SYNTHETIC CDOs

Modelling, Valuation and Risk Management

Credit derivatives have enjoyed explosive growth in the last decade. One of the most important assets in this industry is synthetic Collateralised Debt Obligations (synthetic CDOs). This book describes the state-of-the-art in quantitative and computational modelling of these instruments.

Starting with a brief overview of the structured finance landscape, the book introduces the basic modelling concepts necessary to model and value simple vanilla credit derivatives. Building on this the book then describes in detail the modelling, valuation and risk management of synthetic CDOs. A clear and detailed picture of the behaviour of these complex instruments is built up. The final chapters introduce more advanced topics such as portfolio management of synthetic CDOs and hedging techniques, often not covered in other texts.

Mathematics, Finance and Risk

Editorial Board

SYNTHETIC CDOs

Modelling, Valuation and Risk Management

CRAIG MOUNFIELD

CAMBRIDGE
UNIVERSITY PRESS

Shaftesbury Road, Cambridge CB2 8EA, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi – 110025, India

103 Penang Road, #05–06/07, Visioncrest Commercial, Singapore 238467

Cambridge University Press is part of Cambridge University Press & Assessment,
a department of the University of Cambridge.

We share the University's mission to contribute to society through the pursuit of
education, learning and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9780521897884

First published 2009

A catalogue record for this publication is available from the British Library

Library of Congress Cataloging-in-Publication data
Mounfield, Craig, 1969–
Synthetic CDOs : modelling, valuation and risk management / Craig Mounfield.
p. cm. – (Mathematics, finance and risk)
Includes bibliographical references and index.
ISBN 978-0-521-89788-4 (hbk.)
1. Collateralized debt obligations. I. Title. II. Series.
HG6024.A3M69 2009
332.63´2 – dc22 2008043035

ISBN 978-0-521-89788-4 Hardback

Dedicated to my parents, my wife and my daughter.

Contents

Preface

This is a book about the modelling, valuation and risk management of synthetic collateralised debt obligations (or synthetic CDOs or simply CDOs for short). Synthetic CDOs are an example of a structured credit product. This is a financial product that takes targeted risk for the purpose of achieving targeted returns. Structured credit products utilise two financial engineering technologies: credit derivatives and asset securitisation. Synthetic CDOs have played an increasingly important role in the expansion of the global credit derivatives market which has grown rapidly since the turn of the century. Indeed, it is estimated that by the end of 2006 the total credit derivative notional amount outstanding was over $20 trillion (from virtually zero only a decade earlier). Increased trading volumes naturally led to market participants becoming more sophisticated (in terms of their risk/return characteristics and the strategies they employ) as well as to a commensurate increase in the complexity and subtlety of the products available. This in turn drives the evolution of the mathematical and computational models used to value these products. The objective of this book is to collate, summarise and critically assess the current state-of-the-art in quantitative and computational modelling of synthetic CDOs. The key word here is *modelling*; the book is about mathematical models and their properties. This book is not intended to provide detailed descriptions of the business and economic rationales for trading credit derivatives; there are better resources available that describe this and due reference will be given to these sources. It is meant to provide a detailed quantitative description of the modelling techniques currently employed in the marketplace for characterising synthetic CDOs.

It will be assumed that the technical level and experience of the reader is relatively high. Basic financial concepts will not be described in detail (except insofar as when such detail is necessary). Instead reference will be made to the appropriate resources. The use of financial and technical jargon will hopefully be kept to a minimum, although in a specialised, technical text such as this some jargon is inevitable. The rationale for this approach is to ensure the volume is concise and to the point. It is

intended to describe just enough of the mathematical and computational modelling to enable the reader to understand the relevant issues (along with a discussion of the practical implementation considerations) and help the reader to form their own opinion as to the merits, or otherwise, of the models presented. I will consider the book to be a success if it enables readers to understand the behaviour of models and to build better versions of them. This lean approach will hopefully make the volume attractive to practitioners (who do not always have the time to study a subject in detail) who wish to understand more about the properties of the credit derivative models commonly used in the marketplace. In particular it is envisaged that the volume will be of interest to a range of different types of practitioner.

- Quantitative analysts (quants) and quant developers wanting to understand more about credit modelling and credit derivatives. The book is written with a strong emphasis on models, implementation and understanding of the model behaviour. It is therefore well suited to quants in model validation teams, for example.
- Quantitative risk managers wanting to understand better the models used for valuation, to interpret synthetic CDO risk sensitivities (e.g. spread and correlation sensitivities) and risk manage complex credit portfolios.
- Traders and product controllers seeking a better understanding of the mechanics going on in the black boxes when 'F9' is pressed (and to understand the relative strengths and weaknesses of different models).
- Structurers wanting to understand better the properties of the instruments they are using to construct strategies with specific risk/return characteristics.
- Researchers in academia looking to understand some of the practical issues surrounding the common models used in the marketplace.

The downside to this lean approach is that for less experienced readers the material may at times not give as much explanation as would be liked, or some (basic) concepts are not described fully. However, for the motivated and intelligent reader this should present not a problem but a challenge and (as the author knows from experience) the rewards in terms of deeper understanding are worth the effort.

At the beginning of a project such as writing a book one has a vision as to what the finished product will look like. The vision for this book was that it would be very much model focused, with a strong emphasis on the practical, pragmatic implementation details that are of crucial importance in a live banking environment. This means there is less focus on the 'business' topics of the economics, mechanics and structures of credit derivatives than can be found in other texts. To include this information would have detracted from the core message of models and their properties. Also, when writing a book it is necessary to make compromises and be pragmatic in terms of content. At the beginning of the project one's vision of what will be achieved is vast and expansive. By the end of the project one is simply happy to stumble across the finish line. There are occasions throughout the book

when more detailed analysis of a particular model or scenario would have been very useful indeed to illustrate a particular point further, but due to time constraints was not included. On these occasions it is suggested that the reader build the models and do the analysis themselves as an exercise.

This leads into the next important point about the approach taken in the text. In the modern world of quantitative finance it is almost impossible to develop models of complex derivative trades that are wholly tractable analytically. It is therefore difficult to separate a model's mathematical description from its actual implementation. When it comes to building models suitable for use within a live investment banking environment the devil really is in the details. Full understanding of a model only comes from implementing it, analysing its properties and understanding its weaknesses. An important objective of this volume, therefore, is to provide not only the mathematical descriptions of the models, but also details of the practical implementation issues. To achieve this objective, liberal use is made of pseudo code to illustrate the implementation of an algorithm. The purpose of this code is to allow the reader to convert quickly a description of a model into the programming environment of their choice (although the author is most familiar with C++, and there may appear to be a bias towards the syntax of this language on occasion).

The volume is structured into three distinct sections. Broadly speaking Chapters 1–3 motivate the main topic, synthetic CDOs, and introduce some of the basic modelling tools necessary to describe them. Chapters 4–10 analyse the mathematical and computational modelling techniques applied to synthetic CDOs. Chapters 11–14 look at more advanced topics in the analysis of synthetic CDOs. Each of the chapters can in principle be read in isolation and each is relatively self-contained. However, there is a clear path from chapter to chapter (which reflects the author's own train of thought), particularly in Chapters 4–10. Reading each chapter sequentially will build a clearer and more coherent picture of the subject matter as a whole, but it is by no means a prerequisite.

In the first part of the book we motivate the study of synthetic CDOs by understanding their importance and usage within the broader credit derivatives marketplace. Chapter 1 provides a brief overview of the credit derivatives market in terms of instruments and introduces the crucial concepts of securitisation and tranching which are the basis of CDO technology. In this first section we also provide some of the basic mathematical building blocks necessary for later chapters. Chapter 2 describes the current market standard modelling methodologies for capturing the arrival of default risk of an obligor. This chapter also introduces the concepts and methods used for the modelling of default correlation, which as we will see is one of the most fundamental concepts in the characterisation of synthetic CDOs (and indeed any multi-name credit derivative). The first section of the book ends with a discussion, in Chapter 3, of the valuation models for the simplest and most

vanilla of credit derivatives – credit default swaps or CDSs. The market for single-name default protection CDSs is extremely liquid and a good understanding of the valuation methods for these basic building blocks is a necessary prerequisite for understanding the more complex multi-name products.[1] For a reader already conversant with single-name credit derivatives, the material in Chapters 1–3 will be familiar. Indeed these chapters are only included in order to provide a reference guide to the concepts underpinning the rest of the book.

The second part of the volume, Chapters 4–10, which is its mathematical and computational core, focuses specifically on the valuation and risk analysis of multi-name credit derivatives and synthetic CDOs in particular. Chapter 4 introduces the credit indices that have emerged and evolved over the course of the last few years. The introduction and subsequent trading of these indices has provided enormous impetus to the growth of the credit derivatives market. Chapter 5 then introduces default baskets. In terms of materiality, default baskets are a very small fraction of the overall structured credit marketplace. However, they are the simplest form of multi-name credit derivative and an understanding of their valuation and risk sensitivities can provide substantial insight into the behaviour of more complex synthetic CDOs.

Chapters 6 through 8 develop and analyse the core mathematical models for valuing synthetic CDOs. Chapter 6 describes a number of different methodologies for valuation and, in particular, introduces the current market standard valuation model, the so-called normal copula model. Chapter 7 investigates the fundamental behaviour of the model as certain key parameters are varied systematically. As will be seen in this chapter, the phenomenology of the model is relatively complex and subtle. Chapter 8 analyses the risk sensitivities of the standard market model to variations of input parameters. More importantly this chapter discusses the different risk sensitivity measures such as credit spread 01 (CS01) and value-on-default (VoD) that are necessary to capture and characterise the risk inherent in synthetic CDOs.

The next chapters look at the implications for the standard market model that standardised tranches and the development of a liquid market have had. Initially the market for synthetic CDOs was relatively illiquid and deals were done on a bespoke basis. The introduction of standardised credit indices and the subsequent development of a market for trading tranched exposures to slices of the index provided enormous impetus to the liquidity and volume of trades in single-tranche synthetic CDOs (STCDOs). Eventually the market became sufficiently liquid to allow transparent price discovery for the prices of these standardised index tranches. At this

[1] The main focus of the book is synthetic CDOs. Therefore we will not spend a great deal of time talking about CDSs and other credit derivatives – there are better texts available that describe these products in great detail.

point the role of the standard model changed; it became a mechanism whereby market participants could express and trade their views on default correlation. Chapter 9 introduces the concepts of implied and base correlations that have been developed to capture implied pricing information from market observed prices. As the prices of instruments become transparent in the open market it is crucially important for the standard model to be able to reproduce these prices accurately. Chapter 10 describes some of the different methodologies that have been developed to allow calibration of models of synthetic CDOs to market observed prices (the so-called 'correlation skew').

The final part of the volume, Chapters 11–14, looks at more advanced topics in the characterisation and analysis of synthetic CDOs. Chapter 11 introduces a number of exotic CDOs. Examples include CDOs with asset backed securities as the underlying pool of obligors as well as CDOs with CDOs as the assets in the underlying pool (so called CDO squareds). Correlation trading is the term used to refer to trading strategies designed to exploit the risk/return characteristics of portfolios of CDO tranches. Chapter 12 analyses the risk/return characteristics of a number of popular CDO trading strategies. Chapter 13 considers extending the models developed thus far for a single-tranche position to a portfolio of tranches and assesses how the risk in the tranche portfolio can be quantified and controlled.

Finally, a natural extension of analysing the static (in time) performance of CDO trading and hedging strategies is to look at the through life performance of the trading strategy. In the pricing of simpler derivatives, the value of the derivative is equal to the cost of the dynamic hedging strategy. If a hedging strategy is good at capturing all the risks a position is exposed to then the overall P/L generated from the process of selling the derivative instrument and rebalancing the hedging portfolio as the market risk factors evolve should be small. If the hedging strategy is not adequate there will be significant P/L leakage. Chapter 14 sets up and analyses a simple hedging simulation of synthetic CDO tranches. This chapter is more speculative in nature than previous chapters as it represents the cutting edge of technology applied to the analysis of complex derivative securities.

Acknowledgements

A book is never written in isolation, and it is a pleasure to acknowledge the contribution that a number of individuals have made to the current text. I would like to thank all the people I have worked with in the Model Validation and Risk Management teams of Credit Suisse and Barclays Capital as well as my co-workers at Cheyne Capital Management. A lot of the experience that is encapsulated in this text is a direct result of day-to-day interactions with my colleagues at these institutions. In particular, I would like to thank Dr Niclas Sandstrom of Barclays Capital and Dr Andrea Petrelli of Credit Suisse for their detailed reading of the original manuscript, and for making numerous suggestions as to how it could be improved.

I would also like to thank my editor at CUP, David Tranah (and all the other staff who have contributed to the bringing to fruition of this project), for providing me with an opportunity to write this book. Finally I would like to acknowledge the contribution of my Ph.D. supervisor Professor Sir S. F. Edwards of the Cavendish Laboratory, Cambridge. The scientific training I received under his tutelage has proven to be of enduring value throughout my career. I hope this text reflects some of what I learnt from him.