

Preface

A quiet revival of credit correlation modelling has been taking place steadily in the background for the last few years. No, you don't need to be alarmed yet: we are not going back to pricing the exotic credit derivative products of yesteryears. But a more profound and meaningful shift in derivatives markets has been unfolding in the aftermath of the credit crisis. The products are now much simpler, transparent and very commoditized, but the management of the derivatives business as a whole is much more complex and involves many more subtleties that were not accounted for before. The pricing now transcends the simplistic risk-free valuation assumptions and involves the juxtaposition of (counterparty) credit risk, funding, liquidity and regulatory capital effects on top of the risk-free value. We do not price a trade in isolation, but it is part of a large portfolio of instruments whose valuation depends on the additional pricing adjustments that make up the all-in-price of the trade.

Nowhere is this more true than in credit derivatives' portfolios. The key risks that one has to deal with are the default correlation risk and the likelihood of credit spread jumps when one or multiple counterparties default. Even for a simple portfolio of vanilla, credit default swaps and credit indices computing the credit valuation adjustment when the counterparty defaults necessitate the modelling of default correlation effects between the single-name default events in this problem. The default correlation assumptions translate directly into wrong-way risk and gap risk effects that one needs to account for and manage accordingly. Similarly, the modelling of systemic risk losses that one would incur when trading with a CCP is similar to modelling a waterfall structure of a cash CDO where the joint defaults of the various clearing members, on the one hand, and the loss cushion provided by

the guarantee fund, on the other hand, define the loss variable and the subordination level of this equivalent senior CDO tranche. This is yet another example where default correlation needs to be modelled properly.

With these applications in mind, the objective of this book is primarily to provide the mathematical tools and numerical implementation techniques that are needed to tackle rigourously the credit correlation challenges that we face in the post-crisis XVA modelling era.

Objectives, Audience and Structure

This book is a combination of a lot of the work that I did over the last 15 years on credit correlation modelling and a review of some of the most important mathematical tools and various contributions in the field.

I gave it the title: “Credit Correlation: From FTDs to XVAs—Theory and Practice”. As we are now dealing with more XVA modelling issues (such as CVA, FVA, KVA, CCP and SIMM), many of the mathematical tools and techniques that we have to use time and time again stem from the general “Credit Correlation” modelling toolbox. With that in mind, the book is aimed at both students and researchers interested in credit modelling in general, but also XVA modellers who need solid theoretical foundations in credit correlation (and practical numerical implementation methods) that they would then use to address the various XVA (and credit) challenges that they face. By XVA here, I mean not only the micro-trading desk (standard) derivatives valuation adjustments (CVA, DVA and FVA) but also, more generally, macro-modelling issues when dealing with CCPs, systemic risk and regulatory capital.

Who is the intended audience for this book?

This is a book by a practitioner for practitioners, but it is also aimed at Ph.D. and MSc students in financial engineering, and researchers in mathematical finance.

It is based on solid mathematical foundations needed for credit modelling problems, especially point processes and marked point processes (MPPs), filtration enlargements, and intensities with respect to a given filtration. But its emphasis is geared towards applications to real-life problems, covering numerical implementation issues, calibration approaches and pricing and risk-related topics. The material is at an advanced level and assumes that the reader is familiar with the basics of credit modelling covered, for example, in the books by Schönbucher (2003) and O’Kane (2008).

This is not a collection of papers on credit correlation or a general survey of previously published material, but a coherent book on the topic that builds the theory from the ground up and uses it to solve very specific valuation problems. One of the aims of the book is to provide the reader with the theoretical and practical tools to tackle the challenges encountered when modelling CVA for credit portfolios, including general and specific wrong-way risks and gap risk inherently present in the default-enlarged filtrations used in credit markets.

How does it differ from other books in the field?

This book can be viewed as the next stage of the credit correlation evolution—in the post-crisis era. This would be one level of complexity after the standard textbooks by Schönbucher and O’Kane. The logical sequence, both in terms of coverage and advanced topics, would be: Schönbucher (2003) → O’Kane (2008) → Elouerkhaoui (2017).

What will you learn from this book?

The book by Dominic O’Kane is an excellent reference on standard single-name modelling, Gaussian copula and base correlation for pricing CDOs. The present book builds on that and expands the focus (for the post-crisis era) to cover the (credit) mathematical tools needed for XVA modelling.

Some of the key concepts covered include:

- Filtration enlargements, intensity processes and the generalized Dellacherie formula;
- Poisson shock models à la Marshall-Olkin—both the bottom-up version and the top-down version;
- A special focus on numerical implementation techniques;
- Advanced topics, including: Stochastic Recovery Modelling, Hedging in Incomplete Credit Markets (CR01 vs JTD), and Credit CVA.

I have also included up-to-date material on correlation skew rescaling, expected tranche loss interpolation, relative entropy minimization, CDO-squared and Stein (zero-bias) approximation.

How is this book organized?

Broadly speaking, the book follows a standard organization structure that would mimic the syllabus for a (graduate) course on credit correlation. It is organized in four parts. Part I gives the mathematical tools needed for credit correlation modelling. Part II reviews credit correlation modelling approaches and numerical implementation issues. Part III addresses some advanced credit

correlation topics, including dynamic credit modelling, stochastic recovery and hedging. Part IV gives a prelude to the next credit modelling challenges and describes CVA modelling for credit portfolios.

I start in Chap. 1 by giving the motivation for credit correlation models, namely: (stochastic) spread correlation is not default (event) correlation! This is well known in credit markets, but many XVA modellers do not seem to appreciate the distinction between the two. I give a (chronological) timeline of credit correlation modelling, from the early days of credit with the works of Lando and Duffie-Singleton to dynamic portfolio credit modelling (SPA, GPL, Schonbucher). It is quite fascinating to reflect on how far we have come (over a decade) in our understanding of (subtle) mathematics involved in credit models—many of these subtleties are only limited to credit modelling problems: filtration enlargements, intensity processes, copulas, top-down approaches, etc... I finish the book with a chapter on the next challenge: “CVA modelling”. This is just to give a taste of what is coming up next. All the mathematical tools, credit models and concepts that have been exposed over the various chapters in the book culminate in a “deep” understanding of credit correlation, which is then leveraged to address the new XVA challenges of today and tomorrow.

Description of Contents by Chapter

We give below a detailed description of the contents by chapter.

Chapter 1: Introduction and Context. In this introduction, we start by presenting the main (portfolio) credit derivative contracts that we are interested in. To make precise what is meant by default correlation, in the context of credit portfolio modelling, we take a little detour and construct a toy model based on “correlated intensities”. Our goal is to use this toy example to motivate the need for proper credit correlation modelling by showing that, ultimately, it only generates some second-order effects, which are not directly related to (proper) joint default events dependence. To fully appreciate the breadth and depth of the topic, we also give a brief timeline of default correlation modelling over the last two decades, which highlights the various pieces in the overall credit correlation puzzle that we will put together over the next few chapters.

Part I: Theoretical Tools

Chapter 2: Mathematical Fundamentals. In this chapter, we present the essential mathematical tools needed in the modelling of portfolio credit

derivative products. This includes doubly stochastic Poisson processes, also known as Cox processes; point processes and their intensities, on some given filtration; and copula functions.

Chapter 3: Expectations in the Enlarged Filtration. In this chapter, we derive a formula of the conditional expectation with respect to the enlarged filtration. This is a generalization of the Dellacherie formula. We shall use this key result to compute the expectations that we encounter in the conditional jump diffusion framework. In particular, the conditional survival probability can be computed with our formula. We apply this result in Chap. 4 where conditional survival probability calculations, on the enlarged filtration, are carried out in details.

Chapter 4: Copulas and Conditional Jump Diffusions. Enlarging the economic state-variables' filtration by observing the default process of all available credits has some profound implications on the dynamics of intensities. Indeed, the sudden default of one credit triggers jumps in the spreads of all the other obligors. This is what we refer to as the “Conditional Jump Diffusion” effect. The aim of this chapter is to give a comprehensive and self-contained presentation of the CJD framework. In particular, we derive the default times' density function in the “looping” defaults model, and we study the equivalence between the copula approach and the conditional jumps framework. This is a key result that we will use, in practice, to calibrate non-observable default correlation parameters.

Part II: Correlation Models: Practical Implementation

Chapter 5: Correlation Demystified: A General Overview. This chapter gives a broad overview of default correlation modelling in the context of pricing and risk managing a correlation trading book. We cover both theoretical and practical market aspects, as well as numerical performance issues.

Chapter 6: Correlation Skew: A Black-Scholes Approach. In this chapter, we view the valuation of CDO tranches as an option pricing problem. The pay-off of a CDO tranche is a call-spread on the loss variable. By specifying the distribution of the loss variable at each time horizon, one would be able to value tranches. The standard way of defining this distribution is the base correlation approach. Here, we use a Black-Scholes analogy, and we define an implied volatility for each tranche. Then, given a Black volatility surface, we parameterize the loss distribution with a stochastic CEV model. We show that this parametric form gives a very good fit to the market tranche quotes. In addition, we give an application of the correlation skew Black approach to risk management and hedging.

Chapter 7: An Introduction to the Marshall-Olkin Copula. In this chapter, we study the “Marshall-Olkin” copula model in the context of credit risk modelling. This framework was traditionally used in reliability theory to model the failure rate in multi-component systems. The failure of each component is assumed to be contingent on some independent Poisson shocks. Our aim is to show that MO is a viable alternative to the Gaussian copula. This is done in three steps: (1) we introduce the MO model as the natural extension of a univariate Poisson process, (2) we discuss parametrization and calibration issues, and (3) we compare it with the standard Gaussian copula. We also show that the MO model can be used to reproduce the observed correlation skew in the CDO market. More recently, there has been renewed interest in the MO copula in the context of model risk management (see Morini 2011) and systemic risk modelling (see Gatarek and Jablecki 2016).

Chapter 8: Numerical Tools: Basket Expansions. In the next few chapters, we study some efficient numerical methods for the valuation of large basket credit derivatives. While the approaches are presented in the Marshall-Olkin copula model, most of the numerical techniques are generic and could be used with other copulas as well. The methods presented span a large spectrum of applied mathematics: Fourier transforms, changes of probability measure, numerical stable schemes, high-dimensional Sobol integration and recursive convolution algorithms.

Chapter 9: Static Replication. In principle, the direct (pricing) method requires, for each time step, 2^{n+1} values, corresponding to the set of all possible default combinations; as the size of the underlying basket increases, the number of default configurations explodes exponentially. This significant limitation restricts the applicability of the method to baskets under 10 or 11 credits. As an alternative, we develop a different approach, which is based on a static replication idea.

In this chapter, we describe how this static FTD replication is done: first, we show the relationship between k th-to-default and $(k - 1)$ th-to-default swaps; then, we apply this recursion step-by-step until we arrive at the complete FTD expansion.

Chapter 10: The Homogeneous Transformation. In general, the number of sub-FTDs in the replication formula is a function of n , the size of the basket, and k , the order of the basket default swap. The most time-consuming step in the evaluation is the generation of the sub-FTDs, for all possible combinations. If we had a homogeneous basket, then, for a given subset size l , all the FTD instruments would have exactly the same value, and the pricing

equation would simplify substantially. In particular, the number of sub-FTDs to compute would reduce to one evaluation per l -subset, hence a total of $N(k, n) = k$ FTD evaluations for the whole k th-to-default swap.

The first (natural) approximation that we consider is to transform the original non-homogeneous basket to a homogeneous one while preserving some properties of the aggregate default distribution. In the approach described here, for each default order, we use the corresponding percentile of the aggregate default distribution, and we require that this quantity remains invariant with respect to the homogeneous approximation. We shall see that this transformation is exact for an FTD swap, and that, for higher-order defaults, the approximation gives very good results.

Chapter 11: The Asymptotic Homogeneous Expansion. The transformation, described in the previous chapter, produces a homogeneous portfolio, which mimics some properties of the aggregate default distribution, and can be used for the purposes of basket default swap valuation. By using this homogeneous portfolio, the numerical burden that comes with the pricing of large baskets is eased, and the valuation algorithm is significantly faster. The k^{th} -to-default survival probability $Q_n^{[k]} = \mathbb{P}(\tau^{[k]} > T)$ for the n -dimensional homogeneous portfolio can be computed recursively as:

$$Q_n^{[k]} = \left(\frac{n}{k-1}\right) \left(Q_{n-1}^{[k-1]} - Q_n^{[k-1]}\right) + Q_n^{[k-1]}.$$

Unfortunately, this simple-looking recursion is numerically unstable. As one moves up the recursion tree, the numerical round-off errors propagate rapidly, and the resulting prices are completely erroneous. To address this issue, we take a different route: rather than using the recursive approach, we study the asymptotic behaviour of the homogeneous portfolio. We show, in this chapter, that the solution $Q_n^{[k]}$ admits an asymptotic series expansion, and we explain how to compute each term in the expansion.

Chapter 12: The Asymptotic Expansion. In this chapter, we relax the homogeneous portfolio assumption, and we derive an asymptotic series expansion of the k th-to-default Q-factor in the non-homogeneous case. We also show how to compute the conditional aggregate default distributions that appear in the expansion using the convolution recursion algorithm. The latter and other recursive methods have been traditionally used in actuarial mathematics to evaluate ruin probabilities and insurance premia.

Chapter 13: CDO-Squared: Correlation of Correlation. In this chapter, we analyse the “correlation of correlation” risk in the Marshall-Olkin copula

framework. The valuation of higher-order correlation products such as “CDOs of CDOs” (also known as “CDO-Squared”) is mainly driven by correlation of correlation effects. First, we extend the first-to-default replication method to baskets of basket products. Then, we develop an intuitive methodology for analysing this type of structures. The idea is to model each underlying basket security as a single-name process, and to derive its equivalent intensity process and its decomposition on the MO common market factors. This, in turn, defines the multivariate dependence between the underlying basket securities in the portfolio.

Chapter 14: Second Generation Models: From Flat to Correlation Skew.

In this chapter, we review some popular correlation skew models. We give a brief description of each model and discuss the advantages and limitations of each modelling framework. This includes the stochastic correlation model, local correlation (and random factor loading), the Levy copula and the implied (hazard rate) copula.

The stochastic and local correlation models are so-called second generation models, which extend the Gaussian copula in an attempt to model tranche prices of the entire capital structure at a fixed time horizon. The Levy copula is also an extension of the Gaussian copula, which accounts for the correlation skew, but it also provides some dynamics for the expected loss process. And last but not least, the implied hazard rate copula is a nonparametric copula function, which constructs the implied distribution of the conditioning factor (and the associated conditional single-name probabilities) from the tranche prices directly.

Chapter 15: Third Generation Models: From Static to Dynamic Models.

In this chapter, we review some of the most important dynamic credit models in the literature. We give a brief description of each model and discuss the advantages and limitations of each modelling framework. We also comment on the usefulness of each model for a given family of correlation products. The models discussed include: the Top-Down model (of Giesecke and Goldberg 2005), the Dynamic Generalized Poisson Loss model (of Brigo, Pallavicini, Torresetti 2007), the N+ Model (of Longstaff and Rajan 2008), the Markov Chain Portfolio Loss model (of Schönbucher 2005), and the SPA model (of Sidenius, Piterbarg, Andersen 2008).

Part III: Advanced Topics in Pricing and Risk Management

Chapter 16: Pricing Path-Dependent Credit Products. This chapter addresses the problem of pricing (soft) path-dependent portfolio credit derivatives whose pay-off depends on the loss variable at different time horizons. We review the general theory of copulas and Markov processes, and we

establish the link between the copula approach and the Markov-Functional paradigm used in interest rates modelling. Equipped with these theoretical foundations, we show how one can construct a dynamic credit model, which matches the correlation skew at each tenor, by construction, and follows an exogenously specified choice of dynamics. Finally, we discuss the details of the numerical implementation, and we give some pricing examples in this framework.

Chapter 17: Hedging in Incomplete Markets. In this chapter, we present a methodology for hedging basket credit derivatives with single-name instruments. Because of the market incompleteness due to the residual correlation risk, perfect replication cannot be achieved. We allow for mean self-financing strategies and use a risk-minimization criterion to find the hedge. Managing credit risk is always a fine balance between hedging the jump-to-default exposure (JTD) or the credit spread exposure (CR01). Recently, this credit hedging dilemma (JTD vs CR01) is becoming very topical in the context of managing counterparty credit risk for large derivatives books.

Chapter 18: Min-Variance Hedging with Carry. In this chapter, we present the construction of the Min-Variance Hedging Delta operator used for basket products. Because of the market incompleteness –i.e. we cannot replicate a basket product with its underlying default swaps– min-variance hedging is the best thing that we can hope for. There will always be a residual correlation risk orthogonal to the sub-space of hedging instruments. We also present an extension of the standard MVH optimization to take into account the drift mismatch between the basket and the hedging portfolio. This defines the “Min-Variance Hedging Deltas with Carry”.

Chapter 19: Correlation Calibration with Stochastic Recovery. In this chapter, we expand the base correlation framework by enriching it with Stochastic Recovery modelling as a way to address the model limitations observed in a distressed credit environment. We introduce the general class of conditional-functional recovery models, which specify the recovery rate as a function of the common conditioning factor of the Gaussian copula. Then, we review some of the most popular ones, such as: the Conditional Discrete model of Krekel (2008), the Conditional Gaussian of Andersen and Sidenius (2005) and the Conditional Mark-Down of Amraoui and Hitier (2008). We also look at stochastic recovery from an aggregate portfolio perspective and present a top-down specification of the problem. By establishing the equivalence between these two approaches, we show that the latter can provide a useful tool for analysing the structure of various stochastic recovery model assumptions.

Part IV: The Next Challenge

Chapter 20: New Frontiers in Credit Modelling: the CVA Challenge. In this chapter, we present a general framework for evaluating the (counterparty) credit valuation adjustment for CDO tranches. We shall see that given the “exotic” nature of the CVA derivative pay-off, we will have to leverage a variety of modelling techniques that have been developed over years for the correlation book; this includes default correlation modelling, credit index options’ pricing, dynamic credit modelling and CDO-squared pricing.

References

- L. Andersen, J. Sidenius, Extensions to the Gaussian Copula: Random recovery and random factor loadings. *J. Credit Risk*. **1**(1, Winter 2004/2005), 29–70 (2005)
- S. Amraoui, S. Hitier, Optimal stochastic recovery for base correlation (Working Paper, BNP Paribas, 2008)
- D. Brigo, A. Pallavicini, R. Torresetti, Calibration of CDO tranches with the dynamical generalized-Poisson Loss model. *Risk* (2007)
- Y. Elouerkhaoui, *Credit Correlation: Theory and Practice* (Palgrave, London, 2017)
- D. Gatarek, J. Jablecki, Modeling joint defaults in correlation-sensitive instruments. *J. Credit Risk* **12**(3) (2016)
- K. Giesecke, L. Goldberg, A top down approach to multi-name credit (Working Paper, 2005)
- M. Krekel, Pricing distressed CDOs with base correlation and stochastic recovery (Working Paper, Unicredit, 2008)
- F. Longstaff, A. Rajan, An empirical analysis of the pricing of collateralized debt obligations. *J. Finance*. **63**(2), 529–563 (2008)
- M. Morini, *Understanding and Managing Model Risk: A Practical Guide for Quants, Traders and Validators* (Wiley, Chichester, 2011)
- D. O’Kane, *Modelling Single-Name and Multi-Name Credit Derivatives* (Wiley, Chichester, 2008)
- P.J. Schönbucher, *Credit Derivatives Pricing Models: Models, Pricing and Implementation* (Wiley, Chichester, 2003)
- P.J. Schönbucher, Portfolio losses and the term structure of loss transition rates: A new methodology for the pricing of portfolio credit derivatives (Working Paper, ETH Zurich, 2005)
- J. Sidenius, V. Piterbarg, L. Andersen, A new framework for dynamic credit portfolio loss modeling. *Int. J. Theor. Appl. Financ.* **11**(2), 163–197 (2008)

Credit Correlation

Theory and Practice

Elouerkhaoui, Y.

2017, XXIV, 456 p. 67 illus., Hardcover

ISBN: 978-3-319-60972-0