks. Protection sellers with an aversion to credit tail risk require a premium for selling credit default swaps with high tail risk sensitivity. Hence, corporate credit defaults with high predictive loadings to tail risk have higher expected returns.

We find that credit default swaps in the high tail risk sensitivity quintile earn equal-weighted annual returns 8.1% higher than credit default swaps in the low tail risk portfolio. Past credit tail risk sensitivity is a good predictor of future credit tail risk sensitivity and is rewarded in the cross section of future returns. We find that past credit tail risk sensitivity provides enough divergence in future credit tail risk sensitivity to justify the tail risk premium

[8]Credit default swaps with negative coskewness and low sensitivity to credit tail risk have negatively skewed returns in subsequent periods, which are negative, on average, in each period. These credit default swaps have a negative skewness of -1.75, -0.97, -0.86 and -0.97 for one-, three-, six- and twelve-month out-of-sample returns.

between high and low quintile portfolios.

Other risk factors, such as the exposure to market risk beta, idiosyncratic volatility and coskewness risk, also explain differences in future expected returns of credit default swaps. However, we find that the credit tail risk premium is different from the premiums on market risk, volatility and coskewness.

Concluding Remarks

A measure of tail risk in credit markets is essential to understand the behaviour of credit default swaps. This thesis presents three tail-risk measures based on dynamic power-law models with multiple time-varying tail parameters. The models use univariate and cross-sectional returns of sovereign and corporate credit default swaps to estimate the tail risk at each point of time. Large returns are compatible with the power-law hypothesis, suggesting that the power-law distribution is a reasonable measure of tail risk in credit markets. Tests show that past exposure to extreme event risk has a significant impact on future credit default swap prices and returns. This conclusion holds for sovereign and corporate credit default swaps in Europe and the United States.

Chapter 1 presents the first measure of extreme event risk based on a dynamic power-law model with three time-varying tail parameters (tail threshold, tail exponent, tail percentile). This method combines maximum-likelihood fitting methods with goodness-of-fit tests based on the Kolmogorov-Smirnov statistic, which presents a quantitative traceable method for defining the tail threshold and, consequently, the tail length and percentile. The time-varying power-law approach uses daily returns of 35 univariate time series to inform estimates of tail risk at each point of time. The key finding is that the dynamic power-law is a plausible hypothesis for extreme returns in credit default swap markets. Furthermore, the tail exponents are significantly time-varying for credit default swaps on sovereign debt for all seven European countries and five contract maturities. Results report that the average and most probable value for tail exponents is 3, consistent with the inverse cubic law in other asset classes. The findings are robust for positive and negative returns, normalised returns, absolute returns and various lookback windows.

Chapter 2 investigates the factors affecting changes in the time-varying tail exponent. The difference in two consecutive tail exponents can originate from changes in tail returns, changes of tail length, or both factors simultaneously. Empirical evidence shows that changes in tail exponents most often emerge from simultaneous variations in extreme returns and tail length. The main challenge is to separately quantify changes in the maximum likelihood estimator of the tail exponent when factors coincide. To overcome this problem, we develop a novel tail exponent decomposition method. Our decomposition method quantifies the tail exponent changes due to tail returns and tail length variations separately. The decomposition analysis shows that 81% of the tail exponent variations are due to changes in tail returns given a non-abrupt change in tail length. However, one-third of tail exponent variations concur with larger variations of the tail length. In the presence of abrupt changes in tail length, the variation in extreme returns only explains 22% of the change in the tail exponent. Thus, the major part of tail exponent changes (88%) is due to large tail length variations. We find that longer time series significantly decrease the number of large tail length variations but reduce the model adaptability to recent tail risk changes. Furthermore, we find that the aggregation of returns across the term structure of default substantially reduces the dependency on past data, increases adaptability to current tail risk changes, decreases estimation errors and the number of abrupt changes in tail length. The findings in credit default swap markets are similar to the tail exponent decomposition results in national stock market indices. We conclude that the Kolmogorov-Smirnov method must be treated with caution in the time-varying tail exponent analysis. These new insights motivate us to propose the second tail risk measure based on a smoothing technique for the tail threshold and the use of cross-sectional data.

Chapter 3 presents the second tail risk measure with an application to asset pricing. The dynamic tail risk model in Chapter 1 has two key disadvantages. Firstly, the Kolmogorov-Smirnov method causes abrupt changes in the tail threshold and tail length. This results in changes in the tail exponent, which could not be attributed to changes in extreme returns. To overcome this problem, the new dynamic power-law model incorporates a smoothing technique for the time-varying tail threshold that eliminates tail exponent fluctuations due to abrupt changes in the tail length. Secondly, estimating time-varying tail risk in univariate time series is challenging because of the infrequent nature of extreme events. Measuring tail risk in univariate data relies on long time series, which may include tail events from the distant past (previous crises) with no causality to the current market situation. The new dynamic tail risk model overcomes this difficulty by aggregation of returns across the term structure of default. Aggregated return data significantly reduce the lookback window from several to one year and increase the adaptability of the tail risk measure. We implement the dynamic tail risk estimator using daily returns from ten maturities of US sovereign credit default swaps with 1-year to 30-year tenor. We find that a one-standard-deviation increase in tail risk forecasts an average increase in US sovereign credit default swap spreads of 7.6 bps, which is highly significant. We explore the robustness of the forecasting power of the credit tail risk measure to controlling for 25 alternative predictors. We conclude that increases in tail risk significantly predict increases in credit default swap prices.

Chapter 4 studies the relationship between tail risk and maturities in the cross section of global sovereign credit default swaps. This relationship is referred to as the term structure of tail risk. Tests show that the tail exponent is approximately equal for different credit default swap maturities. This flat term structure of tail risk implies that no maturity dominates the power-law exponent for the aggregated time series. While the implied probability of default increases with maturity, short- and long-dated contracts are similarly exposed to tail risk.

Chapter 5 presents the third dynamic tail risk model, which measures the market-wide tail risk from the cross section of extreme events of individual firms. If the tail risk changes through time, extreme value techniques based on univariate time series (Chapter 1) or aggregated data from a small sample (Chapter 3) are suboptimal in providing reliable tail risk measures. However, if a common process drives tail risk dynamics of firms, the cross section of extreme returns can be used to accurately measure prevailing tail risk in the economy. Tests show that tail risk is highly correlated among sectors (48%) and disjoint sets of firms with a low and high default probability (66%). This high degree of co-movement empirically supports our assumption of common firm-level tail dynamics. Evidence suggests that firms with a high probability of default are similarly exposed to extreme event risk as firms with a lower default probability with no significant difference in the average tail return between both subsets. Therefore, we conclude that firms with low idiosyncratic default risk provide good tail risk hedges if their sensitivity to tail risk is high.

Chapter 6 shows that tail risk has substantial explanatory power for the cross section of expected returns in US corporate credit default swaps. We provide evidence that the cross section of corporate credit default swap returns reflects a premium for tail risk sensitivity. Cross-sectionally, firms that covary highly with tail risk earn average expected annual returns 8.1% higher than credit default swaps with low tail risk covariation. We show that the credit tail risk premium is different from the premiums on market risk, idiosyncratic volatility and coskewness, and robust to controlling for these alternative risk factors. We find that past tail risk sensitivity provides enough divergence in future tail risk sensitivity to justify the tail risk premium between high and low quintile portfolios. We conclude that tail risk is persistent and that protection sellers demand additional compensation for a credit default swap with a high sensitivity to extreme event risk.

In summary, we conclude that tail returns in credit default swap markets are well-characterised by power-law distributions. We provide evidence that tail risk dynamics can influence prices and aggregated returns of sovereign and corporate credit default swaps.

References

Andersen, L. and Andreasen, J. (2000) Jump-diffusion processes: Volatility smile fitting and numerical methods for option pricing *Review of Derivatives Research* **4**(3), pp. 231–262

Ang, A., Chen, J. and Xing, Y. (2006a) Downside risk *The Review of Financial Studies* **19**(4), pp. 1191–1239

Ang, A., Hodrick, R.J., Xing, Y. and Zhang, X. (2006b) The cross-section of volatility and expected returns *The Journal of Finance* **61**(1), pp. 259–299

Atkinson, A.B. and Piketty, T. (2007) *Top incomes over the twentieth century: a contrast between continental european and english-speaking countries* oup Oxford

Axtell, R.L. (2001) Zipf distribution of us firm sizes *science* **293**(5536), pp. 1818–1820

Bai, J. and Wei, S.J. (2012) When is there a strong transfer risk from the sovereigns to the corporates? property rights gaps and cds spreads Technical report National Bureau of Economic Research

Baker, S.R., Bloom, N. and Davis, S.J. (2016) Measuring economic policy uncertainty *The Quarterly Journal of Economics* **131**(4), pp. 1593–1636

Barberis, N. and Huang, M. (2001) Mental accounting, loss aversion, and individual stock returns *the Journal of Finance* **56**(4), pp. 1247–1292

Barberis, N. and Thaler, R. (2003) A survey of behavioral finance *Handbook of the Economics of Finance* **1**, pp. 1053–1128

Barndorff-Nielsen, O.E. and Cox, D.R. (1994) *Inference and asymptotics* Chapman and Hall

Begušić, S., Kostanjčar, Z., Stanley, H.E. and Podobnik, B. (2018) Scaling properties of extreme price fluctuations in bitcoin markets *Physica A: Statistical Mechanics and its Applications* **510**, pp. 400–406

Beirlant, J., Vynckier, P. and Teugels, J.L. (1996) Tail index estimation, pareto quantile plots regression diagnostics *Journal of the American Statistical Association* **91**(436), pp. 1659–1667

Benartzi, S. and Thaler, R.H. (1995) Myopic loss aversion and the equity premium puzzle *The Quarterly Journal of Economics* **110**(1), pp. 73–92

Berndt, A. and Obreja, I. (2010) Decomposing european cds returns *Review of Finance* **14**(2), pp. 189–233

Blanco, R., Brennan, S. and Marsh, I.W. (2005) An empirical analysis of the dynamic relation between investment-grade bonds and credit default swaps *The Journal of Finance* **60**(5), pp. 2255–2281

Bollerslev, T. (1986) Generalized autoregressive conditional heteroskedasticity *Journal of econometrics* **31**(3), pp. 307–327

Bollerslev, T. (1987) A conditionally heteroskedastic time series model for speculative prices and rates of return *The review of economics and statistics* pp. 542–547

Bouchaud, J.P. (2001) Power laws in economics and finance: some ideas from physics

Carhart, M.M. (1997) On persistence in mutual fund performance *The Journal of Finance* **52**(1), pp. 57–82

Chauvet, M. (1998) An econometric characterization of business cycle dynamics with factor structure and regime switching *International Economic Review* pp. 969–996

Chauvet, M. and Piger, J. (2008) A comparison of the real-time performance of business cycle dating methods *Journal of Business & Economic Statistics* **26**(1), pp. 42–49

Clauset, A., Shalizi, C.R. and Newman, M.E. (2009) Power-law distributions in empirical data *SIAM review* **51**(4), pp. 661–703

Clauset, A., Young, M. and Gleditsch, K.S. (2007) On the frequency of severe terrorist events *Journal of Conflict Resolution* **51**(1), pp. 58–87

Cont, R. (2001) Empirical properties of asset returns: Stylized facts and statistical issues *Quantitative Finance* **1**, pp. 223–236

Csorgo, S., Deheuvels, P. and Mason, D. (1985) Kernel estimates of the tail index of a distribution *The Annals of Statistics* pp. 1050–1077

Dacorogna, M.M., Müller, U.A., Pictet, O.V. and de Vries, C.G. (1995) *The distribution of extremal foreign exchange rate returns in extremely large data sets* Tinbergen Institute Netherlands

Dacorogna, M.M. and Pictet, O.V. (1998) Heavy tails in high-frequency financial data *A practical guide to heavy tails: Statistical techniques and applications* pp. 55–78

Danielsson, J., Ergun, L.M., de Haan, L. and de Vries, C.G. (2016) Tail index estimation: Quantile driven threshold selection *Available at SSRN 2717478*

Danielsson, J., de Haan, L., Peng, L. and de Vries, C. (2001) Using a bootstrap method to choose the sample fraction in tail index estimation *Journal of Multivariate Analysis* **2**(76), pp. 226–248

Davis, R. and Resnick, S. (1984) Tail estimates motivated by extreme value theory *The Annals of Statistics* pp. 1467–1487

Dekkers, A.L. and de Haan, L. (1993) Optimal choice of sample fraction *Journal of Multivariate Analysis* **47**, pp. 173–195

Delianedis, G. and Geske, R.L. (2003) Credit risk and risk neutral default probabilities: information about rating migrations and defaults in: *EFA 2003 annual conference paper* 962

Dimitropoulos, X., Krioukov, D., Fomenkov, M., Huffaker, B., Hyun, Y., Riley, G. *et al.* (2007) As relationships: Inference and validation *ACM SIGCOMM Computer Communication Review* **37**(1), pp. 29–40

Dittmar, R.F. (2002) Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns *The Journal of Finance* **57**(1), pp. 369–403

Doyne Farmer, J., Gillemot, L., Lillo, F., Mike, S. and Sen, A. (2004) What really causes large price changes? *Quantitative finance* **4**(4), pp. 383–397

Drees, H. and Kaufmann, E. (1998) Selecting the optimal sample fraction in univariate extreme value estimation *Stochastic Processes and their Applications* **75**(2), pp. 149–172

Engle, R.F. (1982) Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation *Econometrica: Journal of the Econometric Society* pp. 987–1007

Ergun, L.M. (2016) Disaster and fortune risk in asset returns

Ewens, M., Jones, C.M. and Rhodes-Kropf, M. (2013) The price of diversifiable risk in venture capital and private equity *The Review of Financial Studies* **26**(8), pp. 1854–1889

Fama, E.F. (1963) Mandelbrot and the stable paretian hypothesis *The Journal of Business* **36**(4), pp. 420–429

Fama, E.F. and French, K.R. (1993) Common risk factors in the returns on stocks and bonds *Journal of Financial Economics* **33**, pp. 3–56

Fama, E.F. and MacBeth, J.D. (1973) Risk, return, and equilibrium: Empirical tests *Journal of Political Economy* **81**(3), pp. 607–636

Farmer, J.D., Lillo, F. *et al.* (2004) On the origin of power-law tails in price fluctuations *Quantitative Finance* **4**(1), pp. 7–11

Fisher, R.A. (1922) On the mathematical foundations of theoretical statistics *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* **222**(594-604), pp. 309–368

Gabaix, X. (2009) Power laws in economics and finance *Annu. Rev. Econ.* **1**(1), pp. 255–294

Gabaix, X. (2016) Power laws in economics: An introduction *Journal of Economic Perspectives* **30**(1), pp. 185–206

Gabaix, X., Gopikrishnan, P., Plerou, V. and Stanley, E. (2007) A unified econophysics explanation for the power-law exponents of stock market activity *Physica A: Statistical Mechanics and its Applications* **382**(1), pp. 81–88

Gabaix, X., Gopikrishnan, P., Plerou, V. and Stanley, H.E. (2003) A theory of power-law distributions in financial market fluctuations *Nature* **423**(6937), pp. 267–270

Gabaix, X., Gopikrishnan, P., Plerou, V. and Stanley, H.E. (2006) Institutional investors and stock market volatility *The Quarterly Journal of Economics* **121**(2), pp. 461–504

Gabaix, X. and Landier, A. (2008) Why has ceo pay increased so much? *The Quarterly Journal of Economics* **123**(1), pp. 49–100

Galbraith, J.W. and Zernov, S. (2004) Circuit breakers and the tail index of equity returns *Journal of Financial Econometrics* **2**(1), pp. 109–129

Ghashghaie, S., Breymann, W., Peinke, J., Talkner, P. and Dodge, Y. (1996) Turbulent cascades in foreign exchange markets *Nature* **381**(6585), pp. 767–770

Gopikrishnan, P., Meyer, M., Amaral, L.N. and Stanley, H.E. (1998) Inverse cubic law for the distribution of stock price variations *The European Physical Journal B-Condensed Matter and Complex Systems* **3**(2), pp. 139–140

Gopikrishnan, P., Plerou, V., Amaral, L.A.N., Meyer, M. and Stanley, H.E. (1999) Scaling of the distribution of fluctuations of financial market indices *Physical Review E* **60**(5), p. 5305

Gopikrishnan, P., Plerou, V., Liu, Y., Amaral, L.N., Gabaix, X. and Stanley, H.E. (2000) Scaling and correlation in financial time series *Physica A: Statistical Mechanics and its Applications* **287**(3-4), pp. 362–373

Grünwald, P.D. (2007) *The minimum description length principle* MIT press

Gu, G.F., Chen, W. and Zhou, W.X. (2008) Empirical distributions of chinese stock returns at different microscopic timescales *Physica A: Statistical Mechanics and its Applications* **387**(2-3), pp. 495–502

Guillaume, D.M., Dacorogna, M.M., Davé, R.R., Müller, U.A., Olsen, R.B. and Pictet, O.V. (1997) From the bird's eye to the microscope: a survey of new stylized facts of the intra-daily foreign exchange markets *Finance and stochastics* **1**(2), pp. 95–129

de Haan, L. and Resnick, S.I. (1980) A simple asymptotic estimate for the index of a stable distribution *Journal of the Royal Statistical Society: Series B (Methodological)* **42**(1), pp. 83–87

Hall, P. (1982) On some simple estimates of an exponent of regular variation *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 37–42

Hall, P. (1990) Using the bootstrap to estimate mean squared error and select smoothing parameter in nonparametric problems *Journal of multivariate analysis* **32**(2), pp. 177–203

Hall, P., Welsh, A.H. *et al.* (1985) Adaptive estimates of parameters of regular variation *the Annals of Statistics* **13**(1), pp. 331–341

Harvey, C.R. and Siddique, A. (2000) Conditional skewness in asset pricing tests *The Journal of finance* **55**(3), pp. 1263–1295

Hill, B.M. (1975) A simple general approach to inference about the tail of a distribution *The annals of statistics* **3**(5), pp. 1163–1174

Hinloopen, J. and van Marrewijk, C. (2006) Comparative advantage, the rank-size rule, and zipf's law *Available at SSRN 943370*

Jansen, D.W. and De Vries, C.G. (1991) On the frequency of large stock returns: Putting booms and busts into perspective *The review of economics and statistics* pp. 18–24

Jegadeesh, N. and Titman, S. (1993) Returns to buying winners and selling losers: Implications for stock market efficiency *The Journal of Finance* **48**(1), pp. 65–91

Jensen, M.C., Black, F. and Scholes, M.S. (1972) The capital asset pricing model: Some empirical tests

Jessen, H.A. and Mikosch, T. (2006) Regularly varying functions *Publications de L'institut Mathematique* **80**(94), pp. 171–192

Kass, R.E. and Raftery, A.E. (1995) Bayes factors *Journal of the American Statistical Association* **90**(430), pp. 773–795

Kelly, B. (2014) The dynamic power law model *Extremes* **17**(4), pp. 557–583

Kelly, B. and Jiang, H. (2014) Tail risk and asset prices *Review of Financial Studies* **27**(10), pp. 2841–2871

Kolmogorov, A.N. (1933) Sulla determinazione empirica di una legge di distribuzione *Giornale dell'Istituto Italiano degli Attuari* **4**, pp. 83–91

Kraus, A. and Litzenberger, R.H. (1976) Skewness preference and the valuation of risk assets *The Journal of finance* **31**(4), pp. 1085–1100

Krugman, P. (1996) Confronting the mystery of urban hierarchy *Journal of the Japanese and International economies* **10**(4), pp. 399–418

Kuester, K., Mittnik, S. and Paolella, M.S. (2006) Value-at-risk prediction: A comparison of alternative strategies *Journal of Financial Econometrics* **4**(1), pp. 53–89

Kyle, A.S. and Obizhaeva, A.A. (2016) Market microstructure invariance: Empirical hypotheses *Econometrica* **84**(4), pp. 1345–1404

LeBaron, B. and Samanta, R. (2005) Extreme value theory and fat tails in equity markets *Available at SSRN 873656*

Lee, K.E. and Lee, J.W. (2004) Scaling properites of price changes for korean stock indices *arXiv preprint cond-mat/0407418*

Lux, T. (1996) The stable paretian hypothesis and the frequency of large returns: an examination of major german stocks *Applied Financial Economics* **6**(6), pp. 463–475

Makowiec, D. and Gnacinski, P. (2000) Fluctuations of wig-the index of warsaw stock exchange. preliminary studies *arXiv preprint cond-mat/0010190*

Malkiel, B.G. and Xu, Y. (2002) Idiosyncratic risk and security returns *University of Texas at Dallas (November 2002)*

Mandelbrot, B. (1963) The variation of certain speculative prices *The Journal of Business* **36**(4), pp. 394–419

Mantegna, R.N. and Stanley, H.E. (1996) Turbulence and financial markets *Nature* **383**(6601), pp. 587–588

Mason, D.M. (1982) Laws of large numbers for sums of extreme values *The Annals of Probability* pp. 754–764

McCulloch, J.H. (1996) 13 financial applications of stable distributions *Handbook of statistics* **14**, pp. 393–425

McNeil, A.J. and Frey, R. (2000) Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach *Journal of Empirical Finance* **7**(3-4), pp. 271–300

Merton, R. (1987) A simple model of capital market equilibrium with incomplete information *Journal of Finance* **42**(3), pp. 483–510

Merton, R.C. *et al.* (1985) On the current state of the stock market rationality hypothesis

Mu, G.H., Zhou, W.X. *et al.* (2010) Tests of nonuniversality of the stock return distributions in an emerging market *Physical Review E* **82**(6), p. 066103

Muniruzzaman, A. (1957) On measures of location and dispersion and tests of hypotheses in a pare to population *Calcutta Statistical Association Bulletin* **7**(3), pp. 115–123

Pan, J. and Singleton, K.J. (2008) Default and recovery implicit in the term structure of sovereign cds spreads *The Journal of Finance* **63**(5), pp. 2345–2384

Pan, R.K. and Sinha, S. (2007) Self-organization of price fluctuation distribution in evolving markets *EPL (Europhysics Letters)* **77**(5), p. 58004

Pástor, L. and Stambaugh, R.F. (2003) Liquidity risk and expected stock returns *Journal of Political economy* **111**(3), pp. 642–685

Pickands, J. *et al.* (1975) Statistical inference using extreme order statistics *the Annals of Statistics* **3**(1), pp. 119–131

Pitman, E.J. (2018) *Some Basic Theory for Statistical Inference: Monographs on Applied Probability and Statistics* Chapman and Hall/CRC

Plerou, V., Gopikrishnan, P., Amaral, L.A.N., Meyer, M. and Stanley, H.E. (1999) Scaling of the distribution of price fluctuations of individual companies *Physical Review E* **60**(6), p. 6519

Plerou, V., Gopikrishnan, P., Gabaix, X., Stanley, H.E. *et al.* (2004) On the origin of power-law fluctuations in stock prices *Quantitative Finance* **4**(1), pp. C11–C15

Plerou, V., Gopikrishnan, P. and Stanley, H.E. (2005) Quantifying fluctuations in market liquidity: Analysis of the bid-ask spread *Physical Review E* **71**(4), p. 046131

Preis, T., Schneider, J.J. and Stanley, H.E. (2011) Switching processes in financial markets *Proceedings of the National Academy of Sciences* **108**(19), pp. 7674–7678

Press, W.H., Teukolsky, S.A., Flannery, B.P. and Vetterling, W.T. (1992) *Numerical Recipes in C: The Art of Scientific Computing* Cambridge university press

Quintos, C., Fan, Z. and Phillips, P.C. (2001) Structural change tests in tail behaviour and the asian crisis *The Review of Economic Studies* **68**(3), pp. 633–663

Rachev, S. and Mittnik, S. (2000) Stable paretian models in finance

Resnick, S. and Stărică, C. (1997) Smoothing the hill estimator *Advances in Applied Probability* **29**(1), pp. 271–293

Schönbucher, P.J. (2003) *Credit derivatives pricing models: models, pricing and implementation* John Wiley & Sons

Scott, R.C. and Horvath, P.A. (1980) On the direction of preference for moments of higher order than the variance *The Journal of finance* **35**(4), pp. 915–919

Shiller, R.J. (2015) *Irrational exuberance* Princeton university press

Shleifer, A. (2000) *Inefficient markets: An introduction to behavioural finance* OUP Oxford

Smirnov, N. (1948) Table for estimating the goodness of fit of empirical distributions *The annals of mathematical statistics* **19**(2), pp. 279–281

Stoev, S.A., Michailidis, G. and Taqqu, M.S. (2006) Estimating heavy-tail exponents through max self-similarity *arXiv preprint math/0609163*

Stone, M. (1974) Cross-validatory choice and assessment of statistical predictions *Journal of the Royal Statistical Society: Series B (Methodological)* **36**(2), pp. 111–133

Stoyanov, S., Loh, L. and Fabozzi, F.J. (2017) How fat are the tails of equity market indices? *International Journal of Finance & Economics* **22**(3), pp. 181–200

Vuong, Q.H. (1989) Likelihood ratio tests for model selection and non-nested hypotheses *Econometrica: Journal of the Econometric Society* pp. 307–333

Wagner, N. (2003) Estimating financial risk under time-varying extremal return behavior *OR Spectrum* **25**(3), pp. 317–328

Wang, B.H. and Hui, P.M. (2001) The distribution and scaling of fluctuations for hang seng index in hong kong stock market *The European Physical Journal B-Condensed Matter and Complex Systems* **20**(4), pp. 573–579

Wasserman, L. (2013) *All of statistics: a concise course in statistical inference* Springer Science & Business Media

Welch, I. and Goyal, A. (2008) A comprehensive look at the empirical performance of equity premium prediction *The Review of Financial Studies* **21**(4), pp. 1455–1508

Werner, T. and Upper, C. (2004) Time variation in the tail behavior of bund future returns *Journal of Futures Markets: Futures, Options, and Other Derivative Products* **24**(4), pp. 387–398

Yan, C., Zhang, J., Zhang, Y. and Tang, Y. (2005) Power–law properties of chinese stock market *Physica A: Statistical Mechanics and its Applications* **353**, pp. 425–432

Zhang, J., Zhang, Y. and Kleinert, H. (2007) Power tails of index distributions in chinese stock market *Physica A: Statistical Mechanics and its Applications* **377**(1), pp. 166–172

Zhao, L. (2010) A model of limit-order book dynamics and a consistent estimation procedure Ph.D. Thesis Carnegie Mellon University