

Credit valuation adjustments with application to credit default swaps

by

Cara Milwidsky

Submitted in partial fulfillment of the requirements for the degree

Magister Scientiae

to the Department of Mathematics and Applied Mathematics
in the Faculty of Natural and Agricultural Sciences

University of Pretoria
Pretoria

August 2011

Declaration

I, Cara Milwidsky, declare that this dissertation, which I hereby submit for the degree Master of Science at the University of Pretoria, is my own work and has not previously been submitted by me for a degree at this or any other tertiary institution.

Cara Milwidsky
1 September 2011

Abstract

The credit valuation adjustment (CVA) on an over-the-counter derivative transaction is the price of the risk associated with the potential default of the counterparties to the trade. This dissertation provides an introduction to the concept of CVA, beginning with the required backdrop of counterparty risk and the basics of default risk modelling. Right and wrong way risks are central themes of the dissertation. A model for the pricing of both the unilateral and the bilateral CVA on a credit default swap (CDS) is implemented. Each step of this process is explained thoroughly. Results are reported and discussed for a range of parameters. The trends observed in the CDS CVA numbers produced by the model are all justified and the right and wrong way nature of the exposures captured. In addition, the convergence and stability of the numerical schemes utilised are shown to be appropriate. A case study, in which the model is applied to a set of market scenarios, concludes the dissertation. Since the field is far from established, a number of areas are suggested for further research.

Acknowledgements

I would like to thank my supervisor, Prof. Eben Maré, for the valuable advice he has offered throughout the course of my degree and for sharing the experience he has accumulated over the years. Thank you to T. Schelfaut and P. Mayar for finding an easy way out of the characteristic function trap and for taking the time to share it with me. To Ivan Ruscic, thank you for the many hours you have spent discussing ideas and proof reading, as well as for taking a genuine interest in this dissertation. Thank you to both Ivan Ruscic and Glenn Brickhill for giving me the opportunity to implement a CDS CVA model of my choice. To my friends and family, thank you for the tremendous encouragement, patience, support and understanding you have demonstrated. And especially to James for tolerating the piles of papers you so despise.

Cara Milwidsky

1 September 2011

“Courage doesn’t always roar. Sometimes courage is the quiet voice at the end of the day saying, ‘I will try again tomorrow’.”

Mary Anne Radmacher

Contents

I	Introduction, Background and Context	18
1	Introduction	19
1.1	Structure of the Dissertation	19
1.2	Chapter Summary	21
2	Counterparty Credit Risk	23
2.1	Counterparty Credit Exposure	24
2.2	Potential Future Exposure	25
2.3	Counterparty Risk Mitigants	28
2.3.1	Netting Agreements	28
2.3.2	Credit Support Annex	30
2.3.3	Early Termination Provisions	30
2.4	Exposure Management & Credit Default Swaps	31
2.5	Chapter Summary	33
3	Credit Valuation Adjustment	34
3.1	An Introduction to CVA	34
3.1.1	Right and Wrong Way Risk	36
3.2	Unilateral CVA	39
3.3	The Bilateral Case	41
3.3.1	Problems with Hedging DVA	45
3.4	A Simple Example	47
3.4.1	A Formula for Unilateral CVA with Independence	47
3.4.2	Application to a Vanilla Interest Rate Swap	48
3.4.3	A Formula for Implementation	49
3.4.4	The Bilateral CVA on a Swap	50
3.4.5	Numerical Results	51

3.5	Practical Implementation Aspects	54
3.5.1	Central Management of CVA	55
3.5.2	System Challenges	57
3.6	Two Relevant Instruments	58
3.6.1	Contingent Credit Default Swap (CCDS)	58
3.6.2	Credit Extinguisher	59
3.7	Chapter Summary	61
4	Modelling Default Risk	62
4.1	Structural Models	62
4.1.1	The Merton Model	62
4.1.2	First Passage Time Models	64
4.1.3	Further Remarks on the Structural Approach	66
4.2	Reduced Form Modelling	66
4.2.1	The Role of Poisson Processes	66
4.2.2	Time Homogeneous Poisson Process	67
4.2.3	More General Poisson Processes	70
4.3	CDS Pricing	72
4.3.1	General Formula	72
4.3.2	Market Model	75
4.4	Chapter Summary	77
5	Literature Review	79
5.1	Structural Models	79
5.2	Intensity Models	83
5.2.1	Contagion Models	84
5.2.2	Copula Methods	86
5.3	Alternative Approaches	89
5.3.1	A Generalized Affine Model	89
5.3.2	Markov Chain Models	90
5.4	Chapter Summary	92
II	CDS CVA Model Implementation - the Unilateral Case	93
6	Model Background	94
6.1	Credit Default Swaps in the Spotlight	94

6.2	Model Selection	98
6.3	Preamble	101
6.3.1	Notation	101
6.3.2	Assumptions	102
6.3.3	The Stochastic Intensity Process	103
6.4	Model Outline	105
6.5	Chapter Summary	107
7	Calibrating Intensities & Simulating Default Times	108
7.1	Calibrating the Intensity Process	108
7.1.1	Calibration Results	111
7.2	Simulating the Default Times	112
7.2.1	Simulating the Hazard Rates	112
7.2.2	Integrating intensities	117
7.2.3	Determining τ_1 and τ_2	117
7.2.4	Reproducing the Market Survival Curve	123
7.3	Chapter Summary	124
8	Revaluation of the CDS upon Counterparty Default	126
8.1	Survival Probability Preliminaries	127
8.1.1	A Practical Formulation of the Survival Probability	128
8.2	Applying the Fractional Fourier Transform	131
8.2.1	Writing the CDF in the Form of an FRFT Sum	132
8.2.2	The Characteristic Function	134
8.2.3	A Trap	134
8.2.4	Rescued!	140
8.2.5	Parameters Required for the FRFT Implementation	143
8.2.6	A Note on the Limit of the Integrand	154
8.3	CDS Valuation	155
8.3.1	Obtaining the Conditional Copula Function	155
8.3.2	The Survival Probability	158
8.4	CVA Computation	159
8.5	Chapter Summary	160
9	Results of the Model Implementation	162
9.1	Convergence	163

9.2	CVA Results	169
9.2.1	Reference Entity Riskier than Counterparty	169
9.2.2	Counterparty Riskier than Reference Entity	173
9.2.3	A Change in Reference Entity Riskiness	175
9.2.4	Risky Reference Entity	177
9.2.5	Differing Maturities	179
9.3	Chapter Summary	181
 III A Bilateral CVA Model		183
 10 Extending the Model to the Bilateral Case		184
10.1	Notation	185
10.2	Assumptions	186
10.3	Simulating Default Times	188
10.4	Revaluation of the CDS upon Default	190
10.4.1	The Fractional Fast Fourier Transform	191
10.4.2	The Conditional Copula Function	192
10.4.3	CDS Valuation	197
10.4.4	The Bilateral CVA Computation	197
10.5	Chapter Summary	198
 11 Results of the Bilateral Model Implementation		199
11.1	Investor and Counterparty Equally Risky	201
11.2	A Less Risky Investor	203
11.3	Decreasing the Riskiness of the Counterparty	204
11.4	Altering the Reference Entity	206
11.5	The Effect of a Non Zero Investor-Counterparty Default Time Correlation	208
11.6	Chapter Summary	210
 IV Case Study and Conclusion		212
 12 Case Study		213
12.1	Market Data	213
12.2	Model Calibration	216
12.2.1	The Volatility Parameter	216
12.2.2	Correlation	218

12.3 Results	220
12.4 Chapter Summary	221
13 Conclusion	222
13.1 Summary	222
13.2 Directions for Further Research	226
V Appendices	238
A Proofs of the General Counterparty Risk Pricing Formulae	239
A.1 Proof of the Unilateral Formula	239
A.2 Bilateral Formula	241
B Matters Related to CDS Pricing	244
B.1 Bootstrapping the Par CDS Curve	244
B.2 Derivation of the CDS Premium and Default Leg Valuation Formulae . . .	248
B.2.1 CDS Premium Leg (Equation (4.27))	248
B.2.2 CDS Protection (Default) Leg (Equation (4.28))	249
B.3 Derivation of the Constant Hazard Rate	250
C Formulae Related to the CIR Process	252
C.1 CIR Instantaneous Forward Rate	252
C.2 CIR Bond Price	253
C.3 CIR Conditional Expected Value	253
C.4 CIR Conditional Variance	253
D A Basic Introduction to Copulas	254
D.1 Concepts Required to Define Copulas	254
D.2 The Definition of a Copula	255
D.3 Sklar's Theorem	256
D.4 The Uniform Distribution of CDFs	257
D.5 Copula Densities	257
D.6 Copula Conditional Distributions	258
D.7 Proof of the Statement that a Trivariate Gaussian Copula Induces Bivariate Marginals	258

E	Supplementary Information Relating to Chapter 8	260
E.1	Approximating the Integral using the Trapezoidal Rule	260
E.2	Derivation of the New Characteristic Function Formulation	261
E.3	Testing the Probabilities against the CIR Bond Price	262
F	Fourier Transforms	264
F.1	Definition of a Fourier Transform	264
F.2	Characteristic Functions	265
F.3	Fourier Transforms in Finance	265
F.4	Shortcomings of the FFT	265
F.5	The fractional Fourier transform	266
F.6	A procedure for implementing the FRFT	267
G	An Explanation of the Unexpected Non Zero Receiver CVA	268
G.1	Default Time Recap	268
G.2	Survival Probabilities Illustrated	269
G.3	Viewing the Distributions Responsible	271
G.4	Conclusion	273
H	Matlab Code	275
H.1	CIR ++ Simulation	275
H.1.1	Simulating CIR Process	275
H.1.2	Deterministic Shift Computation	277
H.1.3	Simulating Default Times - Unilateral CVA	278
H.2	FRFT	279
H.2.1	Variance of Integrated CIR Process	281
H.3	The function f_j in Unilateral Model	281
H.3.1	Simulating Default Times - Bilateral CVA	282
H.3.2	Conditional Copula Function f at τ_2 in Bilateral Model	283

List of Figures

2.1	An illustration of the concepts PFE and EPE for two different derivatives .	27
2.2	Mechanics of a credit default swap	32
3.1	Mechanics of a vanilla interest rate swap	48
3.2	Swaption values and marginal default probabilities for Case 2	54
3.3	Central management of CVA	55
3.4	Mechanics of a contingent credit default swap	59
4.1	Illustration of the Merton model	65
6.1	Five year CDS spread and credit rating of Lehman Brothers	96
6.2	Diagram of the model outline	106
7.1	Market-implied hazard rates, fitted CIR rates and deterministic shift for the base case	112
7.2	Integrated market-implied hazard rate, CIR rate and deterministic shift for the base case	113
7.3	Simulation of CIR process using transition density	115
7.4	10,000 simulated CIR++ paths for the high risk entity	116
7.5	PDF of bivariate Gaussian copula for various correlations	119
7.6	CDF of the Gaussian copula	120
7.7	Simulation of uniform random variables from Gaussian copula	120
7.8	Ratio of default times correlations	122
7.9	Exact and simulated survival probabilities for each entity	124
8.1	Integrand for various CIR volatility parameter values	135
8.2	Real part of characteristic function for various volatility parameter values .	135
8.3	Aspects of the integrated CIR process characteristic function	136
8.4	Investigations into the denominator without the exponent	137

8.5	Phase of the entire denominator	138
8.6	Investigations into f for $\nu = 0.05$	139
8.7	Real part of characteristic function for various time parameter values	140
8.8	Argument of f for small values of t	141
8.9	Illustration of the Rotation Count Algorithm	142
8.10	Phases of the new Formulation of $\phi(u)$	143
8.11	Imaginary versus real part of $g(u)$	144
8.12	Real part of revised characteristic function	144
8.13	Integrand obtained from the revised characteristic function	145
8.14	Integrand for various time parameter values	148
8.15	Breakdown of integrand for a selection of time parameters	149
8.16	Integrand for various volatility parameter values	149
8.17	Integrand for a selection of volatility parameters with $t = 0.5$	150
8.18	Integrand for μ and y_0 variations	150
9.1	CVA in basis points. Reference entity: intermediate risk; counterparty: low risk.	170
9.2	CDS CVA in basis points. Reference entity: intermediate risk; counter- party: high risk.	174
9.3	CDS CVA - reference entity: low risk; counterparty: high risk	176
9.4	CDS CVA - reference entity: high risk; counterparty: low risk	178
9.5	CDS CVA - reference entity: high risk; counterparty: intermediate risk	179
9.6	CVA for various CDS maturities	180
10.1	Simulation of uniform random numbers from a trivariate Gaussian copula	189
10.2	The conditional distribution of trivariate Gaussian random variables	196
11.1	Illustration of the trade direction from both investor and counterparty per- spectives	200
12.1	Survival probabilities of entities in case study	215
12.2	Market-implied hazard rates, fitted CIR rates and deterministic shift Lehman Brothers 2008	218
12.3	Correlations for case study	219
B.1	Implied hazard rate and survival probability term structure for SOAF at COB 31 August 2010	246
B.2	CDS spreads and survival probabilities for curves in Table B.2	247

G.1	Reference entity survival probabilities for various default times on the same simulation path	270
G.2	Reference Entity Survival Curve at Counterparty Default Time, 0.75 Years, for Various Default Time Correlations.	272
G.3	The components of the survival probability approximation in (G.1).	272

List of Tables

3.1	CVA in ZAR for a long two year IRS with a notional of ZAR1m	52
7.1	CDS term structures for numerical illustrations	111
7.2	CIR++ parameters	111
7.3	Comparison of exact and simulated survival probabilities	123
9.1	Spreads in basis points for entities in the numerical illustrations	162
9.2	Receiver CVA convergence - low ν_1	164
9.3	Payer CVA convergence - low ν_1	165
9.4	Receiver CVA convergence - high ν_1	167
9.5	Payer CVA convergence - ν_1	168
9.6	CDS CVA - reference entity: intermediate risk; counterparty: low risk . . .	171
9.7	CDS CVA - Reference entity: intermediate risk; counterparty: high risk . .	175
9.8	CDS CVA - reference entity: low risk; counterparty: high risk	177
9.9	CVA for various CDS maturities	180
11.1	CDS BCVA - investor, reference entity and counterparty: intermediate risk	202
11.2	CDS BCVA - reference entity and counterparty: intermediate risk; investor: low risk	204
11.3	CDS BCVA - investor and reference entity: intermediate risk; counterparty: low risk	205
11.4	CDS BCVA - reference entity: low risk; investor and counterparty: inter- mediate risk	206
11.5	CDS BCVA - intermediate risk entities; $\rho_{0,2} = -0.6$	208
11.6	CDS BCVA - intermediate risk entities; $\rho_{0,2} = 0.6$	209
12.1	Case study CDS spreads. Source: Bloomberg	214
12.2	BCVA Results in basis points	220

B.1	SOAF spreads for COB 31 August 2010	244
B.2	SOAF spreads, hazard rates and survival probabilities with spread shifts .	247

Glossary

Counterparty credit risk

The possibility that a party to an over the counter (OTC) derivative transaction will fail to perform on its contractual obligations, causing losses to the other party.

Counterparty credit exposure

The counterparty credit exposure on a single derivative at time t is the maximum of zero and the time t value of the derivative.

Counterparty credit adjustment (CCVA)

The portion of the bilateral credit valuation adjustment attributable to the counterparty's default risk.

Credit default swap (CDS)

An agreement to pay a fixed spread on a certain notional at regular intervals (usually quarterly) in exchange for the loss on the principle amount outstanding on a certain reference obligation (such as a bond) in case the entity that issued the bond (the reference entity) defaults.

Credit valuation adjustment (CVA)

The difference between the risk-free value of a derivative or portfolio of derivatives and the fair value of these when the possibility of a counterparty default is taken into account.

Unilateral CVA arises when only the counterparty's risk of default is incorporated into the derivative valuation (one-sided CVA). **Bilateral CVA (BCVA)** arises when the value of the derivative conducted between risk-free counterparties is adjusted by default risk of both counterparties to the trade.

Debt valuation adjustment (DVA)

The portion of the BCVA attributable to the investor's default risk (the own risk portion of the BCVA). This increases the value of the identical derivative conducted between risk-free entities.

Investor

The entity from whose perspective the CVA is computed.

Loss given default (LGD)

The portion of the value of the derivative trade that is not recovered from the assets of an entity when it defaults. If the recovery rate is R , $LGD = 1 - R$.

Mark to market (MTM)

The value of a derivative obtained using current market observable inputs.

Potential future exposure (PFE)

The possible future values that a derivative might assume. The PFE of an instrument is computed by simulating many different scenarios of the main price drivers of the instrument's value at different points in time. It is then revalued at each of these times and the PFE profile is obtained by selecting a percentile (at a confidence level) of the distribution of potential future exposures at each simulation date.

Right way risk

The risk that arises from a negative dependency between the default probability of the counterparty and the value of the derivative traded with the counterparty.

Wrong way risk

The risk that arises from a positive dependency between the default probability of the counterparty and the value of the derivative traded with the counterparty.

Part I

Introduction, Background and Context

Chapter 1

Introduction

A credit valuation adjustment (CVA) is an amendment to the fair price of an over-the-counter derivative transaction to account for the default risk of the counterparties to the trade. It is the price of the counterparty risk inherent in the deal. The focal point of this dissertation is the modelling of the credit valuation adjustment on a credit default swap (CDS). This instrument was chosen due to its relevance as a basic counterparty risk hedge. CDS CVA is also intriguing from a modelling perspective, since the underlying contract has a default risk component in itself. Finally, it illustrates the concepts of right and wrong way risk.

The initial objective is to familiarise the reader with the concept of CVA. This includes an appreciation of the necessity of the adjustment and a consideration for the practical issues associated with its implementation. The second goal of the dissertation is to implement a model for the valuation of CDS CVA and to provide a justification of the results that it produces.

1.1 Structure of the Dissertation

There are five parts to this document. These divisions categorise the various chapters in a manner that is designed to facilitate a clear and concise comprehension of the flow of information and ideas. The first of these parts is entitled *Introduction, Background and Context*. It begins with a description of the key concepts related to counterparty risk, the building block of CVA, in Chapter 2.

This is followed by an introduction to the credit valuation adjustment, the central theme

of the dissertation, in Chapter 3. The related terms, right and wrong way risk, are also introduced. In addition to defining CVA, the chapter explains the basic techniques employed in its measurement. It discusses the problems associated with this relatively novel field, combined with the challenges of institution-wide implementation. The terms *unilateral* and *bilateral CVA* are introduced. Unilateral CVA refers to a single-sided counterparty risk adjustment. Bilateral CVA arises when an investor measures its own risk of default in addition to that of its counterparty. Furthermore, the approaches to interpreting CVA are explored.

Chapter 4 is a discussion of the basic techniques employed in the modelling of default risk. The two main types of methodologies set forth are structural and reduced form (intensity-based) techniques. The focus is on the latter since the CDS CVA model implementation that follows is intensity-based. This is also the market standard approach to default risk modelling. It is applied to the pricing of a credit default swap. Knowledge of this technique is essential for bootstrapping survival probabilities from observed credit spread term structures.

Chapter 5 concludes the first part of the dissertation, with a review of the literature that is devoted to pricing the CVA on a CDS. The concepts introduced in Chapter 4 are requisite background for an understanding of the literature review. The variety of approaches referenced in this chapter are diverse. They illustrate the lack of a single established market model for the valuation of CDS CVA.

The second part of the dissertation is entitled *Model Implementation - the Unilateral Case*. It describes the unilateral model that was selected for implementation from among those discussed in Chapter 5. The first chapter in the section, Chapter 6, provides a description and an outline of the model, along with the reasons for its selection. The following two chapters, Chapters 7 and 8, present a detailed description of the implementation of the model. The former explains the calibration of the stochastic intensity process to the credit spreads observed in the market, in addition to describing the default time simulation procedure. The focus of Chapter 8 is on the revaluation of the CDS at the counterparty default times. A description of the method employed in calculating the CVA from these CDS values concludes the chapter. A range of numerical results, obtained by running the implemented model, are reported and discussed in Chapter 9.

The bilateral extension of the unilateral model is the focus of the third part of the dissertation. Chapter 10, the first of two chapters in the section, provides a description of the approach used to extend the unilateral CDS CVA to one that incorporates the investor's own risk of default. The bilateral CVA on a CDS, for a variety of scenarios, is reported and examined in Chapter 11.

The fourth part of the dissertation is comprised of two chapters. The first of these, Chapter 12, is devoted to a case study. The CVA on a five year CDS is computed at four separate points in time, namely: in 2006 - two years before the occurrence of the global credit crunch; on the last business day prior to the default of Lehman Brothers in 2008; in August 2010, two years subsequent to the crisis and, more recently, in August 2011 during the sovereign debt crisis. Chapter 13 concludes the dissertation and proposes further areas of research. It is followed by the bibliography and the fifth and final part of the document, which contains the appendices. These include information that supplements the understanding of the main text, in addition to background material that is required for the completion of the model implementation.

1.2 Chapter Summary

The focal points of the dissertation were introduced in this chapter. Furthermore, the structure of the document was set forth. The objectives of the remainder of this dissertation are listed below:

- (i) Introduce the concept of CVA, in addition to the background required for the understanding thereof.
- (ii) Provide insight into the term *credit valuation adjustment*, with respect to its measurement, its validity in derivative pricing and its implementation in a financial institution.
- (iii) Present an overview of the models documented in the literature for the pricing of the CVA on a credit default swap.
- (iv) Explain the implementation of a model used for pricing unilateral CDS CVA, along with a justification of the model selection.
- (v) Report and examine the CVA results produced by the model.

- (vi) Extend the model to the bilateral case and discuss the CVA values obtained by implementing such an approach.
- (vii) Demonstrate the application of the bilateral CVA model to a real-world scenario by means of a case study.

Chapter 2

Counterparty Credit Risk

Credit risk is the threat of promised repayments on outstanding instruments, such as loans and bonds, not being received. Counterparty credit risk is the possibility that a party to an over the counter (OTC) ¹ derivative transaction will fail to perform on its contractual obligations, causing losses to the other party [2]. The *credit* risk inherent in financial assets, such as bonds and loan obligations, has long been a part of financial markets. A simple comparison of the yields on otherwise identical bonds issued by two separate corporate entities supports this well-known fact. The pricing of the *counterparty* risk associated with derivatives is, however, a relatively new field [3].

Counterparty credit risk, as well as the accurate measurement and active management thereof, is currently a pertinent subject in finance. The significance of having the ability to quantify the exposure one has to a particular counterparty at a given point in time was highlighted by the 2008 global financial crisis. At the height of this period, the frequent threat of bankruptcy among high profile institutions demonstrated the necessity of having accurate exposure figures readily available. Additionally, the misconception that the counterparty risk on certain entities (deemed to be risk-free) could be overlooked was set aside by the bankruptcy filing of Lehman Brothers² in September 2008.

By the end of 2010, the total outstanding global notional amount of OTC derivatives had reached 601 trillion United States dollars (USD) [4]. A more meaningful statistic is the overall global gross credit exposure (gross credit exposures taking into account legally enforceable bilateral netting agreements) of 3.3 trillion USD [4]. These numbers indicate

¹An OTC transaction is one that is not traded through an exchange [1].

²Lehman Brothers held an investment grade credit rating at the time of its default.

that the counterparty risk inherent in OTC markets cannot be ignored.

The focus of this chapter is on the basic concepts related to counterparty risk and its measurement. For readers unfamiliar with the topic, this background is essential to the understanding of the credit valuation adjustment (CVA) which is introduced in the next chapter. For those already comfortable with the subject of counterparty risk, the chapter serves as a useful reminder of the concepts that are key to understanding CVA.

2.1 Counterparty Credit Exposure

Suppose an investor enters into a derivative transaction with a counterparty. In the event that the counterparty fails to honour its obligation, the investor is required to close out its position with the defaulted counterparty. It is convenient to assume that the investor replaces the derivative in order to keep its market position unchanged³. The loss to the investor is then determined by the replacement cost of the derivative upon default of the counterparty [5]. This, in turn, is a function of whether the derivative has a positive or a negative value to the investor at the default time. Ignoring liquidity constraints, one of the following two situations arises:

- (i) If the value of the derivative is positive to the investor (that is, in-the-money from the investor's perspective):
 - (a) It receives the portion of the value of the derivative that is recovered from the assets of the defaulted counterparty⁴. This portion could be zero and is determined by the amount by which the counterparty's liabilities exceed its assets on default. It is termed the recovery rate.
 - (b) It enters into the same derivative with a new counterparty.
 - (c) It has made a net loss of the portion of the derivative not recovered from the defaulted counterparty. This amount is called the loss given default (LGD).
- (ii) If the value of the derivative is negative to the investor (that is, out-of-the-money from the investor's perspective):

³This assumption permits an intuitive understanding of the loss made by the investor upon default of the counterparty. This is as a result of the fact that the loss is determined by the market value of the derivative position with the counterparty at the default time. The assumption can thus be made without loss of generality.

⁴This may take some time. For large institutions that have entered into many complex deals, it is not uncommon for the time between the default of the counterparty and the determination and settlement of the portion recovered by its creditors to take a few years.

- (a) It closes out the deal by paying the counterparty the amount it owes under the derivative transaction.
- (b) It enters into a new transaction identical to the original one, but with another counterparty. In doing so, it receives an amount equivalent to the loss made in closing out the original transaction.
- (c) It has made zero profit or loss on the trade. Economically, transaction costs aside, the investor's position is unchanged post the default event of the original counterparty.

The scenarios above describe the two possible outcomes, from the investor's perspective, of a counterparty default event. The investor either makes a loss or is unaffected by the counterparty's default. The former occurs when the investor is in-the-money at the time of the default event and the latter is as a result of the investor owing the counterparty money at the default time. Notice the asymmetry in the investor's exposure. Its loss is either zero or positive, but never negative⁵.

The investor's exposure on the derivative contract is therefore the maximum of 0 and the product of the value of the derivative at the counterparty default time and the loss given default (LGD). This is often expressed assuming zero recovery and thus an LGD of 100%. In other words, if we let $V(t)$ denote the value of a single derivative contract at time t , the investor's exposure, on a contract level, is given by:

$$E(t) = \max(V(t), 0). \quad (2.1)$$

2.2 Potential Future Exposure

In measuring an investor's exposure to a particular counterparty, it is not merely the current value of a derivative position that is of concern, but the potential future values that the transaction might assume. Consider, for example, a vanilla interest rate swap between two counterparties in which one pays a fixed interest rate quarterly in exchange for a floating rate of interest⁶. Ignoring bid-offer spreads, the fixed rate is likely to be selected such that the value of the transaction is zero to both parties on the initial trade date. Over time, however, one party will be in-the-money and the other out-of-the-money on the derivative position. This illustrates the point that, while the current counterparty

⁵In other words, its profit is never positive, but may be negative.

⁶For example JIBAR (Johannesburg Interbank Rate) or LIBOR (London Interbank Offer Rate).

exposure on a derivative position may be irrelevant, the potential values it might assume at some time in the future are an important consideration in determining counterparty risk.

In addition, the counterparty to the trade may currently have a low probability of defaulting during the life of the contract. This is likely to change over time. The future value of the derivative, when the credit quality of the counterparty may have deteriorated, is thus relevant.

The potential future exposure (PFE) of an instrument is computed by simulating many different scenarios of the main price drivers of the instrument's value at different points in time. It is then revalued at each of these times and the PFE profile is obtained by selecting a percentile of the distribution of potential future exposures at each simulation date [5]. The magnitude of the percentile is generally a management decision, but it is usually a large value such as the 95th or the 99th percentile.

In most investment banks and large financial institutions, PFEs have traditionally been the main tool for counterparty credit risk management. A counterparty is typically granted a credit line by a credit risk manager, with a certain limit on the institution's exposure to that counterparty. The peak of the PFE percentile over the life of the portfolio of transactions with the counterparty is then measured against the limit to determine if additional trades are permissible.

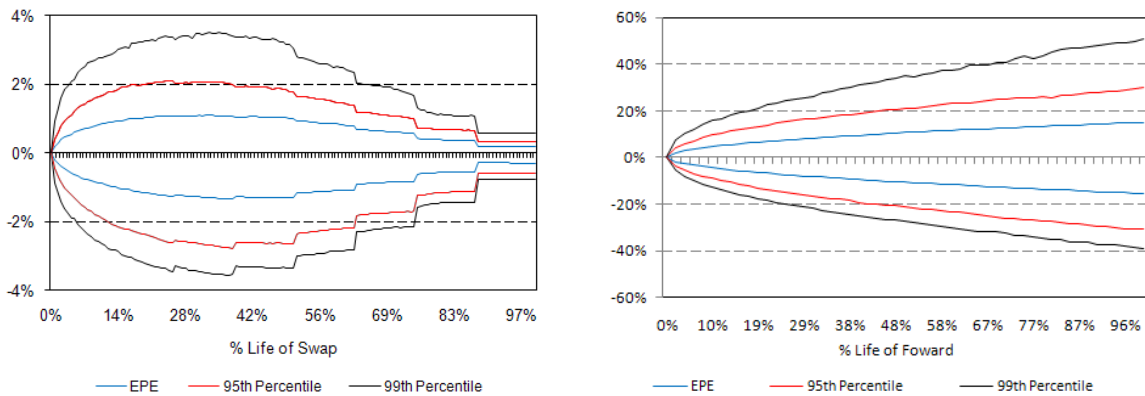
An additional statistic that is commonly employed in the measurement of counterparty risk is the expected positive exposure (EPE) to the counterparty. This is the mean of the positive part of the distribution of simulated values [1]. As demonstrated in the previous section, it is only scenarios in which the derivative has a positive value to the investor (from whose perspective we are measuring the risk) that are relevant. Scenarios in which the position has a negative value to the investor do not have an economic impact on its position as it has to honour its obligation regardless of the default status of its counterparty. Note that the paths in which the counterparty credit exposure is zero are included in the EPE computation. Therefore it is the average of the counterparty credit exposure at each point in time.

The 95th and 99th percentile PFEs and the EPE of a vanilla interest rate swap and

a foreign exchange (FX) forward are illustrated in Figure 2.1 below. The swap exposure peaks roughly a third of the way through its life while the potential loss on the forward increases until the maturity of the instrument.

There are two conflicting factors at play in the swap PFE profile. The first of these is the volatility of the underlying interest rate, which tends to increase the potential future exposure when applied over time. On the other hand, the number of cash flows that are due to occur decreases over time, reducing the potential loss that can be incurred by a counterparty default event. These two factors offset each other. The first part of the PFE tends to increase due to the volatility effect being dominant. It then decreases as a greater number of cash flow dates occur in the past and the number of potential cash flows that would be lost, should the counterparty default, are reduced.

The FX forward exposure, on the other hand, is subject only to the volatility effect since payment on this transaction is due at maturity. There are no intermediate cash flows that occur during the life of the forward. The potential exposure on this transaction thus increases progressively over time.



(a) Two year interest rate swap

(b) One year USD/ZAR foreign exchange forward

Figure 2.1: An illustration of the concepts PFE and EPE for two different derivatives

With reference to Figure 2.1, observe firstly that the potential future exposure increases with the confidence level. This is to be expected since a higher confidence level results in an increase in the possible exposure due to the uncertainty surrounding the future value of the underlying price drivers of the instrument.

We show the PFE percentiles and the EPE for both a long and a short position in the swap and the forward. In effect, we are showing the exposure to both parties on the same graph. The negative values, below the x-axis, are simply a result of the exposure being computed for a long position in both instruments. An investor with a short position would utilise the absolute value of the negative PFE and EPE profiles.

Notice also that, while the swap PFE is relatively symmetrical about the x-axis, the exposure on the FX forward is slightly larger for an investor who is long the forward. The forward contract (whose value was simulated through time) involves the long party paying South African rand (ZAR) in exchange for United States dollars (USD) upon maturity of the contract. The asymmetry in the exposure reflects the greater likelihood, as reflected in historical data, of ZAR (an emerging market currency) depreciating significantly against USD.

A final point that should be noted regarding the computation of counterparty exposure in terms of the PFE percentiles and the EPE, is that these measures are merely a function of the potential value that the derivative transaction might assume over time. The credit quality of the counterparty is not considered. The PFE on two identical instruments, one transacted with a poor credit quality counterparty and the other, with an almost risk-free counterparty, will therefore be identical despite the first transaction being riskier.

2.3 Counterparty Risk Mitigants

There are a number of practices adopted by investment banks and other financial institutions in order to reduce their exposure to a single counterparty. Here we discuss a few of the commonly utilised ones.

2.3.1 Netting Agreements

The discussion thus far has centered around the measurement of the counterparty risk on a single transaction. In practice, however, this risk is generally managed at a counterparty level. If a bank has more than one trade with a particular counterparty and the credit risk is not mitigated in any way, the bank's exposure to the counterparty is the sum of the exposures on each of the individual contracts with the counterparty. Thus, if $E_i(t)$ is

the exposure on the i 'th contract, the counterparty level exposure ($E(t)$) is:

$$E(t) = \sum_i E_i(t) = \sum_i \max(V_i(t), 0). \quad (2.2)$$

The exposure can be significantly reduced by entering into a netting agreement. This is a legally binding contract between two counterparties that allows the aggregation of transactions between the two counterparties in the event of a default [5]. It means that transactions with a negative value can be used to offset those with a positive value. For all trades under the netting agreement, the net counterparty exposure is then:

$$E(t) = \max\left(\sum_i V_i(t), 0\right). \quad (2.3)$$

Compare the position of the summation in (2.2) and (2.3). To appreciate the power of netting, consider an extreme example in which an investor has two opposite, but identical trades with a counterparty. The market risk on the total position will be zero. If the trades are booked under the same netting agreement, the counterparty credit risk will also be zero. However, in the absence of a netting agreement covering both trades, the credit quality of each of the counterparties will be of concern [6]. In the event of a counterparty default, the investor will owe the full amount on the trade that is out-of-the-money from his perspective. Conversely, only recovery will be received on the opposite (in-the-money) transaction.

The ISDA⁷ Master Agreement has become the industry standard for OTC derivative transactions [7]. Under the agreement, trades are documented based on trade confirmations and the Master Agreement acts as an 'umbrella' contract governing each of the individual trade confirmations. The result is a net amount payable by one party to the agreement to the other party. A cross-product netting agreement is a particular type of ISDA that cover trades from several asset classes such as fixed income, foreign exchange and equities [7].

Note that, even when netting agreements are not in place, there are advantages to measuring the exposure on a counterparty level. Consider, for example, an investor who is long an equity forward with a particular counterparty in addition to being short an option on the identical equity with the same counterparty. It is economically impossible for the

⁷International Swaps and Derivatives Association: www.isda.org.

value of both contracts to increase simultaneously. Computing the PFE for each of these trades separately would fail to produce an accurate exposure statistic.

2.3.2 Credit Support Annex

Although a netting agreement can significantly reduce a financial institution's exposure to an entity, the net residual exposure to a particular counterparty can still be quite large. Collateral agreements require counterparties to periodically mark to market their positions and to provide collateral (transfer the ownership of assets) to each other as exposures exceed pre-established thresholds [2]. The terms of such a collateral agreement are usually defined by including an annex to the existing ISDA agreement called a Credit Support Annex (CSA). In addition to the threshold level, the following are defined [2] under the CSA:

1. **Type of collateral:** The securities eligible to be transferred as collateral are defined. The emphasis is on liquidity since the collateral should be easy to monetise (convert to cash) in the event of a counterparty default. A haircut is often applied to various assets. In other words, the full value of the asset is not recognised as collateral.
2. **Minimum transfer amount:** The amount by which the counterparty exposure exceeds the threshold level before a collateral call is made.
3. **Direction of agreement:** The collateral agreement could be unilateral or bilateral. In the former case only one party is required to post margin in the event of extreme market moves. In the latter, both are required to do so, depending on which party is in-the-money.

2.3.3 Early Termination Provisions

There are a number of provisions by which the effective maturity of a trade is shortened, thereby reducing counterparty exposure. Two examples of these are liquidity puts and credit triggers. The former gives the parties to a transaction the right to settle and terminate the trade on pre-specified future dates, while the latter requires trades to be settled should the credit rating of a party fall below a particular level [2]. The events surrounding the 2008 subprime crisis and the role played by rating agencies shattered much of the confidence that was formerly placed in credit ratings. This was due to AAA

ratings being assigned to many of the collateralised debt obligations (CDOs) that were at the root of the crisis [8].

2.4 Exposure Management & Credit Default Swaps

There are a number of credit derivatives that can be employed in the management of counterparty credit risk. Here we explain the mechanism of a single name credit default swap (CDS). An understanding of this basic credit instrument is essential to a discussion on counterparty risk. CDSs are the most commonly traded credit derivatives and can be utilised in the hedging of counterparty credit risk. Note, however, that the hedge cannot be static. This is a result of the value of the derivative on which the counterparty risk is measured, the counterparty exposure, changing over time. A CDS is written on a fixed notional amount.

A CDS is an agreement to pay a fixed spread on a certain notional at regular intervals⁸ in exchange for the loss on the principle amount outstanding on a certain reference obligation (such as a bond) in case the entity that issued the bond (the reference entity) defaults. Typically, default is defined as failure to pay a significant promised cash flow. In terms of CDSs, default generally refers to the mechanism triggering the CDS. This might be failure to pay, bankruptcy or restructuring [9]. Whether or not an event is considered to be a default event is determined by the ISDA Credit Derivatives Determination Committee for the relevant region⁹. Once default has occurred, the premium payments cease. An investor who purchases protection, or pays the spread, is said to own a payer CDS. The counterparty to the trade holds a receiver CDS and is short protection. A CDS is therefore a form of insurance.

Once a counterparty has defaulted, usually only a portion of the amount owing by the entity is recovered by its creditors. We call this portion the *recovery rate* (R) and the fraction that is not recovered is the *loss given default* (LGD) so that:

$$\text{LGD} = 1 - R. \quad (2.4)$$

The payment that will be made by the protection seller to the buyer in the event of default is the product of the *LGD* and the notional amount of the CDS contract. Physical or

⁸usually quarterly.

⁹<http://www.isda.org/credit/>

cash settlement is another specification made in the contract. The convention is rapidly becoming cash settlement. In fact, there are often more CDSs traded on a particular name than there are bonds outstanding for that entity, rendering physical settlement impossible.




Figure 2.2: Mechanics of a credit default swap

The final determination of the actual recovery rate can take many years, particularly for large and complicated bankruptcies. It is also dependent on the level of seniority of the debt that is held. For the purpose of settling CDSs, an auction is held shortly after the default event in order to determine the recovery rate. The CDS recovery rate and the actual recovery rate on derivatives held by the defaulted entity's creditors may be very different. An indication of the market expectation of the CDS recovery rate prior to the auction will be the value at which the defaulted entity's bonds are trading. During this brief period between default and CDS settlement, instruments called recovery swaps (by means of which market participants can bet on the recovery rate) are also sometimes traded. They can be used as an additional means of gauging market expectations of the recovery. The actual auction procedure is described in a document [10] by Markit. As with any market-driven mechanism, the auction may be influenced by supply and demand as well as by technical factors [11].

A final, but integral point to highlight regarding credit defaults swaps, is the presence of two forms of default risk. The underlying reference obligation represents a *credit risk* in the traditional sense. This is because it refers to a bond or loan on which the principle outstanding by the reference entity is known in advance. The protection buyer and seller represent a *counterparty risk* to each other since they have traded an OTC derivative with each other¹⁰. The amount at risk is not known in advance, but depends upon the

¹⁰With credit risk underlying the derivative.

value of the reference entity's credit spreads¹¹ at the default time of the counterparty.

2.5 Chapter Summary

This chapter served as an introduction to counterparty risk and the key concepts with which it is associated. It was intended to provide the background knowledge required for the introduction of credit valuation adjustments in the following chapter. The terms credit and counterparty risk were both introduced. These were followed by a description of the counterparty credit exposure and its measurement by means of the computation of the potential future exposure (PFE). An introduction to counterparty risk mitigants and the credit default swap instrument was then presented.

¹¹CDS spreads are the main price driver. The discount factors will also play a role, as will the exchange rate if the CDS is traded in a foreign currency.

Chapter 3

Credit Valuation Adjustment

The credit valuation adjustment (CVA) is a relatively new concept in modern finance. It is the difference between the risk-free value of a derivative or portfolio of derivatives and the fair value of these when the possibility of a counterparty default is taken into account [12]. Effectively, it assigns a price to the counterparty risk inherent in OTC derivatives. It should incorporate the default risk associated with the entities on both sides of the trade. This chapter serves as an introduction to the concept of CVA and a discussion of the issues associated with it.

3.1 An Introduction to CVA

The reduction in the risk-free value of financial instruments due to repayment risk is not a new concept. The default risk associated with bonds and loan obligations (lending risk) has long been priced into the valuation of these instruments. Two features that set CVA apart from traditional forms of credit risk [3] are:

- (i) **The uncertainty of future exposures:** The replacement cost of the derivative upon default of one of the counterparties depends on the values of the underlying risk drivers of the instrument's value at the default time. This is in contrast to instruments such as bonds where the amount at risk, the face value, is known in advance.
- (ii) **The bilateral nature of the counterparty risk:** Since the value of most derivatives can be either positive or negative to both parties, it is not known in advance which entity will be in-the-money at the time of default. In traditional forms of credit risk, it is known that the borrower owes the lender and the repayment risk is

one-sided.

The focus on precise measurement and active management of counterparty risk is currently permeating global financial activity. Efforts to enhance these functions have been increased since the 2008 credit crisis highlighted the vulnerability of the financial system to counterparty risk. Its management is no longer exclusively the function of back and middle office risk managers. Modern financial institutions recognise the need to dynamically and proactively manage counterparty risk. For this reason, CVA trading desks have been established in a number of foreign banks, with an increasing number of local (South African) banks following suit.

The accurate modelling of CVA is critical to these endeavours. It enables banks and other financial institutions to price counterparty risk directly into transactions, fully accounting for the cost of carrying or hedging the risk [12]. Furthermore, its incorporation into derivative valuation leads to a more comprehensive pricing and risk management approach. Its precision is essential in ensuring that risk is neither under nor overpriced. The danger of undervaluing counterparty risk is obvious. An institution cannot begin to manage the risk it has to a counterparty if this is not correctly valued. Conversely, inflated CVA charges reduce the competitiveness of an institution, costing it business. Several of the benefits associated with accurate CVA pricing within an organisation are listed below.

- (i) **Fair Value Accounting Standards** [12], [13]: International accounting standards require financial institutions to report the fair value of their derivative positions net of credit valuation adjustments. The new accounting standard for financial instruments, IFRS 9 (which will be compulsory from 2013) includes the effect of own credit risk in the valuation of derivative positions¹.
- (ii) **Hedging of Profit and Loss Volatility**: The value of the CVA is sensitive to the underlying market risk factors driving the prices of the instruments on which it is computed as well as to the credit spreads of the counterparties to the trades. Dynamic management of CVA involves hedging both of these aspects, reducing the volatility of the portion of the daily profit and loss that is attributable to CVA. Only when it is accurately modelled can sensitivities be computed and hedges transacted.

¹Note that, in 2007, South Africa adopted International Financial Reporting Standards (IFRS) accounting rules [14]. The new standard will thus be applicable to local financial institutions.

- (iii) **Active Counterparty Risk Management:** Accurate CVA modelling facilitates the establishment of a unit responsible for dedicated management of counterparty risk. Without active CVA management, single name exposures remain unhedged, precluding the possibility of further trades with the counterparty and leaving the institution exposed to a possible default by the counterparty. It is inefficient to create counterparty risk via OTC trading and then neglect to manage it.
- (iv) **Evolution of the Institution's Risk Culture:** CVA is the amalgamation of market and credit risk. It enables the counterparty risk component of an OTC derivative value to be managed as a market risk. The use of CVA will enable institutions to develop an integrated perspective of market and credit risk with the use of consistent metrics [12]. Market risks are generally managed by calculating the sensitivity of an instrument to the underlying risk drivers of its price. When CVA is computed, the counterparty risk can be hedged in a similar manner. It is no longer a number that is computed by a credit risk manager purely for information purposes. Effectively, credit risk becomes a market risk and is managed as such.
- (v) **Reduction in Capital Consumption:** The Basel III framework has a specific focus on improving counterparty risk management. In addition to the default risk capital charge that was required under Basel II, there is a stipulation that institutions hold capital on CVA. The latter is intended to cover the risk of mark to market losses on the expected counterparty risk [15]. Organisations that do not obtain regulatory approval for an internal CVA measurement methodology are heavily penalised by the standardised approach. Particularly relevant is the stipulation regarding capital offset (a reduction in the amount of capital required to be reserved) from the ownership of CDSs. The CVA calculated for capital purposes may only be offset by a CDS position if the latter is used to hedge CVA [15]. In other words, if an institution happens to hold protection referencing one of its counterparties, but the protection was not specifically purchased to hedge the CVA, it cannot be used to offset CVA capital costs. The accurate modelling of CVA is essential for the establishment of a function that hedges counterparty exposure and thereby reduces capital consumption.

3.1.1 Right and Wrong Way Risk

The concepts of right and wrong way risk are integral to any discussion on CVA. In Section 2.1, the relevance of the replacement cost of a derivative upon default of the counter-

party was explained. Wrong way exposure can be understood as a derivative transaction in which there is a significant positive dependency between the counterparty's default probability and the mark to market value of the contract. Therefore, given the default of the counterparty, the replacement cost is generally higher than would otherwise be expected [16] and is likely to be positive.

The classic example utilised in introducing wrong way risk is the case of a company selling put options on its own stock. When the put options become valuable, the share price of the company has declined, implying that its default probability is higher than it was when the options were out-of-the-money. When the strike of the option is extremely low, the trade is similar to the company selling protection on itself.

Wrong way risk is present in all asset classes. Note also that the nature of the risk is dependent on the trade direction. In certain instances, the identification of the type of risk (right or wrong way) is not straightforward and may not be possible based on the information available. The examples listed below, per asset class, illustrate these points.

- (i) **Equities:** A less extreme version of the above example is a transaction in which a company sells put options on a stock, the price of which is highly correlated to its own. For example, the two entities might be in a similar sector and region. The purchaser of the put options has wrong way risk to the seller.
- (ii) **Foreign Exchange:** Eskom, the South African electricity utility, is a parastatal. It could be argued that the default risk of Eskom (as reflected in its CDS spreads) would be correlated with extreme movements in the South African rand (ZAR), which reflects the perceived economic health of the country. An investor who enters into a cross currency swap in which he pays ZAR to Eskom and receives United States Dollars (USD) from Eskom would be exposed to wrong way risk. When ZAR weakens, the USD received from Eskom is worth more in ZAR terms than when the trade was initially conducted. When Eskom is close to default, the investor is therefore likely to be in-the-money. Were the direction of the trade reversed, right way risk would instead arise from the perspective of the investor who is trading with Eskom.

Right and wrong way risk need not always be the result of counterparty risk. Suppose, for example, that a South African investor sold protection in the form of a CDS denominated in USD on a large, dual-listed South African entity such as Anglo

American Plc (AAL). In the event of an AAL default, the investor would be required to settle the amount outstanding on the CDS in USD at a time when ZAR would be likely to have weakened. The correlation that is relevant to the South African investor and gives rise to wrong way risk from his perspective is between the default time of the reference entity and the exchange rate. The purchaser of the USD protection has wrong way risk to his South African counterparty. Note that, if the direction of the trade was reversed (the South African investor sold ZAR protection and purchased USD protection), both wrong way risks would instead be right way exposures.

Both examples in this section pertain to the wrong way risk arising from the correlation between an entity's survival and the currency of the country in which the entity has its foundations. This sort of risk is especially prevalent in emerging markets.

- (iii) **Fixed Income:** An investor who enters into an interest rate swap with a heavily indebted counterparty, in which a floating rate of interest is received from the counterparty at regular intervals, is a form of wrong way risk from the investor's perspective. Unless the indebted counterparty has hedged its interest rate risk, it is likely to have large interest payments due at the same time as the swap moves out-the-money from its perspective.
- (iv) **Commodities:** An investor who is long a gold forward with AAL as the counterparty is expected to have right way exposure since AAL is a gold mining company. On the other hand, AAL may have hedged its exposure to the gold price by shorting forward contracts on it, thereby eliminating any right way risk inherent in the trade. In fact, depending on the level at which the hedges have been transacted and the extent to which the gold mining company has endeavoured to avoid fluctuations in the gold price from affecting its profits, the trade may introduce wrong way risk: when the gold price peaks, AAL may be required to meet margin calls.
- (v) **Credit:** Purchasing protection from an entity on its own sovereign is not uncommon [17] (for example purchasing a CDS referencing the South African government from a South African bank). Clearly, this is a form of wrong way exposure for the protection buyer, as is the purchase of CDSs from any institution whose credit worthiness is highly correlated with that of the underlying reference entity. On the other hand, the sale of protection to such entities is a form of right way risk. This concept will be revisited in depth later in the dissertation when the CVA on a CDS

is modelled.

Since CVA is the price of counterparty risk, it should incorporate both potential future exposure and counterparty default probability. In the event that right or wrong way risk exist, these quantities will be correlated to each other or have some dependence structure linking them. To accurately reflect the price of this risk, a CVA model should be able to capture such dependencies. Right and wrong way exposures complicate counterparty risk pricing, since the assumption of independence between the underlying derivative values and the default times of the counterparty cannot be made. This is elaborated on further in the next section.

3.2 Unilateral CVA

We now consider the one-sided CVA, whereby an investor prices only its counterparty's default risk. This is called *unilateral* CVA. It is sometimes referred to as the CVA of the investor's asset [3], since it affects him only when the derivative or portfolio of derivatives with the counterparty is positive from his viewpoint².

We place ourselves in the probability space $(\Omega, \mathcal{G}, \mathcal{G}_t, \mathbb{Q})$ where \mathcal{G}_t is the filtration modelling the flow of market information up to time t . The symbol \mathbb{E} is used to denote expectation while τ denotes the default time of the counterparty. The discount factor at time t for maturity s is denoted by $D(t, s)$. We call $\Pi^D(t)$ the payoff of a generic defaultable claim at time t with $\Pi(t)$ the value of the equivalent claim with a default-free counterparty. The abbreviation, NPV(t), is the net present value of the residual payoff from t to maturity T . Brigo and Masetti [18] derived the following for the unilateral case:

General Counterparty Risk Pricing formula

At valuation time t , and provided the counterparty has not defaulted before t , the price of the payoff under counterparty risk is

$$\mathbb{E}_t[\Pi^D(t)] = \mathbb{E}_t[\Pi(t)] - \underbrace{\text{LGD} \mathbb{E}_t[1_{\{t < \tau \leq T\}} D(t, \tau) (\text{NPV}(\tau))^+]}_{\text{CVA}} \quad (3.1)$$

where T is the maturity of underlying claim, $\text{LGD} = 1 - R$ is the loss given default of the counterparty and the recovery fraction, R , is assumed to be deterministic.

²This term can be misleading since the derivative might be an asset one day, but become a liability the next day when a cash flow has occurred for example.

For a proof of Equation (3.1), refer to Section A.1 in Appendix A. The formula provides a mathematical demonstration of CVA as being the difference between the value of a derivative with and without the consideration of counterparty risk. The following should be evident from (3.1):

- (i) The unilateral CVA can never attain a negative value. This is due to the fact that a loss (and not a profit) can only ever be made on the counterparty's default. This loss is floored at zero (since a negative loss is a profit). In fact, adjusting the value of the default-free derivative by the CVA is similar to entering into a short call option position, with a strike of zero, on the residual NPV of the derivative at τ (the default time of the counterparty)³. In entering into an OTC derivative transaction, the counterparty is implicitly being sold the option to default if the derivative is out-of-the-money from his perspective. The CVA is the premium of this option that is granted to the counterparty. It knocks in upon a counterparty default event⁴ and has a positive payoff only if the counterparty defaults prior to the maturity date of the underlying trade.
- (ii) The value of the CVA is equivalent to the price of a derivative that pays out the portion of the underlying claim's mark to market that is not recovered in the event of a counterparty default ($\text{LGD} \times \text{NPV}(\tau)$) in the event that the counterparty defaults when he is out-of-the-money. Such a derivative is called a contingent credit default swap (CCDS). It will be discussed further in Section 3.6. In valuing the unilateral CVA, we are thus pricing a CCDS on the underlying claim. The subtraction of the CVA term from the risk-free derivative value can therefore be interpreted as selling protection on the counterparty, where the notional of the protection is the mark to market value of the underlying claim. The value of our position decreases with an increase in the counterparty's default probability (a feature of a short protection position).
- (iii) Notice next that a zero CVA implies that the indicator variable $1_{\{t < \tau \leq T\}}$ is set to zero, implying that the counterparty has a null default probability. In failing to adjust the derivative value by the CVA (and therefore assuming that the CVA is

³To see this, recall that the value of a call option on say S with strike X is given by the expected value of the discounted payoff. In other words, $\mathbb{E}[D(t, T)(S(T) - X)^+]$ where T is the option maturity.

⁴Traditional knock-in options begin to function as normal options once a certain price level of the underlying has been reached.

zero), the counterparty is default-free by implication. This is almost certainly not true.

- (iv) If the default probability of the counterparty is independent of the value of the underlying claim at default, the CVA can be computed as an option on the derivative NPV multiplied by the expected value of the default probability. This is due to the expected value of the product of two independent random variables being equivalent to the product of their expected values [19]. When the default probability is assumed to be deterministic (and independent of the derivative value by implication), the term $1_{\{t < \tau \leq T\}}$ can be factored directly out of the expectation⁵. In a situation in which the assumption of independence between the derivative NPV and the counterparty default probability cannot be made, the complexity of the CVA computation is significantly increased. This is true when right and wrong way exposures are a feature of the counterparty risk.
- (v) The inclusion of the CVA in the valuation of a derivative creates a new hybrid instrument. In the case of a credit derivative underlying, a compound credit derivative is produced. A hybrid is a derivative that combines features of more than one asset class. For example, a vanilla interest rate swap becomes a hybrid instrument when the credit risk of the counterparty is incorporated into its pricing. The two asset classes to which it now belongs are fixed income, as before, and credit derivatives, due to the credit contingent component brought about by the presence of the CVA term. A compound credit derivative is a credit derivative on another credit derivative. The CVA on a CDS is an example of a compound credit derivative.

3.3 The Bilateral Case

Ignoring bid-offer spreads, for two parties to agree on the fair price of an OTC derivative, the value of the default-free contract should be adjusted by both parties' default risk. This is based on the principle that the total CVA in the system is a zero sum game [20]. If entities A and B enter into a derivative transaction, the counterparty CVA that A sees for B is the value of B's own risk.

Bilateral CVA (BCVA) requires a reduction in the value of the CVA in line with the investor's own default risk [21]. This may result in a negative CVA. Its sign depends on:

⁵Since it is no longer stochastic, the value is known.

- (i) The relative riskiness of the investor and the counterparty.
- (ii) The current and potential future values of the derivative whose CVA is being computed, including the effects of right and wrong way risks.

The portion of the CVA attributable to the investor's own risk of default is called the *debt valuation adjustment* or DVA. Note that the part of the BCVA arising from the counterparty risk is not the same as the original unilateral CVA, since the investor's survival is no longer certain in the BCVA valuation. The latter can lose money on the counterparty default only if his own default does not precede that of the counterparty. In this dissertation, the portion of the BCVA attributable to the counterparty's default risk is denoted by the abbreviation CCVA (counterparty CVA). This prevents confusion between CCVA and the unilateral CVA from arising.

The bilateral version of (3.1), the General Counterparty Risk Pricing Formula, can be found in [22]. We state it below, with the proof presented in Section A.2 in Appendix A. The probability space is the same as in the previous section. The investor, from whose perspective the BCVA is computed, is *Name 0*. The counterparty is *Name 2* (*Name 1* is reserved for later use and will refer to the reference entity of a CDS). The investor's default time is thus denoted by τ_0 and the counterparty's default time by τ_2 . Given that T is the maturity of the derivative under consideration, we define the following events ordering the default times⁶:

1. $A = \{\tau_0 \leq \tau_2 \leq T\}$ (The investor defaults before the counterparty and both default prior to the contract maturity.)
2. $B = \{\tau_0 \leq T \leq \tau_2\}$ (The investor defaults before expiry of the contract while the counterparty survives beyond the maturity date.)
3. $C = \{\tau_2 \leq \tau_0 \leq T\}$ (The counterparty defaults before the investor and both default prior to the contract's expiration date.)
4. $D = \{\tau_2 \leq T \leq \tau_0\}$ (The counterparty defaults before the contract's expiration date, while the reference entity survives beyond the maturity date.)
5. $E = \{T \leq \tau_0 \leq \tau_2\}$ (Both entities survive beyond the contract's maturity date and the investor defaults first.)

⁶These events are both mutually exclusive (they cannot occur simultaneously) and exhaustive (they cover all possibilities).

6. $F = \{T \leq \tau_2 \leq \tau_0\}$ (Both entities survive beyond the contract's maturity date and the counterparty defaults first.)

In addition, the stopping time (τ) is defined as:

$$\tau := \min\{\tau_0, \tau_2\}. \quad (3.2)$$

We have now introduced the notation necessary to present the bilateral version of (3.1).

General Bilateral Counterparty Risk Pricing Formula

At valuation time t , and conditional on the event $\{\tau > t\}$, the price of the payoff under bilateral counterparty risk is:

$$\begin{aligned} \mathbb{E}_t[\Pi^D(t)] = & \mathbb{E}_t[\Pi(t)] \\ & + \underbrace{\text{LGD}_0 \cdot \mathbb{E}_t\{1_{A \cup B} \cdot D(t, \tau_0) \cdot [-\text{NPV}(\tau_0)]^+\}}_{\text{DVA}} \\ & - \underbrace{\text{LGD}_2 \cdot \mathbb{E}_t\{1_{C \cup D} \cdot D(t, \tau_2) \cdot [\text{NPV}(\tau_2)]^+\}}_{\text{CCVA}}, \end{aligned} \quad (3.3)$$

where $\text{LGD}_2 = 1 - R_2$ is the loss given default of the counterparty and the counterparty's recovery fraction R_2 is assumed to be deterministic. Similarly, $\text{LGD}_0 = 1 - R_0$ denotes the loss given default of the investor.

The positivity of the second term on the right hand side of (3.3), the DVA portion of the formula, is contingent on the union of events A and B ; that is $A \cup B$. It represents the event that the investor is the first of the two entities to default and that his default time occurs prior to the maturity date of the underlying contract. The positivity of the last term, the CCVA portion of (3.3), is contingent upon the union of events C and D . This union represents the event that the counterparty is both first to default and that his default time occurs prior to the maturity date of the contract. If the counterparty defaults after settlement of the contract, there will be no loss to the investor.

It is intuitive that the value (to the investor) of the risk-free derivative, is increased by his own default risk probability and reduced by the risk of a counterparty default. This is due to the investor's profit being capped at zero in the event of a counterparty default event, but floored at zero in case of his own default. The following observations can be made regarding (3.3):

- (i) The addition of the DVA term in (3.3) is similar to a long put option position, with a strike of zero, on the residual NPV of the underlying derivative contract at the time of the investor's default⁷. This term provides a nonzero contribution only when the investor defaults prior to both the counterparty and the maturity date of the underlying derivative. The DVA can thus be viewed as the premium the investor is required to pay for the option he has received to default in the event that the derivative contract is out-of-the-money from his perspective. The adjustment to the fair value of a derivative due to BCVA is thus the sum of the premium on a long put option (that knocks in when the investor defaults, provided that this is prior to the counterparty defaulting and the maturity of the underlying contract) and the premium on a short call option (that knocks in when the counterparty defaults, provided that this is prior to the investor defaulting and the maturity of the underlying contract).
- (ii) The BCVA is the difference between the CCVA and the DVA, both of which are positive. Therefore, the BCVA could be either positive or negative. Its sign depends on whether the DVA or the CCVA is most valuable. This, in turn, is a function of the relative credit quality of the two entities and the current and expected future values of the underlying claim.
- (iii) The DVA is equivalent to the value of a derivative (a contingent credit default swap) that pays out the portion of the derivative's NPV that is not recovered by the counterparty upon default of the investor (assuming that the counterparty has survived to this point). The payoff of the derivative is $\text{LGD}_0 \times \text{NPV}(\tau_0)$ in the event that the investor defaults when the counterparty is in-the-money. The investor has effectively purchased protection on himself. The value of his assets rises with an increase in his default probability. This is evident from the term $1_{A \cup B}$, which is larger the more likely the investor is to default. As an alternative explanation of this idea, consider the following two scenarios:
- (a) If the investor were to close out the contract when he was out-of-the-money, it would cost him the NPV of the contract.
 - (b) If he defaults, it costs him only $R_0 \times \text{NPV}(\tau_0)$. He thus makes a profit of $\text{LGD}_0 \times \text{NPV}(\tau_0)$ on his own default.

⁷To see this, recall that the value of a put option on a share (S) with strike X is given by the discounted expected value of the payoff. In other words, the discounted value of $\mathbb{E}[X - S]^+$ [23].

Effectively, he has bought protection on himself. The notion of the protection is linked to the NPV of the derivative whose CVA is being computed.

- (iv) Notice lastly that a zero DVA value implies that the indicator variable, 1_{AUB} , is zero. The investor thus has a null default probability for the life of the trade. This is, of course, impossible. Failing to adjust the derivative value by the own risk portion of the BCVA is therefore equivalent to assuming that the investor is risk-free, which is economically incorrect.

When DVA was first introduced into the CVA calculation, it triggered much debate on the moral hazard of attaching an economic value to the gain made by the investor by defaulting when the derivative is out-of-the-money from his perspective. As the concept has become entrenched in CVA computation, institutions can no longer afford to ignore their own default risk since, not only are derivative values required to be adjusted by the BCVA under IFRS 9, but failing to reduce the CVA by the DVA portion reduces an organisation's competitiveness. If its competitors are reducing the counterparty risk charged to clients by their own risk amount (the DVA), it will lose business by charging clients the full unilateral CVA.

An obvious advantage of pricing an asset that increases in value as the default probability of the investor becomes more likely, is that it dampens the impact of credit spread increases by offsetting the associated increases in required reserves [21]. Conversely, the reduction in funding costs when the entity's spreads narrow will be offset by a decrease in the value of its assets, as reflected in the DVA. The own risk portion of the BCVA can thus be seen as the expected cost, to the investor, of margining the trade. It is not simply a profit that he makes as a result of his spreads rising.

3.3.1 Problems with Hedging DVA

The default component of the unilateral CVA is often hedged by purchasing protection, generally in the form of CDSs referencing the counterparty whose CVA is being hedged. However, the adoption of a similar strategy for hedging DVA would result in an investor selling protection on itself. This is not only illegal in most jurisdictions, but it would be next to impossible to obtain a protection buyer who would purchase protection on an entity from the same entity.

Permitting the DVA to remain unhedged results in large profit swings during turbulent

times. During the 2008 credit crisis the daily and weekly profit and loss of the CVA units in the top tier banks (that already had BCVA systems in place) was largely a function of the change in the institution's own spread from one day or week to the next. This is due to the fact that spreads on financial entities were highly volatile at the time. This fluctuation in profit and loss was due to unhedged DVA positions.

An illustration of the effect of unhedged DVA on the earnings of a bank are the 2010 third quarter results of Morgan Stanley. The bank suffered a loss of 91m (91 million) USD, a significant portion of which was attributable to its DVA, reflecting tightening credit spreads. The Financial Times reported that, if the positive BCVA value of 731m USD in the second quarter and the negative BCVA of 660m USD in the third quarter were stripped out of the earnings, the sequential decline would be 20% compared to the reported 50% [24].

A few (unsatisfactory) ways in which the DVA gain, which is made when an institution becomes riskier, could be monetised are listed below:

1. The institution could sell CDS protection on names whose credit spreads are highly correlated to its own [21]. This leaves the idiosyncratic component of the spread, which can make up a large portion of its volatility, unhedged. Additionally, further credit risk would be created by the introduction of this idiosyncratic portion of the spreads of the reference entity underlying the hedge.
2. If a position is unwound, the DVA would be realised. This was the case during the credit crisis when trades with some monoline insurers were unwound as their credit quality decreased [21]. This is likely only once the credit quality of the institution has declined significantly and is thus not a desirable means of recognising the gain.
3. When the institution defaults, it realises the gain [21]. Clearly, this is not a viable means of generating a profit.
4. The institution could purchase its own bonds in the market at a lower price. However, it would still be likely to require funding and would simply have to issue more debt at higher spreads.

Bilateral counterparty risk has become the standard. However, the means of hedging the DVA portion and monetising the value of the asset an institution obtains by synthetically buying protection on itself remains a question. As explained above, pricing a derivative

in the absence of DVA implicitly assumes that the investor is risk-free, which is almost certainly not true. BCVA measurement is necessary for agreement on the fair price between parties. If an investor increases the price charged to the counterparty of a derivative transaction due to the risk of the counterparty defaulting, it is logical for the counterparty to reduce the amount paid on account of the investor's riskiness. The measurement of DVA reflects the economic reality of the investor's default risk.

3.4 A Simple Example

We have now introduced the concept of CVA and discussed both the unilateral and bilateral adjustments in some detail. Before proceeding with the remainder of the chapter, we present a simple example in this section to illustrate CVA practically. We compute both the unilateral and bilateral CVA on a vanilla interest rate swap and demonstrate the effect of these on the fair value of the swap. The adjustment is calculated for the most basic case in which the underlying derivative value and the survival probabilities of both the investor and the counterparty are mutually independent.

3.4.1 A Formula for Unilateral CVA with Independence

Consider the expression for the unilateral CVA in (3.1). Under the assumption of independence between the default probability of the counterparty and the underlying derivative value, the CVA formula can be simplified significantly as follows [25]:

$$\begin{aligned}
 \text{CVA}(t) &= \text{LGD} \mathbb{E}_t[1_{\{t < \tau \leq T\}} D(t, \tau) (NPV(\tau))^+] \\
 &= \text{LGD} \mathbb{E}_t\left[\int_t^T D(t, s) (NPV(s))^+ 1_{\{\tau \in [s, s+ds)\}}\right] \\
 &= \text{LGD} \int_t^T \mathbb{E}_t[D(t, s) (NPV(s))^+ 1_{\{\tau \in [s, s+ds)\}}] \quad (3.4)
 \end{aligned}$$

$$= \text{LGD} \int_t^T \mathbb{E}_t[D(t, s) (NPV(s))^+] \mathbb{E}_t[1_{\{\tau \in [s, s+ds)\}}] \quad (3.5)$$

$$= \text{LGD} \int_t^T \mathbb{E}_t[D(t, s) (NPV(s))^+] \mathbb{Q}_t\{\tau \in [s, s + ds)\} \quad (3.6)$$

$$= \text{LGD} \int_t^T \mathbb{E}_t[D(t, s) (NPV(s))^+] d_s \mathbb{Q}_t\{\tau \leq s\}. \quad (3.7)$$

The notation utilised in the above derivation has all been defined previously, with the exception of T , which is the maturity of the underlying derivative. Note also that $\text{CVA}(t)$

is the value of the CVA at time t . Fubini's theorem (see [26]) was used in switching the time integral with the expectation in (3.4). The transition from (3.4) to (3.5) is permissible due to the assumption of independence between the default time (τ) and the value of the underlying derivative NPV⁸. Finally, (3.6) follows from (3.5) since, for a set A , $\mathbb{E}[1_{\{A\}}(X)] = P(X \in A)$ [27].

Note that here \mathbb{Q}_t denotes the expectation conditional on $\tau > t$. This is perfectly logical since we would not be calculating the CVA value on a trade at time t if the counterparty to the trade had defaulted prior to t .

3.4.2 Application to a Vanilla Interest Rate Swap

A vanilla interest rate swap (IRS) is a derivative in which one party pays a fixed rate of interest on a certain notional amount at regular intervals, usually quarterly, in exchange for a floating rate of interest on the same notional amount [23]. We denote the times at which the fixed for floating payments are scheduled to take place as $T_{a+1}, T_{a+2}, \dots, T_b$ for a swap that starts at time T_a ⁹. Suppose that we are calculating the CVA on an IRS in which the investor, from whose perspective the CVA computation is performed, is long the swap or paying the fixed rate. Figure 3.1 below illustrates the mechanics of a vanilla interest rate swap.




Figure 3.1: Mechanics of a vanilla interest rate swap (IRS). *The direction of the trade is the same as in the CVA computation in this section.*

Consider now the simplified unilateral CVA formula in (3.7). Notice that the term $\mathbb{E}_t[D(t, \tau)(NPV(\tau))^+]$ is merely the value of an option on the NPV of the underlying trade at time t with expiry τ and strike zero. This follows from the fact that the time t value of a generic European option is the expected value of the discounted payoff of the

⁸The expected value of independent random variables is the product of their expectations [19].

⁹Exchanges take place at the end of the payment period.

instrument [23].

A European payer swaption is an option to enter into a long interest rate swap position, upon expiry of the swaption, where the swap fixed rate is the strike of the swaption. In the example in this section, in which the CVA on a payer IRS is being computed, the expectation in (3.7) may therefore be replaced by the expression for the payer swaption price [25] as follows:

$$\text{CVA}(t) = \text{LGD} \int_{T_a}^{T_b} \text{SwaptionP}(t; s, T_b, K, S(t; s, T_b), \sigma_{s, T_b}) d_s \mathbb{Q}_t \{\tau \leq s\}. \quad (3.8)$$

In (3.8) above, $\text{SwaptionP}(t; s, T_b, K, S(t; s, T_b), \sigma_{s, T_b})$ is the price, at time t , of a payer swaption (since the underlying IRS is a payer IRS) with maturity s , strike K , underlying forward swap rate $S(t; s, T_b)$, volatility σ_{s, T_b} and maturity T_b . The forward swap rate at a given time t , in a default-free market, is the one which renders a zero-valued swap at time t .

3.4.3 A Formula for Implementation

In order to utilise the expression for the unilateral CVA on a swap in (3.8), the integral is discretised. Under the assumption that default can occur only at points T_i in the swap payment schedule, the unilateral CVA on a payer swap is computed using the following sum:

$$\begin{aligned} \text{CVA}(t) &= \text{LGD} \sum_{i=a+1}^{b-1} \mathbb{Q}\{\tau \in (T_{i-1}, T_i]\} \text{SwaptionP}(t; K, S_{i,b}(t), \sigma_{i,b}) \\ &= \text{LGD} \sum_{i=a+1}^{b-1} (\mathbb{Q}(\tau > T_{i-1}) - \mathbb{Q}(\tau > T_i)) \text{SwaptionP}(t; K, S_{i,b}(t), \sigma_{i,b}). \end{aligned} \quad (3.9)$$

The expression in (3.9), (which can be found in [25]) is intuitively straightforward. Recall that, in this simple illustration, we have assumed independence between the NPV of the swap and the default time of the counterparty. At time τ , the residual NPV of the IRS is the NPV of a forward starting IRS which begins at τ and matures at time T . The IRS underlying the swaption is thus one that starts at the counterparty default time and has a fixed rate equivalent to the fixed rate of the swap whose CVA we are computing. The credit valuation adjustment is then the sum of the swaption values whose maturities cover the range of potential default times. The independence assumption enables the weights

of the sum to be the probability of defaulting around each of these possible maturities. The unilateral CVA on a payer swap may now be computed using (3.9).

3.4.4 The Bilateral CVA on a Swap

The derivation of the unilateral CVA has been shown in some depth. Since the purpose of this section is merely to present a practical numerical example, the formula for the BCVA on an IRS will be derived as an extension of the unilateral one.

Consider firstly the term in the BCVA formula in (3.3) labeled CCVA. The only difference between this and the expression for the unilateral CVA in (3.1) is the set of values for which the indicator variable assumes a nonzero value (since we have assumed independence between the default time and the swap NPV). The reason for this difference is that the investor is no longer assumed to be risk-free. The CCVA is thus conditional on the investor's survival to the default time of the counterparty (the investor cannot lose money on a counterparty default if the investor defaulted prior to the counterparty).

Recall, from elementary statistics, that for two independent events A and B with probabilities $P(A)$ and $P(B)$ of occurring, the intersection of the events (the probability of both A and B occurring) is the product of their probabilities, $P(A) * P(B)$ [19]. Since the default times of the investor and the counterparty are assumed to be mutually independent, the probability of the counterparty defaulting between times T_{i-1} and T_i is multiplied by the probability of the investor surviving to time T_i . The CCVA portion of the BCVA on the IRS is thus computed using:

$$CCVA(t) = LGD_2 \sum_{i=a+1}^{b-1} \mathbb{Q}(\tau_0 > T_i) (\mathbb{Q}(\tau_2 > T_{i-1}) - \mathbb{Q}(\tau_2 > T_i)) \text{SwaptionP}(t; K, S_{i,b}(t), \sigma_{i,b}). \quad (3.10)$$

Equation (3.10) is almost identical to (3.9), with the additional probability $\mathbb{Q}(\tau_0 > T_i)$, the replacement of τ by τ_2 and of LGD by LGD_2 . Since we are now working with the bilateral case, the default time requires a subscript to indicate to which of the two entities it refers. This notation was introduced in Section 3.3 above.

Turning now to the DVA term in (3.3), the portion involving the swap NPV is $[-NPV(\tau_0)]^+$. Instead of a payer swaption (a call option on the long swap position), we now have a receiver swaption (a put option on the long swap position). Utilising the fact that the swap

NPV and the default times of the two entities are mutually independent, the DVA on the IRS may be computed using:

$$\text{DVA}(t) = \text{LGD}_0 \sum_{i=a+1}^{b-1} \mathbb{Q}(\tau_2 > T_i)(\mathbb{Q}(\tau_0 > T_{i-1}) - \mathbb{Q}(\tau_0 > T_i)) \text{SwaptionR}(t; K, S_{i,b}(t), \sigma_{i,b}), \quad (3.11)$$

where $\text{SwaptionR}(t; K, S_{i,b}(t), \sigma_{i,b})$ is the price, at time t , of a European receiver swaption with maturity s , strike K , underlying forward swap rate $S(t; s, T_b)$, volatility σ_{s, T_b} and maturity T_b .

Combining (3.10) and (3.11), the formula for the computation of the BCVA on a payer interest rate swap, under the assumption of mutual independence between the investor and counterparty default times and the value of the IRS, is:

$$\begin{aligned} \text{BCVA}(t) = & \underbrace{\text{LGD}_2 \sum_{i=a+1}^{b-1} \mathbb{Q}(\tau_0 > T_i)(\mathbb{Q}(\tau_2 > T_{i-1}) - \mathbb{Q}(\tau_2 > T_i)) \text{SwaptionP}(t; K, S_{i,b}(t), \sigma_{i,b})}_{\text{CCVA}} \\ & - \underbrace{\text{LGD}_0 \sum_{i=a+1}^{b-1} \mathbb{Q}(\tau_2 > T_i)(\mathbb{Q}(\tau_0 > T_{i-1}) - \mathbb{Q}(\tau_0 > T_i)) \text{SwaptionR}(t; K, S_{i,b}(t), \sigma_{i,b})}_{\text{DVA}}. \end{aligned} \quad (3.12)$$

3.4.5 Numerical Results

Formulae for the computation of both the unilateral CVA and the BCVA on a vanilla interest rate swap have now been derived under the independence assumptions listed above. The unilateral CVA can be computed using (3.9). The BCVA can be computed using (3.12). We have now arrived at the main objective of this section of the chapter. It is to provide a set of numerical results for the CVA on a simple derivative instrument, an IRS, and to demonstrate the effect of this on the fair value of the instrument.

In Table 3.1 below, we show the CVA for five separate scenarios. These are as follows:

- (i) **Case 1:** The investor and the counterparty are equally risky and the underlying interest rate swap is out-of-the-money (has a negative present value) from the investor's perspective.
- (ii) **Case 2:** The investor and the counterparty are equally risky and the underlying

interest rate swap is at-the-money (has a present value of zero).

- (iii) **Case 3:** The investor and the counterparty are equally risky and the underlying interest rate swap is in-the-money from the investor's perspective.
- (iv) **Case 4:** The investor is riskier than the counterparty and the underlying interest rate swap is at-the-money.
- (v) **Case 5:** The counterparty is riskier than the investor and the underlying interest rate swap is at-the-money.

Recall that the value of the BCVA depends on the current and future values of both the instrument on which it is being calculated and the survival probabilities of the investor and the counterparty. The first three scenarios demonstrate the effect of the underlying swap contract value on the CVA and on the fair IRS value, while the latter two illustrate the effect of the relative riskiness of the parties to the trade.

Scenario	IRS Value	Unilateral CVA	BCVA Numbers			
			CCVA	DVA	BCVA	IRS Fair Value
Case 1	-17k	118	101	1.2k	-1.1k	-16k
Case 2	0	452	398	393	5	-5
Case 3	15k	1.2k	1k	84	1k	14k
Case 4	0	452	391	444	-53	53
Case 5	0	510	450	386	64	-64

Table 3.1: CVA in ZAR for a long two year IRS position with a notional of ZAR1m.

With reference to Table 3.1, consider initially the first three rows. In Case 1, the IRS has a negative value. This results in a DVA that is of much greater significance than the CCVA. The logic behind this is that the default of the counterparty is not of great concern to the investor as the swap is well out-of-the-money. As explained earlier, it is only when the replacement cost is positive to the investor on default of the counterparty that a loss is made due to the counterparty's inability to honour the contract. The probability of this occurring in Case 1 is not very high. As a result, the fair value of the swap is less negative under counterparty risk, since the BCVA is positive. A default by the investor is of greater concern to the counterparty than a default by the counterparty is to the investor.

The situation in Case 1 is reversed in Case 3 when the swap has a large positive value. The risk of the counterparty defaulting is now of far greater significance than that of the

investor defaulting. As a result, the BCVA is now positive, which reduces the fair value of the swap from ZAR15k (ZAR15,000), before the inclusion of CVA, to ZAR14k when the CVA is subtracted from the risk-free value. In Case 2, the swap's fair value is close to zero since it is at-the-money and the investor and the counterparty are equally risky.

The objective behind the scenarios in Cases 1 to 3 was to illustrate the effect of the underlying NPV on the CVA. A common misconception regarding BCVA is that the sign of the adjustment is determined purely by the relative riskiness of the investor and the counterparty. The first three rows of Table 3.1 disprove this fallacy. The situation in which the investor and the counterparty are equally risky will be revisited in Chapter 11 when the BCVA on a CDS is reported.

In the first three scenarios in Table 3.1 both parties were assumed to have a flat CDS spread of 650 basis points (bp)¹⁰ for the sake of simplicity. In Case 4, when the investor is riskier, his spread is assumed to be 100bp higher than that of the counterparty (750bp). In Case 5, the investor's spread is again 650bp and the counterparty's spread is 750bp. The increase in the DVA in Case 4, and in the CCVA in Case 5, reflect the increase in the default probability of the investor and the counterparty in these scenarios respectively. Notice that, whereas the BCVA is close to zero in Case 2 when the IRS is at-the-money and the entities are equally risky, it is clearly negative in Case 4 (resulting in a positive fair trade value) and positive in Case 5 (implying a negative fair trade value).

The graph in Figure 3.2 illustrates the payer swaption values and the default probabilities in the sum in (3.9) for the unilateral CVA in Case 2. The swaption values tend to increase slightly before decreasing. This is a similar phenomenon to the effect we observed in the PFE profile of the swap in Figure 2.1(a) in Section 2.2. The jaggedness of the graph is due to the amount at risk in the event of a counterparty default decreasing when swap payments are exchanged. Note that the assumption made was that swap payments are excluded from the swaption valuation only the day following a date on the swap payment schedule. This is due to the fact that the payment would not occur if the counterparty were to default on the actual cash flow date. In swap valuations that are not linked to counterparty risk computations, the payment would be excluded from the swap price from the day on which it is scheduled to occur. The swaption valuations were completed using a Black-Scholes model with market-implied volatilities of 15-17%.

¹⁰The concept of a CDS term structure, as well as the bootstrapping thereof, will be introduced officially in the next chapter and elaborated on in Section B.1 in Appendix B.

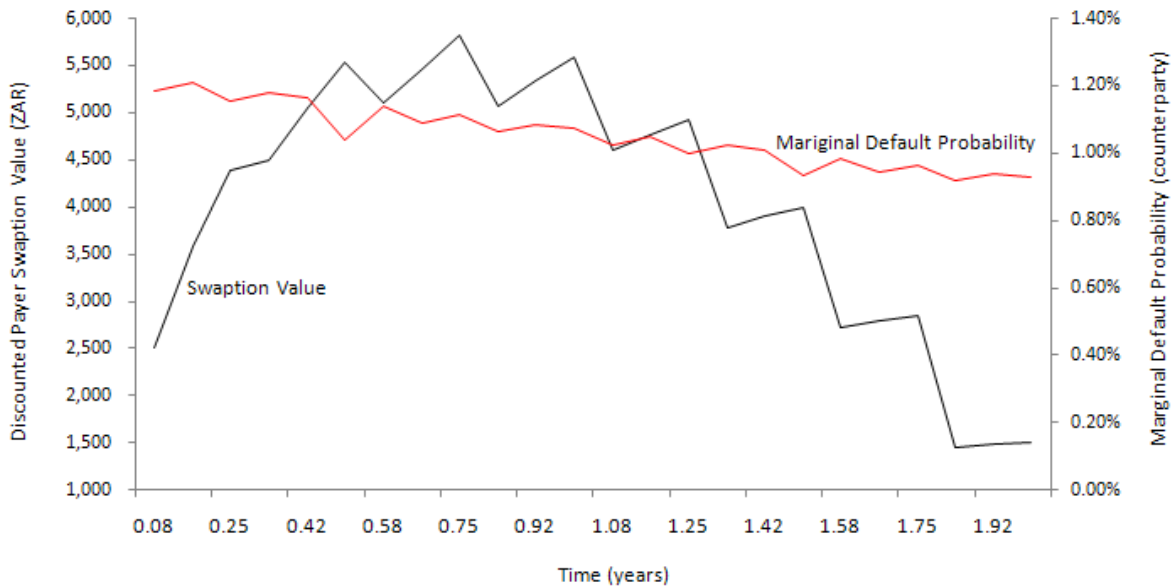


Figure 3.2: Swaption values and marginal default probabilities used in computing the unilateral CVA in Case 2. These are the two factors in the sum in (3.9).

The simplified example presented in this section was intended to provide some practical insight into the credit valuation adjustment that has been introduced and discussed in some detail in this chapter. It also demonstrates the application of the General Counterparty Risk Pricing Formulae (Equations (3.1) and (3.3)), albeit in the simplified case of mutual independence between the underlying NPV and the survival probabilities of the two entities.

3.5 Practical Implementation Aspects

An introduction to CVA would be incomplete without mention of the practical aspects related to its institution-wide implementation. The establishment of a CVA function, whose role it is to manage and hedge counterparty risk, was mentioned briefly earlier in the chapter. This will be explained in greater detail below, along with mention of the challenges related to CVA measurement systems.

3.5.1 Central Management of CVA

Much of the discussion above has centred around the CVA on a particular derivative position. In practice, however, it is more logical for CVA to be centrally managed. A number of offshore banks and an increasing number of local banks have set up CVA desks whose responsibility it is to manage the risk at a counterparty level. Each time an OTC trade is conducted, the particular desk that completes the transaction pays a premium¹¹ to the CVA desk in exchange for receiving the NPV, or the portion of the NPV that is not recovered, in the event of a counterparty default. The CVA desk thus acts as a provider of default insurance to other trading units within the institution. This is illustrated in Figure 3.3. Note that the fluctuation in profit and loss due to the mark to market of the CVA resides with the CVA desk as does the market risk management of the CVA.

Although it is not uncommon for the CVA traders to be allowed to initiate business in the form of credit-linked structures, their main mandate is generally to manage risk already created within the business. Unfortunately, the inter-desk charging can be highly political within an organisation [20].

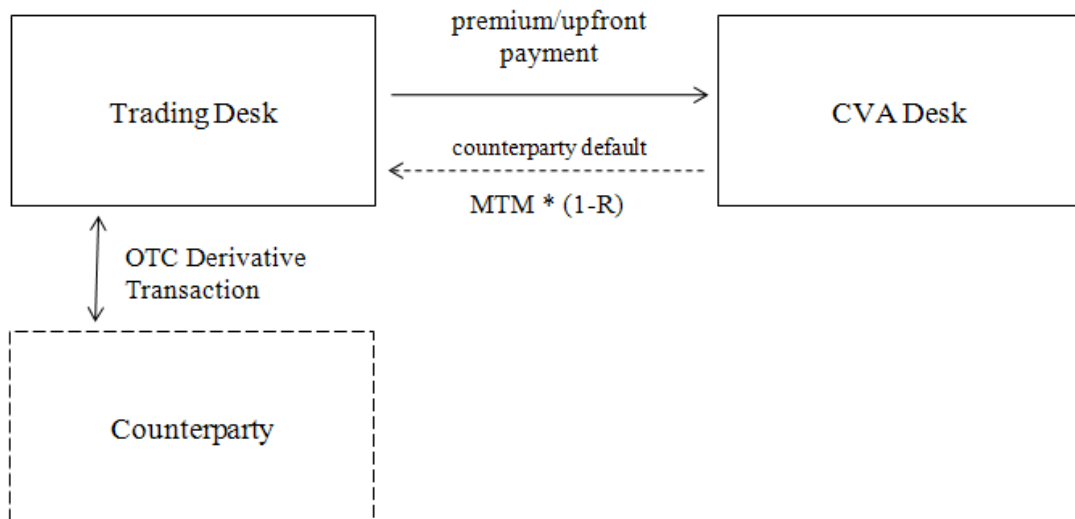


Figure 3.3: Central management of CVA

There are a number of reasons for choosing to manage CVA centrally rather than have

¹¹This could be an upfront premium or could be payable at regular intervals.

each trading unit hedging its own counterparty risk. These include:

- (i) **Counterparty netting:** The total amount owed by (or to) a counterparty will depend on the netting agreements in place. To take advantage of these and hedge only the net exposure to a particular name, the risk has to be managed centrally. For example, a Foreign Exchange (FX) trading unit may have a cross currency swap with Counterparty A, that is out-of-the-money to the FX desk. The Commodities trading unit might have entered into a gold forward with the same counterparty that is in-the-money to the Commodities desk. Combining the exposures on each of these, assuming an ISDA agreement is in place, will reduce the net amount of default risk required to be hedged and therefore reduce hedging costs.

In addition, the marginal contribution of a new trade might have the effect of decreasing the overall CVA to the counterparty rather than increasing it. These cost reductions could be passed onto clients, increasing the competitiveness of an institution. Conversely, the pre-deal check of the net exposure to a particular entity might prevent transactions that would result in too large an exposure to that counterparty, when combined with the current portfolio.

- (ii) **Cross counterparty netting:** In hedging the CVA, it is sometimes necessary to hedge the sensitivity to a particular risk driver, such as an equity price. Suppose a structuring desk is long a forward on equity A while the Equities desk is short the same forward. The sensitivity of the net CVA to the equity price will be different from the sensitivities of each of the individual adjustments. A central CVA desk would have the ability to obtain a holistic view of the sensitivities required to be hedged.
- (iii) **Centre of Expertise:** As explained above, simple instruments belonging to a particular asset class become hybrid instruments when the CVA is included in the fair value. A Fixed Income trader, who has been hired to trade linear interest rate swaps, might be best left to his particular speciality. In addition, the optionality inherent in the CVA can complicate risk management of these products, leading to second order sensitivities¹² on previously straightforward, linear products. These second order hybrid risks may also take the form of cross gamma sensitivities. For example, in the case of the interest rate swap, the rate of change of the PV01

¹²The second derivative of the CVA value with respect to a particular risk driver may exist on instruments that previously only had non zero first order sensitivities.

(the sensitivity of the swap's value to interest rates) of the CVA with respect to credit spreads now exists¹³. Prior to the inclusion of the CVA in the swap price, credit spreads were irrelevant to the valuation and hedging of the swap. A financial institution might not wish to grant a mandate to each trading unit to trade credit-linked instruments such as CDSs and CCDSs. It is therefore logical to have a 'centre of excellence' where these complex risks can be managed by traders hired specifically for that purpose.

3.5.2 System Challenges

CVA measurement systems represent one of the most challenging aspects of implementing a CVA framework, mainly as a result of the modelling complexity [20]. The simple example in Section 3.4 assumed mutual independence between the underlying NPV and the default times of the two relevant entities. However, when right or wrong way risks enter the framework, the pricing becomes significantly more complex. Much of the modelling requires Monte Carlo simulation and cannot be reduced to closed form solutions. The system utilised for CVA is also required to capture the effects of netting and collateral agreements.

The challenge is in the speed of computation. Ideally, an institution should have the ability to enter a potential transaction as a 'what-if' trade and calculate the effect of this on the CVA of the counterparty. Thus, close to real-time computations are important. The sensitivities of the various trading portfolios to each of the relevant risk factors within the portfolio are required for risk management purposes. These include second order effects (second order sensitivities to a particular risk driver). The sensitivities are generally calculated via shifts to the affected risk factors. The revaluation of the CVA is therefore required under many different scenarios, not to mention added stress test scenarios.

Many front office systems already in place cater only for a particular asset class and are capable of obtaining only the present value of the instruments they price, as opposed to the potential future values they might assume. A CVA system is required to price the counterparty risk on instruments belonging to all asset classes, as well the hybrid and compound credit instruments related to the counterparty risk adjustment. A number of institutions attempt to extend the system currently in place for PFE calculations to be able to price the CVA. However, this is not always achievable with the required accuracy

¹³This is called credit interest rate cross gamma.

and speed of a CVA system.

3.6 Two Relevant Instruments

A discussion on CVA would not be complete without mentioning two credit contingent instruments that are often associated with the subject. The final section in this chapter is thus devoted to the introduction of contingent credit default swaps (CCDSs) and credit extinguishers. The former were mentioned in Section 3.2 above. Here we provide a formal definition and explain their use as instruments for hedging counterparty risk.

3.6.1 Contingent Credit Default Swap (CCDS)

A CCDS is a contingent CDS whose notional is contingent on or linked to the present value of an OTC derivative instrument¹⁴. It provides protection, to the party that is long the CCDS, by paying out the replacement cost of the underlying OTC derivative upon default of the reference credit [3].

The mechanics of a CCDS transaction are illustrated in Figure 3.4. Here the trade underlying the CCDS is a cross currency swap. If Bank C defaults when the NPV of the cross currency swap is positive from the perspective of Bank A, Bank B is obligated (under the CCDS contract) to pay A the portion of the NPV that is not recovered from C. In return for this protection, Bank A pays B either an upfront amount, or a regular premium until expiry of the CCDS or until Bank C defaults, whichever occurs first.

It was explained above that the valuation of the unilateral CVA could be seen as the pricing of a CCDS. It should now be clear that the payoff of the CCDS is equivalent to the amount that is lost when the counterparty defaults. That is why the unilateral CVA is equivalent to synthetically selling protection on the counterparty. When CVA is centrally managed in an organisation, the CVA desk is effectively selling CCDSs to the trading desks that face the clients.

The hedging of the counterparty default risk portion of the CVA is usually conducted by means of CDSs. Since these instruments have a fixed notional amount, the hedge

¹⁴Recall that CDSs were introduced in Section 2.4. The difference between a traditional CDS and a CCDS is that the CDS notional is a fixed amount while the CCDS notional is dependent on the value of an underlying derivative.

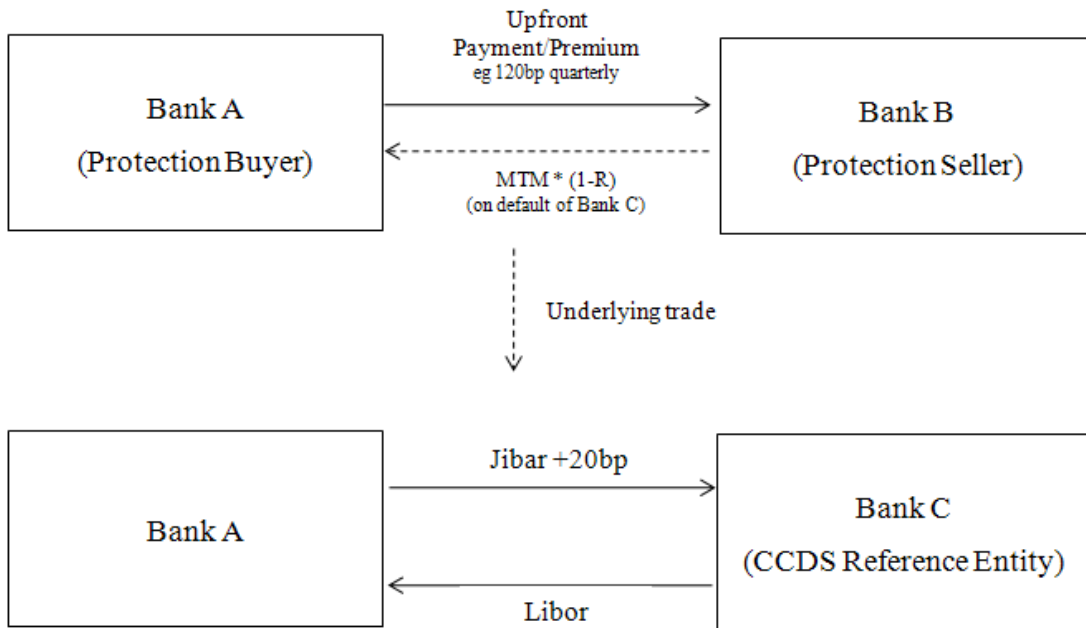


Figure 3.4: Mechanics of a contingent credit default swap with a cross currency swap as the underlying trade. *Note that the cross currency swap may be a synthetic trade.*

is dynamic and has to be adjusted frequently as the NPV of the underlying derivative changes. In contrast, a CCDS provides a perfect hedge for counterparty risk. This may be a static hedge since the CCDS purchased can have exactly the same contract underlying it as the one whose CVA is being hedged. In practice, unless a particular transaction is viewed as being extremely risky, the net CVA for a counterparty would be hedged, assuming a netting agreement exists. The hedging CCDS would then have a hypothetical underlying derivative trade. CCDSs are also sometimes used as a means of hedging the cross gamma exposure of a portfolio or combination of portfolios.

Unfortunately, these instruments are still relatively illiquid in offshore markets and trade even less frequently locally (in South Africa). Their use as hedging instruments for CVA is thus limited at present. Since the CVA is the price of a CCDS, once an institution can value CVA correctly, it can price contingent credit defaults swaps.

3.6.2 Credit Extinguisher

A credit extinguisher is an OTC derivative that has a knock-out feature upon default of the reference credit, with a pre-specified rebate (that may be zero). The typical motiva-

tion for trading an extinguisher is to achieve trade cheapening for the counterparty [3]. To understand this statement, consider Investor A who enters into an OTC trade with Counterparty B. Without the extinguisher, if B defaulted when A was out-of-the-money, A would have to pay the full outstanding NPV to B. However, in the case of an extinguisher, A would owe B nothing (except possibly a rebate). The trade is less risky to A and B would be charged less.

In addition to achieving trade cheapening for the counterparty, A is also forgoing the full amount outstanding on the derivative if A defaults during the life of the trade when it is in-the-money. Effectively, A is foregoing some of the value in the derivative and synthetically selling protection on itself [3]. This is a subtle but important point, particularly with regards to the bilateral hedging discussion in Section 3.3.1 above.

The last point we make with regards to extinguishers is that they eliminate the asymmetry inherent in counterparty risk. To understand this, consider the effect of a default by B both when the trade is extinguishable and when it is not. Assume zero recovery for the sake of convenience. In the absence of an extinguisher:

- (i) If B defaults when the trade is in-the-money to A, A makes a loss since it loses the NPV.
- (ii) If B defaults when the trade has a negative value to A, A makes neither a profit nor a loss.

Therefore, A makes a loss in the first case and is unaffected in the second. On the other hand, B makes a profit in the first and is unaffected in the second. Suppose now that the trade is extinguishable, with zero rebate, upon default of either entity.

- (i) If B defaults when the trade is in-the-money to A, A makes a loss as would have been the case in the absence of the extinguisher.
- (ii) If B defaults when the trade has a negative value to A, A makes a default profit since the negative trade value is cancelled.

The trade is thus symmetric on default. Entity A continues to lose in the scenario in which the NPV is positive. However, A now makes a profit on B's default in the event that the trade is out-of-the-money to A.

Although credit extinguishers are useful instruments in managing the risks associated

with CVA pricing, they are not legal in all jurisdictions [3]. Adding an extinguishable element to a derivative trade is legal in South Africa. However, if the underlying trade is a foreign exchange transaction, exchange control rules prevent it from being permissible. Like CCDSs, extinguishers are not very liquid.

3.7 Chapter Summary

The credit valuation adjustment is the core concept underlying this dissertation. The objective of this chapter has been to introduce the notion of CVA as well as to provide the reader with some insight into the concepts to which it is related. The importance of accurate CVA modelling was explained, followed by an introduction to the terms right and wrong way risk.

The unilateral CVA was then discussed in some depth, together with the General Counterparty Risk Pricing Formula (3.1). The alternative means of viewing the counterparty risk adjustment, namely as a short call option position or as being short protection on the counterparty, were then clarified.

A section on BCVA followed the examination of unilateral CVA. The General Bilateral Counterparty Risk Pricing Formula (3.3) was presented. The concept of own risk or DVA was then defined, along with an explanation of the DVA as a long put option position or synthetically buying protection on oneself. Several concerns surrounding the bilateral measurement of CVA, as opposed to the unilateral version, were then discussed.

In order to demonstrate the use of the General Counterparty Risk Pricing Formulae and to render the CVA concept more tangible, a simple example was presented in which the CVA on a vanilla interest swap was priced. This was followed by an explanation of the central management of CVA within an institution, together with the associated challenges presented by the need for robust systems with sufficient speed and accuracy for CVA pricing and hedging. The chapter concluded with an introduction to CCDSs and credit extinguishers.

Chapter 4

Modelling Default Risk

Broadly speaking, default risk models can be classified as either *structural (firm value)* or *reduced form (intensity)* models. This chapter presents an introduction to both of these approaches. Emphasis is placed on reduced form modelling. Not only is this the technique employed in the CDS CVA model that is implemented later in the dissertation, but it is the general market approach to the pricing of default risk. A discussion on credit modelling would be incomplete without mentioning the structural approach. This is necessary to ensure that the dissertation is reasonably self-contained. Once the basic attributes of each category have been explored, the pricing of a CDS in the reduced form framework will be explained.

4.1 Structural Models

Historically, structural techniques were the first to be applied to default risk modelling. The capital structure of a firm is used directly to model default events. Therefore it is also known as the *Firm Value* approach, as it is based on modelling the behaviour of the total value of the firm's assets [28].

4.1.1 The Merton Model

Structural models originated in 1974 with the paper by Merton [29] in which the life of a firm is linked to its ability to repay its debt. The fundamental accounting equation states that the value of a firm's assets (A) is the sum of its equity (E) and liabilities (L) or:

$$A = E + L. \tag{4.1}$$

Suppose that the firm issues a bond or takes out a loan from the bank to finance its activities. At time T (the maturity of the debt) if the firm is not able to repay the amount it owes, that is if:

$$L > A, \quad (4.2)$$

there has been a default event.

The main assumptions [25] behind the Merton model are:

- (i) The firm has a single liability with face value D .
- (ii) Default of the firm can occur only at time T , the maturity of the debt.
- (iii) The value of the assets of the firm at time t (A_t) follows geometric Brownian motion (GBM) and is therefore lognormally distributed. That is:

$$dA_t = A_t(r - d)dt + \sigma_A A_t dW_t, \quad (4.3)$$

where r is the risk-free rate of interest, d is the constant dividend yield, σ_A is the volatility of the firm's assets and W is a standard Brownian motion under the martingale measure \mathbb{Q} .

The firm will repay its creditors the face value¹ of the debt (D) if the value of its assets exceeds the value of its liabilities at maturity, T . In the event that $A_T < D$ (the debt face value exceeds the asset value at T), the creditors will receive only the value of the assets. It follows that the payoff of the debt at maturity is:

$$L(T) = \min(D, A_T) = D - \max(D - A_T, 0) = D - (D - A_T)^+. \quad (4.4)$$

Equation (4.4) expresses the value of the liabilities at maturity as the sum of a cash amount (the face value of the debt) and the payoff of a short put option position on the assets of the firm with strike D . The value of the liabilities at time t , where $t < T$, must therefore be the sum of the discounted face value and a short position in a put option on the assets of the firm with strike D and maturity T . Using (4.1), the value of the equity of the firm is:

$$E_t = A_t - D(t, T)D + \text{Put}(t, T; A_t, D) = \text{Call}(t, T; A_t, D), \quad (4.5)$$

¹The face value of the debt is the amount required to be paid upon maturity.

where $D(t, T)$ is the discount factor at time t for maturity T ; $\text{Put}(t, T; A_t, D)$ is the present value of a European put option at time t with maturity T , underlying A_t and strike D and $\text{Call}(t, T; A_t, D)$ is the value of the analogous call option. The last equality in (4.5) was obtained using put-call parity².

It is clear from (4.5) that, within the Merton structural framework, the shareholders of an entity effectively have a call option on the value of the firm's assets. Merton then priced the firm's debt using the Black-Scholes formula for European call options. Since this formula is not particularly relevant to the dissertation, the interested reader is referred to [30] where its derivation and the associated underlying assumptions are presented. Within this framework, the risk-neutral probability that the company will default on its debt is $\Phi(-d_2)$ [23], where Φ is the cumulative standard normal distribution function and

$$d_2 = \frac{\ln(A_0/D) + (r - d - \sigma_A^2/2)T}{\sigma_A \sqrt{T}}.$$

The basic structural approach to default risk modelling is illustrated in Figure 4.1. The graph shows a set of simulated asset value paths along which the asset value follows the process in (4.3). The red line in the figure represents the face value of the firm's debt. If the asset value is below this level at maturity of the debt, a default event occurs.

4.1.2 First Passage Time Models

Black and Cox [31] extended Merton's model to account for certain types of bond indenture provisions that are found in practice. These include the effects of safety covenants, subordination arrangements and restrictions on financing dividend and interest payments. Safety covenants require a firm to reimburse its bond holders as soon as the asset value (A_t) reaches a low enough 'safety level' H_t .

Whereas the Merton model involved pricing a European call option, Black and Cox priced a barrier option [25] in which the barrier was the safety level, H_t ³. This is due to the value of the debt being zero (or a recovery amount) once the asset value falls below H_t ;

²This is the relationship between the prices of European calls and puts, both with the same underlying S , strike K and maturity T . If c is the call price and p is the put price, then $c + D(t, T)K = p + S_0$ [23]. This relationship also illustrates that the value of a portfolio that is long c and short p is equivalent to being long a forward contract on S with strike K .

³Recall that a barrier option is one whose payout depends on whether the path of the underlying asset has reached a barrier or pre-determined level [23].




Figure 4.1: Illustration of the Merton model

in other words, the barrier is hit. The first time that the asset value hits the barrier, τ , is defined as:

$$\tau = \inf\{t \geq 0 : A_t \leq H_t\}. \quad (4.6)$$

The title *first passage time models* is derived from the definition in (4.6). Note that, within this framework, default can occur at any time prior to maturity of the debt (T). This is an important enhancement to the basic Merton model in which default can occur only at T [25].

Various researchers have extended the first passage time approach. Developments include: analysis of convertible bonds, a random barrier and random interest rates, study of an optimal capital structure, bankruptcy costs, tax benefits and a constant barrier combined with random interest rates [28].

4.1.3 Further Remarks on the Structural Approach

The basic assumption underlying the structural approach to default risk modelling is that the value of the firm's assets follows geometric Brownian motion. This implies that the default process can be completely monitored based on default-free market information [25]. This is a fundamental difference between structural and reduced form models, as will soon be explained.

The main drawback of traditional structural models is that they cannot be calibrated exactly to observed market prices such as CDS quotes, although Brigo and Tarenghi [32] have derived a structurally based model for CDS valuation that can be calibrated to the term structure of credit spreads. It involves the inclusion of a random barrier and volatility scenarios in the basic Black-Cox first passage time model.

The literature related to firm value models is vast and could be the subject of a dissertation in itself. The fundamental concepts associated with these types of models have been introduced in this section. The information that has been presented is sufficient for an understanding of the references to the structural approach in the literature review in Chapter 5.

4.2 Reduced Form Modelling

In contrast to the description in the previous section, in reduced form models, default is described by means of an exogenous jump process. It is not triggered by basic market observables and is independent of all default-free market information. There is therefore no economic rationale behind its occurrence [25], making it unpredictable. From a practitioner's viewpoint, reduced form models provide a flexible framework and can easily be fitted to the observed term structure of credit spreads [33]. The information presented in this section can be found in [25] unless other references are cited.

4.2.1 The Role of Poisson Processes

The default time in reduced form models is represented by the first jump time of a Poisson process. This type of process is reliant on the number of times a certain physical event has occurred. Insofar as it pertains to default risk modelling, the physical event is the default event and (since a firm can default only once) it is the first occurrence of the event

with which we are concerned.

There are three relevant varieties of Poisson processes. The first of these is time homogenous. It is the simplest of the three. The second is time inhomogeneous and is employed in the basic market model for CDS pricing. The last variety of Poisson process is the Cox process. It occurs when the parameter specifying the distribution of the number of events is both time inhomogeneous and stochastic. In order to explain the basic terms and quantities associated with reduced form models, the time homogenous Poisson process will be utilised. The results obtained are then generalised to the remaining two cases.

4.2.2 Time Homogeneous Poisson Process

Definition (see [19])

Consider a certain type of physical event. Let $M(t)$ denote the number of times such an event occurs in the interval $[0, t]$ and let $P_n(t)$ be the probability of n occurrences in the interval $[0, t]$. Suppose the following properties hold:

- (i) $M(0) = 0$ i.e. no events have occurred at time zero.
- (ii) $P[M(t+h) - M(t) = n | M(s) = m] = P[M(t+h) - M(t) = n]$ for all $0 \leq s \leq t$ and $0 < h$. This is the no memory property of Poisson processes.
- (iii) $P[M(t + \Delta t) - M(t) = 1] = \lambda \Delta t + o(\Delta t)$ (of order Δt) for some constant $\lambda > 0$.
- (iv) $P[M(t + \Delta t) - M(t) \geq 2] = o(\Delta t)$.

Then, for all $t > 0$,

$$\mathbb{Q}[M(t) = n] = \frac{e^{-\lambda t} (\lambda t)^n}{n!}, \quad n = 0, 1, 2, \dots \quad (4.7)$$

In default risk modelling, the event is the default event and $M(t) = 1$ once default has occurred.

Intensity or Hazard Rate

The density function in (4.7) is that of a Poisson random variable with parameter λt . Based on the properties of Poisson random variables [19], the expected value and variance of $M(t)$ are given by:

$$\mathbb{E}[M(t)] = \text{Var}[M(t)] = \lambda t. \quad (4.8)$$

The proportionality constant λ reflects the rate of occurrence or *intensity* of the Poisson process. If (4.8) is rewritten as:

$$\frac{\mathbb{E}[M(t)]}{t} = \frac{\text{Var}[M(t)]}{t} = \lambda, \quad (4.9)$$

it is easier to see that λ can be interpreted as the average arrival rate or the variance per unit of time. When the Poisson process is time homogeneous, the parameter λ is assumed to be constant for all t . Its interpretation remains the same in the more complex cases of time inhomogeneous and stochastic intensities.

Since the event modelled in credit risk applications is the default event, it is clear from (4.9) that the probability of an imminent default by an entity increases with the intensity parameter. For this reason, the terms *reduced form* and *intensity* models are used interchangeably. Additionally, the intensity is frequently referred to as the *hazard rate*.

Survival Probabilities

The probability of an entity with a constant hazard rate λ surviving to time t is equivalent to the number of default events being equal to zero. Making use of (4.7), the survival probability is given by:

$$\mathbb{Q}(M(t) = 0) = \frac{e^{-\lambda t} (\lambda t)^0}{0!} = e^{-\lambda t}. \quad (4.10)$$

The probability of an entity defaulting between times s and t , $s < t$, is therefore:

$$\mathbb{Q}(M(s) = 0) - \mathbb{Q}(M(t) = 0) = e^{-\lambda s} - e^{-\lambda t}. \quad (4.11)$$

Distribution of Jump Times

(An Alternative Derivation of the Survival Probability)

The jump times of the process $M(t)$, where each jump is a unit-jump, are denoted by $\tau^1, \tau^2, \tau^3, \dots$. Then the times between each jump and the subsequent one are independent and identically distributed (i.i.d.) as exponential random variables with parameter λ . That is:

$$\tau^1, \tau^2 - \tau^1, \tau^3 - \tau^2, \dots \text{i.i.d.} \sim \text{exponential}(\lambda), \quad (4.12)$$

with mean $1/\lambda$ (since if $X \sim \text{exponential}(\lambda)$, the expectation of X is $1/\lambda$ [19]). The fact that the waiting times of Poisson processes are exponentially distributed is a well-known

property of the process [25]. The distribution in (4.12) implies that $\lambda\tau^1$ is a standard, unit-mean exponential random variable or:

$$\lambda\tau^1 \sim \text{exponential}(1). \quad (4.13)$$

(Using (4.13), $\mathbb{E}(\lambda\tau^1) = 1 \rightarrow \mathbb{E}(\tau^1) = 1/\lambda$). This is the expectation that would be obtained using (4.12)).

As mentioned previously, the first jump time represents the default time in reduced form models. The probability that an entity survives to time t is thus equivalent to the probability that the first jump time of the Poisson process governing its default process is greater than t . Utilising the cumulative distribution function of an exponentially distributed random variable⁴, the probability of surviving to t is given by:

$$\mathbb{Q}(\tau^1 > t) = 1 - (1 - e^{-\lambda t}) = e^{-\lambda t}. \quad (4.14)$$

The expressions for the default probability in (4.14) (obtained by means of the cumulative distribution function of the exponential random variable) and (4.11) (derived using the the density function of the Poisson random variable) are equivalent.

Default in Small Time Interval

The probability of default (the first jump time of the Poisson process) occurring in an arbitrarily small time interval dt is considered in this section. The objective behind including it is to provide the reader with some intuition into the hazard rate and its relevance in default risk modelling. Consider:

$$P(\tau^1 \in [t, t + dt) | \tau^1 \geq t) = \frac{P(\tau^1 \in [t, t + dt) \cap \tau^1 \geq t)}{P(\tau^1 > t)} = \frac{P(\tau^1 \in [t, t + dt))}{P(\tau^1 > t)}. \quad (4.15)$$

In deriving (4.15), the definition of conditional probability (if A and B are two elements from the same probability space, $P(A|B) = \frac{P(A \cap B)}{P(B)}$) was utilised. The last inequality in (4.15) was obtained by recognising that the probability of τ^1 being simultaneously between t and $t + dt$ and larger than t , is equivalent to the probability of τ^1 being between

⁴The cumulative distribution function of an exponential random variable X , with parameter θ , is given by [19]:

$$F(X; \theta) = \mathbb{Q}(X \leq x) = 1 - e^{-x/\theta}.$$

t and $t + dt$. This is due to the fact that, in order for an entity to default between times t and $t + dt$, it must have survived to time t . Rearranging the last expression in (4.15) and using (4.14), the following is obtained:

$$\frac{P(\tau^1 > t) - P(\tau^1 > t + dt)}{P(\tau^1 > t)} = \frac{e^{-\lambda t} - e^{-\lambda(t+dt)}}{e^{-\lambda t}} \approx \lambda dt. \quad (4.16)$$

Interpreting (4.16), the probability that an entity defaults in a very small space of time dt , given that it has not yet defaulted, is approximately⁵ equal to λdt . This interpretation of the hazard rate aids in the understanding of and perception of the term *intensity* or *hazard rate*.

4.2.3 More General Poisson Processes

The previous section, on time homogeneous Poisson processes, was intended to provide the reader with the intuition behind reduced form or intensity models. The results obtained for the time homogeneous case can all be generalised to both the time inhomogeneous version (in which hazard rates are time-dependent) and to Cox processes (in which hazard rates are both time-dependent and stochastic). In this section, the survival probabilities under these two processes are derived. In pricing credit contingent derivatives, including basic CDSs, survival probabilities are the quintessential constituent.

The Time Inhomogeneous Poisson Process

The standard market model for CDS pricing makes use of a time-varying hazard rate. This enables a term structure of hazard rates to be computed from the quoted CDS spreads in the market, in contrast to the constant credit spread curve produced by the time homogeneous hazard rate. The integrated intensity is defined as:

$$\Lambda(t) := \int_0^t \lambda(s) ds. \quad (4.17)$$

Then the quantity $M(t)$, the number of jumps in the interval $[0, t]$, has a Poisson distribution with parameter $\Lambda(t)$ or:

$$M(t) \sim \text{Poisson}(\Lambda(t)). \quad (4.18)$$

⁵The approximation is due to the fact that we have used the Taylor series expansions of $e^{-\lambda t}$ and $e^{-\lambda(t+dt)}$ and assumed that dt is small enough that terms of order higher than 1 can be ignored.

Analogous to the time homogenous case, the first jump time τ^1 is exponentially distributed so that $\Lambda(\tau^1)$ has an exponential distribution with parameter 1 or:

$$\Lambda(\tau^1) := \xi \sim \text{exponential}(1). \quad (4.19)$$

The expression in (4.19) above is the time inhomogeneous analogue of (4.13). By inverting the equation in (4.19), the expression for the first jump time (the default time) is obtained as:

$$\tau^1 = \Lambda^{-1}(\xi). \quad (4.20)$$

Note that, in inverting Equation (4.19), hazard rates are assumed to be positive. This assumption is discussed in greater detail in Chapter 6.

The survival probability can now be derived (in a similar manner to the time homogeneous case) by employing the cumulative (exponential) distribution function of the first jump time. That is, the probability of surviving to time t is:

$$P(\tau^1 > t) = 1 - P(\tau^1 \leq t) = 1 - (1 - e^{-\Lambda(t)}) = e^{-\int_0^t \lambda(s) ds}. \quad (4.21)$$

Cox Processes

The intensity is now assumed to be both time-varying and stochastic. It is thus assumed to be \mathcal{F}_t -adapted and right continuous, where the filtration \mathcal{F} represents all market observable quantities with the exception of the default events. Therefore, given the default-free market information up to time t , the value of λ is known. Since the hazard rate is stochastic, credit spread volatility may be incorporated in contrast to the time homogeneous and inhomogeneous cases.

The intensity is again denoted by $\lambda(t)$ and the cumulative intensity by $\Lambda(t) = \int_0^t \lambda(s) ds$. Again, the first jump time of the Poisson process is an exponential random variable so that:

$$\Lambda(\tau^1) := \xi \sim \text{exponential}(1). \quad (4.22)$$

Inverting (4.22) leads to the following expression for the first jump time:

$$\tau^1 = \Lambda^{-1}(\xi).$$

Since λ is stochastic, in addition to ξ being random, Cox processes are frequently referred to as *doubly stochastic Poisson processes*.

The survival probability for the case of time-varying and stochastic intensity is then derived as follows [25]:

$$\begin{aligned}
 P(\tau \geq t) &= P(\Lambda(\tau) \geq \Lambda(t)) = P(\xi \geq \int_0^t \lambda(s) ds) & (4.23) \\
 &= E[P(\xi \geq \int_0^t \lambda(s) ds | \mathcal{F}_t)] = E[e^{-\int_0^t \lambda(s) ds}].
 \end{aligned}$$

This last formula is equivalent to the price of a zero coupon bond in an interest rate modelling framework, with the short rate $r(t)$ replaced by $\lambda(t)$. Therefore, the survival probability in default risk modelling is analogous to the discount factor in short rate modelling.

4.3 CDS Pricing

The two main default risk modelling methodologies have been introduced. The pricing of a credit default swap (CDS), the structure of which was introduced in Section 2.4, will be addressed in this section. This is fundamental to a background on default risk modelling, since the survival probabilities used to value credit-contingent instruments are obtained from the CDS spreads quoted in the market.

4.3.1 General Formula

The valuation of a CDS involves the pricing of two legs, the premium and protection (sometimes called default) legs. The former consists of the payments made by the protection buyer to the protection seller. The latter consists of the payoff made by the seller in the event of a reference entity default. The derivation of the CDS valuation formula in this section makes use of [25].

Consider a CDS under which the buyer of protection makes regular payments, at a rate of S , at times $T_{a+1}, T_{a+2}, \dots, T_b$ or until default of the underlying reference entity. In exchange, he receives a single payment of the loss given default (LGD) on the notional amount underlying the trade if the reference entity defaults during the life of the CDS. If the symbol τ represents the reference entity's default time, this last condition can be ex-

pressed as $T_a < \tau \leq T_b$. The discount factor at time t for maturity s is denoted by $D(t, s)$.

The value of the CDS premium leg at time t is then:

$$\begin{aligned} & \text{Premium}_{a,b}(t; S) \\ &= S \sum_{i=a+1}^b D(t, T_i) \alpha_i 1_{\{\tau \geq T_i\}} + SD(t, \tau) (\tau - T_{\beta(\tau)-1}) 1_{\{T_a < \tau < T_b\}}, \end{aligned} \quad (4.24)$$

where α_i is the year fraction $T_i - T_{i-1}$ and $T_{\beta(\tau)-1}$ is the payment date preceding τ . The indicator function $1_{\{\text{condition}\}}$ assumes a value of 1 when the condition in brackets holds and is equal to zero otherwise. The subscript a, b on the left hand side of Equation (4.24) indicates that the next premium payment date is T_{a+1} and that the CDS maturity is T_b . The symbols t and S in the brackets on the left hand side refer to the valuation date and the spread as which the trade was conducted respectively.

The first term in (4.24) is the sum of the payments at each time T_i , discounted to the valuation date t and multiplied by the relevant year fraction. The indicator term represents the probability of the reference entity surviving to T_i since the premium payment will be made only if there has not been a default of the underlying reference entity prior to the payment date.

The second term represents the accrued premium payment that occurs on default. The year fraction is the portion of the year between the default time and the last premium payment prior to default. The indicator function represents the underlying entity's probability of defaulting during the life of the trade. Unlike the full premium payment that is dependent on the reference entity having survived, the accrued premium is paid only when the reference entity defaults between payment times.

The default leg of the CDS is given by:

$$\text{Default}_{a,b}(t; S) = D(t, \tau) \text{LGD} 1_{\{T_a < \tau \leq T_b\}}. \quad (4.25)$$

This is simply the LGD on a notional of 1 discounted from the default time to the valuation date t and multiplied by the probability that the underlying reference entity defaults at a time τ during the life of the CDS contract.

The value of the CDS to the protection seller is:

$$\text{CDS}_{a,b}(t; S) = \text{Premium}_{a,b}(t; S) - \text{Default}_{a,b}(t; S). \quad (4.26)$$

The signs in front of the premium and protection legs in (4.26) are switched when the CDS is valued from the viewpoint of the protection buyer (who is long protection).

The premium leg in (4.24) can be valued by finding its expected value. Obtaining this and performing several calculations and rearrangements (which are shown in Section B.2.1 in Appendix B), the following form of the CDS premium leg is attained:

$$\begin{aligned} & \text{Premium}_{a,b}(0; S) \\ &= S \left[\sum_{i=a+1}^b D(0, T_i) \alpha_i \mathbb{Q}(\tau \geq T_i) - \int_{T_a}^{T_b} D(0, t) (t - T_{\beta t-1}) d_t \mathbb{Q}(\tau \geq t) \right], \end{aligned} \quad (4.27)$$

where interest rates are assumed to be deterministic and thus independent of default times. The probability $\mathbb{Q}(\tau \geq \cdot)$ is the risk neutral survival probability of the underlying reference entity.

Similarly, the value of the default leg in (4.25) simplifies to [25]:

$$\text{Default}_{a,b}(0; S) = -LGD \int_{T_a}^{T_b} D(0, t) d_t \mathbb{Q}(\tau \geq t). \quad (4.28)$$

The derivation of this leg is shown in Section B.2.2 in Appendix B. Note that both (4.27) and (4.28) are model independent expressions [25]. This means that the CDS price derived does not depend on the type of default risk methodology (such as structural or reduced form) that is used to derive the reference entity's survival probabilities.

A *par* CDS is one in which the value of the default leg is set equal to the value of the premium leg so that the CDS has a price of zero at inception. The *fair* or *par* spread is the value of S in (4.24) such that the price in (4.26) is zero. It is thus given by:

$$S_{\text{par}} = \frac{D(t, \tau) LGD 1_{\{T_a < \tau \leq T_b\}}}{\sum_{i=a+1}^b D(t, T_i) \alpha_i 1_{\{\tau \geq T_i\}} + D(t, \tau) (\tau - T_{\beta(\tau)-1}) 1_{\{T_a < \tau < T_b\}}}. \quad (4.29)$$

Substituting (4.27) into (4.24) and (4.28) into (4.25), the par spread is:

$$S_{\text{par}} = \frac{LGD \int_{T_a}^{T_b} D(0, t) d_t \mathbb{Q}(\tau \geq t)}{\sum_{i=a+1}^b D(0, T_i) \alpha_i \mathbb{Q}(\tau \geq T_i) - \int_{T_a}^{T_b} D(0, t) (t - T_{\beta t-1}) d_t \mathbb{Q}(\tau \geq t)}. \quad (4.30)$$

The CDS curves that are quoted in the market are comprised of the par spreads for each tenor on the curve. They can be utilised, together with risk-free rates (for discounting) and the reference entity's assumed recovery rate, to imply the survival probabilities and hazard rates of a given entity.

Recently (at the time of writing), there has been a drive to standardise CDS markets. On 8 April 2009, the CDS Big Bang Protocol [34] went into effect for North American CDSs. Among other amendments to the CDS contracts and conventions, it was agreed that all coupons would be either 100bp or 500bp going forward [34]. This necessitates the exchange of an upfront payment. The Big Bang was followed by the Small Bang [35] in June 2009. This affected European CDS markets in a similar manner to the way in which the Big Bang affected North American trades. European contracts also trade with fixed coupons as a result of the Small Bang, although there is a larger variety of coupon sizes than in North America. This is to minimize the size of the upfront payment [35]⁶. Similar conventional changes were also implemented in Japan and the rest of Asia [36].

The traditional approach, under which CDSs are traded at par, is utilised here. Bloomberg CDS curve quotes remain for par spreads. This makes little difference since the modelling is the same, with a conversion to fixed coupons applied by market participants who have signed up to the new conventions. Much of the CDS trading in South Africa is conducted with offshore counterparties, in which case the conversions may be required locally. Since it is not the focus of this dissertation, the interested reader is referred to a publication by Fitch Solutions on the topic of the conversions [37].

4.3.2 Market Model

The subject of this section is the valuation of the CDS legs in (4.27) and (4.28). As mentioned above, a reduced form model with time-dependent hazard rates is the standard model employed by market participants to imply survival probabilities from market spreads for a given entity [25, 38]. The actual procedure for bootstrapping the credit

⁶The original fixed coupons implemented for European corporates were 25bp, 100bp, 500bp and 1,000bp. Two additional coupons, 300bp and 750bp were added [35].

curve in this environment is described in Section B.1 in Appendix B. In this section, the valuation formulae that would be required for this purpose are presented.

As a first step in pricing CDSs using the market model, hazard rates are assumed to be either piecewise constant or linear [25]. Numerical integration is then performed by discretising the integral in the default leg and in the accrued portion of the premium leg. For this purpose, the premium payment times T_{a+1}, \dots, T_b are employed. Assuming we are standing at time 0, the approximated accrued premium portion of (4.27) is then given by:

$$\begin{aligned} \text{Accrued Premium}_{a,b}(0; S) &= - \int_{T_a}^{T_b} D(0, t)(t - T_{\beta t-1})d_t \mathbb{Q}(\tau \geq t) \\ &\approx \sum_{i=a+1}^b D(0, T_i)(\mathbb{Q}(\tau \geq T_{i-1}) - \mathbb{Q}(\tau \geq T_i)) \frac{\alpha_i}{2}. \end{aligned} \quad (4.31)$$

In line with the JP Morgan approximation [38], $\frac{\alpha_i}{2}$ represents the average accrual from time T_{i-1} to T_i . Substituting the result in (4.31) into the formula for the premium leg in (4.27), the total value of the premium leg at time 0 under the market model for CDS pricing is:

$$\begin{aligned} &\text{Premium}_{a,b}(0; S) \\ &\approx S \left[\sum_{i=a+1}^b D(0, T_i) \alpha_i \mathbb{Q}(\tau \geq T_i) + \sum_{i=a+1}^b D(0, T_i) (\mathbb{Q}(\tau \geq T_{i-1}) - \mathbb{Q}(\tau \geq T_i)) \frac{\alpha_i}{2} \right] \\ &= S \sum_{i=a+1}^b \left[D(0, T_i) \alpha_i e^{-\int_0^{T_i} \lambda(s) ds} + D(0, T_i) (e^{-\int_0^{T_{i-1}} \lambda(s) ds} - e^{-\int_0^{T_i} \lambda(s) ds}) \frac{\alpha_i}{2} \right] \\ &= \text{Premium}_{a,b}(0; S) \{\text{market model}\}. \end{aligned} \quad (4.32)$$

Equation (4.32) was obtained by utilising the survival probability within the time inhomogeneous Poisson process framework which was presented in (4.21). The default leg at

time 0 is similarly approximated by:

$$\begin{aligned}
 \text{Default}_{a,b}(0; S) &= -\text{LGD} \int_{T_a}^{T_b} D(0, t) d_t \mathbb{Q}(\tau \geq t) \\
 &\approx \text{LGD} \sum_{i=a+1}^b D(0, T_i) (\mathbb{Q}(\tau \geq T_{i-1}) - \mathbb{Q}(\tau \geq T_i)) \\
 &= \text{LGD} \sum_{i=a+1}^b D(0, T_i) (e^{-\int_0^{T_{i-1}} \lambda(s) ds} - e^{-\int_0^{T_i} \lambda(s) ds}) \quad (4.33) \\
 &= \text{Default}_{a,b}(0; S) \{\text{market model}\}.
 \end{aligned}$$

Substituting the market model formulae for the premium and default legs ((4.32) and (4.33) respectively) into the expression for the fair spread in (4.30), the market observable par CDS spread is obtained as:

$$\begin{aligned}
 S_{\text{par}} \{\text{market model}\} &= \frac{\text{LGD} \sum_{i=a+1}^b D(0, T_i) (e^{-\int_0^{T_{i-1}} \lambda(s) ds} - e^{-\int_0^{T_i} \lambda(s) ds})}{\sum_{i=a+1}^b [D(0, T_i) \alpha_i e^{-\int_0^{T_i} \lambda(s) ds} + D(0, T_i) (e^{-\int_0^{T_{i-1}} \lambda(s) ds} - e^{-\int_0^{T_i} \lambda(s) ds}) \frac{\alpha_i}{2}]} \quad (4.34)
 \end{aligned}$$

4.4 Chapter Summary

The modelling of default risk generally falls into one of two categories: structural or reduced form methods. Both approaches were introduced in this chapter. The main difference between them is that the default process can be completely monitored based on default-free information in the former framework. It is the result of an exogenous jump process under the latter approach.

Structural models were dealt with relatively briefly. The objective was to provide the reader with a basic background to them, including the main assumptions underlying them. This is due to the fact that a discussion on default risk modelling would be incomplete without making some mention of the structural framework. Additionally, its inclusion in this chapter renders the dissertation reasonably self-contained since the concepts required to understand the references made to structural models in Chapter 5 (the literature review) were all introduced in Section 4.1.

The essential concepts underlying the reduced form or intensity-based approach were

introduced next, by explaining how they are derived from the properties of Poisson processes. An intuitive understanding of the intensity or hazard rate, integral to the reduced form framework, was fostered.

Lastly, the valuation of credit default swaps was examined. CDSs are the most basic derivatives that can be used for hedging default risk and an understanding of their pricing is essential to a background on default risk modelling. Additionally, the related concept of the par spread was introduced. The formula for the model independent CDS price was derived and applied to the determination of the fair spread within the reduced form framework with time-dependent intensities.

Chapter 5

Literature Review

The main focus of this dissertation is on the pricing of the CVA on credit default swaps. This is an intriguing topic from the perspective that there are three entities (two in the unilateral case) whose default risks are relevant. They are the investor (from whose perspective the CDS is valued) the underlying reference entity (whose default triggers the CDS) and the counterparty to the trade. The concept of right and wrong way exposure is particularly relevant due to the presence of credit contagion. In other words, buying protection on an entity whose default probability is expected to increase on default of the counterparty represents a wrong way risk. Similarly, buying the protection from an entity whose default is negatively correlated with the reference entity's is a form of right way risk.

The aim of this chapter is to provide an overview of the models in the literature that have been developed for pricing CDS CVA. The approaches considered are diverse. Despite the fact that there has been much published on credit contagion and default correlation over the years, the best part of the literature devoted specifically to CDS counterparty risk has originated in the last decade. While much of what has been published on the subject is cited in this chapter, it does not claim to be exhaustive.

5.1 Structural Models

The chapter begins with an examination of a set of models inspired by the Merton-type framework. One of the earliest of these and indeed, some of the pioneering work in the field of CDS counterparty risk, is attributable to Hull and White and dates back to 2001¹ [39]. The key feature of their model is a variable called a credit index describing

¹Although the original working paper is dated 2000.

the creditworthiness of each firm. This could be some function of asset values or credit ratings. They argue that the process for the credit index can be transformed to a Wiener process. This assumption is questioned by Chen and Filipović [40] who point out that it implies independence between the credit risk and the risk-free security market.

Hull and White then determine a time-dependent default barrier for each company such that the entity defaults if the index falls below this barrier. The barrier is chosen to ensure that the model is consistent with market-implied default probabilities. The model calibration entails an iterative procedure, involving two integrals that are numerically evaluated, for obtaining the probabilities of the credit index lying within certain bounds given no prior default. Additionally, the iterative scheme is employed in obtaining the value of the barrier. Note that the barrier is required to be non-horizontal in order to provide an exact fit to the observed default probabilities. This implies that the default distribution is no longer stationary.

Correlated diffusion processes are then used to simulate the credit indices and determine whether or not each of the barriers has been breached. Numerical results illustrate that the impact of counterparty risk on CDS values increases as correlation between the counterparty and reference entity increases and the creditworthiness of the counterparty declines.

As discussed in Section 4.1, one of the main drawbacks of traditional structural models is that they are not consistent with market-implied default probabilities. Hull and White manage to overcome this by using the calibration procedure described above. Schönbucher and Schubert [41], however, point out that the proposed calibration mechanism will be numerically expensive and unstable. In addition the model is restricted to a Gaussian dependency structure. This follows by virtue of the fact that it is the Wiener processes driving the credit indices that are correlated.

More recently, **Blanchet-Scalliet and Patras** [42] approached the CDS counterparty risk valuation problem by deriving expressions for the joint default probabilities of the counterparty and reference entity. Much of the work on structural models prior to their paper relied on the computation of the probability that one of the two names, the counterparty or reference entity, defaulted prior to maturity. In contrast, they obtain an exact knowledge of the default times distribution. This is necessary since the CDS is required

to be revalued at the time of the counterparty default to obtain the CVA². As with traditional structural models, the risk-neutral processes of the asset values are modelled using geometric Brownian motion (GBM) and a default is triggered when the Black-Cox time-dependent, deterministic barrier is struck. The introduction of a correlation parameter is achieved by assuming that the Brownian motions driving the GBM processes are correlated. Again, the correlation is restricted to a Gaussian structure.

An assumption worth noting is that, in deriving the value of the CDS premium leg, Blanchet-Scalliet and Patras simplify the computations by assuming that no reference entity default occurs prior to maturity of the contract. They justify this with the argument that the computation of the fee leg, conditional on the hypothesis that no default occurs, is a good first-order approximation to the unconditional default leg for standard spreads and implied default probabilities. However, when entities are in distress, an accurate CVA model is particularly important. At such times spreads are not expected to consist of ‘standard’ values. During the 2008 credit crunch, spreads on many financial entities and sovereigns reached unprecedented levels. Indeed, prior to default, CDSs on Lehman Brothers traded at intraday spreads in excess of 2,000 basis points.

The final form of the counterparty risk-adjusted CDS value that the authors derive involves a double integral on Gaussian densities and the Bessel function. This sort of function is familiar to physicists and (according to Blanchet-Scalliet and Patras) is ‘well-understood, making numerical approximations easy and efficient’. However, a trader in a typical investment bank pricing CVA may be less familiar with the subject matter. The authors postpone the numerical analysis of the problem to later work. The default probabilities are derived only for the unilateral case in which the default risk of two entities is a concern. Lastly, the fact that the authors restrict themselves to an analytically solvable bivariate lognormal model with constant parameters implies that the model would not be able to be used in practice in its current form. Any institution using a model to price CDSs would require that the model have the capability of being marked to market.

An additional structurally based method was published by **Lipton and Sepp** in 2009 [43]. Their approach models the asset price for a single name using the original GBM formulation of Merton-type models as well as an added jump term. The jump diffusion dynamics make use of a time inhomogeneous Poisson process. Instead of using the time-dependent

²See the unilateral counterparty risk pricing formula (3.1).

default barrier (as in Hull and White [39]) to ensure consistency with market-implied default probabilities, a piecewise-constant jump intensity is utilised. This is calibrated to the current CDS term structure for the entity under consideration. The randomness introduced by the jump process indicates that the default process can no longer be monitored completely based on default-free information as with traditional structural models. This adds a reduced form component to the model.

As with the previous structural models discussed, the diffusions³ are correlated. Additionally, jumps are correlated using the ideas in the 1967 paper by Marshall and Olkin [44]⁴. The jump dynamics are independent of the Wiener process. Lipton and Sepp consider two specifications for the jumps, namely: non-random discrete negative jumps and random exponentially distributed jumps. However, they find that when the model is fitted to current market observables, the maximum Gaussian correlation that can be achieved is about 90% with discrete negative jumps and about 50% when negative exponential jumps are utilised. The latter can thus not be used for highly correlated firms in the same industry. This is a drawback since a situation in which a bank sells protection on another financial institution (for example) is not unrealistic. Also, the fact that the correlation, even with negative discrete jumps, can only reach about 90% is a slight limitation. Even though correlations higher than that are not likely in practice, it represents a model constraint.

Lipton and Sepp then develop the pricing equation for the CDS. They describe a numerical scheme that utilises a combination of partial and ordinary differential equations to solve the pricing problem. The model calibration involves a forward induction procedure and a backward induction procedure is adopted for the CDS valuation. They also present an alternative method that utilises the fast Fourier transform. Formulae are derived for the unilateral case. Some empirical results are represented and indicate that the unilateral CVA is determined largely by correlation and spread volatility as expected. Note however, that the spread volatility is implicit rather than explicitly specified⁵.

The final model examined in this section is a more recent one and is attributable to Yi [45]. He utilises an algorithm that was developed in [46] for simulating joint defaults and migrations based on the first passage times of multivariate Brownian motions. In

³Wiener processes in the GBM portion of the model.

⁴This paper derives the multivariate exponential distribution under three types of conditions. Additionally, it obtains the moment generating function of the resulting distribution.

⁵The volatility parameter specified in the model is the volatility of the firm's assets.

contrast to the former models discussed in this section, the GBM process is not employed here, but rather a drifted Brownian motion is introduced. This is comprised of the sum of the initial value of the process, a drift term and a noise term. The Wiener processes are correlated as is the norm in structural models.

The drifted Brownian motions are then simulated through time and the simulated default correlations obtained. The default barrier for each name is computed, using these simulated values, to match the observed default probabilities. In scenarios in which the drifted Brownian motion process is below the default barrier, the conditional marginal survival probabilities are established by employing a square-root finding procedure such as the Newton-Raphson method. An alternative procedure for the simulation just described is also given in [46].

For the CDS CVA, Yi divides the time interval into discrete buckets and the above simulation procedure is applied in a multi-step way, checking for a default in each bucket. The net present value of the CDS is computed at each point based on the simulated default probabilities. If an entity has defaulted, it is excluded from future simulations and its future distance to default set to zero. This procedure is repeated for the required number of simulations.

Numerical results for both bilateral and unilateral CVA are presented. They illustrate the dependence of both on default correlation. Credit spread volatility is, however, not explicitly modelled.

5.2 Intensity Models

Attention is now turned to approaches in the literature that utilise reduced form or intensity-based methodologies for modelling CDS counterparty risk. This section is divided into two parts. The first of these discusses so-called contagion approaches, in which dependence is introduced via a jump in an entity's intensity upon default of a correlated entity. The second subsection describes models in which dependence is introduced by means of a copula.

5.2.1 Contagion Models

This subsection commences with the model put forward by **Jarrow and Yu** in 2001 [47]. They build on the work of Lando [48] where the intensity follows a Cox process⁶ and is a function of various state variables, particularly the risk-free interest rate (but may include others such as stock prices or credit ratings). In this way, dependency between credit and market risk is introduced. This approach was also employed by Jarrow and Yildirim [49] who obtain correlated defaults by virtue of the fact that firms' default intensities depend on the spot interest rate. In addition to this common factor, Jarrow and Yu include a jump process in the set of state variables. It is designed to capture the interdependence between the default processes themselves.

Since mutual dependence among entities in the model creates a looping default process, they make the assumption that certain firms are 'primary' firms and a selection are 'secondary' firms. The default processes of primary firms depend only on macro variables while those of secondary firms depend not only on these, but also on the default processes of the primary firms. Suppose then that Firm A is considered to be a primary firm and B a secondary one. A's intensity is then modelled as a linear function of the risk-free interest rate which follows a one factor short rate process. The intensity of B is modelled in the same way as that of A with an additional term dependent on the default time of A. This term will be 0 before the default of A and a constant factor if A has defaulted. In the case of B's default being positively correlated with A's, the coefficient of the jump term will be positive so that B's riskiness increases upon a default by A.

In applying the model to the valuation of a CDS with counterparty risk, a simple point process is used to model the intensity of the reference entity and a constant interest rate is assumed. Jarrow and Yu consider an idealised CDS in which the protection buyer makes premium payments until maturity regardless of whether or not the counterparty has defaulted. Moreover, the payment by the protection seller to the protection buyer occurs only at maturity of the CDS and not at the time of the reference entity's default. In practice, if a model is to be used to price and hedge CVA, a more accurate representation of reality⁷ is required.

⁶See Section 4.2.3.

⁷Of course, it is possible that a contract may specify that the payment occurs only at maturity of the CDS, but this is not the norm.

Additionally, the way in which the model would be calibrated to the market, and the determination of the magnitude of the jump terms in the hazard rates of secondary firms, is unclear. Schönbucher and Schubert [41] point out that deriving the survival probability for a single obligor is already a major task, thus making calibration very difficult.

Yu [50] builds on the intensity-based approach of Jarrow and Yu [47] by presenting an algorithm for constructing an arbitrary number of default times with intensities dependent on both observed defaults and a common stochastic process. He utilises the so-called total hazard construction of default times and defines the total hazard accumulated by an entity as a sum of its integrated hazard rates. Default times of the firms under consideration are then obtained via a recursive procedure that utilises the same time steps as the sum in the total hazard construction. Default times are determined by generating independent and identically distributed unit mean exponential random variables at each time step.

The advantage of Yu's approach over that of Jarrow and Yu is that the assumption of primary and secondary firms is not required. Mutual dependence is obtained without the looping default that hampered the efforts of Jarrow and Yu. The pricing of a CDS is considered in both the unilateral and bilateral cases. However, the simplified contract assumptions (that the protection buyer pays premium until maturity and the seller compensates the buyer for a default event only at expiration of the CDS) are maintained.

Further development was undertaken by **Leung and Kwok** [51], who formulated the intensities of the entities under consideration in the same manner as Yu [50]. In other words, the correlation between firms was introduced by means of a jump in the intensity of one entity upon default of another. Assuming only two firms, A and B, A's intensity is a linear function of the indicator variable that is equal to 1 upon default of B and 0 otherwise. The deterministic term in this function, as well as the coefficient of the indicator, are assumed to be constant.

Instead of utilising the total hazard construction considered by Yu to remove the looping default phenomenon that complicates computations, Leung and Kwok employ a change of measure that was introduced by Collin-Dufresne, Goldstein and Hugonnier [52]. Collin-Dufresne et al. demonstrated that defaultable claims can always be valued using expected risk-adjusted discounting provided that the expectation is taken under a slightly modified probability measure. This new probability measure assigns a zero probability to paths

along which default occurs prior to maturity of the defaultable claim and is thus only absolutely continuous with respect to the risk-neutral probability measure [52]. Accordingly, Leung and Kwok define a firm-specific probability measure for each of the entities under consideration. This measure assigns a zero probability to paths on which default occurs prior to maturity of the CDS under consideration. Under this measure, the joint density of the default times is obtained and the looping default phenomenon encountered by Jarrow and Yu is no longer a concern.

The CDS is then valued with both unilateral and bilateral counterparty risk, under the assumption that the protection buyer pays a continuous premium at the swap rate until expiry of the CDS or until default of one of the three entities involved. Recall that Jarrow and Yu [47] and Yu [50] both assumed that the premium payments would continue until maturity of the CDS regardless of whether or not a default had occurred prior to that. In addition, the payment by the protection seller to the protection buyer is made at the end of the settlement period after default, rather than the more unrealistic assumption that this occurs on expiration of the CDS. A simplification incorporated into the pricing is that the protection seller can walk away from the contract without the requirement to pay the buyer if the protection buyer defaults prior to the maturity of the CDS. This adds an extinguishable element to the trade⁸ which fails to capture the economic reality of a simple CDS contract.

5.2.2 Copula Methods

In contrast to the contagion models discussed above, a separate branch of CDS counterparty risk modelling has been developed, in which a copula is employed to model the dependence between the default times of the relevant entities. When hazard rates are stochastic, the simplest approach to introducing default correlation would be to correlate the diffusions driving the intensity processes of each firm. It is widely documented that this ‘naïve’ approach produces a low and unrealistic level of correlations across the actual default events. This is discussed, among others, by Brigo and Mercurio [25], Brigo and Chourdakis [53], Hull and White [39] and Jouanin et al. [33]. This limitation is overcome by using a copula to correlate the default times of the entities under consideration.

A copula is an efficient way of introducing dependence since, on the one hand, the default dynamics of a single entity are required to be modelled in a way that is consistent with

⁸Recall the discussion on extinguishers in Section 3.6.

the current term structure of CDS spreads for that name. On the other hand, the correlation between defaults necessitates the use of a dependence structure. Copulas join or couple multivariate distribution functions to their one-dimensional marginal distribution functions [54]. The univariate margins and the dependence structure can be separated, with the latter completely characterised by the copula. In addition, copulas are the most general way to view dependence of random variables. They thus provide a flexible dependence framework that is not limited to linear correlation [25, 54]. An introduction to copulas and the associated concepts can be found in Appendix D.

The first copula-based model under discussion was introduced in a particularly general framework by **Schönbucher and Schubert** [41] in 2001. Under their approach, the dynamics of the intensity processes are not specified, but left to the discretion of the user. The copula is introduced through the specification of the joint distribution of uniform random variables. These are obtained by transforming the intensities or ‘pseudo-intensities’ to which the authors refer⁹. Once one or more of the entities has defaulted, a conditional distribution of the default times is obtained. The authors derive the form of a defaultable bond within the framework they have established. The counterparty risk of a CDS is mentioned, but not explicitly derived.

More recently, **Brigo and Chourdakis** [53] developed a unilateral CDS CVA model that caters both for default time correlation and volatility of the credit spreads of the entities involved. Previous models on the subject of counterparty risk for CDSs focused on the correlation element, the necessity of which is not disputed in the literature. However, volatility of credit spreads has generally been ignored. Even in the structural models discussed above, it was implicit rather than explicitly modelled. In contrast, in the original CVA models (such as in the case of an interest rate swap underlying) volatility was taken into account and the correlations between the default probability of the counterparty and the underlying risk drivers of the derivative’s value were ignored. As discussed in Chapter 3, CVA essentially amounts to valuing an option. Since the price of such an instrument is driven by the volatility of the underlying, the exclusion of spread volatility would be an oversight.

Brigo and Chourdakis employ a CIR++ process (which is a Cox Ingersoll Ross process with a deterministic shift that facilitates calibration to the current term structure)

⁹For a particular entity, the pseudo-intensity coincides with the intensity in the independent case or when information is restricted to that entity alone.

to model the intensity of each entity. Dependence is introduced by means of a copula that links the default times of the reference entity and counterparty. More precisely, the exponential triggers of the uniform random variables of the default times of the two names are linked via the dependence structure. The intensities of each entity are simulated and integrated. Default times are then determined based on uniform random variables that are simulated by means of the copula. Once these have been obtained, it remains to compute the net present value of the CDS in each scenario that has a positive contribution to the CVA¹⁰; that is, in each scenario in which the counterparty defaults prior to both the reference entity and the CDS maturity.

A fractional fast Fourier transform (which is a semi-analytic method) is used to obtain the reference entity survival probabilities from the counterparty default time to the remaining CDS premium payments until expiration. These are conditional survival probabilities, dependent on the counterparty default time and the market information available at that time. Once they have been determined, the calculation of the CDS NPV is straightforward. In each relevant counterparty default scenario, the exposure is computed as the greater of the NPV at default and zero. The results of each scenario are then averaged to obtain the final credit valuation adjustment for the CDS under consideration. Note that the alternative to the Fourier transform would be to utilise a Monte Carlo simulation upon a Monte Carlo simulation. This would be prohibitively time-consuming in practice.

The model by Brigo and Chourdakis was extended by **Brigo and Capponi** [22] to value BCVA. The intensities again follow CIR++ processes, with an additional process for the investor's intensity incorporated. The copula now models a three way dependence between default times. The simulation paths on which the counterparty CVA will give a non zero contribution are now those in which the counterparty defaults, not only prior to the reference entity and CDS maturity, but also prior to the investor.

The DVA, or the investor's own risk portion of the BCVA, is calculated in a similar manner to the CCVA, with the exception that the scenarios in which a revaluation of the CDS is required are those in which the investor defaults prior to the counterparty, reference entity and CDS expiration. The survival probabilities of the reference entity that are used to compute the NPV at the investor default times are then conditional on the investor rather than the counterparty default times and the market observable

¹⁰See the CVA portion of the General Counterparty Risk Pricing Formula in (3.1).

information at those times.

5.3 Alternative Approaches

In this section, CDS counterparty risk pricing models in the literature that do not fall within either the traditional structural or reduced form frameworks are considered. The portion on Markov chain models could have been included in the previous section since these models are all intensity-based. However, they entail a slight paradigm shift from typical reduced form modelling.

5.3.1 A Generalized Affine Model

Chen and Filipović [40] set forth a CDS valuation model in 2003 that incorporated aspects of both the structural and reduced form models available at the time. In their approach, each firm is assigned a credit index that has the same meaning as the one in Hull and White [39] discussed above. Similar to the earlier intensity approaches (such as Jarrow and Yu's [47] discussed at the beginning of the previous section), a common factor, the risk free rate¹¹ is incorporated. For simplicity, a single factor affine model is used for this purpose. Additionally, instead of default being triggered by the credit index of each firm hitting a barrier, an indicator is included that takes on a value of 1 when the entity has defaulted and 0 otherwise.

The theory of positive affine processes is rich and unfamiliar to many of the market practitioners who would be end users of the model. Chen and Filipović derive the fair spread of a CDS under counterparty risk that is obtained via a 7-dimensional affine process. While calibration to current market spreads is mentioned, the actual detail behind this procedure and whether or not it is possible to obtain an exact fit is omitted. In addition, the reference entity is assumed to be a primary firm¹² in that its default probability is unaffected by that of either the counterparty or the investor. This is too simplistic, since a default of one of these entities is expected to have a direct impact on the reference entity, unless the correlation between them is zero. In terms of the CDS contract that is considered, certain assumptions are made that convey an extinguishable element to the trade. Again, this does not accurately reflect the economics of the CVA on a simple CDS.

¹¹The authors point out that further factors could be added.

¹²Recall that the notion of primary and secondary firms was introduced by Jarrow and Yu [47].

5.3.2 Markov Chain Models

Early work that involved the use of Markov transition matrices to compute CDS counterparty risk was presented by **Davis** in 2002 [55]. In his model, hazard rates are assumed to be constant across all maturities. This does not facilitate calibration to an observed term structure of CDS spreads. Building on Davis' formulation, **Walker** [56] introduced time-dependent hazard rates to the model in 2005 to correct for this shortcoming. In his paper, only unilateral counterparty risk was considered. The main attributes of his model will now be highlighted.

Within the framework, each entity is considered to be in one of two states, depending on whether or not it has defaulted. The combination of the two entities (the reference entity and counterparty) can thus be in one of four states. They are: 1) both counterparty and reference entity are in default; 2) counterparty is in default, but reference entity is not; 3) counterparty is not in default, but reference entity is and 4) neither counterparty nor reference entity is in default. The transition rates give the probability per unit of time of moving from one state to the other. It is assumed that joint defaults are not possible. In other words, the transition rate from the state in which neither entity has defaulted to the state in which both have defaulted, is zero. In order to obtain these time-dependent transition rates, the set of Kolmogorov forward differential equations is required to be solved. This is achieved via a perturbation approach.

Dependence between the two entities is determined by the change in the transition rate of one entity when the other has defaulted. For the reference entity this is the difference between:

- (i) the transition rate from the state of neither entity having defaulted to the state in which only the reference entity has defaulted and
- (ii) the transition rate from the state of only the counterparty having defaulted to the state in which both entities have defaulted.

In valuing the CDS under counterparty risk the assumption is made that premiums are paid continuously in order to simplify the computation. Since the CVA is not explicitly modelled, the payment of the NPV on counterparty default to the in-the-money firm is not modelled. As with a number of the other approaches that have been discussed in this chapter, this represents an extinguishable trade that does not reflect the economic realities of counterparty risk. In addition, the volatilities of the underlying credit spreads

are not considered since the hazard rates, although time-dependent, are not stochastic.

In 2009, **Crépey, Jeanblanc and Zargari** [57] put forth a Markov Chain model in which the correlation parameter is modelled via a copula. In their paper, only the unilateral CVA of a payer CDS is considered¹³. There are thus four possible states of the world. One of the shortcomings of the Markov copula approach is that it does not allow for default contagion effects in the sense that a default by the counterparty does not increase the default probability of the reference entity. In order to overcome this weakness, unlike Walker, Crépey et al. assume that joint defaults are possible. They attribute this to the necessity of incorporating wrong way risk into the framework, rather than claiming that simultaneous defaults are a possibility.

The Kolmogorov backward ordinary differential equation (ODE) related to the valuation of the risky CDS is then solved numerically. Since CVA is explicitly modelled here, the simplifying extinguishable assumptions are not made. The computed CVA is restricted to the unilateral adjustment on a payer CDS. In addition, although the hazard rates are time-dependent and able to be calibrated to the CDS term structure, they are not stochastic. Credit spread volatility is thus not modelled.

The authors, **Crépey, Jeanblanc and Zargari**, then extended their model to incorporate stochastic credit spreads [58]. As in Brigo and Chourdakis [53] and Brigo and Capponi [22], the stochastic processes specified are CIR++ processes that can be fitted to the current term structures of CDS spreads. An alternative specification of the intensity process was also considered in the form of an extended CIR process with time-dependent parameters and no deterministic component¹⁴. Crépey et al. report that the CIR++ specification is more robust and accurate than the extended CIR one and computations are quicker.

As in the authors' 2009 publication, the correlation between default times is incorporated by means of a copula and joint defaults are possible in order to incorporate wrong way risk. A semi-explicit formula is derived for the CVA. The formulation is restricted to the unilateral adjustment on a payer CDS.

¹³The CVA is being calculated from the point of view of a risk-free investor who has purchased protection on an underlying reference entity.

¹⁴Recall that the CIR++ model has constant parameters and a deterministic shift that is necessary for calibration to market spreads.

A similar paper to the one described in the previous paragraph is attributable to **Bielecki, Crépey, Jeanblanc and Zargari** [59]. The topic of CVA hedging is considered in addition to pricing. Again, the bilateral case is omitted.

5.4 Chapter Summary

An overview of the various methods employed for pricing CDS counterparty risk was presented in this chapter. The aim of the discussion was to provide the reader with an indication of the ways in which the problem has been approached in the literature.

The first part of the chapter was devoted to the development of structural methodologies. This was followed by a description of reduced form approaches, which was split into two parts. The first of these involved contagion models in which dependence was introduced by means of a jump in the intensity of one entity upon default of another. Models in the second part utilised the copula-based approach for establishing dependence between names. Lastly, a brief summary of two types of models that have been developed, but do not fall within the traditional structural and reduced form frameworks, was presented. These were the generalised affine approach (which could be seen as a combination of structural and intensity models) and Markov chain models (which incorporate reduced form techniques).

The reader should now have a basic understanding of the diverse ways in which the modelling of the CVA on a CDS might be approached. Additionally, the main shortcomings and advantages of these methods should be clear. The subject is relatively new and the literature on the topic is continually expanding. As such, there is not yet a market standard for CDS CVA valuation. This is evident from the variety of methods presented in this chapter.

Part II

CDS CVA Model Implementation - the Unilateral Case

Chapter 6

Model Background

The computation of the CVA on a credit default swap is the main objective of this dissertation. Part II is focused on the selection and implementation of a unilateral CVA model. Additionally, it includes an examination of the results such a model produces. The extension to a bilateral adjustment will be the subject of Part III of the dissertation.

The rationale behind the selection of the model to be implemented, from among those reviewed in Chapter 5, is set forth in this chapter. The introductory details of the model are the subject of the remainder of the chapter. The basic notation required is presented in Section 6.3.1. This is followed by an explanation of the underlying model assumptions in Section 6.3.2. Subsequently, the stochastic intensity process is introduced, followed by an outline of the model implementation steps.

The chapter begins with a discussion of the publicity that CDSs have received over the past two or three years (at the time of writing). This is intended to create some context around them since they are integral to this dissertation.

6.1 Credit Default Swaps in the Spotlight

Credit default swaps have received a great deal of criticism since the start of the 2008 financial crisis, with many blaming them for the carnage that ensued following the default of Lehman Brothers. It should be noted that part of this condemnation has arisen from commentators who fail to understand the differences between simple CDSs and the far more complex asset-backed collateralised debt obligations (CDOs) that were at the root of the crisis, certainly in the early stages [8].

Part of the reason for this misconception was the near default of American International Group (AIG). Two key points stand out from this. The first is that AIG's exposures were in one direction only and of significant scale [7]. The fact that such large unhedged exposures were created is a risk management and regulatory failure and the CDS instrument is not to blame. Significant, single-sided exposures should be avoided in any asset class (be it credit or otherwise) unless the institution with which they reside has sufficient capital to withstand the potentially large losses that may eventually accompany such exposures.

Additionally, much of the exposure of European banks to AIG was not initially collateralised [7]. AIG had a triple A credit rating and was considered to be risk-free by its financial counterparts. This was also true of other investment grade financial institutions. A number of the top tier investment banks had agreements with each other under which counterparty risk charges were overlooked in OTC trades conducted with one another. This notion is particularly problematic given the large concentration of CDS trades within a few of these main institutions. In Europe, the top ten counterparts of each of the large banks surveyed by the ECB in August 2009 accounted for 62-72% of the continent's CDS exposures [7]. Within the bespoke CDS contracts, it was stipulated that collateral would have to be posted by AIG in the event of a credit rating downgrade. By the time that these overdue downgrades arrived, after two consecutive quarters of losses totalling USD 13 billion, AIG had insufficient liquidity to make the USD 20 billion required in collateral calls [7]. It is important to note from this that:

- (i) The credit ratings of the ratings agencies could not be relied upon to manage the default risk as they were not reactive enough. In fact, observe the graph of the Lehman Brothers five year CDS spread along with its credit rating for 2008. This is shown in Figure 6.1. The bank had an investment grade rating on the day it filed for bankruptcy. This was after two consecutive quarterly losses of the largest magnitude in its history¹.
- (ii) No entity is risk-free and counterparty risk valuation is required for all trading counterparts, regardless of their creditworthiness or perceived inability to fail.

More recently, the CDS market was an easy target for blame in the first part of 2010 when the Greek sovereign debt crisis began receiving attention. In March 2010, the Greek prime

¹On 9 June 2008 Lehman Brothers announced a USD 3 billion loss for the second quarter of 2008 and one of USD 4 billion on 11 September 2008 [7].

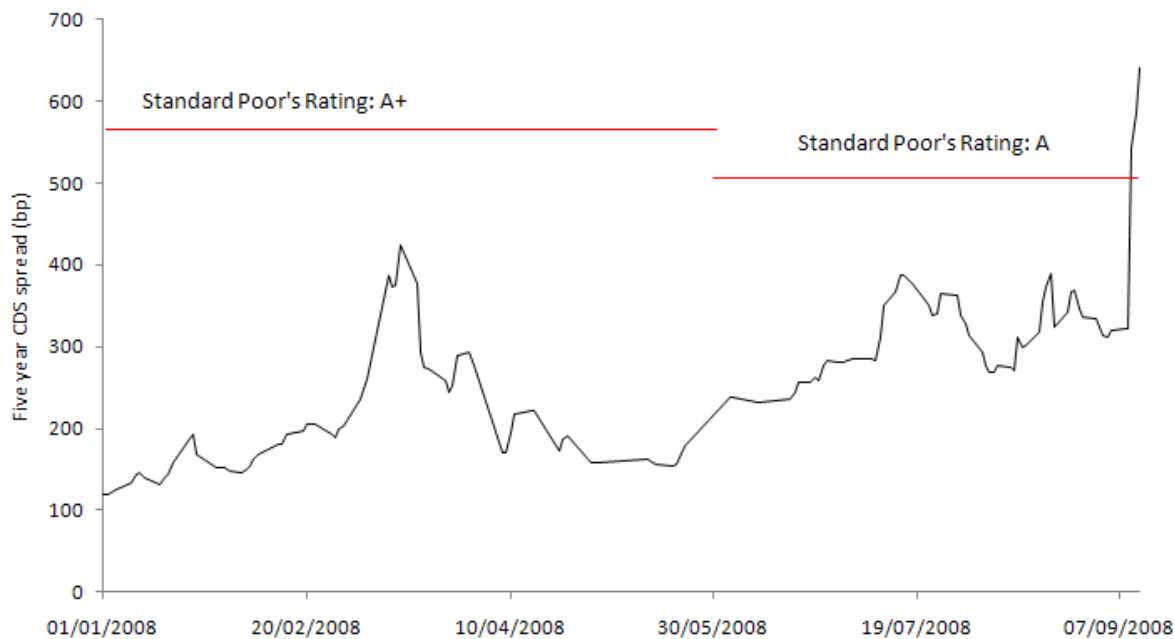


Figure 6.1: Five year CDS spread and credit rating of Lehman Brothers. *Spreads are close of business quotes and were obtained from Bloomberg. Ratings information courtesy of Standard and Poor's.*

minister (George Papandreou) called for a crack down on speculative trading, arguing that exotic bets had driven up Greece's borrowing costs and threatened its effort to ease its debt crisis [60].

Greece's inability to issue debt at the low levels their leaders would have preferred was blamed on credit default swaps. The argument made was that higher CDS spreads, brought about by naked² CDS purchasers speculating on the country's future, drove up bond spreads [62]. A comprehension of the bond CDS basis (see [63]) renders this argument superfluous (see [62] for an exploration of this statement). Furthermore, data from the Depository Trust and Clearing Corporation (DTCC) showed that the notional outstanding on Greek sovereign CDSs was a mere USD 9 billion compared to USD 406 billion on bonds and that the former stayed relatively flat during the first three months of 2010 (when Greek debt was particularly volatile) [64]. Additionally, the hedge funds blamed for the speculative purchases of Greek CDSs were, in fact, net buyers of the country's debt during this time. This is due to many of them having foreseen the mismatch in

²A naked CDS position arises when the purchaser of protection does not own the bond issued by the underlying reference entity of the CDS [61].

the Eurozone, with the weaker member states (Greece, Spain, Portugal and Italy) being priced similarly to stronger states such as Germany many years prior to the crisis. The trade was to purchase CDS protection on these weaker states for a pittance and wait for spreads to eventually reflect the differences in the economies of the weak and strong Eurozone members. As such, many hedge funds (those that held the strategy) were purchasing bonds during the Greek crisis in order to cover long CDS positions [65].

At the end of 2010, a probe by Brussels found no conclusive evidence that CDS trading activity drove up funding costs for European Union member states. It concluded that CDS spreads for the more troubled countries seemed to be low relative to corresponding bond yields, which implied that CDS spreads could not be the cause of high bond yields in these countries [66].

The purpose of the above discussion on the Greek sovereign debt crisis was to demonstrate the misconceptions in the media regarding CDSs and the negative publicity they have received of late. In addition, it illustrated the use of the CDS market as a source of information on the credit quality of issuers. For entities on which CDSs are traded, CDSs and not bonds are now the preferred gauge for an entity's creditworthiness [67]. This avoids the necessity of benchmarking off a perceived risk-free rate. Additionally, issuers tend to have a number of bonds outstanding. These trade at different prices and yields to maturity. Obtaining an overall, market-based sense of an issuer's riskiness is thus difficult looking only at the bond market. On the other hand, the CDS market provides a market-based measure of the credit quality of a wide variety of debt instruments outstanding for a given issuer. This makes the credit quality of an issuer more transparent [67].

Before digressing into a discourse on the merits of credit default swaps, this section is concluded with the following observations:

- (i) The extent to which counterparty risk influenced the financial world during the credit crisis illustrated the importance of CDSs as a tool for the pricing and management of counterparty risk.
- (ii) The ongoing (at the time of writing) sovereign debt crisis illustrates the usefulness of sovereign CDSs for hedging currency and rates CVA to a particular region.
- (iii) CVA is required on positions with all counterparties. No entities can be regarded as risk-free.

- (iv) It is no longer sufficient to consider only the default risk of the underlying reference entity in a CDS trade. The counterparty risk is also relevant. Accurate models for the computation of CDS CVA are thus a requirement.
- (v) Despite populist criticism, CDSs will continue to be traded, particularly as institutions seek to hedge the profit and loss volatility attributable to CVA.

6.2 Model Selection

The model implemented for the valuation of unilateral CDS CVA is based on the one by Brigo and Chourdakis [53]. It was described briefly in Section 5.2.2 of the literature review in Chapter 5 and will be examined in detail in the remainder of Part II. Note that this dissertation aims to be practically relevant, in that the model applied to the CDS CVA valuation problem could realistically be utilised by a financial institution such as an investment bank. It is not sufficient for the pricing problem to be addressed in a purely theoretical manner. With this in mind, the justification for the model selection is discussed below.

- (i) The dependence in the model is introduced by means of a copula that links the uniform random variables of the exponential triggers of the default times of the reference entity and the counterparty (see Appendix D for a brief background to copula functions). The advantages of such a dependence structure are threefold:
 - (a) Copulas are functions that join or couple multivariate distribution functions to their one-dimensional marginal distribution functions [54]. They represent an efficient way of introducing dependence since the univariate margins and the dependence structure can be separated. Thus, the hazard rate processes of each of the entities can be calibrated to the individual CDS term structures observed in the market, regardless of the copula selection [25]. This implies that the model can be marked to market, a minimum requirement for any pricing mechanism to be utilised in practice.
 - (b) Copulas are the most general way of modelling dependence. This generates capacity for a wide variety of dependence structures, extending far beyond the most common choice of linear correlation [25]. The selection of the copula is not determined by the model, leaving scope for extension beyond the Gaussian copula that has been implemented in this dissertation.

- (c) The motivation for introducing dependence between the uniform random variables of the exponential triggers of the default times, rather than between the intensity processes of the individual firms, was mentioned briefly in Chapter 5. Essentially, the most obvious means of introducing dependence into the model would be to correlate the noise processes driving the hazard rates of each of the entities. This has the advantage of keeping the pricing model from the independent case unchanged. However, an explicit derivation of the implied correlation of the default times demonstrates that high default correlations cannot be attained in this framework [33]. In fact, the maximum default correlation that can be attained when hazard rates are perfectly positively correlated is in the same order of magnitude as the default probabilities, typically 0.5 to 3 %. This was demonstrated, among others, by Schlögl [68]. The fact that Brigo and Chourdakis overcame this flaw (inherent in many multiname default risk models) by selecting to introduce dependence between the uniform random of the exponential triggers of the default times, rather than between the hazard rates, is part of the model's appeal.
- (ii) The notions of right and wrong way risk were introduced in Section 3.1.1. They are particularly relevant to the CVA on a credit default swap due to the presence of credit contagion. In the event that the default times between the counterparty and the reference entity demonstrate positive dependence, the survival probability of the latter is expected to decrease upon a counterparty default. This would increase the value (replacement cost) of the underlying CDS from the point of view of an investor who is long protection. The investor's counterparty default loss would thus be larger than it would have been were the two entities completely independent. This is wrong way exposure. Similarly, the loss to an investor who is short protection would increase as correlations decrease.

Since completely independent credit events are highly unlikely in practice, it is vital for a CDS CVA model to capture the effects of right and wrong way risks. The copula linking the exponential triggers of the uniform random of the default times of the entities caters for these risks. The discussion in the previous point indicates that the magnitude of the default risk is appropriately captured by the choice to correlate default times over hazard rates.

- (iii) Credit spread volatility is explicitly modelled by virtue of the fact that we place

ourselves in a Cox process setting³. In other words, the intensity of each entity follows a stochastic process whose volatility can be specified. In the literature review, the inclusion of correlation in favour of volatility in traditional CDS counterparty risk models was discussed. It was pointed out that this was the opposite of the manner in which the CVA of other instruments had been treated in earlier CVA modelling. As such, the model has been selected to ensure that both default time correlation and the volatility of credit spreads are included. Since CVA has an optionality component, it would be an oversight to ignore volatility, the main driver of option prices. Even though CDS options are not observable in the market at present, the model gives users the ability to stress volatilities and manage the associated risk. Once these quotes become observable in the market, it will be a simple matter to facilitate them.

- (iv) The approach implemented in this dissertation can be applied to any basket credit derivatives pricing model. This implies that the methodology could be applied globally within an institution, across all applicable instruments, rather than for the pricing of CDS CVA in isolation. This is an important point to consider for the sake of consistency.
- (v) The survival probabilities required for the revaluation of the CDS upon counterparty default are obtained by means of a fractional fast Fourier transform. This avoids the time-consuming generation of additional Monte Carlo paths at each counterparty default time. Such practical considerations render the model feasible for practical implementation.
- (vi) The survival probabilities are obtained at the counterparty default times and then employed in the valuation of the CDS at these times. Therefore, the CVA on more bespoke CDS-type instruments could also be priced within the framework.
- (vii) The model can be extended to cover BCVA without considerably increasing its complexity. This is the subject of Part III.
- (viii) Lastly, and perhaps most significantly, it is a natural extension of the market model currently used for valuing CDSs. This is due to the fact that it falls within the reduced form framework in which the intensity process for each name is calibrated to the current CDS term structure. End users will thus have an understanding of

³Recall the discussion (in Section 4.2 of Chapter 4) on the various possibilities for modelling the intensity.

the model at the outset. They would also be expected to trust an approach within the paradigm in which they are used to operating.

6.3 Preamble

The motivation behind the model selection has been presented. The remainder of the chapter is devoted to the background necessary for its implementation. The basic notation required for this purpose is set forth in this section, along an explanation of the assumptions underlying the CDS CVA model.

6.3.1 Notation

We represent the default times of the reference entity and the counterparty as τ_1 and τ_2 respectively. The entity from whose perspective the CVA is computed is called the investor. The maturity of the underlying CDS, transacted between the investor and the counterparty, is denoted by T . We place ourselves in the probability space $(\Omega, \mathcal{G}, \mathcal{G}_t, \mathbb{Q})$ where \mathbb{Q} is the risk neutral measure. The following filtrations and sub-filtrations are defined:

- The filtration $(\mathcal{G}_t)_t$ models the information flow of the entire market including credit and defaults.
- The sub-filtration \mathcal{F}_t represents all market observable quantities except for the default events.
- The sub-filtration \mathcal{H}_t contains only the default events.

As stated earlier, we operate in a reduced form or intensity-based framework, with stochastic intensity processes used for both entities' intensities. We are thus in a Cox process setting⁴. The intensities of the reference entity and the counterparty are denoted by the symbols λ_1 and λ_2 respectively. The integrated intensities are then:

$$\Lambda_1(t) = \int_0^t \lambda_1(s) ds \quad \text{and} \quad \Lambda_2(t) = \int_0^t \lambda_2(s) ds \quad (6.1)$$

⁴See Section 4.2.3.

respectively. It follows that the respective integrated intensities at the default times of each of the entities are:

$$\Lambda_1(\tau_1) = \int_0^{\tau_1} \lambda_1(s) ds := \xi_1 \quad \text{and} \quad \Lambda_2(\tau_2) = \int_0^{\tau_2} \lambda_2(s) ds := \xi_2. \quad (6.2)$$

Note that ξ_1 and ξ_2 are standard, unit mean exponential random variables⁵. Assuming invertibility of the integral in (6.2), the default times may be expressed as:

$$\tau_1 = \Lambda_1^{-1}(\xi_1) \quad \text{and} \quad \tau_2 = \Lambda_2^{-1}(\xi_2). \quad (6.3)$$

Finally, the uniform random variables, U_1 and U_2 , are defined as:

$$U_1 := 1 - e^{-\xi_1} \quad \text{and} \quad U_2 := 1 - e^{-\xi_2}. \quad (6.4)$$

The fact that U_1 and U_2 are uniformly distributed follows from the properties of exponential random variables. That is, if x is a standard, unit mean exponential random variable, then $1 - e^{-x}$ has a uniform $[0, 1]$ distribution [19]. The transformation from exponential to uniform random variables is necessary for the implementation of the copula. A basic background to copulas is presented in Appendix D.

6.3.2 Assumptions

The assumptions listed below will be utilised in the quantification of the CDS counterparty risk (the CVA).

- (i) The risk-free interest rate is deterministic. This implies that the default times of both the reference entity and the counterparty are independent of the discount factors in the market. The model can be extended to include stochastic interest rates if this is desired. However, the added complexity would not substantially affect the results since the main drivers contributing to the CDS CVA are the default probabilities of the reference entity and the counterparty and the correlation between their default times. If CDSs on floating rates were to be included, this could potentially have an effect.
- (ii) Intensities are strictly positive, from which we obtain invertibility of the integrated intensity in (6.2)⁶. Aside from the invertibility requirement, an increasing integrated

⁵This should be clear from the background to Poisson processes in Section 4.2.

⁶Recall that, for a function to be invertible, it is required to be one-to-one [26]. If intensities are

intensity implies a reducing cumulative survival probability of the entity over time. Intuitively, this is a logical requirement.

Note that a decreasing cumulative survival probability function does not preclude the possibility of the marginal default probability at a particular point in time being higher than at a later point. For example, when an imminent default is priced into the CDS term structure, spreads are frequently inverted. This represents the fact that the marginal probability of the investor surviving in the short term is lower than in the long term. However, the cumulative probability of surviving in the long term incorporates the marginal probabilities along the entire curve. The cumulative survival probability therefore remains a decreasing function over time.

- (iii) The reference entity and counterparty intensity processes are independent of each other. This is due to the dependence being incorporated by means of the copula linking the uniform random variables of the exponential triggers of the default times of the two entities.
- (iv) The uniform random variables U_1 and U_2 are correlated through the copula function:

$$\mathbb{Q}(U_1 < u_1, U_2 < u_2) = C(u_1, u_2). \quad (6.5)$$

- (v) The copula in the implementation is the Gaussian copula whose correlation parameter ρ is assumed to be constant.
- (vi) The recovery rate, R , is a deterministic constant. The amount recovered on a derivative traded with a defaulted counterparty is thus assumed to be $R \times \max(\text{NPV}(\tau_2), 0)$, where the NPV at τ_2 is measured from the perspective of the investor.

6.3.3 The Stochastic Intensity Process

This section contains an introduction to the stochastic intensity process that is fitted to the hazard rates of each of the entities (the reference entity and the counterparty). An explanation of its applicability to the modelling of hazard rates is included. The calibration of the process is discussed in the next chapter when the actual implementation of the unilateral CDS CVA model is considered.

strictly positive, the integral will be a strictly increasing function and thus one-to-one.

The stochastic intensity is modelled via a CIR++ process. This is the sum of a Cox Ingersoll Ross (CIR) process and a deterministic shift. The deterministic portion of the process is required for the calibration of the intensities to the current CDS term structure for a given name. Deterministic shift extensions of short rate models were first introduced by Brigo and Mercurio [69] in 1998. The intensity in reduced form models is the analogue of the short rate in interest rate models, implying that the theory developed for the latter can be applied to intensity processes [25].

The (stochastic) hazard rate is given by:

$$\lambda_j(t) = y_j(t) + \psi_j(t; \beta_j), \quad t \geq 0, \quad j = 1, 2, \quad (6.6)$$

where ψ is a deterministic function and y follows a CIR process such that:

$$dy_j(t) = \kappa_j(\mu_j - y_j(t))dt + \nu_j\sqrt{y_j(t)}dW_j(t), \quad j = 1, 2. \quad (6.7)$$

Note that $j = 1$ for the reference entity dynamics and $j = 2$ for the counterparty dynamics. The parameter vectors are $\beta_j = (\kappa_j, \mu_j, \nu_j, y_j(0))$, which are all positive deterministic constants. The W are standard Brownian motions⁷. Note that $y_j(0)$ falls under the set of parameters to be calibrated. The following initial condition should thus be satisfied:

$$\psi(0; \beta) = \lambda_0 - y_0. \quad (6.8)$$

An important point to note, regarding the CIR++ model calibrated in the next chapter, is that the Feller condition ($2\kappa\mu > \nu^2$) is not imposed. This condition prevents the CIR process from attaining a zero value⁸. Rather, a positivity constraint is imposed on the shift (ψ) which ensures that the intensity (λ) cannot attain a value of zero. If both this positivity constraint and the Feller Condition were imposed, the model would be over-parameterised and the range of possible CDS option implied volatilities that could be attained would be restricted [57].

The applicability of the CIR++ process to the modelling of intensities is threefold:

- (i) The CIR process is in the class of mean-reverting processes [25]. This implies that, rather than increasing indefinitely, quantities modelled by this process will revert

⁷As discussed under the assumptions, these are independent.

⁸The CIR process cannot become negative. See [25] for further information on this type of process.

back to some long run mean over time. The intensity of an entity is expected to behave in such a manner. In contrast to stock prices that are expected to continue increasing over time, intensities (like interest rates) cannot drift upwards indefinitely.

- (ii) Hazard rates are required to remain positive in order to ensure invertibility of the integrated hazard rate. The CIR process cannot attain negative values. It is thus appropriate.
- (iii) The final justification of the use of the CIR++ process is the deterministic shift portion of the model. This enables calibration to current market spreads. Without this feature, the model could not be marked to market, precluding its use in practice.

6.4 Model Outline

The chapter is concluded with a high level outline of the model implementation steps. These will be addressed in detail in the following two chapters. Recall that the CVA valuation necessitates the computation of the expected value of the replacement cost of the CDS on default of the counterparty, as shown in the CVA term of Equation (3.1).

The methodology employed to achieve this is to determine the relevant counterparty default times and then to obtain the survival probabilities of the reference entity at these times. The CDS can then be revalued at these points, using the calculated survivals, and the replacement cost determined. In the event that the counterparty and reference entity default times demonstrate positive dependency, these survival probabilities are expected to be lower than the pre-default values (the counterparty default). The opposite applies when the correlation is negative.

The fundamental steps in the model implementation are expressed in the diagram in Figure 6.2. The first stage involves stripping the hazard rates from the quoted CDS term structures of the two entities (the counterparty and the reference entity). This procedure is described in Section B.1 in Appendix B. A CIR++ model is then fitted to each of the hazard rate term structures.

Utilising the parameters obtained by means of this calibration procedure, the intensity process of each of the names is simulated through time until maturity of the underlying CDS. Recall that the driving noise processes of these hazard rates are independent of each

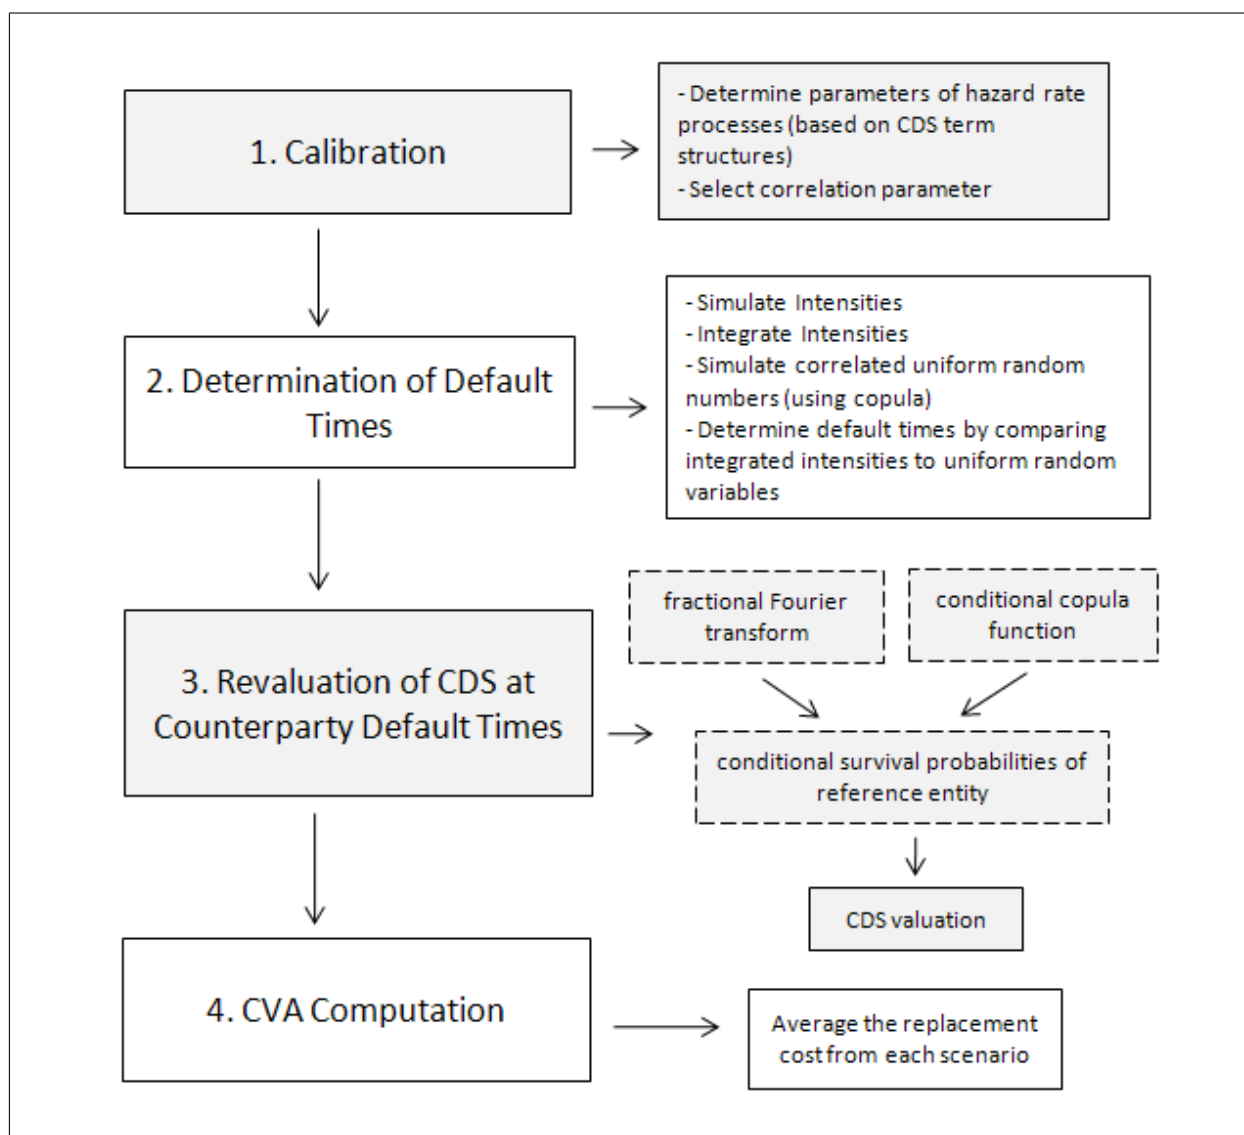


Figure 6.2: Diagram of the model outline

other. Once they have been simulated, they are integrated. The integrals are compared to the uniform random variables that are generated using the copula function and the default times of each of the entities determined. Note that, when the reference entity is the first name to default, the value of the CDS on default of the counterparty will be zero. This is due to the fact that the CDS ceases to exist upon a reference entity default event. Therefore, scenarios that contribute positively to the CVA are those in which the counterparty defaults prior to the reference entity and the CDS maturity.

Once the relevant counterparty default times have been obtained, the next stage is the

revaluation of the CDS at the relevant default times. This is step 3 in Figure 6.2. The main input required for the NPV computation is the survival curve of the reference entity at the counterparty default time (τ_2). Since the discount factor is assumed to be deterministic, the CDS valuation is trivial once these probabilities have been obtained.

A fractional fast Fourier transform (FRFT) (which is a semi-analytic numerical technique [53]) is employed, along with the copula function, in the computation of the survival probabilities. Note that these probabilities are conditional upon the counterparty default time τ_2 and the market information available at τ_2 . The term semi-analytic is used in the description of the FRFT to indicate that a full Monte Carlo simulation is unnecessary. However, an approximation of the integral remains a prerequisite, preventing the methodology from being completely analytical.

The final phase is the averaging of the replacement cost of the CDS in each scenario. Scenarios in which the NPV of the CDS is negative at the counterparty default time, or in which the counterparty does not default prior to the reference entity or the CDS maturity, provide a zero contribution to the CVA value. Scenarios in which the counterparty defaults prior to the reference entity and the NPV of the CDS is greater than zero at τ_2 contribute positively to the unilateral credit valuation adjustment.

6.5 Chapter Summary

The purpose of this chapter was to provide an introduction to the model implementation to follow in Chapters 7 and 8. Since the CVA on a credit default swap is the focus of the remainder of the dissertation, a short discussion on the media hype that surrounded CDSs in the aftermath of the 2008 credit crisis was presented. This was followed by a justification detailing the logic behind the selection of the CDS CVA pricing model. As preparation for the implementation, the basic notation required in subsequent chapters was then introduced, along with the main assumptions underlying the CDS CVA model. This was followed by an introduction to the CIR++ process and its applicability as a hazard rate modelling tool. The chapter was concluded with a brief outline of the model implementation steps. These will become clearer in the next two chapters in which a detailed description of each step is presented.

Chapter 7

Calibrating Intensities & Simulating Default Times

The initial two steps in the model outline in Figure 6.2 are the subject of this chapter. An explanation of the intensity process calibration is presented, prior to a description of the default time simulation procedure.

7.1 Calibrating the Intensity Process

It is assumed that the implied hazard rates of both the reference entity and the counterparty have been bootstrapped from the market-observable par CDS curves. The procedure for bootstrapping a credit spread term structure is described in Section B.1 in Appendix B by way of an example. Below, the calibration of the CIR++ process is presented for a generic entity whose hazard rate dynamics are given by:

$$dy(t) = \kappa(\mu - y(t))dt + \nu\sqrt{y(t)}dW(t), \quad y(0) = y_0 \quad (7.1)$$

$$\lambda(t) = y(t) + \psi(t). \quad (7.2)$$

The elements of the parameter vector $\beta = (\kappa, \mu, \nu, \psi, y_0)$ are positive constants. The generic calibration procedure is then applied to the reference entity and the counterparty hazard rate term structures in turn.

As an initial step, the following integrated quantities are defined:

$$\Lambda(t) = \int_0^t \lambda(s)ds, \quad \Psi(t, \beta) = \int_0^t \psi(s, \beta)ds \quad \text{and} \quad Y(t) = \int_0^t y(s)ds. \quad (7.3)$$

Since we wish to mark the model to market, the survival probabilities implied from market spreads are required to be reproduced by the model. This is analogous to requiring market observable bond prices to be equivalent to model bond prices in short rate models. Using Equation (4.23) in Chapter 4 and the fact that $\Lambda(t) = Y(t) + \Psi(t)$, the model survival probability is given by:

$$\mathbb{Q}(\tau > t)_{model} = \mathbb{E}(e^{-\Lambda(t)}) = \mathbb{E}(e^{-\Psi(t,\beta)-Y(t)}). \quad (7.4)$$

We are required to ensure that

$$\mathbb{Q}(\tau > t)_{model} = \mathbb{E}(e^{-\Psi(t,\beta)-Y(t)}) = e^{-\Lambda^{mkt}(t)}. \quad (7.5)$$

Since Ψ is deterministic, it can be factored out of the expected value sign. Taking logs on both sides of (7.5) leads to:

$$\Psi(t, \beta) = \Lambda^{mkt}(t) + \ln(\mathbb{E}(e^{Y(t)})) = \Lambda^{mkt}(t) + \ln(P^{CIR}(0, t, \beta)), \quad (7.6)$$

where $P^{CIR}(0, t, \beta)$ is the time 0 CIR bond price with maturity t and parameter vector β . The formula for the bond price within the CIR framework can be found in Section C.2 in Appendix C. The last equality in (7.6) follows from the definition of a bond price (survival probability in credit modelling terms) as the expected value of the exponential of the integrated short rate (hazard rate in credit modelling terms) process. Refer to [25] for a more detailed discussion of interest rate theory.

In order to obtain the expression for the deterministic function ψ , we differentiate (7.6) with respect to time to obtain:

$$\psi(t, \beta) = \lambda(t) - f^{CIR}(0, t) = f^M(0, t) - f^{CIR}(0, t), \quad (7.7)$$

where $f^M(0, t)$ is the instantaneous forward rate in the market at time 0 for time t and $f^{CIR}(0, t)$ is the instantaneous forward rate for the CIR process. Recall that the general definition of the instantaneous forward rate is given by $f(t, T) = -\partial \ln P(t, T) / \partial T$ [25]. The valuation formula for the CIR instantaneous forward rate can be found in Section C.1 in Appendix C.

In calibrating short rate models to the market, quoted swaption prices are utilised, along with the term structure of discount rates, to determine the values of the parameters β . In

the same vein, CDS option quotes are ideally required in addition to the term structure of hazard rates for a particular entity, in order to calibrate its CIR++ process. However, CDS option quotes are currently highly illiquid offshore and nonexistent in South Africa. To obtain the CIR++ parameters in the absence of these, we minimise the quantity $\int_0^T \psi(s, \beta)^2 ds$ [53]. This effectively amounts to minimising the amount by which the fitted model departs from the CIR model.

Additionally, the volatility parameter ν can be varied and the hypothetical CDS option volatilities implied by these shifts obtained. A user who is in touch with credit markets can then select the parameter that he or she believes renders the most realistic implied volatility level. There are also certain credit indices¹ on which options are traded. These options are generally illiquid, but could possibly be utilised to gain an indication of realistic implied volatility levels.

Historical CDS spread data is an alternative approach for obtaining the CIR++ hazard rate volatility parameter. The measure with which we are working is then the real world measure, rather than the risk neutral measure that has been assumed up to this point. The use of historical data in the CIR++ calibration procedure is discussed in the case study in Chapter 12.

Note that the following constraints are required to be satisfied in the minimisation of the distance between the market-implied rates and those of the fitted CIR++ process:

- (i) The parameters β must all be positive.
- (ii) The integral Ψ of the deterministic portion is required to be positive.
- (iii) The integral Ψ should be increasing.

The last two criteria ensure that the shift ψ is positive. By implication, the hazard rates produced by the CIR++ process are positive. Recall that the Feller condition, preventing y_t from attaining a zero value, is not imposed. The positivity of the deterministic shift ψ is thus a necessary condition to ensure that hazard rates remain greater than zero. The importance of positive intensities was discussed in Section 6.3.2.

¹A credit index expresses the spread on a basket of entities. CDSs can be traded on these indices and the market for them is often more liquid than for single names. The most well known of the indices are the iTraxx Europe investment grade and crossover indices and the CDX range of North American indices. For further information on these as well as additional indices available see [70] and [71].

7.1.1 Calibration Results

The Republic of South Africa USD denominated debt (SOAF) CDS spread term structure for close of business (COB) 31 August 2010 was introduced in Section B.1 in Appendix B. The results of bootstrapping this curve, along with the two term structures obtained by applying both a 100 basis point (bp) and a 400bp parallel shift to the 31 August spreads, were presented in Table B.2. The same three curves will be utilised in the numerical illustrations that follow. The base curve will be referred to as the low risk case, the 100bp shift scenario as the intermediate case and the 400bp shift scenario as the risky case. For ease of reference, these spreads are repeated in Table 7.1 below.

Term:	1	2	3	4	5	7	10
Low Risk (Base)	81	109	130	144	155	163	170
Intermediate Risk	181	209	230	244	255	263	270
High Risk	481	509	530	544	555	563	570

Table 7.1: CDS term structures for numerical illustrations. *Spreads are quoted in basis points.*

The CIR++ parameters corresponding to the three curves in Table 7.1 are obtained by minimising the distance between the fitted CIR hazard rates and those bootstrapped from the market. They are reported in Table 7.2 below.

Parameter	Fitted Value		
	Low Risk	Intermediate Risk	High Risk
κ	0.500	0.500	0.500
μ	0.026	0.039	0.080
ν	0.050	0.050	0.055
y_0	0.001	0.014	0.054

Table 7.2: CIR++ parameters

The quantities in (7.7), that are used to fit the CIR++ model to the base case curve, are illustrated in Figure 7.1. Linear interpolation is applied to the hazard rates bootstrapped from the market CDS term structure. The integrated quantities in (7.6) are illustrated in Figure 7.2. Notice that the CIR curve (y_t) is upward sloping. This is a result of the long run mean (μ) being larger than the initial value, y_0 . An illustration of the information conveyed in Figure 7.1 for a case in which y_0 is larger than μ can be found in Figure 12.2.

The latter is linked to the Lehman Brothers term structure the Friday prior to its default. Flat interpolation is applied to the Lehman Brothers hazard rates as opposed to the illustrations in this chapter that utilise linear interpolation.




Figure 7.1: Market-implied hazard rates, fitted CIR rates and deterministic shift for the base case (SOAF COB 31 August 2010). *Linear interpolation is applied to hazard rates. Quantities illustrated are those required by equation (7.7).*

7.2 Simulating the Default Times

Once the parameters of the CIR++ process have been obtained for both the reference entity and the counterparty, the default times of each of the entities can be determined. This involves simulating the hazard rates of each of these entities, integrating them and comparing them to the uniform random variables generated by means of the copula [25].

7.2.1 Simulating the Hazard Rates

The generation of the hazard rate, specified by the dynamics in (7.1) and (7.2), is a two-step procedure. The first step is the simulation of the CIR process, the stochastic portion of the CIR++ process. The second is the addition of the deterministic function ψ (see (7.7)) to each of the generated CIR paths. Note that since this function is merely




Figure 7.2: Integrated market-implied hazard rate, CIR rate and deterministic shift for the base case (SOAF COB 31 August 2010). *Linear interpolation is applied to hazard rates. Quantities illustrated are those required by equation (7.6).*

time-varying and not stochastic, the shift added to the CIR rates at each time point will be identical for all simulations.

Simulating the CIR process

The process in (7.1) can be simulated using either the conditional distribution of y or a discretisation scheme. In implementing the CDS CVA model, both options were considered for the sake of comparison. Although the discretisation scheme is generally quicker, the time saved on a 5 year CDS with 100,000 simulations and quarterly time steps is roughly only 30 seconds², not a significant cost for the sake of precision. In addition, employing a discretisation scheme requires finer time steps than when the transition density approach is utilised. This increase in the number of steps is likely to nullify the time saved by using the approximation.

The reason for the time difference between the two simulation schemes is that the transition density-based simulation involves sampling from both a chi-square and either a

²On a Dell Latitude E6400 laptop with a 2.4Ghz dual core processor. In practice, valuations would run on a more powerful computer with possible additions such as multi-threading.

normal or a Poisson distribution. The discretisation scheme, on the other hand, requires only normally distributed random numbers to be sampled. Both approaches are discussed below.

Transition Density Approach

The stochastic differential equation (SDE) for the CIR process cannot be solved analytically in the same way as one might solve the GBM process for example³. However, it is known that the distribution of $y(t)$ given $y(u)$, where $u < t$ is, up to a scale factor, a noncentral chi-square distribution [72]. The procedure for generating paths under the CIR dynamics in (7.1) is presented in Figure 7.3.

The simulation procedure when $4\mu\kappa/\nu^2 > 1$ rests on the representation of a noncentral chi-square distribution, with degrees of freedom parameter greater than one, as a function of an ordinary chi-square distribution and a standard normal distribution. When $4\mu\kappa/\nu^2 \leq 1$, the noncentral chi-square distribution may be sampled by generating a Poisson random variable and then sampling from a standard chi-square random variable. The simulated Poisson random variable features in the degrees of freedom computation of the chi-square random variable that is sampled. Further detail can found in [72].

Discretisation Schemes

The Euler and Milstein schemes are two popular methods used for discretising SDEs. The focus here is on Euler-based schemes as numerical tests have found that Milstein approximations are typically less robust than Euler schemes due to the square root in the CIR SDE. This causes certain sufficient regularity conditions to be violated [73].

The drawback of the basic Euler scheme is that, even when the Feller condition ($2\kappa\mu > \nu^2$) is satisfied, there is no guarantee that simulated values will be positive. Brigo and Alfonsi have derived an implicit scheme, called the *Euler implicit positivity-preserving scheme*, that guarantees that the simulated value of $y(t)$ will be positive when the Feller condition holds. The derivation of this can be found in [74]. We do not dwell on the detail since the availability of the transition density renders discretisation schemes less relevant. Utilising the Euler implicit positivity-preserving scheme, the expression for $y_{t_{i+1}}$ given y_{t_i} , with the

³It is easy to show that if the dynamics of S are given by $dS_t = \mu S_t dt + \sigma S_t dW_t$ with initial condition $S(0) = S_0$, the value of $S(t)$ is given by $S_t = S_0 e^{(\mu - \frac{\sigma^2}{2})t + \sigma\sqrt{t}\varepsilon}$, $\varepsilon \sim N(0, 1)$ [23].

On time grid $0 = t_0 < t_1 < \dots < t_n$ with $d = 4\mu\kappa/\nu^2$

Case 1: $d > 1$
 for $i = 0 : n - 1$
 set $c = \nu^2(1 - e^{-\kappa(t_{i+1}-t_i)})/(4\kappa)$
 set $\phi = y(t_i)(e^{-\kappa(t_{i+1}-t_i)})/c$
 generate $Z \sim N(0, 1)$
 generate $X \sim \chi_{d-1}^2$
 set $y(t_{i+1}) = c[(Z + \sqrt{\phi})^2 + X]$
 end

Case 2: $d \leq 1$
 for $i = 0 : n - 1$
 set $c = \nu^2(1 - e^{-\kappa(t_{i+1}-t_i)})/(4\kappa)$
 set $\phi = y(t_i)(e^{-\kappa(t_{i+1}-t_i)})/c$
 generate $N \sim \text{Poisson}(\phi/2)$
 generate $X \sim \chi_{d+2N}^2$
 set $y(t_{i+1}) = cX$
 end

Figure 7.3: Simulation of CIR process using transition density. *Source: Monte Carlo Methods in Financial Engineering, P. Glasserman [72].*

time grid the same as in Figure 7.3 and $2\kappa\mu > \nu^2$, is as follows:

$$\left(\frac{\nu(W_{t_{i+1}} - W_{t_i}) + \sqrt{\nu^2(W_{t_{i+1}} - W_{t_i})^2 + 4(y_{t_i} + (\kappa\mu - \frac{\nu^2}{2})(t_{i+1} - t_i))(1 + \kappa(t_{i+1} - t_i))}}{2(1 + \kappa(t_{i+1} - t_i))} \right)^2.$$

In the event that the Feller condition does not hold ($2\kappa\mu \leq \nu^2$) the reflected Euler scheme [25] can be employed. The expression for $y_{t_{i+1}}$ is then:

$$| y_{t_i} + \kappa(\mu - y_{t_i})(t_{i+1} - t_i) + \nu\sqrt{y_{t_i}}(W_{t_{i+1}} - W_{t_i}) |. \quad (7.8)$$

Given the presence of the absolute value sign in (7.8) and thus a departure from the original SDE for the sake of ensuring positivity, it is clearly preferable to utilise the transition density approach in cases in which the Feller condition is not satisfied. The expression in (7.8) is shown for the sake of completeness.

The transition density approach has been employed in simulating the hazard rates in all computations in the remainder of the dissertation. The valuation formulae for the

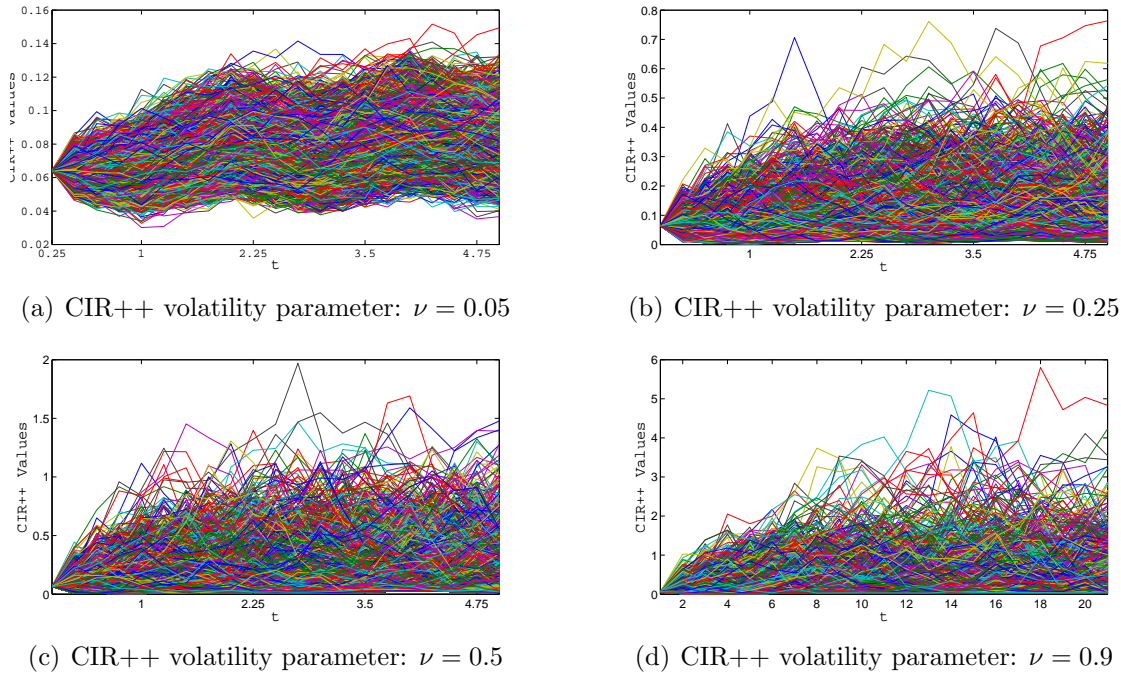


Figure 7.4: 10,000 simulated CIR++ paths for the high risk entity

conditional expectation and variance of the CIR process can be found in Sections C.3 and C.4 in Appendix C respectively. The mean and variance of the simulated CIR values were compared to these as a validation of the simulation procedure.

Once the simulated CIR values have been obtained as detailed above, the deterministic shift is added, resulting in the simulated intensities for the CIR++ process. Figure 7.4 shows simulated paths of the CIR++ process of the high risk entity in Table 7.1 for various values of the volatility parameter ν . Notice how the range of values attained increases with volatility. Furthermore, the quantity of values close to zero increases. The expected value of the process is not a function of the volatility and is required to remain the same regardless of the value of ν . Since the process cannot produce negative hazard rates, the increase in extreme positive values when the volatility is increased cannot be offset by an increase in extreme negative values. There is therefore a greater concentration of values close to zero to maintain the expected value when ν is increased. Note that, although the hazard rates in Figure 7.4(d) appear to reach zero, the minimum value in the graph is 0.0087. The mean reverting nature of the process is clearly visible, particularly in Figure 7.4(a) in which the hazard rate values are less volatile and the shape is more easily observed than for higher volatilities.

7.2.2 Integrating intensities

In order to determine the default times of the two entities, the integrated hazard rates (Λ_1 and Λ_2 in (6.1)) are required. The interpolation method (linear or piecewise constant) that was used to bootstrap the hazard rates from the original CDS term structures and to fit the CIR++ parameters, is utilised in the integration. The time steps at which the integral is calculated are equivalent to the hazard rate simulation points.

7.2.3 Determining τ_1 and τ_2

Recall from Section 6.3.1 that the integrals of the default times, the quantities

$$\Lambda_1(\tau_1) \text{ and } \Lambda_2(\tau_2), \quad (7.9)$$

are standard exponential random variables. Recall also that the variables

$$U_1 = 1 - e^{-\Lambda_1(\tau_1)} \text{ and } U_2 = 1 - e^{-\Lambda_2(\tau_2)}, \quad (7.10)$$

are uniformly $[0,1]$ distributed. Furthermore, the copula $C(U_1, U_2)$ in (6.5) connects the variables U_1 and U_2 .

A summary of the basic theory of copulas is presented in Appendix D. Since the subject of this dissertation is not copulas, only the concepts pertinent to the CDS CVA model implementation, as well as some general background, are discussed in the appendix. The book by Nelson [54] is a valuable reference on the subject and is written in a manner that is accessible to readers with a background in statistics. A shorter reference, that is particularly useful in simulating random numbers using copulas, is a paper aptly entitled *Coping with Copulas* [75].

As mentioned previously, the copula selected in this implementation is the Gaussian copula, although the model can facilitate any differentiable alternative. For the bivariate case, the Gaussian copula is defined as [25]:

$$C(u, v) = \int_{-\infty}^{\Phi^{-1}(u)} \int_{-\infty}^{\Phi^{-1}(v)} \frac{1}{2\pi(1-\rho^2)^{1/2}} e^{-\frac{s^2-2\rho st+t^2}{2(1-\rho^2)}} ds dt, \quad (7.11)$$

where Φ^{-1} is the inverse of the standard normal cumulative distribution function. Readers familiar with statistics should recognise the form of the bivariate normal cumulative dis-

tribution function in (7.11). The bivariate version of the copula is utilised since unilateral CVA is being measured. The single correlation parameter that is relevant is between the uniform random variables associated with the default times of the reference entity and the counterparty. In the extension to the BCVA model in Part III, a trivariate copula function is required.

An illustration of the probability density function (PDF) of the bivariate Gaussian copula is presented in Figure 7.5 for various values of the correlation parameter, ρ . Notice how the shape of the PDF is completely level in the zero correlation case in Figure 7.5(a). As the correlation increases, the PDF values in the quadrants in which the sizes of the two uniform random variables (u and v) are the same, increase. The opposite occurs in the remaining quadrants, with the PDF being virtually zero for extreme opposite values of u and v in Figure 7.5(d). Here the correlation is 0.9. For the definition of a copula PDF see Section D.5 in Appendix D. The cumulative distribution function (CDF) of the bivariate Gaussian copula (7.11) with correlation parameter 0.5 is illustrated in Figure 7.2.3.

Proceeding now with the determination of the default times, we are required to simulate a set of correlated uniform random numbers from the Gaussian copula. These will be transformed into exponential random variables and compared to the integrated intensities in order to determine the default times of each of the entities. The copula random variable generation is a two step procedure. The first is to generate correlated random numbers from the standard normal distribution. The second is to transform these into uniform random variables utilising the cumulative distribution function of the normal distribution. The methodology is presented in Figure 7.7. The Matlab *copularnd* function can be used to carry out the steps in the figure.

Once the uniform random numbers have been generated by means of the copula, it is necessary to transform them into exponential random variables. This is to facilitate a comparison between them and the exponentially distributed integrated intensities. To obtain exponential random variables from the uniform numbers, the inverse of Equation (7.10) above is applied:

$$\Lambda_j(\tau_j) = -\ln(1 - U_j), \quad j = 1, 2. \quad (7.12)$$

It now remains to compare the generated exponential random variables to the integrated intensities that were obtained in Section 7.2.2. Default occurs when the exponential ran-

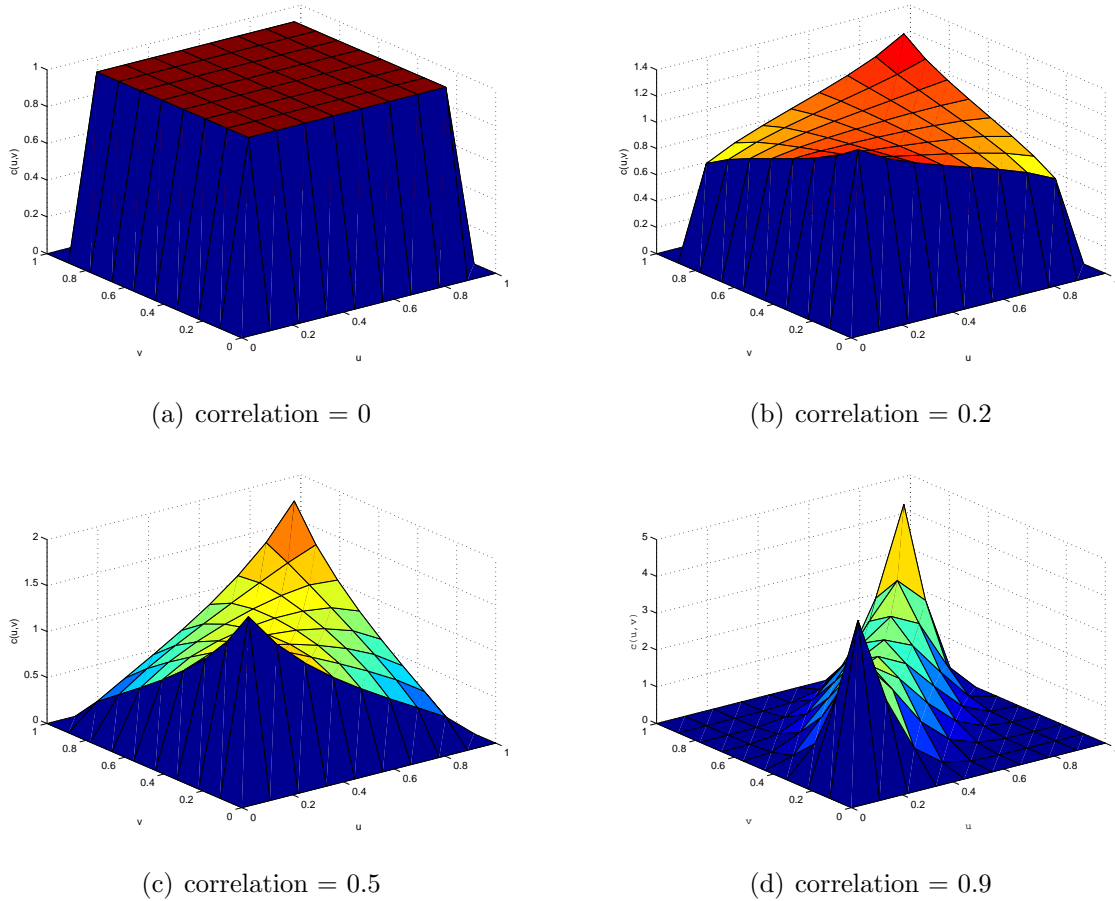


Figure 7.5: PDF of bivariate Gaussian copula for various correlations

dom variable generated by means of the copula is greater than or equal to the integrated intensity obtained by simulating from the CIR++ process and integrating.

As discussed, only scenarios in which τ_1 is greater than τ_2 (the counterparty defaults prior to the reference entity) are significant for calculating the NPV of the CDS on counterparty default. Additionally, it is obvious that τ_2 is required to occur prior to the underlying CDS maturity in order for the counterparty default to be relevant⁴.

Once the relevant counterparty default times have been obtained, the assumption is made that the default event occurs at the next date in the payment schedule in the event of a default occurring between dates in the CDS premium payment schedule. The same approximation was made by Brigo and Chourdakis [53].

⁴The CDS no longer has a replacement cost once it has expired.




Figure 7.6: CDF of the Gaussian copula - correlation parameter set to 0.5.

```

for N simulations

Step 1
  Generate  $x_1 \sim N(0, 1)$  and  $x_3 \sim N(0, 1)$  independently.
  Set  $x_2 = \rho x_1 + \sqrt{1 - \rho^2} x_3$ .

Step 2
  Set  $u_1 = \Phi(x_1)$  and  $u_2 = \Phi(x_2)$ 
  Output  $u_1$  and  $u_2$ 


end
  
```

Figure 7.7: Simulation of uniform random variables from Gaussian copula


In order to provide an intuitive understanding of the concept of default time correlation, the ratio of the low risk entity's default times to those of the high risk entity's are illustrated in Figure 7.8. The two extreme cases of independence and of a default time

correlation of 0.99 are presented in Figures 7.8(a) and 7.8(b) respectively. A hazard rate volatility (ν) of 0.1 was assigned to each of the entities in the simulation. Only those simulations in which both entities defaulted are illustrated on the graphs.

Figure 7.8(a) clearly demonstrates that when the default time correlation is zero, the results are more random. In general, the high risk entity is more likely to default first since it has a higher hazard rate. The percentage of scenarios illustrated in Figure 7.8(a) in which the low risk entity is first to default is higher than in Figure 7.8(b). The ratio of the default times is less than one in such cases. Furthermore, the general level of the ratio of default times is greater for the high correlation case. The final observation, regarding Figure 7.8, is that the number of paths illustrated when the correlation parameter is zero is lower than when it is 0.99. This can be seen by observing the maximum value along the x-axis in both graphs. The reason for this is that there are more paths along which both entities default in the high correlation case. This is despite the fact that the probabilities of each of the entities defaulting remains constant when correlation is varied. The total number of simulations that were run in generating each of the graphs in Figure 7.8 is 100,000.



(a) correlation = 0



(b) correlation = 0.99

Figure 7.8: Ratio of default times of low risk entity to those of high risk entity for simulation paths on which both names default

7.2.4 Reproducing the Market Survival Curve

In order for the calibration and default time simulation procedures to be useful, the survival probabilities observable in the market should be reproduced by the model. For each of the three levels of riskiness in Table 7.1, 100,000 paths of the hazard rates were simulated for a set of values of the volatility parameter, ν . Default times were obtained in the manner outlined above. The survival probability was then computed at each tenor on the CDS curve by counting the number of defaults that had occurred at each time point, dividing this by the total number of simulations to get the default probability and subtracting it from one in order to obtain the survival probability.

The results are presented in Table 7.3. The red values in the table indicate that the simulated survival probability differs from the one obtained from the market by a percentage point. The maximum difference is 1%, indicating that the default time simulation procedure is accurate.

Term:		1	2	3	4	5	7	10
Riskiness	ν	Survival Probability						
Low	Exact	99%	97%	95%	93%	90%	86%	80%
	0.01	99%	97%	95%	93%	90%	86%	79%
	0.05	99%	97%	95%	93%	90%	86%	80%
	0.1	99%	97%	95%	93%	90%	86%	80%
	0.5	99%	97%	95%	92%	90%	85%	80%
	0.9	99%	97%	95%	92%	90%	85%	80%
Intermediate	Exact	98%	95%	91%	88%	84%	78%	69%
	0.01	98%	95%	91%	88%	84%	78%	69%
	0.05	98%	94%	91%	88%	84%	78%	69%
	0.1	98%	95%	91%	88%	84%	78%	69%
	0.5	98%	94%	91%	88%	84%	78%	69%
	0.9	98%	94%	91%	88%	84%	78%	69%
High	Exact	94%	87%	81%	75%	69%	59%	46%
	0.01	94%	87%	81%	75%	69%	59%	46%
	0.05	94%	87%	81%	75%	69%	59%	46%
	0.1	94%	87%	81%	75%	69%	59%	46%
	0.5	94%	87%	81%	74%	69%	59%	46%
	0.9	94%	87%	81%	75%	69%	59%	47%

Table 7.3: Comparison of exact and simulated survival probabilities for each entity in the numerical illustration. *Simulated values were obtained by simulating default times of each name and counting the number of defaults at each time step. Survivals were obtained by subtracting the percentage of defaults from one. Number of Simulations = 100,000.*

The majority of CDS contracts entail quarterly premium payments. In addition, the model implemented in this dissertation assumes that counterparty defaults occur only on dates in the premium payment schedule. An example of the exact and simulated probabilities with time steps of 0.25 years is illustrated in Figure 7.2.4.




Figure 7.9: Exact and simulated survival probabilities for each entity. $\nu = 0.05$. Quarterly time steps are illustrated.

7.3 Chapter Summary

The calibration of the CIR++ process, which was introduced in the previous chapter, was described at the beginning of this one. An explanation of the way in which the results obtained from the calibration procedure are utilised in determining the default times of both the counterparty and the reference entity was then presented. The steps required to obtain the relevant default times of the counterparty from the calibrated hazard rate processes are:

- (i) Simulate K paths of both λ_1 and λ_2 .
 - (a) Simulate K paths of both y_1 and y_2 from the CIR processes of the reference entity and the counterparty respectively.

-
- (b) Add the deterministic shifts (ψ_1 and ψ_2) to each of the CIR paths of the reference entity and counterparty respectively.
 - (ii) Integrate the simulated intensities to obtain Λ_1 and Λ_2 .
 - (iii) Generate two $K \times 1$ vectors of uniform random numbers by means of the copula.
 - (iv) Transform the uniform random numbers to exponential random numbers.
 - (v) Compare these exponential random numbers to the integrated intensities to determine whether and at which times either entity has defaulted along each of the K paths.
 - (vi) Determine where $\tau_2 < \tau_1$ and $\tau_2 < T$. These are the relevant counterparty default times.
 - (vii) Assume that default occurs at the date in the premium payment schedule immediately after the default time if it does not fall on a premium payment date.

The chapter was concluded with a comparison of the survival probability curves obtained from the market and from the default time simulation methodology detailed above. The results demonstrated accuracy in the calibration and default time simulation procedures.

Chapter 8

Revaluation of the CDS upon Counterparty Default

The previous chapter described the execution of Steps 1 and 2 in the diagram in Figure 6.2. The subject of this chapter is the implementation of Steps 3 and 4, the revaluation of the CDS at the relevant counterparty default times and the final determination of the CVA. This, in turn, requires knowledge of the term structure of the reference entity's survival probabilities at each of these default times.

This chapter commences with a derivation of the reference entity survival probability into a form that permits its computation by means of numerical methods in Section 8.1. The components of this derived formulation are then manipulated into expressions that are of practical use. Some of the formulae obtained are not utilised immediately, but are derived in preparation of the implementation discussion to follow. The subsequent section, Section 8.2, addresses the execution of the fractional fast Fourier transform (FRFT). This is employed in the computation of the cumulative distribution of the integrated CIR process, a component of the survival probability curve calculation.

Section 8.3 describes how the output of the FRFT, the results derived in the first section and the copula function are used to compute the term structure of the reference entity's survival probabilities at τ_2 . These are required to value the CDS at each of the relevant simulated counterparty default times. It is straightforward to obtain the CVA once these probabilities and consequently, the NPV on counterparty default, have been computed. The exact derivation of the spread applied to the par CDS premium to adjust for counterparty risk is presented in Section 8.4. As usual, the chapter is concluded with a summary

of the information that has been supplied therein.

8.1 Survival Probability Preliminaries

The results in this section were set forth in [53]. Recall the General Counterparty Risk Pricing Formula which was introduced in Section 3.2 (Equation (3.1)). The CVA term is:

$$\text{CVA}(t) = \text{LGDE} \mathbb{E}_t[1_{\{t < \tau \leq T\}} D(t, \tau) (\text{NPV}(\tau))^+].$$

Since the underlying instrument in this implementation is a CDS, the required CVA is:

$$\text{CVA}(t) = \mathbb{E}_t[1_{\{t < \tau_2 \leq T\}} D(t, \tau_2) (\text{CDS}(\tau_2, S, T))^+], \quad (8.1)$$

where $\text{CDS}(t, S, T)$ is the time t value of a CDS traded at a spread of S and expiring at time T . The counterparty default time is indicated by τ_2 in (8.1) rather than by τ . This permits differentiation between the reference entity and counterparty default times.

Let $\text{CDS}(\tau_2, S, T) = 1_{\{\tau_1 > \tau_2\}} \overline{\text{CDS}}(\tau_2, S, T)$. Then the expected value in (8.1) can be rewritten as:

$$\text{CVA}(t) = \mathbb{E}_t[1_{\{t < \tau_2 \leq T\}} D(t, \tau_2) 1_{\{\tau_1 > \tau_2\}} (\overline{\text{CDS}}(\tau_2, S, T))^+]. \quad (8.2)$$

The indicator functions ($1_{\{t < \tau_2 \leq T\}}$ and $1_{\{\tau_1 > \tau_2\}}$) in (8.2) are incorporated into the calculation by means of the default time simulation procedure in the previous section. Recall that only counterparty default times preceding both reference entity defaults and the CDS maturity were considered relevant.

In order to compute the NPV of the CDS, the probability of the reference entity surviving from the counterparty default time (τ_2) to each of the remaining dates in the premium payment schedule is required. Let T_k be one such remaining payment time. The probability whose computation is necessary is therefore:

$$\begin{aligned} \mathbb{Q}(\tau_1 > T_k | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2) &= \mathbb{Q}(\Lambda_1(\tau_1) > \Lambda_1(T_k) | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2) \\ &= \mathbb{Q}(1 - e^{-\Lambda_1(\tau_1)} > 1 - e^{-\Lambda_1(T_k)} | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2) \\ &= \mathbb{Q}(U_1 > 1 - e^{-Y_1(T_k) - \Psi_1(T_k, \beta)} | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2). \end{aligned} \quad (8.3)$$

The filtration and symbols that feature in the derivation of (8.3) were defined in Sections 6.3.1 and 7.1. Recall that the filtration \mathcal{G}_{τ_2} contains information pertaining to the

entire market, including credit and defaults, at time τ_2 . The filtration will thus uniquely determine the default time τ_2 and hence the value of U_2 , since λ_2 is measurable with respect to \mathcal{G} . The quantity $\Lambda_1(\tau_2)$ is also measurable and contained in the filtration [53].

The reason for expressing the probability in (8.3) as a function of U_1 is that the copula (through which correlation between the default times of the entities is introduced) was defined as a function of the random variables U_1 and U_2 . Note also that the survival probability at τ_2 is a **conditional** survival probability. It is conditional on the default time τ_2 and the market information available at that time. Additionally, it incorporates the dependence between the uniform random variables of the exponential triggers of the default times of the two entities (the reference entity and the counterparty). A positive default time correlation is expected to lead to a reduction in the survival probabilities of the reference entity at τ_2 (subsequent to the counterparty having defaulted). Similarly, an increase in these survivals is expected when the correlation is negative. Therefore, the survival probabilities are reactive with respect to the counterparty default times and the market observables at those times. This feature captures the right and wrong way risks inherent in CDS CVA.

A practical formulation of the expression for the conditional survival probability in (8.3) is derived in the next subsection. Subsequently, attention is focused on each of the two factors of which this expression is comprised.

8.1.1 A Practical Formulation of the Survival Probability

Define the following:

$$P(u_1) := \mathbb{Q}(u_1 > 1 - e^{-Y_1(T_k) - \Psi_1(T_k, \beta)} | \mathcal{G}_{\tau_2}). \quad (8.4)$$

Using (8.4), the probability in (8.3) is rewritten as:

$$E[P(U_1) | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2]. \quad (8.5)$$

The expectation in (8.5) follows from the definition of conditional probability as presented in [76]. Having expressed the survival probability in the form of a conditional expectation, it is formulated as an integral and can thus be evaluated using numerical integration. The Law of the Unconscious Statistician [77], which expresses the expected value of a function as a Riemann-Stieltjes integral (see [78] for an introduction to these integrals), permits

the following form of the survival probability:

$$\mathbb{Q}(\tau_1 > T_k | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2) = E[P(U_1) | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2] = \int_{u_1=\bar{U}_1} P(u_1) dC_{1|2}(u_1, U_2), \quad (8.6)$$

where the conditional distribution of the uniform random variables is defined by

$$C_{1|2}(u_1, U_2) := \mathbb{Q}(U_1 < u_1 | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2). \quad (8.7)$$

The variable \bar{U}_1 is the uniform random variable associated with the reference entity's integrated hazard rate at τ_2 . It is defined formally in (8.12) below. In order to evaluate the integral in (8.6) using numerical methods, the values of both $P(u_1)$ and $C_{1|2}$ for a set of u_1 's are required. Below, each of these two quantities is examined in turn.

A Convenient Expression for $P(u_1)$

Consider first $P(u_1)$. Basic algebra applied to (8.4) leads to the following:

$$\begin{aligned} P(u_1) &= \mathbb{Q}(u_1 > 1 - e^{-Y_1(T_k) - \Psi_1(T_k)} | \mathcal{G}_{\tau_2}) & (8.8) \\ &= \mathbb{Q}(1 - u_1 < e^{-Y_1(T_k) - \Psi_1(T_k)} | \mathcal{G}_{\tau_2}) \\ &= \mathbb{Q}(\ln(1 - u_1) < -Y_1(T_k) - \Psi_1(T_k) | \mathcal{G}_{\tau_2}) \\ &= \mathbb{Q}(Y_1(T_k) + \ln(1 - u_1) < -\Psi_1(T_k) | \mathcal{G}_{\tau_2}) \\ &= \mathbb{Q}(Y_1(T_k) - Y_1(\tau_2) < -\ln(1 - u_1) - Y_1(\tau_2) - \Psi_1(T_k) | \mathcal{G}_{\tau_2}). & (8.9) \end{aligned}$$

The reason for writing $P(u_1)$ in this way is to be able to use a fractional fast Fourier transform (FRFT), which is a semi-analytic method, to calculate the probabilities. To be able to apply the FRFT, the characteristic function of the variable under consideration is required. Since this is known for the integrated CIR process; in other words for

$$Y_1(T_k) - Y_1(\tau_2) = \int_{\tau_2}^{T_k} y_1(t) dt, \quad (8.10)$$

it is convenient to write the expression for $P(u_1)$ in the way we have done in (8.9). The details of the FRFT methodology are discussed in Section 8.2. The alternative would be to use Monte Carlo simulation, which would be significantly more time-consuming since, for each simulation in which the counterparty defaults prior to the reference entity, a new set of paths for the CIR++ process would be required.

A Convenient Expression for $C_{1|2}$

Proceeding with the conditional distribution that was defined in (8.7), the following derivation is performed:

$$C_{1|2}(u_1, U_2) := \mathbb{Q}(U_1 < u_1 | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2) = \mathbb{Q}(U_1 < u_1 | U_2, U_1 > \bar{U}_1), \quad (8.11)$$

where

$$\bar{U}_1 := 1 - e^{-Y_1(\tau_2) - \Psi_1(\tau_2, \beta)}. \quad (8.12)$$

Therefore \bar{U}_1 is the value, at τ_2 , of the uniform random variable associated with the reference entity. The last equality in (8.11) is a result of the following two observations:

- (i) Since the counterparty default time (τ_2) is known, the value of

$$U_2 := 1 - e^{-Y_2(\tau_2) - \Psi(\tau_2, \beta)}$$

is also known.

- (ii) The inequality $\tau_1 > \tau_2$ implies:

$$U_1 = 1 - e^{-Y_1(\tau_1) - \Psi(\tau_1, \beta_1)} > 1 - e^{-Y_1(\tau_2) - \Psi(\tau_2, \beta_1)} = \bar{U}_1. \quad (8.13)$$

The result from Bayesian statistics (which states that if A and B are two elements from the same probability space, the conditional probability of A given B is

$$P(A|B) = \frac{P(A \cap B)}{P(B)},$$

where the symbol \cap represents the intersection of events A and B) is then applied to the result in (8.11). For the condition, $U_1 > \bar{U}_1$,

$$\begin{aligned} C_{1|2}(u_1, U_2) &= \mathbb{Q}(U_1 < u_1 | U_2, U_1 > \bar{U}_1) \\ &= \frac{\mathbb{Q}(U_1 < u_1, U_1 > \bar{U}_1 | U_2)}{\mathbb{Q}(U_1 > \bar{U}_1 | U_2)} \\ &= \frac{\mathbb{Q}(U_1 < u_1 | U_2) - \mathbb{Q}(U_1 < \bar{U}_1 | U_2)}{1 - \mathbb{Q}(U_1 < \bar{U}_1 | U_2)}. \end{aligned} \quad (8.14)$$

The numerator in (8.14) is obtained from the numerator in the previous step by expressing differently the fact that U_1 has a lower bound of \bar{U}_1 and an upper bound of u_1 .

The conditional copula function in (8.14) is computable since the copula $C(u_1, u_2) = \mathbb{Q}(U_1 < u_1, U_2 < u_2)$ exists. The exact computation of the terms in this conditional probability function is examined in Section 8.3.1. The aim of this section was merely to manipulate the distribution into a form that is practically implementable.

Note that, utilising the definition of the conditional distribution of copulas (see Section D.6 in Appendix D), the fraction in (8.14) can be rewritten as:

$$C_{1|2}(u_1, U_2) = \frac{\frac{\partial}{\partial u_2} C(u_1, U_2) - \frac{\partial}{\partial u_2} C(\bar{U}_1, U_2)}{1 - \frac{\partial}{\partial u_2} C(\bar{U}_1, U_2)}. \quad (8.15)$$

Additionally, this form can be useful when the partial derivative of the copula is known in closed form. The independence copula ($C(U_1, U_2) = u_1 u_2$) is one such example.

An integral, which can be evaluated numerically for the computation of the conditional survival probability of the reference entity at τ_2 (the counterparty default time), has now been derived. The integrand has been manipulated into a form that permits the application of an FRFT. Additionally, the conditional distribution of the uniform random variable associated with the reference entity (given the filtration \mathcal{G}_{τ_2} and the information contained therein) has been expressed in a way that permits the use of the copula function that was defined in Section 6.3.2. The application of the FRFT is discussed in the next section. The description is quite lengthy as it aims to provide the reader with an exact understanding of the implementation. Thereafter the CDS valuation will be described and the derivations pertaining to the conditional copula function will be required.

8.2 Applying the Fractional Fourier Transform

We aim to calculate the probability in (8.6). Using (8.9), together with basic calculus and trapezoidal integration, the following approximation (the exact derivation of which is presented in Section E.1 of Appendix E) is obtained:

$$\mathbb{Q}(\tau_1 > T_k | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2) \approx \sum_j \frac{p_{j+1} + p_j}{2} \Delta f_j, \quad (8.16)$$

where

$$p_j = \mathbb{Q}(X < x_j | \mathcal{G}_{\tau_2}) \text{ for a grid } x_j; \quad X := Y_1(T_k) - Y_1(\tau_2)$$

and we have utilised a shortened notation for the conditional copula function, that is:

$$f_j := C_{1|2}(u_j, U_2).$$

For each simulation in which $\tau_2 < \tau_1$ and $\tau_2 < T$, we know the value of τ_2 which, as discussed earlier, will coincide with a date in the premium payment schedule. In order to be able to revalue the underlying CDS at the default times, we require knowledge of the conditional survival probability of the reference entity from τ_2 to each of the remaining dates in the premium payment schedule. As mentioned earlier, we call each of these remaining payment times T_k . The variable T_k can assume a minimum value equal to the second premium payment date and a maximum value of T , the CDS maturity. The minimum T_k in each simulation in which a relevant counterparty default occurs, will depend on the value of τ_2 in that simulation. As an example, consider a quarterly compounding CDS and assume, for simplicity, that the length of each premium period is precisely 0.25 years. If (on a certain simulation path) the counterparty default occurs at time 0.75, the values of T_k for which we require the probabilities in (8.16), are 1, 1.25, 1.5, ..., 4.75, 5.

The aim of this section on fractional fast Fourier transforms is to explain how the probabilities p_j in (8.16) are produced. We are thus concerned with the calculation of the cumulative distribution function (CDF) of the random variable $X = Y_1(T_k) - Y_1(\tau_2)$. In other words, we seek a grid of probabilities $p_j = \mathbb{Q}(X < x_j), j = 1, 2, \dots, N$ where each p_j is associated with an x_j . To do this, we use a fractional fast Fourier transform (FRFT). A brief background to Fourier transforms, fast Fourier transforms and fractional fast Fourier transforms is presented in Appendix F. Only results required in the implementation of this model are presented. The reader is encouraged to read it as the terminology explained there is required for an understanding of the remainder of this section.

8.2.1 Writing the CDF in the Form of an FRFT Sum

The expression for the CDF of a nonnegative random variable X , in terms of the characteristic function $\phi(u)$ of the random variable, is obtained via Fourier inversion. It can be found in [79] or [80] and is given by:

$$F(x) = \mathbb{P}(X \leq x) = \frac{2}{\pi} \int_0^{\infty} \text{Re}(\phi(u)) \frac{\sin ux}{u} du, \quad x > 0. \quad (8.17)$$

Recall that we aim at producing a grid of x 's with associated cumulative probabilities. We use a subscript k for each x where $k = 1, 2, \dots, N$. For $x = x_k$, we aim to rewrite (8.17) in a

form that can accommodate the application of the FRFT. As a first step, we approximate the integral in (8.17) as follows:

$$\frac{2}{\pi} \int_0^\infty \frac{\text{Re}(\phi(u))}{u} \sin(ux_k) du \approx \frac{2}{\pi} \sum_{j=0}^{N-1} \frac{\text{Re}(\phi(u_j))}{u_j} \sin(u_j x_k) \delta w_j. \quad (8.18)$$

where δ is the size of the u step and w_j is the j 'th weight. We have utilised trapezoidal integration in the implementation. The weights in (8.18) would thus be:

$$w_1 = 1, w_2 = 2, w_3 = 2, \dots, w_{N-1} = 2, w_N = 1.$$

Alternative integration schemes were considered. However, tests demonstrated that these did not enhance the accuracy of the numerical approximation. Note that the variable u , utilised in the characteristic function, is not the uniform random variable considered earlier. Since much of the literature employs $\phi(u)$ to represent the characteristic function, we do the same. The meaning of u should be clear from the context.

If we now let the size of the x -step be λ , we can rewrite (8.18) as

$$\frac{2}{\pi} \delta \sum_{j=0}^{N-1} \frac{\text{Re}(\phi(\delta j))}{\delta j} \sin(\delta j \lambda k) w_j. \quad (8.19)$$

Using the well known fact that, for a random variable z ,

$$\sin z = \frac{e^{iz} - e^{-iz}}{2i},$$

we obtain the following:

$$\frac{2}{\pi} \delta \sum_{j=0}^{N-1} w_j \frac{\text{Re}(\phi(\delta j))}{\delta j} \left[\frac{e^{ij\delta\lambda k} - e^{-ij\delta\lambda k}}{2i} \right]. \quad (8.20)$$

This can be rewritten as

$$\frac{2}{\pi} \delta \left(\sum_{j=0}^{N-1} w_j \frac{\text{Re}(\phi(\delta j))}{\delta j 2i} e^{-i2\pi k j \alpha_1} - \sum_{j=0}^{N-1} w_j \frac{\text{Re}(\phi(\delta j))}{\delta j 2i} e^{-i2\pi k j \alpha_2} \right), \quad (8.21)$$

where $\alpha_1 = -\delta\lambda/2\pi$ and $\alpha_2 = \delta\lambda/2\pi$. Comparing the expression in (8.21) to the form of the FRFT sum in Equation (F.8) in Appendix F, it is clear that we can apply two FRFT

processes to obtain the CDF of the variable X . This can be done by letting the function h_j in (F.8) be the following:

$$h_j = w_j \frac{Re(\phi(\delta j))}{\delta j 2i} \quad (8.22)$$

in both terms in (8.21).

8.2.2 The Characteristic Function

We are now required to apply the above result to the integrated CIR process, the characteristic function of which is found on page 89 of [81]. The integrated CIR process $Y = (Y_t, t \geq 0)$ is given by:

$$Y_t = \int_0^t y_s ds. \quad (8.23)$$

Given y_0 (the initial value of the CIR process), the characteristic function of (8.23) is as follows:

$$\phi(u) = \frac{e^{\kappa^2 \mu t / \nu^2} e^{2y_0 i u / (\kappa + \gamma \coth(\gamma t / 2))}}{(\cosh(\gamma t / 2) + \kappa \sinh(\gamma t / 2) / \gamma)^{2\kappa \mu / \nu^2}}, \quad (8.24)$$

where

$$\gamma = \sqrt{\kappa^2 - 2\nu^2 i u}.$$

The dynamics of y were presented in Equation (7.1) of Section 7.1.

8.2.3 A Trap

Unfortunately, there is a small problem with the characteristic function in the form in which it appears in the literature. The expression in (8.24) is substituted into the integrand in (8.17), which is:

$$Re(\phi(u)) \frac{\sin ux}{u}. \quad (8.25)$$

An investigation into the behaviour of the integrand follows. We begin this discussion with a plot of the function being integrated, assuming various values for the CIR volatility. This is shown in Figure 8.1 for the low and high risk entities whose parameters were presented in Table 7.2 in Section 7.1.1. The value of x in the integrand is set equal to μ for each of the entities in the illustration, while the time parameter is assigned a value of 5 years. It is clear from Figure 8.1 that, as the volatility is increased, discontinuities appear in the integrand. There are also discontinuities in the integrand for low volatility parameters, but these are not evident in the diagram.

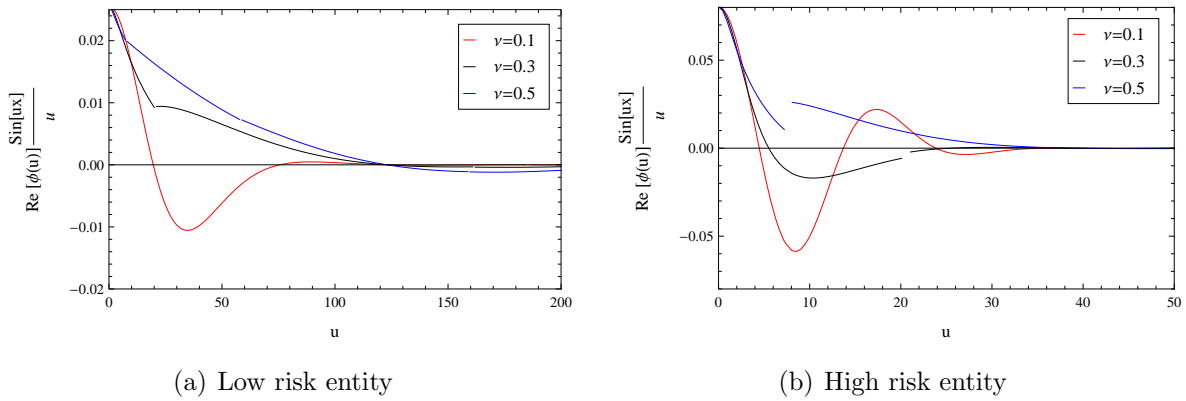


Figure 8.1: Integrand for various CIR volatility parameter values

The discontinuities are more evident in the plot of the riskier entity's integrand. We thus utilise the risky parameters for the remainder of the discussion. Since the $\frac{\sin(ux)}{u}$ portion of the integrand is familiar and known to be continuous, we investigate the real part of the characteristic function for the reason behind the discontinuities in the integrand. As such, observe the plot of the real part of the characteristic function in Figure 8.2, in which the relevant parameters are the same as those used to produce Figure 8.1(b). As expected, the discontinuities are present. We now plot the real and imaginary parts of




Figure 8.2: Real part of characteristic function for various volatility parameter values

both the numerator and the inverse of the denominator of the characteristic function. These are shown in Figure 8.3. Evidently, it is from the denominator that the discontinuity stems. Note that, although it is only the real part of the characteristic function that contributes to the integrand, both the real and imaginary parts of the numerator and denominator will play a role in its value, due to the division of the former by the latter.

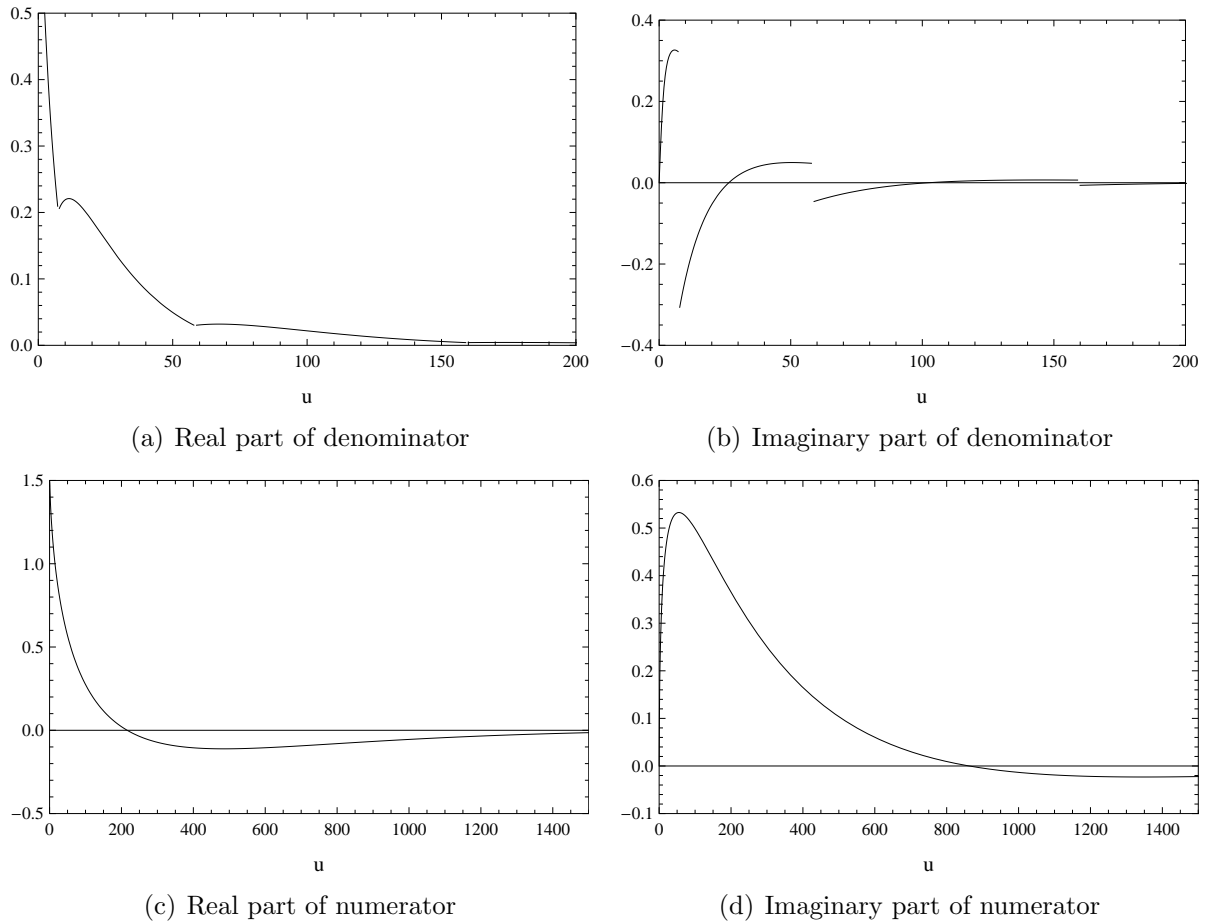


Figure 8.3: Aspects of the integrated CIR process characteristic function

Complex Discontinuities

The problem with the characteristic function is similar to the one that was documented in the Little Heston Trap [82]. However, we have fallen into a slightly different trap. Recall firstly from elementary linear algebra, that a complex number z , can be represented as

$$z = r_z e^{it_z}, \quad t_z \in [-\pi, \pi),$$

where the angle t_z is the phase or argument of z [83]. The fact that the phase is restricted to $t_z \in [-\pi, \pi)$ means that the complex plane is cut along the negative real axis [84]. Most software packages (including Matlab and Mathematica) make this restriction. Whenever the negative real axis of a function is cut, a branch cut, its phase (or argument) changes from $-\pi$ to π , which is a perfectly acceptable shift of 2π . When the function that exhibits the branch cuts is taken to a power of α say, the argument (or phase) changes instead an arbitrary amount from $-\pi\alpha$ to $\pi\alpha$. We no longer have an innocent 2π shift!

Many of the investigative ideas utilised in the remainder of this section were inspired by [85]. To determine the exact problem with the characteristic function, we proceed to examine the phase of the portion of the denominator inside the bracket, which we call f ; that is

$$f(u) := \cosh(\gamma t/2) + \frac{\kappa}{\gamma} \sinh(\gamma t/2), \quad (8.26)$$

where the function γ was defined above. A plot of the argument of the function $f(u)$ is shown in Figure 8.4(a), in which the parameters remain the same as in Figure 8.2 and the volatility parameter takes on a value of 0.5. In addition, we plot the imaginary part of f against the real part in a phase diagram. Note that u takes on a maximum value of only 10 in the phase diagram which is shown in Figure 8.4(b). It is clear from the




Figure 8.4: Investigations into the denominator without the exponent

first part of Figure 8.4 that the phase of f jumps from $-\pi$ to π . The second part of the figure confirms that the function cuts the negative real axis. Since u obtains a maximum value of 10 in Figure 8.4(b), the branch cut illustrated corresponds to the first jump in Figure 8.4(a).

Recall now that the denominator of the characteristic function in (8.24) consists of the function f , to which the graphs in Figure 8.4 apply, along with an exponent of $2\kappa\mu/\nu^2$. The source of the discontinuities is thus the exact phenomenon we described above. We have that $\text{Arg}[\text{denominator}] = \text{Arg}[f]^{\frac{\kappa^2\mu t}{\nu^2}}$ where Arg is the argument or phase of the function. The angle of the denominator thus jumps from $-\pi 2\kappa\mu/\nu^2$ to $\pi 2\kappa\mu/\nu^2$, rather than performing a full revolution. For the parameters utilised in the construction of Figure 8.4, $\pi 2\kappa\mu/\nu^2$ assumes a value of 1.01. As confirmation of our reasoning, we plot the phase of the entire denominator in Figure 8.5 below. This reaffirms the above statements.




Figure 8.5: Phase of the entire denominator

Further Exploration

We have now established the cause of the discontinuities in the integrand. Before searching for a solution, we investigate the problem in more detail. As an initial step, consider once again Figure 8.2, which illustrates the real part of the characteristic function for a range of CIR process volatilities. It appears that the discontinuities are only present when volatilities are high. This might lead us to some conclusion involving the Feller condition not being satisfied.

Delving a bit deeper into the phenomenon, we plot the imaginary versus real parts for a set of ν values and various ranges of the variable u . We find that the smaller ν , the larger the value of u at which the first branch cut occurs. As an example, consider the case in which $\nu = 0.05$. The argument of f , as well as its phase diagram, are plotted in Figure 8.8. The upper limit on u in Figure 8.6(b) is 900, compared to a maximum u value of 10 in the phase diagram for the case in which $\nu = 0.5$ (shown in Figure 8.4(b)). Clearly the first branch cut only occurs somewhere between $u = 700$ and 800 now that the volatility is 0.05. Typically, for very low volatilities, the numerical integration might




Figure 8.6: Investigations into f for $\nu = 0.05$

already have been terminated by the time the first branch cut occurs, since we utilise a discrete Fourier transform in the computation of the CDF of the integrated CIR process. Additionally, the integrand would generally be a great deal closer to zero at this first branch cut than it would be for the first branch cut when volatilities are large, rendering the discontinuity less significant.

For a given set of CIR volatilities and a particular value of t , there must be a value of γ (which is a function of u) at which the first branch cut occurs. We call this value Γ and solve for the value of u that it implies [85]. That is:

$$\begin{aligned} \Gamma &= \sqrt{\kappa^2 - 2\nu^2 i u} \\ \Gamma^2 &= \kappa^2 - 2\nu^2 i u \\ u &= \frac{\kappa^2 - \Gamma^2}{2i\nu^2}, \end{aligned}$$

from which it is clear that the value of u at which the first branch cut occurs is inversely proportional to the size of the volatility ν . This is consistent with the empirical investigations above.

Next, we plot the real part of the characteristic function, $Re[\phi(u)]$, for the risky entity for a set of time parameters and keeping ν constant at 0.5. This is presented in Figure 8.7. The discontinuities appear to be absent when t is small. However, in Fig-





Figure 8.7: Real part of characteristic function for various time parameter values

ure 8.8, we illustrate the phase of the function $f(u)$ for two of the smaller values of t . The branch cuts are now visible. Evidently, this parameter has the same effect as the volatility, in that it is inversely proportional to the value of u for which the first branch cut occurs.

8.2.4 Rescued!

The issue of discontinuous complex valued functions has been addressed in a number of ways in the literature. The aim is to guarantee that the phase of the function f in (8.26), in the denominator of the characteristic function $\phi(u)$ in (8.24), is continuous. This is to ensure that branch cuts are avoided when the exponent $(2\kappa\mu/\nu^2)$ is introduced.


 Figure 8.8: Argument of f for small values of t

The approach that appears to be both the simplest and least likely to introduce errors is the so-called *Rotation Count Algorithm*. Figure 8.9 illustrates an example of the argument of f , in addition to the smooth curve obtained by counting the branch cuts and adjusting the phase by 2π each time there is a jump from $-\pi$ to π . The paper by Lord and Kahl [86], among others, describes the rotation count algorithm. Alternative approaches have proposed keeping track of the number of jumps in the denominator by comparing adjacent points in the numerical integration scheme [86]. However, this leads to a more complex algorithm with greater room for error, particularly if the step size in the numerical integration scheme is not sufficiently small.

Fortunately, there is a significantly simpler solution that comes to us courtesy of T. Schelfaut and P. Mayar [85]. All that is required is to factor the $e^{\gamma t/2}$ term out of the denominator of (8.24) and the discontinuity in the denominator is removed. The resulting formulation of the characteristic function is then:

$$\phi(u) = \frac{e^{\left(\frac{\kappa\mu t}{\nu^2}(\kappa-\gamma)\right)} e^{\left(\frac{2y_0 i u}{\kappa+\gamma \coth(\gamma t/2)}\right)}}{\left[\frac{1}{2}\left(1 + \frac{\kappa}{\gamma} + e^{-\gamma t}\left(1 - \frac{\kappa}{\gamma}\right)\right)\right]^{2\kappa\mu/\nu^2}}, \quad (8.27)$$

where, as before,

$$\gamma = \sqrt{\kappa^2 - 2\nu^2 i u}.$$

The actual derivation of (8.27) appears in Section E.2 of Appendix E. Note that (8.27) is merely a rearrangement of the original characteristic function in (8.24). It has not been altered by any additions to or subtractions from it.




Figure 8.9: Illustration of the Rotation Count Algorithm. *The red line shows the phase of f before the algorithm is applied while the black line contains the adjustment.*

The advantage of the new formulation in (8.27), over the original in (8.24), is that the jumps in the phase of the function f have been shifted to the numerator. However, these are no longer a concern since the numerator does not have an exponent. For the purpose of illustrating this, we define the function g to be the part of the denominator of the newly formulated characteristic function beneath the exponent (the analogue of f in (8.26)):

$$g(u) := \frac{1}{2} \left(1 + \frac{\kappa}{\gamma} + e^{-\gamma t} \left(1 - \frac{\kappa}{\gamma} \right) \right). \quad (8.28)$$

The phase of the function g , as well as that of the newly formulated numerator, are illustrated in Figure 8.10 below.

Next, consider the limit of the function g as u approaches infinity.

$$\begin{aligned} \lim_{u \rightarrow \infty} g(u) &= \lim_{u \rightarrow \infty} \frac{1}{2} \left(1 + \frac{\kappa}{\gamma} + e^{-\gamma t} \left(1 - \frac{\kappa}{\gamma} \right) \right) \\ &= \frac{1}{2} \end{aligned}$$

We can thus be assured that the denominator will not have discontinuities, regardless of the CIR parameter set, since the limit approached by the function g is a real number [85].




Figure 8.10: Phases of the new formulation of $\phi(u)$, the characteristic function of the integrated CIR process.

Three illustrations conclude this section. Firstly, we present a plot of the imaginary versus real parts of the function g in Figure 8.11. The CIR parameters are those that have been employed throughout the section. In addition, a plot of the real part of the new formulation of the characteristic function is illustrated. This is the analogue of the graph that was presented in Figure 8.2, with the distinction that the new characteristic function formulation is now utilised. It is shown in Figure 8.12 below. Lastly, the integrand for both the low and high risk entities (that was plotted in Figure 8.1 at the beginning of this section) is repeated in Figure 8.13 with the aid of the revised formulation of the characteristic function (Equation (8.27)). All three illustrations (Figures 8.11 to 8.13) demonstrate that the original discontinuities have been removed.

The modified formulation of the characteristic function is employed in all computations in the sequel that require $\phi(u)$ as an input.

8.2.5 Parameters Required for the FRFT Implementation

Note that the reason for few references being cited in the remainder of this chapter is that (unless otherwise stated), the ideas are the author's original work. The means of implementing the FRFT process, including the determination of its parameters, is subjective and the product of a great deal of empirical investigative work. Secondly, the means of bypassing the implementation issue associated with the integrand being undefined at the point $u = 0$ (see Section 8.2.6 to follow), was the author's own decision. It is, however,




 Figure 8.11: Imaginary versus real part of $g(u)$


Figure 8.12: Real part of revised characteristic function

likely that the technique has been utilised by other practitioners.

The appropriate form of the integrand in (8.17) has now been obtained. The next step

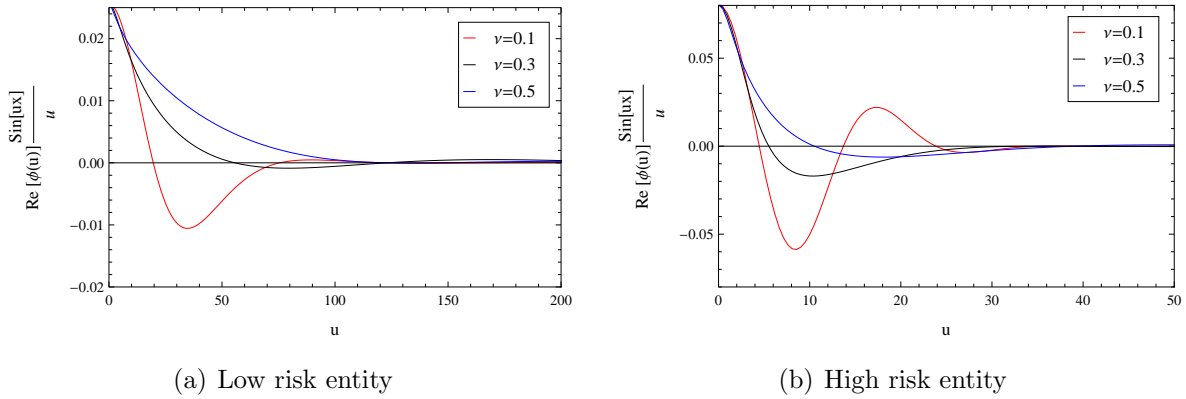


Figure 8.13: Integrand obtained from the revised characteristic function

is to implement the approximation in (8.21) for the calculation of the cumulative distribution function of the integrated CIR process. In other words, we wish to implement the FRFT in order to obtain the p_j 's in (8.16).

Recall firstly that the conditional probability in (8.16) is required for each T_k that takes place after τ_2 , for each simulation in which there is a relevant counterparty default event. We now proceed to explain the parameterisation of the FRFT process for a given T_k in a particular simulation. The output of this is a grid of p_j 's and a grid of x_k 's, both of which are employed in the approximation of the probability in (8.16).

The following parameters are required:

- (i) λ : This is the spacing of the x_k 's (see (8.18) to (8.21)). The parameter thus determines the granularity of the cumulative distribution that is obtained. The value of λ that is selected will contribute to the quality of the approximation of the reference entity's conditional survival probability.
- (ii) δ : This is the distance between successive u_j 's (see (8.18) to (8.21)). Each p_j in the CDF is computed using a set of u_j 's. The value of δ will therefore be a determining factor in the quality of the cumulative probabilities and, in turn, in the quality of the reference entity's conditional survival probability.
- (iii) N : This is both the number of u_j 's in the approximation of each p_j and the number of p_j 's and thus x_j 's in the CDF. It is required to be large enough to ensure sufficient granularity in the approximation of the integrals in (8.16) and (8.17).

- (iv) \mathbf{y}_{τ_2} : This is the spot CIR intensity required as an input to the characteristic function (see (8.27)).

We now discuss each of these parameters in turn.

The Selection of λ

We begin by recollecting a few essential properties of CDF's from elementary statistics (see [87]):

- (i) The CDF of a stochastic variable X is defined by $F(b) = P(X \leq b)$ for every real value of b .
- (ii) For a continuous variable, the cumulative distribution function is continuous and non-decreasing.
- (iii) The minimum value of $F(b)$ is zero for any X while the maximum value that $F(b)$ can attain is one.

Since the integrated CIR process is continuous, it is important to ensure (in the selection of λ) that the x_k 's are granular enough that the approximated survival probability in (8.16) is of an acceptable accuracy. In contrast, the maximum value of x_k (x_N) is required to be sufficiently large that $\mathbb{Q}(X < x_N) = 1$. This is to ensure that the full CDF is obtained. Note also that, since the CIR process assumes a minimum value of zero, the integral of the process cannot be negative. The x grid is thus floored at zero.

We utilise the expected value and variance of the integrated CIR process in the selection of λ . The expression for the expected value, taken from [81], is:

$$\mathbb{E}[Y_t|y_0] = \mu t + \kappa^{-1}(y_0 - \mu)(1 - e^{-\kappa t}). \quad (8.29)$$

The formula for the variance is more complex. We refer the reader to page 16 of [88] where the expression for $\mathbb{E}[Y_t^2]$ is given. Recall the relation from basic statistics which expresses the variance of a variable Y as the difference between the expectation of its squared value and the square of its expectation: $\text{Variance}(Y) = \mathbb{E}[Y^2] - (\mathbb{E}[Y])^2$ [19]. This result is used to obtain the variance of the integrated CIR process as implemented in Matlab (see Appendix H).

The size of λ for every T_k in each relevant simulation is then computed using the following formula:

$$\lambda = \frac{11 * \text{Standard Deviation}(Y) + \mathbb{E}(Y)}{N}, \quad (8.30)$$

where $Y = \int_{\tau_2}^{T_k} y_s ds$.

The decision to multiply the standard deviation by 11 was made by selecting a value that was sufficiently high that the maximum probability attained would be one, but low enough that the CDF would be sufficiently granular for the computation of the survival probability. A number of empirical tests were conducted to confirm this parameterisation and to ensure that it would be generic enough to accommodate a range of volatilities.

The x_j 's are then set equal to: $x_j = (j - 1)\lambda$ for $j = 1, \dots, N$. Note that, since we are using the fractional FFT, λ is selected independently of δ . This attractive property of the FRFT was discussed in Appendix F. It is the reason that a fractional FFT was selected over a standard FFT.

The Determination of δ

The parameterisation of δ (the u spacing) is more complex than that of λ since the integral in (8.17) is being approximated. Therefore, the selection of a δ value should take into account the shape of the integrand, which (for a value of x_k on the x grid) is:

$$\text{Re}(\phi(u)) \frac{\sin ux_k}{u}. \quad (8.31)$$

We proceed to explore the integrand in order to ascertain its behaviour prior to determining a δ selection methodology.

Understanding the Integrand - Varying the Time Parameter

The first parameter whose effect we consider is time t or $T_k - \tau_2$ using our notation. The illustrations below are all for the SOAF COB 31 August 2010 curve that has been utilised throughout this dissertation¹. As a starting point to our investigation into the shape of the integrand, consider the integrand for a selection of time values. This is illustrated in Figure 8.14 below. The parameter x_k is set equal to the long run mean of the CIR process (μ). With reference to Figure 8.14, it appears that the lower the value of the

¹Recall that the curve was presented in Table 7.1, with the CIR++ parameters shown in Table 7.2.




Figure 8.14: Integrand for various time parameter values

time parameter (all other parameters held constant), the longer the integrand takes to converge to zero. The region in which the area under the curve is the most significant is that in which u is small. When u is large, we therefore expect the values of the integrand to have a smaller impact on the valuation of the integral than when it is close to zero. Additionally, observe that the integrand changes rapidly with an increase in t for small u values. In other words, the gradient is steeper. This demonstrates the fact that, as t is increased, a smaller δ will be required to capture the rapidly changing area under the curve.

Consider next plots of the two functions of which the integrand is comprised. These are the real part of the characteristic function (shown in Figure 8.15(b)) and the function $\frac{\sin(ux_k)}{u}$ (graphed in Figure 8.15(a)). Since $\frac{\sin ux}{u}$ is not a function of t , we observe that the reason for the lower maturity integrand taking longer to converge to zero is the real part of the characteristic function differing significantly from the origin. When the time parameter is increased, the real part of the characteristic function converges rapidly, dampening the oscillatory shape of the $\sin(ux)/u$ portion of the integrand.




Figure 8.15: Breakdown of integrand for a selection of time parameters

Understanding the Integrand - Varying the Volatility Parameter

As part of our exploration into the behaviour of the integrand, we now consider the effect of alternative inputs to the CIR volatility parameter (ν). The time parameter is set to 5 in Figure 8.16 below. We observe (from Figure 8.16) that the larger the value of





Figure 8.16: Integrand for various volatility parameter values

the volatility, the less oscillatory the shape of the integrand. We repeat this plot, but for a lower t value of 0.5. This is reported in Figure 8.17. The small time parameter overshadows the differences in volatility of the three graphs. Figure 8.17(b), in which

we have zoomed in on the tail, illustrates that the speed with which zero is attained is still partly affected by the volatility. Note that the size of the integrand for these large u values is considerably lower than for the earlier u values and will be far less significant in determining the value of the integral.

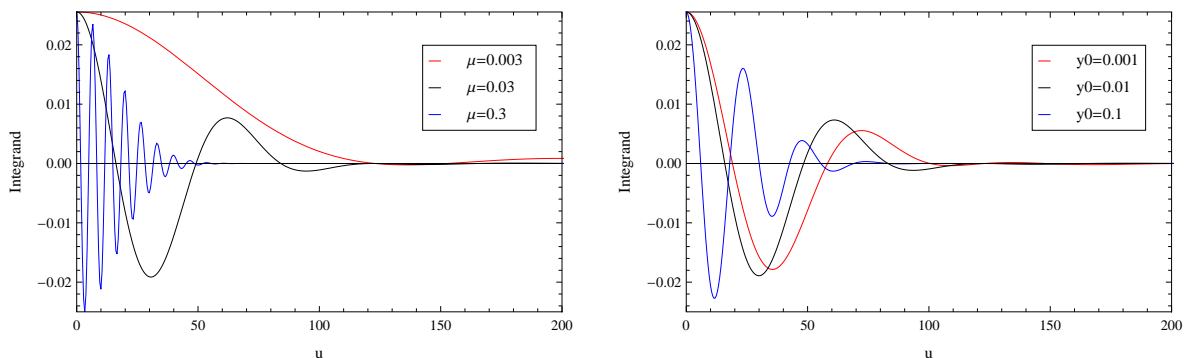

 (a) Plot of integrand - various ν , small t

 (b) Zoom in on integrand for large u values

 Figure 8.17: Integrand for a selection of volatility parameters with $t = 0.5$

Understanding the Integrand - Some Additional Variations

Finally, we illustrate the effect of varying μ (the long run mean) and y_0 (the initial value of the CIR process). These are presented in Figures 8.18(a) and 8.18(b) respectively, with the time parameter set equal to 5. Clearly, the μ parameter has the effect of resulting in


 (a) Effect of μ

 (b) Effect of y_0

 Figure 8.18: Integrand for μ and y_0 variations

a more oscillatory integrand as it is increased. This would require a smaller delta since the area under the integrand changes rapidly, but also converges to zero at a greater pace than for lower values of μ . The effect of altering the initial CIR value is similar to that

of varying μ although significantly less pronounced.

At this point, we have obtained a fairly comprehensive understanding of the integrand and the considerations that have to be taken into account in parameterising a generic model that provides accurate results for an extensive array of CIR parameters. To summarise, in selecting the delta value, we are required to ensure the following:

- (i) The parameterisation process is not too generic to capture the manner in which various parameters affect the choice of δ . There should be sufficient granularity in the set of possible deltas that are available for selection.
- (ii) The value of δ is not so large that the variation in the integrand for small values of u fails to be captured. These early values of u are where the area under the graph is greatest and thus contribute most significantly to the computation of the integral.
- (iii) In contrast, δ ought to be large enough that the integrand has converged to zero by the time u attains a value of $N\delta$.
- (iv) The tolerance level at which the integrand is no longer considered to differ significantly from zero is appropriate for the size of N . A small tolerance combined with a small value of N , will result in a value of δ that is large enough to ensure that the function has converged to zero, but too large to capture the variation in the integrand. Thus the CDF will be inaccurate.

The Delta Selection Methodology

The methodology utilised to select δ is not unique. It was developed by means of a significant quantity of practical testing. For a set value of N , the steps employed in determining δ are listed below:

- (i) **Step 1:** Obtain the expected value of the integrated CIR process, $\int_{\tau_2}^{T_k} y_s ds$ for every T_k in each simulation in which $\tau_2 < \tau_1$ and $\tau_2 < T$.
- (ii) **Step 2:** From the largest and smallest expected values obtained in Step 1, create M equally spaced buckets. The partitions are determined by taking the difference between the largest and smallest expected values and dividing this difference by M . Split the first bucket into a further m partitions. This is due to the presence of more variation in the first bucket than in the others. The δ parameter can thus not be too generic here. Each bucket will be assigned a delta parameter. We will thus have $M + m - 1$ different values of δ .

- (iii) **Step 3:** At this point, each of the expected values computed in Step 1 falls within one of the $M + m - 1$ buckets created in Step 2. For each bucket determine the minimum value of the expected values from Step 1 that lies in the bucket. The values of $T_k - \tau_2$ and y_{τ_2} , corresponding to each of these minimum expected values, are also recorded. These are the two parameters that differentiate the values of $\int_{\tau_2}^{T_k} y_s ds$ from each other for a given set of CIR parameters.
- (iv) **Step 4:** A vector of potential deltas is hard coded. This set is extensive and was selected through a large amount of testing. Now recall that the integrand with which we are concerned is $Re(\phi(u)) \frac{\sin ux_k}{u}$. For each bucket, we set x_k equal to the minimum expected value for the bucket that was obtained in Step 3. Using the corresponding $T_k - \tau_2$ and y_{τ_2} values, the integrand is computed at $u = j\delta$ for $j = 1, 2, \dots, N$ for each of the possible δ values in the vector of potential deltas. If the number of possible delta values is A , we will thus have $M + m - 1$ matrices of size $A \times N$.
- (v) **Step 5:** For each of the $M + m - 1$ buckets, we now select the δ for the bucket from the A possible values. The δ chosen is the smallest of the A possibilities subject to the following:
- (a) The value of the integrand at $u = N\delta$ is not significantly different from zero, based on a pre-defined tolerance of $5E - 5$. This choice of the tolerance level is subjective. It was selected in line with item (iv) in the list above the heading ‘The Delta Selection Methodology’. A large quantity of empirical testing accompanied its selection.
 - (b) The difference between the integrand at $u = N\delta$, $u = (N-1)\delta$ and $u = (N-2)\delta$ is not significantly different from zero, based on the pre-defined tolerance of $5E - 5$. This is to ensure that the integrand is not merely ‘passing through’ zero, but has actually converged.
 - (c) The difference between the integrand at $u = N\delta$ for the selected δ and the previous δ in the vector of potential δ ’s is close to zero. This condition was included after a number of tests presented a case in which the function was zero for a certain delta, but hadn’t actually leveled out yet. The integrand value at $u = j\delta$ was not zero for the delta values before or after it in the vector of potential deltas.

Since each of the reference entity survival probabilities is obtained with the aid of N p_j ’s, testing all N values of the CDF for every T_k in each relevant simulation would be

impractical. Rather, the probabilities computed by means of the FRFT were used to compute the CIR bond price. This was tested against the closed form solution for the bond price (shown in Section C.2 in Appendix C). The exact methodology is explained in Section E.3 of Appendix E. The discussion requires an understanding of the next section and should only be read after Section 8.3.

The Value of N

At the start of Section 8.2.5, it was pointed out that the following parameters were required to implement the FRFT: λ , δ , N and y_{τ_2} . At this point, the selection of the first two has been discussed for a fixed value of N . We now consider the value of N (the number of steps in the FRFT).

Since the model has been implemented in Matlab, which operates well with the use of matrices, a single N value is selected and used throughout. We have employed a default value of 256 steps in the Fourier sum. This was selected to ensure that, even when parameters are badly behaved and maturities are large, the error in the CIR bond price obtained via the probabilities generated using the FRFT is sufficiently small (more on this in Section E.3 of Appendix E). A value of 128 proved to be sufficient for a number of possible parameter sets. However, in certain cases, it did not provide a sufficient level of accuracy. In such cases, in order to ensure that the integrand converges to zero by $u = N\delta$, the size of δ is required to be too large to capture the shape of the integrand. There is thus a trade-off between the sizes of N and δ . The convergence of the CVA is discussed in Section 9 for various choices of N . It turns out that $N = 128$ is sufficient and the isolated cases in the testing in this section, that did not produce satisfactory results when compared with the CIR bond price, do not alter the final CVA value.

The Spot Intensity at Default

It is evident from the characteristic function in (8.27) that the initial spot intensity at τ_2 (y_{τ_2}) is required as an input to the valuation of the integrand in the FRFT sum. Since the p_j 's that we are computing are concerned only with the integrated CIR process, and not the entire CIR++ process, we subtract the deterministic shift ψ_{τ_2} from our simulated intensities to obtain the value of y_{τ_2} . It is this initial spot intensity that is used as an input to the FRFT process.

8.2.6 A Note on the Limit of the Integrand

Before concluding this section on the FRFT implementation, we point out that the first term in the Fourier sum is the value of the integrand at the point $u = 0$. This presents a problem in the Matlab implementation when the built-in *fft* (fast Fourier transform) and *ifft* (inverse fast Fourier transform) functions are utilised. Recall from Section F.6 in Appendix F that two FFT's and one inverse FFT are employed in computing the FRFT.

At the point $u = 0$, the integrand,

$$Re(\phi(u)) \frac{\sin ux}{u},$$

is undefined. The real part of the characteristic function assumes a value of one when $u = 0$. This is easily seen, since the value of $\gamma = \sqrt{\kappa^2 - 2\nu^2 i u}$ is equal to κ at this point. Therefore:

$$\begin{aligned} \phi(0) &= \frac{e^{(\frac{\kappa\mu t}{\nu^2}(\kappa-\kappa))} e^{(\frac{2y_0 i 0}{\kappa+\kappa \coth(\kappa t/2)})}}{[\frac{1}{2}(1 + \frac{\kappa}{\kappa} + e^{-\kappa t}(1 - \frac{\kappa}{\kappa}))]^{2\kappa\mu/\nu^2}} \\ &= \frac{1}{1^{2\kappa\mu/\nu^2}} = 1. \end{aligned}$$

The function $\frac{\sin ux}{u}$, on the other hand, is undefined at $u = 0$. Applying L'Hôpital's rule [26], it is easy to see that²

$$\lim_{u \rightarrow 0^+} \frac{\sin ux}{u} = x.$$

Since $Re(\phi(0))$ is always 1, it follows that:

$$\lim_{u \rightarrow 0^+} Re(\phi(u)) \frac{\sin ux}{u} = x.$$

When the model was implemented in Matlab, the smallest value of u was set equal to the second to avoid an error since the sum in the Fourier transform is evaluated at $u = 0$. This is then corrected, by making use of the limit just described, once the FRFT has been executed.

²This is taken from elementary calculus. Since $\lim_{u \rightarrow 0^+} \frac{\partial}{\partial u} \sin ux = \lim_{u \rightarrow 0^+} \cos ux = x$ and $\lim_{u \rightarrow 0^+} \frac{\partial}{\partial u} u = 1$ the result follows easily.

8.3 CDS Valuation

The first section of this chapter contained some derivations and useful expressions that we stated would be required later. In the second part of the chapter (Section 8.2) we discussed the fractional fast Fourier transform as it would be applied to obtain the cumulative distribution function of the integrated CIR process. In this section, we utilise the results of the previous two sections to obtain the reference entity's survival probabilities at τ_2 . These are required to revalue the underlying CDS upon counterparty default. We focus on the computation of these conditional survival probabilities. The valuation of the CDS is straightforward once they have been obtained.

At this point, for each simulation path in which the counterparty has defaulted prior to the reference entity and CDS maturity, we have a grid of p_j 's and corresponding x_k 's for each CDS premium payment time (T_k) that occurs after τ_2 . Recall from (8.6) that the conditional survival probability for time T_k is given by:

$$\mathbb{Q}(\tau_1 > T_k | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2) = E[P(U_1) | \mathcal{G}_{\tau_2}] = \int_{u=U_1}^1 P(u_1) dC_{1|2}(u_1, U_2).$$

As we pointed out in the previous section, this probability can be approximated by the sum:

$$\sum_j \frac{p_{j+1} + p_j}{2} \Delta f_j, \quad (8.32)$$

where $f_j = C_{1|2}(u_j, U_2)$; $p_j = \mathbb{Q}(X < x_j)$ and $X = Y_1(T_k) - Y_1(\tau_2) = \int_{\tau_2}^{T_k} y_s ds$. Since the computation of the p_j 's has been described, attention is now focused on the values of f_j .

8.3.1 Obtaining the Conditional Copula Function

Recall from Section 8.1, that the expression for the conditional copula function is:

$$C_{1|2}(u_j, U_2) = \frac{\mathbb{Q}(U_1 < u_j | U_2) - \mathbb{Q}(U_1 < \bar{U}_1 | U_2)}{1 - \mathbb{Q}(U_1 < \bar{U}_1 | U_2)}. \quad (8.33)$$

This result is now employed to explain how $f_j = C_{1|2}(u_j, U_2)$ in (8.32) is computed. The quantities U_2 , \bar{U}_1 and u_j are required for this. We explain how each is obtained in turn.

Calculation of U_2

Recall that the expression for U_2 is $U_2 = 1 - e^{-Y_2(\tau_2) - \Psi_2(\tau_2, \beta)}$. Since we have the integrated intensities for both the reference entity and counterparty for each simulation path and since the integrated intensity is $\Lambda(t) = Y(t) + \Psi(t)$, the expression for U_2 is easy to calculate. Note also that we assume that default occurs on the next CDS premium payment date if it is between dates. Since we ensured that the intensity was simulated and integrated at each payment date, no interpolation is required.

Calculation of \bar{U}_1

The expression for \bar{U}_1 is given by $\bar{U}_1 = 1 - e^{-Y_1(\tau_2) - \Psi_1(\tau_2, \beta)}$. The calculation follows in a similar manner to that of U_2 .

Calculation of u_j

From the Fourier transform in the previous section, we have a vector of x_j 's for each CDS premium payment date (T_k) remaining after τ_2 . The number of elements in the vector x will be N , the number of terms in the Fourier sum. A uniform random variable (u_j) corresponding to each x_j is now required. Therefore, we transform the x_j 's into uniform random variables, u_j . Note that the symbol u is used here to denote uniform random variables and it is not the same as the u used in the characteristic function $\phi(u)$ in the previous section. The expression for u_j is given by:

$$u_j = 1 - e^{-x_j - Y_1(\tau_2) - \Psi_1(T_k)}. \quad (8.34)$$

This follows from the fact that the variable X is used to represent $\int_{\tau_2}^{T_k} y_s ds$. Additionally, $\Psi_1(T_k)$ is known at the outset since it is deterministic and $Y_1(\tau_2)$ is known at time τ_2 .

In Section 7.2.2 of Chapter 7, we explained that the integration of the intensities was performed using either piecewise or linear interpolation (depending on which of these was used in fitting the parameters of the CIR++ process). In order to ensure consistency between $\Lambda_1(\tau_2)$ and $Y_1(\tau_2) + \Psi_1(\tau_2)$, we calculate $Y_1(\tau_2)$ and $\Psi_1(T_k)$ using the same method of integration. That is, we obtain the CIR part of the simulated intensities $y_1(t)$ and the deterministic shift $\psi_1(t)$ and integrate in the same way as we integrated the λ 's to obtain Λ .

Now that we have the inputs required for the evaluation of f_j , we turn to the actual

calculation of the conditional copula function. We explain the computation of the term $\mathbb{Q}(U_1 < u_j | U_2)$ in (8.33) and that of the term $\mathbb{Q}(U_1 < \bar{U}_1 | U_2)$ follows in a similar manner.

The Gaussian copula

As discussed previously, a Gaussian copula is employed to link the uniform random variables of the exponential triggers of the default times of the reference entity and counterparty. The bivariate Gaussian copula $C(u, v)$ was presented in Equation (7.11) in Section 7.2.3. For ease of reference we repeat it below:

$$C(u, v) = \int_{-\infty}^{\Phi^{-1}(u)} \int_{-\infty}^{\Phi^{-1}(v)} \frac{1}{2\pi(1-\rho^2)^{1/2}} e^{-\frac{s^2 - 2\rho st + t^2}{2(1-\rho^2)}} ds dt,$$

where ρ is the correlation parameter and Φ^{-1} is the inverse of the standard normal cumulative distribution function.

There is a result in the statistics of multivariate normal distributions that gives the conditional distribution of a vector Y given X where Y and X are jointly multivariate normally distributed. In our application, the normal distribution is bivariate. We thus state the result for scalars X and Y that are bivariate normal (see [19]):

If X and Y are bivariate normally distributed with means μ_1 and μ_2 , standard deviations σ_1 and σ_2 respectively and correlation parameter ρ , the conditional distribution of $Y|X = x$ is normal with a mean of

$$\mu_2 + \frac{\sigma_2}{\sigma_1} \rho (x - \mu_1) \quad (8.35)$$

and a variance of

$$(1 - \rho^2) \sigma_2^2. \quad (8.36)$$

We apply this result to obtain the conditional distribution of $\Phi^{-1}(U_1) | \Phi^{-1}(U_2)$. It follows from Equations (8.35) and (8.36) respectively that the mean and variance of the conditional distribution are given by:

$$\mu_{\text{conditional}} = \rho \Phi^{-1}(U_2) \quad (8.37)$$

and

$$\text{var}_{\text{conditional}} = (1 - \rho^2). \quad (8.38)$$

The probability $\mathbb{Q}(U_1 < u_j | U_2)$ is thus easily calculated using the cumulative normal distribution function, evaluated at the point $\Phi^{-1}(u_j)$, where the mean and variance of the normal random variable are given by the expressions in (8.37) and (8.38) respectively.

As additional confirmation of the methodology used to obtain the conditional copula function, we recall the definition of the copula conditional distribution as presented in Section D.6 in Appendix D:

$$\frac{\partial}{\partial u} C(u, v) = \lim_{\epsilon \rightarrow 0} \frac{C(u + \epsilon, v) - C(u - \epsilon, v)}{2\epsilon}. \quad (8.39)$$

If this expression is approximated as follows:

$$\lim_{\epsilon \rightarrow 0} \frac{C(u + \epsilon, v) - C(u - \epsilon, v)}{2\epsilon} \approx \frac{C(u + \epsilon, v) - C(u - \epsilon, v)}{2\epsilon},$$

we are able (for small ϵ) to numerically evaluate $\mathbb{Q}(U_1 < u_j | U_2)$. Indeed, this approximation was employed to test the conditional distribution implementation described directly above it.

8.3.2 The Survival Probability

We have now explained how to calculate the value of f_j in (8.32). It remains to substitute the p_j 's from the FRFT routine along with the f_j 's into the sum,

$$\sum_j \frac{p_{j+1} + p_j}{2} \Delta f_j,$$

to obtain the conditional survival probability for a particular T_k in a given simulation.

Note that the integral in (8.6) has an upper integration bound of $u = 1$. Since the value of x_N was determined to ensure that p_N is 1, and since $\frac{1+1}{2} = 1$, once j reaches N , the remainder of the integral can be approximated as $f(u = 1) - f(u_N)$.

We now have the conditional survival probabilities for each T_k following τ_2 in the relevant simulations. These are used (in conjunction with the traded spread, the risk free discount curve³ and the direction of the trade) to determine the CDS value at τ_2 in each simulation in which $\tau_2 < \tau_1$ and $\tau_2 < T$. The pricing of CDSs was explained in Section 4.3.

³Recall that interest rates are assumed to be deterministic.

8.4 CVA Computation

For each simulation, we now have either a CDS value, if $\tau_2 < \tau_1$ and $\tau_2 < T$, or a zero value. Recall the discussion on counterparty credit exposure in Section 2.1. The notion of the replacement cost of a derivative, as the maximum of zero and the NPV of the derivative cash flows at the counterparty default time, was introduced. The next step in the CVA computation is thus to discount the CDS values and floor them at zero: $\max(D(0, \tau_2)CDS(\tau_2), 0)$. The sum of these positive values is then divided by the total number of simulations and multiplied by the LGD of the counterparty to obtain the credit valuation adjustment.

The value obtained is a monetary amount. In the numerical illustrations in the next chapter, the CVA is expressed as a spread in basis points. This can be interpreted as the amount to be subtracted from the spread paid on a payer CDS (long protection position) due to the risk of a counterparty default. In the event that we are short protection, it is added to the spread received as compensation for the possibility of the counterparty defaulting when the CDS is in-the-money from the investor's perspective.

Suppose that CDS^D is the value of the CDS when counterparty risk is taken into account and CDS^{ND} is the corresponding counterparty default-free CDS value. Similarly the spread of the defaultable CDS is denoted by S^D , with the corresponding counterparty risk-free spread represented by S^{ND} . We denote the adjustment, in basis points, by S . For a receiver CDS, we thus have $S^D = S^{ND} + S$ and for a long protection position, $S^D = S^{ND} - S$. The reason for the addition of the spread S to the risk-free spread (that we receive on a short protection position) is that the protection we have sold is worth less than it would have been worth if the counterparty were default-free. This is due to the possibility of the counterparty defaulting when the contract is in-the-money from our perspective. Similarly, in the long protection case, we would pay S^{ND} to a risk-free counterparty, but should actually be paying S basis points less for the protection we have purchased, due to the counterparty's risk of default.

Let PL represent the premium leg before it is multiplied by the spread at which the trade was conducted and let DL represent the default (protection) leg of the CDS. For a reminder of the valuation of each of these components, refer to Section 4.3. The adjust-

ment, S , is then obtained as follows for a receiver CDS:

$$\begin{aligned}
 CDS^D &= CDS^{ND} - CVA \\
 [(S^{ND})(PL) - DL] &= [(S^D)(PL) - DL] - CVA \\
 (S^{ND})(PL) - (S^{ND} + S)(PL) &= -CVA \\
 S &= \frac{CVA}{PL}.
 \end{aligned} \tag{8.40}$$

Similarly, for the payer CDS (long protection position) the adjustment is computed using:

$$\begin{aligned}
 CDS^D &= CDS^{ND} - CVA \\
 [DL - (S^{ND})(PL)] &= [DL - (S^D)(PL)] - CVA \\
 -(S^{ND})(PL) + (S^{ND} - S)(PL) &= -CVA \\
 S &= \frac{CVA}{PL}.
 \end{aligned} \tag{8.41}$$

The interpretation of the spread S , that we have derived as a function of the CVA in both (8.40) and (8.41), is thus the counterparty risk adjustment that should be made to the CDS premium with all other inputs to the CDS price remaining the same as in the risk-free counterparty case. Note that the derivations above are the author's own and are intended to ensure that readers are aware of the exact interpretation of the spread in the numerical results reported in the next chapter.

8.5 Chapter Summary

This chapter described the determination of the replacement cost of the CDS along with the CVA computation once this cost has been determined in each simulation. It resumed the implementation discussion at the point in the CDS CVA calculation at which the intensities had been simulated and integrated and the relevant default times determined. Under the premise that the CDS values are readily available once the reference entity's survival probabilities (required in their computation) have been obtained, the majority of the chapter was devoted to the determination of these conditional probabilities. It was divided into four sections.

The first of these, Section 8.1, illustrated the manipulation of the reference entity conditional survival probability into an integral to facilitate numerical integration. Practical formulations of the components of the sum in the numerical integration scheme were then

derived. These permitted the application of an FRFT scheme and the Gaussian copula.

Section 8.2 focused on the actual FRFT implementation, with emphasis on the selection of the parameters required. A flaw in the characteristic function of the integrated CIR process (when it is used in conjunction with an FRFT) as it appears in the literature was identified. It was corrected by means of an algebraic manipulation of the function. Section 8.3 explained how the FRFT output (namely the cumulative distribution of the integrated CIR process) is used, in conjunction with the Gaussian copula function, to obtain the reference entity's survival probabilities at τ_2 . Recall that these are conditional probabilities that depend on the counterparty default times and the information available at these times. The CDS valuation is trivial once they have been computed.

The last section in the chapter detailed the actual CVA calculation from the CDS valuations. The adjustment to the fair spread of a counterparty default-free CDS was explained. The implementation of the CDS CVA model has now been covered in some depth. The next chapter reports numerical results that demonstrate the output of these computations.

Chapter 9

Results of the Model Implementation

Illustrations of the unilateral CVA are presented in this chapter. In each of the scenarios that follow, one of the three entities, whose par CDS curves were presented in Table 7.1, is assigned to the reference entity and one to the counterparty. While some of the parameters considered in this section are unlikely in practice, they have been included for stress testing the model.

For ease of reference, the par CDS term structures of the three entities that are used in the numerical results, are repeated in Table 9.1. The associated hazard rates and survival probabilities can be found in Table B.2 of Appendix B. The CIR++ parameters, obtained by minimising the difference between the fitted CIR model and the hazard rates bootstrapped from the market term structures, are shown in Table 7.2 in Section 7.

Term:	1	2	3	4	5	7	10
Low Risk (Base):	81	109	130	144	155	163	170
Intermediate Risk:	181	209	230	244	255	263	270
High Risk:	481	509	530	544	555	563	570

Table 9.1: Spreads in basis points for entities in the numerical illustrations

The chapter begins with an examination of the convergence of the credit valuation adjustment in basis points, with respect to both the number of Monte Carlo simulations and the number of terms in the FRFT sum. Results for an array of correlation and volatility parameter inputs are discussed thereafter.

9.1 Convergence

In this section, the convergence of the CVA, expressed as the spread, S , that was derived in (8.40) and (8.41) for short and long protection positions respectively, is examined. There are two parameters considered. The first is the number of Monte Carlo simulations required for convergence. The second is the number of steps, N , employed in the FRFT procedure. This parameter was included in the parameterisation discussion in Section 8.2.5. Recall that N is fixed in the implementation of the model in this dissertation and the step sizes (δ and λ) are selected based on this fixed value of N .

The appropriate size of N , the number of terms in the FRFT sum, may be influenced by the reference entity's parameters, since it is a function of the characteristic function of the integrated CIR process associated with the reference entity. The number of Monte Carlo simulations required for convergence could potentially be affected by both entities' parameters. Convergence results are thus shown for a variety of combinations of reference entity and counterparty levels of riskiness.

As discussed in Section 8.2.5, the general level of the CIR parameters affects the shape and speed of convergence of the FRFT integrand. Particularly, the volatility parameter influences the gradient of this function as well as the rate at which it approaches zero. As such, convergence results are shown for reference entity hazard rate volatilities of $\nu_1 = 0.1$, $\nu_1 = 0.8$ and $\nu_1 = 0.9$, which are intended to represent both low and extremely high values of the parameter. Additionally the reference entity is either the low or the high risk entity from Table 9.1.

Table 9.2 below reports the unilateral CVA for a receiver CDS (short protection position). The corresponding information for a payer CDS (long protection position) is presented in Table 9.3. Only negative default time correlation scenarios are illustrated in the former table for ease of readability and to avoid focusing on irrelevant data. Similarly, adjustments relating merely to positive default time correlations are reported in the latter table.

Two cases are presented. In the first, the reference entity is the low risk name, while the counterparty is the intermediate risk entity. In the second, the reference entity is the high risk name and the counterparty remains of intermediate risk.

Receiver CDS (short protection)											
Simulations	$\rho \setminus N$	Case 1 ($\nu_1 = \nu_2 = 0.1$)					Case 2 ($\nu_1 = \nu_2 = 0.1$)				
		Ref: Low; Cpty: Intermediate					Ref: High; Cpty: Intermediate				
		128	256	512	1,024	2,048	128	256	512	1,024	2,048
10,000	-0.99	9	9	8	8	9	36	33	35	33	34
	-0.9	8	8	8	8	8	34	33	34	34	36
	-0.75	8	8	8	8	8	32	31	30	33	33
	-0.5	7	7	7	7	7	24	23	24	25	24
	-0.25	4	4	4	4	4	12	12	12	12	12
20,000	-0.99	8	9	9	8	8	35	34	35	33	33
	-0.9	8	8	8	8	8	34	33	35	34	35
	-0.75	8	8	8	8	8	32	31	30	33	32
	-0.5	7	7	7	7	7	24	24	24	25	24
	-0.25	4	4	4	4	4	12	12	12	12	12
40,000	-0.99	8	9	9	8	8	35	34	34	33	34
	-0.9	8	8	8	8	8	34	33	34	34	35
	-0.75	8	8	8	8	8	32	31	31	32	32
	-0.5	7	7	7	7	7	24	24	24	24	25
	-0.25	4	4	4	4	4	12	12	12	12	12
60,000	-0.99	8	8	8	8	8	34	34	34	33	34
	-0.9	8	8	8	8	8	34	33	34	34	34
	-0.75	8	8	8	8	8	32	31	32	32	32
	-0.5	7	7	7	7	7	24	24	24	24	24
	-0.25	4	4	4	4	4	12	12	12	12	12
80,000	-0.99	8	8	8	8	8	34	34	34	34	34
	-0.9	8	8	8	8	8	34	33	34	34	34
	-0.75	8	8	8	8	8	32	31	32	32	32
	-0.5	7	7	7	7	7	24	24	24	24	24
	-0.25	4	4	4	4	4	12	12	12	12	12
100,000	-0.99	8	8	8	8	8	34	34	34	34	34
	-0.9	8	8	8	8	8	34	33	34	34	34
	-0.75	8	8	8	8	8	32	32	32	32	32
	-0.5	7	7	7	7	7	24	24	24	24	24
	-0.25	4	4	4	4	4	12	12	12	12	12

Table 9.2: Receiver CVA convergence for number of Monte Carlo simulations and number of terms in FRFT sum - low ν_1 . *The labels Ref and Cpty refer to the reference entity and the counterparty respectively, while low, intermediate and high refer to the riskiness of the entity.*

Payer CDS (short protection)											
Simulations	$\rho \setminus N$	Case 1 ($\nu_1 = \nu_2 = 0.1$)					Case 2 ($\nu_1 = \nu_2 = 0.1$)				
		Ref: Low; Cpty: Intermediate					Ref: High; Cpty: Intermediate				
		128	256	512	1,024	2,048	128	256	512	1,024	2,048
10,000	0.25	9	9	9	9	9	16	16	17	15	16
	0.5	20	19	20	20	20	28	31	29	28	28
	0.75	35	34	36	35	36	37	38	38	36	38
	0.9	57	53	52	54	51	30	31	30	29	33
	0.99	80	82	86	79	83	4	4	2	5	3
20,000	0.25	9	9	9	9	9	17	16	16	16	16
	0.5	20	19	19	20	20	30	30	28	29	28
	0.75	35	34	37	35	37	37	40	39	38	38
	0.9	56	54	51	55	52	31	31	31	29	33
	0.99	82	82	84	80	82	3	3	3	4	4
40,000	0.25	9	9	9	9	9	16	16	16	16	16
	0.5	20	19	19	20	19	30	31	29	29	30
	0.75	35	35	36	36	36	38	39	39	38	40
	0.9	54	54	52	54	52	31	32	31	31	33
	0.99	81	82	83	79	83	4	4	4	4	4
60,000	0.25	9	9	9	9	9	16	16	16	16	16
	0.5	20	19	19	20	19	30	31	29	30	30
	0.75	36	35	36	35	36	39	39	39	38	40
	0.9	54	53	53	54	52	32	31	31	31	32
	0.99	81	81	82	79	83	4	4	4	4	4
80,000	0.25	9	9	9	9	9	16	16	16	16	16
	0.5	20	19	19	19	19	30	31	29	30	30
	0.75	36	35	36	35	36	39	39	39	39	39
	0.9	53	53	53	54	52	32	32	31	31	32
	0.99	81	82	81	80	82	4	4	4	4	4
100,000	0.25	9	9	9	9	9	16	16	16	16	16
	0.5	20	19	19	19	19	30	30	30	30	30
	0.75	35	35	36	35	36	39	39	39	39	39
	0.9	54	53	53	55	53	31	32	30	31	31
	0.99	81	82	82	81	82	4	4	4	4	4

Table 9.3: Payer CVA convergence for number of Monte Carlo simulations and number of terms in FRFT sum - low ν_1 . *The labels Ref and cpty refer to the reference entity and counterparty respectively, while low, intermediate and high refer to the riskiness of the entity.*

In this section, the actual levels of the credit valuation adjustment are not discussed. These will be examined in subsequent sections. At present, only convergence is established. With reference to Tables 9.2 and 9.3, we observe the following:

- (i) Both the payer and the receiver CVA converge fairly rapidly as the number of Monte Carlo simulations is increased.
- (ii) The receiver adjustment converges at a more rapid pace than the payer adjustment. This is attributable to hazard rates being floored at zero. There is thus less scope for variation in the value.
- (iii) There is very little variation in CVA for different values of N . For 100,000 simulations, both tables illustrate that the maximum difference between the CVA obtained for the given values of N is 2bp. This is insignificant in practice. Observe also that, as N increases, there is little change in the adjustment, indicating stability in the numerical integration scheme.

We now perform the identical exercise, but with an increase in the hazard rate volatility levels and some variation in the riskiness of both the counterparty and the reference entity. In Tables 9.4 and 9.5, the value of ν_2 is set to 0.5. In the first scenario, the reference entity is the high risk entity, with a hazard rate volatility of 0.9. The counterparty is the low risk entity. In the second case, the low risk entity is the reference entity (whose hazard rate volatility is set to 0.8) and the counterparty is of intermediate risk. Regarding Tables 9.4 and 9.5, the comments are similar to those made for the previous two tables.

Convergence has now been examined for a diverse set of parameters. Some of the choices (such as a reference entity hazard rate volatility of 0.9) are highly unlikely in practice. However they serve to establish the robustness of the numerical methods utilised in the model. The CVA results shown in Tables 9.2 to 9.5 indicate stability in the model output, together with satisfactory convergence.

For the remainder of the dissertation, a value of $N = 256$ is employed for the FRFT scheme and the number of simulations is set to 100,000, unless otherwise stated. These parameter choices are conservative and could be reduced in practice for the sake of enhanced performance.

Receiver CDS (short protection)											
Simulations	$\rho \setminus N$	Case 1 ($\nu_1 = 0.9; \nu_2 = 0.5$)					Case 2 ($\nu_1 = 0.8; \nu_2 = 0.5$)				
		Ref: High; Cpty: Low					Ref: Low; Cpty: Intermediate				
		128	256	512	1,024	2,048	128	256	512	1,024	2,048
10,000	-0.99	30	31	32	31	29	8	8	8	8	8
	-0.9	30	29	28	30	30	8	8	8	8	8
	-0.75	28	28	27	26	26	8	8	7	7	7
	-0.5	22	20	20	20	19	6	6	6	6	6
	-0.25	12	10	12	11	11	4	4	4	4	4
20,000	-0.99	30	31	31	31	30	8	8	8	8	8
	-0.9	29	29	28	29	30	8	8	8	8	8
	-0.75	27	27	26	26	27	8	8	7	7	7
	-0.5	21	19	20	19	19	6	6	6	6	6
	-0.25	12	11	11	11	12	4	4	4	4	4
40,000	-0.99	31	31	30	31	30	8	8	8	8	8
	-0.9	29	29	29	29	30	8	8	8	8	8
	-0.75	27	27	26	26	26	8	7	7	7	7
	-0.5	21	19	20	20	19	6	6	6	6	6
	-0.25	12	12	11	11	11	4	4	4	4	4
60,000	-0.99	31	31	30	31	30	8	8	8	8	8
	-0.9	29	30	29	29	29	8	8	8	8	8
	-0.75	27	27	26	26	26	8	7	7	7	7
	-0.5	21	19	20	20	19	6	6	6	6	6
	-0.25	12	11	12	11	11	4	4	4	4	4
80,000	-0.99	31	31	31	31	31	8	8	8	8	8
	-0.9	30	29	29	30	30	8	8	8	8	8
	-0.75	27	26	26	26	26	8	7	7	7	7
	-0.5	20	19	20	20	19	6	6	6	6	6
	-0.25	12	12	12	11	11	4	4	4	4	4
100,000	-0.99	31	31	31	31	31	8	8	8	8	8
	-0.9	30	29	29	29	29	8	8	8	8	8
	-0.75	27	26	26	26	26	8	8	7	7	7
	-0.5	20	19	20	20	19	6	6	6	6	6
	-0.25	12	12	12	11	11	4	4	4	4	4

Table 9.4: Receiver CVA convergence for number of Monte Carlo simulations and number of terms in FRFT sum - high ν_1 . *The labels Ref and Cpty refer to the reference entity and counterparty respectively, while low, intermediate and high refer to the riskiness of the entity.*

Payer CDS (short protection)											
Simulations	$\rho \setminus N$	Case 1 ($\nu_1 = 0.9; \nu_2 = 0.5$)					Case 2 ($\nu_1 = 0.8; \nu_2 = 0.5$)				
		Ref: High; Cpty: Low					Ref: Low; Cpty: Intermediate				
		128	256	512	1,024	2,048	128	256	512	1,024	2,048
10,000	0.25	12	12	13	12	12	7	6	7	7	7
	0.5	22	25	25	25	23	13	13	13	13	14
	0.75	33	31	37	32	35	25	23	22	24	25
	0.9	37	40	34	36	38	34	32	32	32	33
	0.99	35	35	33	34	35	43	43	40	43	41
20,000	0.25	12	12	13	12	12	7	6	7	7	7
	0.5	23	24	24	24	24	14	14	13	13	14
	0.75	33	33	34	33	34	24	23	22	23	25
	0.9	37	38	36	36	37	33	32	32	32	32
	0.99	36	34	35	32	36	43	43	41	42	41
40,000	0.25	12	12	13	12	12	7	7	7	7	7
	0.5	23	23	23	24	24	14	14	13	14	14
	0.75	34	33	33	34	34	24	23	23	23	24
	0.9	37	37	36	36	38	33	32	32	32	31
	0.99	35	33	35	34	34	42	42	42	43	40
60,000	0.25	12	12	13	12	12	7	7	7	7	7
	0.5	22	23	23	24	24	14	14	13	14	14
	0.75	33	33	33	34	34	24	23	23	23	24
	0.9	37	37	36	37	37	33	32	32	33	31
	0.99	35	34	34	34	33	42	42	42	43	41
80,000	0.25	12	12	13	12	12	7	7	7	7	7
	0.5	22	23	23	24	23	14	14	13	14	14
	0.75	33	33	33	33	34	24	23	23	23	23
	0.9	37	37	37	37	38	33	32	32	32	31
	0.99	35	34	34	35	34	42	42	42	43	42
100,000	0.25	12	12	12	12	12	7	7	7	7	7
	0.5	22	23	23	23	23	14	14	13	14	14
	0.75	33	33	33	33	33	24	23	23	23	23
	0.9	37	37	37	36	38	33	32	31	33	31
	0.99	35	34	34	35	34	42	42	42	43	42

Table 9.5: Payer CVA convergence for number of Monte Carlo simulations and number of terms in FRFT sum - high ν_1 . The labels Ref and Cpty refer to the reference entity and counterparty respectively, while low, intermediate and high refer to the riskiness of the entity.

9.2 CVA Results

In this section, we demonstrate the effect (on the credit valuation adjustment) of volatility, correlation, CDS maturity and the riskiness level of both the reference entity and the counterparty. We are considering the spread S , derived in (8.40) and (8.41). The counterparty hazard rate volatility (ν_2) is set to 10%.

Before examining the results, recall the following basic principles:

- (i) The holder of protection (payer CDS) benefits from a general increase in the par CDS spreads of the underlying reference entity. This is because the protection becomes more valuable when the probability of the entity on which it is written surviving until maturity of the contract decreases. In other words, the protection buyer is paying a lower premium than would be required were the contract written subsequent to the spread increase.
- (ii) Similarly, the seller of protection benefits from a general reduction in the par CDS spreads of the underlying reference entity.
- (iii) The higher (lower) the replacement cost of the CDS at τ_2 , the larger (smaller) the CVA on the CDS.

The parameters ν_1 (the reference entity hazard rate volatility in the CIR++ framework) and ρ (the correlation between the default times of the reference entity and the counterparty) are varied in the scenarios below.

The first of these parameters was selected due to the fact that the CVA is seen as an option on the mark to market value of the CDS. Volatility of the underlying is a driving factor in option pricing and the reference entity hazard rate is a determining factor in the mark to market of the CDS. The correlation parameter was selected due to the fact that a CCDS with a CDS as the underlying is fundamentally a correlation trade. The importance of capturing the sensitivity to right and wrong way risk in a CDS CVA model (and hence of the correlation effect) should be clear, based on earlier discussions.

9.2.1 Reference Entity Riskier than Counterparty

Consider a 5 year CDS with a quarterly premium payment frequency. The reference entity is the intermediate risk name (as described at the beginning of the chapter) and

the counterparty is the low risk entity. The five year par spread at which we price the counterparty risk-free CDS is thus 255bp. Recall from Table B.2 in Appendix B that (based on the par CDS quotes) the counterparty has a 10% probability of defaulting prior to maturity of the CDS.

Table 9.6 below contains the CVA for both a short and a long protection position. The correlation, ρ , is varied along the columns of the table, while the value of ν_1 is varied along the row dimension. Note that the loss given default (LGD) of 75% has been included in the calculation of the fair spread adjustment in the table¹. This value of the LGD is based on the Bloomberg SOAF recovery rate of 25%. The results in Table 9.6 are illustrated graphically in Figure 9.1 for a selection of ν_1 values.




Figure 9.1: CVA in basis points. Reference entity: intermediate risk; counterparty: low risk.

A few patterns are observed from Table 9.6 and Figure 9.1. Firstly, the CVA of the payer CDS (long protection position) tends to increase with correlation. This is to be expected since, on default of the counterparty, we expect reference entity spreads to increase (decrease) when default times are positively (negatively) correlated (a reference entity default is now more likely). An increase in spreads will cause the payer CDS to rise in value, resulting in a greater loss to the holder of protection at the time of the

¹See Equation (3.1) in Chapter 3 for a reminder of the role played by the LGD in the CVA computation.

Receiver (short protection)													
ν_1 / ρ	-0.99	-0.9	-0.7	-0.5	-0.3	-0.1	0	0.1	0.3	0.5	0.7	0.9	0.99
0.05	8	8	8	7	5	1	0	0	0	0	0	0	0
0.1	8	8	8	7	5	2	0	0	0	0	0	0	0
0.2	8	8	8	7	5	2	1	0	0	0	0	0	0
0.3	8	8	8	7	5	2	1	0	0	0	0	0	0
0.4	8	8	8	7	5	2	1	0	0	0	0	0	0
0.5	8	8	8	7	5	2	1	0	0	0	0	0	0
0.6	8	8	8	7	5	2	1	0	0	0	0	0	0
0.7	8	8	8	7	5	2	1	0	0	0	0	0	0
0.8	8	8	7	7	5	2	1	0	0	0	0	0	0
0.9	8	8	8	7	5	2	1	0	0	0	0	0	0
Payer (long protection)													
ν_1 / ρ	-0.99	-0.9	-0.7	-0.5	-0.3	-0.1	0	0.1	0.3	0.5	0.7	0.9	0.99
0.05	0	0	0	0	0	0	1	3	9	15	22	24	9
0.1	0	0	0	0	0	0	1	3	9	15	21	24	11
0.2	0	0	0	0	0	1	2	3	8	14	20	24	18
0.3	0	0	0	0	0	1	2	3	8	13	20	24	21
0.4	0	0	0	0	1	1	2	3	8	13	19	24	25
0.5	0	0	0	0	1	1	2	3	8	13	19	25	29
0.6	0	0	0	0	1	1	2	3	7	13	19	27	30
0.7	0	0	0	1	1	1	2	3	7	13	19	27	32
0.8	0	0	0	1	1	1	2	3	8	13	20	28	34
0.9	0	0	0	1	1	1	1	3	8	13	20	29	35

Table 9.6: CDS CVA in basis points for a range of correlations and reference entity hazard rate volatilities. Reference entity: intermediate risk; counterparty: low risk.

counterparty default. For this reason, the CVA on the payer CDS tends towards zero for negative correlations.

In contrast, the receiver CDS (short protection position) is of greater value to the seller of protection when spreads decline. The receiver CDS CVA thus tends towards zero for positive correlations. Clearly, purchasing protection on an entity that is positively correlated with the counterparty is a form of wrong way risk, as is the sale of protection to a counterparty whose default time is highly negatively correlated with the reference entity's default time. The results in Table 9.6 and Figure 9.1 show that these risks are captured by the model.

Notice next that when the volatility of the reference entity's hazard rate (ν_1) is par-

ticularly low, the payer CVA begins decreasing for very large correlations. The reason for this is that the scenarios providing a non zero contribution to the CVA are those in which the counterparty defaults prior to the reference entity. Since (in this example) the counterparty is less risky than the reference entity, we expect the reference entity's default time to occur before the counterparty's in general and particularly when the hazard rate volatility of the former is low.

The decrease in CVA, when ν_1 is particularly low, is thus due to the decrease in the number of paths in the Monte Carlo simulation on which the counterparty is the first of the two entities to default. The integrated hazard rate of the reference entity, that is used to determine its default time (see Section 7.2), tends to fall within a smaller band when the volatility is low (see the simulated hazard rates for various volatilities in Figure 7.4), preventing a counterparty default event from occurring first. Notice that, in line with this explanation, for a default time correlation of 99%, the payer CVA is particularly small when ν_1 assumes a value of 0.05. Bear in mind that this scenario, in which the reference entity hazard rate volatility is half of the counterparty's and the default time correlation is nearly unity, is extremely unlikely in practice.

This phenomenon provides us with an added layer in our comprehension of the CVA. That is, the computed CVA does not only decrease when the replacement cost of the underlying contract is reduced, but also depends on the likelihood of the event that the counterparty defaults prior to the reference entity. It points to the joint probabilities of default between the reference entity and the counterparty, a factor that increases the complexity of the CVA on credit derivatives when compared to these computations performed on other asset classes. When BCVA is modelled, the joint default probabilities of the investor and the counterparty influence the CVA, so that the timing of entities' defaults becomes relevant regardless of the asset class considered. For now, however, the focus is on unilateral CVA. Part III of the dissertation deals with the bilateral case.

Observe next that the receiver CDS CVA appears to be capped. In contrast, the payer CVA rises to a significantly higher value. This is attributable to the fact that the receiver CDS increases in value as the survival probabilities of the reference entity increase (spreads drop). Survival probabilities for the intermediate risk name, the reference entity in this case, are already relatively high and are capped at one². They cannot increase

²Alternatively, spreads are floored at zero.

indefinitely, much like a put option being bounded in value.

Lastly, notice that the CVA values around the zero correlation point increase with the level of ν_1 . The lack of correlation between the counterparty and reference entity default times leads us to expect the conditional survival probabilities of the reference entity (conditional on the counterparty default time and the market information available at that time) to be similar to the unconditional probabilities. However, there will be an increase in the number of simulations in which the reference entity survival curve, at the time of the counterparty default, differs from the current forward curve when the volatility of the hazard rate is increased. This reflects the optionality in the CVA value, since volatility of the underlying risk factor drives option prices.

9.2.2 Counterparty Riskier than Reference Entity

Consider the same CDS contract as described in Section 9.2.1, with the exception that the counterparty is the high risk entity from Table 9.1. The reference entity remains the intermediate risk name. It may seem peculiar to purchase protection from an entity that is a fair amount riskier than the one against whose default you are ensuring. However, over time, counterparties may deteriorate and the relative credit qualities of the reference entity and the counterparty may change, a fact of which institutions that purchased protection from Lehman Brothers in 2006 would be well aware.

Based on its par CDS curve, the counterparty now has a 31% probability of defaulting prior to maturity of the CDS contract (see Table B.2 in Appendix B). For the sake of comparison with the previous example, the value of ν_2 (the counterparty hazard rate volatility) remains at 0.1. The results are illustrated in Table 9.7 and Figure 9.2, with the same format as in the previous table employed.

Examining the results in Table 9.7 and Figure 9.2, we observe firstly that the general level of the counterparty risk adjustment is significantly higher than in the previous example. This is in line with our expectations, since the counterparty is now far riskier than in the previous case, even though the level of riskiness of the reference entity has remained the same.

In contrast to the table in the previous scenario, in which the payer CVA decreases for high correlations and low reference entity hazard rate volatilities, it increases in Ta-




Figure 9.2: CDS CVA in basis points. Reference entity: intermediate risk; counterparty: high risk.

ble 9.7. This is a result of the fact that the counterparty is now significantly riskier than the reference entity. Thus, in the extreme case of a 99% correlation, the payer CVA attains its maximum value when the volatility is at its lowest. The counterparty almost always defaults prior to the reference entity in the low volatility case, resulting in a greater number of scenarios that contribute to the non zero part of the CVA. In other words, the reason for the low volatility payer CVA decreasing in the previous example when extreme positive correlations are attained, is reversed in this example.

An interesting observation is that the receiver CVA increases by a small amount when extreme positive correlations are attained. This was not observable in the previous example, in which the reference entity was the intermediate risk name and the counterparty was the low risk one. It appears counterintuitive, since we expect the positive correlation to result in very low conditional survival probabilities for the reference entity once the counterparty has defaulted. Importantly, the payer CVA is still increasing for these high correlations. Our intuition is thus correct and the wrong way risk that we wish to capture by the model is clearly present.

The reason for the small increase in the receiver CDS CVA lies in the way in which the correlation is incorporated into the model. Since the observed values are small, the

Receiver (short protection)													
ν_1 / ρ	-0.99	-0.9	-0.7	-0.5	-0.3	-0.1	0	0.1	0.3	0.5	0.7	0.9	0.99
0.05	30	30	28	22	13	4	0	0	0	0	0	0	3
0.1	31	31	28	22	13	4	1	0	0	0	0	0	3
0.2	30	30	28	22	14	6	2	1	0	0	0	0	3
0.3	30	30	27	21	14	7	3	1	0	0	0	1	3
0.4	31	30	27	21	14	7	3	1	0	0	0	1	3
0.5	30	29	26	21	14	7	3	1	0	0	0	1	3
0.6	30	29	26	21	14	7	3	1	0	0	0	1	2
0.7	29	29	26	21	14	7	3	1	0	0	0	1	2
0.8	29	29	26	21	14	6	2	0	0	0	0	1	2
0.9	30	29	26	21	14	6	2	0	0	0	0	0	2
Payer (long protection)													
ν_1 / ρ	-0.99	-0.9	-0.7	-0.5	-0.3	-0.1	0	0.1	0.3	0.5	0.7	0.9	0.99
0.05	0	0	0	0	0	1	3	9	23	40	64	107	154
0.1	0	0	0	0	0	2	4	9	23	40	63	105	149
0.2	0	0	0	1	1	3	5	10	22	38	59	96	131
0.3	0	0	1	1	2	4	6	10	22	36	55	87	112
0.4	0	1	1	2	3	5	6	10	21	35	53	81	99
0.5	1	1	2	2	3	5	6	9	20	33	51	77	92
0.6	1	1	2	3	3	5	6	9	19	33	50	75	89
0.7	1	2	2	3	3	5	5	8	19	33	49	74	88
0.8	2	2	2	3	3	4	5	8	20	33	50	76	89
0.9	2	2	2	3	3	4	5	8	20	33	50	76	90

Table 9.7: CDS CVA in basis points for a range of correlations and reference entity hazard rate volatilities - Reference entity: intermediate risk; counterparty: high risk

scenarios in which they occur are rather extreme and the explanation is lengthy, we refer the interested reader to Appendix G for the discussion.

9.2.3 A Change in Reference Entity Riskiness

At this point, we have shown the CVA on a 5 year CDS, both when the counterparty is riskier than the reference entity and when it is less risky. We have also observed the effect on the adjustment of altering only the counterparty's risk level. In this section, the counterparty remains the risky entity. However, the reference entity is changed from being the intermediate risk name (as in the previous example) to the low risk entity. The counterparty's default probability is thus the same as before. However, the reference entity's chance of defaulting prior to maturity of the CDS contract has decreased from 16% to 10%. The value of ν_2 remains 0.1. Results are presented in Table 9.8 and Figure 9.3.

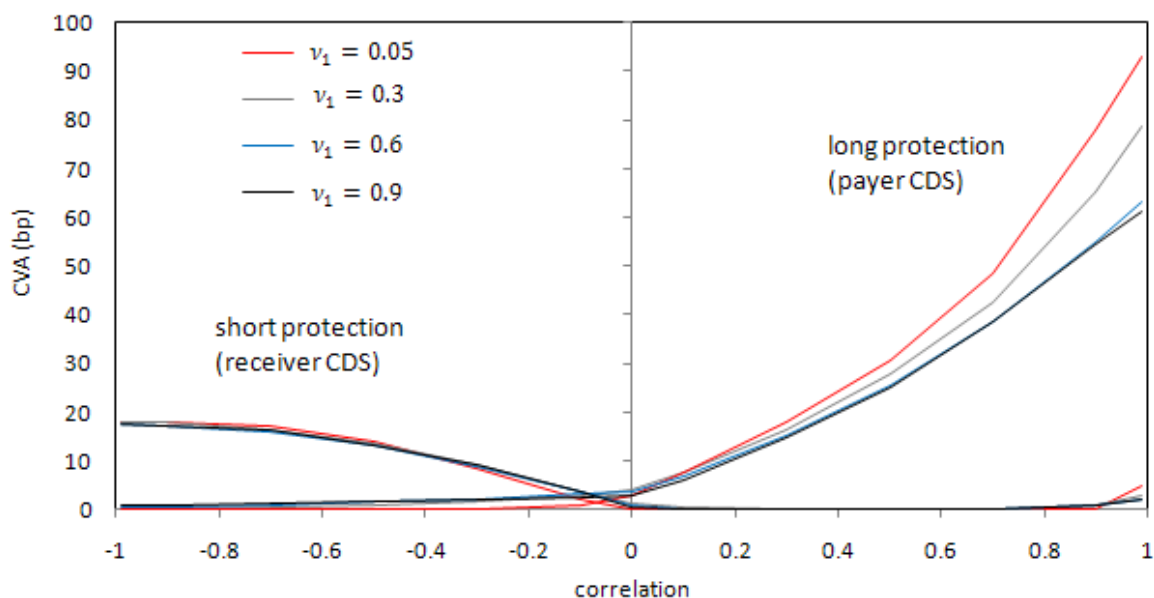


Figure 9.3: CVA in basis points - reference entity: low risk; counterparty: high risk

As expected, the general level of the CVA is lower than in the previous example, reflecting the reduction in the riskiness of the reference entity. This corresponds to the fact that the expected replacement cost of the contract will be lower than in the previous example, due to the lower default risk of the entity on which it is written. It is like an option written on a stock whose value is larger the higher the initial stock value.

Also note that the small increase in the receiver CVA, when correlations are extremely high, is slightly larger for very low values of ν_1 in this example than in the previous one. This reflects the greater differential in counterparty and reference entity spreads (see Appendix G).

In effect, we have now considered both elements of the CVA. The first is the potential future value of the contract, which would traditionally have been calculated via the PFE. The second is the level of riskiness of the counterparty. We see that both of these have an effect. In addition, right and wrong way risk are captured by the model, elements that a simple combination of PEAS and expected default probabilities would fail to reflect.

Receiver (short protection)													
ν_1 / ρ	-0.99	-0.9	-0.7	-0.5	-0.3	-0.1	0	0.1	0.3	0.5	0.7	0.9	0.99
0.05	18	18	17	14	9	2	0	0	0	0	0	0	5
0.1	18	18	17	14	9	3	0	0	0	0	0	0	4
0.2	18	18	17	14	9	3	1	0	0	0	0	1	3
0.3	18	18	17	14	9	4	1	0	0	0	0	1	3
0.4	18	18	17	14	9	4	1	0	0	0	0	1	2
0.5	18	18	16	13	9	4	1	0	0	0	0	1	2
0.6	18	17	16	13	9	4	1	0	0	0	0	1	2
0.7	17	17	16	13	9	4	1	0	0	0	0	1	2
0.8	18	17	16	13	9	3	1	0	0	0	0	1	2
0.9	18	17	16	13	9	4	0	0	0	0	0	1	2
Payer (long protection)													
ν_1 / ρ	-0.99	-0.9	-0.7	-0.5	-0.3	-0.1	0	0.1	0.3	0.5	0.7	0.9	0.99
0.05	0	0	0	0	0	1	3	8	18	31	48	78	93
0.1	0	0	0	0	0	1	3	8	18	31	48	77	94
0.2	0	0	0	0	1	2	4	8	17	29	45	72	89
0.3	0	0	0	1	2	3	4	8	16	28	42	65	79
0.4	0	0	1	1	2	3	4	8	16	27	40	61	73
0.5	0	0	1	1	2	3	4	7	16	26	39	57	66
0.6	1	1	1	2	2	3	4	7	15	26	39	55	63
0.7	1	1	1	2	2	3	3	7	15	25	38	54	62
0.8	1	1	1	2	2	2	3	7	15	25	38	54	61
0.9	1	1	1	2	2	3	3	6	15	25	38	54	61

Table 9.8: CDS CVA - reference entity: low risk; counterparty: high risk

9.2.4 Risky Reference Entity

In this final combination of default risk levels, the reference entity is the risky name from Table 9.1. The CVA, with the low and intermediate risk entities assuming the role of counterparty, is shown in Figures 9.4 and 9.5 respectively. In Section 9.2.1, the reason for the short protection CVA being capped was mentioned. The purpose of this scenario is to demonstrate that the differential between the payer and receiver CVA is not as marked when the reference entity is highly risky. Hazard rates are large and thus have more scope to decrease prior to reaching the floor of zero.

Regarding Figures 9.4 and 9.5, the observable patterns have all been discussed in previous scenarios. We notice that:

- (i) The general level of the counterparty valuation adjustment is lower for the low risk

counterparty (Figure 9.4) than for the intermediate risk counterparty (Figure 9.5), as expected.

- (ii) For extreme positive correlations, the payer CVA begins decreasing. Since the reference entity is much riskier than the counterparty in both examples, the likelihood of a reference entity default occurring prior to a counterparty default minimises the possible loss associated with the latter event. This phenomenon is more pronounced in Figure 9.4 than in Figure 9.5, due to the higher default probability of the counterparty in the second case.




Figure 9.4: CVA in basis points - reference entity: high risk; counterparty: low risk

The situation (in which the reference entity is much riskier than the counterparty) is one that is likely to occur in practice. It is plausible that an investor might hedge its default risk to an extremely risky entity by purchasing protection from a low risk counterparty. It thus illustrates the significance of the receiver CDS CVA, which is sometimes considered to be less noteworthy than the adjustments applied to long protection positions. The reason for this perception is partly due to the fact that positive default time correlations are more likely in practice than negative correlations. However, this does not mean that negative correlations are impossible.




Figure 9.5: CVA in basis points - reference entity: high risk; counterparty: intermediate risk

9.2.5 Differing Maturities

Since all the numerical results up to this point have applied to credit default swaps with maturities of five years, we now compare the differences in CVA for differing CDS maturities. The adjustments for a 1 year, 2 year, 5 year and 10 year par CDS, with an intermediate risk reference entity and a low risk counterparty, are shown in Table 9.9 below and illustrated in Figure 9.6. Although the 5 year CDS CVA was illustrated in Section 9.2.1, it is included for the sake of comparison.

Consistent with the increased risk assumed with an increase in the time for which a position is held, as well as the fact that the credit spread curve is upward sloping, the credit valuation adjustment increases with maturity when all other model inputs remain unchanged.




Figure 9.6: CVA in basis points for various CDS maturities - reference entity: intermediate risk; counterparty: low risk

ρ	CDS Maturity							
	Receiver (short protection)				Payer (long protection)			
	1 year	2 year	5 year	10 year	1 year	2 year	5 year	10 year
-0.99	1	3	8	22	0	0	0	0
-0.9	1	3	8	22	0	0	0	0
-0.7	1	3	8	19	0	0	0	0
-0.5	1	3	7	15	0	0	0	0
-0.3	1	2	5	9	0	0	0	0
-0.1	0	1	2	3	0	0	0	1
0	0	0	0	1	0	1	1	2
0.1	0	0	0	0	1	2	3	5
0.3	0	0	0	0	2	6	9	12
0.5	0	0	0	0	5	11	15	19
0.7	0	0	0	0	9	17	21	27
0.9	0	0	0	0	11	18	24	30
0.99	0	0	0	0	2	6	11	16

Table 9.9: CVA in basis points for various CDS maturities - reference entity: intermediate risk; counterparty: low risk

9.3 Chapter Summary

The objective of this chapter was to exhibit the output produced by the unilateral CDS CVA model, the implementation of which was described in Chapters 6 to 8. The chapter was divided into two parts, the first of which focused on convergence. Both the number of steps in the fractional fast Fourier transform and the quantity of Monte Carlo simulations were included in the analysis. The results reported indicated stability, robustness and convergence.

The second part of the chapter presented the unilateral CVA for a variety of scenarios, which were selected to illustrate various attributes of the adjustment. The reference entity hazard rate volatility and the correlation between the default times of the counterparty and the reference entity were varied in each scenario³. The observations made are listed below.

- (i) The general level of the unilateral CVA increases (decreases) with an increase (decrease) in:
 - (a) the riskiness of the counterparty.
 - (b) the riskiness of the reference entity.
 - (c) the maturity of the underlying credit default swap.
- (ii) **Payer CDS (long protection position)**
 - (a) CVA tends to increase with a rise in default time correlation.
 - (b) For very high default time correlations, CVA decreases with an increase in correlation when the counterparty is less risky than the reference entity. This is due to the increased likelihood of the reference entity defaulting prior to the counterparty.
 - (c) Under large positive default time correlations, the CVA is highest (lowest) for very small ν_1 values when the reference entity is less risky (riskier) than the counterparty. This is a function of the number of scenarios in which the reference entity defaults prior to the counterparty.
- (iii) **Receiver CDS (short protection position)**

³The reasons for selecting these two parameters in the stress testing were discussed at the beginning of Section 9.2.

- (a) CVA tends to decrease with an increase in correlation.
- (b) The exception to this, shown in Sections 9.2.2 and 9.2.3, occurred when the reference entity was less risky than the counterparty and correlations were close to unity. This unlikely situation was analysed in Appendix G.
- (c) CVA values are capped due to the fact that spreads (hazard rates) are floored at zero. When the reference entity is the high risk entity, this is less pronounced, since there is more potential for a decrease in very high hazard rates.

The CVA results produced by the model are consistent with the assumptions that have been made. Importantly, the adjustments reported in this chapter illustrate that the model has captured the nature of the right and wrong way exposures inherent in CDS counterparty risk. The directions of the changes in CVA values when the default risk levels of the two entities (the counterparty and the reference entity) were modified were intuitively accurate.

Part III

A Bilateral CVA Model

Chapter 10

Extending the Model to the Bilateral Case

The concept of bilateral CVA (BCVA) was introduced in Chapter 3, along with an explanation of how it differs from the unilateral adjustment. The investor is the entity from whose perspective the CVA is computed. The entity on the opposite side of the trade is referred to as the counterparty. Essentially, the unilateral CVA computation assigns a zero probability of defaulting to the investor. Only the default risks of the counterparty and the reference entity (whose default triggers the underlying CDS contract) are considered significant. In contrast, the BCVA computation takes into account the default potential of all three parties to the trade (the investor, the reference entity and the counterparty).

From the bilateral version of the General Counterparty Risk Pricing Formula (Equation (3.3) in Chapter 3) the value of the BCVA is as follows:

$$\begin{aligned}
 \text{BCVA}(t) &= \underbrace{\text{LGD}_2 \cdot \mathbb{E}_t\{1_{C \cup D} \cdot D(t, \tau_2) \cdot [\text{NPV}(\tau_2)]^+\}}_{\text{CCVA}} \\
 &\quad - \underbrace{\text{LGD}_0 \cdot \mathbb{E}_t\{1_{A \cup B} \cdot D(t, \tau_0) \cdot [-\text{NPV}(\tau_0)]^+\}}_{\text{DVA}}. \tag{10.1}
 \end{aligned}$$

The notation used in (10.1) was defined in Section 3.3 of Chapter 3. The signs preceding the terms in (10.1) are the opposite of those in (3.3). The reason for this is that the BCVA in (10.1) is expressed on a stand-alone basis. This was being subtracted from the risk-free value of a derivative in (3.3), in order to obtain its risky value (the value of the derivative transacted between risk-free entities less the BCVA).

Briefly, the first term in (10.1) is the portion of the BCVA that is attributable to the counterparty's risk of default. It is referred to as the CCVA in this dissertation since it differs from the original unilateral adjustment. This is due to the relevance of the investor's survival probability to the BCVA calculation (in contrast to the unilateral case). At τ_2 (the counterparty default time), the filtration \mathcal{G}_{τ_2} contains information pertaining to the investor's hazard rate in addition to those of the counterparty and the reference entity as in the unilateral case. This will be discussed in greater detail in Section 10.4.

The DVA term accounts for the investor's own default risk. The CCVA (DVA) seen from the investor's perspective is the DVA (CCVA) from the counterparty's viewpoint. The DVA term is subtracted from the CCVA term, since the counterparty risk faced by the investor is reduced by an amount attributable to its own risk. The various ways in which the components of the BCVA (particularly the DVA) may be understood, were discussed in some detail in Chapter 3. In this chapter, the focus is on an explanation of the modifications to the unilateral CDS CVA model already implemented, in order to compute BCVA.

In the following section, the notation that was pertinent within the unilateral framework is supplemented with that which is required for the bilateral extension. The subsequent section, Section 10.2, reviews the chief assumptions made in the bilateral model. These are merely extensions of those outlined in Section 7.2. The alterations made to each main step of the CDS CVA computation, to account for the extension from a unilateral to a bilateral adjustment, constitute the remainder of the chapter. These core steps were summarised in Figure 6.2.

10.1 Notation

In Section 6.3.1, the notation relevant to the unilateral model was introduced. The subscript 1 was employed when referring to the reference entity, while the subscript 2 was linked to the counterparty. The additional subscript 0 will be associated with the investor. The default times of the investor, the reference entity and the counterparty are thus denoted by τ_0 , τ_1 and τ_2 respectively. Their respective intensities are represented by the symbols λ_0 , λ_1 and λ_2 .

The integrated intensities are then:

$$\Lambda_0(\tau_0) = \int_0^{\tau_0} \lambda_0(s)ds := \xi_0, \quad \Lambda_1(\tau_1) = \int_0^{\tau_1} \lambda_1(s)ds := \xi_1 \quad \text{and} \quad \Lambda_2(\tau_2) = \int_0^{\tau_2} \lambda_2(s)ds := \xi_2, \quad (10.2)$$

where the default times (obtained by inverting the equations in (10.2)) are:

$$\tau_0 = \Lambda_0^{-1}(\xi_0), \quad \tau_1 = \Lambda_1^{-1}(\xi_1) \quad \text{and} \quad \tau_2 = \Lambda_2^{-1}(\xi_2). \quad (10.3)$$

In addition, the uniform random variables,

$$U_0 := 1 - e^{-\xi_0}, \quad U_1 := 1 - e^{-\xi_1} \quad \text{and} \quad U_2 := 1 - e^{-\xi_2}, \quad (10.4)$$

are defined in the same manner as in Section 6.3.1.

Since there are now two possible entities whose default events are significant to the revaluation of the underlying CDS, the notation $\bar{U}_1 = 1 - e^{-\Lambda_1(\tau_2)}$ of Part II, is replaced by the more general

$$\bar{U}_{i,j} = 1 - e^{-\Lambda_i(\tau_j)} \quad i = 0, 1, 2, \quad j = 1, 2. \quad (10.5)$$

This is the notation that was utilised in [22]. When the relevant defaulting party is the investor, the particular cases of (10.5) that are pertinent are $\bar{U}_{1,0} = 1 - e^{-\Lambda_1(\tau_0)}$ and $\bar{U}_{2,0} = 1 - e^{-\Lambda_2(\tau_0)}$, the uniform random variables related to the two surviving entities at the investor's default time. Similarly, when the counterparty defaults prior to the reference entity, the investor and the CDS maturity, the significant instances of (10.5) are $\bar{U}_{1,2} = 1 - e^{-\Lambda_1(\tau_2)}$ and $\bar{U}_{0,2} = 1 - e^{-\Lambda_0(\tau_2)}$.

10.2 Assumptions

The assumptions underlying the unilateral CVA model were outlined in Section 6.3.2. These continue to hold with the following additions:

- (i) The intensity processes of all three entities (investor, reference entity and counterparty) are independent of each other.
- (ii) The uniform random variables U_0 , U_1 and U_2 are correlated through the trivariate copula function:

$$\mathbb{Q}(U_0 < u_0, U_1 < u_1, U_2 < u_2) = C(u_0, u_1, u_2). \quad (10.6)$$

In the unilateral case, a bivariate copula was required, since the default times of only two entities (the reference entity and the counterparty) were relevant. The addition of the investor's default time necessitates the introduction of a trivariate copula function.

- (iii) The copula in the implementation is the Gaussian copula. The correlation parameters ($\rho_{0,1}$, $\rho_{1,2}$ and $\rho_{0,2}$) are assumed to be constant. These represent the correlation between the uniform random variables associated with the default times of the investor and the reference entity, the reference entity and the counterparty and the investor and the counterparty respectively. Thus, there are two additional correlation parameters.

The vector $\mathbf{u} := [u_0, u_1, u_2]$. The definition of the trivariate Gaussian copula (see [89]) is:

$$C(u_0, u_1, u_2) = \int_{-\infty}^{\Phi^{-1}(u_0)} \int_{-\infty}^{\Phi^{-1}(u_1)} \int_{-\infty}^{\Phi^{-1}(u_2)} \frac{1}{(2\pi)^{3/2} |\boldsymbol{\Sigma}|^{1/2}} e^{-\mathbf{u}'\boldsymbol{\Sigma}^{-1}\mathbf{u}/2}, \quad (10.7)$$

where the covariance matrix, $\boldsymbol{\Sigma}$, is defined as

$$\boldsymbol{\Sigma} = \begin{pmatrix} 1 & \rho_{0,1} & \rho_{0,2} \\ \rho_{0,1} & 1 & \rho_{1,2} \\ \rho_{0,2} & \rho_{1,2} & 1 \end{pmatrix}.$$

The symbol ' indicates the transpose of a vector or matrix and $||$ signifies the matrix determinant¹. The correlation matrix $\boldsymbol{\Sigma}$ is required to be positive definite [90]. Mathematically, this ensures that the matrix is invertible. Positive definiteness can be established by ensuring that the eigenvalues of the matrix are positive [89] or that its determinant is positive [83]. Practically, the combinations of correlations implied by a matrix that is not positive definite are inconsistent with each other. For example, if $\rho_{0,1} = \rho_{1,2} = -0.9$, it would be inconsistent for $\rho_{0,2}$ to be 0.6.

Note that a trivariate Gaussian copula induces bivariate marginals. The proof of this well-known result is presented Section D.7 of Appendix D.

- (iv) The recovery rates of the investor, the reference entity and the counterparty, denoted by R_0 , R_1 and R_2 respectively, are assumed to be deterministic constants. The

¹Refer to [83] for a definition of the determinant of a matrix.

associated loss given default parameters are LGD_0 , LGD_1 and LGD_2 .

In addition to the above assumptions, recall from Section 6.3.2 that the risk-free interest rate is assumed to be deterministic, precluding the possibility of a dependence structure between this rate and the default times of any of the entities that feature in the BCVA model.

The intensity process of the investor is a CIR++ process, the calibration of which was explained thoroughly in Section 7.1. The intensities of the two entities relevant to the unilateral CDS CVA model were assumed to be strictly positive and calibrated in such a way as to ensure this. Similarly, the investor's intensity remains greater than zero at all times. Since the calibration of the intensity process of each entity is performed independently of the intensities of the other entities, this procedure is not altered by the extension to the bilateral case. There is merely an additional process (belonging to the investor) that is calibrated.

The second step in the model outline in Figure 6.2 is the subject of the next section in this chapter.

10.3 Simulating Default Times

Within the unilateral framework, Section 7.2 explained the procedure for simulating the default times of the entities and identifying counterparty default times that were relevant for the CDS revaluation. The steps were as follows:

- (i) Simulate K paths of both λ_1 and λ_2 .
- (ii) Integrate the simulated intensities to obtain Λ_1 and Λ_2 .
- (iii) Generate two $K \times 1$ vectors of uniform random numbers via the copula $C(u_1, u_2)$.
- (iv) Transform the uniform numbers into exponential random numbers.
- (v) Compare these exponential random numbers to the integrated intensities to determine whether and at which times either entity has defaulted along each of the K paths.
- (vi) The relevant counterparty default times (those at which the replacement cost of the underlying credit default swap is required to be computed) are the ones that precede

both a reference entity default and the CDS maturity date.

In the bilateral case, the process is slightly more computationally intensive. However, it is merely a trivariate extension of the bivariate unilateral version. Each of the steps corresponding to the ones listed above are discussed in turn:

- (i) Simulate K paths of each of the three entities' intensities, namely λ_0 , λ_1 and λ_2 . The procedure for this was explained in much detail in Section 7.2.1.
- (ii) Integrate the simulated intensities in order to obtain Λ_0 , Λ_1 and Λ_2 .
- (iii) Generate three $K \times 1$ vectors of uniform random numbers using the trivariate Gaussian copula that was defined in (10.7). The simplest method is to utilise the Matlab *copularnd* function. Alternatively, the procedure in Figure 10.1 can be invoked [91].

for K simulations

Step 1

Generate $\mathbf{Z} \sim N_3(\mathbf{0}, \Sigma)$ where the subscript 3 indicates a trivariate distribution.

- (i) Perform a Cholesky decomposition of Σ to obtain the Cholesky factor $\Sigma^{1/2}$.
- (ii) Generate a vector $\mathbf{X} = (X_1, X_2, X_3)'$ of independent standard normal variates.
- (iii) Set $\mathbf{Z} = \Sigma^{1/2}\mathbf{X}$.

Step 2

Set $U_0 = \Phi(Z_1)$, $U_1 = \Phi(Z_2)$ and $U_2 = \Phi(Z_3)$ (following our notation).

Output U_0 , U_1 and U_2 .

end

Figure 10.1: Simulation of uniform random numbers from a trivariate Gaussian copula

- (iv) Transform the uniform random numbers, that were generated in the previous step, into standard exponential random numbers using the equation

$$\Lambda_j(\tau_j) = -\ln(1 - U_j), \quad j = 0, 1, 2. \quad (10.8)$$

- (v) For each of the three entities, compare the generated exponential random numbers to the integrated intensities obtained from the simulated hazard rates. Default occurs when the exponential random number is greater than or equal to the integrated intensity [25].
- (vi) Compare the default times of the three entities in order to obtain two sets. The first is the collection of default times that will be used to calculate the DVA portion of the BCVA. It includes investor default times that precede both reference entity and counterparty defaults, as well as the CDS maturity. The second set contains the times at which the counterparty defaults prior to the other two entities and to the CDS maturity date. The revaluation of the CDS at these dates will be used in the CCVA computation.

10.4 Revaluation of the CDS upon Default

Step 3 of the model outline in Figure 6.2 has now been reached. As mentioned previously in the dissertation, the valuation of a credit default swap amounts to the computation of the survival probabilities of the CDS reference entity at the valuation date. Once these have been determined, the pricing of the CDS is trivial. The focus in this section is therefore on the calculation of the survival probabilities of the reference entity upon a relevant counterparty or investor default event. CDS pricing was covered in Section 4.3.

In the unilateral model, there was a set of times corresponding to the significant counterparty default times. In the bilateral model, there are two sets of default events. The first of these contains the relevant investor default times. The CDS values as these points will impact the value of the DVA, the second term in (10.1). The revaluation of the CDS at the set of relevant counterparty default times will contribute to the size of the first term in (10.1), the CCVA.

The expression for the conditional survival probability of the reference entity, based on the counterparty default time and the information available in the market at that point in time, was derived in Chapter 8. The probability of surviving to a time T_k in the CDS premium payment schedule that occurs after τ_2 was then approximated in (8.16) as:

$$\mathbb{Q}(\tau_1 > T_k | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2) \approx \sum_j \frac{p_{j+1} + p_j}{2} \Delta f_j, \quad (10.9)$$

where $f_j = C_{1|2}(u_j, U_2) := \mathbb{Q}(U_1 < u_j | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2)$. The remainder of the notation was defined and utilised extensively in Chapter 8.

Since there are three entities relevant to the BCVA computation, the bivariate copula of Part II was extended to a trivariate copula in Section 10.2. The function f_j in (10.9) is thus replaced by:

$$f_j = C_{1|0,2}(u_j, U_2) := \mathbb{Q}(U_1 < u_j | \mathcal{G}_{\tau_2}, U_1 > \bar{U}_{1,2}, U_0 > \bar{U}_{0,2}), \quad (10.10)$$

when the CDS is being revalued under a counterparty default event, and by:

$$f_j = C_{1|2,0}(u_j, U_0) := \mathbb{Q}(U_1 < u_j | \mathcal{G}_{\tau_0}, U_1 > \bar{U}_{1,0}, U_2 > \bar{U}_{2,0}), \quad (10.11)$$

when the relevant default time belongs to the set of investor defaults [22]. Comparing the conditional copula distribution in (10.10) to the expression for $C_{1|2}(u_j, U_2)$ in the unilateral model, the additional information relating to the survival of the investor is incorporated into the former by the inequality, $U_0 > \bar{U}_{0,2}$. The notation $\bar{U}_{i,j}$ was defined in equation (10.5).

The remainder of this section of the chapter is divided into two parts. The first of these focuses on the computation of the p_j 's in (10.9). The second part is centred on the computation of the conditional copula functions; that is, the f_j 's in (10.10) and (10.11).

10.4.1 The Fractional Fast Fourier Transform

The implementation of and justification behind the fractional fast Fourier transform, which was employed in the determination of the p_j 's in (10.9), was explained thoroughly in Chapter 8 for the unilateral model. These probabilities represent the cumulative distribution function of the integrated CIR process of the reference entity. The inputs to the FRFT process were related to the counterparty default times only insofar as the values of the reference entity process were required at τ_2 .

Therefore, no alteration to the FRFT implementation is required in extending the unilateral model to the bilateral case. When the default time at which the CDS value is being computed belongs to the set of investor default times, the spot CIR intensity that is input to the FRFT process is the one prevailing at τ_0 . Similarly, for a counterparty default event, the spot CIR intensity of the reference entity at τ_2 is required. We now

proceed to the computation of the f'_j s in (10.10) and (10.11).

10.4.2 The Conditional Copula Function

Consider the conditional copula function, $C_{1|0,2}(u_j, U_2)$, in (10.10). When the CDS is being revalued at a significant counterparty default time, this is the expression that will be utilised for f_j in (10.9). Similarly, in the event of a relevant investor default, the conditional copula function that is represented by f_j in (10.9) is $C_{1|2,0}(u_j, U_0)$ (see (10.11)).

It is not immediately obvious how these two functions will be computed in practice. We illustrate the manipulation of (10.10) into a form that is practical to implement. This was demonstrated in [22]. The manipulation of (10.11) follows in a similar manner.

Consider the conditional copula function:

$$C_{1|0,2}(u_j, U_2) := \mathbb{Q}(U_1 < u_j | \mathcal{G}_{\tau_2}, U_1 > \bar{U}_{1,2}, U_0 > \bar{U}_{0,2}). \quad (10.12)$$

Since U_2 is known at τ_2 , the expression in (10.12) may be rewritten as:

$$\begin{aligned} C_{1|0,2}(u_j, U_2) &= \mathbb{Q}(U_1 < u_j | U_2, U_1 > \bar{U}_{1,2}, U_0 > \bar{U}_{0,2}) \\ &= \frac{\mathbb{Q}(U_1 < u_j, U_1 > \bar{U}_{1,2} | U_2, U_0 > \bar{U}_{0,2})}{\mathbb{Q}(U_1 > \bar{U}_{1,2} | U_2, U_0 > \bar{U}_{0,2})}. \end{aligned} \quad (10.13)$$

The fraction in (10.13) follows from the definition of conditional probability and has been utilised a number of times in this dissertation.

We wish to express the conditional copula function in terms of the trivariate copula in (10.6) and its bivariate marginals. This is similar to the approach that was pursued in Section 8.1.1 for the unilateral model. Consider firstly the numerator in (10.13). It can

be manipulated as follows:

$$\begin{aligned} \text{Numerator} &= \mathbb{Q}(U_1 < u_j, U_1 > \bar{U}_{1,2} | U_2, U_0 > \bar{U}_{0,2}) \\ &= \frac{\mathbb{Q}(U_1 < u_j, U_1 > \bar{U}_{1,2}, U_0 > \bar{U}_{0,2} | U_2)}{\mathbb{Q}(U_0 > \bar{U}_{0,2} | U_2)} \end{aligned} \quad (10.14)$$

$$= \frac{\underbrace{\mathbb{Q}(U_1 < u_j, U_0 > \bar{U}_{0,2} | U_2)}_{\text{term } a} - \underbrace{\mathbb{Q}(U_1 < \bar{U}_{1,2}, U_0 > \bar{U}_{0,2} | U_2)}_{\text{term } b}}{\underbrace{\mathbb{Q}(U_0 > \bar{U}_{0,2} | U_2)}_{\text{term } c}}. \quad (10.15)$$

Equation (10.14) again follows from the definition of conditional probability. The numerator in (10.15) expresses the fact that the values that may be assumed by the random variable U_1 are floored at $\bar{U}_{1,2}$ and capped at u_j . The terms in (10.15) have been labelled for ease of reference. We now consider the computation of each of these terms in turn.

$$\text{term } a = \mathbb{Q}(U_1 < u_j, U_0 > \bar{U}_{0,2} | U_2) \quad (10.16)$$

$$= \mathbb{Q}(U_1 < u_j | U_2) - \mathbb{Q}(U_1 < u_j, U_0 < \bar{U}_{0,2} | U_2) \quad (10.17)$$

$$= \left. \frac{\partial C_{1,2}(u_j, u_2)}{\partial u_2} \right|_{u_2=U_2} - \left. \frac{\partial C(\bar{U}_{0,2}, u_j, u_2)}{\partial u_2} \right|_{u_2=U_2}. \quad (10.18)$$

Equation (10.17) is an alternative expression to (10.16) of the fact that the values that may be taken on by the random variable U_0 are floored at $\bar{U}_{0,2}$ and capped at 1. Since the random variable is less than or equal to unity with probability one, $\mathbb{Q}(U_1 < u_1 | U_2) = \mathbb{Q}(U_1 < u_1, U_0 < 1 | U_2)$. The second step in the manipulation of term a , (10.18), follows from the definition of the copula conditional distribution in Section D.6 in Appendix D.

Similarly:

$$\begin{aligned} \text{term } b &= \mathbb{Q}(U_1 < \bar{U}_{1,2}, U_0 > \bar{U}_{0,2} | U_2) \\ &= \mathbb{Q}(U_1 < \bar{U}_{1,2} | U_2) - \mathbb{Q}(U_1 < \bar{U}_{1,2}, U_0 < \bar{U}_{0,2} | U_2) \end{aligned} \quad (10.19)$$

$$= \left. \frac{\partial C_{1,2}(\bar{U}_{1,2}, u_2)}{\partial u_2} \right|_{u_2=U_2} - \left. \frac{\partial C(\bar{U}_{0,2}, \bar{U}_{1,2}, u_2)}{\partial u_2} \right|_{u_2=U_2}, \quad (10.20)$$

and, trivially:

$$\text{term } c = \mathbb{Q}(U_0 > \bar{U}_{0,2} | U_2) = 1 - \mathbb{Q}(U_0 \leq \bar{U}_{0,2} | U_2) = 1 - \left. \frac{\partial C_{0,2}(\bar{U}_{0,2}, u_2)}{\partial u_2} \right|_{u_2=U_2}. \quad (10.21)$$

Substituting (10.18), (10.20) and (10.21) into (10.15) and simplifying, the numerator of the conditional copula in (10.13) is given by:

Numerator

$$= \frac{\frac{\partial C_{1,2}(u_j, u_2)}{\partial u_2} \Big|_{u_2=U_2} - \frac{\partial C(\bar{U}_{0,2}, u_j, u_2)}{\partial u_2} \Big|_{u_2=U_2} - \frac{\partial C_{1,2}(\bar{U}_{1,2}, u_2)}{\partial u_2} \Big|_{u_2=U_2} + \frac{\partial C(\bar{U}_{0,2}, \bar{U}_{1,2}, u_2)}{\partial u_2} \Big|_{u_2=U_2}}{1 - \frac{\partial C_{0,2}(\bar{U}_{0,2}, u_2)}{\partial u_2} \Big|_{u_2=U_2}}. \quad (10.22)$$

Turning now to the denominator in (10.13), we assign labels to the terms of which it is comprised:

$$\text{Denominator} = \mathbb{Q}(U_1 > \bar{U}_{1,2} | U_2, U_0 > \bar{U}_{0,2}) = \frac{\underbrace{\mathbb{Q}(U_1 > \bar{U}_{1,2}, U_0 > \bar{U}_{0,2} | U_2)}_{\text{term } d}}{\underbrace{\mathbb{Q}(U_0 > \bar{U}_{0,2} | U_2)}_{\text{term } e}}. \quad (10.23)$$

Since term e is identical to term c , we focus on term d . The techniques employed in simplifying this term have all been used previously in this section and require no additional explanation.

$$\begin{aligned} \text{term } d &= \mathbb{Q}(U_1 > \bar{U}_{1,2}, U_0 > \bar{U}_{0,2} | U_2) \\ &= \mathbb{Q}(U_0 > \bar{U}_{0,2} | U_2) - \mathbb{Q}(U_0 > \bar{U}_{0,2}, U_1 < \bar{U}_{1,2} | U_2) \\ &= [1 - \mathbb{Q}(U_0 < \bar{U}_{0,2} | U_2)] \\ &\quad - [\mathbb{Q}(U_1 < \bar{U}_{1,2} | U_2) - \mathbb{Q}(U_0 < \bar{U}_{0,2}, U_1 < \bar{U}_{1,2} | U_2)] \\ &= 1 - \mathbb{Q}(U_0 < \bar{U}_{0,2} | U_2) - \mathbb{Q}(U_1 < \bar{U}_{1,2} | U_2) \\ &\quad + \mathbb{Q}(U_0 < \bar{U}_{0,2}, U_1 < \bar{U}_{1,2} | U_2) \\ &= 1 - \frac{\partial C_{0,2}(\bar{U}_{0,2}, u_2)}{\partial u_2} \Big|_{u_2=U_2} - \frac{\partial C_{1,2}(\bar{U}_{1,2}, u_2)}{\partial u_2} \Big|_{u_2=U_2} \\ &\quad + \frac{\partial C(\bar{U}_{0,2}, \bar{U}_{1,2}, u_2)}{\partial u_2} \Big|_{u_2=U_2}. \end{aligned} \quad (10.24)$$

Next, we substitute (10.24) into the expression for the denominator in (10.23). This, together with the numerator in (10.22), is then substituted into the conditional copula in (10.13). Terms c and e cancel, and the resultant expression (which is the same as the

one derived in [22]) is:

$$\begin{aligned}
 & C_{1|0,2}(u_j, U_2) \\
 &= \frac{\frac{\partial C_{1,2}(u_j, u_2)}{\partial u_2} \Big|_{u_2=U_2} - \frac{\partial C(\bar{U}_{0,2}, u_j, u_2)}{\partial u_2} \Big|_{u_2=U_2} - \frac{\partial C_{1,2}(\bar{U}_{1,2}, u_2)}{\partial u_2} \Big|_{u_2=U_2} + \frac{\partial C(\bar{U}_{0,2}, \bar{U}_{1,2}, u_2)}{\partial u_2} \Big|_{u_2=U_2}}{1 - \frac{\partial C_{0,2}(\bar{U}_{0,2}, u_2)}{\partial u_2} \Big|_{u_2=U_2} - \frac{\partial C_{1,2}(\bar{U}_{1,2}, u_2)}{\partial u_2} \Big|_{u_2=U_2} + \frac{\partial C(\bar{U}_{0,2}, \bar{U}_{1,2}, u_2)}{\partial u_2} \Big|_{u_2=U_2}}.
 \end{aligned} \tag{10.25}$$

Similarly, the expression for f_j that applies in the event of an investor default, (10.11), is expressed as:

$$\begin{aligned}
 & C_{1|2,0}(u_j, U_0) \\
 &= \frac{\frac{\partial C_{0,1}(u_0, u_j)}{\partial u_0} \Big|_{u_0=U_0} - \frac{\partial C(u_0, u_j, \bar{U}_{2,0})}{\partial u_0} \Big|_{u_0=U_0} - \frac{\partial C_{0,1}(u_0, \bar{U}_{1,0})}{\partial u_0} \Big|_{u_0=U_0} + \frac{\partial C(u_0, \bar{U}_{1,0}, \bar{U}_{2,0}, u_2)}{\partial u_0} \Big|_{u_0=U_0}}{1 - \frac{\partial C_{0,2}(u_0, \bar{U}_{2,0})}{\partial u_0} \Big|_{u_0=U_0} - \frac{\partial C_{0,1}(u_0, \bar{U}_{1,0})}{\partial u_0} \Big|_{u_0=U_0} + \frac{\partial C(u_0, \bar{U}_{1,0}, \bar{U}_{2,0})}{\partial u_0} \Big|_{u_0=U_0}}.
 \end{aligned} \tag{10.26}$$

Implementation of the Conditional Copula Function

In Section 8.3.1, we explained how the conditional distribution function of the bivariate Gaussian distribution could be used to determine the copula conditional distribution. The methodology that was outlined there can be applied to the bivariate marginals in (10.25) and (10.26). For the trivariate terms, a generalisation of the methodology is required.

We begin with the conditional distribution of the multivariate normal distribution. This can be found in [89]. We apply it to the specific case of a trivariate normal distribution, as outlined in Figure 10.2 below.

To compute the probability $\mathbb{Q}(U_0 < u_0, U_1 < u_1 | U_2)$, we apply formulae (10.27) and (10.28) in Figure (10.2) to obtain the mean and covariance matrix respectively of the distribution of $[\Phi^{-1}(U_0), \Phi^{-1}(U_1)] | \Phi^{-1}(U_2)$ with $\mu = \mathbf{0}$ and

$$\Sigma = \begin{pmatrix} 1 & \rho_{0,1} & \rho_{0,2} \\ \rho_{0,1} & 1 & \rho_{1,2} \\ \rho_{0,2} & \rho_{1,2} & 1 \end{pmatrix}.$$

Suppose that the 3-dimensional random vector $\mathbf{X} = [X_1, X_2, X_3]$ has a normal distribution with mean

$$\mu = \begin{pmatrix} \mu_1 \\ \mu_2 \\ \mu_3 \end{pmatrix}$$

and covariance matrix

$$\Sigma = \begin{pmatrix} \sigma_{11} & \sigma_{12} & \sigma_{13} \\ \sigma_{21} & \sigma_{22} & \sigma_{23} \\ \sigma_{31} & \sigma_{32} & \sigma_{33} \end{pmatrix};$$

that is

$$\mathbf{X} \sim N_3(\mu, \Sigma).$$

If μ is partitioned as:

$$\mu = \begin{pmatrix} \underline{\mu}_1 \\ \underline{\mu}_2 \end{pmatrix}, \text{ where } \underline{\mu}_1 = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} \text{ and } \underline{\mu}_2 = \mu_3,$$

and Σ is partitioned as:

$$\Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix},$$

where

$$\Sigma_{11} = \begin{pmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{21} & \sigma_{22} \end{pmatrix}, \Sigma_{12} = \begin{pmatrix} \sigma_{13} \\ \sigma_{23} \end{pmatrix}, \Sigma_{21} = (\sigma_{31} \quad \sigma_{32}) \text{ and } \Sigma_{22} = \sigma_{33},$$

then the distribution of $\mathbf{X} = [X_1, X_2]$ conditional on $X_3 = x_3$ is normal with mean

$$\bar{\mu} = \underline{\mu}_1 + \Sigma_{11}\Sigma_{22}^{-1}(x_3 - \underline{\mu}_2) \quad (10.27)$$

and covariance matrix

$$\bar{\Sigma} = \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21}. \quad (10.28)$$

Figure 10.2: The conditional distribution of trivariate Gaussian random variables

The probability can then be calculated by evaluating the bivariate cumulative distribution function, with mean $\mathbf{0}$ and covariance matrix $\bar{\Sigma}$ (in (10.28)), at the point $[\Phi^{-1}(u_0), \Phi^{-1}(u_1)]$.

In the event of an investor default, when we wish to compute the probability $\mathbb{Q}(U_1 < u_1, U_2 < u_2 | U_0)$, we partition the covariance matrix, Σ , as follows:

$$\Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix},$$

where

$$\Sigma_{11} = \sigma_{11}, \quad \Sigma_{12} = \begin{pmatrix} \sigma_{12} & \sigma_{13} \end{pmatrix}, \quad \Sigma_{21} = \begin{pmatrix} \sigma_{21} \\ \sigma_{31} \end{pmatrix}, \quad \Sigma_{22} = \begin{pmatrix} \sigma_{22} & \sigma_{23} \\ \sigma_{32} & \sigma_{33} \end{pmatrix}.$$

Applying these results to the trivariate terms in (10.25) and (10.26), along with those in Section 8.3.1 to the bivariate terms, the computation of the f_j 's is straightforward.

10.4.3 CDS Valuation

The reference entity's survival curve is obtained at each relevant counterparty default time by computing the probability $\mathbb{Q}(\tau_1 > T_k | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2, \tau_0 > \tau_2)$ for each T_k that occurs in the CDS premium payment schedule subsequent to τ_2 . Using this survival curve (together with the spread at which the original CDS trade was transacted, the risk-free discount factors and the reference entity's recovery rate) the CDS is revalued at τ_2 .

Similarly, at each τ_0 in the set of relevant investor default times, the survival probabilities $\mathbb{Q}(\tau_1 > T_k | \mathcal{G}_{\tau_0}, \tau_1 > \tau_0, \tau_2 > \tau_0)$ are used to compute the reference entity's survival curve at τ_0 . This is employed in the revaluation of the CDS at τ_0 .

10.4.4 The Bilateral CVA Computation

We have now obtained a set of CDS values that correspond to the relevant investor default times, as well as a set for the significant counterparty default times. Step 4 of the model outline in Figure 6.2, *CVA Computation*, remains. This is straightforward since the CDS values at the default times are known.

- (i) Discount the CDS values to the date at which the CVA is being computed. In other words, obtain the present values.
- (ii) Sum all of the positive CDS present values in the set of counterparty defaults and divide these by the total number of simulations. This is to obtain the CCVA portion of the BCVA. Multiply this result by the LGD of the counterparty.
- (iii) Obtain the sum of all of the negative CDS present values in the set of investor defaults. Divide this value by the total number of simulations. Multiply the result by the investor's LGD. This is the DVA portion of the BCVA.

- (iv) Now add the results from Steps (ii) and (iii). This is the BCVA. Note that the value obtained in (iii) is negative. The DVA is thus being subtracted from the CCVA.
- (v) Express the CVA in basis points using the result that was derived in (8.40) in Section 8.4 for a receiver CDS. Utilise the result in (8.41) for a long protection position.

10.5 Chapter Summary

The unilateral model was explained in detail in Part II. The extension of the model to the bilateral case was the subject of this chapter. The main steps in the valuation of the bilateral adjustment are the same as in the unilateral case and were outlined in the diagram in Figure 6.2. The chief modelling and implementation differences that were described, are listed below.

- (i) There are now three entities whose hazard rate processes are simulated. There are therefore three par CDS term structures that are stripped in order to obtain hazard rate term structures to which CIR++ processes are fitted. This is in contrast to the two hazard rate processes in the unilateral model.
- (ii) Since there is an additional name involved, a trivariate copula (rather than a bivariate copula) is necessary to link the uniform random variables of the exponential triggers of the default times of the three entities.
- (iii) In determining the counterparty default times that are relevant to the CVA computation, the additional requirement that τ_2 occurs prior to τ_0 (the investor default time) exists.
- (iv) In addition to the relevant counterparty default times, there is a set of investor default times.
- (v) In revaluing the underlying CDS at the significant default times, the conditional copula function is derived from the trivariate copula and depends on which of the two entities, the investor or the counterparty, has defaulted.
- (vi) The CVA is the difference between two expected values, the CCVA and the DVA, rather than a single expected value as in the unilateral case.

Chapter 11

Results of the Bilateral Model Implementation

This chapter contains various examples of the bilateral credit valuation adjustment on a credit default swap. The results were obtained by means of the methodology detailed in the previous chapter. The CDS BCVA model is merely an extension of the unilateral version. Since the features of the unilateral CVA were discussed thoroughly in Chapter 9, the focus of this chapter is on the characteristics particular to the bilateral adjustment.

There are now three separate default time correlation parameters to consider:

- (i) between the investor and the counterparty, denoted by $\rho_{0,2}$;
- (ii) between the investor and the reference entity, denoted by $\rho_{0,1}$; and
- (iii) between the counterparty and the reference entity, denoted by $\rho_{1,2}$.

It is not possible to vary all three of these parameters in a two dimensional matrix. The investor-counterparty correlation ($\rho_{0,2}$) remains constant within each table of values presented below since it is not the main risk driver underlying the BCVA value. A blank cell in these grids indicates that the combination of correlation parameters implied for that cell does not produce a positive definite correlation matrix. The results are reported from the point of view of the investor.

Note that when an investor is short protection (receiver CDS), the counterparty to the trade is long protection (payer CDS). The counterparty to the payer CDS in this scenario is the investor. Similarly, the counterparty to the trade of an investor who is long protection is short protection. This simple, but important concept, illustrated in Figure 11.1,

should be borne in mind when the results below are scrutinised. It implies that the DVA on a payer CDS is equivalent to the CCVA on a receiver CDS, with the investor assuming the role of counterparty.

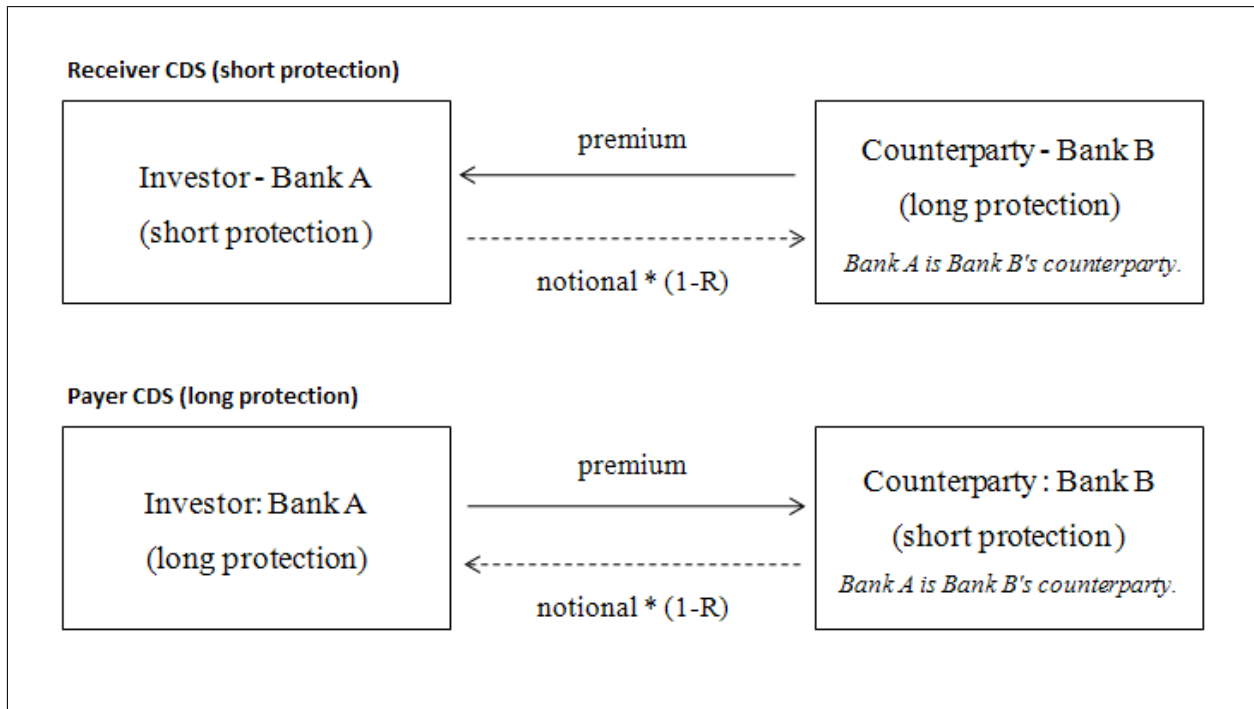


Figure 11.1: Illustration of the trade direction from both investor and counterparty perspectives

The effect of the following on the BCVA is investigated in this chapter:

- (i) Equal investor and counterparty levels of risk (Section 11.1).
- (ii) A change in the investor's riskiness (Section 11.2).
- (iii) A change in the counterparty's riskiness (Section 11.3).
- (iv) A change in the reference entity's riskiness (Section 11.4).
- (v) A change in the correlation between the default times of the investor and the counterparty (Section 11.5). Note that the investor-reference entity and reference entity-counterparty default time correlations are varied in all of the tables of results reported below.

In the scenarios below, each of the three entities (the investor, the reference entity and the counterparty) is either the low or the intermediate risk entity that featured in the

examples in Chapter 9. Their CDS term structures were displayed in Table 9.1. In the following four sections, the parameter $\rho_{0,2}$ is set to zero. Thereafter, in Section 11.5, an example is presented in which this correlation is adjusted. All hazard rate volatilities (ν_0 , ν_1 and ν_2) are set to 0.1.

11.1 Investor and Counterparty Equally Risky

We begin by considering an example in which the investor, the reference entity and the counterparty are all of intermediate risk. This an approximation of the situation in which the three entities have similar risk profiles. Table 11.1 below reports the BCVA for this scenario for both a short and a long protection position and a range of possible inputs to the parameters $\rho_{0,1}$ and $\rho_{1,2}$. Recall that the trade direction is expressed from the point of view of the investor.

A common misconception, regarding BCVA, is that the sign of this adjustment is purely a function of the credit quality of the investor relative to that of the counterparty. This section is designed to illustrate the inaccuracy of such a perception when the underlying trade is a CDS. Recall that the concept was introduced in Section 3.4.5. There, the BCVA on a vanilla interest rate swap was computed for varying levels of moneyness and equal investor and counterparty risk levels. In addition to addressing the misconception just described, the example in this section isolates the effect of the trade direction (long or short protection) on the bilateral adjustment that is applied to a CDS.

Note that, since the CVA is charged to the counterparty, a negative adjustment indicates that the investor is required to pay an additional spread to the counterparty or to subtract a spread from the amount that is charged to the counterparty. In this case, the size of the DVA exceeds the value of the CCVA. This, in turn, increases the value of the equivalent trade conducted between risk-free entities.

Consider now the adjustments in Table 11.1. We begin with a discussion of the receiver CDS BCVA, shown in the first half of the table. Since the investor is short protection in this transaction, the counterparty is long protection. Notice firstly that the receiver BCVA is the largest for small (very negative) values of $\rho_{1,2}$. This corresponds to the fact that the CCVA portion of the receiver BCVA is greatest in this region. As we move across the table from left to right (smallest to largest $\rho_{1,2}$ values) the size of the BCVA decreases

Receiver (short protection)							
$\rho_{0,1} / \rho_{1,2}$	-0.9	-0.6	-0.3	0	0.3	0.6	0.9
-0.9			7	0	0		
-0.6		12	7	0	0	0	
-0.3	13	12	7	0	0	0	0
0	11	10	5	-1	-1	-1	-1
0.3	-1	-2	-6	-12	-12	-12	-13
0.6		-17	-20	-26	-28	-28	
0.9			-41	-47	-51		
Payer (long protection)							
$\rho_{0,1} / \rho_{1,2}$	-0.9	-0.6	-0.3	0	0.3	0.6	0.9
-0.9			-13	-11	1		
-0.6		-12	-12	-10	2	16	
-0.3	-7	-7	-7	-5	6	20	40
0	0	0	0	1	13	28	49
0.3	0	0	0	1	12	27	52
0.6		0	0	1	12	27	
0.9			0	1	13		

Table 11.1: CDS BCVA in basis points - investor, reference entity and counterparty: intermediate risk

in value. This follows the same reasoning as the decrease in the unilateral receiver CVA with an increase in the correlation between the default times of the counterparty and the reference entity (since the parameter $\rho_{0,1}$ remains constant within a particular row).

Continuing to focus on the receiver BCVA portion of the table, consider next the large negative values when $\rho_{0,1}$ is high. Recall that the DVA is subtracted from the CCVA. The absolute size of this DVA is equivalent to the CCVA on a payer CDS with the investor assuming the role of counterparty. Since the unilateral payer CVA generally increases in value with the correlation between the default times of the counterparty and the reference entity, the DVA value increases with $\rho_{0,1}$. This results in a large negative BCVA. Scanning a particular column in the first half of Table 11.1 (in which the value of $\rho_{1,2}$ remains constant), we observe a reduction in the BCVA amount (moving down the column from top to bottom), due to the increase in the absolute size of the DVA.

The payer CDS BCVA (reported in the second half of Table 11.1) is similar to the receiver portion of the table, with the $\rho_{1,2}$ and $\rho_{0,1}$ axes switched and the signs of the values reversed. This is attributable to the counterparty and the investor possessing the same

risk profiles. This example, in which all three entities are of intermediate risk, clearly illustrates the effect of the trade direction on the sign of the BCVA. Were this adjustment merely a function of the relative riskiness of the investor and the counterparty, the values in Table 11.1 would all be zero. Note also that the recovery parameters are the same for all three entities. If these differed for the investor and the counterparty, the positive values would not necessarily mirror the negative ones¹. The large difference in value between the maximum positive and negative payer and receiver BCVA is a function of the unilateral payer CVA attaining much larger values than the receiver, unless the reference entity is exceptionally risky (see the discussions on this in Sections 9.2.1 and 9.2.4).

Regarding the payer BCVA, the large positive values when $\rho_{1,2}$ is high, are attributable to the CCVA. This is due to the fact that the unilateral CVA on a long protection position increases with the correlation between the counterparty and the reference entity default times. In the same vein, scanning a particular row of the payer CVA portion of Table 11.1 (from left to right) reveals an increase in the BCVA.

The large negative payer values when $\rho_{0,1}$ is small, are attributable to the DVA (which is subtracted from the CCVA) being larger than the CCVA in this area. Recall that the DVA is equivalent to the CCVA portion of the receiver BCVA, with the investor assuming the role of counterparty. Casting the eye down a particular column of the second part of Table 11.1, the effect of the correlation between the reference entity and the investor results in an increase in the payer BCVA.

11.2 A Less Risky Investor

Suppose now that the reference entity and the counterparty remain intermediate risk entities, as in the previous section. The investor is substituted with the low risk entity. The BCVA for this scenario is reported in Table 11.2 below.

A comparison of the results in Tables 11.1 and 11.2 reveals a general increase in the level of the BCVA in the latter. This is attributable to the decrease in the DVA due, in turn, to the reduction in the riskiness level of the investor. The decrease in DVA is most prominent when $\rho_{0,1}$ is large in the receiver CDS case and when $\rho_{0,1}$ is small in the event

¹See (10.1) for a reminder of the role played by the recovery (or, equivalently, the loss given default) parameter in the computation of BCVA.

Receiver (short protection)							
$\rho_{0,1} / \rho_{1,2}$	-0.9	-0.6	-0.3	0	0.3	0.6	0.9
-0.9			7	0	0		
-0.6		13	7	0	0	0	
-0.3	14	13	7	0	0	0	0
0	12	11	6	-1	-1	-1	-1
0.3	4	4	-1	-8	-8	-8	-8
0.6		-5	-9	-17	-17	-17	
0.9			-14	-21	-23		
Payer (long protection)							
$\rho_{0,1} / \rho_{1,2}$	-0.9	-0.6	-0.3	0	0.3	0.6	0.9
-0.9			-7	-5	6		
-0.6		-7	-7	-5	6	23	
-0.3	-4	-4	-4	-2	9	24	46
0	0	0	0	2	13	28	50
0.3	0	0	0	2	13	28	53
0.6		0	0	1	13	29	
0.9			0	1	13		

Table 11.2: CDS BCVA in basis points - reference entity and counterparty: intermediate risk; investor: low risk

that the investor is long protection. These are the regions in which the riskiness level of the investor has the greatest effect on the DVA since it is largest in these areas. For more neutral investor-reference entity default time correlations, the increase in BCVA is less prominent.

An additional point to consider is that, since the investor is less risky than in the previous section, its survival provability is now greater. The effect of this on the CCVA reinforces the BCVA increase in this scenario, which is driven predominantly by the lower DVA.

11.3 Decreasing the Riskiness of the Counterparty

In this section, we again alter the scenario in Section 11.1. However, rather than changing the riskiness of the investor (as in Section 11.2), it is the counterparty that is modified from being the intermediate risk entity to the low risk name. The investor and the reference entity remain of intermediate risk. The results are reported in Table 11.3 below.

A comparison of Tables 11.1 and 11.3 reveals a general reduction in the size of the BCVA

Receiver (short protection)							
$\rho_{0,1} / \rho_{1,2}$	-0.9	-0.6	-0.3	0	0.3	0.6	0.9
-0.9			4	0	0		
-0.6		7	4	0	0	0	
-0.3	8	7	4	0	0	0	0
0	6	5	2	-2	-2	-1	-2
0.3	-6	-6	-9	-13	-13	-13	-13
0.6		-22	-24	-28	-29	-29	
0.9			-46	-50	-52		
Payer (long protection)							
$\rho_{0,1} / \rho_{1,2}$	-0.9	-0.6	-0.3	0	0.3	0.6	0.9
-0.9			-14	-12	-5		
-0.6		-13	-13	-11	-3	6	
-0.3	-7	-7	-7	-6	1	9	13
0	0	0	0	1	8	17	21
0.3	0	0	0	1	8	18	24
0.6		0	0	1	8	18	
0.9			0	1	8		

Table 11.3: CDS BCVA in basis points - investor and reference entity: intermediate risk; counterparty: low risk

in the latter due to the decrease in the default probability of the counterparty. The direction of the differences between the values in these tables is the same as it would be if unilateral CVA were being computed, since the sign in front of the CCVA is positive and the investor remains constant in the two scenarios.

The reduction in BCVA for the short protection position (which is reported in the first half of the table) is most pronounced for extreme negative correlations between the default times of the reference entity and of the counterparty. The converse applies to the payer trade. This follows from the correlation impact on the CVA that was discussed in Chapter 9. The reduction in the default risk level of the counterparty demonstrates the greatest effect in the regions in which the correlation significantly impacts the credit valuation adjustment.

In addition to the counterparty's default probability affecting the CCVA portion of the BCVA, the increased survival probability reinforces the general reduction in BCVA from Table 11.1 to Table 11.3. This is due to the intensifying effect of the counterparty's survival probability on the size of the DVA.

11.4 Altering the Reference Entity

The effects (on the BCVA) of reducing the level of riskiness of both the investor and the counterparty in isolation have been considered. The consequence of a change in the reference entity's level of default risk will now be examined. As in Section 11.1, the investor and the counterparty are both intermediate risk entities. The reference entity is modified from intermediate to low risk. The results are reported in Table 11.4 below.

Receiver (short protection)							
$\rho_{0,1} / \rho_{1,2}$	-0.9	-0.6	-0.3	0	0.3	0.6	0.9
-0.9			4	0	0		
-0.6		7	4	0	0	0	
-0.3	8	7	4	0	0	0	0
0	6	5	3	-1	-1	-1	0
0.3	-3	-4	-6	-10	-10	-9	-9
0.6		-18	-19	-23	-23	-23	
0.9			-46	-50	-51		
Payer (long protection)							
$\rho_{0,1} / \rho_{1,2}$	-0.9	-0.6	-0.3	0	0.3	0.6	0.9
-0.9			-8	-6	3		
-0.6		-7	-7	-5	4	18	
-0.3	-4	-4	-4	-3	6	20	46
0	0	0	0	1	10	23	51
0.3	0	0	0	1	10	24	51
0.6		0	0	1	9	23	
0.9			0	0	9		

Table 11.4: CDS BCVA in basis points - reference entity: low risk; investor and counterparty: intermediate risk

The effect of adjusting the reference entity's level of riskiness is not as straightforward as for the other two entities. It is predominantly a single term in the BCVA that is affected by a change in the default risk level of either the counterparty or the investor. However, both the CCVA and the DVA are reduced when the reference entity's riskiness is decreased. Since the DVA is subtracted from the CCVA, it is not clear at the outset whether the alteration of the reference entity's attributes will result in an increase or a reduction in the BCVA.

Comparing Tables 11.1 and 11.4, consider firstly the receiver CDS or short protection

position, as seen from the perspective of the investor. In general, when the CCVA portion of the BCVA is dominant, we expect a decrease in BCVA. This is evidenced by a reduction for very negative values of $\rho_{1,2}$ and for all values in the first quadrant in the table of results. In this first quadrant, the DVA is generally zero, or close to zero, and it is the CCVA that determines the size of the BCVA. Similarly, in the fourth quadrant, the receiver CVA is larger in Table 11.4 than in Table 11.1. Here the first term in the BCVA is generally zero, or approaching zero, and the reduction in the DVA is the influential factor.

For a $\rho_{0,1}$ value of 0.99, combined with a $\rho_{1,2}$ value of -0.6, there is a large reduction in the BCVA when the reference entity's level of riskiness is reduced. Since the CCVA is likely to be non zero at this point, a reduction in the BCVA is not impossible. However, it would be expected that the reduction in DVA would offset it to some extent, given that $\rho_{0,1}$ is 0.99 and the DVA is thus at its largest. Contrary to this expectation, the decrease in the BCVA is -6bp. The reason for this is that the DVA increases by a small amount for this combination of parameters, rather than decreasing as expected. When the reference entity is modified from intermediate to low risk, the number of scenarios in which the reference entity is the first to default in the 99% correlation case is reduced, resulting in a slight increase in DVA. In other words, the DVA increases because the investor is more likely to default before the reference entity. Thus, instead of DVA decreasing for this extreme combinations of correlations, it increases.

Turning now to the change in the payer BCVA values, we compare the results that are reported in the second half of Table 11.4 to the corresponding numbers in Table 11.1. All of the first quadrant values are higher than when the reference entity was the intermediate risk name. This is expected since the DVA portion of the BCVA is dominant in the first quadrant, with the CCVA portion expected to be zero or approaching zero. Similarly, the converse holds in the fourth quadrant in which the BCVA decreases.

When $\rho_{0,1}$ attains a value of 0.99 and $\rho_{1,2}$ is set to -0.6, there is a 6bp increase in the payer BCVA. Notice that the investor and the counterparty are equally risky. As such, the payer and receiver values mirror each other, with the $\rho_{1,2}$ and $\rho_{0,1}$ axes reversed and the sign of the reported BCVA numbers switched. The reason for such a large increase in BCVA (resulting from a decrease in DVA not being offset by a reduction in CCVA) is therefore attributable to the same effect as the 6bp reduction in receiver BCVA for a $\rho_{0,1}$ value of 0.99 combined with a $\rho_{1,2}$ value of -0.6. That is, the reduction in the reference

entity riskiness results in an increase in the payer CCVA since there are fewer scenarios in which the reference entity is the first to default.

11.5 The Effect of a Non Zero Investor-Counterparty Default Time Correlation

The parameter $\rho_{0,2}$ (which represents the correlation between the default times of the investor and the counterparty) has been set to zero in generating the results presented thus far. The assumption is relaxed in this section. All three entities are of intermediate risk as in Section 11.1 and their hazard rate volatilities (ν_1 , ν_2 and ν_3) remain 0.1. The BCVA when $\rho_{0,2}$ is set to -0.6 is reported in Table 11.5. The results for a $\rho_{0,2}$ value of 0.6 are presented in Table 11.6.

Receiver (short protection)							
$\rho_{0,1} / \rho_{1,2}$	-0.9	-0.6	-0.3	0	0.3	0.6	0.9
-0.9					-1	15	
-0.6				-11	1	17	40
-0.3			-7	-6	6	21	44
0		0	0	1	13	28	
0.3	0	0	0	2	13		
0.6	0	0	0	2			
0.9		0	0				
Payer (long protection)							
$\rho_{0,1} / \rho_{1,2}$	-0.9	-0.6	-0.3	0	0.3	0.6	0.9
-0.9					0	0	
-0.6				0	0	0	0
-0.3			8	0	0	0	0
0		11	6	-2	-2	-2	
0.3	0	0	-6	-13	-13		
0.6	-15	-17	-22	-29			
0.9		-40	-45				

Table 11.5: CDS BCVA in basis points - investor, reference entity and counterparty: intermediate risk; $\rho_{0,2} = -0.6$

An initial comparison of the results in Tables 11.5 and 11.6 with those in Table 11.1 (the zero investor-reference entity correlation case) reveals that the parameter $\rho_{0,2}$ has less impact on the CDS BCVA than either of the remaining two correlations ($\rho_{0,1}$ and $\rho_{1,2}$). This is to be expected since the spreads of the reference entity determine the CDS values.

Receiver (short protection)							
$\rho_{0,1} / \rho_{1,2}$	-0.9	-0.6	-0.3	0	0.3	0.6	0.9
-0.9		-11	-11				
-0.6	-10	-10	-10	-7			
-0.3	-5	-5	-6	-4	6		
0		0	0	1	11	23	
0.3			0	0	10	24	42
0.6				0	7	22	42
0.9					5	23	
Payer (long protection)							
$\rho_{0,1} / \rho_{1,2}$	-0.9	-0.6	-0.3	0	0.3	0.6	0.9
-0.9		10	5				
-0.6	11	10	5	0			
-0.3	11	10	6	0	0		
0		7	4	-1	0	0	
0.3			-6	-11	-10	-7	-5
0.6				-23	-24	-22	-23
0.9					-41	-43	

Table 11.6: CDS BCVA in basis points - investor, reference entity and counterparty: intermediate risk; $\rho_{0,2} = 0.6$

Therefore, the correlation of the default times of the reference entity with the defaulting counterparties is of greatest significance. It is in modelling the correlation between the investor and counterparty default times with those of the reference entity that the wrong/right way nature of the exposure is captured.

The differences between the BCVA in the $\rho_{0,2} = -0.6$ (Table 11.5) and $\rho_{0,2} = 0$ (Table 11.1) scenarios are practically insignificant. More relevant is the modification from a zero correlation to a (fairly large) positive one (Table 11.6). Recall the discussion regarding default time correlations in Chapter 7. As the correlation between the default times of two entities increases in value, the number of paths along which both entities default is expected to increase. The total percentage of defaults per entity remains constant however. There is therefore a decline in the number of relevant investor and counterparty default times when the correlation between them is increased to 0.6.

This intuition corresponds to the results presented. For the payer CDS, there tends to be an increase in the BCVA in the first quadrant, corresponding to the decrease in DVA (since the DVA is subtracted from the CCVA, a lower DVA results in a larger BCVA).

The bilateral adjustment is driven by the DVA in this first quadrant. Similarly, there is a general reduction in BCVA in the fourth quadrant since the CCVA is the more prominent of the two terms of which the BCVA is comprised in this quadrant. The most significant changes in the BCVA occur when the values of $\rho_{0,1}$ and $\rho_{1,2}$ are large and positive. It is in these regions that the BCVA in the independent case was the largest. The absolute change on a large value is expected to outsize the absolute change on a small value. Additionally, in these regions, the default times between all three entities are highly correlated. This reduces the number of relevant default times of each name.

The changes in the receiver CDS BCVA when $\rho_{0,2}$ is increased from 0 to 0.6 are almost identical to the payer CDS case, with the axes reversed and the signs switched. The reasoning for this symmetry (when the investor and counterparty are equally risky) was explained in Section 11.1.

11.6 Chapter Summary

The BCVA values were reported in this chapter for various combinations of riskiness of the three relevant entities, namely the investor, the counterparty and the reference entity. The chapter illustrated the following:

- (i) The sign of the BCVA depends on both the direction of the CDS trade and on the relative riskiness levels of the investor and the counterparty.
- (ii) A decrease (increase) in the investor's par spreads (and thus hazard rate term structure) in isolation results in a general increase (reduction) in BCVA levels. This is attributable to the reduction (increase) in the DVA values that are subtracted from the CCVA.
- (iii) A decrease (increase) in the counterparty's par spreads (and thus hazard rate term structure) in isolation results in a general reduction (increase) in BCVA levels. This is attributable to the lower CCVA portion of the BCVA.
- (iv) The result of altering the riskiness level of the reference entity can be either an increase or a decrease in the BCVA, depending on whether the CCVA or the DVA decreases the most. This, in turn, depends on the riskiness levels of these entities and the size of the correlation parameters $\rho_{0,1}$ and $\rho_{1,2}$.

- (v) The correlation between the default times of the investor and the counterparty ($\rho_{0,2}$) has a smaller effect on the BCVA than the other two correlation parameters. They are more significant in capturing the wrong/right way risk of the CDS trade. A change in $\rho_{0,2}$ can result in a small increase or decrease in BCVA depending on the size of the correlation parameters $\rho_{0,1}$ and $\rho_{1,2}$. The change is most significant when all three correlation parameters assume extreme positive values.

Part IV

Case Study and Conclusion

Chapter 12

Case Study

In May 2005, Barclays Bank Plc acquired a 60% stake in the ABSA Group, one of South Africa's largest financial services organisations [92]. Suppose that Barclays wishes to hedge a portion of its country risk to ABSA by purchasing protection referencing the South African government offshore debt (SOAF). The BCVA on a five year par CDS is computed by assuming that the contract originates at four possible points in time. The results are reported from the perspective of Barclays (BACH), the protection buyer, with the following institutions each assuming the role of counterparty: Lehman Brothers (LEH), American International Group (AIG) and J.P. Morgan (JPM).

This is the penultimate chapter of the dissertation. The first section contains details of the market data that is obtained for the CVA calculation. It is followed by an explanation of the model calibration. Thereafter, in Section 12.3, the CVA in each scenario considered is reported and discussed. Finally, there is a brief summary of the chapter in Section 12.4.

12.1 Market Data

The 2008 global financial crisis was a catalyst in the recognition of default risk. We compute the CVA on the 5 year par CDS described above on the eve of the Lehman Brothers default. In addition, the results are compared to those that would have been obtained for the same counterparty and investor combinations in 2006, two years prior to the crisis and in 2010, two years later¹. In addition, the CVA is computed more recently, near the beginning of August 2011, in the midst of the sovereign debt crisis.

¹It goes without saying that only two potential counterparties are considered in 2010.

The relevant par CDS curves, reported in Table 12.1 below, were obtained from Bloomberg and are the recorded market closing prices on the following days: 15 September 2006, 12 September 2008, 31 August 2010 and 5 August 2011.

Term	AIG				BACR				JPM			
	2006	2008	2010	2011	2006	2008	2010	2011	2006	2008	2010	2011
6m	-	940	155	80	-	-	88	81	-	-	-	-
1y	2	1,170	198	114	2	67	89	80	7	81	64	42
2y	-	-	255	164	3	91	101	113	-	-	76	61
3y	6	918	292	199	4	109	112	141	11	110	86	81
4y	-	857	324	237	-	-	-	-	-	127	99	97
5y	11	858	353	268	7	133	133	173	15	128	113	108
7y	14	713	351	284	-	-	-	-	20	133	111	123
10y	16	699	350	299	13	136	140	193	26	130	112	130

Term	LEH				SOAF			
	2006	2008	2010	2011	2006	2008	2010	2011
6m	-	973	-	-	-	-	-	-
1y	12	1,128	-	-	17	61	81	54
2y	-	-	-	-	33	115	109	80
3y	20	817	-	-	43	152	130	107
4y	-	702	-	-	-	-	144	130
5y	23	642	-	-	60	189	155	146
7y	27	582	-	-	72	214	163	167
10y	37	545	-	-	85	231	170	188

Table 12.1: Case study CDS spreads. Source: Bloomberg

The implied survival probabilities were stripped from the spreads in Table 12.1 using flat left interpolation of the hazard rates. Refer to Section B.1 in Appendix B for a description of this procedure. The choice of flat left interpolation is appropriate given that the AIG and LEH 2008 curves are both inverted. In general, flat interpolation is preferred in such cases, with linear interpolation likely to result in negative hazard rates² [25].

The SOAF recovery rate reported on Bloomberg is 25%. The assumed recovery for all other entities involved is 40%. Note that, although an actual recovery rate for LEH has been determined, we assume that the CVA is priced pre-default and thus include the entity under the 40% recovery assumption. The survival probabilities obtained on each of

²These do, in fact, occur when linear interpolations is applied in stripping the 2008 LEH curve.

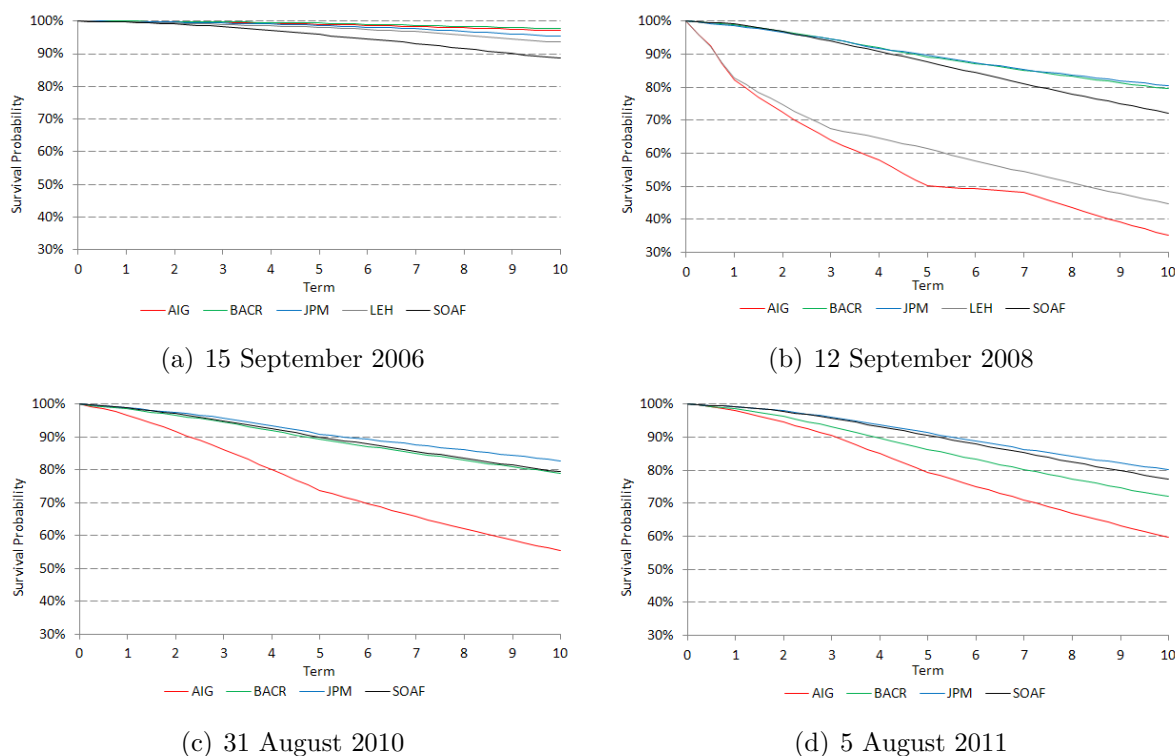


Figure 12.1: Survival probabilities of entities in case study. *Par CDS curves were stripped assuming flat interpolation of hazard rates.*

the pertinent days are illustrated in Figure 12.1. Note that, in each graph in the figure, the minimum value on the y-axis is kept constant to facilitate comparison between the different dates.

With reference to Table 12.1 and Figure 12.1, notice firstly that the default risk in 2006 was significantly understated in comparison to later years. In 2008, spreads were elevated in general due to the global financial crisis. Two years later, they had still not returned to 2006 levels. This illustrates a new awareness of the risks of default and an increased risk aversion among market participants. There is effectively a new benchmark in the CDS market. The entity in Table 12.1 whose default was considered most likely in 2006 is the Republic of South Africa. This reflects the risk aversion that has long been associated with emerging markets. It is ironic that, at the time of writing, the sovereign debt crisis is occurring in the developed world.

The concept of a CDS curve inversion has been discussed elsewhere in the dissertation. This is normally understood to mean that the one-year contract is trading above the

five-year contract [93]. The LEH and AIG 2008 curves are textbook illustrations of the phenomenon. This is, of course, appropriate given that the former defaulted the following day while the latter was almost certain to default without government intervention. In Figure 12.1(b), there is a clear distinction between the survival probabilities of LEH and AIG and the remaining three names in the case study. This demonstrates the idiosyncratic risk priced into the market despite the systemic risk illustrated by the general reduction in default probabilities between 2006 and 2008.

Lastly, observe that, although the shape of the AIG curve has normalised and survivals have increased significantly since 2008, these are still lower in 2010 and 2011 than for the other three entities whose survival curves are illustrated.

12.2 Model Calibration

A CIR++ process was fitted to each of the hazard rate term structures obtained from each of the par CDS curves in Table 12.1. Since flat left interpolation was utilised in stripping the curves, it is again employed in fitting the CIR++ processes. This procedure was discussed in Section 7.1 and will not be repeated here. The two parameters that deserve attention are the volatilities, ν , and the correlations, ρ .

12.2.1 The Volatility Parameter

Due to the absence of a CDS options market, there is a free parameter in the CIR++ process calibration. This was discussed in Section 7.1. In this chapter, we utilise the real world measure in obtaining the parameter, ν , by fitting the process to historical hazard rates.

A simple Euler discretisation of the CIR stochastic differential equation (SDE) was criticised in Section 7.2.1. This was due to the possibility of negative hazard rates being obtained if such a discretisation scheme were to be employed in the Monte Carlo simulation of the hazard rate. For the sake of simplicity and, since the discretisation methodology employed in this section is purely for the purpose of calibration rather than simulation,

we rewrite the CIR SDE in (7.1), using a Euler discretisation [94], as:

$$\begin{aligned}
 y_{t+\delta t} - y_t &= \kappa(\mu - y_t)\delta t + \nu\sqrt{y_t}\epsilon_t \\
 \Rightarrow \frac{y_{t+\delta t} - y_t}{\sqrt{y_t}} &= \frac{\kappa(\mu - y_t)}{\sqrt{y_t}} + \nu\epsilon_t,
 \end{aligned} \tag{12.1}$$

where $\epsilon_t \sim N(0, \delta t) \Rightarrow \nu\epsilon_t \sim N(0, \nu^2\delta t)$.

For each name relevant to the case study, the steps below outline the determination of the historical hazard rates for the entity and their application to the computation of the volatility parameter, ν .

- (i) The daily five year par CDS spreads of the entity are obtained for the two years ending on the date at which the CVA is computed. The five year point is selected as it is generally the most liquid contract that is traded in the CDS market.
- (ii) These credit spreads are then required to be converted into hazard rates. For this purpose, we make the simplifying assumption of constant hazard rates and interest rates. This implies that the hazard rate on each day may be computed as

$$\lambda = \frac{1}{t_c} \ln \left(\frac{S * t_c}{\text{LGD}} + 1 \right), \tag{12.2}$$

where t_c is the time between CDS premium payments and S is the spread. The derivation of (12.2) is straightforward and can be found in Section B.3 in Appendix B.

- (iii) The historical hazard rates computed in the previous step are then used to calculate the observed values of $\frac{y_{t+\delta t} - y_t}{\sqrt{y_t}}$ in (12.1). On each day, the hazard rate from the previous day is subtracted from the current rate and the difference divided by the square root of the value from the previous day.
- (iv) Next, the term $\frac{\kappa(\mu - y_t)}{\sqrt{y_t}}$ is computed on each day using the values of μ and κ that are determined as per the guidelines in Section 7.1.
- (v) For each day, the difference between the quantities in Steps (iii) and (iv) above is determined.
- (vi) The value of the parameter, ν , is set equal to the annualised standard deviation of these differences. Note that, since the deterministic part of the CIR++ model is

constant for each five year point, its inclusion in the above calibration procedure would not alter the value of ν . This is the justification behind the focus on the CIR portion of the CIR++ process.

The market-implied hazard rates, fitted CIR rates and deterministic shift, the three quantities relevant to the fitting of the CIR++ model, were illustrated in Figure 7.1 of Section 7.1.1 for linear hazard rate interpolation. In this section, we show the results for the 2008 LEH curve as an illustration of the use of flat hazard rate interpolation.




Figure 12.2: Market-implied hazard rates, fitted CIR rates and deterministic shift for the Lehman Brothers' curve on COB 12 September 2008. *Flat left interpolation is applied to hazard rates. Quantities illustrated are those required in equation (7.7) for the CIR++ process calibration.*

12.2.2 Correlation

A Gaussian copula, introduced in Section 10.2, is used to correlate the uniform random variables U_0 , U_1 and U_2 , which are associated with the default times of the investor, the reference entity and the counterparty respectively. In the absence of a market from which to imply the correlation parameters, we again resort to historically observed rates.

Since there is no actual history of default time correlations for the entities that feature in the case study, spread correlation is employed as a proxy data source. Spread correlation is an attractive source of correlation data [90]. The movement in spreads is an indication of the expected default times of the entities. In addition, it has a constant maturity and is relatively clean when available. The five year point refers to a liquid instrument [90] for the entities relevant to the case study.

Since the correlation calculated between spreads is dependent on actual levels, we utilise the correlation between the log returns of the five year spreads of the relevant entities. The period of the historical data is the two years leading up to the day on which the CVA is computed.

The correlations obtained for the relevant pairs of entities are illustrated in Figure 12.3. Notice that the systemic risk priced into the credit market has increased significantly since




Figure 12.3: Correlations for case study

2006. In fact, for all of the entity pairs illustrated in Figure 12.3, with the exception of BACR AIG, correlation has increased since 2008. This points to an awareness of systemic risk that was brought about by the crisis of 2008. The decrease in the correlation between Barclays and AIG from 2008 to 2010 and 2011 is expected, given that a default was priced

into the AIG spreads in 2008.

12.3 Results

The BCVA is computed on 15/09/2006, 12/09/2008, 31/08/2010 and 5/08/2011 respectively for a five year par CDS originated on these dates. The results are reported in Table 12.2 from the perspective of Barclays (the investor). The underlying reference credit is SOAF.

Counterparty	2006	2008	2010	2011
American International Group	0.06	25.01	10.42	10.53
J.P. Morgan	0.25	6.90	5.93	6.47
Lehman Brothers	0.31	24.75	-	-

Table 12.2: BCVA results in basis points

The first column of Table 12.2 illustrates the lack of awareness of default risk in the credit markets in 2006. Aside from the assumption that certain counterparties were deemed to be risk-free and the CVA on these entities ignored, the results would have been too small to be significant had they been calculated, due to the low default probabilities in the market. Note that, since the correlations utilised in the determination of the CVA in Table 12.2 are all positive, the DVA portion of the bilateral adjustment is essentially insignificant.

In the 2008 numbers, there is a marked difference between the CVA that is reported with J.P. Morgan as the counterparty and for the other two scenarios. This is to be expected in light of the par CDS spreads that were reported in Section 12.1 for the three counterparties. The fact that LEH defaulted the following day points to the importance of collateral management. Without day 1 collateral being posted by LEH, the inclusion of CVA into the CDS spread would have been irrelevant since the first premium (adjusted for CVA) was only due at the end of the quarter. In addition, if the CVA is not a day 1 charge, the concept of trading in points upfront³ can be extended from the reference entity to the counterparty (or the investor in the event that the DVA with the risky investor is the more relevant term in the BCVA). As a point of interest, the SOAF five year spread

³When a reference credit has particularly high spreads, the running premium or a portion of the running premium is present valued and required, by the seller of protection, as an upfront payment [95].

widened by 13.5bp between COB 12/09/2008 and COB 15/09/2008 (the former was a Friday).

There is not a significant difference between the 2010 and 2011 CVA values. Importantly, the CVA when AIG is the counterparty is higher than when J.P. Morgan is the counterparty, due to the greater risk of default that is priced into the credit spreads of the former.

12.4 Chapter Summary

The purpose of this chapter was to apply the CDS BCVA model that has been implemented, to actual market scenarios. Although the calibration methodology utilised in the case study is not a unique means of determining the appropriate free parameters, it represents a practical illustration of one possible solution to the calibration challenge.

The results obtained demonstrate that the model adapts to the prevailing market conditions. The inclusion of CVA into pricing does not, however, replace traditional credit risk mitigants such as collateralization unless the CVA is incorporated into the CDS trade as an upfront fee. Even if this is the case, collateralisation may involve simulating the potential trade values and utilising a higher confidence interval than the expected value for calculating the collateral required to be posted. This is a more cautious approach that may be adopted for high risk counterparties. The effects of collateralisation should then be incorporated into the CVA computation.

Chapter 13

Conclusion

The attention that is being received by CVA at the time of writing (in both offshore and in South African financial markets), as well as the time and resources being devoted to it, are empirical evidence of its relevance. Financial institutions that have not (at the very least) embarked on a strategy for its implementation are in danger of losing ground to their competitors. Aside from the economic implications of inaccurately pricing counterparty risk, Basel III regulations impose severe capital penalties on institutions that do not manage counterparty risk dynamically and have accurate models for its valuation. This was discussed in Chapter 3.1.

Counterparty risk in the over the counter derivatives market is a reality. The 2008 global financial crisis was a harsh reminder that no corporate entities are risk-free. More recently, the debacle surrounding sovereign credit risk in 2010 and 2011 has erased the final fragments of comfort that were left intact after the Lehman Brothers' default. Not even CVA on sovereigns can be overlooked.

This final chapter is divided into two sections. The first is a summary of the dissertation. The conclusions drawn therein are highlighted. The second section is a discussion pertaining to potential extensions of the CDS CVA model that has been implemented. Further research areas are also proposed.

13.1 Summary

The central application of CVA modelling in this dissertation is to credit default swaps. The fundamental building blocks of default risk pricing were presented in such a way as

to ensure that the document is fairly self-contained, even to readers unfamiliar with credit derivatives markets.

Chapters 1 through 12 were categorised into four main parts. The first of these served as a broad introduction to the required concepts. It was composed of five chapters. The first four targeted Objectives (i) and (ii) of the dissertation. These were set forth in Section 1.2. The fifth and final chapter of Part I addressed Objective (iii).

Chapter 1 presented an outline of the structure of the document, as well as stating the chief objectives of the dissertation. Both credit and counterparty risk were introduced and discussed in Chapter 2. Core concepts, that form the background required for an introduction to CVA, were reviewed. These included counterparty credit exposure, potential future exposure and a primer on credit risk mitigants. Finally, the credit default swap instrument was introduced.

CVA was the subject of Chapter 3. The term *credit valuation adjustment* was defined and its relevance and advantages motivated. Additionally, the importance of modelling it accurately was highlighted. Right and wrong way risk, buzz words in modern finance and concepts integral to the accurate management and measurement of CVA, were described. A thorough discussion of unilateral CVA was followed by a similar discourse on bilateral adjustments, including the difficulties associated with own risk management. In order to demonstrate the application of the General Counterparty Risk Pricing Formulae and to provide readers with a tangible example, the CVA on a simple vanilla interest rate swap was computed. Since the chapter was intended to provide a complete overview of CVA, a section was then devoted to the practical aspects of its implementation in a financial institution (such as an investment bank). The advantages of managing CVA centrally were an integral part of this discussion. Two instruments associated with CVA concluded the chapter. These were presented, not merely as a matter of interest, but due to the fact that CVA pricing is equivalent to the valuation of one of these derivatives (a contingent credit default swap).

The mathematics associated with default risk modelling was introduced in Chapter 4. Both structural and reduced form techniques were presented. The area of focus was the latter, due both to its presence in the CDS CVA model that followed and market conventions. The pricing of credit default swaps (CDSs) was then explained, together with a

reference to Section B.1 in Appendix B in which the mechanism for bootstrapping hazard rates from market-observable spreads was described.

The final chapter in Part I presented a review of the literature devoted to CDS CVA valuation. A variety of methods, including their advantages and disadvantages, were discussed. The fact that the field is relatively new was evident from the wide array of approaches that are utilised, and by the absence of an agreed upon market model such as the one that exists for CDS pricing.

Part II encompassed Chapters 6 to 9. It was devoted to the implementation of a model for pricing the unilateral CVA on a CDS. The first four chapters in this section accomplished the fourth dissertation objective (as stated in Section 1.2). Chapter 9 targeted Objective (v).

In Chapter 6, the justification of the model that was selected from the literature was presented, together with its underlying assumptions. A diagram illustrating the four principal steps employed in pricing the unilateral CDS CVA completed the chapter. The implementation of Steps 1 and 2 in the diagram were the subject of Chapter 7, with Chapter 8 devoted to the latter two steps.

The calibration of the CIR++ processes to the intensities of the relevant entities was explained at the beginning of Chapter 7. Various simulation paths of a specific CIR++ process were illustrated for a set of volatility parameters in order to provide a visualisation of the hazard rates generated by this type of process. The simulation of the default times of the counterparty that are relevant to the CVA computation was the topic of the second part of the chapter. The results of tests conducted to ensure that the correct market-implied survival probabilities were obtained from the simulation of these calibrated processes were reported.

The majority of Chapter 8 was devoted to the determination of the survival probabilities of the reference entity at the counterparty default times. This encompassed both a conditional copula function and a fractional fast Fourier transform (FRFT). The motivation behind the use of the latter was included in an introduction to Fourier theory in Appendix F. The implementation of the FRFT was described in detail, including an investigation into the integrand, the selection of the parameters required in the numerical

integration scheme and a trap in the characteristic function of the integrated CIR process as it appears in the literature. Additionally, manipulation of the conditional copula function into a form in which the cumulative normal distribution function could be applied, was explained. The determination of the CVA, which is straightforward once the survival probabilities have been computed, concluded the chapter.

Chapter 9 completed Part II. The unilateral CDS CVA obtained from the model, whose implementation was the subject of Chapters 7 and 8, was reported in this chapter for combinations of entities of varying risk levels. The key parameters, correlation and reference entity volatility, were varied and the model stress tested. The effect of the CDS maturity on the CVA was also investigated. The results presented were analysed thoroughly. From a convergence, stability and robustness perspective, the model's performance was pleasing. The results were satisfactory, with trends and comparisons of the CVA obtained for various parameter choices both intuitive and justifiable.

Particularly noteworthy was the presence of right and wrong way risk in the CVA results reported. In general, the payer CDS (long protection) CVA increased in conjunction with a rise in the positive correlation between the uniform random variables of the exponential triggers of the default times of the reference entity and the counterparty. The receiver CDS (short protection) CVA increased with a reduction in correlation. The relative levels of the payer and receiver CVA were discussed. The effect of the reference entity's volatility featured in these explanations, since the CVA has an option-like structure. In particular, the zero correlation CVA increased with an increase in this parameter. The changes in the model output resulting from modifications to the risk levels of the two entities were also discussed. A summary of the conclusions drawn from the CVA values reported in Chapter 9 was presented in Section 9.3.

The unilateral CDS CVA model was extended into the bilateral realm in Part III. This accomplished Objective (vi). The assumptions pertaining to the extension, as well as the additional mathematics required, were set forth in Chapter 10. Numerical illustrations of the BCVA were presented in Chapter 11, accompanied by similar analyses and conclusions to those in Chapter 9. The examination of the CDS BCVA was not as straightforward as the unilateral results. This is due to the former being comprised of two terms (the CCVA and the DVA). The effects of each of these on the BCVA may differ since the DVA is subtracted from the CCVA. The changes in the bilateral model output resulting from

alterations to the risk levels of the three entities (the investor, counterparty and reference entity) were examined for an array of inputs to the parameters $\rho_{0,1}$ (the investor-reference entity default time correlation) and $\rho_{1,2}$ (the reference entity-counterparty correlation). The directions in the changes of the BCVA values when the riskiness levels of the entities were modified in isolation were intuitively justifiable. Note also that the effect (on the BCVA values) of changing the size of the correlation parameter linking the default times of the investor and the counterparty ($\rho_{0,2}$) was less significant than the effect of altering the remaining two correlations ($\rho_{0,1}$ and $\rho_{1,2}$). This is to be expected since the latter two parameters capture the wrong and right way risks inherent in the trade as they connect the underlying reference entity's default probabilities to those of the defaulting entities. These were analysed in the text and summarised in Section 11.6.

Finally, Chapter 12 (the first in Part IV) presented a case study that demonstrated the application of the CDS BCVA model to real world scenarios. The effects of the changes in default and systemic risk perceptions from 2006 to the time of writing were demonstrated. The results obtained illustrated the point that CVA cannot replace traditional credit risk mitigants such as collateralisation unless it is charged upfront. Even then, collateral still has a place and should be incorporated into the CVA calculation. Chapter 12 achieved the seventh and final objective of the dissertation (as stated in Section 1.2).

13.2 Directions for Further Research

The literature devoted to CVA and CVA-related topics is vast and is being updated at a rapid pace. The relevance of the subject and the fact that methodologies and ideas are not yet properly established render it an exciting and dynamic area of research. The model for CDS CVA, of which a large portion of this dissertation was comprised, relies on a number of assumptions. Its implementation methodology is not unique. Three simplifications that deserve attention are:

- (i) The copula connecting the uniform random variables related to the default times of the relevant entities is a Gaussian copula. Experimentation with alternative copulas could be considered. The advantage of the Gaussian copula is that the correlation between each pair of entities is a single parameter that requires calibration. Even with this advantage, the copula parameterisation is not straightforward. A less familiar copula with added parameters would create additional complexities. Additionally, the normal distribution is widely used in financial modelling and is well understood

by many practitioners. As such, an intuitive understanding of the modelling techniques and the effects of the correlation parameter is more easily established than for an arbitrary copula. An alternative, that requires the estimation of a single additional parameter (the degrees of freedom) is the t-copula (see for example [96]). The transition from the Gaussian to the t-copula in default risk correlation modelling is similar to the use of a t-distribution in place of a Gaussian distribution in estimating value at risk in equity markets (see for example [97]). The t-distribution has the same characteristics as the normal distribution with the addition of leptokurtosis (fat tails), rendering extreme events more likely.

- (ii) Although double (multiple) defaults are impossible, the potential default of the reference entity during the time taken for the deal to be closed out and replaced upon counterparty default, should be incorporated into the CDS CVA model. In the bilateral version, the consideration of both the investor and the reference entity defaulting during this close-out period should be investigated.
- (iii) The CDS CVA model implemented in this dissertation does not take collateral into account. An extension, that takes a CSA into account, should be considered. This is discussed further below.

The points raised in items (i) to (iii) above are simplifications particular to the CDS CVA model that has been implemented in this dissertation. We now consider further research areas related to CVA in general.

In revaluing the CDS at the time of default of the counterparty (investor), we apply a conditional trivariate copula function to the computation of the reference entity's survival probabilities. The CDS replacement cost is then determined based on these survivals. In reality, the BCVA at τ_2 (τ_0) would be reflected in the replacement cost of the CDS. In other words, if the surviving party replaced the CDS, a BCVA charge would be built into the price of the new trade. This is a consideration that requires further research. Brigo and Morini [98] have begun work on the concept of a substitution closeout amount, based on the ISDA documentation. They take the risk of default of the survived counterparty into account in determining the BCVA. A simple example is illustrated in their paper. Further research into the effect of the substitution closeout amount on the CVA of basic derivatives is required, prior to extending the results to the more complex case of the CVA on a CDS.

Note that the assumption made in this dissertation and in much of the literature on CVA to date, has been that the market is liquid. As such, the derivative whose CVA is computed can be replaced immediately, without liquidity constraints. This is often unlikely, particularly in the South African CDS market, which is extremely illiquid. The final research area that we suggest encompasses collateral, CVA, liquidity and funding costs into the derivative valuation methodology. The association between these concepts renders their joint consideration appropriate.

In [99], Morini and Prampolini analyse the liquidity component in a derivative transaction. Additionally, they examine the effect of the survival probabilities of the investor and of the counterparty on their funding costs and benefits. They build on the work of Piterbarg [100], who considers the difference in discounting rates that are used for CSA and non-CSA versions of the same derivative. Morini and Prampolini illustrate their ideas by means of simplified examples. They find that the introduction of liquidity by means of modifying the discount rate, along with the explicit modelling of the BCVA, can result in double counting.

The research related to these ideas and the link between the liquidity, funding, BCVA and collateral of a derivative is in its infancy. As yet there is not agreement between market participants as to how each of these factors should be incorporated into a derivative's fair value. Further work is required in this area, along with extensions to less simple examples than those that have been considered to date.

Bibliography

- [1] G. Cesari, J. Aquilina, N. Charpillon, Z. Filipović, G. Lee, and I. Manda, *Modelling, Pricing and Hedging Counterparty Credit Exposure. A Technical Guide*. Berlin, Germany: Springer, 2009.
- [2] E. Canabarro and D. Duffie, *Measuring and Marking Counterparty Risk, Asset/liability Management of Financial Institutions*. London, United Kingdom: Euromoney books, 2003.
- [3] Y. Tang and B. Li, *Quantitative Analysis, Derivatives Modeling, and Trading Strategies in the Presence of Counterparty Risk for the Fixed-Income Market*. Singapore: World Scientific Publishing Co. Pte. Ltd, 2007.
- [4] *OTC derivatives market activity in the second half of 2010*, Bank for International Settlements, May 2011.
- [5] M. Pykhtin and S. Zhu, “A guide to modelling counterparty credit risk,” *GARP Risk Review*, vol. 37, 2007.
- [6] S. Alavian, J. Ding, P. Whitehead, and L. Laudicina, “Counterparty valuation adjustment (CVA),” *Working paper, Barclays Capital, Nomura, Deutsche Bank and Nomura*, 2009.
- [7] *Credit Default Swaps and Counterparty Risk*, European Central Bank, 2009. [Online]. Available: www.ecb.europa.eu
- [8] J. Hull, “The credit crunch of 2007: What went wrong? Why? What lessons can be learned?” *University of Toronto*, May 2009.
- [9] G. Chaplin, *Credit Derivatives Risk Management, Trading and Investing*. England, United Kingdom: John Wiley & Sons Ltd, 2005.
- [10] *Credit Event Auction Primer*, Markit. [Online]. Available: www.markit.com/cds

-
- [11] M. Barrachin, C. Liu, and O. Casey, “Hammer it home,” *Markit Magazine*, Spring 2010.
- [12] B. Boettcher, “Price check on aisle 5 - how credit valuation adjustment can help banks enhance the pricing and management of credit risk,” *Think*, December 2009.
- [13] C. Davidson, “Exposing exposure,” *Risk*, March 2010.
- [14] D. Kolitz, *Early Adoption of IFRS Enhances South African Position as Global Player in the Accounting Field*, South African Institute of Chartered Accountants, 2007. [Online]. Available: www.saica.org.za
- [15] *Basel III: A Global regulatory framework for more resilient banks and banking systems*, Bank for International Settlements, 2010.
- [16] C. Redon, “Wrong-way risk modelling,” *Asia Risk*, May 2006.
- [17] M. Pengelly, “Self referencing CDS risk,” *Risk Magazine*, June 2010.
- [18] D. Brigo and M. Masetti, “Chapter 10,” in *Counterparty Credit Risk Modelling, Risk Management Pricing and Regulation*. London: Risk Books, 2005.
- [19] L. Bain and M. Engelhardt, *Introduction to Probability and Mathematical Statistics*. The United States of America: Brooks/Cole, 1992.
- [20] *Credit Valuation Adjustment and challenges of implementation*, Sungard Adaptiv and InteDelta, 2010. [Online]. Available: <http://www.intedelta.com/content/report-download-0?reportID=16logID=48825>
- [21] J. Gregory, “Being two-faced over counterparty credit risk,” *Risk*, vol. 22(2), 2009.
- [22] D. Brigo and A. Capponi, “Bilateral counterparty risk valuation with stochastic dynamical models and application to credit default swaps,” *Working paper*, 2008. [Online]. Available: www.defaultrisk.com
- [23] J. Hull, *Options, Futures and other Derivatives*. New Jersey, United States of America: Prentice Hall, 2006.
- [24] *Morgan Stanley gets risky*, FT Alphaville, October 2010. [Online]. Available: www.ft.com/alphaville

- [25] D. Brigo and F. Mercurio, *Interest Rate Models - Theory and Practice with Smile, Inflation and Credit*. Germany: Springer, 2006.
- [26] J. Stewart, *Calculus, Early Transcendentals*, 4th ed. California, United States of America: Brooks/Cole Publishing Company, 1999.
- [27] K. Imai, "Pol 571: Expectation and functions of random variables," *Princeton University*, March 2006. [Online]. Available: <http://imai.princeton.edu/teaching/files/Expectation.pdf>
- [28] T. Bieliński, M. Jeanblanc, and M. Rutkowski, "Modeling credit risk." Cape Town, South Africa: AIMS Summer School in Mathematical Finance, February 2008.
- [29] R. Merton, "On the pricing of corporate debt: The risk structure of interest rates," *The Journal of Finance*, vol. 3, 1974.
- [30] T. Björk, *Arbitrage Theory in Continuous Time*. New York, United States of America: Oxford University Press, 2004.
- [31] F. Black and J. Cox, "Valuing corporate securities: Some effects of bond indenture provisions," *The Journal of Finance*, vol. 31(2), 1976.
- [32] D. Brigo and M. Tarenghi, "Credit default swap calibration and counterparty risk valuation with a scenario based first passage model," *Working paper, Banca IMI, Italy and Ecole Nationale des Ponts et Chaussées, France*, 2005. [Online]. Available: www.damianobrigo.it/cdsscenario.pdf
- [33] J.-F. Jouanin, G. Rapuch, G. Riboulet, and T. Roncalli, "Modelling dependence for credit derivatives with copulas," *Working paper, Crédit Lyonnais, France*, 2001.
- [34] *The CDS Big Bang: Understanding the Changes to the Global CDS Contract and North American Conventions*, Markit, March 2009. [Online]. Available: www.markit.com/CDS
- [35] *The CDS Small Bang: Understanding the Global Contract and European Contract Changes*, Markit, July 2009. [Online]. Available: www.markit.com/CDS
- [36] *Forthcoming CDS Convention Changes for Japan and Asia*, Markit, December 2009. [Online]. Available: www.markit.com/CDS
- [37] *Upfront and Running Credit Default Swap Conversion without a Big Bang*, Fitch Solutions, April 2009.

-
- [38] *Par Credit Default Swap Spread Approximation from Default Probabilities*, JPMorgan Securities Inc., October 2001.
- [39] J. Hull and A. White, “Valuing credit default swaps II: Modeling default correlations,” *Journal of Derivatives*, vol. 8, 2001.
- [40] L. Chen and D. Filipović, “Pricing credit default swaps under default correlations and counterparty risk,” *Working paper, Princeton University*, 2003.
- [41] P. J. Schönbucher and D. Schubert, “Copula-dependent default risk in intensity models,” *Working paper, Bonn University, Germany*, 2001.
- [42] C. Blanchet-Scalliet and F. Patras, “Counterparty risk valuation for CDS,” *Working paper, Université de Lyon and Université de Nice et CNRS*, 2008.
- [43] A. Lipton and A. Sepp, “Credit value adjustment for credit default swaps via the structural default model,” *The Journal of Credit Risk*, vol. 5(2), 2009.
- [44] A. Marshall and I. Olkin, “A multivariate exponential distribution,” *Journal of American Statistical Association*, vol. 62(317), 1967.
- [45] C. Yi, “Simulating joint defaults and migration II: Credit value adjustment for credit default swaps,” *Working paper, RBC Capital Market, Toronto, Canada*, 2010.
- [46] ———, “Simulating multiple defaults and migration I: Algorithms and default correlations,” *Working paper, Bank of Montreal, Canada*, 2009.
- [47] R. Jarrow and F. Yu, “Counterparty risk and the pricing of defaultable securities,” *The Journal of Finance*, vol. 56, 2001.
- [48] D. Lando, “On Cox processes and credit risky securities,” *Review of Derivatives Research*, vol. 2, 1998.
- [49] R. Jarrow and Y. Yildirim, “A simple model for valuing default swaps when both market and credit risk are correlated,” *Journal of Fixed Income*, vol. 11, 2002.
- [50] F. Yu, “Correlated defaults in intensity-based models,” *Working paper, University of California, Irvine*, 2005.
- [51] S. Leung and Y. Kwok, “Credit default swap valuation with counterparty risk,” *The Kyoto Economic Review*, vol. 74(1), 2005.

- [52] P. Collin-Dufresne, R. Goldstein, and J. Hugonnier, “A general formula for valuing defaultable securities,” *Econometrica*, vol. 72(5), 2004.
- [53] D. Brigo and K. Chourdakis, “Counterparty risk for credit default swaps: Impact of spread volatility and default correlation,” *International Journal of Theoretical and Applied Finance*, vol. 12, no. 7, 2009.
- [54] R. Nelson, *An Introduction to Copulas*. New York: Springer, 1999.
- [55] M. Davis, “Correlation of default rates and multi-asset products,” 2002. [Online]. Available: www2.ma.ic.ac.uk/~mdavis/docs/Davis_GD2002
- [56] M. Walker, “Credit default swaps with counterparty risk: A calibrated Markov model,” *Working paper, University of Toronto*, 2005.
- [57] S. Crépey, M. Jeanblanc, and B. Zargari, “Counterparty risk on a CDS in a Markov chain copula model with joint defaults,” *Working paper, Université d’Evry Val d’Essonne, France and CRIS Consortium*, 2009.
- [58] —, “Counterparty risk on a CDS in a model with joint defaults and stochastic spreads,” *Working paper, Université d’Evry Val d’Essonne, France, Europlace Institute of Finance, Sharif University of Technology, Iran and CRIS Consortium*, 2010.
- [59] T. Bielecki, S. Crépey, M. Jeanblanc, and B. Zargari, “Valuation and hedging of CDS counterparty exposure in a Markov copula model,” *Working paper, Illinois Institute of Technology, Université d’Evry Val d’Essonne, Europlace Institute of Finance, CRIS Consortium and Sharif University of Technology*, 2010.
- [60] *Greek Leader wants to Restrict Speculative Trades*, New York Times, March 2010. [Online]. Available: www.nytimes.com
- [61] *Naked CDS: Exposed*, Derivative Dribble, August 2009. [Online]. Available: <http://derivativedribble.wordpress.com/2009/08/10/naked-cds-exposed/>
- [62] *On the non-existent basis of a (Greek) CDS ban*, FT Alphaville, March 2010. [Online]. Available: www.ft.com/alphaville
- [63] M. Choudhry, “The credit default swap basis: illustrating positive and negative basis arbitrage trades,” July 2006. [Online]. Available: www.yieldcurve.com/Mktresearch/files/Choudhry_BasisTrade_Jul06_Logo.pdf

- [64] *Regulation: EU comes out against naked CDS shorting*, Euromoney, April 2010. [Online]. Available: www.euromoney.com
- [65] *The benefits of naked CDS*, FT Alphaville, March 2010. [Online]. Available: www.ft.com/alphaville
- [66] *Brussels Gives Sovereign CDS Trades All-clear*, Financial Times, December 2010. [Online]. Available: www.ft.com
- [67] *What Critics Don't Get About Credit Default Swaps*, The Atlantic, March 2010. [Online]. Available: www.theatlantic.com
- [68] E. Schlögl, "Default correlation modelling," 2004. [Online]. Available: <http://www.qgroup.org.au/SFMW/CreditDerivativesSchlogl2.pdf>
- [69] D. Brigo and F. Mercurio, "On deterministic-shift extensions of short-rate models," *Working paper, Banca IMI, Italy*, 2001.
- [70] *Markit Credit Indices*, Markit, February 2010. [Online]. Available: www.markit.com
- [71] *Markit Credit Indices, A Primer*, Markit, October 2010. [Online]. Available: www.markit.com
- [72] P. Glasserman, *Monte Carlo Methods in Financial Engineering*. United States of America: Springer-Verlag, 2004.
- [73] L. Andersen, P. Jäckel, and C. Kahl, *Simulation of Square-Root Processes, forthcoming in Encyclopedia of Quantitative Finance*. John Wiley and Sons, 2010.
- [74] D. Brigo and A. Alfonsi, "Credit default swaps calibration and option pricing with the SSRD stochastic intensity and interest-rate model," *Finance and Stochastics*, vol. 9(1), 2005.
- [75] T. Schmidt, "Coping with copulas," in *Copulas - from Theory to Applications in Finance*. England: Risk Books, 2006.
- [76] *Lecture 8: Conditional expectation*, Department of Statistics, University of Wisconsin - Madison, March 2011. [Online]. Available: <http://www.stat.wisc.edu/doksum/STAT709/n709-8.pdf>
- [77] B. Ringnér, "The Law of the Unconscious Statistician," *Centre for Mathematical Sciences, Lund University, Sweden*, January 2009.

- [78] R. Bartle, *The Elements of Real Analysis*, 2nd ed. United States of America: John Wiley & Sons, Inc., 1964.
- [79] M. Broadie and Ö. Kaya, “Exact simulation of stochastic volatility and other affine jump diffusion processes,” *Operations Research*, vol. 54(2), 2006.
- [80] G. Laue, “Remarks on the theory of characteristic functions of non-negative random variables,” *The Indian Journal of Statistics*, vol. 45(1), 1983.
- [81] W. Schoutens, *Lévy Processes in Finance, Pricing Financial Derivatives*. England: Wiley, 2003.
- [82] H. Albrecher, P. Mayer, W. Schoutens, and J. Tistaert, “The Little Heston Trap,” *Wilmott Magazine*, 2007.
- [83] J. Fraleigh and R. Beauregard, *Linear Algebra*, 3rd ed. United States of America: Addison-Wesley Publishing Company, Inc., 1995.
- [84] C. Kahl and P. Jäckel, “Not-so-complex logarithms in the Heston model,” *Wilmott Magazine*, 2005.
- [85] T. Schelfaut, personal communications, September 2010.
- [86] R. Lord and C. Kahl, “Why the rotation count algorithm works,” *Tinbergen Institute discussion paper*.
- [87] A. Steyn, C. Smit, S. du Toit, and C. Strasheim, *Modern Statistics in Practice*, 1st ed. Pretoria, South Africa: J.L. van Schaik, 1994.
- [88] D. Dufresne, “The integrated square-root process,” *Working paper, University of Melbourne*, 2001.
- [89] R. Johnson and D. Wichern, *Applied Multivariate Statistical Analysis*, 5th ed. New Jersey, United States of America: Pearson Education, Prentice-Hall, Inc., 2002.
- [90] G. Chaplin, *Credit Derivatives: Risk Management, Trading and Investing*. England: John Wiley and Sons, 2005.
- [91] A. McNeil, R. Frey, and P. Embrechts, *Quantitative Risk Management, Concepts, Techniques, Tools*. United States of America: Princeton University Press, 2005.

- [92] “The Barclays: Absa takeover approved,” May 2005. [Online]. Available: <http://moneyvibes.co.uk/2005/05/09/barclays-absa/>
- [93] *A CDS inversion unlike any other*, FT Alphaville, July 2011. [Online]. Available: www.ft.com/alphaville
- [94] R. Seydel, *Tools for Computational Finance*, 3rd ed. The Netherlands: Springer, 2006.
- [95] *Glossary - Points upfront*, Creditflux, August 2011. [Online]. Available: www.creditflux.com/Glossary/Points-upfront/
- [96] J. Chan and D. Kroese, “Efficient estimation of large portfolio loss probabilities in t-copula models,” *The University of Queensland*, November 2009.
- [97] C. Milwidsky and E. Maré, “Value at risk in the South African equity market: a view from the tails,” *South African Journal of Economic and Management Sciences*, September 2010.
- [98] D. Brigo and M. Morini, “Dangers of bilateral counterparty risk: the fundamental impact of closeout conventions,” *King’s College, London and Bocconi University, Milan*, November 2010. [Online]. Available: www.defaultrisk.com
- [99] M. Morini and A. Prampolini, “Risky funding with counterparty and liquidity charges,” *Risk*, March 2011.
- [100] V. Piterbarg, “Funding beyond discounting: collateral agreements and derivatives pricing,” *Risk*, February 2010.
- [101] *Conditional Expectation and Martingales*, Department of Statistics, The University of Chicago, March 2011. [Online]. Available: <http://galton.uchicago.edu/lalley/Courses/313/Martingales.pdf>
- [102] K. Matsuda, “Introduction to option pricing with Fourier transform: Option pricing with exponential lévy models,” *Working paper, University of New York*, 2004.
- [103] P. Carr and D. Madan, “Option valuation using the fast Fourier transform,” *Journal of Computational Finance*, vol. 2, 1998.
- [104] K. Chourdakis, “Option pricing using the fractional FFT,” *Journal of Computational Finance*, vol. 8, 2005.

- [105] D. Bailey and P. Swartztrauber, “The fractional Fourier transform and applications,” *SIAM Review*, vol. 33(3), 1991.

Part V

Appendices

Appendix A

Proofs of the General Counterparty Risk Pricing Formulae

The General Counterparty Risk Pricing Formula was introduced in Chapter 3. Since this result is fundamental to the dissertation, the proofs of both the unilateral and bilateral versions of the formula are presented below. These were introduced in Equations (3.1) and (3.3) respectively. The proof of the former can be found in [25] with the latter appearing in [22]. The notation employed in this appendix was introduced in Chapter 3.

A.1 Proof of the Unilateral Formula

Recall that $\Pi^D(t)$ is the payoff of a generic defaultable claim at time t . In this section, defaultable refers to the default risk associated with the counterparty. $\text{CashFlow}(u, s)$ denotes the net cash flows of the claim between times u and s , discounted back to u . Then:

$$\text{NPV}(\tau) = \mathbb{E}_\tau\{\text{CashFlow}(\tau, T)\} \quad (\text{A.1})$$

and:

$$\begin{aligned} \Pi^D(t) = & 1_{\{\tau > T\}} \text{CashFlow}(t, T) + \\ & + 1_{\{\tau \leq T\}} [\text{CashFlow}(t, \tau) + D(t, \tau)(R(\text{NPV}(\tau))^+ - (-\text{NPV}(\tau))^+)]. \end{aligned} \quad (\text{A.2})$$

The expression for the payoff of the defaultable claim in (A.2) is comprised of two parts. The first of these consists of the net cash flows of the claim between times t and T discounted to time t in the event that there is not a counterparty default event prior to T (the maturity of the claim). The second part of Equation (A.2) deals with the payoff of

the claim in the event that the counterparty defaults prior to maturity of the underlying claim. The payoff in this case is the sum of the following two quantities:

- (i) The net cash flows of the claim between time t and the counterparty default time (τ) discounted to t .
- (ii) The discounted value of the expected cash flows of the claim between times τ and T . This is the quantity in brackets: $(R(\text{NPV}(\tau))^+ - (-\text{NPV}(\tau))^+)$. When the expectation of these net cash flows is positive (the first term in brackets), the quantity in brackets is the portion of the NPV that is recovered on default. When their expectation is negative (the second term in brackets), the quantity in brackets is equal to their expectation.

The expected value of the payoff of the defaultable claim in Equation (A.2) is the price of the claim under counterparty risk. The General Counterparty Risk Pricing Formula was introduced in Equation (3.1) for the unilateral case. It is repeated below for ease of reference:

$$\mathbb{E}_t[\Pi^D(t)] = \mathbb{E}_t[\Pi(t)] - \underbrace{\text{LGD} \mathbb{E}_t[1_{\{t < \tau \leq T\}} D(t, \tau) (\text{NPV}(\tau))^+]}_{\text{CVA}}. \quad (\text{A.3})$$

Thus, we are required to prove that the expectation of the right hand side of Equation (A.2) is equivalent to the right hand side of (A.3).

Note firstly that:

$$\Pi(t) = \text{CashFlow}(t, T) = 1_{\{\tau > T\}} \text{CashFlow}(t, T) + 1_{\{\tau \leq T\}} \text{CashFlow}(t, T). \quad (\text{A.4})$$

The expression for the default-free claim in (A.4) is permissible since the events $\{\tau > T\}$ and $\{\tau \leq T\}$ are both mutually exclusive and exhaustive. Utilising (A.4) and the fact that expectation is a linear operator [19], the terms inside the expectation on the right hand side of (A.3) can be written as:

$$\begin{aligned} & 1_{\{\tau > T\}} \text{CashFlow}(t, T) + 1_{\{\tau \leq T\}} \text{CashFlow}(t, T) \\ & + \{(R - 1)[1_{\{\tau \leq T\}} D(t, \tau) (\text{NPV}(\tau))^+]\} \\ & = 1_{\{\tau > T\}} \text{CashFlow}(t, T) + 1_{\{\tau \leq T\}} \text{CashFlow}(t, T) \\ & + R 1_{\{\tau \leq T\}} D(t, \tau) (\text{NPV}(\tau))^+ - 1_{\{\tau \leq T\}} D(t, \tau) (\text{NPV}(\tau))^+. \end{aligned} \quad (\text{A.5})$$

Consider now the expectation of the second and fourth terms in (A.5):

$$\mathbb{E}_t[1_{\{\tau \leq T\}} \text{CashFlow}(t, T) - 1_{\{\tau \leq T\}} D(t, \tau) (\text{NPV}(\tau))^+] \quad (\text{A.6})$$

$$= \mathbb{E}_t\{1_{\{\tau \leq T\}} [\text{CashFlow}(t, T) + D(t, \tau) E_\tau[\text{CashFlow}(\tau, T)] - D(t, \tau) (E_\tau[\text{CashFlow}(\tau, T)])^+]\} \quad (\text{A.7})$$

$$= \mathbb{E}_t\{1_{\{\tau \leq T\}} [\text{CashFlow}(t, \tau) - D(t, \tau) (E_\tau[\text{CashFlow}(\tau, T)])^-]\} \quad (\text{A.8})$$

$$= \mathbb{E}_t\{1_{\{\tau \leq T\}} [\text{CashFlow}(t, \tau) - D(t, \tau) (E_\tau[-\text{CashFlow}(\tau, T)])^+]\} \quad (\text{A.9})$$

$$= \mathbb{E}_t\{1_{\{\tau \leq T\}} [\text{CashFlow}(t, \tau) - D(t, \tau) (-\text{NPV}(\tau))^+]\}. \quad (\text{A.10})$$

The following were utilised in the above derivation:

- (i) Expectation is a linear operator [19].
- (ii) The Tower Property of conditional expectation [101]: $E_t[E_\tau[\cdot]] = E_t[\cdot]$ with $t < \tau$. This was employed in obtaining the second term in the expectation in (A.7).
- (iii) For a function f , $f^+ - f^- = f^+ - (-f)^+$ [25] (used in deriving (A.8) from (A.7)).
- (iv) For a variable X , $\min(0, X) = \max(-X, 0)$ (used in obtaining (A.9) from (A.8)).
- (v) The fact that $\text{NPV}(\tau) = \mathbb{E}_\tau\{\text{CashFlow}(\tau, T)\}$ (see (A.1)) was used in the last step.

Taking the expectation of (A.5) and replacing the 2nd and 4th terms with (A.10), the expectation of the right hand side of (A.2) is obtained (reusing the fact that expectation is a linear operator), as required.

A.2 Bilateral Formula

Consider now the proof of the General Bilateral Counterparty Risk Pricing Formula which was derived in [22]. The formula is repeated below for ease of reference:

$$\begin{aligned} \mathbb{E}_t[\Pi^D(t)] = & \mathbb{E}_t[\Pi(t)] \\ & + \underbrace{\text{LGD}_0 \cdot \mathbb{E}_t\{1_{A \cup B} \cdot D(t, \tau_0) \cdot [-\text{NPV}(\tau_0)]^+\}}_{\text{DVA}} \\ & - \underbrace{\text{LGD}_2 \cdot \mathbb{E}_t\{1_{C \cup D} \cdot D(t, \tau_2) \cdot [\text{NPV}(\tau_2)]^+\}}_{\text{CCVA}}. \end{aligned} \quad (\text{A.11})$$

Using the notation defined in the previous section and in Section 3.3, the discounted payoff of a generic defaultable claim ($\Pi^D(t, T)$) at time t with maturity T is:

$$\begin{aligned} \Pi^D(t, T) &= 1_{E \cup F} \text{CashFlow}(t, T) \\ &\quad + 1_{C \cup D} [\text{CashFlow}(t, \tau_2) + D(t, \tau_2)(R_2(\text{NPV}(\tau_2)))^+ - (-\text{NPV}(\tau_2))^+] \\ &\quad + 1_{A \cup B} [\text{CashFlow}(t, \tau_0) + D(t, \tau_0)((\text{NPV}(\tau_0)))^+ - R_0(-\text{NPV}(\tau_0))^+]. \end{aligned} \quad (\text{A.12})$$

We begin the proof by stating the following:

$$\begin{aligned} \Pi(t, T) &= \text{CashFlow}(t, T) \\ &= 1_{A \cup B} \text{CashFlow}(t, T) + 1_{C \cup D} \text{CashFlow}(t, T) + 1_{E \cup F} \text{CashFlow}(t, T). \end{aligned} \quad (\text{A.13})$$

The expression in (A.13) is permissible since the set of events ($A \cup B$, $C \cup D$ and $E \cup F$), defined in Section 3.3, form a complete set. They are both mutually exclusive and exhaustive. Since expectation is a linear operator [19], the right hand side of (A.11) can be rewritten as:

$$\mathbb{E}_t \{ \Pi(t, T) + \text{LGD}_0 \cdot 1_{A \cup B} \cdot D(t, \tau_0) \cdot [-\text{NPV}(\tau_0)]^+ - \text{LGD}_2 \cdot 1_{C \cup D} \cdot D(t, \tau_2) \cdot [\text{NPV}(\tau_2)]^+ \}. \quad (\text{A.14})$$

Substituting the expression for $\Pi(t)$ in (A.13) into (A.14), we obtain the following:

$$\begin{aligned} &\mathbb{E}_t [1_{A \cup B} \text{CashFlow}(t, T) + (1 - R_0)1_{A \cup B} D(t, \tau_0) [-\text{NPV}(\tau_0)]^+ \\ &\quad + 1_{C \cup D} \text{CashFlow}(t, T) + (R_2 - 1)1_{C \cup D} D(t, \tau_2) [\text{NPV}(\tau_2)]^+ \\ &\quad + 1_{E \cup F} \text{CashFlow}(t, T)] \\ &= \mathbb{E}_t [1_{A \cup B} \text{CashFlow}(t, T) + (1 - R_0)1_{A \cup B} D(t, \tau_0) [-\text{NPV}(\tau_0)]^+ \\ &\quad + \mathbb{E}_t [1_{C \cup D} \text{CashFlow}(t, T) + (R_2 - 1)1_{C \cup D} D(t, \tau_2) [\text{NPV}(\tau_2)]^+] \\ &\quad + \mathbb{E}_t [1_{E \cup F} \text{CashFlow}(t, T)]. \end{aligned} \quad (\text{A.15})$$

We now consider each of the three expected values in (A.15) in turn. The expression inside the first expectation is:

$$\begin{aligned} &1_{A \cup B} \text{CashFlow}(t, T) + (1 - R_0)1_{A \cup B} D(t, \tau_0) [-\text{NPV}(\tau_0)]^+ \\ &= 1_{A \cup B} \text{CashFlow}(t, T) + 1_{A \cup B} D(t, \tau_0) [-\text{NPV}(\tau_0)]^+ - R_0 1_{A \cup B} D(t, \tau_0) [-\text{NPV}(\tau_0)]^+. \end{aligned} \quad (\text{A.16})$$

If we condition on the information available in the market at the default time of the investor (τ_0), the expectation of the expression in (A.16) is as follows:

$$\begin{aligned}
 & \mathbb{E}_{\tau_0}[1_{A \cup B} \text{CashFlow}(t, T) + 1_{A \cup B} D(t, \tau_0)[-NPV(\tau_0)]^+ - R_0 1_{A \cup B} D(t, \tau_0)[-NPV(\tau_0)]^+] \\
 &= \mathbb{E}_{\tau_0}[1_{A \cup B} [\text{CashFlow}(t, \tau_0) + D(t, \tau_0) \text{CashFlow}(\tau_0, T) + D(t, \tau_0)(-\mathbb{E}_{\tau_0}[\text{CashFlow}(\tau_0, T)])^+ \\
 &\quad - R_0 D(t, \tau_0)[-NPV(\tau_0)]^+]] \\
 &= 1_{A \cup B} [\text{CashFlow}(t, \tau_0) + D(t, \tau_0) \mathbb{E}_{\tau_0}[\text{CashFlow}(\tau_0, T)] \\
 &\quad - R_0 D(t, \tau_0)[-NPV(\tau_0)]^+] \\
 &= 1_{A \cup B} [\text{CashFlow}(t, \tau_0) + D(t, \tau_0)(\mathbb{E}_{\tau_0}[\text{CashFlow}(\tau_0, T)])^+ \\
 &\quad - R_0 D(t, \tau_0)[-NPV(\tau_0)]^+] \tag{A.17}
 \end{aligned}$$

$$= 1_{A \cup B} [\text{CashFlow}(t, \tau_0) + D(t, \tau_0)(NPV(\tau_0))^+ - R_0 D(t, \tau_0)[-NPV(\tau_0)]^+]. \tag{A.18}$$

In the above derivation, the first equality follows from the fact that, since $\tau_0 < T$ under the event $A \cup B$, we have that:

$$1_{A \cup B} \text{CashFlow}(t, T) = 1_{A \cup B} [\text{CashFlow}(t, \tau_0) + D(t, \tau_0) \text{CashFlow}(\tau_0, T)]. \tag{A.19}$$

Applying the Tower Rule [101] (since $t < \tau_0$), the first term in (A.15) is obtained by taking the expectation of (A.18) with respect to t . That is:

$$\mathbb{E}_t[1_{A \cup B} [\text{CashFlow}(t, \tau_0) + D(t, \tau_0)(NPV(\tau_0))^+ - R_0 D(t, \tau_0)[-NPV(\tau_0)]^+]]. \tag{A.20}$$

A similar derivation (the detail of which is presented in [22]) can be shown for the second term in (A.15), resulting in the following expression for this term:

$$\mathbb{E}_t[1_{C \cup D} [\text{CashFlow}(t, \tau_2) + D(t, \tau_2) R_2 (NPV(\tau_2))^+ - D(t, \tau_2)[-NPV(\tau_2)]^+]]. \tag{A.21}$$

Substituting the expressions in (A.20) and (A.21) into the places of the first and second expectation in (A.15) respectively, we see that the result obtained is equal to the expected value of the right hand side of (A.12). Therefore, the result is proven.

Appendix B

Matters Related to CDS Pricing

The first section of this appendix describes the procedure for bootstrapping the market-observed term structure of par credit spreads belonging to a particular entity. This process is applied in order to determine the implied term structure of hazard rates and, consequently, the entity's survival probabilities. Deriving this information from CDS spreads is fundamental to the pricing of credit contingent instruments since the market-implied probabilities of survival are required in their valuation. The subject of Section B.2 is the derivation of the expressions for the CDS premium leg in Equation (4.27) and the default leg in Equation (4.28). The formula for the hazard rate under the assumption of constant hazard rates and risk-free interest rates is presented in the final section of this appendix. The result is utilised in Section 12.2.1 of the case study.

B.1 Bootstrapping the Par CDS Curve

For illustrative purposes, the credit spread curve referencing the Republic of South Africa USD denominated debt (SOAF) for close of business (COB) 30 August 2010 is utilised. This curve was obtained from Bloomberg and is presented in Table B.1 below. Note that Bloomberg SOAF quotes reflect a 25% recovery rate assumption.

Term (years)	1	2	3	4	5	7	10
Spread (bp)	81	109	130	144	155	163	170

Table B.1: SOAF spreads for COB 31 August 2010. Source: Bloomberg

Bootstrapping a CDS curve entails the acquisition of the hazard rate term structure that implies a zero fair value for a set of CDSs with maturities equal to the tenors quoted on the CDS curve. The fair value of a CDS with maturity T is the price of the CDS with the premium set equal to the par spread for tenor T . At the outset, a method of interpolation of the market hazard rate is selected. This is generally either piecewise constant or linear interpolation. The relative merits of the two are discussed briefly in [25]. In general, when the term structure of spreads is inverted or CDSs are trading at distressed levels, piecewise constant calibration is the more appropriate of the two interpolation schemes. Both choices are illustrated in this section.

The process underlying the market model for CDS pricing is a time inhomogeneous Poisson process. This was pointed out in Sections 4.2.3 and 4.3.2. The survival probabilities for time inhomogeneous hazard rates were presented in the former section. The formula for the computation of the par spread for a given CDS was derived in the latter section and can be found in Equation (4.34). Note that the specification of a time inhomogeneous Poisson process implies that a term structure of deterministic hazard rates can be derived from quoted CDS spreads. The market model valuation formulae for the premium and protection legs of the CDS were presented in Equations (4.32) and (4.33) in Section 4.3.2 respectively.

A quarterly CDS premium payment frequency is assumed. The values of T_i in (4.32) and (4.33) are thus 0.25, 0.5, 0.75, ..., T . The bootstrapping procedure for the SOAF curve in Table B.1 is as follows:

- (i) Assume firstly that the hazard rates at time 0 and the first tenor for which a spread is quoted are identical. For the SOAF curve in this illustrative example, the first tenor is 1 year. Regardless of the interpolation method assumed for the remainder of the curve, flat left interpolation is utilised from time 0 to the first tenor (1 year).
- (ii) Price a 1 year CDS with a spread of 81bp. Solve for the value of the 1 year hazard rate that implies a zero valuation for the CDS (the default and premium legs are equal in size). It is evident from (4.32) and (4.33) that the integration of the hazard rate is required to obtain the necessary survival probabilities.
- (iii) The next tenor on the SOAF curve is the 2 year point. Price a 2 year CDS using the 1 year hazard rate already obtained, together with an assumed 2 year hazard rate. Either linear or piecewise interpolation can be utilised between the 1 and 2

year points. The chosen interpolation method will be employed in the integration of the hazard rates in the computation of the survival probabilities (as required in the valuation formulae of both the default and premium legs).

- (iv) Solve for the exact value of the 2 year hazard rate that renders the 2 year CDS price equal to zero when a spread of 109bp is utilised in the computation of the premium leg.
- (v) Repeat steps (iii) and (iv) for the remainder of the tenors in the spread curve until the full hazard rate term structure has been obtained. Using this term structure of hazard rates, the CDS prices for CDSs with the same maturities as curve tenors should be zero when the par spreads at the respective points in the curve are inserted into the premium leg valuation. In other words, when the hazard rates (obtained by means of the bootstrapping procedure just described) are used to imply the par spreads for CDSs with maturities equal to the tenors in the spread curve, formula (4.34) should produce the quoted CDS spreads.

The implied hazard rates and survival probabilities for the SOAF curve in Table B.1 are illustrated in Figures B.1(a) and B.1(b) for linear and piecewise constant interpolation of the hazard rates respectively.

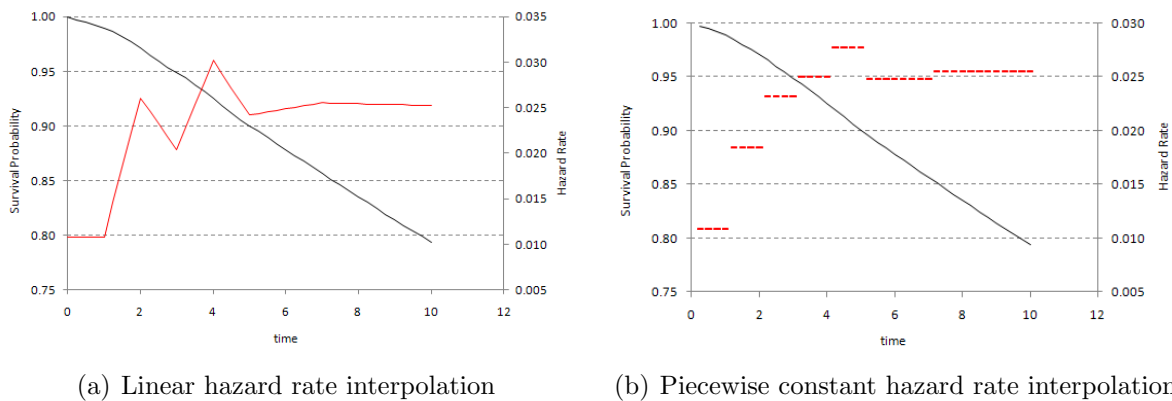


Figure B.1: Implied hazard rate and survival probability term structure for SOAF at COB 31 August 2010

In order to obtain curves of varying levels of riskiness, 200bp and 500bp parallel shifts were applied to the SOAF curve in Table B.1. The market-implied hazard rates and survival probabilities for these curves were then obtained using the bootstrapping procedure described in this section, with the recovery assumption remaining 25%. The three CDS

curves, along with the corresponding (linearly interpolated) hazard rates and survival probabilities are reported in Table B.2 below. The spreads and survivals are illustrated in Figure B.1. In this graph, the dotted curves represent the survival probabilities (shown on the primary axis) while the solid lines display the corresponding CDS spreads (with values shown on the secondary vertical axis). These three curves, that represent varying levels of riskiness, are utilised in the numerical illustrations throughout the dissertation.

Term\ Shift	Spread(bp)			Hazard Rate			Survival Probability		
	0bp	100bp	400bp	0bp	100bp	400bp	0bp	100bp	400bp
1	81	181	481	0.011	0.024	0.064	99%	98%	94%
2	109	209	509	0.026	0.032	0.080	97%	95%	87%
3	130	230	530	0.020	0.037	0.074	95%	91%	81%
4	144	244	544	0.030	0.039	0.085	93%	88%	75%
5	155	255	555	0.024	0.041	0.078	90%	84%	69%
7	163	263	563	0.026	0.038	0.079	86%	78%	59%
10	170	270	570	0.025	0.039	0.080	79%	69%	46%

Table B.2: SOAF spreads, hazard rates and survival probabilities with spread shifts




Figure B.2: CDS spreads and survival probabilities for curves in Table B.2

B.2 Derivation of the CDS Premium and Default Leg Valuation Formulae

The purpose of this section is to present derivations of the premium and default legs in Equations (4.27) and (4.28) respectively. The derivations can be found in [25].

B.2.1 CDS Premium Leg (Equation (4.27))

The formula for the valuation of the CDS premium leg was presented in Equation (4.24). It is repeated below for ease of reference (with notation defined in Section 4.3):

$$\begin{aligned} & \text{Premium}_{a,b}(t; S) \\ &= S \sum_{i=a+1}^b D(t, T_i) \alpha_i 1_{\{\tau \geq T_i\}} + SD(t, \tau) (\tau - T_{\beta(\tau)-1}) 1_{\{T_a < \tau < T_b\}}. \end{aligned} \quad (\text{B.1})$$

The premium leg is valued by taking the expectation of (B.1). The valuation is performed without loss of generality at time 0 ($t = 0$). Since the expected value is a linear operator [19]:

$$\begin{aligned} & \text{Premium}_{a,b}(0; S) \\ &= S \left[\sum_{i=a+1}^b \mathbb{E}[D(0, T_i) \alpha_i 1_{\{\tau \geq T_i\}}] + \mathbb{E}[D(0, \tau) (\tau - T_{\beta(\tau)-1}) 1_{\{T_a < \tau < T_b\}}] \right] \\ &= S \left[\sum_{i=a+1}^b D(0, T_i) \alpha_i \mathbb{E}[1_{\{\tau \geq T_i\}}] + \mathbb{E} \left(\int_0^\infty D(0, t) (t - T_{\beta(t)-1}) 1_{\{T_a < \tau < T_b\}} 1_{\tau \in [t, t+dt]} \right) \right] \\ &= S \left[\sum_{i=a+1}^b D(0, T_i) \alpha_i \mathbb{Q}(\tau \geq T_i) + \int_{t=T_a}^{T_b} \mathbb{E}[D(0, t) (t - T_{\beta(t)-1}) 1_{\{T_a < \tau < T_b\}} 1_{\tau \in [t, t+dt]}] \right] \quad (\text{B.2}) \\ &= S \left[\sum_{i=a+1}^b D(0, T_i) \alpha_i \mathbb{Q}(\tau \geq T_i) + \int_{t=T_a}^{T_b} D(0, t) (t - T_{\beta(t)-1}) \mathbb{E}[1_{\{T_a < \tau < T_b\}} 1_{\tau \in [t, t+dt]}] \right] \\ &= S \left[\sum_{i=a+1}^b D(0, T_i) \alpha_i \mathbb{Q}(\tau \geq T_i) + \int_{t=T_a}^{T_b} D(0, t) (t - T_{\beta(t)-1}) \mathbb{Q}(\tau \in [t, t + dt]) \right] \\ &= S \left[\sum_{i=a+1}^b D(0, T_i) \alpha_i \mathbb{Q}(\tau \geq T_i) - \int_{t=T_a}^{T_b} D(0, t) (t - T_{\beta(t)-1}) d_t \mathbb{Q}(\tau \geq t) \right]. \end{aligned} \quad (\text{B.3})$$

The following were utilised in the above derivation:

- (i) Fubini's theorem (see [26]) in switching the time integral with the expectation in (B.2).
- (ii) The discount factor and the default time are assumed to be independent (since the expectation of the product of these two quantities was used interchangeably with the product of their expectations).
- (iii) The deterministic nature of the discount factor implied that the expectation of this variable was simply the value of the variable itself.
- (iv) For a set A , $\mathbb{E}[1_{\{A\}}(X)] = P(X \in A)$ [27].

The expression in (B.3) is the same as Formula (4.27), as required.

B.2.2 CDS Protection (Default) Leg (Equation (4.28))

The default leg in this section is derived in a similar manner to the premium leg in the previous one. The formula for the valuation of the protection leg was presented in Equation (4.25). It is repeated below for ease of reference:

$$\text{Default}_{a,b}(t; S) = D(t, \tau) \text{LGD} 1_{\{T_a < \tau \leq T_b\}}. \quad (\text{B.4})$$

Taking the expectation of the right hand side of (B.4) and assuming again that the valuation is performed at time 0 ($t = 0$), the derivation of the formula for the valuation of the default leg is as follows:

$$\begin{aligned} \text{Default}_{a,b}(0; S) &= \mathbb{E}[D(0, \tau) \text{LGD} 1_{\{T_a < \tau \leq T_b\}}] \\ &= \text{LGD} \mathbb{E}\left[\int_0^\infty 1_{\{T_a < \tau \leq T_b\}} D(0, t) 1_{\{\tau \in [t, t+dt)\}}\right] \\ &= \text{LGD} \mathbb{E}\left[\int_{T_a}^{T_b} D(0, t) \mathbb{E}[1_{\{\tau \in [t, t+dt)\}}]\right] \\ &= \text{LGD} \mathbb{E}\left[\int_{T_a}^{T_b} D(0, t) \mathbb{Q}(\tau \in [t, t+dt))\right] \\ &= -\text{LGD} \int_{T_a}^{T_b} D(0, t) d_t \mathbb{Q}(\tau \geq t). \end{aligned} \quad (\text{B.5})$$

The techniques used in deriving (B.5) were all utilised in the derivation of the premium leg (Equation (B.3)). The final formulation of the default leg in (B.5) is identical to the one in (4.28), as required.

B.3 Derivation of the Constant Hazard Rate

In this final section of Appendix B, the derivation of the hazard rate in Equation (12.2) in Section 12.2.1 is shown. Assume that the following parameters are constant:

- (i) The hazard rate, λ .
- (ii) The risk-free interest rate, r_f .
- (iii) The time between CDS premium payments, t_c .

The market model valuation formula for the premium leg of a CDS was shown in (4.32) in Section 4.3.2. It is repeated below for ease of reference:

$$\begin{aligned} & \text{Premium}_{a,b}(0; S) \{\text{market model}\} \\ &= S \sum_{i=a+1}^b \left[D(0, T_i) \alpha_i e^{-\int_0^{T_i} \lambda(s) ds} + D(0, T_i) (e^{-\int_0^{T_{i-1}} \lambda(s) ds} - e^{-\int_0^{T_i} \lambda(s) ds}) \frac{\alpha_i}{2} \right]. \end{aligned} \quad (\text{B.6})$$

The notation employed in (B.6) was defined in Section 4.3.2. If the partial payment portion of the premium leg is omitted then, under the assumptions made above, the premium leg reduces to the following:

$$\begin{aligned} \text{Premium}_{a,b} \{\text{simplified}\} &= S \sum_{i=a+1}^b e^{-r_f T_i} t_c e^{-\lambda T_i} \\ &= S \sum_{i=a+1}^b t_c e^{-t_c \times i(r_f + \lambda)}. \end{aligned} \quad (\text{B.7})$$

The market model valuation formula for the default leg can be found in (4.33) in Section 4.3.2. It is repeated below for ease of reference:

$$\text{Default}_{a,b}(0; S) \{\text{market model}\} = \text{LGD} \sum_{i=a+1}^b D(0, T_i) (e^{-\int_0^{T_{i-1}} \lambda(s) ds} - e^{-\int_0^{T_i} \lambda(s) ds}). \quad (\text{B.8})$$

Under the assumptions listed at the beginning of this section, the default leg reduces to the following:

$$\begin{aligned}
 \text{Default}_{a,b}(0; S)\{\text{simplified}\} &= \text{LGD} \sum_{i=a+1}^b e^{-r_f T_i} (e^{-\int_0^{T_i-1} \lambda(s) ds} - e^{-\int_0^{T_i} \lambda(s) ds}) \\
 &= \text{LGD} \sum_{i=a+1}^b e^{-r_f t_c \times i} (e^{-\lambda t_c \times i} - e^{-\lambda t_c \times (i-1)}) \\
 &= \text{LGD} \sum_{i=a+1}^b e^{-(r_f + \lambda) t_c \times i} (1 - e^{\lambda t_c}). \tag{B.9}
 \end{aligned}$$

Since the default and premium legs on a par CDS are equal, we set (B.7) equal to (B.9) and solve for the value of λ that this implies:

$$\begin{aligned}
 S \sum_{i=a+1}^b t_c e^{-t_c \times i (r_f + \lambda)} &= \text{LGD} \sum_{i=a+1}^b e^{-(r_f + \lambda) t_c \times i} (1 - e^{\lambda t_c}) \\
 e^{\lambda t_c} &= \left(\frac{S \times t_c}{\text{LGD}} + 1 \right) \\
 \lambda &= \frac{1}{t_c} \ln \left(\frac{S \times t_c}{\text{LGD}} + 1 \right),
 \end{aligned}$$

as required.

Appendix C

Formulae Related to the CIR Process

The objective behind this appendix is to set forth CIR-related formulae that are referenced in the main text.

The CIR dynamics are specified as:

$$dy(t) = \kappa(\mu - y(t))dt + \nu\sqrt{y(t)}dZ(t), \quad y(0) = y_0, \quad (\text{C.1})$$

with

$$h = \sqrt{\kappa^2 + 2\mu^2}. \quad (\text{C.2})$$

The formulae below follow.

C.1 CIR Instantaneous Forward Rate

Under the dynamics in (C.1), the instantaneous forward rate at time 0 for time t is calculated as follows:

$$f^{CIR}(0, t, \beta) = 2\kappa\mu \frac{e^{th} - 1}{2h + (\kappa + h)(e^{th} - 1)} + y_0 \frac{4h^2 e^{th}}{[2h + (\kappa + h)(e^{th} - 1)]^2} \quad (\text{C.3})$$

C.2 CIR Bond Price

The price at time t of a zero-coupon bond, maturing at time T is:

$$P(t, T) = A(t, T)e^{-B(t, T)y(t)}, \quad (\text{C.4})$$

where the formulae for $A(t, T)$ and $B(t, T)$ are specified as:

$$A(t, T) = \left[\frac{2he^{(k+h)(T-t)/2}}{2h + (k+h)(e^{(T-t)h} - 1)} \right]^{2\kappa\mu/\nu^2} \quad (\text{C.5})$$

$$B(t, T) = \frac{2(e^{(T-t)h} - 1)}{2h + (k+h)(e^{(T-t)h} - 1)}. \quad (\text{C.6})$$

C.3 CIR Conditional Expected Value

Given the dynamics of the CIR process in Equation (C.1), the conditional expectation of $y(t)$ given $y(s)$, $s < t$, is obtained by utilising:

$$E[y(t)|y(s)] = y(s)e^{-\kappa(t-s)} + \mu(1 - e^{-\kappa(t-s)}). \quad (\text{C.7})$$

C.4 CIR Conditional Variance

The conditional variance of $y(t)$ given the value of $y(s)$ ($s < t$), where the dynamics of y were specified in Equation (C.1), is computed as follows:

$$\text{Var}[y(t)|y(s)] = y(s)\frac{\sigma^2}{\kappa}(e^{-\kappa(t-s)} - e^{-2\kappa(t-s)}) + \mu\frac{\sigma^2}{2\kappa}(1 - e^{-\kappa(t-s)})^2. \quad (\text{C.8})$$

Appendix D

A Basic Introduction to Copulas

The basic concepts related to copulas are presented in this appendix. It is intended only to provide a rudimentary background to the topic, the main purpose of which is to ensure that the employment of copulas in the CDS CVA model in this dissertation is understood. The references used throughout are [54] and [25], unless otherwise stated. The initial definitions are presented for a bivariate copula since the two variable case permits an intuitive understanding of the required concepts. The extension to higher orders is straightforward. It is presented where necessary for a comprehension of the trivariate copula, which was employed in the BCVA model to link the uniform random variables associated with the default times of the three entities.

The intuition behind copulas is uncomplicated: they are functions that join or couple multivariate distribution functions to their one-dimensional, uniform margins.

D.1 Concepts Required to Define Copulas

An understanding of the terms *grounded* and *2-increasing* are necessary background to the exact definition of a bivariate copula function. The former in turn requires the definition of *volume*.

Volume definition: Let S_1 and S_2 be non empty subsets of $\overline{\mathbb{R}}$ (the extended real line $[-\infty, \infty]$) and let H be a function such that the domain of H is given by $S_1 \times S_2$. Let B be a rectangle, all of whose vertices are in the domain of H . The the H -volume of B is:

$$V_H(B) = H(x_2, y_2) - H(x_2, y_1) - H(x_1, y_2) + H(x_1, y_1). \quad (\text{D.1})$$

A 2-increasing function can now be defined.

Definition of a 2-increasing function: A 2-place real function H is *2-increasing* if $V_H(B) \geq 0$ for all rectangles B whose vertices lie in the domain of H .

The remaining concept that requires a definition prior to defining the copula function is a grounded function.

Definition of a grounded function: Suppose that the smallest elements of S_1 and S_2 are a_1 and a_2 respectively. We say that a function H from $S_1 \times S_2$ into \mathbf{R} is *grounded* if

$$H(x, a_2) = 0 = H(a_1, y)$$

for all (x, y) in $S_1 \times S_2$.

D.2 The Definition of a Copula

Definition of a Subcopula: A *two-dimensional subcopula* is a function C' with the following properties:

- (i) The domain of C' is $S_1 \times S_2$ where S_1 and S_2 are subsets of $\mathbf{I} = [0, 1]$ containing 0 and 1;
- (ii) C' is grounded and 2-increasing;
- (iii) For every u in S_1 and every v in S_2 ,

$$C'(u, 1) = u \text{ and } C'(1, v) = v.$$

Definition of a Copula: A *two-dimensional copula* is a two-dimensional subcopula whose domain is \mathbf{I}^2 . This means that a copula is a function C from \mathbf{I}^2 to \mathbf{I} with the following properties:

- (i) For every u, v in \mathbf{I} ,

$$C(u, 0) = 0 = C(0, v) \tag{D.2}$$

(the function is grounded)

and

$$C(u, 1) = u \text{ and } C(1, v) = v \tag{D.3}$$

(property 3 in the definition of a subcopula).

(ii) For every u_1, u_2, v_1, v_2 in \mathbf{I} such that $u_1 \leq u_2$ and $v_1 \leq v_2$,

$$C(u_2, v_2) - C(u_2, v_1) - C(u_1, v_2) + C(u_1, v_1) \geq 0 \quad (\text{D.4})$$

(the function C is 2-increasing).

D.3 Sklar's Theorem

The foundation of most applications of copulas to statistics is Sklar's theorem. It explains the relationship between multivariate distribution functions and their univariate margins. For a bivariate copula, it reads as follows:

For any joint distribution function $H(x, y)$ with margins F and G , there exists a copula $C(u, v)$ (i.e. a joint distribution function on 2 uniforms) such that

$$H(x, y) = C(F(x), G(y)). \quad (\text{D.5})$$

When H is the n -dimensional distribution function with margins F_1, F_2, \dots, F_n , then there exists an n -copula C such that for all \mathbf{x} in $\overline{\mathbb{R}}^n$:

$$H(x_1, x_2, \dots, x_n) = C(F_1(x_1), F_2(x_2), \dots, F_n(x_n)). \quad (\text{D.6})$$

Note that C contains pure dependence information. Note also that the copula function is an n -dimensional function, as opposed to representing dependence by means of a mere scalar. Equation (D.6) can also be expressed as:

$$C(u_1, u_2, \dots, u_n) = H(F_1^{-1}(u_1), F_2^{-1}(u_2), \dots, F_n^{-1}(u_n)), \quad (\text{D.7})$$

where $F_1(x_1) = u_1, F_2(x_2) = u_2, \dots, F_n(x_n) = u_n$. The form in (D.7) shows that any known joint distribution function H can be used to define a copula.

Next, it is shown why the cumulative distribution function (CDF) of a random variable is uniformly distributed.

D.4 The Uniform Distribution of CDFs

Since it has been pointed out in this appendix that copulas are joint distributions on n uniforms and since Sklar's theorem expresses the copula as a function of the univariate margins (see (D.6)), it is important to understand the link between the CDF and a uniform random variable. This is demonstrated below.

Consider a transformation of variable X by its cumulative distribution function $F_X(X)$. To this end, set $U = F_X(X)$ and assume that F_X is invertible. An expression for $F_U(U)$ at a point u in the interval $[0, 1]$ is now obtained. From the definition of a cumulative distribution function, it follows that:

$$F_U(u) = \mathbb{Q}(U \leq u) = \mathbb{Q}(F_X(X) \leq u) = \mathbb{Q}(F_X(X) \leq F_X(F_X^{-1}(u))). \quad (\text{D.8})$$

The second equality in (D.8) is a result of the choice to set $U = F(X)$. The final equality in (D.8) is permissible since it was assumed that F is invertible. Therefore:

$$\mathbb{Q}(F_X(X) \leq F_X(F_X^{-1}(u))) = \mathbb{Q}(X \leq F_X^{-1}(u)) = F_X(F_X^{-1}(u)) = u.$$

So, $F_U(u) = u$, which is characteristic of the uniform random variable in $[0, 1]$. Thus $U = F_X(X)$ is a uniform random variable. Since F_X is one-to-one, U contains the same information as X .

Clearly, all random variables can be transformed by their cumulative distribution functions to obtain uniform random variables that contain the same information as the starting variables. Therefore the link between the univariate marginals F and G in Sklar's theorem and the uniform random variables in copulas has been established. The definitions of two remaining concepts that are referenced in the dissertation follow. The chapter is then concluded with a short proof of the statement in Section 10.2 that a trivariate Gaussian copula induces bivariate marginals.

D.5 Copula Densities

The copula density is defined as [75]:

$$c(\mathbf{u}) := \frac{\partial^d C(u_1, u_2, \dots, u_d)}{\partial u_1 \partial u_2 \dots \partial u_d}. \quad (\text{D.9})$$

D.6 Copula Conditional Distributions

The copula conditional distribution for the bivariate case is presented in [75] as follows:

$$\begin{aligned}
 P(U_2 \leq u_2 | U_1 = u_1) &= \lim_{\delta \rightarrow 0} \frac{P(U_2 \leq u_2, U_1 \in (u_1 - \delta, u_1 + \delta))}{P(u_1 - \delta, u_1 + \delta)} \\
 &= \frac{C(u_1 + \delta, u_2) - C(u_1 - \delta, u_2)}{2\delta} \\
 &= \frac{\partial}{\partial u_1} C(u_1, u_2).
 \end{aligned} \tag{D.10}$$

Extending this to a trivariate copula:

$$\begin{aligned}
 P(U_2 \leq u_2, U_3 \leq u_3 | U_1 = u_1) &= \lim_{\delta \rightarrow 0} \frac{P(U_2 \leq u_2, U_3 \leq u_3, U_1 \in (u_1 - \delta, u_1 + \delta))}{P(u_1 - \delta, u_1 + \delta)} \\
 &= \frac{C(u_1 + \delta, u_2, u_3) - C(u_1 - \delta, u_2, u_3)}{2\delta} \\
 &= \frac{\partial}{\partial u_1} C(u_1, u_2, u_3).
 \end{aligned} \tag{D.11}$$

D.7 Proof of the Statement that a Trivariate Gaussian Copula Induces Bivariate Marginals

This section of the appendix is included in order to substantiate the well-known claim in Section 10.2 that the two-dimensional marginal distributions of the trivariate Gaussian copula are themselves Gaussian. The proof is straightforward and can also be found in [22]. Suppose that the 3×3 correlation matrix associated with the trivariate copula is represented by the symbol Σ . The proof for the marginal distribution of U_1 and U_2 is provided below. The marginal copulas linking the remaining combinations of pairs of the

three uniform random variables (U_1 , U_2 and U_3) can be derived in a similar manner.

$$\mathbb{Q}(U_1 < u_1, U_2 < u_2) = \mathbb{Q}(U_1 < u_1, U_2 < u_2, U_3 < 1) \quad (\text{D.12})$$

$$= C_{\Sigma}(u_1, u_2, 1) \quad (\text{D.13})$$

$$= \lim_{b \rightarrow 1^-} \Phi_{\Sigma}(\Phi^{-1}(u_1), \Phi^{-1}(u_2), \Phi^{-1}(b)) \quad (\text{D.14})$$

$$= \lim_{b \rightarrow 1^-} \mathbb{Q}(X_1 < \Phi^{-1}(u_1), X_2 < \Phi^{-1}(u_2), X_3 < \Phi^{-1}(b)) \quad (\text{D.15})$$

$$= \mathbb{Q}(X_1 < \Phi^{-1}(u_1), X_2 < \Phi^{-1}(u_2)).$$

$$\lim_{b \rightarrow 1^-} \mathbb{Q}(X_3 < \Phi^{-1}(b) | X_1 < \Phi^{-1}(u_1), X_2 < \Phi^{-1}(u_2)) \quad (\text{D.16})$$

$$= \Phi_{\Sigma_{1,2}}(\Phi^{-1}(u_1), \Phi^{-1}(u_2)) \quad (\text{D.17})$$

$$= C_{1,2}(u_1, u_2). \quad (\text{D.18})$$

The derivation is explained below:

- (i) The equality in (D.12) follows from the fact that the uniform random variable U_3 has an upper bound of 1.
- (ii) The expression in (D.13) is a result of the definition of a copula function. The subscript Σ indicates that the correlation matrix for the three variables is Σ .
- (iii) Expression (D.14) follows from the definition of a trivariate Gaussian copula. Additionally, the limit is necessary since the inverse of the cumulative standard normal distribution is undefined at the point 1. The superscript ‘-’ in the limit points to the fact that the limit is taken from the left hand side (since $U_3 < 1$). The symbol Φ represents the standard cumulative normal distribution function, with Φ^{-1} representing its inverse.
- (iv) The expression in (D.16) utilises the well-known result from elementary statistics (employed in numerous places in the dissertation) that for two events A and B , $P(B) \cdot P(A|B) = P(A \cap B)$.
- (v) The symbol $\Sigma_{1,2}$ represents the portion of the correlation matrix containing the correlation associated with U_1 and U_2 .
- (vi) The last line (D.18) employs the definition of the bivariate Gaussian copula.

Appendix E

Supplementary Information Relating to Chapter 8

The derivations of Equations (8.16) and (8.27) in Chapter 8 were relegated to this appendix. They are presented in Sections E.1 and E.2 respectively. Additionally, the discussion pertaining to the testing of the probabilities produced by the FRFT procedure against the close form solution of the CIR bond price follows in Section E.3.

E.1 Approximating the Integral using the Trapezoidal Rule

The Trapezoidal Rule, for approximating the integral of a function $f(x)$ from $x = a$ to $x = b$, is as follows [26]:

$$\int_a^b = \sum_j \frac{[f(x_{i-1}) + f(x_i)]}{2} \Delta x. \quad (\text{E.1})$$

We apply this rule (E.1) to

$$\int_{u_1=\bar{u}_1}^1 P(u_1) dC_{1|2}(u_1, U_2),$$

the expression for the condition survival probability that was derived in Equation (8.6). This results in

$$\int_{u_1=\bar{u}_1}^1 P(u_1) dC_{1|2}(u_1, U_2) \approx \sum_j \frac{[P(u_{j-1}) + P(u_j)]}{2} \Delta C_{1|2}(u_1, U_2). \quad (\text{E.2})$$

From (8.9),

$$P(u_j) = P(X < x_j | \mathcal{G}_{\tau_2}) = p_j,$$

where $X = Y_1(T_k) - Y_1(\tau_2)$ and $x_j = -\ln(1 - u_j) - Y_1(\tau_2) - \Psi_1(T_k)$. In addition, if $f_j = C_{1|2}(u_j, U_2)$, the sum in (E.2) can be expressed as:

$$\mathbb{Q}(\tau_1 > T_k | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2) \approx \sum_j \frac{p_{j+1} + p_j}{2} \Delta f_j,$$

as required.

E.2 Derivation of the New Characteristic Function Formulation

This section demonstrates the derivation of the form of the characteristic function in Equation (8.27) of Section 8.2.4.

The original form of the characteristic function, taken from [81] was:

$$\phi(u) = \frac{e^{\kappa^2 \mu t / \nu^2} e^{2y_0 i u / (\kappa + \gamma \coth(\gamma t / 2))}}{(\cosh(\gamma t / 2) + \kappa \sinh(\gamma t / 2) / \gamma)^{2\kappa \mu / \nu^2}}, \quad (\text{E.3})$$

where

$$\gamma = \sqrt{\kappa^2 - 2\nu^2 i u}.$$

Using the fact that

$$\cosh x = \frac{1}{2}(e^x + e^{-x}) \quad \text{and} \quad \sinh x = \frac{1}{2}(e^x - e^{-x}),$$

the denominator in (E.3) is rewritten as:

$$\begin{aligned} \text{Denominator} &= \left[\frac{1}{2}(e^{\gamma t / 2} + e^{-\gamma t / 2}) + \frac{\kappa}{\gamma} \frac{1}{2}(e^{\gamma t / 2} - e^{-\gamma t / 2}) \right]^{2\kappa \mu / \nu^2} \\ &= \left[\frac{1}{2} \left(1 + \frac{\kappa}{\gamma} \right) e^{\gamma t / 2} + \frac{1}{2} \left(1 - \frac{\kappa}{\gamma} \right) e^{\gamma t / 2} \right]^{2\kappa \mu / \nu^2} \\ &= \left[e^{\gamma t / 2} \left(\frac{1}{2} \left(1 + \frac{\kappa}{\gamma} \right) + \frac{1}{2} \left(1 - \frac{\kappa}{\gamma} \right) e^{-2\gamma t / 2} \right) \right]^{2\kappa \mu / \nu^2} \\ &= e^{\frac{\gamma t \kappa \mu}{\nu^2}} \left[\frac{1}{2} \left(1 + \frac{\kappa}{\gamma} + e^{-\gamma t} \left(1 - \frac{\kappa}{\gamma} \right) \right) \right]^{2\kappa \mu / \nu^2}. \end{aligned}$$

Changing the sign of the $e^{\frac{\gamma t \kappa \mu}{\nu^2}}$ term, putting it in the numerator and using the fact that

$$e^{\frac{\kappa^2 \mu t}{\nu^2}} e^{-\frac{\gamma t \kappa \mu}{\nu^2}} = \frac{\kappa \mu t}{\nu^2} (\kappa - \gamma),$$

the following is obtained:

$$\phi(u) = \frac{\exp\left(\frac{\kappa \mu t}{\nu^2} (\kappa - \gamma)\right) \exp\left(\frac{2y_0 i u}{\kappa + \gamma \coth(\gamma t/2)}\right)}{\left[\frac{1}{2}\left(1 + \frac{\kappa}{\gamma} + e^{-\gamma t}\left(1 - \frac{\kappa}{\gamma}\right)\right)\right]^{2\kappa \mu / \nu^2}},$$

as required.

E.3 Testing the Probabilities against the CIR Bond Price

This section provides an explanation of the exact methodology employed to test the probabilities computed by means of the FRFT in Section 8.2 against the CIR bond price (which has a closed form solution that can be found in Section C.2 of Appendix C). It is assumed that the deterministic shift portion of the CIR++ process is zero since the cumulative distribution of the integrated CIR process is being tested (the stochastic part of the CIR++ process).

Recall from Section 8.3 that the conditional distribution f_j is given by:

$$C_{1|2}(u_j, U_2) = \frac{\frac{\partial}{\partial u_2} C(u_j, U_2) - \frac{\partial}{\partial u_2} C(\bar{U}_1, U_2)}{1 - \frac{\partial}{\partial u_2} C(\bar{U}_1, U_2)}.$$

Assuming independence between U_1 and U_2 , the bond price can be calculated. The independence copula is given by $C(u, v) = uv$. Therefore:

$$C_{1|2}(u_j, U_2) = \frac{u_j - \bar{U}_1}{1 - \bar{U}_1}. \quad (\text{E.4})$$

The following quantities are required to compute the conditional copula function in (E.4):

$$\begin{aligned} \bar{U}_1 &= 1 - e^{-Y_1(\tau_2) - \Psi_1(\tau_2)} \\ u_j &= 1 - e^{-x_j - Y_1(\tau_2) - \Psi_1(T_k)}. \end{aligned} \quad (\text{E.5})$$

Assume that we are standing at time τ_2 . This facilitates comparison with the current bond price. The deterministic portion also falls away since the CIR (and not the CIR++) bond price is being computed. Therefore, the above expressions become:

$$\bar{U}_1 = 1 - e^0 = 1 - 1 = 0 \quad (\text{E.6})$$

$$u_j = 1 - e^{-x_j - 0} = 1 - e^{-x_j}. \quad (\text{E.7})$$

Lastly, the following sum is computed using (E.6) and (E.4):

$$\sum_j \frac{p_{j+1} + p_j}{2} \Delta f_j. \quad (\text{E.8})$$

The sum in (E.8) is compared against the CIR bond price with the matching maturity. In testing the cumulative distribution function (CDF) generated via the FRFT against the CIR bond price, the largest difference considered acceptable was in the order of E^{-3} . Thinking about this in terms of survival probabilities, it is an acceptable level of accuracy in practice. Differences in survivals due to interpolation methods can be larger.

Appendix F

Fourier Transforms

The subject of this appendix is a very brief introduction to the concepts in Fourier theory required to implement the CDS CVA model. The paper by Matsuda [102] provides a useful background to Fourier theory and the functions associated with Fourier transforms. Various other references are cited below.

F.1 Definition of a Fourier Transform

The Fourier transform (FT) of a function $g(t)$ to a function $G(\omega)$ is defined as [102]:

$$G(\omega) = \int_{-\infty}^{\infty} e^{i\omega t} g(t) dt. \quad (\text{F.1})$$

The inverse Fourier transform is [102]:

$$g(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-i\omega t} G(\omega) d\omega, \quad (\text{F.2})$$

where $i = \sqrt{-1}$. A more general definition exists, but for the purpose of understanding the application in this dissertation, the above definition will suffice. Note that t is the time domain while ω is the angular frequency domain. The FT thus transforms the function $g(t)$ from the time domain to the angular frequency domain. The inverse FT performs the opposite on the function $G(\omega)$.

The definition of a characteristic function is required next.

F.2 Characteristic Functions

Let X be the random variable with probability density function $\mathbb{P}(x)$. The characteristic function $\phi(\omega)$ with $\omega \in \mathbb{R}$ is defined as the FT of the probability density function $\mathbb{P}(x)$ [102]. Employing the definition of a Fourier transform in (F.1),

$$\phi(\omega) = \int_{-\infty}^{\infty} e^{i\omega x} \mathbb{P}(x) dx = \mathbb{E}[e^{i\omega x}]. \quad (\text{F.3})$$

F.3 Fourier Transforms in Finance

There has been much work, particularly in recent years, on the use of FTs in finance. A large portion of this work is associated with option pricing and the Heston stochastic volatility model. In 1999, Carr and Madan [103] illustrated how the fast Fourier transform (FFT) could be applied to option valuation, a major step forward in the use of Fourier theory in finance. The FFT is an efficient algorithm for approximating the continuous Fourier transform with its discrete counterpart by computing sums of the form:

$$\sum_{j=0}^{N-1} e^{-i\frac{2\pi}{N}(j)(k)} h_j, \quad k = 0, 1, \dots, N-1 \quad (\text{F.4})$$

for a vector $\mathbf{h} = (h_j)_{j=0}^{N-1}$ [104]. The output of the FFT procedure is a new vector $\mathbf{f} = (f_k)_{k=0}^{N-1}$. Each element f_k of this vector corresponds to a sum of the form (F.4), computed at a pre-determined value of $x = x_k$ [104]. The paper by Carr and Madan [103] is useful for understanding the background on how FFTs are applied in finance.

F.4 Shortcomings of the FFT

The integral approximated by the sum in (F.4) is:

$$\int_0^{\infty} e^{-ixu} g(u) du, \quad (\text{F.5})$$

for a function g . When implementing the FFT, we select the grid size of vector u in (F.5) which we will denote as δ . If the grid size of x is denoted by λ , then it is easy to see that δ and λ are inversely related. Since $x_k = x_0 + \lambda k$ and $u_j = j\delta$, we would approximate the

integral in the Fourier transform as

$$\begin{aligned}
 \int_0^{\infty} e^{-ixu} g(u) du &\approx \sum_{j=0}^{N-1} e^{-ix_k u_j} g_j w_j \delta \\
 &= \sum_{j=0}^{N-1} e^{-i(x_0+k\lambda)(j\delta)} g_j w_j \delta \\
 &= \sum_{j=0}^{N-1} e^{-i\lambda\delta jk} (e^{-ix_0 j\delta} g_j w_j \delta)
 \end{aligned} \tag{F.6}$$

where \mathbf{w} is vector of weights associated with the numerical integration technique selected. Comparing the expression in (F.6) with (F.4), we see that, if we set $h_j = e^{-ix_0 j\delta} g_j w_j \delta$, we have that

$$\lambda\delta = \frac{2\pi}{N}. \tag{F.7}$$

Clearly, in order to obtain a small λ i.e. a fine grid across the x 's, δ has to be increased, at the expense of accuracy of the numerical integration. Alternatively, N , the size of the input vector could be increased leading to a waste of computational time. This leads to the usefulness of the fractional Fourier transform.

F.5 The fractional Fourier transform

Chourdakis [104] extended the FFT option pricing methodology of Carr and Madan [103] to utilise the fractional FFT (FRFT). The FRFT computes sums of the form:

$$D_k(\mathbf{h}, \alpha) = \sum_{j=0}^{N-1} e^{-i2\pi k j \alpha} h_j. \tag{F.8}$$

Comparing the sum in (F.8) to the FFT in (F.4), it is clear that the FFT is a special case of the FRFT with $\alpha = 1/N$. Equating the sums in F.6 and F.8 with $h_j = e^{-ix_0 j\delta} g_j w_j \delta$ as above, the following is obtained:

$$2\pi\alpha = \lambda\delta. \tag{F.9}$$

So the x and u spacings, λ and δ can be chosen independently. This is particularly useful in the application in this dissertation: the results from the FRFT are employed in a numerical integration formula. Thus the spacing has to be sufficiently small to ensure accurate survival probabilities. Additionally, the spacing of the size of λ determines the granularity of the cumulative distribution of the integrated CIR process. A trade-off

between the sizes of the parameters λ and δ is therefore undesirable.

F.6 A procedure for implementing the FRFT

Bailey and Swartztrauber [105] derived an algorithm for computing the FRFT. It involves three $2N$ -point FFT procedures (2 FFTs and one inverse FFT). The following two $2N$ -point vectors have to be defined:

$$\mathbf{y} = \left((h_j e^{-i\pi j^2 \alpha})_{j=0}^{N-1}, (0)_{j=0}^{N-1} \right) \quad \mathbf{z} = \left((e^{i\pi j^2 \alpha})_{j=0}^{N-1}, (e^{i\pi(N-j)^2 \alpha})_{j=0}^{N-1} \right).$$

The FRFT is then calculated using:

$$D_k(\mathbf{h}, \alpha) = (e^{-i\pi k^2 \alpha})_{k=0}^{N-1} \odot F_k^{-1}(\mathbf{w}), \quad 0 \leq k < N. \quad (\text{F.10})$$

In (F.10), \mathbf{w} is the $2N$ -point sequence defined by $w_k = F_k(\mathbf{y})F_k(\mathbf{z})$; the symbol \odot denotes element by element multiplication; F_k represents an ordinary FFT procedure and F_k^{-1} is an inverse FFT. Note that the inverse FFT results in a $2N$ long vector. The last N points are discarded. At first glance, it may seem counterintuitive to be creating the $2N$ point vectors \mathbf{y} and \mathbf{z} and then discarding the last N points of the result. The reason for this is explained in [105] and is based on the requirement to satisfy the property for a circular convolution. The reader may also be wondering from whence the j^2 terms originate. The solution lies in the fact that $2jk$ can be written as $2jk = j^2 + k^2 - (k-j)^2$. More information is contained in [105].

Appendix G

An Explanation of the Unexpected Non Zero Receiver CVA

In the examination of the unilateral CVA in Sections 9.2.2 and 9.2.3, we observed a slight positive receiver (short protection position) CVA for extremely high positive default time correlations. The reference entity was the intermediate risk name in the former section and the low risk entity in the latter section, while the counterparty was the high risk name in both sections. The results under discussion are in Table 9.7 and Figure 9.2 in Section 9.2.2 and in Table 9.8 and Figure 9.3 in Section 9.2.3. The situations in which this occurs are extreme and the actual CVA values reported are small. However, a thorough analysis of the model renders further investigation into these few counterintuitive results necessary.

G.1 Default Time Recap

Recall firstly the discussion on default time correlation in Section 7.2, with particular emphasis on Figure 7.8. In this figure, the ratio of τ_1 to τ_2 was plotted on two separate graphs, one in which the counterparty-reference entity default time correlation was 0 and one in which it was 0.99. The reference entity in the illustration was the low risk name, while the counterparty was the high risk entity. The following observations were made in Section 7.2 and are pertinent to the present discussion:

- (i) The ratio of the default time of the low risk entity to that of the high risk entity is generally greater than one. The exact expected value depends on the par CDS spreads of each name.

- (ii) When the correlation between the names is zero, the ratio tends to be fairly random, with a large amount of volatility exhibited (see Figure 7.8(a)). On the other hand, there is much less variance in the ratio when the correlation is extremely high (see Figure 7.8(b)).
- (iii) There are more paths along which both entities default when the correlation is extremely high than when it is zero. Thus, the likelihood of a default of the low risk name is significantly higher when the high risk entity has defaulted and the correlation between the default times is close to unity. The timing of the default is, however, less variable.

These statements provide the key to understanding the non zero receiver CVA that is under investigation.

G.2 Survival Probabilities Illustrated

It has been pointed out elsewhere in the dissertation that, in pricing a CDS, the survival probabilities of the reference entity are the key inputs. Once these are known, the actual valuation of the instrument is straightforward. When these survival probabilities increase, the value of a short protection position increases, while the holder of said protection loses money on a mark to market basis. The converse applies for a reduction in survival probabilities.

The fact that the receiver CVA is not zero for extremely high default time correlations, in the examples in Sections 9.2.2 and 9.2.3, implies that the survival probabilities of the reference entity at the counterparty default times, when the CDS value (replacement cost) is computed, are higher than expected in certain cases. The payer CVA at these points is still increasing. There must therefore be a significant number of very low survival curves, as would be expected¹.

With the low risk name playing the role of the reference entity and the high risk entity playing the role of the counterparty, we examine the results of the intermediate computations made by the model. It appears that the scenarios contributing to the non zero receiver CVA under discussion, are those in which the counterparty default occurs close to the maturity date of the underlying CDS.

¹Scenarios that result in a negative (positive) CDS mark to market at the counterparty default time contribute to the payer (receiver) CVA if a receiver CDS is being priced.

As such, for a given simulation path, we compute the survival probability curve of the reference entity on default of the counterparty, assuming that the counterparty defaults at various different times along the path. The correct inputs to the survival probability computation that differ with default time, $T - \tau_2$, $y_1(\tau_2)$, \bar{U}_1 and U_2 , are utilised in the calculation of the reference entity survivals at each of these default times. The term structure of survivals, that is used to calculate the replacement cost of the CDS at the counterparty default time, is illustrated in Figure G.1 for a variety of assumed default times along the same simulation path and a default time correlation parameter of 0.99. Note that this is only one possible simulation path. Each curve is computed assuming that we are standing at the default time, as explained in Chapter 8.




Figure G.1: Reference entity survival probabilities for various default times on the same simulation path. For each point, T_k , on the x-axis, the corresponding value on the y-axis is $\mathbb{Q}(\tau_1 > T_k | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2)$ i.e. the probability of the reference entity surviving to T_k , given the value of τ_2 and the market information available at that time.

With reference to Figure G.1, it should be noted that, although the initial survival probabilities for each default time appear to be unity, they are, in fact, all less than 1 and are all decreasing in T_k . Now recall that the reference entity is the low risk entity and

the counterparty is the high risk entity in this example. The reason for the increase in receiver CDS values when the default time is close to CDS maturity is clear from the graph. Since the correlation between default times is so high, the reference entity survival probability curve is effectively pricing in a default (the survival probability curve decreases significantly in value over time), but not immediately.

This is in line with the statements above regarding default time correlation. That is, the extremely high correlation renders a reference entity default in the future more likely. However, the likelihood of it occurring in the immediate future is extremely small given the differential in riskiness between the two entities. Recall that hazard rates are not correlated.

In Figure G.2 below, we illustrate the reference entity survival curve for various correlations, assuming a default time of 0.75. Notice how the 99% correlation curve initially displays larger values than the other curves. However, in the long term, its values are clearly the lowest, indicating a much larger probability of a reference entity default, given that the counterparty has defaulted. The effect that was illustrated in Figure G.1 for the 99% case is apparent here. For lower correlations, it is not as pronounced, due to the greater volatility in the ratio of default times for these correlations.

G.3 Viewing the Distributions Responsible

In Chapter 8, the computation of the survival probabilities of the reference entity, at the counterparty default time, was explained thoroughly. The equation used to approximate these, for each payment time T_k that occurs after τ_2 , was given in Equation (8.16) as:

$$\mathbb{Q}(\tau_1 > T_k | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2) \approx \sum_j \frac{p_{j+1} + p_j}{2} \Delta f_j, \quad (\text{G.1})$$

where

$$p_j = \mathbb{Q}(X < x_j | \mathcal{G}_{\tau_2}) \text{ for a grid } x_j; \quad X := Y_1(T_k) - Y_1(\tau_2)$$

and

$$f_j := C_{1|2}(u_j, U_2).$$

In Figure G.3(a), we plot the two components, $\frac{p_{j+1} + p_j}{2}$ and Δf_j , of the sum in (G.1) for an example in which the survival probability, computed using the plotted values, is highest


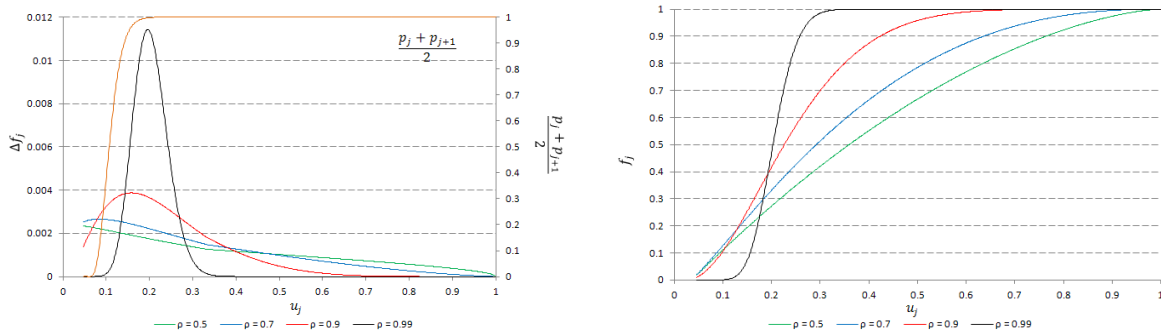


Figure G.2: Reference Entity Survival Curve at Counterparty Default Time, 0.75 Years, for Various Default Time Correlations. For each point, T_k , on the x-axis, the corresponding value on the y-axis is $\mathbb{Q}(\tau_1 > T_k | \mathcal{G}_{\tau_2}, \tau_1 > \tau_2)$ i.e. the probability of the reference entity surviving to T_k , given the value of τ_2 and the market information available at that time³.

when the default time correlation is 99%. Note that the same p_j values are used for all correlations since this variable is not dependent on correlation. The actual values of f_j used to compute the Δf_j 's are then illustrated in Figure G.3(b).



(a) $\frac{p_{j+1} + p_j}{2}$ (primary axis) and Δf_j (secondary axis)

(b) f_j

Figure G.3: The two components of the survival probability approximation in (G.1). The example shown is one for which the 99% correlation survival probability is higher than for the other correlations.

With reference to Figure G.3, we notice that the higher than expected survival prob-

ability in the 99% correlation case is as a result of the distribution, Δf_j , peaking only once the cumulative distribution function of the integrated CIR process (the p_j 's) is close to unity. The graph of f_j has a steep gradient in the 99% correlation case. This is a result of the fact that the standard deviation of the conditional copula function is inversely proportional to the absolute size of the correlation between default times. It accounts for the large peak and tiny tails of the Δf_j curve in the 99% case when compared to the others.

This leads us to conclude that the choice of the Gaussian copula for modelling default time correlation plays a role in the unexpected non zero receiver CVA under discussion. It would be pertinent to investigate alternatives, starting with the t-copula, which exhibits fatter tails than the Gaussian copula, and is still very tractable.

G.4 Conclusion

The positive receiver CVA, when the reference entity is much less risky than the counterparty and default time correlations are extremely large and positive, is clearly a result of fact that it is default time correlations (more precisely correlation between the uniform random variables of the exponential triggers of the default times) the we are considering as opposed to correlating hazard rates. The reasons for this choice were discussed in Chapters 5 and 6. This could be seen as an attempt to capture economic reality. However, it fails to capture market sentiment, which would almost certainly result in a reduction in survival probabilities of the reference entity if it was perceived to be highly correlated, in any way, with the defaulted counterparty.

This might be seen as a flaw in the model. One way of overcoming it would be to include a jump parameter, dependent on the counterparty default time and the default time correlation. On the other hand, one might argue that the wrong way risk desired in the model is captured by virtue of the fact that the payer CVA increases despite the receiver CVA increasing.

In addition, it is important to keep the results in perspective. The situations in which this non zero receiver CVA occurs are rather extreme and are included for the purposes of stress testing the model. A situation in which it was believed that two entities had a 99% default time correlation and, yet, exhibited such a differential in spreads, is improbable. From the perspective of the low risk entity, it effectively amounts to purchasing protec-

tion on himself, which is cause for an increase in receiver CVA from the high risk entity's perspective.

Appendix H

Matlab Code

The unilateral and bilateral CVA models were both implemented Matlab. The code was written to take advantage of the fact that Matlab is most efficient in terms of speed when matrix operations are used in favour of loops. Due to the magnitude of the computational requirements, not all loops were avoidable. Since there is a large amount of Matlab code, a selection is shown below.

H.1 CIR ++ Simulation

```
function Results=Simulating_CIR_pp(Times_for_Sim,k,mu,v,NoSim,Exact_1_Numerical_2,y_0,intensity_interp,
    t_axis_intensity,y_axis_intensity)

CIR_Values=Simulating_CIR(Times_for_Sim,k,mu,v,NoSim,Exact_1_Numerical_2,y_0);

phi=Calculating_Deterministic_Shift_CIR_pp([0 Times_for_Sim],k,mu,v,t_axis_intensity,y_axis_intensity,
    intensity_interp,y_0);

CIR_pp_Values=CIR_Values+repmat(phi,NoSim,1);

Results=CIR_pp_Values;

end
```

H.1.1 Simulating CIR Process

```
function Results=Simulating_CIR(time,k,mu,v,NoSim,Exact_1_Numerical_2,y_0)

%The parameters are for a CIR++ Process where  $\lambda=y+\phi$  and the dynamics of  $y$  are  $dy=k(\mu-y)dt+v\sqrt{y}dW$ 
%T is maturity in years and t_step is time between steps, also in years
%The parameter Exact_1_Numerical_2 should be set to 1 if you wish to simulate by sampling from the
    transition density and 2 if you wish to use a discretised scheme
tStart123=tic;
```

```

%y_0 is the initial value

if Exact_1_Numerical_2==2 && 2*k*mu<=v^2
    button=questdlg('The positivity constraint is not satisfied. It is recommended that you rather use exact
    simulation than numerical. Do you want to continue using numerical? Press YES to do so or press NO to change
    to exact.','Continue using numerical simulation?');
    if strcmp(button,'No')
        Exact_1_Numerical_2=1;
    end
end

Sim_Values=zeros(NoSim,size(time,2)+1);

Sim_Values(:,1)=y_0*ones(NoSim,1);

delta_t=[time(1,1) diff(time,[],2)];

if Exact_1_Numerical_2==1          %Exact simulation from pg 124, Glasserman book

    d=4*mu*k/(v^2);

    if d>1
        c=(v^2*(ones(size(delta_t))-exp(-k*delta_t))/(4*k));
        z=randn(NoSim,size(time,2));
        X=random('chi2',d-1,NoSim,size(time,2));
        for i=1:(size(time,2))
            lambda=(Sim_Values(:,i)*exp(-k*delta_t(1,i)))/c(1,i);
            Sim_Values(:,i+1)=c(1,i).*(z(:,i)+sqrt(lambda)).^2+X(:,i));
        end
    end

    if d<=1 %Checked mean and variance of results at t_step and T for 100,000
        sims according to formulae on pg 66 of Brigo and Mercurio book
        c=(v^2*(ones(size(delta_t))-exp(-k*delta_t))/(4*k));
        for i=1:(size(time,2))
            lambda=(Sim_Values(:,i)*exp(-k*delta_t(1,i)))/c(1,i);
            N=random('poiss',lambda/2);
            X=random('chi2',d*ones(NoSim,1)+2*N);
            Sim_Values(:,i+1)=c(1,i).*X;
        end
    end

end

if Exact_1_Numerical_2==2
    NormRV=randn(NoSim,size(time,2));
    if 2*k*mu>v^2
        %This is the Euler Implicit Positivity-Preserving Scheme
        %Taken from 'Credit Default Swaps Calibration and Option Pricing
        with the SSRD Stochastic Intensity and IR Model', Brigo and Alfonsi, 2005, pg 13

        denom=2*(ones(size(delta_t))+k*delta_t);          %Tested mean and std dev of this same as for exact case.
        for i=1:(size(time,2))
            Sim_Values(:,i+1)=((v*sqrt(delta_t(1,i))*NormRV(:,i)+sqrt(v^2*(delta_t(1,i))*NormRV(:,i).^2
            +4*(Sim_Values(:,i)+(k*mu-0.5*v^2)*delta_t(1,i)).*(1+k*delta_t(1,i))*ones(NoSim,1))))/denom(1,i)).^2;
        end
    end
end

```

```

    end
  else
    for i=1:(size(time,2))

      %Sim_Values(:,i+1)=Sim_Values(:,i)+k*(mu*ones(NoSim,1)-max(Sim_Values(:,i),0))*t_step+v*NormRV(:,i)
      .*sqrt(max(Sim_Values(:,i),0))*sqrt(t_step);
      Sim_Values(:,i+1)=abs(Sim_Values(:,i)+k*(mu*ones(NoSim,1)-Sim_Values(:,i))*delta_t(1,i)+v*NormRV(:,i)
      .*sqrt(Sim_Values(:,i))*sqrt(delta_t(1,i)));
    end
  end

end

end

Results=Sim_Values;

toc(tStart123)

end

```

H.1.2 Deterministic Shift Computation

```
function Results=Calculating_Deterministic_Shift_CIR_pp(time,k,mu,v,t_axis_intensity,y_axis_intensity
,intensity_interp,y_0)
```

%See page 102 of Brigo and Mercurio book for info on the CIR++ model. Also see the paper 'On deterministic-shift extensions of short-rate models' by Brigo and Mercurio, 2001.

```

f_CIR=Forward_Rate_CIR_Process(time,k,mu,v,y_0);

if strcmp(intensity_interp,'linear')
  f_MKT_t0=interp1(t_axis_intensity, y_axis_intensity, time(1,:), intensity_interp);
else
  f_MKT_t0=zeros(1,size(time,2));
  for i=1:size(time,2)
    Indicator_for_Before=(t_axis_intensity>=time(1,i));
    A=cumsum(Indicator_for_Before);
    A=sum(((A==1)).*sort(randperm(size(t_axis_intensity,2))));
    if A==0
      f_MKT_t0(1,i)=y_axis_intensity(:,size(y_axis_intensity(:,2)));
    else
      f_MKT_t0(1,i)=y_axis_intensity(:,A);
    end
  end
end

end

phi=f_MKT_t0-f_CIR;

Results=phi;

end

```

H.1.3 Simulating Default Times - Unilateral CVA

```

function [Result]=Simulating_Default_Times(correlation, Times_for_Integral,NoSim,Maturity,
    Integral_Intensities1,Integral_Intensities2) %Note: we use a Gaussian copula here.

%NB: This outputs the default time of the counterparty if it is a)before the default time of the reference entity
and b)before maturity of the CDS.
%For info on simulating default time see page 778 of Brigo and Mercurio book.
%Also see algorithm on pg 9 of the paper 'Modelling dependence for credit derivatives with copulas' by Jouanin,
Rapuch, Riboulet and Roncalli.

% U=rand([NoSim,2]);

U=copularnd('Gaussian',[1 correlation;correlation,1],NoSim); %The Gaussian copula is calculated by Matlab using the
line: u = normcdf(mvnrnd(zeros(1,d),Rho,n)). This follows from C(u1,u2)=normsdist(normsinv(u1),normsinv(u2)) since
normsdist(x)=u and thus normsinv(u)=x.

E=-log(ones(NoSim,2)-U);

ExprV1=repmat(E(:,1),1,size(Integral_Intensities1,2));
ExprV2=repmat(E(:,2),1,size(Integral_Intensities2,2));

Time1=Integral_Intensities1>=ExprV1;
Time2=Integral_Intensities2>=ExprV2;

Cum_Sum_Default1=cumsum(cumsum(Time1,2),2);
Cum_Sum_Default2=cumsum(cumsum(Time2,2),2);

Def_Time_1=Cum_Sum_Default1==1;
Def_Time_2=Cum_Sum_Default2==1;

% time=t_step*ones(NoSim,T/t_step);
% time=cumsum(time,2);
time=repmat(Times_for_Integral(:,2:size(Times_for_Integral,2)),size(Def_Time_1,1),1);

Tau1=sum(Def_Time_1.*time,2);
Tau2=sum(Def_Time_2.*time,2);

Default_Times=[Tau1 Tau2];

Ref_Entity=(Default_Times(:,1)==0)*1000000;
Ref_Entity_Default_Times=Ref_Entity+Default_Times(:,1);

Default_CP_before_ref=Default_Times(:,2)<Ref_Entity_Default_Times;
Default_CP_before_Maturity=Default_Times(:,2)<=Maturity;
Result=Default_Times(:,2).*Default_CP_before_ref.*Default_CP_before_Maturity; %So this vector houses the CP
default time if this is during the life of the trade and it is before the ref entity default.

%Below is the code for implementing the copula simulation manually. It can be tested that
the correlation between the uniform RV's generated in this way and the uniform RV's generated
from the copula function in the first line is the same.

% corrMat=[1 correlation;correlation 1];
% R=chol(corrMat,'lower');
```

```

% X=randn(2,NoSim);
% X=R*X;
% g=corrcoef(X(1,:),X(2,:))
%
% Unif=normcdf(X);
% h=corrcoef(Unif(1,:),Unif(2,:))

end

```

H.2 FRFT

```

function [Result]=FRFT_Step5a_without_Delta(k,mu,v,Spot_Intensities_Tau2,Integration_Rule,NoDefaults,Time_to_Ts,x1_p0,
      N,Delta) %This implements is the Fractional Fourier Transform to calculate the probabilities
      in Step 5a.

%This bit is to ensure that we don't waste time on rows that don't matter.
if x1_p0==0
    OriginalNoDefaults=size(Time_to_Ts,1);
    Spot_Intensities_Tau2=Remove_zero_rows((Spot_Intensities_Tau2+1000*ones(size(Spot_Intensities_Tau2)))
      .*(Time_to_Ts>0));
    Spot_Intensities_Tau2=Spot_Intensities_Tau2-1000*ones(size(Spot_Intensities_Tau2));
    Numbers=(Time_to_Ts>0).*transpose(sort(randperm(NoDefaults)));
    Numbers=Remove_zero_rows(Numbers);
    Delta=Remove_zero_rows(Delta.*repmat((Time_to_Ts>0),1,size(Delta,2)));
    Time_to_Ts=Remove_zero_rows(Time_to_Ts);
    NoDefaults=size(Time_to_Ts,1);
end
Expected_Values=Integrated_CIR_Expected_Value(k*ones(size(Time_to_Ts)),mu*ones(size(Time_to_Ts)),Time_to_Ts
  ,Spot_Intensities_Tau2);
Std_Dev=sqrt(Integrated_CIR_Variance(k,mu,Time_to_Ts,Spot_Intensities_Tau2,v));

Lambda_Multiple=11;
Lambda=(Lambda_Multiple*Std_Dev+Expected_Values)/N;
Lambda_Matrix=repmat(Lambda,1,N);

clear Lambda Std_Dev Expected_Values

x_Matrix=(repmat((sort(randperm(N))-ones(1,N)),NoDefaults,1)).*Lambda_Matrix;

if x1_p0==1
    Result=x_Matrix;
    return
end

Time_to_Ts_Matrix=repmat(Time_to_Ts,1,N);

clear Time_to_Ts

Spot_Intensities_Matrix=repmat(Spot_Intensities_Tau2,1,N);

clear Spot_Intensities_Tau2

```



```

if strcmp(Integration_Rule,'Trapezoidal')
    w=ones(1,N);
    w(1,1)=0.5;
    w(1,N)=0.5;
end

w_Matrix= repmat(w,NoDefaults,1);

u_Matrix= repmat(Delta,1,N).*repmat((sort(randperm(N))-ones(1,N)),NoDefaults,1);
u_Matrix(:,1)=u_Matrix(:,2); %Temp because of error in denomin at zero.

t_Start=tic;

% function_Matrix=(w_Matrix.*real(Charac_Func_CIR_New_Formulation(k,mu,v,Spot_Intensities_Matrix,u_Matrix
    ,Time_to_Ts_Matrix)))/(u_Matrix*2i);
function_Matrix=(w_Matrix.*real(Charac_Func_CIR_New_Formulation(k,mu,v,Spot_Intensities_Matrix,u_Matrix
    ,Time_to_Ts_Matrix)))/(u_Matrix*2i);

End_function=toc(t_Start)

clear w_Matrix Spot_Intensities_Matrix Time_to_Ts_Matrix u_Matrix

%Now implement FRFT

Delta_Matrix=repmat(Delta,1,N);

clear Delta

t_Start1=tic;

j_Matrix=repmat((sort(randperm(N))-ones(1,N)),NoDefaults,1);
Sum1=Fractional_Fast_Fourier_Transform(function_Matrix,j_Matrix,-1,Lambda_Matrix,Delta_Matrix);
Sum2=Fractional_Fast_Fourier_Transform(function_Matrix,j_Matrix,1,Lambda_Matrix,Delta_Matrix);

clear j_Matrix Lambda_Matrix function_Matrix

Result=real(2/pi*Delta_Matrix.*((Sum1-Sum2)+w(1,1)*x_Matrix)); %Take real here
    because some results have an imaginary part due to machine error in the region of E-18.
End_FRFT=toc(t_Start1)

% if NoDefaults<OriginalNoDefaults

    clear Sum1 Sum2 x_Matrix Delta_Matrix w

    Result1=zeros(OriginalNoDefaults,N);
    for i=1:NoDefaults
        Result1(Numbers(i,1),:)=Result(i,:);
    end
    Result=Result1;

% end

end

```

H.2.1 Variance of Integrated CIR Process

```
function [Result]=Integrated_CIR_Variance(k,mu,t,y_0,v) %See pg 16 of 'The Integrated Square-Root Process' by
                                                    Daniel Dufresne

Alpha=-k;
Beta=-Alpha*mu;

EY=Integrated_CIR_Expected_Value(k,mu,t,y_0);
EYSq=y_0.^2/Alpha^2+2*y_0*Beta/Alpha^3+ones(size(y_0))*Beta^2/Alpha^4-y_0*v^2/Alpha^3-ones(size(y_0))*5*Beta*v^2
    ./((2*Alpha^4)+t.*(2*y_0*Beta/Alpha^2+2*Beta^2/Alpha^3*ones(size(t))-Beta*v^2/Alpha^3*ones(size(t)))
    +t.^2*Beta^2/Alpha^2+exp(t*Alpha).*(-2*y_0.^2/Alpha^2-4*y_0*Beta/Alpha^3-2*Beta^2/Alpha^4*ones(size(t))
    +2*Beta*v^2/Alpha^4*ones(size(t))+t.*(-2*y_0*Beta/Alpha^2-2*Beta^2/Alpha^3*ones(size(t))-2*y_0*v^2/Alpha^2
    -2*Beta*v^2/Alpha^3*ones(size(t))))+exp(2*t*Alpha).*(y_0.^2/Alpha^2+2*y_0*Beta/Alpha^3+Beta^2/Alpha^4*ones(size(t))
    +y_0*v^2/Alpha^3+Beta*v^2/(2*Alpha^4)*ones(size(t)));
Result=(EYSq-EY.^2).*(t>0);
end
```

H.3 The function f_j in Unilateral Model

```
function [Result]=Step5c_Conditional_Distributions(u, U2, U1bar,correlation)

% [row_U2_U1_bar ~]=size(U2_and_U1bar); %Both of these depend on the simulation number
%
% if row_U2_U1_bar>2 %U2 and U1bar are column vectors
% U2=U2_and_U1bar(:,1);
% U1bar=U2_and_U1bar(:,2);
% else
% U2_and_U1bar=transpose(U2_and_U1bar);
% U2=U2_and_U1bar(1,:);
% U1bar=U2_and_U1bar(2,:);
% end

[NoSim N]=size(u);

Mean_Conditional_Dist=correlation*norminv(U2);
Mean_Conditional_Dist(isinf(-Mean_Conditional_Dist))=0;

Var_Conditional_Dist=(ones(NoSim,1)-correlation^2*ones(NoSim,1)).*(U2>0);

Mean_Conditional_Dist= repmat(Mean_Conditional_Dist,1,N);
Var_Conditional_Dist= repmat(Var_Conditional_Dist,1,N);

U1bar= repmat(U1bar,1,N);

First_Term_Top=normcdf(norminv(u),Mean_Conditional_Dist,sqrt(Var_Conditional_Dist));
Second_Term_Top=normcdf(norminv(U1bar),Mean_Conditional_Dist,sqrt(Var_Conditional_Dist));

f=(First_Term_Top-Second_Term_Top)./(ones(NoSim,N)-Second_Term_Top);

Result=f;
```

end

H.3.1 Simulating Default Times - Bilateral CVA

```
function [Result]=Simulating_Default_Times(correlation0_1,correlation0_2,correlation1_2,
    Times_for_Integral,NoSim,Maturity, Integral_Intensities1,Integral_Intensities2, Integral_Intensities0)
    %Note: we use a Gaussian copula here. This can be extended later. The copula thus requires only one input
    which is the correlation%Note: we use a Gaussian copula here. This can be extended later. The copula thus
    requires only one input which is the correlation

    %NB: This outputs the default time of the counterparty if it is a)before the default time of the reference entity and
    b)before maturity of the CDS.
    %For info on simulating default time see page 778 of Brigo and Mercurio book.
    %Also see algorithm on pg 9 of the paper 'Modelling dependence for credit derivatives with copulas'
    by Jouanin, Rapuch, Riboulet and Roncalli.

    %    U=rand([NoSim,2]);

    U=copularnd('Gaussian',[1 correlation0_1 correlation0_2;correlation0_1 1 correlation1_2; correlation0_2
        correlation1_2 1],NoSim);    %The Gaussian copula is calculated by Matlab using the line: u =
        normcdf(mvnrnd(zeros(1,d),Rho,n)). This follows from C(u1,u2)=normsdist(normsinv(u1),normsinv(u2))
        since normsdist(x)=u and thus normsinv(u)=x.

    E=-log(ones(size(U))-U);

    ExpRV0= repmat(E(:,1),1,size(Integral_Intensities0,2));
    ExpRV1= repmat(E(:,2),1,size(Integral_Intensities1,2));
    ExpRV2= repmat(E(:,3),1,size(Integral_Intensities2,2));

    Time0=Integral_Intensities0>=ExpRV0;
    Time1=Integral_Intensities1>=ExpRV1;
    Time2=Integral_Intensities2>=ExpRV2;

    Cum_Sum_Default0=cumsum(cumsum(Time0,2),2);
    Cum_Sum_Default1=cumsum(cumsum(Time1,2),2);
    Cum_Sum_Default2=cumsum(cumsum(Time2,2),2);

    Def_Time_0=Cum_Sum_Default0==1;
    Def_Time_1=Cum_Sum_Default1==1;
    Def_Time_2=Cum_Sum_Default2==1;

    time=repmat(Times_for_Integral(:,2:size(Times_for_Integral,2)),size(Def_Time_1,1),1);

    Tau0=sum(Def_Time_0.*time,2);
    Tau1=sum(Def_Time_1.*time,2);
    Tau2=sum(Def_Time_2.*time,2);

    Default_Times=[Tau0 Tau1 Tau2];

    Ref_Entity=(Default_Times(:,2)==0)*1000000;
    Ref_Entity_Default_Times=Ref_Entity+Default_Times(:,2);
```

```

Inv_Temp=(Default_Times(:,1)==0)*1000000;
Inv_Temp=Inv_Temp+Default_Times(:,1);

Default_CP_before_ref=Default_Times(:,3)<Ref_Entity_Default_Times;
Default_CP_before_inv=Default_Times(:,3)<Inv_Temp;
Default_CP_before_Maturity=Default_Times(:,3)<=Maturity;
Relevant_CP=Default_Times(:,3).*Default_CP_before_ref.*Default_CP_before_Maturity.*Default_CP_before_inv; %So this
    vector houses the CP default time if this is during the life of the trade and it is before the ref entity default.
    Where the CP doesn't default or the default of the CP is after the ref entity the entry in the vector is 0.

clear Inv_Temp;

CP_Temp=(Default_Times(:,3)==0)*1000000;
CP_Temp=CP_Temp+Default_Times(:,3);
Default_Inv_before_ref=Default_Times(:,1)<Ref_Entity_Default_Times;
Default_Inv_before_cp=Default_Times(:,1)<CP_Temp;
Default_Inv_before_Maturity=Default_Times(:,1)<=Maturity;
Relevant_Inv=Default_Times(:,1).*Default_Inv_before_ref.*Default_Inv_before_Maturity.*Default_Inv_before_cp;

clear CP_Temp;

Indicator_CP=Relevant_CP>0;
Indicator_Inv=Relevant_Inv>0;

Tau=Relevant_CP+Relevant_Inv;

Result=[Tau Indicator_Inv Indicator_CP];

end
    
```

H.3.2 Conditional Copula Function f at τ_2 in Bilateral Model

```

function [Result]=Step5c_Conditional_Distributions_trivariate_try_again(term_d,u, U2, Ubar_0_2,Ubar_1_2,
    correlation0_1, correlation0_2, correlation1_2)

Sigma1_1=[1 correlation0_1;correlation0_1 1];
Sigma1_2=[correlation0_2; correlation1_2];
Sigma2_1=[correlation0_2 correlation1_2];
%Sigma2_2=1;

Cond_Var_Tri=Sigma1_1-Sigma1_2*Sigma2_1;
mean_tri_u1_bit=correlation1_2*norminv(U2);
mean_tri_u0_bit=correlation0_2*norminv(U2);

normalised_u_for_tri=norminv(u)-repmat(mean_tri_u1_bit,1,size(u,2));
normalised_ubar_0_2_for_tri=norminv(Ubar_0_2)-mean_tri_u0_bit;

vector=[repmat(normalised_ubar_0_2_for_tri,size(u,2),1) reshape(normalised_u_for_tri,size(u,1)*size(u,2),1)];

%opts = statset('tolfun',1e-4);
    
```

```

% rho=Cond_Var_Tri(1,2)/(sqrt(Cond_Var_Tri(1,1))*sqrt(Cond_Var_Tri(2,2)));

tic

%rho=Cond_Var_Tri(1,2)/(sqrt(Cond_Var_Tri(1,1))*sqrt(Cond_Var_Tri(2,2)));
%term_b=mvncdf(vector,0,(Cond_Var_Tri));

rho=Cond_Var_Tri(1,2)/(sqrt(Cond_Var_Tri(1,1))*sqrt(Cond_Var_Tri(2,2)));

if rho==0

    term_b=normcdf(vector(:,1)/sqrt(Cond_Var_Tri(1,1))).*normcdf(vector(:,2)/sqrt(Cond_Var_Tri(2,2)));

else

    %test1=mvncdf(vector,0,(Cond_Var_Tri));

    b1=vector(:,1)/sqrt(Cond_Var_Tri(1,1));
    b2=vector(:,2)/sqrt(Cond_Var_Tri(2,2));

    toc

    if rho > 0
        p1=normcdf(min([b1 b2],[],2));
    %    p1(any(isnan(b),2)) = NaN;
    else
        p1 = normcdf(b1)-normcdf(-b2);
        p1(p1<0) = 0; % max would drop NaNs
    end

    %clear vector

    no_steps=500;
    highlimit=sign(rho).*pi./2;
    lowlimit=asin(rho);

    if highlimit < lowlimit
        lowlimit=sign(rho).*pi./2;
        highlimit=asin(rho);
    end

    step_size=(highlimit-lowlimit)/no_steps;

    theta=[lowlimit sort(randperm(no_steps))*step_size+ones(1,no_steps)*lowlimit];

    sintheta = sin(theta);
    sintheta = repmat(sintheta,size(b1,1),1);

    cossqtheta = cos(theta).^2;
    cossqtheta = repmat(cossqtheta,size(b1,1),1);

    % clear theta
    b1=repmat(b1,1,size(sintheta,2));
  
```

```

b2= repmat(b2,1,size(sintheta,2));

integrand =2*( exp(-((b1.*sintheta - b2).^2 ./ cossqtheta + b1.^2)/2));

% clear b1 b2 sintheta cossqtheta

integrand(:,1)=0.5*integrand(:,1);
integrand(:,no_steps+1)=0.5*integrand(:,no_steps+1);

term_b=sum(integrand,2)*step_size/2;

if (sign(rho).*pi/2) < asin(rho)
    term_b=-term_b;
end

term_b=pi-term_b/(2*pi);

% clear integrand p1

end

toc

term_b=reshape(term_b,size(u,1),size(u,2));

term_a=normcdf(norminv(u),repmat(correlation1_2*norminv(U2),1,size(u,2))
    ,repmat(sqrt(1-correlation1_2^2),size(u,1),size(u,2)));

term_c=normcdf(norminv(Ubar_1_2),correlation1_2*norminv(U2),repmat(sqrt(1-correlation1_2^2),size(U2,1),1));
term_c=repmat(term_c,1,size(u,2));

term_e=normcdf(norminv(Ubar_0_2),correlation0_2*norminv(U2),repmat(sqrt(1-correlation0_2^2),size(U2,1),1));
term_e=repmat(term_e,1,size(u,2));

% normalised_ubar_1_2_for_tri=norminv(Ubar_1_2)-mean_tri_u1_bit;
% term_d=repmat(mvncdf([normalised_ubar_0_2_for_tri normalised_ubar_1_2_for_tri],0,(Cond_Var_Tri)),1,size(u,2));

% term_f=term_c;
% term_g=term_d;

Result=(term_a-term_b-term_c+term_d)./(ones(size(term_e))-term_e-term_c+term_d);

end

```