Correlated default and parameter risk

Von der Wirtschaftswissenschaftlichen Fakultät der
Gottfried Wilhelm Leibniz Universität Hannover
zur Erlangung des akademischen Grades

Doktor der Wirtschaftswissenschaften
— Doctor rerum politicarum —

genehmigte Dissertation
von

Diplom–Ökonom Martin Schmelzle
geboren am 25. August 1977 in Hannover

2018

Referent: Professor Dr. Maik Dierkes
Korreferent: Professor Dr. Daniel Rösch (Universität Regensburg)

Tag der Promotion: 19. September 2018

To my family
You are awesome

Abstract

Financial markets are larger, more diverse, and more dynamic than ever before. A deep understanding of the constantly evolving markets is key to financial institutions seeking new business opportunities, regulatory authorities tasked with supervising them, and investors trying to profit from them. This cumulative dissertation presents three contributions to two crucial aspects of modern financial markets. The first aspect is directed towards one of the most influential innovations in recent financial history—markets for correlated default risk. Credit risks are disaggregated into separate tradable units that subsequently are repackaged and redistributed according to differing credit quality and investors demand. As the financial crisis unfolded, these new types of securities revealed a multitude of risks that sellers and buyers of credit risks had apparently not been aware of—at least not to full extent. The second aspect is geared towards the management of the fundamental and oftentimes intertwined risks driving asset prices. Spectacular failures of financial institutions remind the financial profession to keep up with prudential and robust risk management systems to serve the needs of individual institutions and to ensure a resilient financial system.

The first chapter starts from the premise that prices for credit contingent contracts in markets for correlated default risk are attributable to idiosyncratic, sectoral and systemic risk components. In contrast to a sophisticated three-jump model, an alternative and much simpler calibration approach is presented which yields virtually identical risk-neutral stress event intensities for the three risk factors. The main contribution is an extensive econometric analysis of the three latent intensity processes.

The second chapter develops a novel methodology to extract model-free moments from the default loss distribution contained in markets for correlated default risk. Key to this approach is the 'spanning' argument in that future payoffs may be synthesized, or 'spanned', from a portfolio of plain vanilla options across the state space. While this argument had been successfully applied to equity, bond, foreign exchange, interest rate swap or commodity markets, work on credit markets seem nonexistent. The main con-

tribution is to fill this void and to provide a first extensive empirical analysis of model-free loss moments with applications.

The third chapter is concerned with the hedging of possible estimation errors in the parameters of risk management models—termed parameter risk. A growing body of literature ascertains that sampled risk measures are, on average, underestimating true risks. That is, risk measures themselves are subject to estimation risk which may threaten the solvency of individual institutions. One remedy is to provide additional risk capital to buffer unfortunate consequences resulting from parameter risk. This, however, comes with social costs since risk capital required to cushion parameter risk is not available for future investments or lending activities. The main contribution is to develop a framework that enables financial institutions to hedge this parameter risk instead. By virtue of diversification, the pooling of hedging activities potentially reduces systemic risks and thus benefits the whole society.

The contributions in this dissertation seek to deepen the understanding of correlation markets and to advance the risk management of financial institutions.

Keywords: stress event intensities, model-free moments of default loss distributions, hedging parameter risk

Zusammenfassung

Finanzmärkte sind größer, vielfältiger und dynamischer als jemals zuvor. Ein tiefes Verständnis der sich fortlaufend entwickelnden Märkte ist wesentlich für Finanzinstitutionen auf der Suche nach neuen Geschäftsfeldern, Aufsichtsbehörden mit der Aufgabe diese zu überwachen und Investoren welche von ihnen profitieren möchten. Diese kumulative Dissertation besteht aus drei Beiträgen zu zwei zentralen Aspekten moderner Finanzmärkte. Der erste Aspekt bezieht sich auf eine der einflussreichsten Innovationen der jüngeren Geschichte—Märkte für korrelierte Ausfallrisiken. Kreditrisiken werden in einzelne handelbare Einheiten zerlegt und in der Folge neu zusammengesetzt und ihrer Kreditqualität und Investorennachfrage entsprechend weiterveräußert. Als die Finanzkrise ihren Lauf Lauf genommen hat, zeigte sich zusehends, dass diese neuen Wertpapiere eine Vielzahl von Risiken beinhalten, deren sich die Verkäufer und Käufer nicht im Klaren waren—zumindest nicht in vollem Umfang. Der zweite Aspekt zielt auf das Management der fundamentalen und oft miteinander verflochtenen Risiken, welche die Preise von Vermögenswerten beeinflussen. Spektakuläre Zusammenbrüche von Finanzinstitutionen erinnern den Berufsstand der Finanzwirte daran, vorsichtigkeitsorientierte und robuste Risikomanagementsysteme zu entwickeln, welche den Bedürfnissen der Finanzinstitute dienen und ein widerstandsfähiges Finanzsystem sicherstellen.

Das erste Kapitel startet von der Prämisse, dass Preise für kreditabhängige Kontrakte auf Märkten für korrelierte Ausfallrisiken idiosynkratischen, sektoralen und systemischen Risikokomponenten zugeordnet werden können. Im Gegensatz zu einem hochentwickelten Dreisprungmodell, wird ein alternativer und wesentlich einfacherer Kalibrierungsansatz vorgeschlagen, welcher praktisch identische risikoneutrale Intensitäten von Stressereignissen für die drei Risikofaktoren liefert. Der Hauptbeitrag besteht in einer umfassenden ökonometrischen Analyse von den drei latenten Intensitätsprozessen.

Das zweite Kapitel entwickelt eine neue Methodologie, um modellfreie Momente von Kreditausfallverlustverteilungen zu extrahieren, welche den Märkten von korrelierten Ausfallrisiken innewohnen. Der Schlüssel zu diesem

Ansatz liegt in dem „spanning“ Argument mit dem zukünftige Auszahlungsprofile mithilfe eines Portfolios von einfachen Optionen über den Zustandsraum synthetisiert, oder „aufgespannt“, werden können. Während dieses Argument erfolgreich auf Aktien-, Anleihen-, Fremdwährungs-, Zinsswap- oder Rohstoffmärkte angewendet werden konnte, so scheinen Kreditmärkte bislang nicht betrachtet worden zu sein. Der Hauptbeitrag liegt darin, dieses Vakuum zu füllen sowie eine erste umfangreiche empirische Analyse von modellfreien Momenten samt einigen Anwendungen durchzuführen.

Das dritte Kapitel befasst sich mit der Absicherung von möglichen Schätzfehlern in den Parametern von Risikomanagementmodellen—genannt Parameterrisiko. Zunehmend mehr Literaturbeiträge stellen fest, dass auf Stichproben basierende Risikomaße, im Mittel, wahre Risiken unterschätzen. Das heißt, die Risikomaße selbst sind mit Schätzrisiken behaftet, was letztlich die Zahlungsfähigkeit von einzelnen Finanzinstituten bedrohen könnte. Eine Gegenmaßnahme besteht darin, zusätzliches Risikokapital bereitzustellen, um unerwünschte Folgen aufgrund von Parameterrisiken abfedern zu können. Dies könnte jedoch zu Wohlfahrtsverlusten führen, da Risikokapital nun als Puffer für Parameterrisiken benötigt wird und so nicht für zukünftige Investitionen oder Kreditvergabe verfügbar ist. Der Hauptbeitrag liegt in der Entwicklung eines Rahmenwerks, welches es Finanzinstitutionen ermöglicht diese Parameterrisiken zu einem Bruchteil der sonst benötigten Kosten abzusichern. Aufgrund von möglichen Diversifikationseffekten kann ein Bündeln von Absicherungskontrakten ferner zu einer Reduzierung von systemischen Risiken führen und so der gesamten Gesellschaft förderlich sein.

Die Beiträge in dieser Dissertation beabsichtigen das Verständnis für Korrelationsmärkte zu vertiefen und das Risikomanagement von Finanzinstitutionen weiter voranzubringen.

Schlagwörter: Intensität von Stressereignissen, modellfreie Momente von Verlustverteilungen, Absichern von Parameterrisiko

Contents

List of Figures

List of Tables

Executive summary

The past decades have witnessed an extraordinary flow of new financial instruments and innovations. New security designs, advancing information technology and scientific breakthroughs in finance theory paved the way to profound changes in the structure of financial markets and institutions (Merton, 1995). At the core of this astounding development have been sophisticated and Nobel prize winning mathematical analyses of derivative contracts (MacKenzie, 2006). In essence, derivative contracts are probabilistic bets contingent on some future state of the economy or specific event. Accordingly, the value of a contingent contract is 'derived' from what happens to some underlying reference asset or other contractual state of affairs.

The markets in derivatives such as options and futures were tiny until the 1970s, when economic, regulatory and technological conditions laid the basis for an astounding growth in derivatives markets. By mid 2017, the notional amount of outstanding contracts of the global market for over-the-counter (OTC) derivatives were totaling US$ 511 trillion (BIS Quarterly Review, March 2018). To put this number into perspective, the market capitalization of all equity in the world just crossed the US$ 75 trillion mark (Bloomberg world exchange market capitalization index) and the global debt securities markets for corporates added to US$ 55 trillion (BIS Quarterly Review, March 2018).

Derivatives present a cornucopia of new investment opportunities and diverse means of managing risks for financial institutions and market participants. However, derivative contracts are Janus faced. On the one side, they serve hedging purposes which reduces risk and contributes to social welfare. On the other side, when used for speculation, they potentially create new risks which would not exist otherwise, thus reducing social welfare (Stulz, 2004). Not surprisingly, derivative securities played an ambiguous role in the recent financial crisis (Acharya et al., 2009; Stout, 2011). In particular, derivatives whose payoffs are linked to the credit quality or default of one or more debt securities came on the radar of academics, financial researchers and policy makers (Stulz, 2010). There are basically two major types of credit derivatives: credit default swaps and collateralized debt obligations (Partnoy and Skeel, 2007; O'Kane, 2008). A credit default swap (CDS) is a contract where private parties bet on a specific debt issuer's credit event like failure to pay or restructuring. A collateralized debt obligation (CDO) consists of a pool of underlying debt securities whose capital structure is sliced into tranches and redistributed with differing credit qualities. In the common case of synthetic CDOs, these instruments are actually backed by CDS instead of the underlying reference bonds. The value of positions in credit derivatives is highly dependent on the linkages across the reference names issuing debt. Hence markets trading credit derivatives are commonly referred to as markets for correlated default risk (Collin-Dufresne, 2009).

The notion of correlated defaults in the study of financial markets trading credit portfolio derivatives cannot be overemphasized. Moreover, the extremely fast development of active markets in indexed and tranched credit derivatives allows practitioners, academics and policy makers to inquire into the structure of correlated default risk to make informed decisions and draw valuable conclusions. Default correlation and likewise clustering and risk contagion, are commonly attributed to the joint exposure of packaged reference entities to observable and non-observable systematic risk factors (Azizpour et al., 2018). What exactly drive correlated defaults is, however, still an open issue within the finance profession. Correlation markets convey a lot of information, and as recent history shows, provide an important tool to elicit default and correlation expectations from market participants. For instance, the risk of losses confined to bad economic states had been severely underestimated by investors (Coval et al., 2009). Structured finance products like CDOs are highly sensitive to assumptions regarding probabilities of default, recovery rates and the correlation of defaults. In reverse, market quotes of these kinds of products allow financial researchers to gather information in how market participants value correlated default risk along the capital structure and value systemic risks. In this dissertation, the first two chapters are devoted to markets of correlated default risk.

Along with the profound changes in the structure of financial markets and financial institutions comes an enormous increase in the ability to spread and manage risks (Rajan, 2006). These revolutionary changes make it easier than ever before for investment management firms, corporations and financial institutions to cater specific risk profiles by either stripping undesired risks or take on more risks that suit their portfolios (Froot et al., 1993). Risk management is a rich topic and of tremendous importance to investors, financial managers, regulators and policy makers . Typically, risk management focuses on issues pertaining to risk capital, deposit insurance or the performance of financial institutions (Merton, 1995). On a firm level, risk management is also closely tied to optimal capital budgeting and capital structure decisions

(Froot and Stein, 1998). A number of notorious incidents involving Barings Brothers and Co. Ltd., Orange County, California or Metallgesellschaft AG's US oil trading subsidiary, MG Refining & Marketing Inc., and the more recent collapses of Bear Stearns Companies, Inc., Lehman Brothers Holdings Inc., or Washington Mutual, Inc., make it clear that proper risk management systems are of utmost importance for the well-being of financial institutions and to avoid significant social welfare losses and severe economic distress. The finance profession also has learned some lessons from the 1998 failure of Long-Term Capital Management (LTCM) which is said to have nearly blown up the world's financial system (Jorion, 2000; Stulz, 2008).

A key role in risk management plays risk capital, i.e., the equity investment that backs obligations to liability holders and maintains the firm's credit quality (Merton and Perold, 2005; Erel et al., 2015). The more risk capital a firm provisions, the more remote are financial distress and the likelihood of bankruptcy. However, risk capital is costly since this capital is not available for further investment or lending activities and consequently might hinder social welfare. For the most part, the literature on risk management and risk capital is centered around value-at-risk (VaR) like measures of risk (Kupiec, 2001). As a measure of downside risk, the VaR maps a real number to the downside of the profit and loss random variable of a portfolio at a given confidence level. Alternative risk measures like conditional value-at-risk, expected shortfall or tail conditional expectations have been developed to address some shortcomings of the VaR, among others, they are capable to quantify the losses beyond a given threshold or confidence level that might be encountered in the tail (Acerbi and Tasche, 2002; Rockafellar and Uryasev, 2002). Effective risk management encompasses many key steps and the definition of an appropriate measure of risk is only but one of them. Applying risk management tools requires the accurate measurement of risk, i.e., the involved risks have to be estimated from available data. However, estimation uncertainties arise through limited and noisy data, biased estimators, or calibration problems with multiple local minima. This estimation risk might translate to biased risk measures and hence to biased risk capital (Jorion, 1996). In an illuminating case study Marshall and Siegel (1997) consider the divergence of VaR estimates from a number of risk management consulting firms, all asked to employ the J.P. Morgan RiskMetrics™ model. Despite the risk management system's vendors all use exactly the same model (i.e., there is no model uncertainty) and are provided with the exact same input data, there is wide variation in the VaR estimates from the four vendors completing the entire reference portfolio ranging from US$ 3.85 million to US$ 6.14 million. Thus, in last consequence, an estimation error might lead to a situation where a firm does not hold sufficient capital reserves to cushion against an adverse economic event or operational incident. The third chapter of this dissertation is exploring possibilities to hedge parameter risk due to possible estimation errors in risk management applications. A brief synopsis of each chapter follows.

Chapter 1: A simple econometric approach for modeling stress event intensities (joint work with Rainer Jobst, Daniel Rösch and Harald Scheule)
The composition of credit contingent contracts in markets for correlated default risk is an important issue in finance. Previous research suggests that the spreads of a credit default swap index are attributable to idiosyncratic, sectoral and global shocks in the economy. Using a sophisticated three-jump model Longstaff and Rajan (2008) are able to infer the risk-neutral intensities of the three latent risk factors. These stress event intensities reveal market expectations about the stochastic waiting times until the occurrence of such events. This kind of information is crucial to risk management purposes like scenario based stress testing activities or the modeling of tail risks. They also provide daily available information to policy makers and regulatory authorities regarding the likelihood of economic downturns.

In this study, we develop an alternative and computationally fast 'back of the envelope' calibration approach to these stress event intensities. We show that the resulting stress event intensities are hardly distinguishable from the fully calibrated three-jump model and indicate average risk-neutral waiting times of 0.9 for idiosyncratic, 30.5 for sectoral, and 217 years for systemic default events.

This is the first study in the literature that applies advanced econometric techniques to the three latent risk factor processes. For one, we use a vector autoregressive approach with exogenous risk factors (VARX) and multivariate generalized conditional heteroskedastic error terms (MGARCH). Contrary to the independence assumption in the three-jump model, our flexible specification enables to control for first and second order dependencies in the moments of the intensities. To this end, we first estimate the conditional mean of the daily intensity innovations. Next, the time variation of the residuals is modeled via dynamic conditional correlations (DCC) with a multivariate t-distribution to accommodate for fat-tailed residuals. Our findings identify a stock market index, a market volatility index, a market skewness index, and treasury yields as the explaining economic risk factors of the latent stress event intensities.

Chapter 2: Expectations in markets for correlated default risk (joint work with Rainer Jobst, Daniel Rösch and Harald Scheule) The majority of prevailing studies on correlated default risk is relying on certain parametric assumptions about the stochastic default or loss process. However, this exposes the financial analyst to substantial model and parameter risk, since the true stochastic nature of defaults is not observable. Fortunately, though, using the path breaking work on 'spanning' future terminal payoffs via a continuum of plain vanilla options across the state space, we are able to extract risk-neutral moments of the default loss distribution in a model-independent way. Simple arbitrage conditions in the credit contingent contracts reveal the latent martingale pricing measure. This enables us to unlock the information content hidden in correlation markets free from model misspecification and calibration risk.

This study is first in applying the spanning argument to credit markets and thus closes an important gap in the literature on model-free approaches to derivatives markets. We apply this novel methodology to an extensive data set of credit portfolio derivatives covering calm periods and times of severe economic distress. The model-free estimates of the default loss distribution disclose a number of revealing insights into the risk assessment of correlated default risk from market participants. The implied expectations on aggregate losses in the credit portfolio have been relatively low before the crisis with values around 2%. During the heights of the financial crisis these values peak up to 13% and nearly revert to pre-crisis levels when the crisis settles down. This is in sharp contrast to the higher order moments which document persistent changes in the pricing of correlated default risk. For instance, the loss distribution widens dramatically during the crisis and after the crisis the loss variance remains on a significantly elevated level compared to pre-crisis levels. This lasting change in the risk assessment is also mirrored by the time evolution of the skewness and kurtosis of the loss distribution.

The model-free estimation of expectations from correlation markets offers a wide application spectrum and provides numerous possibilities for empirical inquiries. First applications include the inference of the total risk-neutral default loss distribution leveraging the partial information contained in the model-free moments. Second, we are able to obtain model-free estimates for the implicit risk-neutral default correlation by equating the index variance to the constituents of the reference portfolio. Third, we explore the driving forces for the daily innovations of the default loss moments and identify a number of observable risk factors from the equity, credit, and fixed income markets.

Chapter 3: Hedging parameter risk (joint work with Arndt Claußen and Daniel Rösch) The precise assessment of financial risk with risk measures such as value-at-risk or expected shortfall is of pivotal importance for successful risk management. Unfortunately, though, risk measures themselves are subject to estimation risk. If estimated from sample data they are essentially estimates of risk—not known with certainty. Treating an estimate as if it were the true risk measure typically yields to an underestimation of the true risk. This predicament triggered an expanding body of literature which, in essence, provides means to add some 'conservatism' to uncertain estimates or provision supplemental capital buffers as safety factors for estimation risk.

Ignorance of possible errors in the estimation of risk model parameters—termed parameter risk—comes with important consequences. However, even if financial institutions are fully aware of the possibly harmful impact of uncertain estimates, they may be reluctant to the idea of providing additional and costly risk capital. This study marks the first attempt in the literature to develop a framework that allows financial institutions to hedge parameter risk in place of fully provisioning for 'conservative' risk model parameters. This framework enables the involved contract parties to uniquely determine the fair pricing for parameter risk protection. The underlying idea is reminiscent to a swap deal. The protection buyer pays an insurance premium upfront in exchange for a specified payoff contingent on the future loss of the protected portfolio. The range of parameter risk protection is confined by the statistical point estimate of the used risk measure and some additional conservatism. Instead of fully provisioning for this difference, the protection premium is derived from the statistical expectation of this difference which is only a tiny fraction of the full amount. The additional conditional risk capital thus provides an enlarged buffer against unfavorable events thereby greatly reducing potential negative effects on financial institutions' behavior. An empirical application of the framework to credit risk modeling reveals that, subject to contract type and rating grade, a financial institution may pay protection premiums for about 460 to 1430 years in comparison to a one time provision of additional risk capital to the difference between the VaR using the estimates and the VaR for the conservative parameters.

The well-being of individual institutions is also of great concern for prudential regulatory authorities and policy makers. Now, if one contract party is selling parameter risk protection to more than one counterparty it may profit from arising beneficial diversification effects. That is when a number of protection buyers are pooled the overall risk of the portfolio is potentially smaller than the sum of the individual risks. In light of recent financial turmoils, enhancing the resilience of the financial system as a whole is a most important topic on the agenda of international supervision. The inclusion of a macro perspective thus helps to measure and control systemic risks and cope with the well-being of individual institutions at the same time.

Each of the three chapters add original contributions to the financial literature. Needless to say, each of the projects leave room for additional analyses and further research as they had to be declared 'finished' at some point due to naturally occurring limits in space and time. Though the individual parts consider unique and seemingly unrelated topics, they might be fruitfully combined in future work. For instance, one might derive model-free idiosyncratic, sectoral and global default risk components in dividing the model-free total loss distribution into three segments along the capital structure. These components might then further be used to infer model-free risk-neutral idiosyncratic, sectoral and systemic default correlations. Next steps include the connection of these risk-neutral entities to the modeling of real-world phenomena to learn new aspects about pricing kernels and the risk perceptions of market participants. Furthermore, this perspective might deliver new insights which in turn help deepening the understanding needed for the successful hedging of parameter risk.

In conclusion, this dissertation provides genuinely new theoretical advancements and revealing empirical insights that might be fruitful to leave old, beaten paths and venture forward to explore new ones within the fascinating and ever changing field of modern finance.

All is flux, nothing stays still.
—Heraclitus (535–475)

Chapter 1

A simple econometric approach for modeling stress event intensities

The content of this chapter is published as:

Jobst, Rainer, Daniel Rösch, Harald Scheule and Martin Schmelzle (2015), 'A simple econometric approach for modeling stress event intensities', *The Journal of Futures Markets* **35**(4), 300–320. DOI: 10.1002/fut.21695

Abstract

This paper introduces a simple, non-parametric way of inferring risk-neutral credit stress event intensities for idiosyncratic, sectoral, and global shocks contained in market credit spreads. We provide an econometric analysis of the implied latent stress event dynamics. A vector autoregressive regression model with exogenous variables finds that these intensities can be related to an observable stock market index, the market volatility, the volatility skew, and treasury yields.

Keywords: stress event intensities, portfolio credit derivatives, systemic risk

JEL: C51, G01, G28

Online available at: `https://doi.org/10.1002/fut.21695`

Chapter 2

Expectations in markets for correlated default risk

The content of this chapter refers to the working paper:

Jobst, Rainer, Daniel Rösch, Harald Scheule and Martin Schmelzle (2018), 'Expectations in markets for correlated default risk', Working Paper, Universität Regensburg and University of Technology Sydney.

Abstract

We develop a new methodology to derive model-free risk-neutral expectations of the latent credit portfolio loss distribution embedded in markets for correlated default risk. We document persistent changes in the higher order risk assessment of aggregate debt in the aftermath of the financial crisis. The distributional moments are used (i) to infer the total risk-neutral default loss distribution, (ii) to discuss the implicit default correlation structure, and (iii) to regress these on common risk factors from the equity, credit and fixed income markets.

Keywords: default loss distribution, portfolio credit derivatives, default correlation, model-free expectations

JEL: G13, G32, C14

2.1 Introduction

Increased trading activity in credit derivative markets has been paralleled by an increasing academic interest to gain important insights about potential sources of expected correlated default risk, relying on certain parametric assumptions about the underlying stochastic default or loss process. The true underlying stochastic nature of defaults is not known with certainty, which may lead to model misspecification, model bias and parameter risk. Originating from the options pricing literature, an active area of research is devoted to the model-free extraction of information contained in the prices of derivatives. These activities have subsequently been expanded to other derivatives markets. Little attention has been paid to the market segment of credit risk. The main contribution of this paper is to identify expectations of the marginal densities of the stochastic default loss process without imposing any distributional assumptions. Simple conditions of no arbitrage in contingent claims written on the credit portfolio loss allow us to infer the latent martingale pricing measure. Doing so, our novel methodology fills an important gap in the literature and enables the model-independent inference of the total risk-neutral loss distribution, implied default correlations, and the identification of economic risk drivers of the default loss moments.

Using daily data on the tranched and untranched CDX North American Investment Grade Index (CDX.NA.IG, or CDX for short) from September 2005 to February 2016, we provide model-free moment estimates for the underlying latent default loss distribution. The market for CDX index products is by far the biggest for structured corporate debt and correlation dependent credit derivatives in the world. According to the CFTC, in the first quarter of 2016, the overall gross notional amount outstanding for the CDX.NA.IG indices was around $1,203.3 billion ($309.7 billion) for the untranched (tranched) CDX indices, thus the tranched index markets make up more than 20% of the total CDX index markets.[1] The swap reports further reveal some insights about the composition of major market participants. The untranched CDX market is dominated by the reporting swap dealers with a share of roughly 72% of the total gross notional outstanding, the remaining 28% may be attributed to buyside firms. On the tranched CDX market, swap dealers play an even more dominant role, accounting for roughly 89% of the total market share. Sophisticated investors make extensive use of CDS index products for a variety of reasons. For example, more conservative investors use them as part of yield

[1]To put the numbers into perspective, the US corporate bond market is at $8,352.5 billion outstanding notional, with an average daily trading volume of around $32.9 billion ($20.1 IG and $12.8 HY) in the first quarter of 2016, see `https://www.sifma.org/`. The North American indexed (untranched) and tranched CDS index market is at $1,896.8 billion outstanding notional, with an average daily trading volume of around $71.7 billion ($53.9 IG and $17.8 HY), see `http://www.cftc.gov/MarketReports/SwapsReports/`. We observe much smaller risk outstanding and much higher liquidity (i.e., dollar trading volumes) in the CDS index market in comparison to the corporate bond market. Credit default swap (CDS) indices are now the most liquid instruments in the credit markets.

enhancement strategies. Swap dealers or hedge funds engage in correlation trades, take on more idiosyncratic risks or exploit cross market linkages and basis trades. In particular, tranched CDS index markets allow investors to take fine-grained positional views on the distribution of default losses in the indexed reference pool.

In order to imply the risk-neutral distributional expectations up to arbitrary order from CDS index markets, we describe in detail the methodology in linking today's price of power loss contracts with payoff at a future terminal date to a continuum of contingent claims, written on the portfolio loss along the capital structure of the credit index maturing at the same future date—the expected tranche losses. These expected tranche losses are inferred under no arbitrage conditions from a set of CDS index and tranche quotes. The risk-neutral moments then follow by an integration over the expected tranche losses, using a shape preserving interpolation function, thus ensuring absence of arbitrage from zero to the maximum possible loss of the credit portfolio. Given this procedure, we calculate the first four centered moments of the latent default loss distribution of the CDX credit portfolio, i.e., the expected loss, the default loss variance, the skewness and the kurtosis.

Our empirical analysis of the CDX credit index discloses a number of distinctive features. We find that the market implied expectations on aggregate losses in the underlying reference pool have been relatively low pre-crisis with values around 2%. This market perception changes drastically with the outbreak of the financial crisis where the expected losses reach peak values up to 13%. In the aftermath of the crisis, the expected losses revert to levels slightly above the values recorded before the crisis. The loss variance is highly correlated to the expected loss and exhibits roughly the same time series behavior. Pre-crisis the loss distribution is quite narrow and widens dramatically during the crisis. Contrary to the expected loss, however, we document a persistent change in the risk assessment of correlated defaults after the crisis. The loss variance remains on a significantly elevated level than what have been observed pre-crisis. The skewness and kurtosis are both highly negatively correlated to the expected loss and the default variance. Initially, the skewness and the kurtosis seem quite high which is by construction to accommodate the extremely low expected losses and narrow loss distributions on a bounded domain. Both moments exhibit a sustained decline with the onset of the crisis. Moreover, the skewness and kurtosis show almost no indications of reversion to pre-crisis levels. This is a very distinctive feature and provides further evidence of the persistent changes in the risk perception of credit market participants.

Given the historical evolution of these multi-name credit moments over the sample period, covering more than ten years, our new methodology provides contributions for three distinct applications with regard to the structure and the economic behavior of correlated default risk over time. Therefore, we wish to be able to unfold the valuable embedded information and focus on three initial questions. First, we provide risk-neutral loss distributions. The payoffs of the power loss contracts depend on the cumulative loss due to defaults in the credit index, thus the expectation of the market participants regarding future losses directly translate into the shape and tail behavior of the underlying loss distribution. Unlike for calibrated credit portfolio models, the total loss distribution cannot be readily calculated from distributional moments alone without assuming some additional structure.[2] Here, we infer the default loss distributions over our complete sample period by borrowing well established techniques from the theory of orthogonal polynomials. Specifically, we series expand the loss distribution using the Jacobi orthogonal polynomials with corresponding beta distribution as basis weighting function. To the best of our knowledge, this is the first application of beta series expansions applied to credit portfolio losses. The resulting loss distributions sharpen our intuition about the nature of the correlation dependent expected defaults within the underlying reference pool.

Second, we provide risk-neutral default correlations for the tranched and untranched CDX index markets. Default correlation determines the proportion of total portfolio risk that can be attributed to each segment of the capital structure. In other words, market quotes for the CDX index reveal risk perceptions and risk preferences for different parts of the distribution of default losses in the pool of reference assets. Considering a simple Bernoulli correlation measure, we find that the implied correlations are decidedly higher than what would be expected from real-world applications to portfolio credit risk, which suggests a strong risk premium. This is in close analogy to typical findings in the equity derivatives literature, where—most of the time—implied risk-neutral correlations remain substantially higher than realized correlations, see, e.g., Driessen et al. (2013).

Third, to gain some insight into the main economic drivers for the changes of loss distribution moments, we regress these on a number of observable risk factors from the equity, credit and fixed income markets. More specifically, we run a number of univariate time series regressions for all four daily moments. We find adjusted R^2 of 72.2% for the expected loss, 46.7% for the default loss variance, 64.3% for the skewness, and 62.4% for the kurtosis of the default

[2] There is a considerable body of literature dealing with loss distributions within real-world and risk-neutral applications. The majority of these studies is model based and predominantly within intensity or structural frameworks, compare, e.g., Duffie and Gârleanu (2001), Das et al. (2007), Longstaff and Rajan (2008), Duffie et al. (2009), Azizpour et al. (2011), Chava et al. (2011), Giesecke and Kim (2011), and Giesecke et al. (2015) among many others.

loss distributions across our sample. The CDX–CDS index basis as a proxy for credit market liquidity is the single most important determinant of the daily variation in the default loss moments. The other statistically significant variables are related to the ratio of realized upside to downside variance of the CDX index, a yield spread, a financial sector CDS index, and the CBOE VIX index. The results suggest that liquidity is one key pricing factor in indexed credit derivatives.

Our work is motivated by the growing interest in model-free applications to derivative securities. Specifically, our methodology draws from several contributions in the options pricing literature, where arbitrary future payoffs may be decomposed or replicated by holding a portfolio of properly weighted plain vanilla call and put options. In essence, the replication of some payoff function is a simple consequence of the important findings due to Breeden and Litzenberger (1978) in relating the risk-neutral density function or state-price density to options over a continuum of strike prices along all possible future states of the underlying asset. Pioneering work along these lines mainly geared towards the fair pricing of future realized variance or—equivalently—inferring model-free implied variance, compare, e.g., Neuberger (1994), Carr and Madan (1998), and Britten-Jones and Neuberger (2000). These fundamental developments have paved the way for an expanding diversity of possible applications. The volatility contracts have subsequently been complemented by higher order power contracts to derive skewness and kurtosis swaps (Bakshi et al., 2003). The outstanding role of risk premia is studied in, e.g., Bakshi and Madan (2006), Carr and Wu (2009), Bollerslev et al. (2011), Kozhan et al. (2013), and Bondarenko (2014). Other important contributions to model-free applications include portfolio selection (DeMiguel et al., 2013; Kempf et al., 2015), expected stock returns (Conrad et al., 2013; Bali et al., 2017), implied beta (Chang et al., 2012; Baule et al., 2016), and option implied correlations as well as correlation risk premia (Driessen et al., 2009; Buss and Vilkov, 2012; Buraschi et al., 2014) among many others.

Given the generality of model-free approaches, other derivative market segments have been explored as well. For instance, carry trades, exchange rate predictability and foreign exchange correlation risk premia are studied in Jurek (2014), Della Corte et al. (2016), Chen (2017), and Mueller et al. (2017). Commodity and exchange traded funds (ETF) markets are investigated in, e.g., Prokopczuk et al. (2017) and Tee and Ting (2017). Choi et al. (2017) study US Treasury bond variance risk premia and interest rate swap markets are addressed in Mele et al. (2015). Furthermore, the practical relevance finds expression in dozens of model-free benchmark indices in the industry.[3] This paper contributes to the model-free literature in applying well established payoff decomposition principles to markets of correlated default risk in analogy to the insights first gained in the equity markets. We hereby close a significant gap in the model-free extraction of distributional information expected by the market participants in the derivatives markets.

Our paper is also related to the literature on inquiring into the nature of expected correlated default risk. Commonly, economic drivers for correlated or clustered defaults are mainly attributed to the joint exposure of firms to observable and latent (i.e., frailty) systematic risk factors, and contagion effects in that the default of one firm has a direct impact on the health of related firms. Beginning with the influential article from Longstaff and Rajan (2008), a rapidly growing body of empirical literature is emerging to extract market expectations regarding the nature and degree of default clustering, contagion and correlated default risk across the underlying assets within a reference pool, see Bhansali et al. (2008), Coval et al. (2009), Eckner (2010), Azizpour et al. (2011), Collin-Dufresne et al. (2012), Longstaff and Myers (2014), Seo and Wachter (2016), and Azizpour et al. (2018) among others. This paper differs from this strand of literature in that most of these studies draw their conclusions based on parametric inference. We approach correlated default risk from a different, i.e., model-free, perspective and thus contribute further valuable insights to this area of research.

The remainder of the paper is organized as follows. Section 2.2 develops the theoretical foundation on the model-free derivation of risk-neutral expectations for portfolio loss distributions from markets for correlated default risk. Section 2.3 describes the data and implementation details to calculate the moment estimates from tranched and untranched CDS index quotes. Section 2.4 presents the results of three empirical applications. Section 2.5 provides additional robustness analyses with respect to the calculation of implied moments. Section 2.6 concludes.

[3]For example, the Chicago Board Options Exchange (CBOE) maintains model-free indices for the future expected variance of the S&P 500 (VIX), the VIX of VIX (VVIX), the expected skewness of the S&P 500 (SKEW), US Treasury volatility (TYVIX), interest rate swap volatility (SRVIX), and many more. See `http://www.cboe.com/products/vix-index-volatility/volatility-indexes` for an exhaustive list of volatility related indices for the US. Many major capital centers across the world disseminate similar model-free indices.

2.2 Synthesizing risk-neutral expectations from multi-name credit markets

This section exhibits our approach for synthesizing risk-neutral distributional expectations, derived from markets for correlated default risk. In order to establish these model-free estimators of credit portfolio losses, we first review the essential insights gained from equity markets that future terminal payoffs can be spanned via a portfolio of plain vanilla options. Next, we set out to apply these insights to determine risk-neutral expectations of credit portfolio loss distributions, embedded in market quotes from multi-name credit markets.

2.2.1 Payoffs, power contracts and spanning expectations

Consider a market free of arbitrage in a risky asset viewed from time t to some fixed maturity T with $T \geq t$. Furthermore, assume the existence of a continuum of call and put options deriving their values from the underlying asset, striking prices K and payoff at terminal time T. Following Carr and Madan (1998) and Bakshi and Madan (2000), we know that this market setting permits investors to manufacture any smooth payoff function $f(F_T)$ of the terminal futures price F_T, where $f(.)$ is some arbitrary measurable and twice-differentiable function. Accordingly, any payoff $f(.) \in \mathscr{L}^2$ can be spanned algebraically by taking static positions in the market

$$f(F_T) = f(\kappa) + f'(\kappa)(F_T - \kappa) + \int_0^{\kappa} f''(K)(K - F_T)^+ \, dK + \int_{\kappa}^{\infty} f''(K)(F_T - K)^+ \, dK$$

with some nonnegative separation strike κ.[4] The absence of arbitrage and taking expectations under a corresponding martingale measure reflecting present market beliefs, yields the time t price of the hypothetical claim

$$\mathbb{E}_t\Big[f(F_T)\Big] = f(\kappa) + f'(\kappa)(F_t - \kappa) + \int_0^{\kappa} f''(K) P_{t,T}(K) \, dK + \int_{\kappa}^{\infty} f''(K) C_{t,T}(K) \, dK$$

since from risk-neutral valuation we have $\mathbb{E}_t\Big[(K - F_T)^+\Big] = P_{t,T}(K)$ for put options and $\mathbb{E}_t\Big[(F_T - K)^+\Big] = C_{t,T}(K)$ for call options. Breeden and Litzenberger

[4]Alternatively, we may span arbitrary payoff functions via characteristic functions or moment generating functions as is first shown in Bakshi and Madan (2000) since the two approaches are completely interchangeable. A concrete application is discussed in Appendix 2.A.1 for the case of credit portfolio loss distribution moments.

(1978) show that a unique risk-neutral density function exists, $\varphi_{t,T}(.)$, encoding future uncertainty about the underlying asset that may be recovered from a continuum of plain vanilla option prices

$$\varphi_{t,T}(F_T) = \left.\frac{\partial^2 C_{t,T}(K)}{\partial K^2}\right|_{K=F_T} = \left.\frac{\partial^2 P_{t,T}(K)}{\partial K^2}\right|_{K=F_T}$$

This in turn means that for any payoff being integrable against the risk-neutral density (i.e., has finite expectation $\int_0^\infty \left|f(F_T)\right| \varphi_{t,T}(F_T)\,dF_T < \infty$), it follows

$$\mathbb{E}_t\left[f(F_T)\right] = \int_0^\infty f(K)\varphi_{t,T}(K)\,dK$$

From a purely statistical perspective, these quantities simply resemble the raw (non-centered) moments with respect to some governing probability distribution.

Common examples of future terminal cash flows are, e.g., the identity payoff

$$f(F_T) := F_T$$

and the logarithmic return payoff

$$f(F_T) := \ln\frac{F_T}{F_t}$$

To synthesize distributional expectations of F_T (i.e., moments of the logarithmic return density function), Bakshi et al. (2003) define power payoffs where the first order payoff function $f(F_T)$ is taken to the n-th power $f(F_T) := f(F_T)^n$. Hence, these higher order contracts, or power contracts, are contingent claims which pay off the amount of $f(F_T)^n$ at maturity T for $n = 2, 3, 4$.

With most being theoretical entities not traded on real markets, the curvature of the payoff functions can be reproduced by properly weighting positions of traded options. Thus, the expectations of power contracts under a suitable pricing measure imply no arbitrage prices of most contingent claims. Lastly, no assumptions regarding the stochastic process driving the futures prices of the underlying asset need to be made.

2.2.2 Risk-neutral expectations of credit portfolio loss distributions

Synthetic CDS credit indices have become one of the most significant financial innovations of the last decades, deriving their values from the payoffs of aggregate losses due to defaults in the underlying reference portfolio. With the proliferation of efficient CDS markets, the price discovery process of these standardized contingent claims provides condensed information about correlated default risks. As such, these market expectations are well suited to enrich our knowledge about the latent portfolio loss distribution.

We work on a filtered probability space $(\Omega, \mathscr{F}, \mathbb{Q})$, where Ω is the set of possible market scenarios, equipped with a right-continuous filtration of complete σ-algebras $(\mathscr{F}_t)_{t\geq 0}$ representing the investors' observable information flow, and $\mathbb{Q}$ is a martingale pricing measure reflecting investors' market beliefs in assigning probabilities to $\mathscr{F}$. We assume that, on $(\Omega, \mathscr{F}, \mathbb{Q})$, for each defaultable reference asset i, $i \in 1, \ldots, N$, there is an increasing and nonexplosive sequence of N default times $0 < \tau_1 < \cdots < \tau_N < \infty$ of a $(\mathscr{F}_t)$-adapted counting process $(C_t)_{t\geq 0}$, where $C_t := \sum_{i=1}^{N} \mathbb{I}_{\{\tau_i \leq t\}}$.[5] According to the τ_i, the non-decreasing cumulative default loss process $(L_t)_{t\geq 0}$ counts the number of defaults in the reference credit portfolio, where by definition $L_0 = 0$. Normalizing by the number of constituents $\sum_{i=1}^{N} w_i = 1$ and accounting for the recovery rates R_i after the default of asset i, the losses at time t are

$$L_t = \sum_{i=1}^{N} w_i \, (1 - R_i) \, \mathbb{I}_{\{\tau_i \leq t\}} \tag{2.1}$$

Hence, depending on the choice of recovery rates, the maximum attainable loss may be less than the total notional of the reference portfolio. For example, assuming a constant recovery rate of 40% for all obligors (which is the market convention for most investment grade CDS indices), then the maximum loss due to defaulted assets at some future time T is at 60% of the portfolio notional. Thus, we have

$$0 \leq L_T \leq L_{max}$$

Next, as our aim is to gain valuable information about the total loss distribution of a reference portfolio, we examine the link of tranched and untranched CDS indices to expected losses.[6] In analogy to a capital structure of a firm, a

[5] A filtration representing the flow of observable information is the crucial element of any credit portfolio risk model, as it determines the key properties of the model. Here, we do not make additional assumptions about the filtration, as we use market embedded information about the total loss distribution in a model-independent way. For detailed discussions about filtrations in common bottom-up and top-down model specifications, we refer to Ehlers and Schönbucher (2009) and Giesecke (2009).

[6] The "untranched" CDS index itself can naturally also be regarded as a tranche with attachment point 0% and detachment point 100% (being properly adjusted by the assumed recovery rate for the index).

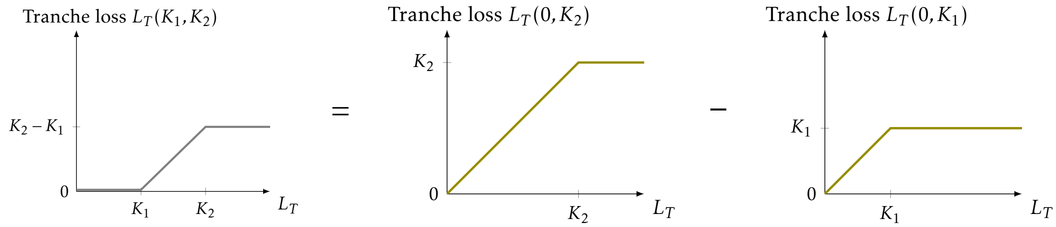

Figure 2.1: Decomposition of CDS index tranche losses. The left hand side of this figure plots the loss profile on a tranche with attachment point K_1 and detachment point K_2 as a function of the terminal underlying portfolio loss L_T. The right hand side of the equation decomposes the tranche loss payoff of $L_T(K_1,K_2)$ into the difference of two equity (base) tranches with strikes K_2 and K_1.

CDS index is tranched into different risk buckets with increasing seniority, starting with an equity tranche absorbing the first occurring losses, over mezzanine tranches up to the most secure super-senior tranche. The expected loss of a tranche $L_T\left(K_1,K_2\right)$ with attachment point K_1 and detachment point K_2 is

$$\mathbb{E}_t\left[L_T\left(K_1,K_2\right)\right] = \frac{\mathbb{E}_t\Big[\min(L_T,K_2)\Big]-\mathbb{E}_t\Big[\min(L_T,K_1)\Big]}{K_2-K_1} \tag{2.2}$$

As depicted in Figure 2.1, the expected tranche loss from K_1 to K_2 is simply a linear combination of equity expected tranche losses that are at the base of the capital structure, i.e., have attachment points equal to zero

$$\mathbb{E}_t\left[L_T\left(K_1,K_2\right)\right] = \frac{K_2\,\mathbb{E}_t\Big[L_T(0,K_2)\Big]-K_1\,\mathbb{E}_t\Big[L_T(0,K_1)\Big]}{K_2-K_1}$$

With the whole capital structure split into adjacent tranches, the expected portfolio loss is simply the sum of the expected tranche losses scaled by the tranche width, i.e.,

$$\mathbb{E}_t\left[L_T\right] = \sum_{j\geq 1}\mathbb{E}_t\Big[L_T(K_{j-1},K_j)\Big](K_j-K_{j-1})$$

For more in-depth treatments of expected tranche losses, including no arbitrage conditions, boundary conditions and model bounds we refer to O'Kane (2008).

Now, let $\psi_{t,T}(K)$ denote the expected tranche loss function

$$\psi_{t,T}(K) = \mathbb{E}_t\Big[\min(L_T,K)\Big] = \int_0^{L_{max}} \min(L_T,K)\,\varphi_{t,T}(L_T)\,dL_T$$

which is the expected loss of an equity tranche with width K. Using the argument in Breeden and Litzenberger (1978), we may infer the risk-neutral density function of the total portfolio loss distribution by twice differentiating the negative of the expected tranche loss[7]

$$\varphi_{t,T}(L_T) = -\frac{\partial^2 \psi_{t,T}(K)}{\partial K^2}\bigg|_{K=L_T}$$

We now apply the useful properties of expected tranche losses to risk-neutral expectations of power loss contracts which are key to our methodology

$$\mathbb{E}_t\Big[f(L_T)\Big] = \int_0^{K_{max}} f(K)\varphi_{t,T}(K)\,dK = \int_0^{K_{max}} f(K)\left(-\frac{\partial^2 \psi_{t,T}(K)}{\partial K^2}\right)dK \tag{2.3}$$

A detailed derivation in terms of generic spanning is in Appendix 2.A.2. Given these preliminaries, our main result is

Proposition 2.1. *Under no arbitrage, the time t risk-neutral expectations of power loss contracts with identity payoff $f(L_T) := L_T$ and order $n \geq 1$ of a credit portfolio with fixed horizon T can be synthesized from the market quotes of (equity) expected tranche losses across all strikes $0 \leq K \leq K_{max}$*

$$\mathbb{E}_t\Big[L_T^n\Big] = n\,\mathbb{E}_t\left[L_T\right]K_{max}^{n-1} - n(n-1)\int_0^{K_{max}} K^{n-2}\,\psi_{t,T}(K)\,dK \tag{2.4}$$

Proof. See Appendix 2.A.1. □

The theoretical link in Equation (2.4), in spanning power loss contracts with a continuum of expected tranche losses provides the basis for inferring the statistical expectations of the underlying credit portfolio loss distribution.[8] This means the centered moments of the loss distribution are then readily available by standard moment conversion formulas.[9]

---

[7] Using the relation $\min(L,K) = L - (L-K)^+$ makes the analogy to the approach in equity markets more evident, since a tranche resembles the well-known structure of a call spread written on the credit portfolio loss.

[8] Our methodology is not affected in the case of one or more defaults happening for the credit index version currently under consideration. There are, however, two points to consider. First, a default reduces the notional of the index (tranche). This affects the actual attachment and detachment points of the index tranches now reflecting the recovered amount of the defaulted entities and the number of remaining names in the index. Second, though the spreads will continue to be quoted using the standard attachment and detachment points, index calculations are subject to the actual attachment and detachment points. A detailed example of the tranche mechanics following a default event is provided in Markit (2014). How the adjusted attachment and detachment points influence the empirical implementation of our methodology is discussed in Section 2.3.2.

[9] Compare Theorem 1 in Bakshi et al. (2003) for an analogical application of power contracts to recover higher order statistical moments of equity return distributions.

2.3 Data and calculation of market implied credit portfolio loss distribution moments

Following the theoretical foundation of our model-free approach to credit portfolio loss distributions, we now examine various methodological issues when applied to empirical data. We begin by outlining the CDS index data used, and then move on to detail our methodology.

2.3.1 Data

The CDS index data used in this study is from the most liquid synthetic credit derivative index backed by US corporate CDS, the CDX North American Investment Grade Index (CDX.NA.IG, or CDX for short). This CDS index is administered by Markit Group Limited and references the 125 most liquid single-name CDS in terms of traded volume. Every six months, on 20th March and September, the equally weighted index is revised or "rolled" over to a new version, which then constitutes the current on-the-run series, whilst the previous version goes off-the-run.[10] CDX trading began on October 21, 2003 and led to well established market quotations for both tranched and untranched credit index products. Besides the CDX index referencing the entire underlying CDS pool, the CDX allows investors to take exposure or to hedge specific portions, or tranches, of the capital structure of the index in exchange for periodic coupon payments due on 20th March, June, September and December until maturity. The riskiness of a tranche is mainly determined by the lower striking level of the tranche, the attachment point, and the upper strike of the tranche, the detachment point. Initially, the attachment points of the CDX tranches had been at the 0%, 3%, 7%, 10%, 15% and 30% levels. In light of the financial crisis, the trading and quoting conventions underwent several amendments, and currently the set of attachment points is at the 0%, 3%, 7% and 15% levels.[11] Furthermore, while the on-the-run CDX index is rolled over every half-year, trading in the tranched CDX indices rolls over once a year, beginning with the CDX series 13. This means, from CDX.NA.IG.13 onwards, the odd numbered series carry additional information about the tranched CDX markets alongside the untranched CDX markets. Detailed expositions of tranche mechanics are provided in Longstaff and Rajan (2008), Coval et al. (2009), and Collin-Dufresne et al. (2012) *inter alia*.

To extract market implied expectations regarding the latent credit index portfolio loss distribution, we collect daily data for the on-the-run series

[10] See the Markit (2015) index rules for a detailed description of the index roll inclusion and exclusion criteria and new series creation process.

[11] See Markit (2014) for historical tranche coupons.

of the five year CDX.NA.IG credit index from Markit for the period from September 2005 to February 2016 (series 5 to 25). This extensive data set covers economic upturns and economic downturns including the financial crisis with major credit events such as the collapse of Lehman Brothers in late 2008 and the S&P downgrade of US sovereign debt in August 2011. The CDS markets experienced extended periods of exponential growth preceding the financial crisis.[12] During the financial crisis however, trading activities in correlation markets dropped sharply and suffered severe liquidity constraints. In addition, deal based price discovery became temporarily discontinuous for a number of tranches in which the CDX series 11 to 13 were on-the-run. Dealer quotes for the on-the-run CDX index series are available throughout. Remarkably, the market in CDX series 9 tranches remained highly liquid during these stressed times and consequently tranche spreads continued to be quoted actively. Section 2.B details our approach on how to deal with these non-trivial data gaps in order to obtain continuous time series for all tranches of the CDX on-the-run series.

The historical behavior of the CDX index and index tranches for the various attachment levels is shown in Figure 2.2. As mentioned above, trading and quotation conventions have changed repeatedly, therefore we express all available data in runnings spreads to facilitate a comparison of the data over time and over the different attachment points.[13]

During the financial crisis, all index and tranche spreads show a sharp increase and reach their maximum in the last quarter of 2008 and first quarter of 2009, respectively, following the defaults of Lehman Brothers and the CDX index members Fannie Mae, Freddie Mac and Washington Mutual. After this, the spreads remain significant above their pre-crises levels, except for the equity tranche—its minimum appears in 2014. After the financial crises, we see two steep increases. The first is in August 2011 where US debt is downgraded. The second is at the end of our time series, where regulatory and liquidity aspects are jointly responsible for an increase of the index and tranche spreads, see Boyarchenko et al. (2016) for a detailed discussion. Table 2.1 provides corresponding summary statistics for the levels and logarithmic differences. As expected for tranche products, all measures of location and dispersion are

[12]Compare, e.g., Collin-Dufresne (2009) and Augustin et al. (2014) for more general credit market overviews.

[13]Compare, e.g., Doctor and Singh (2010) for a practitioner-oriented summary of the main changes regarding CDS markets. These are primarily changes in the CDS documentation introduced via the Big Bang and Small Bang protocols, the trading and quoting conventions, to improve netting and clearing of CDS contracts as well as the standardization of the conversion of upfront (plus fixed coupon) to spreads all running and *vice versa*, and issues pertaining to central clearing via CDS clearing houses.

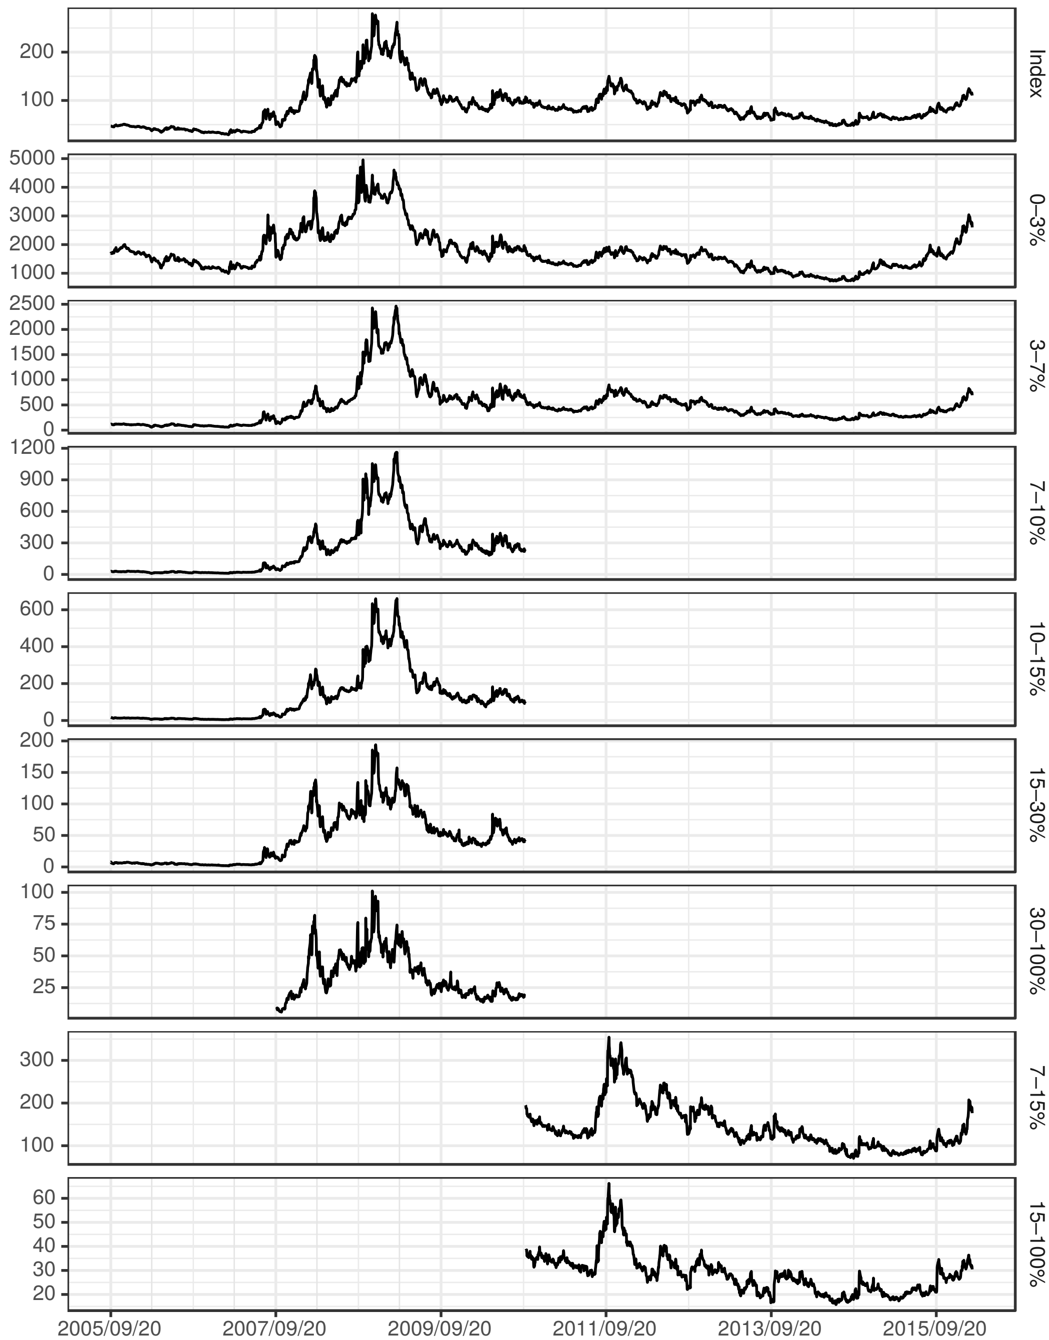

Figure 2.2: Time series of tranched and untranched CDX index spreads. This figure graphs the historical evolution of the quoted on-the-run five year tranched and untranched CDX North American Investment Grade credit index. The spreads are all running in basis points. Index roll dates are indicated via the vertical grid lines. The sample period is from September 20, 2005 to February 26, 2016.

Table 2.1: Summary statistics for the levels and logarithmic differences of the tranched and untranched CDX index spreads. Entries report the summary statistics for the quoted spreads of the on-the-run five year tranched and untranched CDX North American Investment Grade credit index. The spreads are all running in basis points. The upper panel is based on index spread levels. We also report the corresponding statistics for the log differences in the lower panel. IQR denotes the interquartile range defined as the difference between the 75% and 25% quantiles. The sample is from September 20, 2005 to February 26, 2016.

	Correlations																	
	0–3%	3–7%	7–10%	10–15%	15–30%	30–100%	7–15%	15–100%	Mean	Std. Dev.	IQR	Min.	Median	Max.	Skew.	Kurt.	Serial corr.	*N*
Index	0.852	0.953	0.968	0.962	0.979	0.932	0.915	0.916	87.600	42.176	41.244	28.875	81.775	279.667	1.479	2.917	0.996	2,600
0–3%		0.795	0.879	0.866	0.884	0.854	0.572	0.599	1737.699	729.936	674.617	728.399	1580.054	4956.321	1.545	2.538	0.996	2,600
3–7%			0.991	0.982	0.901	0.750	0.922	0.870	477.871	385.415	338.629	56.972	402.614	2463.886	2.274	6.621	0.997	2,600
7–10%				0.985	0.925	0.807	–	–	238.692	251.789	297.343	9.111	218.996	1161.774	1.413	1.670	0.996	1,242
10–15%					0.931	0.835	–	–	125.249	141.958	156.602	4.321	102.646	660.958	1.644	2.412	0.996	1,242
15–30%						0.968	–	–	45.523	42.117	68.833	1.639	39.935	193.978	0.863	0.065	0.993	1,242
30–100%							–	–	34.751	18.451	27.757	5.703	28.257	101.107	0.863	0.228	0.978	752
7–15%								0.883	142.740	53.316	64.062	70.420	129.702	354.352	1.296	1.555	0.993	1,358
15–100%									28.432	8.361	10.757	15.799	27.743	66.206	1.211	2.122	0.990	1,358
Δ Index	0.844	0.880	0.851	0.832	0.807	0.740	0.912	0.839	0.000	0.032	0.030	−0.197	−0.000	0.290	0.827	10.630	0.101	2,599
Δ 0–3%		0.848	0.747	0.674	0.619	0.414	0.847	0.775	0.000	0.029	0.029	−0.229	−0.000	0.196	0.264	7.733	0.155	2,599
Δ 3–7%			0.921	0.833	0.716	0.532	0.919	0.830	0.001	0.045	0.042	−0.277	−0.001	0.339	0.773	7.953	0.113	2,599
Δ 7–10%				0.885	0.777	0.625	–	–	0.002	0.064	0.062	−0.328	−0.001	0.460	0.454	6.143	0.105	1,241
Δ 10–15%					0.811	0.656	–	–	0.001	0.068	0.067	−0.335	−0.003	0.459	0.462	5.026	0.100	1,241
Δ 15–30%						0.747	–	–	0.001	0.074	0.072	−0.375	−0.001	0.500	0.600	6.691	0.025	1,241
Δ 30–100%							–	–	0.001	0.082	0.093	−0.332	0.001	0.539	0.614	4.946	0.006	751
Δ 7–15%								0.840	−0.000	0.040	0.041	−0.141	−0.001	0.363	1.894	15.363	−0.011	1,357
Δ 15–100%									−0.000	0.038	0.040	−0.166	−0.001	0.346	1.344	11.873	−0.022	1,357

monotone decreasing along the capital structure. Overall, the correlations of tranche and index spreads are very high. Only the equity tranche shows a rather different behavior. Here, the correlations are substantially smaller, especially for the new attachments 7–15% and 15–100%.

2.3.2 Calculation of market implied moments

As shown in Equation (2.4), the model-free expectations of power loss contracts are defined as integrals of expected tranche losses over a continuous collection of strikes, ranging from zero to the maximum possible loss of the underlying reference portfolio. This poses two empirical challenges. First, the capital structure of the CDX credit index is divided into a finite number of consecutive attachment and detachment points; consequently, only corresponding strikes are actually traded in the market. Hence, we need to complete the set of available strikes to attain a continuous range on $[0, K_{max}]$. Second, following CDX index market conventions, expected tranche losses are implicitly embedded in the quotes of tranched and untranched CDS index spreads. Thus, we need to extract the implicit information contained in tranche quotes and translate the available information into equity expected tranche losses obeying no arbitrage conditions. These implementation issues will be discussed in detail in the following.

Availability of strike prices and numerical integration To evaluate the integrand in Equation (2.4), we need expected tranche losses across a continuous set of strike prices over the bounded range $[0, K_{max}]$. Given the tranching conventions for the CDX index products, expected tranche losses are only traded at the quoted attachment points along the capital structure of the CDX. Following common practice in the equity literature (e.g., Jiang and Tian, 2005, 2007; Carr and Wu, 2009), we apply a curve fitting method preserving the no arbitrage properties of expected tranche losses and interpolate along available attachment points. Figure 2.3 illustrates expected tranche losses at various attachment points typical to the CDX index. In addition, we know that a zero width tranche has zero loss, i.e., $\psi_{t,T}(0) = 0$. Furthermore, we know that the expected tranche loss of an equity tranche with a width greater than or equal to the maximum portfolio loss L_{max} must be equal to the expected loss of the portfolio, i.e., $\psi_{t,T}(K) = \mathbb{E}_t[L_T]$ for $K \geq L_{max}$. Moreover, given that a tranche loses no more than its width, i.e., $\psi_{t,T}(K) \leq K$, and using $\psi_{t,T}(0) = 0$, we obtain that $\left.\frac{\partial \psi_{t,T}(K)}{\partial K}\right|_{K=0} \leq 1$. Thus, the $\psi_{t,T}(K)$ curve must be a monotone increasing mapping to available strikes with slope being bounded between one and zero, or put differently, $\frac{\partial \psi_{t,T}(K)}{\partial K} = \mathbb{P}(L_T > K)$ is the probability

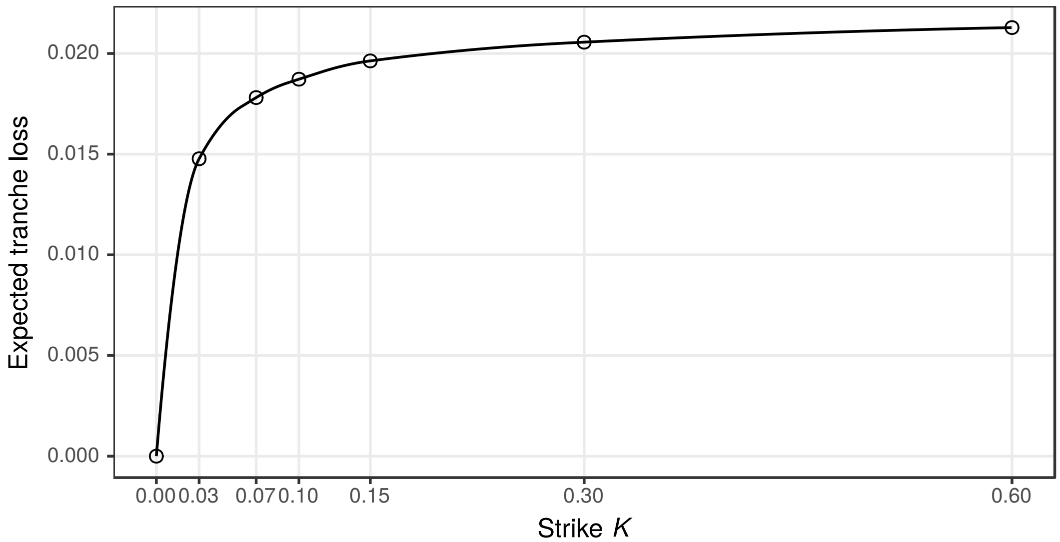

Figure 2.3: Monotone and concave interpolating spline function. This figure graphs an exemplary monotone and concave shape preserving interpolating spline function to obtain a smooth continuum of equity expected tranche losses from zero losses to K_{max} given a set of attachment points determined via the calibration instruments.

of having a loss greater than K. To ensure that the density of the portfolio loss distribution is positive everywhere, we further require that $\frac{\partial^2 \psi_{t,T}(K)}{\partial K^2} \leq 0$ which strengthens the monotonicity condition, by implying a globally concave mapping of $\psi_{t,T}(K)$ to available strikes. These observations complete the set of requirements for the interpolation function.

Thus, in order to prevent arbitrage, we need an interpolation scheme capable of addressing these monotonicity and concavity requirements. Here, we employ the monotone and concave quadratic spline interpolation method introduced in Schumaker (1983) using slopes approximated by harmonic means (Butland, 1980).[14] The involved expressions are all analytical, do not involve any optimization and, thus, are fast to evaluate. As shown in Figure 2.3, the Schumaker spline function exhibits a smooth continuum of expected tranche losses given the set of available strike prices.

Next, we turn our attention to the numerical evaluation of the integral in (2.4) and make a few remarks in comparison to equity applications. First, we are not concerned with the issue of truncation errors, since we know the upper and lower limits of the integration range and do not have to extrap-

[14]Other choices of interpolation schemes applied to expected tranche losses are discussed below in Section 2.5.1.

olate to either side. Second, we have no discretization and no integration errors as we analytically evaluate the involved integrals leveraging the known piecewise constant spline coefficients.[15] Third, though a number of five to seven available strikes may seem few, model-free applications to single-name equities limit the calculation to a minimum of at least three (e.g., Carr and Wu, 2009) or four available strikes (e.g., Conrad et al., 2013). These strikes have to be interpolated and extrapolated to both sides on $[0,+\infty]$ (predominantly within the model based domain of Black–Scholes implied volatilities) and are typically evaluated on a discretized grid. In contrast, our domain of integration is fully supported by the data. In addition, we analytically evaluate the interpolation function directly working in the problem space, while adhering to the no arbitrage conditions.

No arbitrage calibration of expected tranche losses Up to now, we presumed full knowledge about the expected tranche losses at time T and tranche attachments embedded in market quotes. Here, we illustrate how these entities may be extracted from a set of index and tranche spreads. In fact, we need a full expected tranche loss surface across the periodic coupon payments until maturity in the time dimension and the strike dimension.

Consider standard tranched index pricing via equating default and premium legs viewed from time 0 to maturity at time T

$$s_{0,T}(A,D) = \frac{\int_0^T D(0,t)\, d\,\mathbb{E}_0\left[L_t(A,D)\right] - U_{0,T}(A,D)}{\sum_{i=1}^{N_T} \Delta_i D(0,t_i)\Big(1 - \Big(\mathbb{E}_0\left[L_{t_i}(A,D)\right] + \mathbb{E}_0\left[R_{t_i}(A,D)\right]\Big)\Big)} \tag{2.5}$$

where the numerator (the default leg) and denominator (the default risky annuity) are linear functions of the expected tranche loss $\mathbb{E}_0\left[L_t(A,D)\right]$ and expected tranche recovery $\mathbb{E}_0\left[R_t(A,D)\right]$ at each point in time for the attachment point with strike A and detachment point with strike D. The number of payment dates until T is denoted by N_T, $D(0,t)$ are the time 0 prices of zero bonds paying one dollar at time t (assuming that interest rates and default times are independent), and $\Delta_i = t_i - t_{i-1}$ is the time difference from the current to the previous payment date. Given tranche spreads $s_{0,T}(A,D)$ and initial upfront payments $U_{0,T}(A,D)$ to enter the positions (depending on the quoting convention), then we may back out the market implied dependence structure in the form of the expectations $\mathbb{E}_0\left[L_t(A,D)\right]$ and $\mathbb{E}_0\left[R_t(A,D)\right]$ in

[15]The general idea of analytical spline integration is outlined in Section 2.C of the Appendix.

a model-independent way (given minimal interpolation assumptions, see below).[16]

The index tranche loss payoff as a function of the portfolio loss L_t at time t is

$$L_t(A,D) = \frac{\min\Big((L_t - A)^+, D - A\Big)}{D - A}$$

and the index tranche recovery payoff

$$R_t(A,D) = \frac{\min\Big((R_t - (1-D))^+, D - A\Big)}{D - A}$$

reduces the outstanding notional for the calculation of the premium payments by the amount recovered in the case of default.[17]

To back out the desired $\psi_{0,T}(K)$, we may reformulate the pricing formula in (2.5) in terms of the equity expected tranche losses $\psi_{0,T}(K)$ using the relation in (2.2). For simplification, we assume the expected tranche losses increase linearly in time between $\psi_{0,0}(K) = 0$ to $\psi_{0,T}(K)$, in order to calculate the cash flows on the payment dates before maturity T. Following the industry standard, we use the deposit-swap curve in order to discount the cash flow streams in the loss compensation and premium legs. Specifically, we collect daily USD LIBOR deposit rates as well as USD ISDA FIX swap rates from the Thompson Reuters Datastream database and bootstrap the term structure of zero bonds based on a piecewise constant forward rate assumption to match the coupon payment dates. Then, given a set of calibration instruments, we minimize the difference of tranche spreads observed in the marketplace $s_{0,T}^{Market}(A,D)$ to the tranche spreads implied by the corresponding expected tranche losses $s_{0,T}^{\psi}(A,D)$

$$\arg \min_{\psi(3\%,7\%,\ldots,100\%)} \sum_{A,D} w_{A,D} \Big(s_{0,T}^{Market}(A,D) - s_{0,T}^{\psi}(A,D)\Big)^2 \tag{2.6}$$

[16]During the actual calibration exercises, we discretize the involved intregrals and additionally account for accrued premiums and follow market practices in evaluating the risky annuity and the default leg by assuming that defaults occur midway between payment dates.

[17]For the super senior tranche with attachment point A below $L_{max} = (1-R)$, each default in the portfolio reduces the notional by $\frac{1}{N} - \frac{(1-R)}{N}$. Since “the loss goes to the equity tranche, the recovery to the super senior” Markit (2014, p. 16), for practical applications we may simply scale the portfolio recovery by the tranche width to obtain the “expected” recovery for the super senior tranche. For the [0%,100%] tranche, i.e., the CDX index, Equation (2.5) naturally corresponds to the more traditional single purpose index pricing formula, as the tranche width then covers the whole capital structure. For the remaining tranches other than the super senior, the amount recovered is zero, as expected.

with some possible weighting $w_{A,D}$.[18] To include as much available market data as possible, we also consider the index spread as a [0%,100%] tranche. Most importantly, the calibration of the $\psi_{0,T}(K)$ is carried out with no arbitrage constraints. This means, in order to maintain monotonicity, concavity and time consistency, we require

$$
\begin{aligned}
0 \leq \psi_{0,t}(K) &\leq \mathbb{E}_0\left[L_t\right] \\
\psi_{0,t_i}(K) &\geq \psi_{0,t_{i-1}}(K) \\
\psi_{0,t}(K_{j-1}) &\leq \psi_{0,t}(K_j) \\
\psi_{0,t}(K_{j+1}) &\leq \psi_{0,t}(K_j) + (K_{j+1} - K_j)\frac{\psi_{0,t}(K_j) - \psi_{0,t}(K_{j-1})}{K_j - K_{j-1}}
\end{aligned}
$$

For more details on the optimization procedure we refer to Brigo et al. (2010). As is mentioned in Section 2.2.2, occurring default events in the index reference pool do not impair our methodology. However, in these cases care must be taken (i) to the calibration of the expected tranche losses to market quotes, and (ii) to the interpolating spline functions which need to be set up accordingly with the adjusted attachment and detachment points.

Figure 2.4 plots the time series of daily root mean squared error (RMSE) in basis points from fitting expected tranche losses to the index and index tranches. As shown, initially the calibration yields a virtually perfect fit with RMSE hardly distinguishable from zero. This is a major improvement in comparison to sophisticated parametric pricing models like the three-jump intensity based model introduced in Longstaff and Rajan (2008) where the authors report RMSE around two basis points during calm periods and exceeding five basis points during stressed times like the May 2005 correlation crisis with peaks up to 19 basis points within their sample period from October 2003 to October 2005.[19] This is until the super senior tranche became more liquid and Markit started to disseminate corresponding dealer quotes on September 24, 2007 with the inception of CDX series 9. With the super senior quotes now available, the capital structure is completed and the fit reflects that the super senior expected tranche loss $\psi_{0,T}(100\%)$ now must reprice the super senior tranche and the index simultaneously. Though, given the

[18] In Section 2.5.2 we consider a number of alternative objective loss functions to assess the robustness on implied moments.

[19] Furthermore, Longstaff and Rajan (2008) do not include the index in their objective function. Rather, the index level is used as an implicit calibration instrument to provide additional structure during their calibration exercises in that the three default components add up to the index level. Finally, the authors do not consider the super senior tranche probably due to limited liquidity in this remaining part of the capital structure within their proprietary data set.

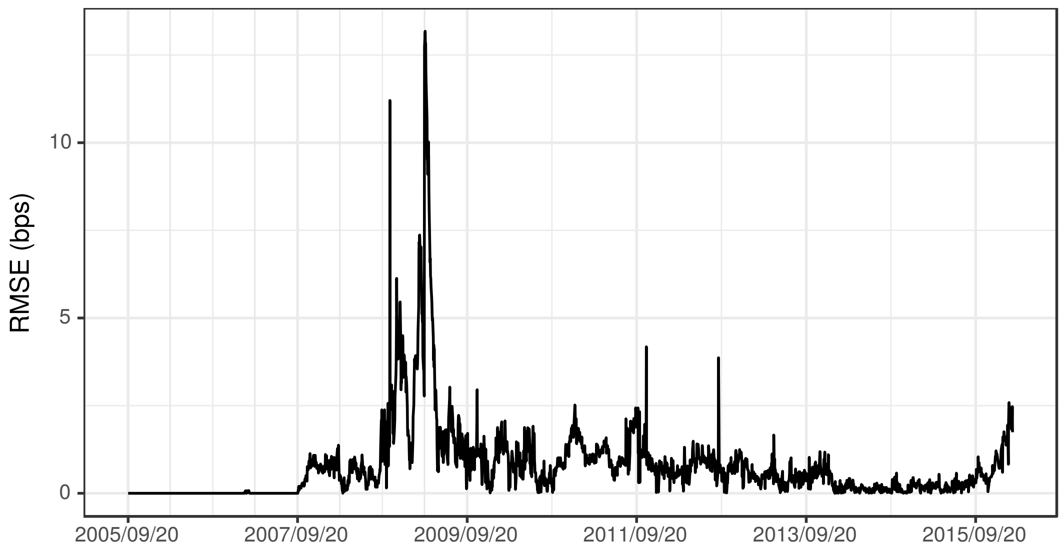

Figure 2.4: Calibration quality. This figure graphs the daily evolution of root mean squared error (RMSE) from fitting expected tranche losses to the CDX index calibration instruments. All RMSE are in basis points. From early series 9 onwards (i.e., 09/24/2007), the calibration instruments include the super senior tranches. Index roll dates are indicated via the vertical grid lines. The sample is from September 20, 2005 to February 26, 2016.

still small magnitudes of the RMSE, this may nevertheless be considered a near perfect fit. The overall RMSE of the benchmark calibration is 1.548 basis points whereas the mean and median of the daily RMSE are 0.772 and 0.489 basis points, respectively. Not surprisingly, however, during the financial crisis the calibration quality is somewhat diminished and exhibits a few spikes reaching about 13 basis points. This turbulent phase is also the main contributor to the overall RMSE. After the crisis, at times, we observe small peaks in the RMSE which are potentially due to minor violations of the no arbitrage conditions in the market data. In conclusion, the mispricing of the index and index tranche spreads is extremely small.

Comparison to equity applications In conclusion, the outlined procedure shares many similarities to model-free equity applications. This means, before evaluating the involved integrals over a continuum of (inter- and extrapolated) call and put options, the underlying raw data comes under scrutiny regarding valid no arbitrage bounds, monotonicity and convexity restrictions, see, e.g., Aït-Sahalia and Duarte (2003) or Andersen and Bondarenko (2007). Hence the market implied expected payoffs $\mathbb{E}_t\left[(K-F_T)^+\right] = P_{t,T}(K)$ and $\mathbb{E}_t\left[(F_T-K)^+\right] =$

$C_{t,T}(K)$ are filtered or smoothed accordingly. Furthermore, the expected total portfolio loss $\mathbb{E}_t[L_T]$ enters the calculation of power loss contracts. This entity is recovered from liquid tranche quotes under no arbitrage restrictions, which in principle is no different to the approach of regressing an implied forward level on the validity of put-call parity around the at-the-money level to mitigate a number of possible empirical issues. Common issues are, for instance, time synchronicity issues with corresponding options or issues pertaining to cost of carry (e.g., expected future dividends or equity repo rates) calculations, see, among many others, Aït-Sahalia and Lo (1998). In fact, it is the expectation $\mathbb{E}_t[F_T]$ that implicitly enters the calculation of equity power contracts, since under risk-neutrality and the money market as numéraire, the future is a martingale, or $\mathbb{E}_t[F_T] = F_t$.

Historical evolution of the market implied moments of the credit portfolio loss distribution Figure 2.5 exhibits the first four centered moments for the five year losses of the CDX index portfolio. Before the onset of the financial crisis in August 2007, the market implied expected losses were very low at around 2%. We also observe quite small variances and pronounced values for the skewness and kurtosis coefficients. One reason for these seemingly extreme values for the third and forth moments is that the loss variable is bounded between $[0, L_{max}]$. This naturally leads to heavily right skewed and leptokurtic loss distributions, when the members of the portfolio have relatively low idiosyncratic default probabilities, but the probability of simultaneous losses, which also may be classified as systemic risk, is relevant.[20] Our higher moments are consistent with the loss distributions depicted in, e.g., Collin-Dufresne et al. (2012), who analyze the CDX index from series 3 to series 10 via a structural model, and in Longstaff and Rajan (2008), who analyze the CDX series 1 to 5. As a consequence, the risk-neutral probability of zero or one default is very high before the financial crisis.

With the outbreak of the financial crisis, we see quite pronounced changes in all of the four implied moments. The expected loss increases to 13%, with a mean of around 6% during the financial crisis, which is three times higher than pre-crisis. Taking an assumed recovery rate of 40% into account, this means that the implied risk-neutral default probability for the CDX reaches a

[20]The shape and asymmetry of a univariate distribution is commonly measured by its skewness coefficient. It is well known that the textbook interpretation of kurtosis as a measure for peakedness, or tail weight of a distribution, is somewhat confounded when applied to asymmetric distributions. That is, higher skewness inevitably leads to higher kurtosis. A higher kurtosis follows from a movement of probability mass from the "shoulders" of a distribution into its tails and center which in turn are affected by its skewness. These multifarious subtleties are discussed in, e.g., Jones et al. (2011) and the references therein.

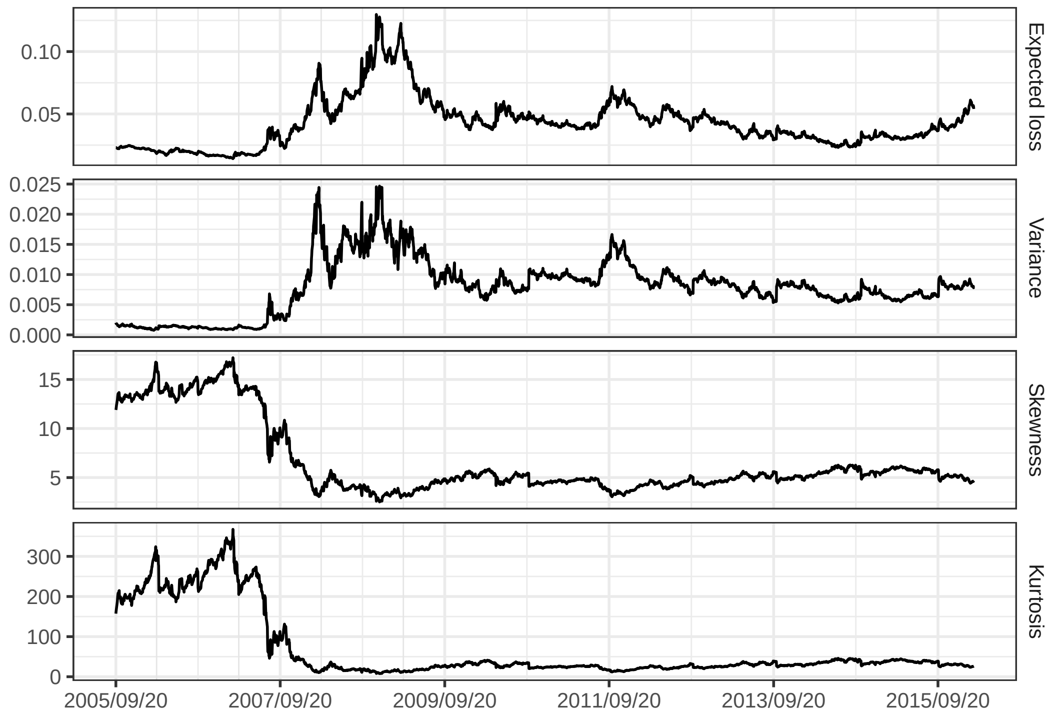

Figure 2.5: CDX credit portfolio loss distribution moments. This figure graphs the daily implied model-free risk-neutral expected loss, loss variance, skewness, and kurtosis of the five year tranched and untranched CDX North American Investment Grade credit index. Index roll dates are indicated via the vertical grid lines. The sample is from September 20, 2005 to February 26, 2016.

maximum of over 20%. For comparison, the physical five year default rate from Moody's for all investment grades in the period of 1983 to 2016 is around 1%, with a cohort specific maximum for the year 1986 of 2.3%, see Ou (2017).

The increase of the second moment during the financial crisis is significantly stronger than for the first moment. The mean of the second moment during the financial crisis is 1.2%, which is eight times higher than the pre-crisis value. The maximum value is 2.5%. The variability of the expected loss and the variance is very high during the financial crisis, in contrast to the pre-crisis period, where both exhibit a relatively stable behavior.

In contrast to the first and second moment, the decrease of the skewness and kurtosis have already commenced by mid 2007. Here, the skewness drops from its maximum of 17 to around 8 and the kurtosis reduces from 367 to around 50. This is well before we see a reaction in one of the first two moments. During the financial crisis, skewness and kurtosis still decrease and reach their minimum values of 2.5 and 7.9, respectively.

Table 2.2: Summary statistics for the levels and logarithmic differences of the CDX credit portfolio loss distribution moments. Entries report the summary statistics for the implied model-free risk-neutral default loss distribution moments of the on-the-run five year CDX North American Investment Grade credit index. The upper panel is based on levels. We also report the corresponding statistics for the log differences in the lower panel. IQR denotes the interquartile range defined as the difference between the 75% and 25% quantiles. The sample is from September 20, 2005 to February 26, 2016.

	Correlations												
	Var	Skew	Kurt	Mean	Std. Dev.	IQR	Min.	Median	Max.	Skew.	Kurt.	Serial corr.	N
$\mathbb{E}$	0.912	−0.675	−0.601	0.042	0.020	0.020	0.014	0.040	0.130	1.389	5.567	0.996	2,600
Var		−0.817	−0.748	0.008	0.005	0.004	0.001	0.008	0.025	0.523	3.699	0.993	2,600
Skew			0.988	6.497	3.663	1.557	2.515	5.057	17.219	1.545	3.839	0.999	2,600
Kurt				66.437	83.592	17.855	7.871	30.064	367.462	1.826	4.785	0.999	2,600
$\Delta\mathbb{E}$	0.743	−0.898	−0.883	0.000	0.031	0.030	−0.171	−0.000	0.253	0.830	12.210	0.090	2,599
ΔVar		−0.852	−0.922	0.001	0.051	0.042	−0.441	−0.001	0.528	1.101	17.026	0.061	2,599
ΔSkew			0.987	−0.000	0.025	0.021	−0.283	0.001	0.186	−1.455	20.062	0.050	2,599
ΔKurt				−0.001	0.050	0.041	−0.536	0.001	0.401	−1.438	20.145	0.059	2,599

After the financial crisis, the first and second moment recover from the market turbulences but still remain at a higher than pre-crisis level. The post-crisis mean of the expected loss is 3.9% and 0.8% for the variance. On the other hand, the third and fourth moments are substantially smaller than pre-crisis, with mean values of 4.9 and 28.1, respectively. All implied measures suggest that the risk assessment in the CDX market differs noticeably between pre- and post-crisis periods.

If we calculate the realized coefficient of variation for all four moments during the pre-crisis, crisis and post-crisis periods, we note that all moments, except for the variance, have virtually the same pre- and post-crisis level. The highest variation is observed during the financial crisis. This is not true for the second moment. Here, the relative variation is highest pre-crisis with 0.54, and decreases to a post-crisis level of 0.24. The coefficient of variation is 0.40 for the financial crisis.

Approaching the end of the time series, the expected losses increase, however, we see little or no effects on the higher moments. In contrast, the variance increases and skewness as well as kurtosis decrease with the inception of the financial crisis or the US downgrade in August 2011. This supports the findings of Boyarchenko et al. (2016) in that the increase of credit index spreads at the end of 2015 and start of 2016 is mainly driven by other reasons than the two preceding crises.

Table 2.2 reports summary statistics for all higher moments and their log differences. The high mean values of skewness and kurtosis are mainly driven from the time period before 2007. The correlations between the four moments are high, especially for the third and fourth moments, where we observe a near perfect linear dependence. This still holds true for the log differences. On the equity markets, the correlation between option implied skewness and kurtosis is roughly as large as in the credit market.[21]

2.4 Empirical applications

Synthesizing model-free expectations from markets of correlated default risk offers a plethora of possible applications and provides a rich source for empirical inquiries. Here, we focus on three initial questions. First, what can we learn from the distributional moments of the latent underlying loss process regarding the total risk-neutral default loss distribution? Second, what can

[21] Given the data provided in Chang et al. (2013) for the S&P 500 from January 1996 through December 2007, the correlation of the implied skewness to the kurtosis is −0.915. The corresponding correlations of the implied variance to the skewness (0.053) and variance to the kurtosis (−0.216) are, however, decidedly lower.

we learn about the implicit risk-neutral correlation structure embedded in the tranched and untranched CDX index markets? Third, may the changes of risk-neutral default loss moments be attributed to common risk factors?

2.4.1 Credit portfolio loss distributions

Markets for correlated default risk provide forward looking estimates for the expected future portfolio loss, variance and higher order moments of the latent default counting loss process on the underlying reference portfolio. Ideally, we wish to be able to augment our partial observations in the form of the market implied model-free distributional moments and derive the total risk-neutral portfolio loss distribution. In the realm of equity derivatives applications, there is a substantial body of literature to extract the underlying risk-neutral density function given the market participants' expectations of the future asset price process in the form of liquid call and put options.[22] In the realm of credit derivatives applications, the existing literature seems quite scarce with respect to non-parametric approaches to the estimation of full credit portfolio loss distributions. One notable exception is the approach taken by Vacca (2005), who introduces the principle of maximum entropy to derive portfolio loss distributions consistent with tranche quotes on CDS indices.[23] Here, we take another route and apply the model-free loss distribution moment estimates to orthogonal series expansions of probability density functions and discuss implications of the loss distribution moments on the risk-neutral loss distribution.[24]

An expansion of probability density and distribution functions into a series of orthogonal functions is a well-established tool when the only available information on the variable of interest is limited to its moments. In our approximation of the loss distribution, we focus on generalized Gram–Charlier expansions based on one of the "classical" polynomials with their corresponding weight functions which are the Hermite, Laguerre and Jacobi orthogonal

[22] For recent applications and comprehensive reviews of the literature, we refer to Monteiro et al. (2008), Haven et al. (2009), Monnier (2013), and Schlögl (2013) among many others.

[23] Vacca (2005) uses the actual tranche quotes as moment constraints. This may lead to kinky loss distributions as, in this case, the Lagrange multipliers yield exponential piecewise linear interpolants between the tranche price constraints. Future applications could apply our model-free (non-centered) distributional moments leading to polynomial interpolants and consequently to smooth loss density functions.

[24] Given the knowledge of a series of loss distribution moments, the moment generating function, or equivalently the characteristic function, one may be inclined to directly infer the total loss distribution via Laplace or Fourier inversion methods. It turns out that this direct approach seems unsuitable without further investigation, compare Rompolis and Tzavalis (2008) for similar conclusions in options pricing applications.

polynomials.[25] Here we consider smooth continuous representations of the loss distribution, discrete approximates may be found by continuity corrections, see Provost et al. (2009). In the sequel, our layout borrows from Oakley (1990), Reinking (2002), and Provost (2005).

For a probability density function $f(x)$, defined almost everywhere on the interval (a,b) and for which $\int_a^b [f(x)]^2\,dx < \infty$ is said to be square integrable over the interval (a,b), i.e., $f(x) \in \mathscr{L}^2$, the formal series expansion may be written as

$$f(x) = w(x)\sum_{k=0}^{\infty} d_k P_k(x) \tag{2.7}$$

A system of (generic) polynomials $P_k(x)$ with degree exactly k is said to be orthogonal on the interval (a,b) with respect to an admissible weight function $w(x)$ that is non-negative and fulfills $\int_a^b w(x)\,dx > 0$, if

$$\int_a^b w(x)P_nP_m(x)\,dx = \begin{cases} 0 & \text{if } m \neq n \\ h_n \neq 0 & \text{if } m = n \end{cases}$$

The constants d_k are determined by

$$d_k = \frac{1}{h_k}\int_a^b f(x)P_k(x)\,dx$$

Now, if $f(x) = w(x)\sum_{k=0}^{\infty} d_k P_k(x)$ for all $x \in (a,b)$ and the inner product $\int_a^b w(x)[P_k(x)]^2\,dx = h_k$, then d_k is simply a linear combination of the moments of X, for all k. To see this, we may apply the useful fact that the orthogonal polynomials are linear combinations in the non-negative powers of x, i.e., $P_k(x) = \sum_{i=0}^{k} p_{ik}x^i$, where p_{ik} denotes the corresponding coefficients of x^i in the polynomial $P_k(x)$ with degree k. Therefore, we obtain the equivalent representation of the integral in d_k

$$\int_a^b f(x)P_k(x)\,dx = \sum_{i=0}^{k} p_{ik}m_i$$

[25] Fully generalized Gram–Charlier expansions, based on Fourier techniques to express the probability density function of a random variable in terms of any other, are developed in Berberan-Santos (2007). Gram–Charlier series expansions are also intimately related to the Edgeworth and Cornish–Fisher approximations, compare, e.g., Oakley (1990) for details and historical accounts.

where m_i denotes the i-th raw (non-centered) moment of X. Using this important relation, we may express the coefficients d_k in the equivalent form

$$d_k = \frac{1}{h_k} \sum_{i=0}^{k} p_{ik} m_i$$

Thus, given the knowledge of the coefficients of the polynomial $P_k(x)$ and the raw (non-centered) moments of the probability distribution, the constants d_k may be easily calculated. It is worth mentioning, that given a sequence of moments and the range of the random variable is finite, this is a sufficient condition to define a unique density function.[26] For any practical applications of the orthogonal series expansion in (2.7), the infinite sum will be truncated at some point n

$$f(x) = w(x) \sum_{k=0}^{n} d_k P_k(x) + \epsilon(n,x) \tag{2.8}$$

where $\epsilon(n,x)$ represents the resulting error in the approximation.[27]

In order to apply the outlined methodology to credit portfolio loss distributions, one must carefully decide on the set of orthogonal polynomials corresponding to some of the most important of probability density functions as weighting functions. One obvious approach in the selection of the polynomials is to relate the domain of the random variable to the interval of orthogonality of the polynomial family. Another simple possibility is a visual inspection of the histogram of a population sample and a comparison of its shape to the curvature of the approximating base density functions.

For distributions with infinite support, the Hermite polynomials orthogonal over the infinite interval $(-\infty,\infty)$ are a natural choice. The corresponding weight function is the normal distribution defined on $(-\infty,\infty)$. The use of Hermite polynomials has been well studied in a variety of financial applications and gives rise to the common Gram–Charlier type A series.

For a random variable with support on the real half-line, the (generalized) Laguerre polynomials orthogonal over the semi-infinite interval $[0,\infty)$ with

---

[26] This is not necessarily true for random variables with support on the semi-infinite or infinite domain. These circumstances are commonly referred to as "the moment problem". Compare Section 2b.5 in Rao (1973) for details.

[27] Interestingly, the coefficients $d_k = \frac{1}{h_k}\mathbb{E}\left[P_k(x)\right]$ minimize the weighted mean squared error of the truncated series over the interval (a,b). Thus, for any orthogonal polynomial the d_k may be computed independent of the number of polynomials used in the truncated series. In other words, the coefficients for the finite series (2.8) are identical to those from the infinite series in (2.7). Compare Oakley (1990) for a formal derivation of error estimates due to series truncation.

the gamma distribution (also defined on $[0,\infty))$ as weight function is often suitable. Random variables on the domain $[a,\infty)$ or $(-\infty,a]$ may be easily mapped to $[0,\infty)$.

Finally, a random variable assuming compact support on (a,b) lends itself to the use of Jacobi polynomials orthogonal over the finite interval $(-1,1)$ or the more convenient transformed representation of the Jacobi polynomials orthogonal over the finite interval $(0,1)$. The corresponding weight function is the beta distribution defined on $[0,1]$. Via linear affine variable transformations, arbitrary intervals (a,b) may be easily mapped to the interval of orthogonality of the respective orthogonal polynomial.

Here, based on the expected shape of the unknown loss density and given the domain of the underlying portfolio loss variable in relation to the interval of orthogonality, we consider the beta distribution an appropriate candidate as underlying base density to series expand credit portfolio loss distributions around our model-free moments. The domain of the loss process on $[0,L_{max}]$ is transformed accordingly to the support of the beta on $[0,1]$. Besides the natural mapping of the loss variable to the Jacobi polynomials, beta distributions have a longstanding tradition in the modeling of correlated default risk, compare, e.g., Frey and McNeil (2002, 2003).

In the literature, there is no clear consensus on the choice of truncation point for the series expansion. Too few moments may miss critical information about the shape, potential multimodality and the tail behavior. A number too high may introduce or accumulate numerical errors in the estimation of sample based moments. Referring to the integrand in Equation (2.4), the expected tranche losses are increasingly weighted with growing order of the moment along the strike dimension up to the maximum attainable portfolio loss. Furthermore, if the spread quotes are contaminated with noise, this will eventually add up to unreliable estimates for very high degrees. To strike a good balance to capture the characteristics of the loss distribution and possible numerical issues for very high order moment estimates, we find the first eight distributional moments to be a good compromise in our empirical analysis. Additionally, we average the daily moment estimates over the calendar week to reduce the number of dates in the sample. Given that here we are particularly interested in the historical behavior and are discussing economical implications, this approach has the beneficial side effect of smoothing the daily variation in the data and minimizing possible weekday effects (Bakshi et al., 2003).

Figure 2.6 plots the beta series expanded risk-neutral loss distributions based on the model-free moments estimates. We can clearly identify three different regimes, one pre-crisis, one financial crisis, and one post-crisis. At the start of our sample (pre-crisis), the loss distribution is multimodal and

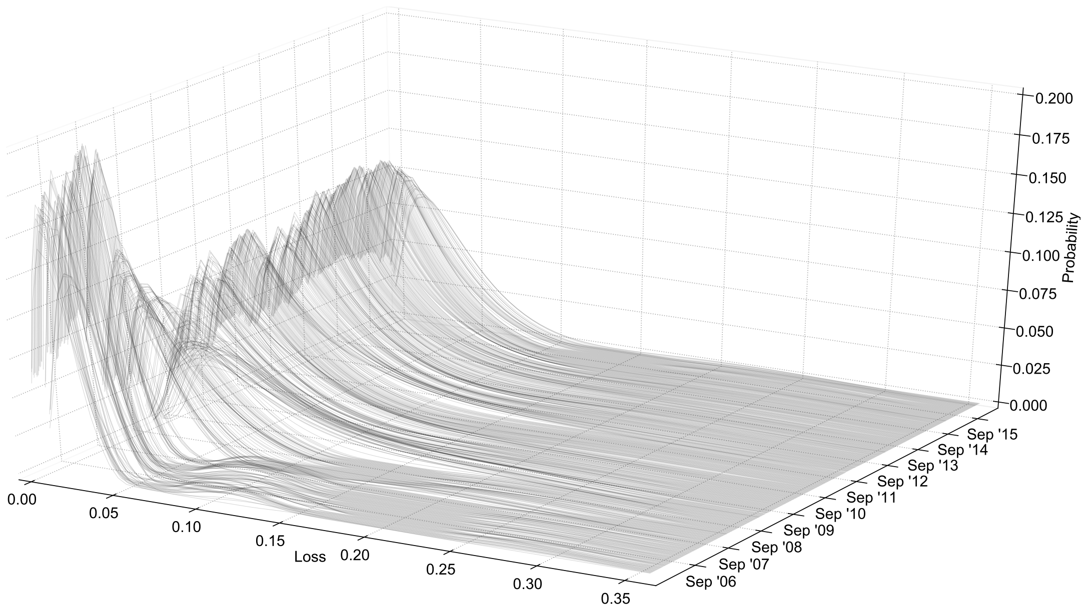

Figure 2.6: CDX credit portfolio loss distributions. This figure graphs the risk-neutral default loss distributions of the on-the-run five year CDX North American Investment Grade credit index for portfolio losses ranging from 0.00 to 0.35. The latent total loss distributions are recovered via Jacobi orthogonal series expansions around the distributional raw moments derived in Section 2.3.2 and represent smooth continuous approximations to the true discrete loss distributions. We use weekly averages of the daily implied model-free default loss moments with a five year horizon. The sample is from September 20, 2005 to February 26, 2016.

the risk can be virtually decomposed into idiosyncratic up to around 5% and systemic risk, up to around 15%. The probabilities for losses over this level are nearly zero. Our pre-crisis implied loss distributions exhibit great similarities to those of the three-jump model from Longstaff and Rajan (2008).

With the outbreak of the financial crisis in 2007 the shape of the implied loss distribution significantly changes. The loss distributions become more unimodal and the probability for high losses substantially increases. Coval et al. (2014) conclude that the prior non-observance of possible state-contingent model errors is the main reason for these changes after this "learning" event. In order to fit the senior tranche spreads after the start of the financial crisis, Collin-Dufresne et al. (2012) add catastrophic jump parameter in their model specification. This also indicates a pronounced change in the market beliefs about future losses in correlated credit markets.

After the financial crisis the shape of the loss distributions remains mainly unimodal, but the probability for high portfolio losses is much greater than pre-crisis. This clearly shows that the market participants learned from the financial crisis and adjusted their risk assessments.

2.4.2 Default correlations

For tranched CDS markets, default correlation determines the proportion of total portfolio risk that can be attributed to each tranche of the capital structure, or by the same token, the relation of a tranche spread to the index spread. This means a liquid market in tranched portfolio credit derivatives is tantamount to a liquid market in correlated default risk and provides unique insights into market embedded risk perceptions or risk preferences for different parts of the default loss distribution.

Speaking of default correlation, we are particularly interested in modeling the phenomenon that the likelihood of one reference asset defaulting on its credit obligations is affected by the contemporaneous defaults of other reference assets.[28] To model this default correlation, we consider a Bernoulli or one trial binomial correlation measure. In fact, the portfolio loss process given in (2.1) models the binomial events $\mathbb{I}_{\{\tau_i \leq T\}}$, where τ_i is the random default time of firm i. The indicator function takes the value of one, if the entity defaults before maturity T, and zero otherwise. We write the risk-neutral probability of default for firm i viewed from time 0 over time T as

$$\mathbb{E}\left[\mathbb{I}_{\{\tau_i \leq T\}}\right] = 1 - Q(\tau_i \leq T) = \pi_i$$

where the risk-neutral survival probabilities $Q(\tau_i \leq T)$ are determined by calibrating them to their individual CDS spreads. The variance of a Bernoulli binomial event that firm i will default or not is $\pi_i(1-\pi_i)$. Furthermore, let $\mathbb{E}\left[\mathbb{I}_{\{\tau_i \leq T\}}\mathbb{I}_{\{\tau_j \leq T\}}\right] = 1 - Q(\tau_i \leq T \cap \tau_j \leq T) = \pi_{ij}$ be the joint probability of default for two firms i and j. Hence, the joint (linear) default dependence coefficient of $\mathbb{I}_{\{\tau_i \leq T\}}$ and $\mathbb{I}_{\{\tau_j \leq T\}}$ is

$$\rho\left(\mathbb{I}_{\{\tau_i \leq T\}}, \mathbb{I}_{\{\tau_j \leq T\}}\right) = \rho_{ij} = \frac{\pi_{ij} - \pi_i \pi_j}{\sqrt{\pi_i(1-\pi_i)}\sqrt{\pi_j(1-\pi_j)}} \tag{2.9}$$

where the term $\pi_{ij} - \pi_i \pi_j$ denotes the covariance between two firms i and j, i.e., $\mathrm{Cov}\left(\mathbb{I}_{\{\tau_i \leq T\}}, \mathbb{I}_{\{\tau_j \leq T\}}\right)$. Next, we know that for any linear combination

[28] For an overview of default correlation and empirical evidence of historical default correlations, compare the discussions provided by Lucas (1995) and Albanese et al. (2013).

$Z = \alpha_0 + \sum_{i=1}^{N} \alpha_i X_i$, the variance is given by

$$\mathrm{Var}(Z) = \sum_{i=1}^{N} \alpha_i^2 \,\mathrm{Var}(X_i) + 2 \sum_{i,i<j} \sum_{j} \alpha_i \alpha_j \,\mathrm{Cov}(X_i, X_j)$$

In order to understand the relationship between the cross-section of the individual CDX index constituents to the index, let $\mathrm{Var}(Z)$ be the model-free risk-neutral variance of the CDX credit index loss distribution $\mathrm{Var}\,(L_T)$ and X_i be the probabilities of default π_i. Using (2.9) this then results in

$$\mathrm{Var}\,(L_T) = \sum_{i=1}^{N} \alpha_i^2 \pi_i (1-\pi_i) + 2 \sum_{i,i<j} \sum_{j} \alpha_i \alpha_j \rho_{ij} \sqrt{\pi_i (1-\pi_i)} \sqrt{\pi_j (1-\pi_j)}$$

In order to proceed, we implicitly assume that the expected loss of the sum of individual CDS losses equals the expected CDX index total loss. Without frictions in the market, this relation is a valid assumption. However, due to liquidity constraints, limits to arbitrage or other transaction related issues on either on the side of the CDS index market or the individual CDS market, this equality may be disturbed. Therefore, we adjust the individual CDS spreads not only to match the maturity of the CDX, but also to control for the CDX–CDS index basis (henceforth referred to as the CDX basis), see O'Kane (2011) for details. Furthermore, while contractual terms specify an assumed recovery rate of 40% for the CDX index, we have to consider the firm i specific consensus expected recovery rates R_i for the individual CDS. Given that in the CDX index there are $N = 125$ names, one default reduces the notional by $\frac{1}{125} = 0.008$. Thus, we set the loss fraction for firm i to $\alpha_i = 0.008 \times (1 - R_i)$. In addition, given that all pairwise default correlations are identical, $\rho_{ij} := \rho$, we finally have

$$\rho = \frac{\mathrm{Var}\,(L_T) - \sum_{i=1}^{N} \alpha_i^2 \pi_i (1-\pi_i)}{2 \sum_{i,i<j} \sum_j \alpha_i \alpha_j \sqrt{\pi_i (1-\pi_i)} \sqrt{\pi_j (1-\pi_j)}}$$

For an empirical assessment of market implied risk-neutral default correlations, we collect CDS spread level and market consensus recovery rate data for the on-the-run CDX series constituents from Markit. The overall minimum (maximum) of the single-name recovery rates in the sample is at 17.12% (56.75%), both the mean and median stay close to 40%. Figure 2.7 plots the historical evolution of the estimated risk-neutral default correlations.

The black series denotes the implied default correlations adjusted for the CDX basis. Prior to the financial crisis, the default correlations stay relatively

Figure 2.7: CDX default correlations. This figure plots the implied risk-neutral default correlations in the cross-section of the on-the-run five year CDX North American Investment Grade credit index. Using the implied model-free default loss variance of the CDX index, we infer the pairwise constant Bernoulli binomial correlation measures by equating the index variance to the individual names in the underlying reference pool. The black (gray) time series depicts the default correlations that have (not) been adjusted for the CDX–CDS index basis. Index roll dates are indicated via the vertical grid lines. The sample is from September 20, 2005 to February 26, 2016.

low around 10% with a minimum below 7% and then increase dramatically up to a maximum of 82%.[29] This high level of expected default correlation is confirmed by anecdotal evidence from the simultaneous default event of the Federal National Mortgage Association (Fannie Mae) and the Federal Home Loan Mortgage Corporation (Freddie Mac) on 09/07/2008. In the aftermath of the crisis, the default correlations appear to level off to a new regime hovering

[29]Following a similar approach Longstaff and Rajan (2008) equate the implied default loss variance from their calibrated three-jump model to the single-name constituents of the CDX.NA.IG series 1 through series 5. Their time series of default correlations are implied under the simplifying assumption of a constant recovery rate of 50% for the reference assets in the index. Furthermore, the authors do not explicitly assume a recovery rate for the index, the recoveries are modeled separately for each series and implicit within their pricing approach. The resulting values for the recoveries range from 48.8% to 58.6% for the five CDX series under consideration. Keeping these simplifying assumptions and model based idiosyncrasies in mind, our time series of default correlations might be considered a direct model-free continuation of the time series depicted in their Figure 5 as the numbers roughly share the same magnitudes.

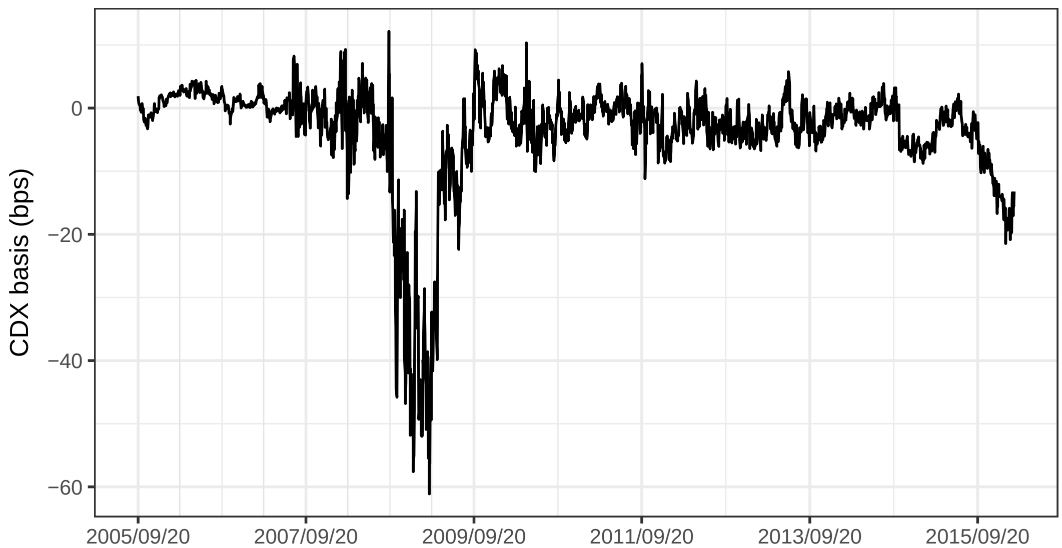

Figure 2.8: CDX–CDS index basis. This figure graphs the time series of the daily CDX–CDS index basis defined as difference of the CDX index level to the theoretical net asset value of the single-name constituents in the underlying reference pool. The CDX basis is measured in basis points. Index roll dates are indicated via the vertical grid lines. The sample is from September 20, 2005 to February 26, 2016.

around 40% to 50%, indicating that market participants may have reassessed the risk-neutral probability of two firms defaulting simultaneously.

As mentioned above, the CDX basis plotted in Figure 2.8 may have a pronounced effect on the calculation of the default correlations. Therefore, we additionally graph the default correlations (in gray) in Figure 2.7, not adjusting for the CDX basis. As is visible in the graph, the two series deviate considerably during times where the (absolute) CDX basis indicates strong limits of CDX index to CDS constituents arbitrage mechanisms, liquidity issues or market supply and demand effects.

Overall, our results sharpen the common intuition that a higher variance of the total loss distribution is directly related to a higher pairwise default correlation in the names of the underlying reference pool and lowers potential diversification benefits. In conclusion, the findings may shed some light on the understanding of the important issue of implicit name dependence structure and the degree of default clustering in CDS markets. Considering observable risk factors and frailty, Nickerson and Griffin (2017) report average pairwise bond default correlations around 10% pre-crisis which increase by roughly 25% up to 12.5% incorporating information from the crisis period. Contrasting their real-world estimates to our risk-neutral default correlations suggest a substantial correlation risk premium.[30] The high degree of risk-neutral correlation risk, though, may be mainly attributed to liquidity related issues rather than being a reflection of fundamental correlation risk *per se*.[31] Further indications may be found in Azizpour et al. (2018). Breaking down the sources of corporate default clustering the authors show that, besides idiosyncratic risk and joint exposure to known or unknown systematic risk factors, contagion effects play a major role in explaining clustered default risk premiums after controlling for expectation components. In addition, the risk premium may be positively related to increasing ambiguity surrounding the state of the underlying economy as shown in Benzoni et al. (2015). The implications of default correlation and correlation risk premia are left for future research.

2.4.3 Determinants of market implied credit portfolio loss distribution moments

In this section, we regress our four moments on different risk proxies to determine the reasons for their daily innovations. In contrast to Collin-Dufresne et al. (2001), Blanco et al. (2005), Zhang et al. (2009) and others, who analyze individual CDS or bond spreads, our estimated moments enable us to examine the relationships between the market implied moments and risk proxies from different markets on an aggregated basis. As in Azizpour et al. (2011), who analyze the pricing of risk in the CDX market, we use, *inter alia*, the CBOE VIX as proxy for the risk at the equity market and a bond yield spread for the default risk premium at the bond market.

We fit the daily logarithmic differences of the implied expected loss, variance, skewness and kurtosis instead of the levels to control for the serial correlation of the endogenous variables, since all endogenous variables exhibit unit roots. Specifically, we fit the log changes via univariate first-order autoregressive models and apply robust standard errors to yield conclusions on statistical significance. We analyze financial variables from the credit, equity and fixed income markets to explain our moments and control for the role date effect via a dummy.

Table 2.3 presents the resulting parameter estimates for the univariate AR(1) models. All exogenous variables, except for the CBOE VIX for the

[30] Credit rating agencies' assumptions of default correlation for structured products have been around 1% pre-crisis and increased up to 3% post-crisis which would indicate an even more pronounced correlation risk premium, see Nickerson and Griffin (2017).

[31] See, e.g., Acharya et al. (2015) for a thorough discussion on linking funding liquidity, market liquidity and liquidity risk to correlation risk with a view on the May 2005 automobile correlation crisis.

Table 2.3: Univariate time series regressions for the logarithmic differences of the CDX credit portfolio loss distribution moments. This table reports the estimated coefficients and adjusted R^2 of each AR(1) model for the model-free risk-neutral expected loss, loss variance, skewness, and kurtosis. The adjusted R^2 for all regressions with a single exogenous variable is given in parentheses. All variables are in log differences but the CDX basis, where we consider the first differences.

Variable	$\Delta\mathbb{E}$	ΔVar	ΔSkew	ΔKurt
Lagged endogenous variable	0.128**	0.037**	0.121**	0.136**
	(0.819%)	(0.372%)	(0.253%)	(0.345%)
CDX basis	0.680**	1.126**	−0.603**	−1.219**
	(46.907%)	(37.331%)	(48.635%)	(47.932%)
CDX volatility skewness	0.017**	0.018**	−0.011**	−0.021**
	(28.233%)	(20.283%)	(27.372%)	(26.349%)
US corporate high vs. BBB yield	0.345**	0.403**	−0.228**	−0.454**
	(26.269%)	(16.867%)	(23.875%)	(22.619%)
US financial sector CDS index	0.116**	0.110**	−0.081**	−0.151**
	(18.809%)	(13.426%)	(19.937%)	(18.614%)
CBOE VIX	0.074**	0.038*	−0.039**	−0.067**
	(41.458%)	(25.044%)	(36.744%)	(34.753%)
Adjusted R^2	72.160%	46.659%	64.291%	62.407%

*, and ** denote significance at the 5%, and 1% levels, respectively

second moment, are statistically significant at a level of 1% and the estimated parameters show the expected signs.

The contribution of the lagged endogenous variables to the adjusted R^2 is not relevant for any of the moments, which means that the changes of the previous moment have no explanatory power for the actual moment changes.

The most important exogenous variable for all four moments is the CDX basis. According to the discussion in Junge and Trolle (2015), the CDX basis is a measure of the liquidity between the CDX and CDS markets and a measure of the limits to arbitrage between these two markets. We calculate the CDX basis liquidity measure as the relative difference of the CDX index spread to the constituent single-name CDS implied index level. As can be seen in Figure 2.8, the CDX basis becomes extremely negative during the financial crisis. Boyarchenko et al. (2016) confirm that changes of liquidity preferences and liquidity concentration have the strongest impact on the widening of the CDX basis during the second half of 2015 and the first quarter of 2016. We estimate a positive sign for the first two moments and a negative for the third and fourth, which means that a less negative CDX basis is related to more

Table 2.4: Correlations of the exogenous variables. This table reports the correlations for all considered exogenous variables of the AR(1) models. All variables are in log differences but the CDX basis, where we consider the first differences.

Variable	(1)	(2)	(3)	(4)	(5)
CDX basis (1)	1.000	0.330	0.083	0.095	0.498
CDX volatility skewness (2)		1.000	0.236	0.174	0.341
US corporate high vs. BBB yield (3)			1.000	0.380	0.406
US financial sector CDS index (4)				1.000	0.276
CBOE VIX (5)					1.000

implied risk in the CDX market. The CDX basis also represents the influence of arbitrage trading activities respectively limitations on the CDX market.

As another risk measure for the credit market we use the realized CDX volatility skewness. This measure is defined as the realized upside to downside variation of the CDX index spread returns. In good times, the downside variation, measured as the sum of all quadratic variations of negative index returns in the last month (lower risk), dominates the upside variation and, as a consequence, the CDX volatility skewness is relatively low. The opposite is true for bad times. The realized CDX volatility skewness is relevant for all moments, especially for the first one.

As a proxy for the default risk premium at the bond market, we use the difference between high-yield and BBB effective yields indices of US corporate bonds from Bank of America Merrill Lynch. Our analysis shows that this risk proxy is relevant for all four moments and the explanatory power is slightly smaller to that of the CDX volatility skewness.

The changes of five year US financial sector credit default swap index, as a proxy of the overall credit risk in the financial sector has only moderate explanatory power.

The CBOE VIX index, as a proxy for fear at the equity market, have only small explanatory power, when we consider all five exogenous variables in our autoregressive models. Surprisingly, the CBOE VIX index is not significant for the second moment at a level of 1%. In a univariate setting, the adjusted R^2 of the CBOE VIX index is the second highest for all four moments. As can be seen in Table 2.4 the changes of the CBOE VIX index is the risk factor which has the highest correlations to all other considered risk factors. In consequence, the explanatory power of this risk proxy is reduced in our autoregressive models with five exogenous variables. Another explanation could be that the equity market recovers much faster after the financial crisis than the credit market, therefore, risk proxies from the credit market are potentially more important.

In addition to the five risk proxies, we also consider Fama–French factors as measures for equity market risk, and the TED and LIBOR–OIS spreads as common measures for liquidity risk. These risk factors have no relevant additional explanatory power.

The influence of the exogenous variables is particularly relevant for the expected loss with an adjusted R^2 over 72%. In contrast to this, the adjusted R^2 of the AR(1) model for the log changes of the variance is substantially smaller than for all other moments.

2.5 Robustness analyses

In this section, we conduct a number of additional robustness tests to verify the validity of our approaches taken in Section 2.3.2. Specifically, we investigate alternative interpolating functions to complete the set of available strikes and the choice of objective loss functions to calibrate expected tranche losses.

2.5.1 Interpolating spline functions

Market data of tranched index spreads is limited to a finite number of strikes along the loss dimension. Consequently, we need to complete the set of available strikes via some interpolation procedure to compute the integrals in Proposition 2.1. Here, we discuss the impact of alternative interpolating spline functions relative to the Schumaker (1983) shape preserving quadratic benchmark spline (SPQS) which is at least $\mathscr{C}^1$ differentiable. Guided by the no arbitrage requirements discussed earlier, we consider as an alternative the shape preserving rational spline function (SPRS) by Cai and Judd (2012) which is $\mathscr{C}^\infty$ on each interval and at least $\mathscr{C}^1$ globally and ensures monotonicity and concavity if the sample points fulfill these criteria. In order to assess the implications of less restrictive spline functions, we further investigate the popular local monotone cubic spline (LMCS) by Fritsch and Butland (1984) which is at least $\mathscr{C}^1$ and ensures monotonicity as well as a plain vanilla cubic spline (CS) which is $\mathscr{C}^2$ and ensures neither monotonicity nor concavity. The SPQS, SPRS and LMCS are all local, i.e, the piecewise constant spline coefficients are determined solely by the information given by the respective subintervals being unaware of data farther away and involve only analytical expressions. On the other hand the CS is global in nature, is determined by all available data and requires solving a tridiagonal system of equations.

In Figure 2.9, we plot the resulting curves of the different spline functions under consideration given the same data points as in Figure 2.3. We note that

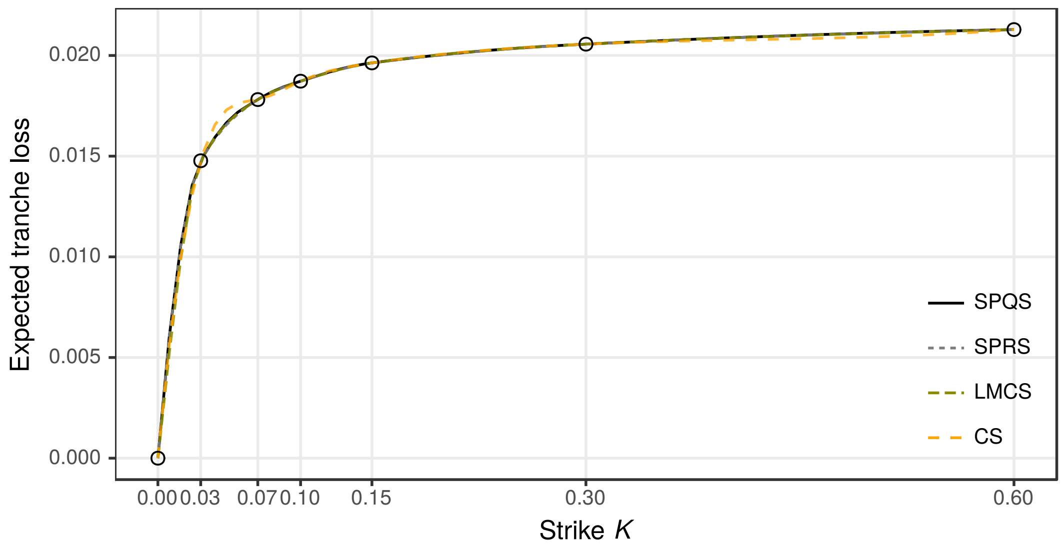

Figure 2.9: Alternative interpolating spline functions. In addition to the Schumaker (1983) benchmark spline (SPQS) this figure graphs the shape preserving rational spline (SPRS), the local monotone cubic spline (LMCS) and the plain vanilla cubic spline (CS) functions. The sample points are as in Figure 2.3.

the monotonicity (LMCS) and shape preserving (SPQS and SPRS) splines are hardly distinguishable by visual inspection. The cubic spline is incapable of matching the curvature of the sample expected tranche losses without overshooting and thus will violate basic no arbitrage conditions. We check the no arbitrage conditions by inspecting the first and second derivatives of the spline functions. All first derivatives are globally positive which indicates that all splines are globally monotone for the given data set. The second derivatives happen to be globally negative for the SPQS, SPRS and LMCS splines resulting in globally concave curves. This behavior is expected for the SPQS and SPRS, for the LMCS spline this behavior is not guaranteed though. The CS spline reveals changing signs for the second derivatives and, in consequence, is not obeying the concavity restriction.

Next, we employ the four splines and calculate the first four corresponding loss distribution moments. Table 2.5 summarizes the resulting figures relative to the Schumaker (1983) spline which is expressed in levels. We can see that the absolute relative error of the SPRS is not exceeding 0.008, thus the resulting moments are in close alignment to the benchmark spline. The monotonicity preserving LMCS is deviating slightly more than the SPRS, but is still quite close to the SPQS. Unsurprisingly, the cubic spline is way off and produces unreliable estimates for the loss distribution moments.

Table 2.5: Relative deviations from the benchmark spline. Entries reports the relative errors of the alternative shape preserving rational spline (SPRS), the local monotone cubic spline (LMCS) and the plain vanilla cubic spline (CS) relative to the benchmark Schumaker (1983) spline function (SPQS). The entries for the SPQS are expressed in levels. The relative error is defined as $\frac{\text{Alternative}-\text{SPQS}}{\text{SPQS}}$.

	SPQS (levels)	SPRS	LMCS	CS
$\mathbb{E}$	0.021	–	–	–
Var	0.002	0.005	0.022	0.038
Skew	8.611	−0.004	−0.022	0.107
Kurt	106.592	−0.008	−0.020	0.241

Overall the results indicate that as long as the no arbitrage constraints are met, the resulting moments do not deviate too much. This clearly excludes the CS spline. Though the LMCS compares relatively well, it is by construction not guaranteed to obey the additional concavity constraint. In conclusion, the best results are ensured by employing spline functions which adhere to both the monotonicity and concavity constraints.

2.5.2 Objective loss functions

In Equation (2.6), we minimize the sum of squared errors (SSE) of the tranched and untranched index spreads $s_{0,T}^{Market}(A,D)$ to the spreads implied by the corresponding expected tranche losses $s_{0,T}^{\psi}(A,D)$. This natural choice of objective loss function is of course only one amongst many other possible options.[32] Here, we assess the impact of choosing different objective functions on implied loss distribution moments. Specifically, we consider three alternatives in addition to the benchmark objective function in (2.6). These are discussed in detail as follows.

Relative errors In Equation (2.6) the calibration instruments have decreasing weights from the index spread up to the most senior tranche. This means that the nonlinear least square regression is minimizing spread differences with varying magnitudes. To address this issue a common alternative is to give the calibration instruments roughly equal weights by minimizing the relative spread differences, i.e., the weights in (2.6) are then given as $w_{A,D} = \frac{1}{s_{0,T}^{Market}(A,D)^2}$.

[32] The importance of different loss function in option pricing is substantiated in, e.g., Christoffersen and Jacobs (2004).

Table 2.6: Correlations of the benchmark moments to alternative objective loss functions. This table reports the correlations of the resulting default loss moments from the alternative objective loss functions with respect to the benchmark calibration in Equation (2.6).

	$\mathbb{E}$	Var	Skew	Kurt
Relative errors	1.000	0.997	1.000	1.000
Index only	0.999	0.985	1.000	1.000
Super senior only	1.000	0.997	1.000	1.000

Index only With the inception of CDX series 9, super senior quotes become available. This completes the capital structure and the [0–100%] expected tranche loss now must reprice the CDX index spread as well as the [30–100%] super senior tranche (or the [15–100%] tranche from CDX series 15 onwards). Typically, the untranched CDX index is traded much more heavily than the end of the capital structure. In cases of market frictions and limits to arbitrage, the two spreads might carry different market expectations which cannot be handled with a single [0–100%] expected tranche loss. To examine this issue, we continue the approach from the CDX series 5 to 8 and do only consider the index and tranche spreads but the super senior even though it may be available later on.

Super senior only In contrast to the index only approach, we also consider the case where we omit the index spread if the super senior tranche is present.

The influences of different objective loss functions on the default loss moments are depicted in Figure 2.10. Overall, we find that the resulting moments are hardly distinguishable from each other. From Table 2.6, this is backed by extremely high correlations of the default loss moments with respect to the benchmark calibration in Equation (2.6).

Comparing the first four moments, the loss variance seems to be a little sensitive regarding the choice of objective function. The relative errors and super senior only calibration are quite similar and apparently extract essentially the same information. Further, the resulting variance estimates tend to be slightly higher than the remaining approaches. The variance estimates resulting from the index only calibration tend to be on the lower end. The benchmark calibration is mostly in between the other approaches.

While, for most of the time during our sample period, the alternatives exhibit virtually synchronous behavior with the benchmark approach, we observe a temporary departure from this close alignment around early 2009

Figure 2.10: CDX credit portfolio loss distribution moments with alternative objective loss functions. This figure graphs the daily implied model-free risk-neutral expected loss, loss variance, skewness, and kurtosis of the five year tranched and untranched CDX North American Investment Grade credit index for four different objective loss functions. Besides the benchmark SSE calibration, we plot the resulting moments for the relative error, the index only and the super senior only (if available) objective functions. Index roll dates are indicated via the vertical grid lines. The sample is from September 20, 2005 to February 26, 2016.

when the crisis seems to slowly settle down. This might indicate that during this phase of extremely high uncertainty surrounding the overall state of the economy, the index and the super senior tranche are not well aligned, possibly revealing temporary market frictions and limits to arbitrage. A further reason may be traced back to a disagreement about default and correlation risk (Broer, 2018). There seems to be little disagreement about the default rates in the underlying reference pool as per the first moment. The second moment clearly exhibits temporarily diverging behavior, which might be a reflection of differences in beliefs of low correlation vs. high correlation investors.

2.6 Conclusion

We provide a novel methodology to extract model-free risk-neutral expectations of the latent credit portfolio loss distribution contained in credit index spreads. This new approach has the advantage of disclosing the information content of portfolio credit derivatives free from potential model misspecification, model bias or parameter risk. This research closes an important gap in the literature on model-free extraction of information embedded in the derivatives markets and contributes to the empirical literature on correlated default risk, which has been predominantly model based.

Our empirical analysis of the CDX North American Investment Grade credit index reveals a number of important findings. The aggregate view on expected losses in the pool has been negligible before the onset of the financial crisis. During the heights of the crisis market participants expect severe losses with peaks reaching 13%. Post-crisis, these expectations revert to levels slightly higher than pre-crisis. For the second moment of the loss distribution, we document a persistent change in the risk assessment of correlated default risk. The loss variance increases dramatically during the crisis, but contrary to the expected loss, the variance remains at a much higher levels afterwards. This is paralleled by the time evolution of the skewness and kurtosis which drops sharply to account for the persistent increase in expected losses and default variance.

Furthermore, we show three first empirical applications. First, the model-free default loss moments allow us to infer the total latent risk-neutral loss distribution. The resulting distributions reveal a much more complete picture of the aggregate portfolio default risk than the moments alone. In particular, with widening loss distributions it is evident that investors expect higher losses in more senior parts of the capital structure as had previously been the case. Before the crisis, the main source of risk had been perceived as idiosyncratic risk. The model-free default loss distributions could serve as important indicators for policy makers and regulatory authorities to inform monetary policy and the management of business cycles.

Second, in relating the single-name constituents to the index variance, we imply model-free risk-neutral default correlations. These default correlations may serve as a useful additional tool to inquire into the very nature of correlated default risk and correlation risk premia.

Third, the model-free loss moments may be related to common observable risk factors from the equity, credit and fixed income markets. A natural extension could be to consider more elaborate dependency structures and examine the moments in a multivariate setting.

Given these promising results, the model-free methodology provides a valuable and versatile tool to imply market expectations from correlation dependent credit derivatives. It may serve investors in debt markets, sellers and buyers of credit protection, and highly sophisticated swap dealer or hedge funds to refine their investment, hedging and risk management strategies. Furthermore, this tool allows policy makers to learn about the economic state of aggregate corporate debt on a daily basis, without being subject to model risk.

Our study leaves room for further research. One important issue is the relation to *ex ante* real-world default loss expectations and *ex post* realized losses in the reference pool. This may provide further essential insights into the structure of the default pricing kernel or stochastic discount factor in credit markets. Other natural venues for future research are empirical applications to high-yield CDS index markets or cross market linkages.

2.A Proofs

2.A.1 Proof of Proposition 2.1

For the proof of Proposition 2.1, we pursue two completely interchangeable approaches. For one, we employ the generic spanning approach outlined in Carr and Madan (1998). For the second approach, we build on the insight first developed in Bakshi and Madan (2000) in that arbitrary payoff functions may equally be spanned by characteristic functions or moment generating functions. As is shown, both approaches are in fact equivalent, since they both derive via the Breeden and Litzenberger (1978) argument.

Proof of Proposition 2.1: The spanning approach. Let $f(L_T)$ denote a smooth payoff function of the terminal portfolio loss L_T. By the sifting and symmetry properties of a Dirac delta function, $\delta(.)$, we have

$$\begin{aligned} f(L_T) &= \int_0^{K_{max}} f(K)\delta(L_T - K)\,dK \\ &= \int_0^{\kappa} f(K)\delta(L_T - K)\,dK + \int_{\kappa}^{K_{max}} f(K)\delta(L_T - K)\,dK \\ &= \int_0^{\kappa} f(K)\delta(K - L_T)\,dK + \int_{\kappa}^{K_{max}} f(K)\delta(L_T - K)\,dK \end{aligned}$$

with some nonnegative separation strike κ. Integration by parts yields

$$f(L_T) = f(K)\mathbb{I}_{\{L_T<K\}}\Big|_0^{\kappa} - \int_0^{\kappa} f'(K)\mathbb{I}_{\{L_T<K\}}\,dK + f(K)\mathbb{I}_{\{L_T\geq K\}}\Big|_{\kappa}^{K_{max}} - \int_{\kappa}^{K_{max}} f'(K)\mathbb{I}_{\{L_T\geq K\}}\,dK$$

$$=f(\kappa)\mathbb{I}_{\{L_T<\kappa\}}-\int_0^\kappa f'(K)\mathbb{I}_{\{L_T<K\}}\,dK+f(\kappa)\mathbb{I}_{\{L_T\geq\kappa\}}-\int_\kappa^{K_{max}} f'(K)\mathbb{I}_{\{L_T\geq K\}}\,dK$$

using the two facts that the loss L_T can never go below zero or exceed the maximum loss at K_{max}. Integrating by parts once again implies

$$\begin{aligned}
f(L_T) &= f(\kappa)\mathbb{I}_{\{L_T<\kappa\}} - f'(K)(K-L_T)^+\Big|_0^\kappa + \int_0^\kappa f''(K)(K-L_T)^+\,dK \\
&\quad + f(\kappa)\mathbb{I}_{\{L_T\geq\kappa\}} - f'(K)(L_T-K)^+\Big|_\kappa^{K_{max}} + \int_\kappa^{K_{max}} f''(K)(L_T-K)^+\,dK \\
&= f(\kappa)\left[\mathbb{I}_{\{L_T<\kappa\}}+\mathbb{I}_{\{L_T\geq\kappa\}}\right] + f'(\kappa)\left[(L_T-\kappa)^+-(\kappa-L_T)^+\right] \\
&\quad + \int_0^\kappa f''(K)(K-L_T)^+\,dK + \int_\kappa^{K_{max}} f''(K)(L_T-K)^+\,dK \\
&= f(\kappa)+f'(\kappa)(L_T-\kappa) + \int_0^\kappa f''(K)(K-L_T)^+\,dK + \int_\kappa^{K_{max}} f''(K)(L_T-K)^+\,dK
\end{aligned} \tag{2.10}$$

Next, we assume the separation strike κ to take the value zero which is where the value of the payoff function vanishes. Remember that for every new trading day possible defaults in the credit index, i.e., $L_t > 0$, are actually reflected in adjusted attachment and detachment points. Subsequently, we may assume $\kappa = 0$ for all t until maturity (that is until the credit index is completely wiped out).

Now, as a first case, we consider the identity payoff function $f(L_T) := L_T$ with derivatives $f'(L_T) = 1$ and $f''(L_T) = 0$. This leads to

$$\begin{aligned}
f(L_T) &= \underbrace{f(0)}_{=0} + \underbrace{f'(0)(L_T-0)}_{=1\times L_T} + \int_0^{K_{max}} 0\times(L_T-K)^+\,dK \\
&= L_T
\end{aligned}$$

Taking expectations, we get

$$\mathbb{E}_t[L_T]=\mathbb{E}_t[L_T]$$

which is, of course, by construction.

Next, we consider the quadratic loss payoff $f(L_T) := L_T^2$ with derivatives $f'(L_T) = 2L_T$ and $f''(L_T) = 2$. This leads to

$$f(L_T) = \underbrace{f(0)}_{=0} + \underbrace{f'(0)(L_T-0)}_{=0\times L_T} + \int_0^{K_{max}} 2(L_T-K)^+\,dK$$

Here, up to now, we are exclusively dealing with plain vanilla (put and) call option payoffs. Hence, since we model the market quotes in the form of (equity) expected tranche losses, we use the relation

$$\min(L_T, K) = L_T - (L_T - K)^+$$

and substitute the call payoffs accordingly. Thus,

$$\begin{aligned} f(L_T) &= \int_0^{K_{max}} 2\Big[L_T - \min(L_T, K)\Big]\, dK \\ &= \int_0^{K_{max}} 2L_T - 2\min(L_T, K)\, dK \\ &= 2L_T \int_0^{K_{max}} dK - 2\int_0^{K_{max}} \min(L_T, K)\, dK \end{aligned}$$

Taking expectations and remembering that $\mathbb{E}_t\Big[\min(L_T, K)\Big] = \mathbb{E}_t[L_T] - \mathbb{E}_t\Big[(L_T - K)^+\Big] = \psi_{t,T}(K)$, we finally get

$$\mathbb{E}_t\Big[L_T^2\Big] = 2\,\mathbb{E}_t\,[L_T]\,K_{max} - 2\int_0^{K_{max}} \psi_{t,T}(K)\, dK$$

For the cubic loss payoff $f(L_T) := L_T^3$ with derivatives $f'(L_T) = 3L_T^2$ and $f''(L_T) = 6L_T$, we have

$$\begin{aligned} f(L_T) &= \int_0^{K_{max}} 6K(L_T - K)^+\, dK \\ &= 6L_T \int_0^{K_{max}} K\, dK - 6\int_0^{K_{max}} K \min(L_T, K)\, dK \\ &= 3L_T K_{max}^2 - 6\int_0^{K_{max}} K \min(L_T, K)\, dK \end{aligned}$$

or

$$\mathbb{E}_t\Big[L_T^3\Big] = 3\,\mathbb{E}_t[L_T]K_{max}^2 - 6\int_0^{K_{max}} K\psi_{t,T}(K)\, dK$$

Finally, considering the quartic loss payoff $f(L_T) := L_T^4$ with derivatives $f'(L_T) = 4L_T^3$ and $f''(L_T) = 12L_T^2$, we have

$$f(L_T) = \int_0^{K_{max}} 12K^2(L_T - K)^+\, dK$$

$$= 12L_T \int_0^{K_{max}} K^2\, dK - 12 \int_0^{K_{max}} K^2 \min(L_T, K)\, dK$$
$$= 4L_T K_{max}^3 - 12 \int_0^{K_{max}} K^2 \min(L_T, K)\, dK$$

or

$$\mathbb{E}_t\left[L_T^4\right] = 4\,\mathbb{E}_t[L_T] K_{max}^3 - 12 \int_0^{K_{max}} K^2 \psi_{t,T}(K)\, dK$$

Thus, for any $n \geq 1$, we are led to the conjecture

$$\mathbb{E}_t\left[L_T^n\right] = n\,\mathbb{E}_t\left[L_T\right] K_{max}^{n-1} - n(n-1) \int_0^{K_{max}} K^{n-2}\, \psi_{t,T}(K)\, dK$$

□

Proof of Proposition 2.1: The moment generating function approach. Let the function

$$\mathcal{M}_L(\omega) = \mathbb{E}_t\left[e^{\omega L_T}\right] = \int_0^{L_{max}} e^{\omega \ell}\, dF(\ell)$$

if it exists over an interval of ω about the origin on the real line, denote the moment generating function of the portfolio loss L_T. Note, for functions of the portfolio loss $f(L_T)$, the moment generating function would simply be $\mathbb{E}_t\left[e^{\omega f(L_T)}\right]$. Here, we are concerned with the portfolio loss as such, hence we implicitly assume $f(L_T) := L_T$.

Unfortunately we have no knowledge about the latent density of the portfolio loss, but using Breeden and Litzenberger (1978) we have

$$\mathcal{M}_L(\omega) = \int_0^{K_{max}} e^{\omega K} \left(-\frac{\partial^2 \psi_{t,T}(K)}{\partial K^2}\right) dK \tag{2.11}$$

Additionally, we know the limiting behavior of the equity expected tranche losses

$$\lim_{K \to 0} \psi_{t,T}(K) = 0 \qquad \lim_{K \to K_{max}} \psi_{t,T}(K) = \mathbb{E}_t\left[L_T\right]$$
$$\lim_{K \to 0} \frac{\partial \psi_{t,T}(K)}{\partial K} = 1 \qquad \lim_{K \to K_{max}} \frac{\partial \psi_{t,T}(K)}{\partial K} = 0$$

Applying integration by parts to Equation (2.11) yields

$$\mathcal{M}_L(\omega) = -e^{\omega K}\, \frac{\partial \psi_{t,T}(K)}{\partial K}\bigg|_0^{K_{max}} + \omega \int_0^{K_{max}} e^{\omega K} \frac{\partial \psi_{t,T}(K)}{\partial K}\, dK$$

$$= 1 + \omega \int_0^{K_{max}} e^{\omega K} \frac{\partial \psi_{t,T}(K)}{\partial K} dK$$

Substituting the boundary terms and integrating by parts once again gives

$$\begin{aligned} \mathcal{M}_L(\omega) &= 1 + \omega e^{\omega K} \psi_{t,T}(K)\Big|_0^{K_{max}} - \omega^2 \int_0^{K_{max}} e^{\omega K} \psi_{t,T}(K)\, dK \\ &= 1 + \mathbb{E}_t\left[L_T\right] \omega e^{\omega K_{max}} - \omega^2 \int_0^{K_{max}} e^{\omega K} \psi_{t,T}(K)\, dK \qquad (2.12) \end{aligned}$$

Differentiating $\mathcal{M}_L(\omega)$ n times and evaluating at zero eventually yield the n-th raw (non-centered) moment of the loss distribution

$$\mathbb{E}_t\left[L_T^n\right] = \left. \frac{d^n \mathcal{M}_L(\omega)}{d\,\omega^n} \right|_{\omega=0}$$

Thus

$$\begin{aligned} \mathbb{E}_t\left[L_T\right] &= \mathbb{E}_t\left[L_T\right] \\ \mathbb{E}_t\left[L_T^2\right] &= 2\,\mathbb{E}_t\left[L_T\right] K_{max} - 2 \int_0^{K_{max}} \psi_{t,T}(K)\, dK \\ \mathbb{E}_t\left[L_T^3\right] &= 3\,\mathbb{E}_t\left[L_T\right] K_{max}^2 - 6 \int_0^{K_{max}} K\, \psi_{t,T}(K)\, dK \\ \mathbb{E}_t\left[L_T^4\right] &= 4\,\mathbb{E}_t\left[L_T\right] K_{max}^3 - 12 \int_0^{K_{max}} K^2\, \psi_{t,T}(K)\, dK \end{aligned}$$

or more generally

$$\mathbb{E}_t\left[L_T^n\right] = n\,\mathbb{E}_t\left[L_T\right] K_{max}^{n-1} - n(n-1) \int_0^{K_{max}} K^{n-2}\, \psi_{t,T}(K)\, dK$$

Here, we work with the moment generating function. One could of course equally span the characteristic function via a continuum of expected tranche losses. Defining the characteristic function and using Breeden and Litzenberger (1978) we have

$$\begin{aligned} \phi_L(\omega) &= \mathcal{F}\left[\varphi_{t,T}(L_T)\right](\omega) = \mathbb{E}_t\left[e^{i\omega L_T}\right] \\ &= \int_0^{K_{max}} e^{i\omega K} \left(-\frac{\partial^2 \psi_{t,T}(K)}{\partial K^2} \right) dK \end{aligned}$$

where $\mathcal{F}[.]$ denotes Fourier transform. Non-centered or raw moments then naturally follow via

$$\mathbb{E}_t\left[L_T^n\right] = i^{-n}\phi_L^{(n)}(0) = i^{-n}\left.\frac{d^n\phi_L(\omega)}{d\,\omega^n}\right|_{\omega=0}$$

A very useful property of characteristic functions is that they always exist, even when the probability density function or moment generating function may not exist in all cases. The derivation of power loss contracts using ϕ_L is virtually identical to the moment generating function approach and is thus omitted.

□

Remember that in empirical applications the time t expectation of the terminal time T portfolio loss, $\mathbb{E}_t[L_T]$, is inferred by fitting to available index and/or super senior tranche spreads, see Section 2.3.2 for details.

2.A.2 Proof of Equation (2.3)

Proof of Equation (2.3): Using the simple relation of call and put options (i.e., put-call parity) to the minimum function employed in the expected tranche losses

$$\begin{aligned}(K-L_T)^+ &= \underbrace{(L_T-K)^+}_{L_T-\min(L_T,K)} - L_T + K \\ &= K-\min(L_T,K)\end{aligned}$$

we may restate Equation (2.10) in terms of tranche losses

$$f(L_T) = f(\kappa) + f'(\kappa)(L_T-\kappa) + \int_0^{\kappa} f''(K)\Big[K-\min(L_T,K)\Big]\,dK + \int_{\kappa}^{K_{max}} f''(K)\Big[L_T-\min(L,K)\Big]\,dK$$

Given a continuum of expected tranche losses, no arbitrage and integrating by parts implies

$$\begin{aligned}\mathbb{E}_t\Big[f(L_T)\Big] = f(\kappa) &+ f'(\kappa)(\mathbb{E}_t[L_T]-\kappa) \\ &+ f'(K)\Big[K-\psi(K)\Big]\Big|_0^{\kappa} - \int_0^{\kappa} f'(K)\Big[1-\psi'(K)\Big]\,dK \\ &+ f'(K)\Big[\mathbb{E}_t[L_T]-\psi(K)\Big]\Big|_{\kappa}^{K_{max}} - \int_{\kappa}^{K_{max}} f'(K)\Big[-\psi'(K)\Big]\,dK\end{aligned}$$

where $\psi(K) := \psi_{t,T}(K)$ and $\psi'(K) := \frac{\partial \psi_{t,T}(K)}{\partial K}$ to simplify notation. Evaluating at the upper and lower points and substituting the limits of $\psi(K)$ the first order derivative terms outside the integrals cancel and yield to

$$\mathbb{E}_t\Big[f(L_T)\Big] = f(\kappa) - \int_0^{\kappa} f'(K)\Big[1 - \psi'(K)\Big]\, dK + \int_{\kappa}^{K_{max}} f'(K)\Big[\psi'(K)\Big]\, dK$$

Integrating by parts once again gives

$$\mathbb{E}_t\Big[f(L_T)\Big] = f(\kappa) - f(K)\Big[1 - \psi'(K)\Big]\Big|_0^{\kappa} + \int_0^{\kappa} f(K)\Big[-\psi''(K)\Big]\, dK + f(K)\,\psi'(K)\Big|_{\kappa}^{K_{max}} - \int_{\kappa}^{K_{max}} f(K)\psi''(K)\, dK$$

where $\psi''(K) := \frac{\partial^2 \psi_{t,T}(K)}{\partial K^2}$. Evaluating again and substituting the limits of $\psi'(K)$ the resulting terms outside the two integrals cancel and we are left with

$$\mathbb{E}_t\Big[f(L_T)\Big] = -\int_0^{\kappa} f(K)\psi''(K)\, dK - \int_{\kappa}^{K_{max}} f(K)\psi''(K)\, dK$$

which yields the result in (2.3)

$$\mathbb{E}_t\Big[f(L_T)\Big] = \int_0^{K_{max}} f(K)\Big(-\psi''(K)\Big)\, dK$$

□

2.B Data preparation for the CDX.NA.IG series 11–13

At the heights of the financial crisis the price discovery process for a number of index tranches is temporarily impaired for the on-the-run series 11 to 13. In contrast, tranche spreads for the off-the-run series 9 are available allover and are much more liquid than the corresponding on-the-run index tranches during this time period. Untranched index spreads for the on-the-run series, though, are available throughout and exhibit greater liquidity than the corresponding series 9 index spreads. To mitigate the issue of these non-trivial data gaps for the tranched on-the-run index, we use the available information from series 9 and translate the daily spread innovations to the respective on-the-run series. The idea is to fill the missing value of $s_{\text{on-the-run},t+1}$ at time $t+1$ with

$$\hat{s}_{\text{on-the-run},t+1} = s_{\text{on-the-run},t} \times \frac{s_{9,t+1}}{s_{9,t}} \times c$$

Table 2.7: Example of missing value handling. This table displays an illustrative example of the replacing algorithm for data gaps. The missing values of the on-the-run series 11 are replaced with and without constant c. All calculations in this paper are performed with the latter.

Date	Series 9 off-the-run	Series 11 on-the-run	Series 11 without c	Series 11 with c
11/07/2008	4618.65	3570.56	3570.56	3570.56
11/10/2008	4488.09		3469.63	3475.85
11/12/2008	4627.38		3577.31	3590.15
11/13/2008	4639.88	3606.31	3586.97	3606.31

where $s_{\text{on-the-run},t} := s_{0,T}(A,D)$ at time t is the last available, $s_{\text{on-the-run},t+u}$ is the next available tranche spread of the on-the-run series and $s_{9,t}$ with $s_{9,t+u}$ denote the corresponding off-the-run tranche spreads from series 9. To ensure that we match the next available tranche spread of the on-the-run series $s_{\text{on-the-run},t+u}$, we multiply the tranche spread innovations of series 9 with a constant c defined as

$$c = \left(\frac{s_{\text{on-the-run},t+u}}{s_{\text{on-the-run},t}} \times \frac{s_{9,t}}{s_{9,t+u}} \right)^{\frac{1}{u}}$$

For the following missing value $\hat{s}_{\text{on-the-run},t+2}$, we proceed by taking $\hat{s}_{\text{on-the-run},t+1}$ as the next starting point until there are no more data gaps in the data set.

Table 2.7 exhibits a brief illustrative example with and without the constant c for the time period from 11/07/2008 to 11/13/2008. Our approach employs all available tranche information and adjusts for the higher spreads (i.e., higher risk) of the off-the-run series 9. As is shown in Figure 2.11, the spreads of the equity tranche of series 9 are substantially higher than the respective on-the-run series which we account for with the constant c.

During the time period in question, four credit events occurred in series 9. Fannie Mae and Freddie Mac were placed into conservatorship on 09/07/2008, Washington Mutual was placed into receivership on 09/25/2008 (and filed for Chapter 11 bankruptcy on the next day) and CIT Group filed for bankruptcy protection under Chapter 11 on 11/01/2009. Following a credit event in a constituent of the CDX index, a new version of the index is issued after setting the recovery price in an auction.[33] That is, the final auction price will determine the new attachment and detachment points of version 2 of the

[33] The results of credit event auctions are published on the Credit Event Fixings website, see `http://www.creditfixings.com/`.

Figure 2.11: CDX series 11 to 13 fixed. This figure graphs the CDX running spreads for series 9, series 10, the fixed series 11 through 13, and series 15 during the crisis period. To fix non-trivial data gaps (orange) in the on-the-run series (gray) during the liquidity crisis, we use the highly liquid "crisis" series CDX.NA.IG.9 (olive).

index. In the case of Fannie Mae and Freddie Mac the recovery prices are pretty high with 91.51% and 94%, respectively. Thus, the adjustments to the attachment and detachment points had been negligible. The auction prices for Washington Mutual and CIT Group are 57% and 68.125%, respectively. Hence, these two credit events had a little greater impact on the new attachment and detachment points though the recovery rates are well above the contractual recovery rate of 40% for the CDX index.

2.C Integration using interpolating spline functions

Given an interpolating spline function $s(x)$ as a valid approximation to the true function $f(x)$, it is straightforward to derive analytical solutions to integrals on arbitrary intervals in $[a, b]$

$$\int_a^b f(x)\,dx \approx \int_a^b s(x)\,dx = \sum_{i=0}^{n-1} \int_{x_i}^{x_{i+1}} s_i(x)\,dx$$

where $[a, b]$ is partitioned into $n-1$ consecutive subintervals $[x_i, x_{i+1}]$, n knots of the spline function and $a = x_0 < x_1 < \cdots < x_n = b$.[34] On any interval $[x_i, x_{i+1}]$ the interpolating function is written as (for generality, we consider a cubic spline function)

$$s_i(x) = a_i + b_i(x - x_i) + c_i(x - x_i)^2 + d_i(x - x_i)^3$$

Thus given the piecewise constant coefficients, a_i, b_i, c_i, d_i, according to some spline function, it follows from the interval additivity of definite integration that the integrals are the sum of the piecewise analytical solutions over the whole range $[a, b]$. The piecewise solutions are

$$x\left(a_i - b_i x_i + c_i x_i^2 - d_i x_i^3\right) + \frac{x^2}{2}\left(b_i - 2c_i x_i + 3d_i x_i^2\right) + \frac{x^3}{3}\left(c_i - 3d_i x_i\right) + \frac{x^4}{4} d_i \Bigg|_{x=x_i}^{x=x_{i+1}}$$

Besides the common case $\int_a^b f(x)\,dx$, we are further interested in

$$\int_a^b x^k f(x)\,dx \approx \int_a^b x^k s(x)\,dx = \sum_{i=0}^{n-1} \int_{x_i}^{x_{i+1}} x^k s_i(x)\,dx$$

for arbitrary order $k \geq 0$. Consequently, the analytical integral of any subinterval on $[x_i, x_{i+1}]$ is given by the evaluation of the generalized antiderivative

$$\frac{x^{1+k}}{1+k}\left(a_i - b_i x_i + c_i x_i^2 - d_i x_i^3\right) + \frac{x^{2+k}}{2+k}\left(b_i - 2c_i x_i + 3d_i x_i^2\right) + \frac{x^{3+k}}{3+k}\left(c_i - 3d_i x_i\right) + \frac{x^{4+k}}{4+k} d_i \Bigg|_{x=x_i}^{x=x_{i+1}}$$

For piecewise quadratic splines like the Schumaker (1983) interpolation function, set $d_i = 0$. The approach is equally applicable for other spline functions like rational spline interpolation, the resulting antiderivatives are a bit more involved though.

[34] This is the use case for our calculations. Cases for arbitrary intervals with $a > x_0$ or $b < x_n$ are a matter of bisecting the corresponding intervals i and setting $[a, x_{i+1}]$ and $[x_i, b]$ for the lower and upper bounds, respectively. Intervals of $[a, b]$ outside $[x_0, x_n]$ will require some form of extrapolation which is outside the scope of this paper.

Chapter 3

Hedging parameter risk

The content of this chapter refers to the working paper:

Claußen, Arndt, Daniel Rösch and Martin Schmelzle (2017), 'Hedging parameter risk', Working Paper, Leibniz Universität Hannover and Universität Regensburg.

Abstract

The accurate measurement and effective control of financial risk are of crucial importance to risk managers and regulators. However, risk measures are potentially affected by errors in the estimation of model parameters from limited samples, leading to parameter risk. The key contribution of this paper is the formulation of a general framework to hedge this parameter risk. Applying the new framework to credit portfolio modeling, we highlight the importance of parameter risk, model type, estimation methods, and diversification effects.

Keywords: estimation error, parameter risk, hedging

JEL: C13, G13, G32

3.1 Introduction

The assessment of real-world financial risk (e.g., value-at-risk or expected shortfall) usually involves the estimation of model parameters. In practical applications, the statistical estimates of those parameters are typically treated as if they were the true (i.e., known) values of unknown parameters. As a result, the effects of estimation errors resulting from insufficient statistical data or the incapability to accurately estimate the model parameters are ignored.

However, what is optimal in the absence of estimation risk is not necessarily optimal if estimation risk and parameter uncertainty is just neglected (Klein et al., 1978). Thus it is imperative for rational risk managers and prudential regulators to take account of this uncertainty when assessing risk model outcomes.

Pioneering work on the impact of estimation risk on investment decisions and portfolio selection include Kalymon (1971), and Klein and Bawa (1976). Barry and Brown (1985) and Coles and Loewenstein (1988) study effects upon market equilibrium and asset pricing in the presence of uncertain parameters and estimation risk. Common approaches to address estimation risk are Bayesian methods like diffusive priors (Barry, 1974) or informative priors formed via asset pricing models (Pástor, 2000). Other prominent approaches include 'robust' estimation principles (Garlappi et al., 2007).[1]

First efforts devoted to the issue of estimation errors on statistical measures of risk exposures in risk management are due to Jorion (1996). Since then, a steadily growing body of literature is emerging. The impact and assessment of parameter and model uncertainty on capital adequacy and risk capital calculations is studied by, e.g., Bao and Ullah (2004), Christoffersen and Gonçalves (2005), Kerkhof et al. (2010), Tarashev (2010), Alexander and Sarabia (2012), Bion-Nadal and Kervarec (2012), Embrechts et al. (2013), Lönnbark (2013), Embrechts et al. (2015), Fröhlich and Weng (2015), and Bignozzi and Tsanakas (2016) among others.[2] Whereas the meaning of parameter uncertainty in the literature is evident, the notion of model uncertainty is somewhat ambiguous. The uncertainty in the estimates of parameters within a model might also be understood as model uncertainty. Regardless the point of view and even if it were possible to neglect model uncertainty issues due to, e.g., regulatory guidelines, the need to specify unknown parameters within a given model class remains an important issue.[3] In essence, the literature provides orientation to add some 'conservatism' to estimates from statistical inference

[1]The role of parameter uncertainty in portfolio choice, asset pricing and asset allocation is further investigated in, e.g., Barberis (2000), Veronesi (2000), Xia (2001), Maenhout (2004), and Kan and Zhou (2007). Other studies pertaining to uncertainty in model parameters include option pricing (Bunnin et al., 2002), credit spreads (Korteweg and Polson, 2010), bond portfolios (Feldhütter et al., 2012), and mortgage securitization (Rösch and Scheule, 2014) *inter alia*. The majority of these studies, however, do not focus on risk management issues as such.

[2]There is also a fast growing literature considering the accuracy and backtesting of risk predictions, see, e.g., Berkowitz and O'Brien (2002), Berkowitz et al. (2009), Escanciano and Olmo (2011), Gouriéroux and Zakoïan (2013), Boucher et al. (2014), and Du and Escanciano (2017) *inter alia*.

[3]Therefore the present paper is not concerned with the hedging of *model risk* as such.

or otherwise provision additional capital buffers to cushion for estimation uncertainties.[4]

This paper represents, to the best of our knowledge, the first attempt to formulate a framework that allows financial institutions to hedge parameter risk instead of fully provision for 'conservative' parameters. Here, the term *hedge* is used in its broadest sense, as an offsetting position to potential losses, and *parameter risk* is the possibility for errors in the parameters. The framework is general, modest in assumptions, and enables contract parties to uniquely determine the fair pricing for parameter risk protection. Such a protection becomes increasingly expedient for circumstances characterized by a high chance for parameter error. This might be particularly pronounced for risk models where estimated risk measures, on average, approach the true, but unknown, risk measure from below, compare, e.g., Figlewski (2004) and Lönnbark (2010). Typically, dependent on the risk measure under consideration, implications become worse with asymmetric distributions and higher tail cutoffs, where a few, but severe, underestimations harm much more than many small ones.

We illustrate our framework applied to credit risk modeling.[5] Firstly, we conduct a simulation study. Employing a common credit portfolio model allowing for closed form hedge premium pricing, we find that (i) a prespecified interval of model parameters can be covered for a small periodic premium payment, (ii) a protection buyer (seller) has an incentive to engage in such a hedge, if the true parameters are underestimated (overestimated), respectively, (iii) the premium is quite robust to the specification of the credit risk model, (iv) a parameter risk hedge to some degree implicitly covers model risk with respect to point estimates from alternative assumptions about the true default generating process, and (v) a protection seller engaging in more

[4]Risk based capital requirements seek to lower the probability of liquidation. However, Hellmann et al. (2000) demonstrate that capital requirements, while putting bank equity at risk and inducing prudent bank behavior, at the same time adversely affect banks' franchise values, hence stimulating gambling incentives. A similar argument is made by Keppo et al. (2010), who show that higher regulatory capital requirements may postpone recapitalization, thereby actually increasing the institution's probability of default. Consequently, additional capital burdens to achieve a greater level of safety may negatively affect financial institutions' behavior.

[5]Estimation errors of parameters are documented to be particularly pronounced in credit risk, Crouhy et al. (2000) for instance find evidence that industry credit models yield estimates substantially less accurate in comparison to market risk models. Furthermore, Gordy and Heitfield (2010) show that the true model parameters are—on average—underestimated. Other applications of uncertainty in the parameters to credit risk modeling include, e.g., Löffler (2003), Yamai and Yoshiba (2005), Tarashev (2010), and Bernard et al. (2017) among others.

than one hedge contract may diversify parameter risk, and decrease its risk of extreme losses. Secondly, an application of our framework to historical one year default rates reveals that the periodic premium of the protection buyer insuring against one weighted standard deviation of the estimates is about 0.02 to 1.66 basis points. Or, subject to contract type and rating grade, about 460 to 1430 years of paying a premium equals the one-off difference between the VaR using the estimates and the VaR for the conservative parameters. We find strong empirical support that in practical applications high rated financial instruments are substantially more prone to parameter errors than lower rated instruments, and—as a result—are more costly in relative terms to be hedged against parameter risk.

For practical applications, parameter risk hedging hinges on institutions willing to sell protection and are capable to steadily provide protection even during periods of high uncertainty surrounding the state of the economy. Prospects could be federal institutions or state funds who themselves are less vulnerable to get into financial distress. In fact, there is ample evidence that financial firms in distress have already been backed by federal agencies, however, without any formal framework, see, e.g., Jorion (2000) and Veronesi and Zingales (2010). The creation of a 'Parameter Risk Hedge Funds' administered by, e.g., the major regulatory, treasury, deposit protection or central bank agencies, might be a suitable candidate to provide protection against uncertainty with respect to parameter errors.[6] Unlike taxpayer bailouts, such a funds could mutualize risks among financial firms through an emergency authority. In addition, a protection seller gets into the unique position to diversify parameter errors across different parameter risk hedge contracts.

The rest of the paper is organized as follows. Section 3.2 introduces our framework to hedge parameter risk. In Section 3.3, we illustrate the key concepts of our framework, and discuss real world implications. Section 3.4 summarizes our findings and concludes.

3.2 Hedging of parameter risk

Let $(\Omega,\mathcal{F},\mathbb{P})$ be an atomless probability space and $L^0(\Omega,\mathcal{F},\mathbb{P})$ a set of $\mathcal{F}$-measurable $\mathbb{P}$-a.s. finitely valued random variables on that probability space.

[6]Federal or regulatory agencies might already maintain relationships to financial institutions and closely monitor them. Parameter risk hedge funds do not necessarily introduce additional issues pertaining to asymmetrical information or excessive risk taking. Discussions on the problems of agency issues from a regulatory or supervisory perspective include incentives to underestimate VaR (Cuoco and Liu, 2006), and decrease the quality of risk management systems (Daníelsson et al., 2002).

Definition 3.1 (Risk of a security). *Let $\mathcal{M} \subseteq L^0(\Omega, \mathcal{F}, \mathbb{P})$ be a convex cone, in which any random variable $L(\boldsymbol{Y}, \boldsymbol{\theta}) \in \mathcal{M}$, with cumulative (discrete) distribution function $G(\ell, \boldsymbol{\theta}) = \mathbb{P}(L(\boldsymbol{Y}, \boldsymbol{\theta}) \leq \ell)$, $\ell \in [0,1]$, and probability density (mass) function $g(\ell, \boldsymbol{\theta})$, describes a loss of a risky security (e.g., loan, bond, stock, portfolio). $\boldsymbol{Y}$ is a random vector and $\boldsymbol{\theta}$ is a vector of model parameters. Then $L(\boldsymbol{Y}, \boldsymbol{\theta})$ models the* risk *(e.g., market risk, credit risk)* of the security.

Quantitative tools to aggregate the risk of a security and to express the riskiness of financial positions in one key figure are *risk measures*. A risk measure is a function that maps risk to the real numbers. Therefore any function $\mathcal{R} : \mathcal{M} \to \mathbb{R}$ is a risk measure. Statistical risk models often require a number of model parameters $\boldsymbol{\theta}$ to be specified. In practical applications $\boldsymbol{\theta}$ is unknown and must be determined, e.g., via expert judgments, calibration to market data or estimation from past observations. However, the inferred parameters do not necessarily match the true underlying model parameters under a classical statistical perspective.

Definition 3.2 (Estimates and parameter error). *Given the parameters according to Definition 3.1 are unknown, then any specified parameters are* estimates $\hat{\boldsymbol{\theta}}$ *and the (p-norm) distance $\left\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\right\|_p$, with real-valued $p \geq 1$, is the* parameter error.

Definition 3.3 ((Positive/negative) effective parameter error). *For a given risk measure $\mathcal{R}$, an estimate $\hat{\boldsymbol{\theta}}$ implies an* effective *parameter error if*

$$\mathcal{R}(L(\boldsymbol{Y}, \boldsymbol{\theta})) \neq \mathcal{R}(L(\boldsymbol{Y}, \hat{\boldsymbol{\theta}})),$$

and is called positive *effective parameter error if*

$$\mathcal{R}(L(\boldsymbol{Y}, \boldsymbol{\theta})) - \mathcal{R}(L(\boldsymbol{Y}, \hat{\boldsymbol{\theta}})) > 0,$$

otherwise it is called negative *effective parameter error.*

All estimates lead almost surely to effective parameter errors. In practical applications, a negative effective parameter error—equivalent to an overestimation of the quantified risk of a security—may imply a reduced possibility of a default, which will be costly for a financial institution due to misallocation of capital. On the contrary, a positive effective parameter error—equivalent to an underestimation of the quantified risk of a security—may reduce the cost of capital in the short run, but in the long run it might increase the possibility of a default because the financial institution will not have enough provisions to cover incurred losses.

Definition 3.4 (Parameter risk). *The possibility of an effective parameter error is called* parameter risk.

Remark 3.1. *Knight (1921) defines the terms 'risk' and 'uncertainty' as random variation according to a known or unknown stochastic law, respectively, whereas the type of desired or undesired outcome remains unspecified. In Definition 3.4 we introduce parameter 'risk', since for the definition of a parameter error we assume the model is known with certainty.*

Financial institutions are particularly hurt from positive effective parameter errors, therefore current literature calls to adjust for this particular form of parameter risk by add ons, compare, e.g., Tarashev (2010). However, if these parameters—adjusted for parameter risk—are employed and serve as a basis for the allocation of capital reserves, the provisions may be too costly. Instead of fully providing additional capital reserves, a financial institution may benefit from engaging in a *hedge* of parameter risk.

Definition 3.5 (Payoff of parameter risk hedge)**.** *Given a risk measure $\mathcal{R}$, an estimate $\hat{\boldsymbol{\theta}}$ is* hedged *to the level of $\boldsymbol{\theta}_h$ for*

$$\mathcal{R}(L(\boldsymbol{Y},\hat{\boldsymbol{\theta}})) \leq \mathcal{R}(L(\boldsymbol{Y},\boldsymbol{\theta}_h)),$$

if a protection seller pays the protection buyer the (possibly discounted) payoff

$$p_c = \begin{cases} 0 & \text{if} \quad L(\boldsymbol{Y},\boldsymbol{\theta}) \leq \mathcal{R}(L(\boldsymbol{Y},\hat{\boldsymbol{\theta}})) \\ L(\boldsymbol{Y},\boldsymbol{\theta}) - \mathcal{R}(L(\boldsymbol{Y},\hat{\boldsymbol{\theta}})) & \text{if} \quad \mathcal{R}(L(\boldsymbol{Y},\hat{\boldsymbol{\theta}})) \leq L(\boldsymbol{Y},\boldsymbol{\theta}) \leq \mathcal{R}(L(\boldsymbol{Y},\boldsymbol{\theta}_h)) \\ c\,(\mathcal{R}(L(\boldsymbol{Y},\boldsymbol{\theta}_h)) - \mathcal{R}(L(\boldsymbol{Y},\hat{\boldsymbol{\theta}}))) & \text{if} \quad \mathcal{R}(L(\boldsymbol{Y},\boldsymbol{\theta}_h)) < L(\boldsymbol{Y},\boldsymbol{\theta}), \end{cases}$$

with $c \in \{0,1\}$.

The contractual parameter c in Definition 3.5 defines the shape of the payoff structure above the hedge level $\boldsymbol{\theta}_h$ restricted to two polar edge cases. Choosing $c = 1$ the hedge contract can be modeled as a call spread on the underlying state variable $L(\boldsymbol{Y},\boldsymbol{\theta})$. Here, a protection seller offers to bear losses beyond a certain threshold $\mathcal{R}(L(\boldsymbol{Y},\hat{\boldsymbol{\theta}}))$, but subject to a cap $\mathcal{R}(L(\boldsymbol{Y},\boldsymbol{\theta}_h))$ and, as such, is basically selling mezzanine protection. For $c = 0$ the hedge contract resembles an up and out call option. After breaking the upper barrier nothing is paid to the protection buyer resulting in a much cheaper fee to be paid for this type of protection. Of course, in principle, c could also be modeled more generally within a functional relation to the loss variable $L(\boldsymbol{Y},\boldsymbol{\theta}_h)$ to allow for more diverse payoff shapes above the hedge level. However, one should keep in mind that values for $c > 1$ would result in an upwards parallel shift of the horizontal leg above the hedge level with unclear economic interpretation. Likewise, values $c < 0$ seem unreasonable since this would lead to a negative cash flow in case of a breached hedge level. That is, we may interpret $c = 1$

and $c = 0$ as natural lower and upper bounds of protection above the hedge level, respectively.

For both characteristics the compensating payment is restricted by $\mathcal{R}(L(\boldsymbol{Y}, \boldsymbol{\theta}_h)) - \mathcal{R}(L(\boldsymbol{Y}, \hat{\boldsymbol{\theta}}))$. Thus, by effectively truncating the payoff the protection seller faces no unpredictable and heavy-tailed worst case scenarios.

Remark 3.2. *The economic interpretations of the two alternatives for c in Definition 3.5 should be well considered. A severe economic state described by an extreme realization of $\boldsymbol{Y}$ could imply a large loss. In the case of $c = 1$, the contract will payoff regardless whether parameter risk is causal to the loss. Alternatively, one could argue that a compensation for a possible parameter risk does not seem adequate ($c = 0$) and it is then preferable to liquidate the financial institution.*

Remark 3.3. *Our framework intends to adjust possibly misspecified parameters upwards to gain additional safety. Alternatively, protection could be provided in the reverse direction. In this case, the attachment point for the resulting payoff profile would be described by $\mathcal{R}(L(\boldsymbol{Y}, \boldsymbol{\theta}_h))$ and would be confined to the right by the estimated parameter as an upper level of protection, i.e., $\mathcal{R}(L(\boldsymbol{Y}, \hat{\boldsymbol{\theta}}))$. Specifically, cash flows are now tied to a scenario where an involved party has strong belief that their parameter estimates are effectively too conservative and may spare a fraction of some 'parameter risk buffer'. Though this scenario is conceivable, it seems hard to imagine that it could—in any form—be beneficial for a sound financial system.*

Based on these preliminary observations, the fair premium for bearing parameter risk then directly follows via integrating over the loss probabilities given $\boldsymbol{Y}$ and the true model parameters $\boldsymbol{\theta}$. For bearing the parameter risk from the protection buyer, the protection seller receives a fair hedge fee, or, equivalently, a fair hedge premium, defined as the expectation of the payoff, $\mathbb{E}[p_c]$, taken under some appropriate probability measure (see Remark 3.4).

Theorem 3.1 (Fair fee f_c for a hedge of parameter risk)**.** *The* fair fee $f_c = \mathbb{E}[p_c]$ *of a parameter risk hedge according to Definition 3.5 is given by*

$$f_0 = f_1 - (D - A)\,\mathbb{P}(D < L(\boldsymbol{Y}, \boldsymbol{\theta})), \tag{3.1}$$

$$f_1 = D - A - \int_A^D \mathbb{P}(L(\boldsymbol{Y}, \boldsymbol{\theta}) \leq \ell)\, d\ell, \tag{3.2}$$

where $A = \mathcal{R}(L(\boldsymbol{Y}, \hat{\boldsymbol{\theta}}))$ and $D = \mathcal{R}(L(\boldsymbol{Y}, \boldsymbol{\theta}_h))$.

Proof. The payoff p_1 corresponds to the difference of a long call with strike $A = \mathcal{R}(L(\boldsymbol{Y}, \hat{\boldsymbol{\theta}}))$ and short call with strike $D = \mathcal{R}(L(\boldsymbol{Y}, \boldsymbol{\theta}_h))$. The contractual payment equals

$$p_1 = (L(\boldsymbol{Y}, \boldsymbol{\theta}) - A)^+ - (L(\boldsymbol{Y}, \boldsymbol{\theta}) - D)^+.$$

Due to the linearity of the expectation operator we simply calculate

$$\begin{aligned}\mathbb{E}\Big[(L(\boldsymbol{Y},\boldsymbol{\theta})-A)^{+}\Big] &= \int_A^1 (\ell - A)\, g(\ell,\boldsymbol{\theta})\, d\ell = \int_A^1 \ell\, g(\ell,\boldsymbol{\theta})\, d\ell - A\int_A^1 g(\ell,\boldsymbol{\theta})\, d\ell \\ &= \ell\, G(\ell,\boldsymbol{\theta})\Big|_A^1 - \int_A^1 G(\ell,\boldsymbol{\theta})\, d\ell - A\, G(\ell,\boldsymbol{\theta})\Big|_A^1 \\ &= 1 - A - \int_A^1 G(\ell,\boldsymbol{\theta})\, d\ell = 1 - A - \int_A^1 \mathbb{P}(L(\boldsymbol{Y},\boldsymbol{\theta}) \le \ell)\, d\ell,\end{aligned}$$

which follows by integration by parts and resembles a call option with strike A. Subtraction leads to Equation (3.2). The payoff p_0 equals p_1 without $\mathcal{R}(L(\boldsymbol{Y},\boldsymbol{\theta}_h)) - \mathcal{R}(L(\boldsymbol{Y},\hat{\boldsymbol{\theta}}))$ if $\mathcal{R}(L(\boldsymbol{Y},\boldsymbol{\theta}_h)) < L(\boldsymbol{Y},\boldsymbol{\theta})$. Due to the linearity of the expectation operator Equation (3.1) holds. □

Remark 3.4. $\mathbb{E}[p_c]$ *in Theorem 3.1 denotes expectation with respect to at least two probability measures. If it were possible to construct a hedging portfolio through the trading of liquid assets which would perfectly replicate the parameter risk hedge payoff profile, the expectation could be taken under an 'equivalent martingale measure'. If there is no market absent of arbitrage for the payoff profile under consideration, then 'risk neutral' pricing cannot be applied. Thus, the expectation will include some additional risk premium to compensate the protection seller for not being able to fully eliminate all risk. Hence 'real world' pricing will take place under a physical probability measure.*

The protection seller pays on average $\mathbb{E}[p_c]$ without necessarily requiring knowledge about the true parameter $\boldsymbol{\theta}$. However, in practical applications $\boldsymbol{\theta}$ is inevitably unknown and fair fees according to Theorem 3.1 are not computable. For that reason, we introduce a *contractual* fee for a parameter risk hedge.

Definition 3.6 (Contractual fee f_c^* for a parameter risk hedge)**.** *The* contractual fee f_c^* *is defined as the expectation of the parameter risk hedge payoff for* $L(\boldsymbol{Y},\boldsymbol{\theta}_h)$ *instead of* $L(\boldsymbol{Y},\boldsymbol{\theta})$. f_c^* *accords with* f_c *in Theorem 3.1 where* $\boldsymbol{\theta}$ *is replaced by* $\boldsymbol{\theta}_h$.

The Definition 3.6 implies three advantageous consequences. First, given a hedge range from $\mathcal{R}(L(\boldsymbol{Y},\hat{\boldsymbol{\theta}}))$ to $\mathcal{R}(L(\boldsymbol{Y},\boldsymbol{\theta}_h))$ the contractual fee f_c^* is deterministic. Second, the contractual fee f_c^* trivially equals the fair fee f_c if the hedge level $\boldsymbol{\theta}_h$ matches the latent $\boldsymbol{\theta}$. And, third, $\boldsymbol{\theta}_h$ is merely a conservative estimate (or, the result of an add on to an estimate) for $\boldsymbol{\theta}$ and as result the hedge framework does not induce any new parameter risk.

Remark 3.5. *A closer look at Equation (3.2) reveals familiar similarities to 'classical' derivatives pricing. Initially considering* $\boldsymbol{\theta}$*, we obtain the fair fee previously* *discussed assuming the true parameters are known. In addition, considering $\boldsymbol{\theta}_h$ subject to specified contract terms instead of $\boldsymbol{\theta}$, we arrive at what we label the contractual fee. In conclusion, replacing $\boldsymbol{\theta}$ with $\hat{\boldsymbol{\theta}}$, ultimately yields to conventional call spread pricing.*

Remark 3.6. *A protection seller gets into the unique position to directly diversify parameter risk across different hedge contracts. If a parameter hedge contract does (not) fulfill the condition*

$$\int_A^D \mathbb{P}(L(\boldsymbol{Y}, \boldsymbol{\theta}_h) \leq \ell)\, d\ell \;\; > \;\; \int_A^D \mathbb{P}(L(\boldsymbol{Y}, \boldsymbol{\theta}) \leq \ell)\, d\ell, \tag{3.3}$$

then it follows from Equation (3.2) and Definition 3.6 that the fair fee f_1 is greater than (less than or equal to) the contractual fee f_1^* and, as result, the protection seller receives on average less than (more than or equal to) what is paid. This condition, given by the areas under the cumulative distributions functions ranging from A to D, may be interpreted as an underestimation of the true extreme losses which triggers the payout of a hedge contract.

In real world applications the condition in Equation (3.3) cannot be verified for a specific hedge contract, since $\boldsymbol{\theta}$ is essentially unknown. However, if a protection seller engages in more than one hedge contract and the contractual terms therein are not perfectly correlated, then there is a positive likelihood that some contracts will fulfill the condition in Equation (3.3) and others will not. Thus, there is a chance that the underestimation of the true risk, represented by the condition in Equation (3.3), is offset—at least to some extent—with an overestimation of the true risk of some other hedge contract. As result, a protection seller is also able to diversify the parameter risk across several protection buyers.

In conclusion, the parameter risk hedge framework, in particular Definition 3.5 and Theorem 3.1, is independent from the model type, chosen hedge interval, and holds for any risk measure.

3.3 Application

3.3.1 Parameter risk hedge premiums

To begin with, we apply the parameter risk hedge framework to the asymptotic single risk factor (ASRF) model. This model has clear economic interpretations and offers closed form solutions for parameter risk hedge premiums. Underpinning the Basel internal ratings-based (IRB) approach (BCBS, 2006), the ASRF model is well known to academics and practitioners. Gordy and

Heitfield (2010) and Tarashev (2010) demonstrate the importance of dealing with parameter risk in this model.[7]

Losses on a homogeneous credit portfolio are assumed to be driven by a systematic risk factor, and each asset represents an infinitesimal share of this portfolio. Idiosyncratic risk disappears owing to full diversification. The distribution of credit portfolio loss in a given period is modeled by

$$L^G(Y,[\rho,\pi]) = \Phi\left(\frac{\Phi^{-1}(\pi) - \sqrt{\rho}\,Y}{\sqrt{1-\rho}}\right), \quad Y \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1), \tag{3.4}$$

where $\pi \in (0,1)$ is an unconditional probability of loan or bond default, $\rho \in (0,1)$ is the asset (return) correlation, Y is a standard normally distributed systematic risk factor, and Φ is the standard normal distribution function (with Φ^{-1} denoting its inverse).[8]

Given the analytical tractability of the Gaussian ASRF, hedge premiums are straightforward to derive and can be expressed in closed form.

Corollary 3.1 (Parameter risk hedge in the Gaussian ASRF)**.** *The Equation (3.2) from Theorem 3.1 for the Gaussian ASRF can be easily expressed by the difference of two bivariate standard normal cumulative distribution functions*

$$\mathbb{E}\left[p_1^G\right] = \Phi_2\left(-\Phi^{-1}(A), \Phi^{-1}(\pi), \varrho\right) - \Phi_2\left(-\Phi^{-1}(D), \Phi^{-1}(\pi), \varrho\right), \tag{3.5}$$

with

$$\varrho = -\sqrt{1-\rho}, \quad A = \mathcal{R}\Big(L^G(Y,[\hat{\rho},\hat{\pi}])\Big), \quad \textit{and} \quad D = \mathcal{R}\Big(L^G(Y,[\rho_h,\pi_h])\Big),$$

where $\Phi_2(\cdot,\cdot;\varrho)$ denotes the bivariate standard normal cumulative distribution function with correlation parameter ϱ.

Additionally $\mathbb{P}\Big(D \le L^G(Y,[\rho,\pi])\Big)$ in Equation (3.1) can be expressed in closed form, leading to

$$\mathbb{E}\left[p_0^G\right] = \mathbb{E}\left[p_1^G\right] - (D-A)\,\Phi\left(\frac{\varrho\,\Phi^{-1}(D) + \Phi^{-1}(\pi)}{\sqrt{\rho}}\right). \tag{3.6}$$

Proof. Using Equation (30c) in Andersen and Sidenius (2005). □

[7] Section 3.A in the Appendix further illustrates parameter risk in the ASRF.

[8] For simplicity, we skip the notion of a nonzero recovery rate. The superscript G distinguishes the original Gaussian ASRF from non-Gaussian generalizations in later sections.

In the sequel, we consider the common value-at-risk (VaR) and conditional value-at-risk (CVaR) measures. For the Gaussian ASRF these read

$$\mathrm{VaR}_{\alpha}^{G}(\rho,\pi) = \Phi\left(\frac{\Phi^{-1}(\pi) - \sqrt{\rho}\,\Phi^{-1}(1-\alpha)}{\sqrt{1-\rho}}\right),$$

and

$$\mathrm{CVaR}_{\alpha}^{G}(\rho,\pi) = \frac{1}{1-\alpha}\Phi_2\left(\Phi^{-1}(\pi), \Phi^{-1}(1-\alpha), \sqrt{\rho}\right).$$

Next, we illustrate general properties and magnitudes of hedge premiums for different contract types and risk measures. During our case studies, we focus on the essentials of the parameter risk hedge framework and refrain from discussing possible influences due to risk preferences of economic agents. That is, we consider zero risk premiums and undiscounted payoffs.

For the asset correlation, we assume $\rho = 20\%$, which lies in between the lower and upper bounds found in the Basel Accords. The probabilities of default are $\pi \in \{0.2\%, 1\%, 5\%\}$, roughly corresponding to historical one year default rates of Moody's Baa, Ba, and B rated corporates (compare Table 3.4). We further assume that the estimation leads to an underestimation of 25% for each parameter.[9] Since we know the true parameters, we can calculate the fair hedge fee, when the true parameter is hedged, i.e., the hedge range is defined from $\hat{\boldsymbol{\theta}}$ to $\boldsymbol{\theta}_h = \boldsymbol{\theta}$.

Table 3.1 provides the results for three different parameter settings and for the risk measures VaR and CVaR for the confidence levels $\alpha = 99\%$ and $\alpha = 99.9\%$. The expected hedge premiums for the barrier option like payoff structure f_0 is less than the corresponding expected call spread option payoff f_1, which, from Equation (3.1), is by construction. Otherwise, the two payoff types share similar behaviors, thus we will focus on the case $c = 0$ for further discussions.

The resulting fair hedge premiums range from 0.17 to 16.73 basis points.[10] For instance, for $\alpha = 99.9\%$ and the medium default risk bucket ($\pi = 1\%$) an estimated VaR of 9.01% is hedged to the level of 14.553%. For this hedge level the protection buyer has to pay a fee of 0.929 bp, which is equivalent to $596 \approx \frac{14.553\% - 9.01\%}{0.929\,\mathrm{bp}}$ years of paying a fee instead of increasing capital for achieving the same degree of safety.

[9] From the analysis presented in Figure 3.3 of the Appendix, we learn that such an underestimation for the shortest time horizon with $T = 7$ years and lowest probability of default occurs in around 30% of the cases.

[10] In comparison, e.g., the FDIC charges *total base assessment rates* around 7–77.5 basis points annually for all risk categories, compare FDIC (2011).

Table 3.1: Comparison of fair hedge premiums and risk measures for different parameter settings, contractual types, and confidence levels. Entries report the resulting risk measures $\mathcal{R}$, i.e., VaR and CVaR, alongside with their corresponding fair hedge premiums $f_c = \mathbb{E}[p_c]$, i.e., the hedge parameter $\boldsymbol{\theta}_h$ equals $\boldsymbol{\theta}$, for two contract types $c \in \{0,1\}$, confidence levels $\alpha \in \{99\%, 99.9\%\}$ and the hedge range is from $\hat{\boldsymbol{\theta}}$ to $\boldsymbol{\theta}_h = \boldsymbol{\theta}$. For the Gaussian ASRF model the relevant vector of model parameter is $\boldsymbol{\theta} = [\rho, \pi]$. The first two columns display the underlying true parameter vectors $\boldsymbol{\theta}$, while the next two columns represent possible estimates of the parameters $\hat{\boldsymbol{\theta}}$ with an assumed 25% underestimation of the true model parameters.

True [%]		Estimates [%]		Risk measures /	VaR		CVaR	
ρ	π	$\hat{\rho}$	$\hat{\pi}$	Premiums	$\alpha = 99\%$	$\alpha = 99.9\%$	$\alpha = 99\%$	$\alpha = 99.9\%$
20	0.2	15	0.15	$\mathcal{R}(\boldsymbol{\theta})$ [%]	1.9954	4.7187	3.1410	6.4159
				$\mathcal{R}(\hat{\boldsymbol{\theta}})$ [%]	1.2490	2.7378	1.8780	3.6486
				f_0 [bp]	0.4628	0.2604	0.3844	0.1685
				f_1 [bp]	1.2091	0.4585	0.8032	0.2638
20	1.0	15	0.75	$\mathcal{R}(\boldsymbol{\theta})$ [%]	7.5251	14.5525	10.5129	18.1436
				$\mathcal{R}(\hat{\boldsymbol{\theta}})$ [%]	4.8354	9.0102	6.6139	11.1908
				f_0 [bp]	2.1321	0.9294	1.5535	0.5582
				f_1 [bp]	4.8218	1.4837	2.9210	0.8083
20	5.0	15	3.75	$\mathcal{R}(\boldsymbol{\theta})$ [%]	24.9575	38.4422	30.8119	43.8506
				$\mathcal{R}(\hat{\boldsymbol{\theta}})$ [%]	17.0061	26.3358	21.0410	30.3641
				f_0 [bp]	8.7783	2.8725	5.6220	1.5995
				f_1 [bp]	16.7296	4.0832	9.2661	2.1091

Given the parameter constellations in Table 3.1 the fair hedge premiums decrease with higher α. While the true VaR roughly multiply with 2.5 ($\pi = 0.2\%$), 2 ($\pi = 1\%$), and 1.5 ($\pi = 5\%$) by increasing α from 99% to 99.9%, the fair hedge premiums (in absolute terms) decrease by 0.60, 0.45 and 0.30, respectively. This behavior is expected since with α getting larger, actual payoffs from the parameter risk hedge become less probable.

By definition, the VaR are less than CVaR for all confidence levels, however, the fair hedge premiums for CVaR are consistently lower than they are for the VaR. Since the CVaR is equal to VaR with a higher confidence level, the smaller hedge fee for the CVaR follows directly from the results for a higher α.

For the same parameter error (in all cases each parameter is underestimated by 25%) the risk buckets with lowest probability of default react more sensitive. The relative error is 60% (55% for $\pi = 1\%$, 45% for $\pi = 5\%$). Therefore the lower risk buckets are apparently more prone to parameter errors. However, the relative premium $f_c/\mathcal{R}(\boldsymbol{\theta})$ is with 0.213%, 0.283%, 0.352% smaller, respectively. Thus, although higher rated risk buckets are more prone to parameter errors, the true hedge fee is (relatively) smaller for the same relative magnitude of parameter errors.

Next, we analyze how the fair hedge premiums f_c relate to contractual fees f_c^*. These premiums are deterministic with respect to a chosen hedge range, however, the hedge range is stochastic because it depends on the random realization of the estimates $\hat{\boldsymbol{\theta}}$ and standard errors $\hat{\sigma}(\hat{\boldsymbol{\theta}})$. For each case, we sample credit losses according to Equation (3.4). During each iteration, we employ the estimation method described in Duellmann et al. (2010) providing analytical solutions for the estimates $\hat{\boldsymbol{\theta}} = [\hat{\rho}, \hat{\pi}]$ with corresponding standard errors $\hat{\sigma}(\hat{\boldsymbol{\theta}})$. For this illustration, we define the hedge parameter as the sum of the estimate and its standard error weighted by a factor $\kappa = 75\%$. Therefore, for each single realization, we obtain different fair and contractual hedge premiums.

Figure 3.1 shows the realization of the distributions for the VaR sorted in ascending order, and the corresponding fair fees and contractual fees for two contract types $c \in \{0, 1\}$. The true parameters are $[\rho, \pi] = [20\%, 1\%]$ and the time horizons are 15 and 30 years.

In the left panel, we plot the distributions of the hedge range given by the difference of the lower attachment point A fixed by the VaR at the estimates as well as the upper detachment point D confined by the VaR at the hedged parameter at confidence level $\alpha = 99.9\%$. The vertical, dashed gray lines indicate the intersections where the estimated and hedged VaR equal the true VaR. These crossing points separate the simulated distributions in each subplot into three different areas labeled (I), (II), and (III).

(I) $A \leq D \leq \text{VaR}(\boldsymbol{\theta})$. From a hedging perspective, the product is beneficial, since the true risk is clearly underestimated (intersection from D to $\text{VaR}(\boldsymbol{\theta})$).

(II) $A \leq \text{VaR}(\boldsymbol{\theta}) \leq D$. This case describes a situation where the true risk is—on average—slightly overhedged. Without the hedge product, however, the true risk would remain underestimated.

(III) $\text{VaR}(\boldsymbol{\theta}) \leq A \leq D$. The true risk is overestimated.

In the right panel, we plot the contractual fee f_c^* in comparison to the resulting fair fee f_c for the two contractual types $c \in \{0, 1\}$. The passage from case (I) to (II) occurs, if $f_c = f_c^*$. That is, in this particular case, when the upper hedge range D correspond to the true VaR, then the hedge premium—irrespective of the contractual type—naturally equals the fair fee.

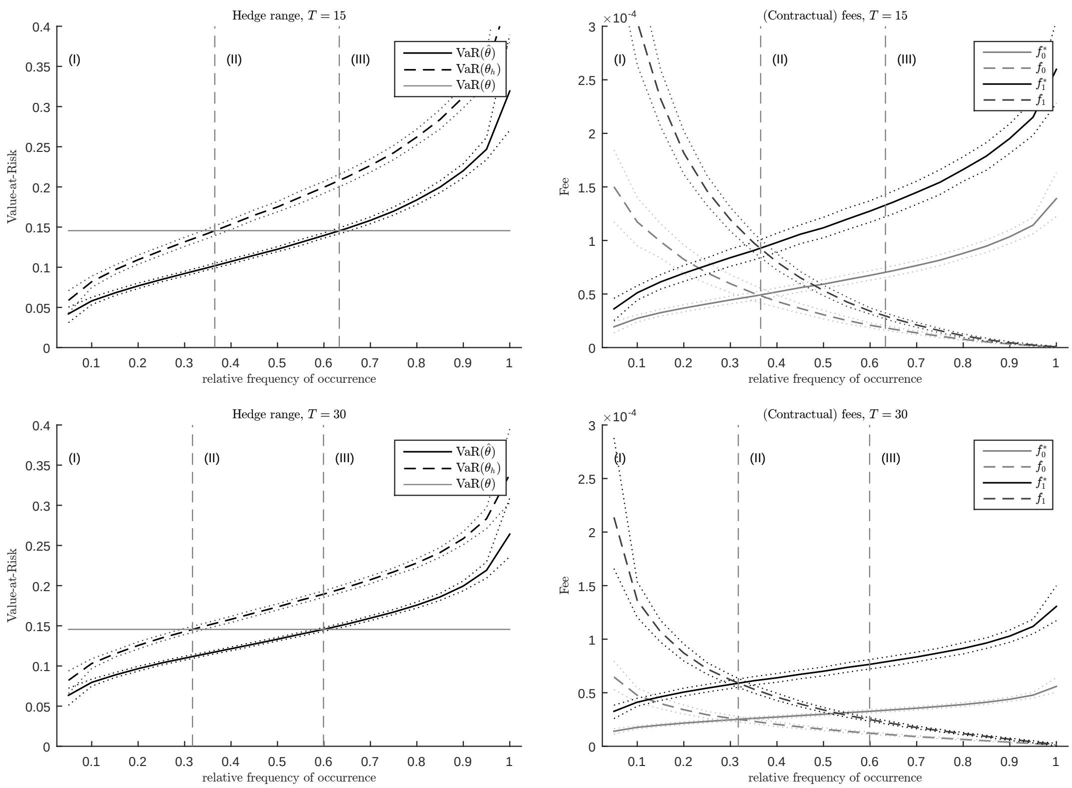

Figure 3.1: Effects of parameter risk on hedge premiums. The left figure shows the distributions of A given by $\text{VaR}(\hat{\boldsymbol{\theta}})$ (solid lines) and D given by $\text{VaR}(\boldsymbol{\theta}_h)$ (dashed lines) in comparison to the true $\text{VaR}(\boldsymbol{\theta})$ (horizontal lines) at confidence level $\alpha = 99.9\%$ for $\boldsymbol{\theta} = [\rho, \pi]$ with $\rho = 20\%$ and $\pi = 1\%$. The right figure shows the contractual fee f_c^* in comparison to the resulting fair fee f_c (i.e., the average expected payoff from the product) for the two contract types $c = 0$ (gray) and $c = 1$ (black). The upper two figures are for $T = 15$, and the lower two figures are for $T = 30$. The vertical, dashed gray lines indicate the intersecting points where the estimated and hedged VaR equal the true VaR. The simulation is repeated 10^4 times and results binned into 20 equally spaced intervals. The graphs are sorted by A, all other data points are adjusted accordingly. The solid and dashed lines plot the mean values, the dotted lines depict the corresponding 10% and 90% quantiles, respectively.

(I) The fair fee (expected payoff) is greater than the contractual fee. Recalling from the left panel that a parameter hedge is necessary, the protection buyer profits in two respects. Protection is needed, and the net present value for the protection buyer is positive.

(II) A parameter hedge is necessary, however, this comes at the cost of a slight overprotection. Thus, the net present value is negative.

(III) Any unit of extra protection will disproportionately drive associated cost.

The three cases each make up roughly one third of the relative frequency of occurrences. With T getting smaller, we observe a right shift of these divisions. This shift implies, that for higher parameter risk a hedge becomes increasingly reasonable for the protection buyer.

For the contractual type $c = 1$, the product is more vulnerable to parameter misspecifications viewed from both sides of protection, which becomes evident through the much steeper gradients in the right panel in comparison to $c = 0$.

Given the generality of our framework, the same kind of analysis conducted here could be applied to any model and risk type. With such a procedure the sensitivity of a specific model to parameter risk becomes assessable.

3.3.2 Diversification of parameter risk

As demonstrated by the analysis presented in Figure 3.1, the product translates the realized parameter risk of an underlying portfolio to a positive or negative difference of the fair fee f_c to the contractual fee f_c^*. Moreover, as argued at the end of Section 3.2, a protection seller may diversify this risk by engaging in more than one hedge contract with different protection buyers. A requirement for such diversification is that the hedge contracts are at least to some extent independent of each other.

We extend the ASRF framework following Pykhtin and Dev (2002) to demonstrate this diversification effect. The systematic single risk factor Y in Equation (3.4) now turns into a sectoral risk factor. Thus Y^j is composed of a super-systematic factor Y^*, representing the overall macro-economy, and a sector j specific risk factor U^j, according to

$$Y^j = \sqrt{\delta}\, Y^* + \sqrt{1-\delta}\, U^j,$$

where Y^* and U^j are i.i.d. standard normal. This extension allows modeling credit risky portfolios covering different sectors and the linear dependence

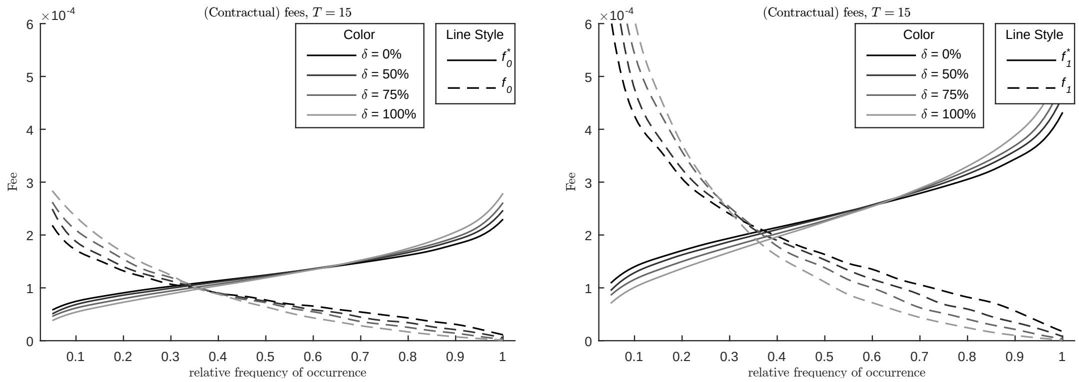

Figure 3.2: Diversification of parameter risk. This figure shows the contractual fee $f_c^* = f_c^{*,1} + f_c^{*,2}$ in comparison to the resulting fair fee $f_c = f_c^1 + f_c^2$ (i.e., the average expected payoff from the product) for the two contract types $c = 0$ (left graph) and $c = 1$ (right graph), for a protection seller engaging in two hedge contracts. $f_c^{*,i}$ (f_c^i) is the contractual (fair) fee of hedge contract $i \in \{1,2\}$. Both underlying portfolios differ sector wise and the dependence between these sectors is modeled by δ, a larger value of δ indicates a higher correlation and both sectors are perfectly (un)correlated for $\delta = 1$ ($\delta = 0$). All other parameters are assumed to be equal, i.e., $\boldsymbol{\theta} = [\rho, \pi]$ with $\rho = 20\%$ and $\pi = 1\%$. The simulation is repeated 10^4 times and results binned into 20 equally spaced intervals. The graphs are sorted by A (not depicted), all other data points are adjusted accordingly.

structure between the sectors is explicitly described by $\delta \in [0,1]$. If δ equals one, all sectors are perfectly correlated and are only driven by the super-systematic factor, while if δ equals zero, the super-systematic factor has no impact on the firms and they are only jointly affected by the sector-specific risk factor and become independent from each other.

Next, we assume that the protection seller engages in two hedge contracts, covering two credit risky portfolios. These portfolios only differ in sectors, all other settings being equal. The true parameters of each portfolio is $\rho = 20\%$ and $\pi = 1\%$. However, again these parameters are unknown and have to be estimated based on observable losses for $T = 15$ years using the analytical estimation method outlined in Duellmann et al. (2010). For both contracts, the hedge parameter is the sum of the estimated parameters and their standard errors weighted by $\kappa = 75\%$ and the VaR is calculated at confidence level $\alpha = 99.9\%$. For each sample, we calculate the contractual and fair fee for each contract. The protection seller receives $f_c^* = f_c^{*,1} + f_c^{*,2}$ and has to pay on average $f_c = f_c^1 + f_c^2$, where $f_c^{*,i}$ and f_c^i is the contractual and fair fee of hedge contract $i \in \{1,2\}$. We repeat this sampling procedure 10^4 times for $\delta \in \{0\%, 50\%, 75\%, 100\%\}$ and Figure 3.2 summarizes the resulting fees with similar interpretation as in Figure 3.1.

Table 3.2: Diversification benefits. This table reports beneficial diversification effects in case of imperfect correlation from a super-systematic factor to a sectoral risk component. The case $\delta = 1$, i.e., perfect correlation, serves as the base case in relation to values $\delta < 1$. The diversification benefit $\mathcal{D}_c$ is given by $\frac{\mathcal{A}_{(\delta<1)} - \mathcal{A}_{(\delta=1)}}{\mathcal{A}_{(\delta=1)}}$, where $\mathcal{A}$ is the (absolute) area between the fair and contractual fees for two contract types $c = 0, 1$. The computations of $\mathcal{D}_c$ are based on the simulation results summarized in Figure 3.2.

δ	$\mathcal{D}_0$ [%]	$\mathcal{D}_1$ [%]
0.75	−11.38	−11.37
0.50	−25.73	−28.56
0.00	−42.87	−49.03

If δ equals one, both portfolios are solely driven by the super-systematic factor and result in the same default rates. Hence, all estimated parameters and hedge parameters are the same, leading to exactly the same fees. Therefore the figure equals the graph presented in the upper-right hand graph in Figure 3.1, with all values simply doubled. However, with $\delta < 1$ the resulting default rates increasingly differ and, as a result, the corresponding hedge fees.

To quantify and compare potential diversification benefits for different degrees of δ, we consider perfect correlation with $\delta = 1$ as a reference case (i.e., no diversification). The differences from the fair fees (dashed lines) to the corresponding contractual fees (solid lines) depicts the actual parameter error, i.e., the higher the difference, the higher the error. To aggregate the errors over the relative frequency of occurrences, we calculate the (absolute) areas $\mathcal{A}_\delta$ confined by these two lines. Next, we compare the resulting areas from the remaining δ and relate them to the base case given by $\delta = 1$. The results of the diversification benefit $\mathcal{D}_c = \frac{\mathcal{A}_{(\delta<1)} - \mathcal{A}_{(\delta=1)}}{\mathcal{A}_{(\delta=1)}}$ are summarized in Table 3.2.

All six cases exhibit the potential for diversification effects. Even for a relatively high correlation with $\delta = 0.75$, the average (absolute) deviances from f_c to f_c^* diminish by more than 11% for both contract types. This gets more pronounced with decreasing values for δ up to the edge case with $\delta = 0$, where we observe a reduction of premium differences of nearly 50%. Furthermore, for contract type $c = 1$ and the scenario under consideration, there seems to be more potential for diversification effects in comparison to $c = 0$ with decreasing δ. This is important given that the fee is for $c = 1$ more sensitive to parameter misspecification as discussed above.

In conclusion, we see with lower correlation between sectors that the difference of the fair fee and contractual fees decreases. As a result the protection seller can diversify the parameter risk across two protection buyers. Some degree of independence of the premium contracts may additionally emerge, if the protection buyer employs different modeling approaches, differ in implementation details or have differing data quality.

3.3.3 Model sensitivity

One advantage of our parameter risk hedge framework is the applicability of any risk model with tractable loss distributions. To analyze alternative distributional assumptions, we consider generalized ASRF specifications by using more sophisticated distributions for the factors. Empirical asset pricing literature has documented the stylized fact that asset returns possess heavier tails than predicted by the normal distribution (Cont, 2001). Therefore, we compare the Gaussian to the Student-t and the normal inverse Gaussian (NIG) copula model in the framework of homogeneous portfolios. It is advantageous that in all cases, the main model parameters are the same, i.e., the asset (return) correlation ρ and the probability of default π, and share the same interpretation. Besides these commonalities, we treat ν and α representing the heavy-tailedness of the Student-t and NIG distribution as hyperparameters.[11] For both of these loss distributions, the Gaussian occurs as a limit case, letting $\alpha, \nu \to \infty$.

For the Student-t copula model the credit portfolio loss in a given period is modeled by

$$L^T(X, Y, [\rho, \pi]) = \Phi\left(\frac{\sqrt{\frac{X}{\nu}}\, T_\nu^{-1}(\pi) - \sqrt{\rho}\, Y}{\sqrt{1-\rho}}\right), \quad X \overset{\text{i.i.d.}}{\sim} \mathcal{X}_\nu^2, \quad Y \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1),$$

where $\mathcal{X}_\nu^2$ is the chi-squared distribution with ν degrees of freedom and T_ν^{-1} is the inverse of the Student-t distribution, see, e.g., Hamerle and Rösch (2005).

[11] Note, that in the Gaussian copula model, the location and scale parameter of a normal distribution are set to have mean zero and unit variance. Hence, these parameters are also hyperparameters. Standardizing the distributions, and thus reducing the number of parameters to be fitted, is not only convenient but also necessary to allow for the interpretation of ρ as a correlation coefficient in valid statistical terms.

For the NIG copula model the credit portfolio loss in a given period is modeled by

$$L^{\text{NIG}}(M,[\rho,\pi]) = F_{\text{NIG}\left(\sqrt{\frac{1-\rho}{\rho}}\right)}\left(\frac{F^{-1}_{\text{NIG}\left(\frac{1}{\sqrt{\rho}}\right)}(\pi) - \sqrt{\rho}M}{\sqrt{1-\rho}}\right), \quad M \overset{\text{i.i.d.}}{\sim} F_{\text{NIG}(1)},$$

where following the notation of Kalemanova et al. (2007) $F_{\text{NIG}(s)} := F_{\text{NIG}}\left(x; s\alpha, s\beta, -s\frac{\beta\gamma^2}{\alpha^2}, s\frac{\gamma^3}{\alpha^2}\right)$ is the normal inverse Gaussian distribution function. To reduce the number of distribution parameters, β is set to zero which makes the distribution symmetric and setting $\gamma = \sqrt{\alpha^2 - \beta^2}$ makes the distribution having zero mean and unit variance.

To investigate the sensitivity of fees to model choice, we now look at a case with fixed hedge levels. The hedge range can be defined without specification of a model and corresponding parameters. Fees, on the other hand, depend on a model with the corresponding model parameters. Therefore we can calibrate model parameters to a given hedge range. To explore this relation of fair premiums to our models under consideration, we recall the results from Table 3.1 for the Gaussian model and analyze the case given by $\rho = 20\%$ and $\pi = 1\%$ to be consistent with previous sections. With these parameters, the VaR for the Gaussian benchmark yields to VaR = 14.5525% at the $\alpha = 99.9\%$ confidence level. The estimates, assumed to be again underestimated by 25%, are set to $\hat{\rho} = 15\%$ and $\hat{\pi} = 0.75\%$ yielding $\text{VaR}(\hat{\boldsymbol{\theta}}) = 9.0102\%$. This will fix our hedge range and serves as a benchmark to which we will calibrate the other two models. The Student-t and NIG model both have three parameters. By the construction of the ASRF, the probability of default π equals the expectation of the three distributions under consideration. This allows us, given π and VaR to back out precisely the biunique ρ_{impl} implied by the hedge range from the Gaussian benchmark and given a degree of heavy-tailedness ν (α) for the Student-t (NIG) model, respectively. Hence, for a given set of tail measures for the two model alternatives, we infer ρ_{impl} by solving

$$\text{VaR}^G_\alpha(\rho,\pi) = \text{VaR}^T_\alpha(\rho_{impl},\pi,\nu) = \text{VaR}^{\text{NIG}}_\alpha(\rho_{impl},\pi,\alpha) = 14.5525\%,$$

and calculate the corresponding fees f_c for the Student-t and NIG model. We emphasize that for pricing purposes, we only need the hedge parameter and the hedge range from 9.0102% to 14.5525% provided by the Gaussian model.

The range of tail measures we consider is guided by reported values from the literature and own estimations based on physical loss data. For instance, Hull and White (2004) report four degrees of freedom showing a good fit to iTraxx tranche data and Kalemanova et al. (2007) find a $\alpha = 0.4794$ calibrated to iTraxx tranches. These represent 'risk neutral' parameters implied from market expectations which are known to describe much heavier tails in the return distributions than what is typically observed for physical return distributions. Based on rating performance data from Standard & Poor's, Hamerle and Rösch (2005) find approximately 33 degrees of freedom for an application to the rating grade BB, which confirms that 'real world' asset return distributions inferred from historical default rates are not as heavy-tailed as typically result from calibration to market data.

Fitting the two models under consideration to historical one year default rates, we find that for the Student-t model the estimated degrees of freedom ν are about 25, while for the NIG shape parameter α, we find a value about 5.[12] Furthermore, it turns out that the fitted π are quite stable and hardly distinguishable from each other, which provides empirical support holding π fixed according to the value obtained via the Gaussian model. Moreover, we find that, in general, the $\text{VaR}(\boldsymbol{\theta}_h)$ from the Gaussian model cover all $\text{VaR}(\hat{\boldsymbol{\theta}})$ of the other models with more sophisticated factor distributions. Thus, a parameter risk hedge with a Gaussian hedge implicitly insures against model risk with respect to point estimates from the alternative model specifications.

In Table 3.3, we report the implied correlations ρ_{impl} as well as the resulting model based premiums f_c for the prespecified hedge range. The last two columns report the price ratio $Q_c = \frac{f_c - f_c^G}{f_c^G}$ of the Student-t or NIG model to the Gaussian case (denoted by G) for both contractual types $c \in \{0, 1\}$. For the Student-t and the NIG model we consider seven different tail measures each. These range from typical parameter values calibrated to market expectations, over parameters typically estimated from historical loss data, and finally to parameter values approaching the Gaussian limit, letting $\alpha, \nu \to \infty$.

The implied ρ_{impl} shown in Table 3.3 are decreasing relative to the Gaussian benchmark case if the shape parameter for both alternatives, the Student-t and NIG models, are getting smaller leading to heavier tails than the normal. This indicates that the shape parameter and the parameter ρ are interrelated since they both influence the skewness of the loss distributions.

Comparing the relative changes Q_c to the Gaussian model, we see different signs for the two model alternatives. The heavier the tails for a Student-t

[12]The Student-t and NIG model parameters are estimated with a correlated binomial model instead of the ASRF, compare Section 3.3.4 for the data and methodology. Estimation results for the model alternatives are reported in Section 3.C of the Appendix.

Table 3.3: Model sensitivity of fair fees. This table reports correlation parameters ρ_{impl} for the Student-t and NIG model implied from a Gaussian benchmark model (denoted by G) given by $\pi = 1\%$, $\rho = 20\%$, and $\text{VaR}^G_{\alpha=99.9\%}(\rho, \pi) = 14.5525\%$. Since the expected value equals the probability of default in all three models, the ρ_{impl} are the biunique outcomes for given π and VaR for seven prespecified degrees of heavy tailedness encoded by ν and α. The fourth and fifth column report the fair premiums for two contractual types $c \in \{0, 1\}$ in basis points. The last two columns report the relative difference to the benchmark Gaussian fair fees $Q_c = \frac{f_c - f_c^G}{f_c^G}$. The entries are grouped into three panels. The upper panel reports the values from the Gaussian benchmark model. The lower two panels each report the results of a Student-t and NIG copula model for different sets of heavy tailedness, respectively.

Model		ρ_{impl} [%]	π [%]	VaR_α [%]	f_0 [bp]	f_1 [bp]	Q_0 [%]	Q_1 [%]
Gauss	$\nu = \alpha$							
	∞	20	1	14.5525	0.9294	1.4837	0	0
Student-t	ν							
	100	17.8379	1	14.5525	0.9307	1.4850	0.1390	0.0871
	50	15.7001	1	14.5525	0.9364	1.4907	0.7516	0.4708
	25	11.5140	1	14.5525	0.9606	1.5148	3.3495	2.0983
	20	9.4812	1	14.5525	0.9800	1.5343	5.4429	3.4097
	15	6.2099	1	14.5525	1.0262	1.5804	10.4085	6.5203
	10	0.2286	1	14.5525	1.1956	1.7498	28.6358	17.9388
	5	0.0010	0.1807[a]	14.5525	0.4750	1.0293	−48.8920	−30.6283
NIG	α							
	10	19.6820	1	14.5525	0.9044	1.4587	−2.6909	−1.6856
	5	18.8177	1	14.5525	0.8411	1.3954	−9.5008	−5.9517
	4	18.2446	1	14.5525	0.8025	1.3568	−13.6541	−8.5536
	3	17.1772	1	14.5525	0.7369	1.2911	−20.7199	−12.9799
	2	14.9725	1	14.5525	0.6217	1.1760	−33.1091	−20.7411
	1	10.1662	1	14.5525	0.4274	0.9816	−54.0172	−33.8389
	0.5	6.5030	1	14.5525	0.2823	0.8365	−69.6314	−43.6204

[a] For the Student-t copula model with $\nu = 5$, it is not possible to reach the true VaR = 14.5525% without lowering π even for zero correlation.

copula model, fair premiums are increasing.[13] We find the reverse for the NIG model. Here, the heavier the tails of the factors become, fair premiums are decreasing.

However, for both models and parameter settings under consideration, we do not observe price differences exceeding about ±10% for the same hedge range and plausible parameter constellations.[14] Overall, we can conclude that the fair fees are quite robust on these alternative model specifications for the given hedge interval.

3.3.4 Historical defaults, data quality, and estimation methods

We now apply our framework to historical one year default rates from the Moody's annual default study 2015 (Ou et al., 2015). The use of one year default rates is supported by, e.g., the requirements of the current Basel Accords stating that "PD estimates must be a long-run average of one-year default rates for borrowers in the grade" (BCBS, 2006). We consider a 30 year time span ranging from 1985 to 2014. The number of rated companies documents an increase from 1,608 in 1985 to 5,340 in 2014. On average, we observe 3,590 rated companies per year. Table 3.4 provides descriptive statistics of the default rates for each rating grade.

The mean of the historical one year default rates increases with rating grade, while the ratio of the mean and the standard deviation decreases. This observation supports the previous finding that high rated debt may be more prone to parameter risk than lower rated debt.

In contrast to the asymptotic ideal of the ASRF in the simulation study above, real world portfolios are not perfectly fine grained. Therefore we consider a Gaussian finite homogeneous pool model. Additionally, in each rating grade, there is at least one year with zero defaults, which makes the asymptotic case inapplicable. Another data driven consequence is that we employ two different estimation methods, since the low number of default data for high rated risk buckets is too sparse for conventional statistical inferences.

[13]For $\nu = 5$, we observe a departure from this pattern. The change of sign in this particular case is due to the inherent limitations of the Student-t model dynamics being incapable to meet the target VaR without simultaneously lowering the probability of default. Thus, with a shifted expected loss in comparison to the other parameter settings, we are essentially observing a fundamentally different model.

[14]Given that we are primarily concerned with the risk management of 'real world' credit losses during our analysis, we consider $\nu \approx 25$ and $\alpha \approx 5$ being plausible parameter constellations. The much lower values of ν and α, typical for calibrated values from market expectations, would not fit to the historical default rates any more as they rather tend to the Gaussian limit.

Table 3.4: Summary statistics of one year corporate default rates. This table reports descriptive statistics for historical one year default rates in percent from the Moody's annual default study 2015 (Ou et al., 2015). The sample period is from 1985 to 2014.

	Aaa	Aa	A	Baa	Ba	B	Caa–C
Mean	–	0.043	0.048	0.191	1.097	4.754	18.802
SD	–	0.166	0.114	0.319	1.204	4.179	13.249
Min.	–	–	–	–	–	–	–
Med.	–	–	–	0.033	0.732	4.268	15.201
Max.	–	0.687	0.518	1.147	4.932	16.104	61.905
ø	113	418	836	786	485	724	228
#	–	6	14	46	151	762	1,035

ø (#) denote the average (total) number of rated (defaulted) debt

Guidance how to distinguish between low default and 'normal' default debt pools comes from, e.g., the FCA (2016) in BIPRU 4.3.95 "[...] a firm's internal experience of exposures of a type covered by a model or other rating system is 20 defaults or fewer [...]". For this reason, we apply the *most prudent estimation principle* detailed in Pluto and Tasche (2005) for the rating grades Aaa to A. For the remaining rating grades Baa to Caa–C, we apply the maximum likelihood estimation method outlined in Frey and McNeil (2003).

For the low default rating grades we infer the probability of default as an upper confidence bound for a given confidence level γ such that the maximum value of $\hat{\pi}$ from the set of all admissible values of π fulfills the inequality

$$1 - \gamma \leq \mathbb{P}[\text{No more then } k \text{ defaults observed}].$$

Here, we infer the upper bound of a confidence interval, $\overline{\pi}_\gamma$, with the extension of the most prudent estimation principle to correlated default events in a multi-period setting. One restriction is that this estimation principle solely calibrates the probability of default. Thus we have to set ρ. In BCBS (2006) ρ is floored and capped by 12% and 24%, respectively. That is why we determine the most prudent probabilities of default for these two values of ρ. To define the hedge range we apply the estimation principle to two different confidence levels of γ. As a proxy to an $\hat{\pi}$ estimate, we consider a value of $\gamma = 50\%$ and π_h is approximated by $\gamma = 90\%$ to add some conservatism to our estimate and serves as the hedge parameter. Table 3.5 summarizes the estimated $\overline{\pi}$ as functions of the confidence level γ. The third and fourth column show the estimates for π corresponding to two correlation assumptions for

Table 3.5: Contractual fees for historical one year default rates for low default risk buckets. This table reports most prudent estimates for low default risk buckets Aaa to A. Upper confidence bounds for probability of default estimates for two confidence levels $\gamma \in \{50\%, 90\%\}$ are denoted by $\overline{\pi}_\gamma$. Contractual fees f_c^* are for contract types $c \in \{0,1\}$. The sample period is from 1985 to 2014.

Rating	ρ [%]	$\overline{\pi}_{50\%}$ [%]	$\overline{\pi}_{90\%}$ [%]	$\mathcal{R}_{\alpha=99.9\%}$	$\mathcal{R}(\overline{\pi}_{50\%})$ [%]	$\mathcal{R}(\overline{\pi}_{90\%})$ [%]	f_0^* [bp]	f_1^* [bp]
Aaa	12	0.0566	0.0818	VaR	1.7699	2.6549	0.1175	0.1512
				CVaR	2.6511	3.1240	0.0005	0.0185
	24	0.0560	0.1119	VaR	3.5398	4.4248	0.0699	0.1553
				CVaR	4.4890	6.6875	0.0703	0.1461
Aa	12	0.0566	0.0879	VaR	1.4354	1.6746	0.0167	0.0394
				CVaR	1.7332	2.2842	0.0160	0.0353
	24	0.0611	0.1165	VaR	2.6316	4.3062	0.1681	0.3149
				CVaR	3.9173	5.9978	0.0748	0.1413
A	12	0.0600	0.0987	VaR	1.1962	1.7943	0.0678	0.1140
				CVaR	1.5943	2.2577	0.0237	0.0485
	24	0.0646	0.1306	VaR	2.6316	4.4258	0.1948	0.3696
				CVaR	3.9179	6.3013	0.0929	0.1747

each rating grade for two different levels of γ. These two upper confidence bounds proxy the parameter estimate and hedge parameter. They also span the hedge range shown in column six and seven for both the VaR and CVaR at confidence level $\alpha = 99.9\%$. The last two columns report the contractual fees f_c^* for a parameter risk hedge in basis points.

Table 3.6 summarizes the results of the maximum likelihood approach (Frey and McNeil, 2003) for the rating grades Baa to Caa–C. The second and third column show the estimates for ρ and π for each rating class and their corresponding standard errors $\hat{\sigma}(\hat{\rho})$ and $\hat{\sigma}(\hat{\pi})$ in parenthesis. We hedge the estimates by a weighted standard deviation with $\kappa = 0.75$. Assuming that the estimator follows a normal distribution, the true parameter would then be less then or equal to this hedge level with a probability of 77.34%. The fifth and sixth column show the VaR and CVaR for a confidence level at 99.9% at the estimates $\hat{\boldsymbol{\theta}}$ and the hedge level $\boldsymbol{\theta}_h$. Finally, the seventh and eights column report the contractual fees f_c^* for a parameter hedge in basis points.

In Table 3.5 and Table 3.6 the estimates $\hat{\pi}$ are increasing with decreasing rating grades and compare well to the average default rates reported in Table 3.4. Due to the most prudent estimation principle, the probabilities of default for the first three ratings grades do not differ substantially and are

Table 3.6: Contractual fees for historical one year default rates. This table reports maximum likelihood estimates $\hat{\boldsymbol{\theta}} = [\hat{\rho}, \hat{\pi}]$ of each risk bucket for the rating grades Baa to Caa–C. The standard errors (in parentheses) for the parameter estimates are obtained by inverting the negative Hessian of the log-likelihood at the estimates. Risk measures VaR and CVaR at the estimate $\hat{\boldsymbol{\theta}}$ and hedge level $\boldsymbol{\theta}_h = \hat{\boldsymbol{\theta}} + 0.75\,\hat{\sigma}(\hat{\boldsymbol{\theta}})$ are evaluated for a confidence level of $\alpha = 99.9\%$. Contractual fees f_c^* are for contract types $c \in \{0, 1\}$. The sample period is from 1985 to 2014.

Rating	$\hat{\rho}$ [%]	$\hat{\pi}$ [%]	$\mathcal{R}_{\alpha=99.9\%}$	$\mathcal{R}(\hat{\boldsymbol{\theta}})$ [%]	$\mathcal{R}(\boldsymbol{\theta}_h)$ [%]	f_0^* [bp]	f_1^* [bp]
Baa	14.7791**	0.1952***	VaR	3.4351	5.4707	0.2448	0.4432
	(6.3334)	(0.0657)	CVaR	4.5558	7.3357	0.1434	0.2407
Ba	13.9788***	1.1101***	VaR	11.5464	15.4639	0.3932	0.7657
	(4.4327)	(0.2522)	CVaR	13.9618	18.7691	0.2238	0.3928
B	23.4588***	4.9322***	VaR	43.2320	51.9337	0.7803	1.6501
	(5.1883)	(1.0612)	CVaR	49.2976	58.5884	0.3973	0.7555
Caa–C	13.8124***	17.9317***	VaR	60.9649	67.9825	0.8052	1.4774
	(3.8477)	(2.0877)	CVaR	65.0747	72.2711	0.3858	0.6766

*, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively

increasing for the upper correlation level. For the rating grades Baa to Caa–C the asset correlation is also estimated and, with the exception of rating grade B, we observe that the asset correlations $\hat{\rho}$ are higher for higher rated risk buckets. For these rating classes the two risk measures at the estimates $\mathcal{R}(\hat{\boldsymbol{\theta}})$ are strictly increasing with declining credit quality.[15] However, we do not see such behavior for the low default rating grades. The reasoning behind this is, that both risk measures in the correlated binomial model takes into account the absolute number of entities in each rating grade under consideration. Comparing Table 3.4, the highest rating grade has 113 entities, Aa rating grade has 418 entities and A has 836. With a greater number of entities idiosyncratic risk increasingly diversifies and consequently the risk measures decrease. However, if all three rating grades had the same number of entities, we would observe an increasing effect in the risk measures for these rating grades.

[15]The overall magnitudes of the risk measures may seem somewhat high, however, as noted above we do not consider nonzero recovery rates. On average, annual recovery is roughly about 38–42%, however, it is well documented that recovery rates tend to be quite volatile over time and are subject to the kind of provided bond collateral (secured, unsecured, subordinated), see, e.g., Altman et al. (2005).

Given the sparsity of default events in the higher rated risk buckets, estimated standard errors are large in relation to the parameter estimates. By comparison, the (relative) standard errors for $\hat{\rho}$ are more pronounced than for $\hat{\pi}$. Note, however, the parameter sensitivity in π of the hedge premiums is about twice than for ρ.[16] Thus, there is no clearcut distinction between the importance of the parameters under consideration.

For the resulting risk measures, the values for the hedge parameters are higher than the values for the estimates by construction. If a financial firm were to interpret the hedge level as conservative parameter estimates and fully provide capital reserves to reduce the probability of (unexpected) losses, this would come at high costs, e.g., compare for the Baa rated risk bucket a $\mathrm{VaR}_\alpha(\boldsymbol{\theta}_h) = 5.4707\%$ to a $\mathrm{VaR}_\alpha(\hat{\boldsymbol{\theta}}) = 3.4351\%$. If, instead, the same firm engages in a parameter hedge, it could achieve the same security level for a periodic contractual premium payment amounting to $f_0^* = 0.2448\,\mathrm{bp}$ or $f_1^* = 0.4432\,\mathrm{bp}$.[17] Further, we see that the f_c^* are consistently higher for VaR than for CVaR, which is in line with the findings in Section 3.3.1.[18] The calculated contractual hedge fees—based on the hedge level determined by the parameter estimation errors—tend to result in higher absolute premiums with lower rating grades. We find, however, higher relative premiums $f_c^*/\mathcal{R}(\boldsymbol{\theta}_h)$ for higher rating grades. The relative premiums (in percent, $c = 0$) for the Baa through Caa–C rating grades for the VaR are 0.0447, 0.0254, 0.0150 and 0.0118, respectively. Thus, we conclude, higher rated risk buckets are more prone to parameter risk than lower rated risks buckets. The same observation holds also for CVaR and the two contractual payoff types under consideration. This observation is in contrast to the findings from Section 3.3.1, where we assume the same (relative) error leading to higher relative fees for higher rating grades are smaller.

Summarizing, we find strong empirical support that in practical applications high rated financial instruments are more prone to parameter risk than lower rated instruments, and—as a result—are more costly in relative

[16] Compare Section 3.B in the Appendix for an illustration.

[17] Though we regard A a low default risk bucket, we also estimate the corresponding maximum likelihood estimates (standard errors) yielding $\hat{\rho} = 20.0577\%$ (13.2795%), $\hat{\pi} = 0.0567\%$ (0.0310%). The estimated correlation is well in between the 12% and 24% bounds we consider for the low default estimations. The hedge range is from $\mathrm{VaR}_\alpha(\hat{\boldsymbol{\theta}}) = 1.9139\%$ to $\mathrm{VaR}_\alpha(\boldsymbol{\theta}_h) = 4.1866\%$. The upper hedge level is now more than double than the risk measure at the estimate. The hedge premiums would yield to $f_0^* = 0.3153\,\mathrm{bp}$ and $f_1^* = 0.5276\,\mathrm{bp}$.

[18] For Aaa rated debt, however, we observe a peculiarity. The differences from f_c^* for CVaR are rather high since both, the ratio of the hedge levels and the absolute range are quite pronounced. The fee for f_0^* is extremely low, whereas for f_1^* it seems more reasonable. This is traced back to the fact that f_0^* merely insures the range from 2.6511% to 3.1240% which can only happen if three entities default.

terms to be hedged against parameter risk. Further, we exemplify different estimation methods and highlight the significant effects a departure from the ASRF ideal might have on real world credit portfolios.

3.4 Conclusion

The possibility that inferred model parameters deviate from the true model parameters—termed parameter risk—may significantly affect calculated risk measures. We introduce a framework allowing economic agents to hedge this parameter risk. It is easily applicable to arbitrary types of risk, is independent from the model type, chosen hedge interval, and works for any risk measure.

We show that the possibility to hedge potential parameter errors up to a certain level may drastically reduce costs in comparison to provisions. This may have important economic implications in light that additional capital requirements may influence institutional's behavior and potentially provide adverse incentives. Further, we demonstrate that prospective protection sellers may get into the unique position to diversify parameter risk. This may have significant policy implications as the consequences from beneficial diversification might provide positive macro side effects to an economy as a whole.

Future applications could include other loss categories than credit risk like, e.g., market risk or operational risk. Additionally, it would be of great relevance to compare different estimation methods and survey feasible approaches to identify appropriate hedge levels. The framework may further be a promising addition to the toolset for model validation, discriminate between parameter and model risk or analyze effects of hedged parameters on stress testing.

3.A Parameter risk in the asymptotic single risk factor model

Underpinning the Basel internal ratings-based (IRB) approach (BCBS, 2006), the ASRF model is well known to academics and practitioners. Its foundation and derivation is given by Vašíček (1987, 1991, 2002) and for rigorous treatments of this model compare Gordy (2000, 2003).

To gain intuition about parameter risk in the ASRF model, we perform the following simulation study. For the asset correlation, we assume $\rho = 20\%$, which lies in between the lower and upper bounds found in the Basel Accords. The probabilities of default are $\pi \in \{0.2\%, 1\%, 5\%\}$, roughly corresponding to historical one year default rates of Moody's Baa, Ba, and B rated corporates (compare Table 3.4). To analyze effects of differing sample sizes, we consider three time horizons ranging from a short period of $T = 7$ years, over $T = 15$, up to a time period of $T = 30$ years. For each case, we sample credit losses according to Equation (3.4). During each iteration, we employ the estimation method described in Duellmann et al. (2010) providing analytical solutions for the estimates $\hat{\boldsymbol{\theta}} = [\hat{\rho}, \hat{\pi}]$ with corresponding standard errors $\hat{\sigma}(\hat{\boldsymbol{\theta}})$. Based on these results, we define the hedge parameter as the sum of the estimate and its standard error weighted by a confidence factor $\kappa = 75\%$ (arguments on the choice and definition of such a hedge parameter follow below). For each estimate $\hat{\boldsymbol{\theta}}$ and hedge parameter $\boldsymbol{\theta}_h$ we calculate the value-at-risk (VaR) at the 99.9% confidence level. This procedure is repeated 10^5 times for each parameter set and time horizon. Figure 3.3 summarizes the results.

The upper and middle panel illustrate how parameter risk affects estimates of the asset correlation $\hat{\rho}$ and the probability of default $\hat{\pi}$ for three different assumptions of $\pi \in \{0.2\%, 1\%, 5\%\}$ and time horizons $T \in \{7, 15, 30\}$, respectively. Each subplot depicts from the left to right the estimates ($\hat{\theta}/\theta$, dark gray), standard errors ($\hat{\sigma}(\hat{\theta})/\theta$, gray) and hedge parameters (θ_h/θ, light gray) normalized by the true parameter $\theta \in \{\rho, \pi\}$. The lower panel displays the normalized VaR for the estimates ($\text{VaR}(\hat{\boldsymbol{\theta}})/\text{VaR}(\boldsymbol{\theta})$, dark gray) and the hedge parameters ($\text{VaR}(\boldsymbol{\theta}_h)/\text{VaR}(\boldsymbol{\theta})$, light gray).

In each repetition, the sampled loss is a random event. Therefore the estimates $\hat{\boldsymbol{\theta}}$, standard errors $\hat{\sigma}(\hat{\boldsymbol{\theta}})$, hedge parameters $\boldsymbol{\theta}_h$ and value-at-risk are stochastic, and their distributions describe the parameter risk. These parameter risk distributions are displayed in the form of Tukey box plots where the circles depict the mean value in Figure 3.3. For instance, for the shortest period of 7 years and the lowest probability of default of 0.2% the median (mean) of the asset correlations' parameter risk distribution is roughly 82% (87%) of the true value. The lower and upper quartiles range between 56% and 112%. More precisely, in 66% of all repetitions, the estimate $\hat{\rho}$ results in an underestimation of the true asset correlation ρ.

With decreasing time span T, the range of parameter estimates, their standard errors for both the correlation and probability of default as well as resulting VaR increase. The true parameters are—in the median—underestimated. On average, however, the inferred probabilities of default overestimate the true probability of default, which is caused by few, but severe, outliers. Whereas the correlation estimates are below the true parameter, both on average and in the median. In these cases, the estimated VaR figures underestimate the value-at-risk, both on average and in the median. The described effects are particularly pronounced for seemingly safe risk buckets linked to low probabilities of default. Summarizing, the ASRF model displays considerable parameter risk, at least on average, for short time horizons and low probabilities of default.

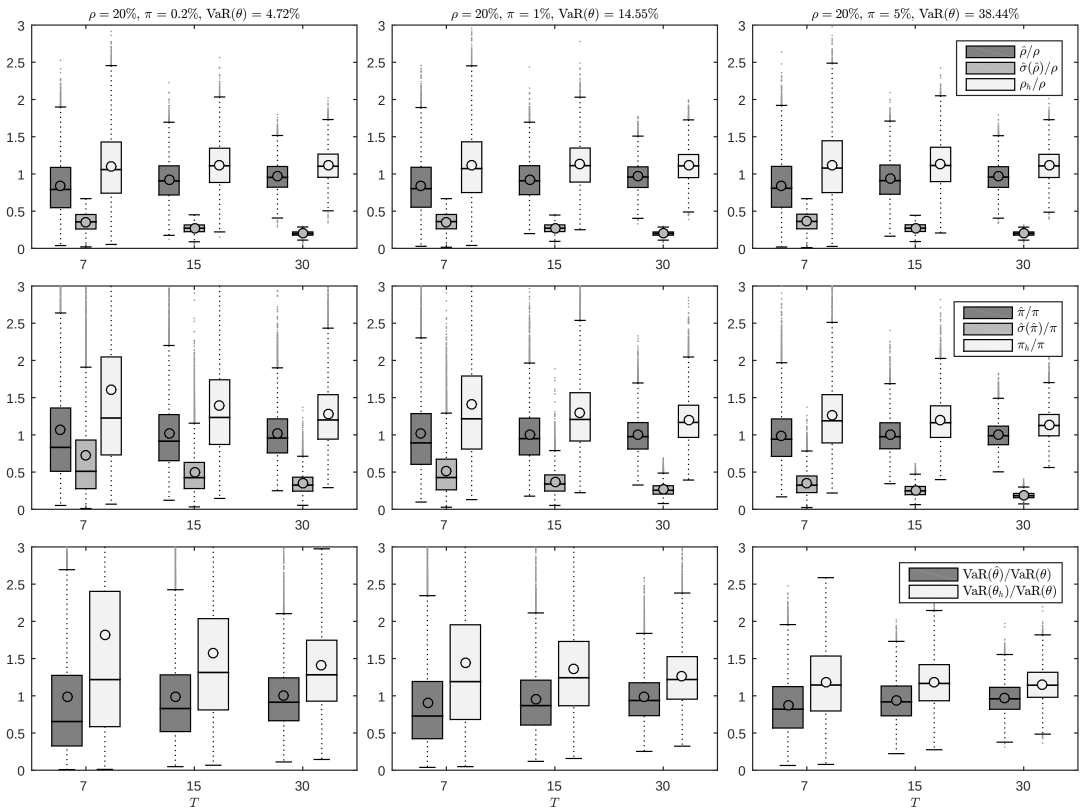

Figure 3.3: Impact of parameter risk on probability of default, correlation, and value-at-risk. This figure presents how parameter risk affects estimates of the correlation ρ (first row), probability of default π (second row), and each with the corresponding value-at-risk measure (third row) at confidence level $\alpha = 99.9\%$. For all three columns, the true correlation is given by $\rho = 20\%$. The true probabilities of default are increasing from the left column with $\pi = 0.2\%$, $\pi = 1\%$ in the middle column, and $\pi = 5\%$ in the right column. Each subplot individually reports the normalized estimation results for three time periods given by $T \in \{7, 15, 30\}$ years. Hedge parameters are defined according to Equation (3.7) by $\boldsymbol{\theta}_h = \hat{\boldsymbol{\theta}} + \kappa\, \hat{\sigma}(\hat{\boldsymbol{\theta}})$, where $\hat{\boldsymbol{\theta}}$ denote the estimate and $\hat{\sigma}(\hat{\boldsymbol{\theta}})$ denote the corresponding standard error given the estimation approach in Duellmann et al. (2010). The confidence factor is set to $\kappa = 0.75$. The simulation is repeated 10^5 times. The boxes are confined by the lower and upper quartiles and contain horizontal lines which draw the median values. Vertical extensions depict 1.5 times the interquartile range from the median, values outside these fences (whiskers) are considered outliers and plotted individually with a dot. Additionally, the circles within the box plots depict the mean values.

Table 3.7: Number of exceedances for value-at-risk at estimates and hedge parameters. This table reports the relative number of exceedances (%NoE) for given confidence levels, i.e., $\frac{\text{NoE [\%]}}{1-\alpha}$, with $\alpha \in \{99\%, 99.9\%\}$. The true correlation is $\rho = 20\%$. The upper panel reports the %NoE given the NoE are based on the value-at-risk at the estimates VaR($\hat{\boldsymbol{\theta}}$). The lower panel reports the %NoE given the NoE are based on the value-at-risk at the estimates plus one standard error times the confidence factor $\kappa = 0.75$, i.e., VaR($\boldsymbol{\theta}_h$) is determined via Equation (3.7). The simulation is repeated 10^6 times.

	$\alpha = 99\%$			$\alpha = 99.9\%$		
π [%]	$T = 7$	$T = 15$	$T = 30$	$T = 7$	$T = 15$	$T = 30$
Estimates						
0.2	4.5298	2.3582	1.6118	18.3455	5.9541	2.8357
1.0	4.5246	2.3475	1.6008	18.3932	5.9074	2.8150
5.0	4.5208	2.3536	1.6144	18.3488	5.8908	2.8131
Hedge, $\kappa = 0.75$						
0.2	2.3785	1.2106	0.8991	8.3860	2.4541	1.2302
1.0	2.2872	1.1594	0.8643	8.1334	2.3124	1.1748
5.0	2.1807	1.1043	0.8357	7.7846	2.2053	1.1223

Next, we illustrate the effects of parameter risk for individual scenario realizations. This is particularly relevant from the perspective of a single financial institution, since they have to deal with one sole VaR estimation. Such an estimate could be seen as the result of a single realization of the VaR distributions for the estimates depicted in the lower panel of Figure 3.3 by the dark gray box plots. Therefore we examine whether a simulated loss for subsequent periods under the true default generating process is greater than the anticipated VaR or not. These number of exceedances (NoE) are then accumulated. If the true VaR were anticipated at all iterations, the NoE must equal $1 - \alpha$ by definition of the VaR. We compare the relative NoE (%NoE) with the NoE permitted by the true model. Hence, if the ratio is greater than one, the VaR is insufficient to cover the losses encoded by α. Table 3.7 reports the %NoE for confidence levels $\alpha = 99\%$ and $\alpha = 99.9\%$ for the estimates and hedge parameters.

We find that the probability of default has a limited impact on the %NoE. This is in contrast to the time dimension. Reported %NoE in the upper panel based on the estimates show a substantial increase in the %NoE for decreasing T. This increase in %NoE gets worse the further in the tail one is trying to predict, e.g., the %NoE for $T = 7$ is roughly 2.5 (6.5) times higher than the %NoE at $T = 30$ at the $\alpha = 99\%$ ($\alpha = 99.9\%$) level, respectively. For example, the 18-fold tearing of the true allowed %NoE for $T = 7$ interprets as follows. Given the chosen estimation method and an estimation horizon of seven years, the resulting estimated VaR lead to 18 times more breachings of the intended VaR level than it would be the case if instead using the true parameters. This %NoE implies a VaR of an effective confidence level of 98.16% instead of 99.9%. That is, an estimation approach, which would imply an estimation approach leading in 1/3 of the cases to the true VaR, in 1/3 of the cases to an underestimation by 25% and in 1/3 of the cases to an overestimation, results in the same %NoE. This essentially signifies that an equiprobable overestimation is by no means able to even out an equiprobable underestimation particularly in the rare cases.

This analysis is further used to motivate an appropriate hedge level. Given the risk measure under consideration is monotone increasing in the parameters (for this chosen model type), the hedge parameter is defined as

$$\boldsymbol{\theta}_h = \hat{\boldsymbol{\theta}} + \kappa\, \hat{\sigma}(\hat{\boldsymbol{\theta}}), \tag{3.7}$$

where κ is a confidence factor describing the influence of estimation error $\hat{\sigma}(\hat{\boldsymbol{\theta}})$ on the hedge level. We find $\kappa = 0.75$ to be a good candidate to encode the influence of estimation errors on hedge parameters to reach a fair level of conservatism. That is, with the resulting hedge level, the %NoE are roughly cut by half and for $T = 30$ is almost approaching the $1 - \alpha$ value. Though our choice may seem rather *ad hoc*, more elaborate methods are conceivable (compare, e.g., Tarashev (2010)). For instance, the confidence factor could be inferred from the distribution of (robust) standard errors, or could be calibrated in a QIS like impact study from regulatory authorities. The FCA (2016) in BIPRU 4.3.88 explicitly demands that “A firm must add to its estimates a margin of conservatism that is related to the expected range of estimation errors. Where methods and data are less satisfactory and the expected range of errors is larger, the margin of conservatism must be larger”. Similarly, the confidence factor could also be seen as reflecting regulators’ confidence in the employed model, compare, e.g., BCBS (2013). For our analysis, we rely on Equation (3.7) for its analytical tractability and its clearcut interpretation.

3.B Parameter sensitivity of fair hedge premiums and relation of fair hedge premiums to contractual fees

The isolines in Figure 3.4 connect pairs of parameter estimates sharing the same fair hedge premium f_c for the two risk measures VaR and CVaR at confidence level $\alpha = 99.9\%$, as well as the two contract types $c \in \{0, 1\}$. A $-45°$

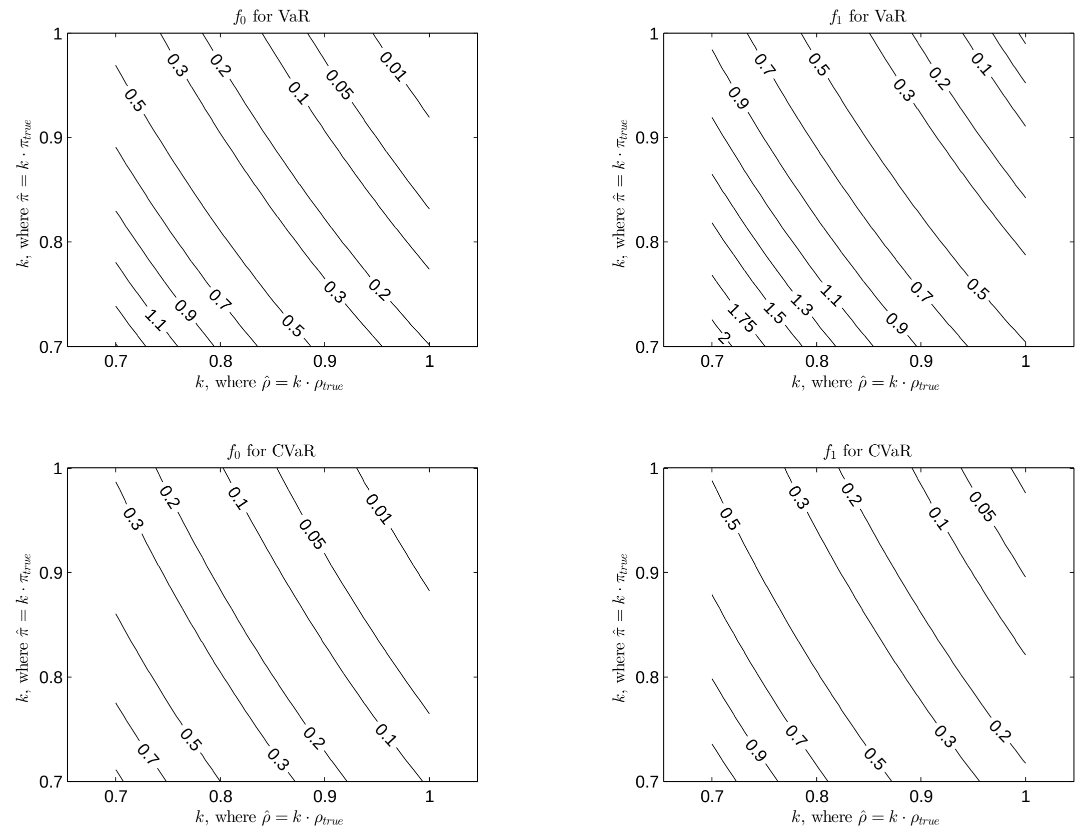

Figure 3.4: Parameter sensitivity of fair hedge premiums. This figure plots the fair hedge premiums $f_c, c \in \{0,1\}$ in basis points, for the risk measures VaR and CVaR, given the true parameter constellation $\rho = 20\%$, $\pi = 1\%$, and confidence level $\alpha = 99.9\%$. The estimate $\hat{\rho}$ vary along the abscissa, while the estimate $\hat{\pi}$ vary along the ordinate.

isoline would represent a situation where the two estimates $[\hat{\rho}, \hat{\pi}]$ have equal influence on the fair premiums, however, we can see that the sensitivity in π is about twice as high than it is for ρ. This pattern is stable for all four cases, otherwise the contours share pretty similar and stable shapes, only the levels change, i.e., f_1 are strictly higher than for f_0 for both risk measures as expected.

To get an impression how the fair hedge premiums with known true values f_c relate to contractual fees f_c^* we look at following setup. For example for $\alpha = 99.9\%$, we assume $\rho = 20\%$ and $\pi = 1\%$. This results to a $\text{VaR}_\alpha^G(\rho, \pi) = 14.55\%$. Further assume an estimate $[\hat{\rho}, \hat{\pi}] = [15\%, 0.75\%]$ implying a $\text{VaR}_\alpha^G(\hat{\rho}, \hat{\pi}) = 9.01\%$.

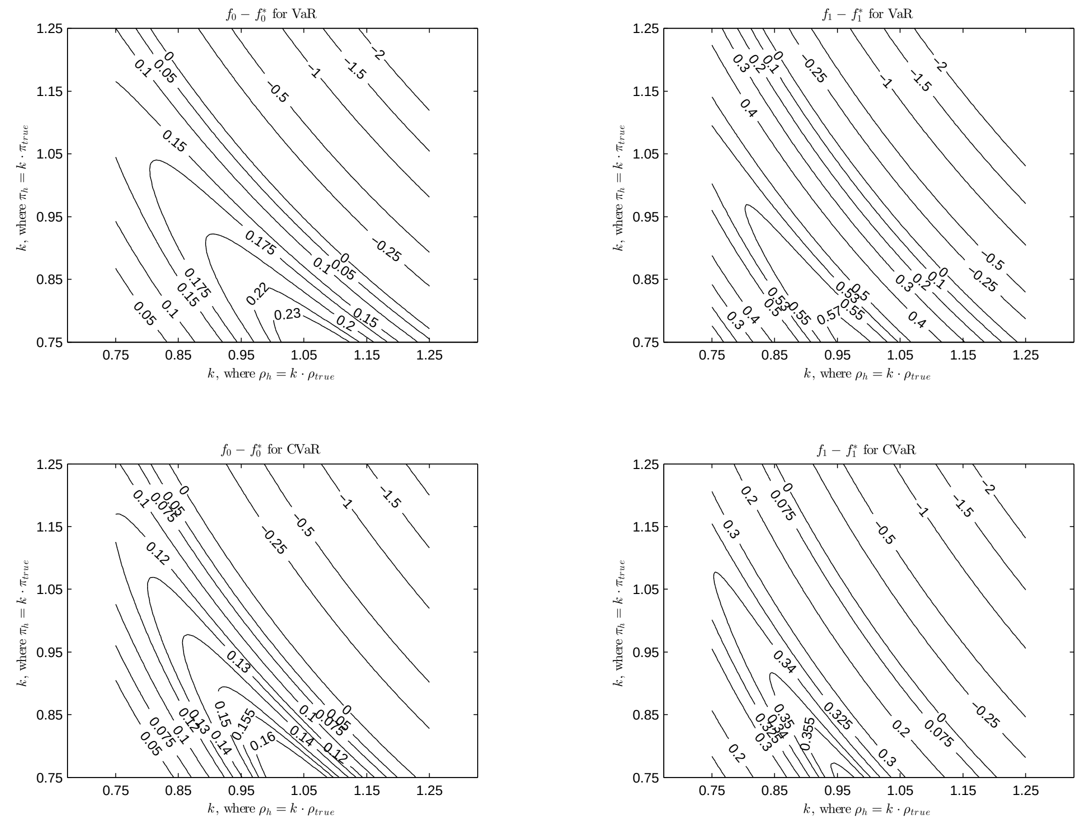

Figure 3.5: Relation of fair hedge premiums to contractual fees. This figure plots the differences from fair hedge premiums f_c to contractual fees f_c^* for the risk measures VaR and CVaR and for both payoff types $c \in \{0, 1\}$ given the true parameter constellation $\rho = 20\%$, $\pi = 1\%$, and confidence level $\alpha = 99.9\%$ in basis points. The assumed estimates are given by $\hat{\rho} = 15\%$ and $\hat{\pi} = 0.75\%$, while the hedge parameter ρ_h vary along the abscissa and π_h along the ordinate.

The results of this setup are depicted in Figure 3.5. Clearly, if $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}_h$, then $f_c = f_c^* = 0$, because there is no hedge contract. Further, in cases with a positive effective error and $\boldsymbol{\theta}_h = \boldsymbol{\theta}$, we see $f_c = f_c^* > 0$. Even if $\boldsymbol{\theta}_h \neq \boldsymbol{\theta}$, it is possible that $f_c = f_c^* > 0$. But in this case, of course, the fee is not the true fee, because the hedge level changed. For $\hat{\boldsymbol{\theta}} \leq \boldsymbol{\theta}_h \leq \boldsymbol{\theta}$ the difference is $f_c - f_c^* \geq 0$, while for $\boldsymbol{\theta}_h \geq \boldsymbol{\theta}$ the difference is $f_c - f_c^* \leq 0$. Further, we note a positive maximum for the difference $f_c - f_c^*$ around $\hat{\rho} < \rho_h < \rho$ for small π_h.

The protection buyer has a strong incentive to engage in a hedge, if $f_c - f_c^* \geq 0$, which holds particularly for the quadrant $\hat{\boldsymbol{\theta}} \leq \boldsymbol{\theta}_h \leq \boldsymbol{\theta}$. The incentive for selling protection is diametral, i.e., is highest for $f_c - f_c^* \leq 0$ in the area described by $\boldsymbol{\theta}_h \geq \boldsymbol{\theta}$.

3.C Alternative model specifications

For model inference, we use the same maximum likelihood approach and data detailed in Section 3.3.4. In a first step, we fit the model parameter ρ and π simultaneously with the hyperparameter ν and α. For the Student-t model the estimated degrees of freedom ν is about 25, while for the NIG shape parameter α, we find a value about 5. This results clearly indicate that the asset return distributions inferred from historical default rates are not as heavy tailed as typically result from calibration to market data. Further, we note that the estimates of the hyperparameter and the correlation parameter ρ are highly correlated. A possible explanation could be that they both heavily influence the skewness of the loss distributions. For this reason, we fix the parameters to $\nu = 15$ and $\alpha = 3$. Thus, we force the models to exhibit slightly more heavier tails than supported by empirical data to better exemplify the impact of heavy tailed models. Then, subject to these fixed hyperparameter, we repeat the estimation of ρ and π. Finally, based on the estimation errors of the parameters we then define the hedge level $\boldsymbol{\theta}_h = \hat{\boldsymbol{\theta}} + 0.75\,\hat{\sigma}(\hat{\boldsymbol{\theta}})$, obtain the corresponding $\text{VaR}_{\alpha=99.9\%}$ and calculate the contractual fees. The results of this estimation procedure are reported in Table 3.8.

The Gaussian, Student-t and NIG models show considerable variation in the resulting VaR for each risk bucket. With decreasing rating grade the relative differences between the VaR are decreasing, too. Overall, the Student-t and NIG models typically lead to higher VaR which seems reasonable, since with heavier tails, *ceteris paribus* the VaR should increase. For the speculative rating grades B and Caa–C, we see an exception for the Student-t model.[19] In general, the parameters appear to be plausibly estimated. Correlation estimates $\hat{\rho}$ around zero for the first two risk buckets Baa and Ba, may seem somewhat unusual, are, however, well supported by empirical evidence, see, e.g., Hamerle and Rösch (2005). Further, we find that, in general, the $\text{VaR}(\boldsymbol{\theta}_h)$ from the Gaussian model cover all $\text{VaR}(\hat{\boldsymbol{\theta}})$ of the other models with more sophisticated factor distributions. Thus, a parameter risk hedge with a Gaussian hedge implicitly insures against model risk with respect to point estimates from the alternative model specifications. Analyzing the contractual premiums of a parameter risk hedge for real default data for a number of loss models show that the prices are quite robust, share similar properties and consistent behavior. On the other hand, different models typically coincide with different hedge levels or parameter ranges.

[19] This exception mirrors the issues with the Student-t model reported in Table 3.3 and briefly discussed in footnote 13.

Table 3.8: Contractual fees for historical one-year default rates for alternative model specifications. This table reports $\hat{\boldsymbol{\theta}} = [\hat{\rho}, \hat{\pi}]$ of each risk bucket for the rating grades Baa to Caa–C for three different ASRF model specifications. The first row for each rating category shows the results for the Gaussian copula model as the limiting case for both the Student-t and NIG models. The second row shows the results for the Student-t model with fixed degrees of freedom $\nu = 15$. The third row shows the results for the NIG model with fixed shape parameter $\alpha = 3$. Based on the parameter estimates, the VaR at the estimates and the VaR at the hedge level $\boldsymbol{\theta}_h = \hat{\boldsymbol{\theta}} + 0.75\,\hat{\sigma}(\hat{\boldsymbol{\theta}})$ are reported for the confidence level $\alpha = 99.9\%$. The last two columns show the contractual fees f_c^* in basis points. The sample period is from 1985 to 2014.

Rating	Model	$\hat{\rho}$ [%]	$\hat{\pi}$ [%]	VaR($\hat{\boldsymbol{\theta}}$) [%]	VaR($\boldsymbol{\theta}_h$) [%]	f_0^* [bp]	f_1^* [bp]
Baa	$\nu = \alpha = \infty$	14.7792**	0.1952***	3.4351	5.4707	0.2448	0.4432
	$\nu = 15$	0.0001	0.2588***	5.9796	7.3791	0.0698	0.2016
	$\alpha = 3$	16.3660***	0.2015**	4.5802	7.8880	0.3286	0.6575
Ba	$\nu = \alpha = \infty$	13.9788***	1.1101***	11.5464	15.4639	0.3932	0.7657
	$\nu = 15$	0.0100	1.1795***	11.7526	17.3196	0.8886	1.4139
	$\alpha = 3$	15.1828***	1.1403***	14.4330	20.2062	0.5055	1.0715
B	$\nu = \alpha = \infty$	23.4588***	4.9322***	43.2320	51.9337	0.7803	1.6501
	$\nu = 15$	15.0279**	4.8368***	40.6077	49.4475	0.8637	1.7323
	$\alpha = 3$	25.0968***	5.1423***	52.3481	63.1215	0.8435	1.9123
Caa–C	$\nu = \alpha = \infty$	13.8124***	17.9317***	60.9649	67.9825	0.8052	1.4774
	$\nu = 15$	11.3962***	17.9661***	60.5263	67.5439	0.8323	1.5203
	$\alpha = 3$	14.2876***	17.7821***	65.3509	73.2456	0.7836	1.5424

*, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively

Bibliography

Acerbi, Carlo and Dirk Tasche (2002), 'On the coherence of expected shortfall', *Journal of Banking & Finance* **26**(7), 1487–1503.

Acharya, Viral V., Stephen Schaefer and Yili Zhang (2015), 'Liquidity risk and correlation risk: A clinical study of the General Motors and Ford downgrade of May 2005', *Quarterly Journal of Finance* **5**(2), 1550006.

Acharya, Viral V., Thomas Cooley, Matthew Richardson and Ingo Walter (2009), 'Manufacturing tail risk: A perspective on the financial crisis of 2007–2009', *Foundations and Trends® in Finance* **4**(4), 247–325.

Aït-Sahalia, Yacine and Andrew W. Lo (1998), 'Nonparametric estimation of state-price densities implicit in financial asset prices', *The Journal of Finance* **53**(2), 499–547.

Aït-Sahalia, Yacine and Jefferson Duarte (2003), 'Nonparametric option pricing under shape restrictions', *Journal of Econometrics* **116**(1–2), 9–47.

Albanese, Claudio, David Li, Edgar Lobachevskiy and Gunter Meissner (2013), 'A comparative analysis of correlation approaches in finance', *The Journal of Derivatives* **21**(2), 42–66.

Alexander, Carol and José María Sarabia (2012), 'Quantile uncertainty and value-at-risk model risk', *Risk Analysis* **32**(8), 1293–1308.

Altman, Edward I., Brooks Brady, Andrea Resti and Andrea Sironi (2005), 'The link between default and recovery rates: Theory, empirical evidence, and implications', *The Journal of Business* **78**(6), 2203–2228.

Andersen, Leif and Jakob Sidenius (2005), 'Extensions to the Gaussian copula: random recovery and random factor loadings', *Journal of Credit Risk* **1**(1), 29–70.

Andersen, Torben G. and Oleg Bondarenko (2007), Construction and interpretation of model-free implied volatility, NBER Working Paper No. 13449, National Bureau of Economic Research.

Augustin, Patrick, Marti G. Subrahmanyam, Dragon Yongjun Tang and Sarah Qian Wang (2014), 'Credit default swaps: A survey', *Foundations and Trends® in Finance* **9**(1–2), 1–196.

Azizpour, Shahriar, Kay Giesecke and Baeho Kim (2011), 'Premia for correlated default risk', *Journal of Economic Dynamics & Control* **35**(8), 1340–1357.

Azizpour, Shahriar, Kay Giesecke and Gustavo Schwenkler (2018), 'Exploring the sources of default clustering', *Journal of Financial Economics*, forthcoming.

Bakshi, Gurdip and Dilip Madan (2000), 'Spanning and derivative-security valuation', *Journal of Financial Economics* **55**(2), 205–238.

Bakshi, Gurdip and Dilip Madan (2006), 'A theory of volatility spreads', *Management Science* **52**(12), 1945–1956.

Bakshi, Gurdip, Nikunj Kapadia and Dilip Madan (2003), 'Stock return characteristics, skew laws, and the differential pricing of individual equity options', *The Review of Financial Studies* **16**(1), 101–143.

Bali, Turan G., Jianfeng Hu and Scott Murray (2017), Option implied volatility, skewness, and kurtosis and the cross-section of expected stock returns, Working Paper, Georgetown University, Singapore Management University and Georgia State University.

Bao, Yong and Aman Ullah (2004), 'Bias of a value-at-risk estimator', *Finance Research Letters* **1**(4), 241–249.

Barberis, Nicholas (2000), 'Investing for the long run when returns are predictable', *The Journal of Finance* **55**(1), 225–264.

Barry, Christopher B. (1974), 'Portfolio analysis under uncertain means, variances, and covariances', *The Journal of Finance* **29**(2), 515–522.

Barry, Christopher B. and Stephen J. Brown (1985), 'Differential information and security market equilibrium', *Journal of Financial and Quantitative Analysis* **20**(4), 407–422.

Baule, Rainer, Olaf Korn and Sven Saßning (2016), 'Which beta is best? On the information content of option-implied betas', *European Financial Management* **22**(3), 450–483.

BCBS (2006), Basel II: International convergence of capital measurement and capital standards: A revised framework – comprehensive version, Basel Committee on Banking Supervision, Bank for International Settlements.

BCBS (2013), Foundations of the Proposed Modified Supervisory Formula Approach, BCBS Working Papers 22, Basel Committee on Banking Supervision, Bank for International Settlements.

Benzoni, Luca, Pierre Collin-Dufresne, Robert S. Goldstein and Jean Helwege (2015), 'Modeling credit contagion via the updating of fragile beliefs', *The Review of Financial Studies* **28**(7), 1960–2008.

Berberan-Santos, Mário N. (2007), 'Expressing a probability density function in terms of another PDF: A generalized Gram–Charlier expansion', *Journal of Mathematical Chemistry* **42**(3), 585–594.

Berkowitz, Jeremy and James O'Brien (2002), 'How accurate are value-at-risk models at commercial banks?', *The Journal of Finance* **57**(3), 1093–1111.

Berkowitz, Jeremy, Peter Christoffersen and Denis Pelletier (2009), 'Evaluating value-at-risk models with desk-level data', *Management Science* **57**(12), 2213–2227.

Bernard, Carole, Ludger Rüschendorf, Steven Vanduffel and Jing Yao (2017), 'How robust is the value-at-risk of credit risk portfolios?', *The European Journal of Finance* **23**(6), 507–534.

Bhansali, Vineer, Robert Gingrich and Francis A. Longstaff (2008), 'Systemic credit risk: What is the market telling us?', *Financial Analysts Journal* **64**(4), 16–24.

Bignozzi, Valeria and Andreas Tsanakas (2016), 'Parameter uncertainty and residual estimation risk', *Journal of Risk and Insurance* **83**(4), 949–978.

Bion-Nadal, Jocelyne and Magali Kervarec (2012), 'Risk measuring under model uncertainty', *The Annals of Applied Probability* **22**(1), 213–238.

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

Bollerslev, Tim, Michael Gibson and Hao Zhou (2011), 'Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities', *Journal of Econometrics* **160**(1), 235–245.

Bondarenko, Oleg (2014), 'Variance trading and market price of variance risk', *Journal of Econometrics* **180**(1), 81–97.

Boucher, Christophe M., Jón Daníelsson, Patrick S. Kouontchou and Bertrand B. Maillet (2014), 'Risk models-at-risk', *Journal of Banking & Finance* **44**, 72–92.

Boyarchenko, Nina, Pooja Gupta, Nick Steele and Jacqueline Yen (2016), Trends in credit market arbitrage, Staff Reports No. 784, Federal Reserve Bank of New York.

Breeden, Douglas T. and Robert H. Litzenberger (1978), 'Prices of state-contingent claims implicit in option prices', *The Journal of Business* **51**(4), 621–651.

Brigo, Damiano, Andrea Pallavicini and Roberto Torresetti (2010), *Credit Models and the Crisis: A Journey into CDOs, Copulas, Correlations and Dynamic Models*, John Wiley & Sons, Chichester, West Sussex.

Britten-Jones, Mark and Anthony Neuberger (2000), 'Option prices, implied price processes, and stochastic volatility', *The Journal of Finance* **55**(2), 839–866.

Broer, Tobias (2018), 'Securitization bubbles: Structured finance with disagreement about default risk', *Journal of Financial Economics* **127**(3), 505–518.

Bunnin, F. O., Y. Guo and Y. Ren (2002), 'Option pricing under model and parameter uncertainty using predictive densities', *Statistics and Computing* **12**(1), 37–44.

Buraschi, Andrea, Robert Kosowski and Fabio Trojani (2014), 'When there is no place to hide: Correlation risk and the cross-section of hedge fund returns', *The Review of Financial Studies* **27**(2), 581–616.

Buss, Adrian and Grigory Vilkov (2012), 'Measuring equity risk with option-implied correlations', *The Review of Financial Studies* **25**(10), 3113–3140.

Butland, Judy (1980), A method of interpolating reasonable-shaped curves through any data, *in* 'Proceedings of Computer Graphics 80', Online Publications Ltd., Northwood Hills, Middlesex, UK, 409–422.

Cai, Yongyang and Kenneth L. Judd (2012), 'Dynamic programming with shape-preserving rational spline Hermite interpolation', *Economics Letters* **117**(1), 161–164.

Carr, Peter and Dilip Madan (1998), Towards a theory of volatility trading, *in* R. A.Jarrow, ed., 'Volatility: New Estimation Techniques for Pricing Derivatives', Risk Books, London, 417–427.

Carr, Peter and Liuren Wu (2009), 'Variance risk premiums', *The Review of Financial Studies* **22**(3), 1311 –1341.

Chang, Bo Young, Peter Christoffersen and Kris Jacobs (2013), 'Market skewness risk and the cross section of stock returns', *Journal of Financial Economics* **107**(1), 46–68.

Chang, Bo-Young, Peter Christoffersen, Kris Jacobs and Gregory Vainberg (2012), 'Option-implied measures of equity risk', *Review of Finance* **16**(2), 385–428.

Chava, Sudheer, Catalina Stefanescu and Stuart Turnbull (2011), 'Modeling the loss distribution', *Management Science* **57**(7), 1267–1287.

Chen, Shu-Hsiu (2017), 'Carry trade strategies based on option-implied information: Evidence from a cross-section of funding currencies', *Journal of International Money and Finance* **78**(Supplement C), 1–20.

Choi, Hoyong, Philippe Mueller and Andrea Vedolin (2017), 'Bond variance risk premiums', *Review of Finance* **21**(3), 987–1022.

Christoffersen, Peter and Kris Jacobs (2004), 'The importance of the loss function in option valuation', *Journal of Financial Economics* **72**(2), 291–318.

Christoffersen, Peter and Sílvia Gonçalves (2005), 'Estimation risk in financial risk management', *Journal of Risk* **7**(3), 1–28.

Coles, Jeffrey L. and Uri Loewenstein (1988), 'Equilibrium pricing and portfolio composition in the presence of uncertain parameters', *Journal of Financial Economics* **22**(2), 279–303.

Collin-Dufresne, Pierre (2009), 'A short introduction to correlation markets', *Journal of Financial Econometrics* **7**(1), 12–29.

Collin-Dufresne, Pierre, Robert S. Goldstein and Fan Yang (2012), 'On the relative pricing of long-maturity index options and collateralized debt obligations', *The Journal of Finance* **67**(6), 1983–2014.

Collin-Dufresne, Pierre, Robert S. Goldstein and J. Spencer Martin (2001), 'The determinants of credit spread changes', *The Journal of Finance* **56**(6), 2177–2207.

Conrad, Jennifer, Robert F. Dittmar and Eric Ghysels (2013), 'Ex ante skewness and expected stock returns', *The Journal of Finance* **68**(1), 85–124.

Cont, Rama (2001), 'Empirical properties of asset returns: stylized facts and statistical issues', *Quantitative Finance* **1**, 223–236.

Coval, Joshua D., Jakub W. Jurek and Erik Stafford (2009), 'Economic catastrophe bonds', *American Economic Review* **99**(3), 628–666.

Coval, Joshua D., Kevin Pan and Erik Stafford (2014), Capital market blind spots, Working Paper, Harvard University.

Crouhy, Michel, Dan Galai and Robert Mark (2000), 'A comparative analysis of current credit risk models', *Journal of Banking & Finance* **24**(1–2), 59–117.

Cuoco, Domenico and Hong Liu (2006), 'An analysis of VaR-based capital requirements', *Journal of Financial Intermediation* **15**(3), 362–394.

Daníelsson, Jón, Bjørn N. Jorgensen and Casper G. de Vries (2002), 'Incentives for effective risk management', *Journal of Banking & Finance* **26**(7), 1407–1425.

Das, Sanjiv R., Darrell Duffie, Nikunj Kapadia and Leandro Saita (2007), 'Common failings: How corporate defaults are correlated', *The Journal of Finance* **62**(1), 93–117.

Della Corte, Pasquale, Tarun Ramadorai and Lucio Sarno (2016), 'Volatility risk premia and exchange rate predictability', *Journal of Financial Economics* **120**(1), 21–40.

DeMiguel, Victor, Yuliya Plyakha, Raman Uppal and Grigory Vilkov (2013), 'Improving portfolio selection using option-implied volatility and skewness', *Journal of Financial and Quantitative Analysis* **48**(6), 1813–1845.

Doctor, Saul and Harpreet Singh (2010), CDS v2.0: The new architecture of the CDS market, Europe Credit Research, J.P. Morgan Securities Ltd., London.

Driessen, Joost, Pascal J. Maenhout and Grigory Vilkov (2009), 'The price of correlation risk: Evidence from equity options', *The Journal of Finance* **64**(3), 1377–1406.

Driessen, Joost, Pascal J. Maenhout and Grigory Vilkov (2013), Option-implied correlations and the price of correlation risk, Working Paper, Tilburg University, INSEAD and Goethe University Frankfurt.

Du, Zaichao and Juan Carlos Escanciano (2017), 'Backtesting expected shortfall: Accounting for tail risk', *Management Science* **63**(4), 940–958.

Duellmann, Klaus, Jonathan Küll and Michael Kunisch (2010), 'Estimating asset correlations from stock prices or default rates—Which method is superior?', *Journal of Economic Dynamics & Control* **34**(11), 2341–2357.

Duffie, Darrell, Andreas Eckner, Guillaume Horel and Leandro Saita (2009), 'Frailty correlated default', *The Journal of Finance* **64**(5), 2089–2123.

Duffie, Darrell and Nicolae Gârleanu (2001), 'Risk and valuation of collateralized debt obligations', *Financial Analysts Journal* **57**(1), 41–59.

Eckner, Andreas (2010), Risk premia in structured credit derivatives, Working Paper, Stanford University.

Ehlers, Philippe and Philipp J. Schönbucher (2009), 'Background filtrations and canonical loss processes for top-down models of portfolio credit risk', *Finance and Stochastics* **13**(1), 79–103.

Embrechts, Paul, Bin Wang and Ruodu Wang (2015), 'Aggregation-robustness and model uncertainty of regulatory risk measures', *Finance and Stochastics* **19**(4), 763–790.

Embrechts, Paul, Giovanni Puccetti and Ludger Rüschendorf (2013), 'Model uncertainty and VaR aggregation', *Journal of Banking & Finance* **37**(8), 2750–2764.

Erel, Isil, Stewart C. Myers and James A. Read (2015), 'A theory of risk capital', *Journal of Financial Economics* **118**(3), 620–635.

Escanciano, J. Carlos and Jose Olmo (2011), 'Robust backtesting tests for value-at-risk models', *Journal of Financial Econometrics* **9**(1), 132–161.

FCA (2016), Prudential sourcebook for Banks, Building Societies and Investment Firms, Release 4, Financial Conduct Authority, United Kingdom.

FDIC (2011), 'Federal Deposit Insurance Corporation (FDIC) Part II, 12 CFR Part 327, Assessments, Large Bank Pricing', *Federal Register* **76**(38), 10672–10733.

Feldhütter, Peter, Linda Sandris Larsen, Claus Munk and Anders B. Trolle (2012), Keep it simple: Dynamic bond portfolios under parameter uncertainty, Working Paper, London Business School, University of Southern Denmark, Copenhagen Business School, Swiss Finance Institute and École Polytechnique Fédérale de Lausanne.

Figlewski, Stephen (2004), Estimation error in the assessment of financial risk exposure, Working Paper, NYU Stern School of Business, New York.

Frey, Rüdiger and Alexander J. McNeil (2002), 'VaR and expected shortfall in portfolios of dependent credit risks: Conceptual and practical insights', *Journal of Banking & Finance* **26**(7), 1317–1334.

Frey, Rüdiger and Alexander J. McNeil (2003), 'Dependent defaults in models of portfolio credit risk', *Journal of Risk* **6**(1), 59–92.

Fritsch, Fred N. and Judy Butland (1984), 'A method for constructing local monotone piecewise cubic interpolants', *SIAM Journal on Scientific and Statistical Computing* **5**(2), 300–304.

Fröhlich, Andreas and Annegret Weng (2015), 'Modelling parameter uncertainty for risk capital calculation', *European Actuarial Journal* **5**(1), 79–112.

Froot, Kenneth A., David S. Scharfstein and Jeremy C. Stein (1993), 'Risk management: Coordinating corporate investment and financing policies', *The Journal of Finance* **48**(5), 1629–1658.

Froot, Kenneth A. and Jeremy C. Stein (1998), 'Risk management, capital budgeting, and capital structure policy for financial institutions: An integrated approach', *Journal of Financial Economics* **47**(1), 55–82.

Garlappi, Lorenzo, Raman Uppal and Tan Wang (2007), 'Portfolio selection with parameter and model uncertainty: A multi-prior approach', *The Review of Financial Studies* **20**(1), 41–81.

Giesecke, Kay (2009), Portfolio credit risk: Top-down versus bottom-up approaches, *in* R.Cont, ed., 'Frontiers in Quantitative Finance: Volatility and Credit Risk Modeling', John Wiley & Sons, Hoboken, NJ, 251–267.

Giesecke, Kay and Baeho Kim (2011), 'Risk analysis of collateralized debt obligations', *Operations Research* **59**(1), 32–49.

Giesecke, Kay, Konstantinos Spiliopoulos, Richard B. Sowers and Justin A. Sirignano (2015), 'Large portfolio asymptotics for loss from default', *Mathematical Finance* **25**(1), 77–114.

Gordy, Michael B. (2000), 'A comparative anatomy of credit risk models', *Journal of Banking & Finance* **24**(1–2), 119–149.

Gordy, Michael B. (2003), 'A risk-factor model foundation for ratings-based bank capital rules', *Journal of Financial Intermediation* **12**(3), 199–232.

Gordy, Michael B. and Erik Heitfield (2010), Small-Sample Estimation of Models of Portfolio Credit Risk, *in* M.Kijima, C.Hara, K.Tanaka and Y.Muromachi, eds, 'Recent Advances in Financial Engineering 2009', World Scientific, Singapore, 43–63.

Gouriéroux, Christian and Jean-Michel Zakoïan (2013), 'Estimation-adjusted VaR', *Econometric Theory* **29**(4), 735–770.

Hamerle, Alfred and Daniel Rösch (2005), 'Misspecified copulas in credit risk models: How good is Gaussian?', *Journal of Risk* **8**(1), 41–58.

Haven, Emmanuel, Xiaoquan Liu, Chenghu Ma and Liya Shen (2009), 'Revealing the implied risk-neutral MGF from options: The wavelet method', *Journal of Economic Dynamics & Control* **33**(3), 692–709.

Hellmann, Thomas F., Kevin C. Murdock and Joseph E. Stiglitz (2000), 'Liberalization, moral hazard in banking, and prudential regulation: Are capital requirements enough?', *American Economic Review* **90**(1), 147–165.

Hull, John C. and Alan D. White (2004), 'Valuation of a CDO and an *n*-th to default CDS without Monte Carlo simulation', *The Journal of Derivatives* **12**(2), 8–23.

Jiang, George J. and Yisong S. Tian (2005), 'The model-free implied volatility and its information content', *The Review of Financial Studies* **18**(4), 1305–1342.

Jiang, George J. and Yisong S. Tian (2007), 'Extracting model-free volatility from option prices', *The Journal of Derivatives* **14**(3), 35–60.

Jones, M. C., J. F. Rosco and Arthur Pewsey (2011), 'Skewness-invariant measures of kurtosis', *The American Statistician* **65**(2), 89–95.

Jorion, Philippe (1996), 'Risk^2: Measuring the risk in value at risk', *Financial Analysts Journal* **52**(6), 47–56.

Jorion, Philippe (2000), 'Risk management lessons from Long-Term Capital Management', *European Financial Management* **6**(3), 277–300.

Junge, Benjamin and Anders B. Trolle (2015), Liquidity risk in credit default swap markets, Working Paper, École Polytechnique Fédérale de Lausanne and Swiss Finance Institute.

Jurek, Jakub W. (2014), 'Crash-neutral currency carry trades', *Journal of Financial Economics* **113**(3), 325–347.

Kalemanova, Anna, Bernd Schmid and Ralf Werner (2007), 'The normal inverse Gaussian distribution for synthetic CDO pricing', *The Journal of Derivatives* **14**(3), 80–94.

Kalymon, Basil A. (1971), 'Estimation risk in the portfolio selection model', *Journal of Financial and Quantitative Analysis* **6**(1), 559–582.

Kan, Raymond and Guofu Zhou (2007), 'Optimal portfolio choice with parameter uncertainty', *Journal of Financial and Quantitative Analysis* **42**(03), 621–656.

Kempf, Alexander, Olaf Korn and Sven Saßning (2015), 'Portfolio optimization using forward-looking information', *Review of Finance* **19**(1), 467–490.

Keppo, Jussi, Leonard Kofman and Xu Meng (2010), 'Unintended consequences of the market risk requirement in banking regulation', *Journal of Economic Dynamics & Control* **34**(10), 2192–2214.

Kerkhof, Jeroen, Bertrand Melenberg and Hans Schumacher (2010), 'Model risk and capital reserves', *Journal of Banking & Finance* **34**(1), 267–279.

Klein, Roger W., Lawrence C. Rafsky, David S. Sibley and Robert D. Willig (1978), 'Decisions with estimation uncertainty', *Econometrica* **46**(6), 1363–1387.

Klein, Roger W. and Vijay S. Bawa (1976), 'The effect of estimation risk on optimal portfolio choice', *Journal of Financial Economics* **3**(3), 215–231.

Knight, Frank Hyneman (1921), *Risk, Uncertainty, and Profit*, number 16 *in* 'Hart, Schaffner and Marx Prize Essays, no. 31', Houghton Mifflin, Boston.

Korteweg, Arthur and Nicholas Polson (2010), Corporate credit spreads under parameter uncertainty, Working Paper, Stanford University and University of Chicago.

Kozhan, Roman, Anthony Neuberger and Paul Schneider (2013), 'The skew risk premium in the equity index market', *The Review of Financial Studies* **26**(9), 2174–2203.

Kupiec, Paul (2001), 'Estimating credit risk capital: What's the use?', *The Journal of Risk Finance* **2**(3), 17–34.

Löffler, Gunter (2003), 'The effects of estimation error on measures of portfolio credit risk', *Journal of Banking & Finance* **27**(8), 1427–1453.

Longstaff, Francis A. and Arvind Rajan (2008), 'An empirical analysis of the pricing of collateralized debt obligations', *The Journal of Finance* **63**(2), 529–563.

Longstaff, Francis A. and Brett W. Myers (2014), 'How does the market value toxic assets?', *Journal of Financial and Quantitative Analysis* **49**(2), 297–319.

Lönnbark, Carl (2010), 'A corrected value-at-risk predictor', *Applied Economics Letters* **17**(12), 1193–1196.

Lönnbark, Carl (2013), 'On the role of the estimation error in prediction of expected shortfall', *Journal of Banking & Finance* **37**(3), 847–853.

Lucas, Douglas J. (1995), 'Default correlation and credit analysis', *The Journal of Fixed Income* **4**(4), 76–87.

MacKenzie, Donald (2006), 'Is economics performative? Option theory and the construction of derivatives markets', *Journal of the History of Economic Thought* **28**(1), 29–55.

Maenhout, Pascal J. (2004), 'Robust portfolio rules and asset pricing', *The Review of Financial Studies* **17**(4), 951–983.

Markit (2014), Markit credit indices: A primer, White Paper, Markit Group Limited.

Markit (2015), Markit CDX high yield & Markit CDX investment grade index rules, White Paper, Markit Group Limited.

Marshall, Chris and Michael Siegel (1997), 'Value at risk: Implementing a risk measurement standard', *The Journal of Derivatives* **4**(3), 91–111.

Mele, Antonio, Yoshiki Obayashi and Catherine Shalen (2015), 'Rate fears gauges and the dynamics of fixed income and equity volatilities', *Journal of Banking & Finance* **52**, 256–265.

Merton, Robert C. (1995), 'Financial innovation and the management and regulation of financial institutions', *Journal of Banking & Finance* **19**(3), 461–481.

Merton, Robert C. and André Perold (2005), 'Theory of risk capital in financial firms', *Journal of Applied Corporate Finance* **6**(3), 16–32.

Monnier, Jean-Baptiste (2013), 'Risk-neutral density recovery via spectral analysis', *SIAM Journal on Financial Mathematics* **4**(1), 650–667.

Monteiro, Ana Margarida, Reha H. Tütüncü and Luís N. Vicente (2008), 'Recovering risk-neutral probability density functions from options prices using cubic splines and ensuring nonnegativity', *European Journal of Operational Research* **187**(2), 525–542.

Mueller, Philippe, Andreas Stathopoulos and Andrea Vedolin (2017), 'International correlation risk', *Journal of Financial Economics* **126**(2), 270–299.

Neuberger, Anthony (1994), 'The log contract', *The Journal of Portfolio Management* **20**(2), 74–80.

Nickerson, Jordan and John M. Griffin (2017), 'Debt correlations in the wake of the financial crisis: What are appropriate default correlations for structured products?', *Journal of Financial Economics* **125**(3), 454–474.

Oakley, Steven James (1990), Orthogonal polynomials in the approximation of probability distributions, PhD Thesis, The University of Arizona, Tucson.

O'Kane, Dominic (2008), *Modelling Single-name and Multi-name Credit Derivatives*, John Wiley & Sons, Chichester, West Sussex.

O'Kane, Dominic (2011), 'Force-fitting CDS spreads to CDS index swaps', *The Journal of Derivatives* **18**(3), 61–74.

Ou, Sharon (2017), Annual default study: Corporate default and recovery rates, 1920–2016, Data Report No. 1059749, Moody's Investors Service.

Ou, Sharon, John Kennedy, Zhi Zeng and Albert Metz (2015), Annual default study: Corporate default and recovery rates, 1920–2014, Special Comment, Moody's, New York.

Partnoy, Frank and David A. Jr. Skeel (2007), 'The promise and perils of credit derivatives', *University of Cincinnati Law Review* **75**, 1019–1052.

Pástor, Ľuboš (2000), 'Portfolio selection and asset pricing models', *The Journal of Finance* **55**(1), 179–223.

Pluto, Katja and Dirk Tasche (2005), 'Thinking positively', *Risk* **18**(8), 72–78.

Prokopczuk, Marcel, Lazaros Symeonidis and Chardin Wese Simen (2017), 'Variance risk in commodity markets', *Journal of Banking & Finance* **81**, 136–149.

Provost, Serge B. (2005), 'Moment-based density approximants', *The Mathematica Journal* **9**(4), 727–756.

Provost, Serge B., Min Jiang and Hyung-Tae Ha (2009), 'Moment-based approximations of probability mass functions with applications involving order statistics', *Communications in Statistics – Theory and Methods* **38**(12), 1969–1981.

Pykhtin, Michael and Ashish Dev (2002), 'Credit risk in asset securitisations: an analytical model', *Risk* **15**(5), S16–S20.

Rajan, Raghuram G. (2006), 'Has finance made the world riskier?', *European Financial Management* **12**(4), 499–533.

Rao, C. Radhakrishna (1973), *Linear Statistical Inference and its Applications*, Wiley Series in Probability and Statistics, 2nd edn, John Wiley & Sons, New York.

Reinking, J. Todd (2002), 'Orthogonal series representation of probability density and distribution functions', *The Mathematica Journal* **8**(3), 473–483.

Rockafellar, R. Tyrrell and Stanislav P. Uryasev (2002), 'Conditional value-at-risk for general loss distributions', *Journal of Banking & Finance* **26**(7), 1443–1471.

Rompolis, Leonidas S. and Elias Tzavalis (2008), 'Recovering risk neutral densities from option prices: A new approach', *Journal of Financial and Quantitative Analysis* **43**(4), 1037–1053.

Rösch, Daniel and Harald Scheule (2014), 'Forecasting mortgage securitization risk under systematic risk and parameter uncertainty', *Journal of Risk and Insurance* **81**(3), 563–586.

Schlögl, Erik (2013), 'Option pricing where the underlying assets follow a Gram/Charlier density of arbitrary order', *Journal of Economic Dynamics & Control* **37**(3), 611–632.

Schumaker, Larry L. (1983), 'On shape preserving quadratic spline interpolation', *SIAM Journal on Numerical Analysis* **20**(4), 854–864.

Seo, Sang Byung and Jessica A. Wachter (2016), Do rare events explain CDX tranche spreads?, NBER Working Paper No. 22723, National Bureau of Economic Research.

Stout, Lynn A. (2011), 'Derivatives and the legal origin of the 2008 credit crisis', *Harvard Business Law Review* **1**, 1–38.

Stulz, René M. (2004), 'Should we fear derivatives?', *Journal of Economic Perspectives* **18**(3), 173–192.

Stulz, René M. (2008), 'Risk management failures: What are they and when do they happen?', *Journal of Applied Corporate Finance* **20**(4), 39–48.

Stulz, René M. (2010), 'Credit default swaps and the credit crisis', *Journal of Economic Perspectives* **24**(1), 73–92.

Tarashev, Nikola A. (2010), 'Measuring portfolio credit risk correctly: Why parameter uncertainty matters', *Journal of Banking & Finance* **34**(9), 2065–2076.

Tee, Chyng Wen and Christopher Ting (2017), 'Variance risk premiums of commodity ETFs', *The Journal of Futures Markets* **37**(5), 452–472.

Vacca, Luigi (2005), 'Unbiased risk-neutral loss distributions', *Risk* **18**(11), 97–101.

Vašíček, Oldřich Alfons (1987), Probability of loss on loan portfolio, KMV Corporation, San Francisco, California.

Vašíček, Oldřich Alfons (1991), Limiting loan loss probability distribution, KMV Corporation, San Francisco, California.

Vašíček, Oldřich Alfons (2002), 'Loan portfolio value', *Risk* **15**(12), 160–162.

Veronesi, Pietro (2000), 'How does information quality affect stock returns?', *The Journal of Finance* **55**(2), 807–837.

Veronesi, Pietro and Luigi Zingales (2010), 'Paulson's gift', *Journal of Financial Economics* **97**(3), 339–368.

Xia, Yihong (2001), 'Learning about predictability: The effects of parameter uncertainty on dynamic asset allocation', *The Journal of Finance* **56**(1), 205–246.

Yamai, Yasuhiro and Toshinao Yoshiba (2005), 'Value-at-risk versus expected shortfall: A practical perspective', *Journal of Banking & Finance* **29**(4), 997–1015.

Zhang, Benjamin Yibin, Hao Zhou and Haibin Zhu (2009), 'Explaining credit default swap spreads with the equity volatility and jump risks of individual firms', *The Review of Financial Studies* **22**(12), 5099–5131.