The Valuation and Hedging of Default-Contingent Claims in Multiple Currencies

Gavin Kenneth Truter

Programme in Advanced Mathematics of Finance,
University of the Witwatersrand,
Johannesburg.

A dissertation submitted to the Faculty of Science in fulfilment of the requirements for the degree of Master of Science.

19 April 2012

Declaration

I declare that this dissertation is my own, unaided work. It is being submitted for the degree of Master of Science in the University of the Witwatersrand, Johannesburg. It has not been submitted before for any degree or examination in any other University.

____

19 April 2012

Abstract

This dissertation examines the pricing of the same credit risk in two currencies, and hence the valuation of credit-contingent foreign exchange products. Such pricing hinges upon the dependence of the credit risk and the foreign exchange rate. We recall the reduced-form model proposed by Ehlers (2007), which allows credit-currency dependence through correlation between the Brownian motions driving the default intensity and the exchange rate, and through a jump in the exchange rate at the default time. Four basic specifications of this model are considered. Two of these specifications have not previously appeared in the literature and one of these, based on a lognormal process for the default intensity, proves to be especially useful and tractable. The problem of hedging defaultable claims in one currency with similar claims in another is briefly considered, and it is shown that hedging against the default event and against credit spread movements are not in general equivalent.

Acknowledgements

Standard Bank provided me with generous financial support while I was studying towards an M.Sc., though of course this dissertation does not reflect the bank's views.

Most of all, I would like to thank my family, John, Gill and Michael, for their encouragement and assistance, and James Tayor for his kind support and supervision. I would also like to thank David Taylor for his role in supervising me, everyone who answered my many questions (Tom, Diane, Craig, Nicolette and Monique Jeanblanc) and those who helped with proofreading (Evashun and Michael). Lastly I would like to thank my friends for their encouragement, and especially Warrick for many discussions. I am very grateful to all of you.

Contents

List of Figures

Chapter 1

Introduction and Literature Review

1.1 Introduction

This dissertation investigates how the values of default-contingent claims depend upon the currency in which the claims are denominated.

The following stylised example illustrates why this dependence should exist. Consider a security that promises a payment of a million US dollars to its holder if and when a particular large South African parastatal defaults, should that occur in the next year. Now consider a similar instrument that pays instead eight million South African rand. We suppose that at the moment one dollar costs eight rand exactly, so the two contracts pay the same amount converted at today's exchange rate. However, a default by the parastatal would likely be related to (cause or be caused by) a deterioration in the South African economic situation, which would itself likely lead to a weakening of the rand. Thus it is probable that in the event of default, a payment of \$1,000,000 will be more valuable than a payment of R8,000,000. So it seems that the dollar-denominated security should be more expensive today than the rand-denominated one.

Of course this is hardly no-arbitrage pricing, but it gives the correct result and illustrates the intuitive basis for the pricing differences that will be quantified in the rest of this dissertation. In particular, we will see that these pricing differences are primarily a result of dependence between the exchange rate and the credit risk.

In this dissertation we concentrate on the pricing of defaultable payments of fixed amounts, which allows us to value defaultable bonds and credit default swaps denominated in each currency, vulnerable foreign exchange forwards and cross currency swaps, and similar linear instruments. We will also briefly consider the pricing of vulnerable currency options. We work with two currencies (which we call domestic and foreign) and suppose that only one market participant may default – we do not consider bilateral credit risk or the risk that a seller of default protection may default on the protection payment. The credit risk is modelled exclusively using a reduced-form approach, and the models for the default intensity that we consider in detail are all single-factor diffusions. Nonetheless, finding accurate and fast procedures to price default-contingent claims in multiple currencies proves to be challenging. The author is unaware of any previous thorough examination of specific, tractable, reduced-form models for a multiple-currency market with default; hopefully this dissertation will partially fill that gap.

The remainder of this chapter consists of three sections. The first discusses the structural and reduced-form approaches to credit risk modelling in order to contextualise our use of reduced-form models. The second gives an overview of the existing literature on the pricing of credit-contingent products in multiple currencies, and the last explains the structure of the rest of the dissertation.

1.2 Single-Currency Credit Risk Modelling

There are two main approaches to modelling credit risk. The first, known as the firm's value or structural approach, originated with Black & Scholes [5] and Merton [37]. In this approach, one models the total value of a firm's assets. Default occurs when this process declines to a barrier representing the value of the firm's debt, so that default occurs when the firm becomes effectively bankrupt. The structural approach is attractive in that it explains the default event and provides a link between the prices of the equity and the debt of a firm. On the other hand, such models assume that the firm's asset value is continuously observable, which is commonly accepted to be untrue. Also, calibrating such models to market data can be difficult, especially for short maturities – see Schönbucher [40], Section 9.6.

The reduced-form or intensity-based approach to credit risk modelling was originally contributed to by, amongst others, Jarrow & Turnbull [27], Lando [32], Madan & Unal [36] and Duffie & Singleton [13]. In contrast to structural models, in reduced-form models default is assumed to be an exogenous random event (though the likelihood of the default occurring in the next instant may be linked to other variables, like the firm's equity value). Reduced-form models make no attempt to explain the default event and so have no economic justification; but this sometimes leaves them free enough to describe the market pricing of credit risk more accurately than structural models.

Jarrow & Protter [26] argue that structural and reduced-form models are closely linked: the same model can appear to be structural or reduced-form depending on the amount of information that we assume is available to the modeller. Structural models correspond to assuming that the modeller has the same information as the firm's management; reduced-form models correspond to assuming that the modeller has only the information available to the market. Seen from this perspective, it seems that for our purposes (the pricing of risky debt and credit derivatives) a reduced-form approach is more useful. We will thus use reduced-form models exclusively.

1.3 Literature

Most dynamic models of multiple-currency markets with default use a structural approach. There are several papers in this vein, but the only published paper known to the author is that of Barnhill & Maxwell [1], who measure, by simulation, the risk of a multi-currency fixed-income portfolio. Kafetzaki-Boulamatsis & Tasche [29] value the equity of a firm with operations in several currencies using an approximation originally conceived to price basket options. Warnes & Acosta [44] price the debt of a firm with domestic operations but debt issued in both domestic and foreign currency. Yigitbasioglu [45] applies a basic model to the pricing of convertible bonds with currency risk. Chan-Lau & Santos [9] price the zero-coupon, foreign currency debt of a firm with several models for the foreign currency value of the firm. Tasche [43] obtains relationships between the probabilities of default and the correlations of the asset values of several debtors whose assets and debt are in multiple currencies.

A separate vein of papers investigates the differences between credit spreads on bonds issued by the same corporation in two different currencies. Jankowitsch & Pichler [25] show that, roughly, if the exchange rate and survival indicator random variables are uncorrelated under the domestic risk-neutral measure then the credit spreads in the two currencies are equal. They use a multi-curve splines model proposed by Houweling et al. [19] to test the hypothesis that the exchange rate and survival are independent and find strong evidence against this hypothesis. They conclude that in estimating credit spread curves, bonds issued in different currencies by the same corporation cannot be pooled without taking into account the effect of currency-default dependence. A similar, unpublished paper is that of Sener & Kenc [41]. In another unpublished paper, Chan-Lau [8] presents a simple rule-of-thumb to convert credit spreads between currencies based on a statistical analysis of the percentage of exchange rate variance that is explained by changes in macroeconomic risk, for which sovereign credit default swap rates are used as a proxy. While these papers are concerned with credit spreads in different currencies, they focus on the conversion of credit spread curves rather than on no-arbitrage pricing in a dynamic model.

For our purposes, by far the most significant work on the subject is contained in the doctoral thesis of Philippe Ehlers [14] (another part of that thesis examines a framework for modelling portfolio credit risk). An early version of the material on credit risk in multiple currencies is also available under the names of Ehlers and his supervisor Philipp Schönbucher [15].[1]

This work is the most complete and sophisticated existing account of the reduced-form pricing of credit risk in a multi-currency market. It consists of three parts, the first two of which we cover in some detail. The first part is a general model that we present in Chapter 3 and use throughout. The model allows for dependence between the credit risk and the exchange rate both through correlation of the default intensity and exchange rate processes and through a jump in the exchange rate at the default time. The second part is an affine version of this general model that allows for tractable pricing even with multi-factor models for the default intensity and interest rates. This version is recapitulated in Chapter 9. We examine two specifications of this version in detail (the Hull-White and Alternative CIR models of Chapters 5 and 8); more complicated specifications seem difficult to work with.

The third part of Ehlers' thesis is an empirical investigation of the difference between credit default swap (CDS) rates on Japanese companies when the swaps are denominated in US Dollars and in Japanese Yen. Ehlers considered twenty-five Japanese debtors referenced by CDSs traded in both USD and JPY in the period 2000-2005. In many cases the JPY CDS rates were around 20% lower than the corresponding USD rates, and Ehlers is able to statistically reject (for each name) the hypothesis that the USD and JPY CDS rates are noisy observations of the same underlying rate. Ehlers then sets out to test whether the observed rate differences can be reproduced without a devaluation of the Yen upon default. Using a database of CDS, interest and exchange rates, and an assumption about the relationship between the real-world and USD measures, Ehlers estimates the parameters of a reasonably simple model[2] without devaluation upon default using only the USD CDS rates and the interest and exchange rate data. He finds that the correlations between the default intensities and the USDJPY exchange rate are quite low and that the CDS rate differences implied by the model are much smaller than those observed in the market, thus rejecting the idea that we can simply set the devaluation upon

[1] Shortly before this dissertation was to be submitted, the author became aware that Philippe Ehlers has written a book called *Credit Derivatives: Pricing and Modelling of Credit Portfolio and Credit Hybrid Derivatives*. While we were not able to obtain a copy of this book, it appears to have similar contents to Ehlers' thesis.

[2] The model used is the Alternative CIR model of Chapter 8 with stochastic but independent interest rates.

default to zero (at least with the other historically-estimated parameters).

In an unpublished paper, Li [33] produces some results similar to Ehlers' and our own, and also examines the use of copulas to model the dependence of survival and the exchange rate at some future time, an approach that we do not consider.

1.4 The Structure of the Dissertation

Chapter 3 presents Ehlers' general model of a multi-currency market with default risk. Since the mathematics underlying this model is rather technical, we recall some of the requisite details in the preceding Chapter 2; the subjects covered include predictability, compensators, stochastic functions, random measures and point processes. Chapter 4 presents in detail a more basic version of the general model – it is this basic model that we will most often use.

The next three chapters present three specifications of the basic model. The specifications differ in what process we assume for the default intensity: Chapter 5 uses a Hull-White Gaussian process, Chapter 6 a Cox-Ingersoll-Ross square-root process and Chapter 7 the lognormal process of the Black-Karasinski interest rate model. Another basic model, due to Ehlers, that uses a Cox-Ingersoll-Ross process for the default intensity and a non-standard exchange rate process, is examined in Chapter 8.

The last three main chapters cover rather diverse topics. Chapter 9 recounts the affine version of Ehlers' general model. Chapter 10 briefly examines the pricing of vulnerable options and the use of more sophisticated processes for the exchange rate (local and stochastic volatility). Lastly, Chapter 11 presents some results on the hedging of positions in defaultable claims denominated in one currency with similar claims denominated in another currency.

Chapter 12 concludes the dissertation.

Chapter 2

Mathematical Preliminaries

This chapter recalls various results, definitions and notation used in the pricing framework of Chapter 3. These results are collected from Jacod & Shiryaev [24], Liptser & Shiryaev [35] and Karatzas & Shreve [30]. More accessible references are Shreve [42] and Klebaner [31].

We work on a complete probability space $(\Omega, \mathcal{F}, P)$. Our time index set will be $\mathbb{R}_+ = [0, \infty)$ and we will consider only real-valued processes. The corresponding definitions and results for a time index set $[0, T^*]$ and $\mathbb{R}^n$-valued processes should be either obvious or easily accessible.

2.1 Filtrations, Measurability and Random Times

We denote the Borel σ-algebra on any topological space A by $\mathcal{B}(A)$ or $\mathcal{B}_A$, and abbreviate $\mathcal{B}(\mathbb{R}_+)$ to $\mathcal{B}_+$.

We denote filtrations by blackboard bold characters, and the individual σ-algebras by the corresponding calligraphic characters; for example, $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$. As usual, when given a filtration $\mathbb{F}$, we define the filtration $\mathbb{F}^-$ by

$$\mathcal{F}_t^- = \mathcal{F}_{t-} = \begin{cases} F_0 & \text{if } t = 0 \\ \bigvee_{s \in [0,t)} \mathcal{F}_s & \text{if } t > 0 \end{cases}$$

where by $\bigvee_{s \in [0,t)} \mathcal{F}_s$ we mean the smallest σ-algebra containing every element of $\mathcal{F}_s$ for all $s \in [0, t)$. We also define the σ-algebra $\mathcal{F}_\infty = \bigvee_{t \geq 0} \mathcal{F}_t$.

We denote the natural filtration of any process X by $\mathbb{F}^X$. We also assume that all our filtrations satisfy the usual conditions.

From now on we assume that we are given a filtration $\mathbb{F}$, and work on the fixed filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, P)$.

A *stochastic process* is a real-valued function on $\Omega \times \mathbb{R}_+$ such that, for each fixed $t \geq 0$, $X(t)$ is a random variable (i.e. $\omega \mapsto X(\omega, t)$ is $\mathcal{F}$-measurable). A process

X is said to be *adapted* if $X(t)$ is $\mathcal{F}_t$-measurable for each $t \geq 0$. X is said to be *measurable* if it is $(\mathcal{F} \otimes \mathcal{B}_+)$-measurable. X is said to be *progressively measurable* or *progressive* if, for each $t \geq 0$, its restriction to $\Omega \times [0, t]$ is $(\mathcal{F}_t \otimes \mathcal{B}[0, t])$-measurable. If a process is progressively measurable then it is measurable and adapted, and any process that is measurable and adapted has a progressive modification.[1] Also, any adapted process that is right- or left-continuous (as defined in the next paragraph) is progressive.

A process is said to be right-continuous if its paths are right-continuous almost surely, i.e. if

$$\{\omega \in \Omega : t \mapsto X(\omega, t) \text{ (for } t \in \mathbb{R}_+) \text{ is right-continuous}\}$$

is in $\mathcal{F}$ and has probability one. Similarly, processes are said to be left-continuous, RCLL (right-continuous with left-hand limits), continuous or non-decreasing if their paths satisfy those properties almost surely.

We denote by $\mathcal{O}$ the *optional* σ*-algebra* on $\Omega \times \mathbb{R}_+$: the σ-algebra generated by all adapted RCLL processes. A process X is said to be optional if it is $\mathcal{O}$-measurable. Any optional process is progressive. Any RCLL adapted process is optional, as is any left-continuous adapted process.

We will assume that all our processes are measurable. Then if a process is adapted, we may replace it with its progressive modification, so for us 'adapted' and 'progressive' will be synonymous. Also all our processes will also be either RCLL or left-continuous, so 'adapted' and 'progressive' are the same as 'optional'.

We denote by $\mathcal{P}$ the *predictable* σ*-algebra* on $\Omega \times \mathbb{R}_+$: the σ-algebra generated by all left-continuous adapted processes. A process is said to predictable if it is $\mathcal{P}$-measurable. It can be shown that $\mathcal{P}$ is a subset (and can be a strict subset) of $\mathcal{O}$; thus any predictable process is optional, and there are optional processes that are not predictable. A predictable process is adapted to $\mathbb{F}^-$, i.e. if X is a predictable process then $X(t)$ is $\mathcal{F}_{t-}$-measurable for each $t \geq 0$. So, intuitively, the values that a predictable process takes are known an instant ahead of time.

If the filtration $\mathbb{F}$ is the natural filtration of a Brownian motion, then the optional and predictable σ-algebras coincide.

We denote the left-continuous version of any RCLL process X by X_-; this left-continuous version is defined by

$$X_-(t) = X(t-) = \begin{cases} X(t) & \text{if } t = 0 \\ \lim_{s \uparrow t} X(s) & \text{if } t > 0. \end{cases}$$

[1] A process Y is said to be a modification of a process X if $X(t) = Y(t)$ almost surely for each $t \geq 0$.

Since X_- is left-continuous, it is predictable.

A *random time* is a $[0, \infty]$-valued random variable. A random time T is called a *stopping time* if $\{T \leq t\} \in \mathcal{F}_t$ for each $t \geq 0$, which we interpret as meaning that we observe the random time when it arrives.

Let T be a stopping time. We define

$$\mathcal{F}_T = \{A \in \mathcal{F} : A \cap \{T \leq t\} \in \mathcal{F}_t \text{ for all } t \geq 0\}.$$

Then $\mathcal{F}_T$ is a σ-algebra consisting of all the events that occur before or at the random time T. If the random time T almost surely equals the deterministic time $t \in \mathbb{R}_+$, then $\mathcal{F}_T = \mathcal{F}_t$.[2]

A random time T is said to be *predictable* if its indicator process N, defined by

$$N(\omega, t) = \begin{cases} 1 & \text{if } t \geq T(\omega) \\ 0 & \text{otherwise,} \end{cases}$$

is predictable. So, roughly, a predictable random time is one that we know will arrive an instant before it does. Every predictable time is a stopping time. A random time T is a predictable time if and only if it has an *announcing sequence*, i.e. there is an increasing sequence $(T_n)_{n \in \mathbb{N}}$ of stopping times with $\lim T_n = T$ a.s. and $T_n < T$ for all $n \in \mathbb{N}$ on $\{T > 0\}$.[3]

A stopping time T is called *totally inaccessible* if $P(T = S < \infty) = 0$ for all predictable times S. Colloquially, a stopping time T is totally inaccessible if there are no times that we know will arrive an instant before they do and that have any chance of coinciding with T.

Let X be a process and T a random time. Then we define the *stopped process* X^T by $X^T(\omega, t) = X(\omega, \min\{t, T(\omega)\})$. If X is optional (predictable) and T is a stopping time then X^T is also optional (predictable, respectively).

A process X is said to have a property $\mathfrak{p}$ *locally* if there exists an increasing sequence $(T_n)_{n \in \mathbb{N}}$ of stopping times, depending on both $\mathfrak{p}$ and X, such that $\lim T_n = \infty$ a.s. and that for each $n \in \mathbb{N}$ the stopped process X^{T_n} has the property $\mathfrak{p}$.

2.2 Stochastic Integrals

Here we recall some basic facts about Itô integrals with respect to local martingales and, in particular, Brownian motion.

[2]If $T = \infty$ identically then $\mathcal{F}_T = \mathcal{F} \supseteq \mathcal{F}_\infty$ (using our earlier notation). Due to this ambiguity, one sometimes finds the definitions $\mathcal{F}_{\infty-} = \bigvee_{t \geq 0} \mathcal{F}_t$ and $\mathcal{F}_\infty = \mathcal{F}$; then the notation $\mathcal{F}_T$ is completely unambiguous. We will not use this convention.

[3]We denote the natural numbers $\{1, 2, \ldots\}$ by $\mathbb{N}$.

Let X be a local martingale. We can define the stochastic Itô integral process $H \cdot X$, where

$$(H \cdot X)(t) = \int_0^t H(s)dX(s),$$

for any process H that is predictable and locally bounded. This integral process is null at zero, adapted, and a local martingale. Like X, $H \cdot X$ is RCLL and it jumps only at the jump times of X.

Let W be a standard, one-dimensional Brownian motion. Stochastic integrals with respect to W can be defined for a much wider class of integrands than the locally bounded, predictable processes. In fact, the integral process $H \cdot W$ is defined for any process H that is measurable and adapted and satisfies

$$\int_0^t H^2(s)ds < \infty \text{ a.s. for all } t \geq 0.$$

This integral process is null at zero, adapted and continuous, and is a local martingale. If H also satisfies

$$E\left[\int_0^t H^2(s)ds\right] < \infty \text{ for all } t \geq 0$$

then $H \cdot W$ is a martingale, and $(H \cdot W)(t)$ is square-integrable for each $t \geq 0$.

Throughout this dissertation, when we write an integral or differential we will implicitly assume that the processes involved allow the existence of these integrals – sufficient conditions for this will seldom be given explicitly.

2.3 Compensators

A process X is said to be integrable if $\sup_{t \geq 0} E[|X(t)|] < \infty$. A process is said to be of integrable variation if its variation process is integrable; since the variation process is non-decreasing, a process is of integrable variation if and only if the expectation of its variation over $\mathbb{R}_+$ is finite.

We denote by $\mathcal{A}$ the class of real-valued processes that are null at zero, are adapted, have RCLL paths a.s., and are of integrable variation. We denote by $\mathcal{A}_{\text{loc}}$ the class of processes locally in $\mathcal{A}$.

Let $X \in \mathcal{A}_{\text{loc}}$. Then there exists a predictable process $A \in \mathcal{A}_{\text{loc}}$ such that $X - A$ is a local martingale. The process A is unique up to an evanescent set and is called the *compensator* (or *predictable compensator* or *dual predictable projection*) of X. This compensator also has the property that for any predictable process H such that $H \cdot X \in \mathcal{A}_{\text{loc}}$, $H \cdot A$ is the compensator of $H \cdot X$.

We denote by $\mathcal{A}^+$ the class of real-valued processes X that are null at zero, are adapted, have RCLL, non-decreasing paths, and are integrable. We denote by $\mathcal{A}^+_{\text{loc}}$ the class of processes that are locally in $\mathcal{A}^+$.

If X is a process that is null at zero and has non-decreasing paths, then X and its variation coincide: the variation of X over $[0,t]$ is $X(t)$, for any $t \geq 0$. Thus such an X is integrable if and only if it is of integrable variation. From this we can see that $\mathcal{A}^+$ is the set of non-decreasing processes in $\mathcal{A}$: $\mathcal{A}^+ \subseteq \mathcal{A}$ and $\mathcal{A}^+_{\text{loc}} \subseteq \mathcal{A}_{\text{loc}}$.

Let $X \in \mathcal{A}^+_{\text{loc}}$. Then X has a compensator A, and in addition we have $A \in \mathcal{A}^+_{\text{loc}}$, $E[X(T)] = E[A(T)]$ for any stopping time T, and

$$E\left[\int_0^\infty H(t)dX(t)\right] = E\left[\int_0^\infty H(t)dA(t)\right]$$

for any non-negative predictable process H. If X is in $\mathcal{A}^+$, then its compensator A is also in $\mathcal{A}^+$ and $X - A$ is a uniformly integrable martingale.

The indicator process N of a stopping time T, given by $N(t) = I_{\{t \geq T\}}$, is in $\mathcal{A}$, and hence has a compensator.[4] This compensator is continuous if and only if T is totally inaccessible.

2.4 Random Measures and Point Processes

A point process is a random sequence of times. Each time corresponds to an event, such as the arrival of a vehicle at a ferry terminal. We are interested in the times when these events occur. We may also be interested in some other characteristic of the times, such as the weights of the arriving vehicles. These pairs of times and characteristics form a marked point process (the characteristics are *marks* to the *points* of occurrence).

For point processes, especially marked ones, the most elegant way of describing the process is a random measure. In the case of an unmarked point process the random measure gives, for any set of time $A \subseteq \mathbb{R}_+$, the number of events that occur in A. For a marked point process, where the possible marks (characteristics) are some set E, the random measure gives, for any $A \subseteq \mathbb{R}_+$ and $Z \subseteq E$, the number of events that occurred in A and had marks in Z.

In this section we briefly discuss general random measures, their compensators and the specific case of marked point processes following Jacod & Shiryaev [24] and Liptser & Shiryaev [34].

We assume that we are given a measurable space $(E, \mathcal{E})$ that is a Blackwell space (a separable metric space whose Borel sets do not contain a separable proper sub-σ-

[4] Throughout the dissertation we use I_A for the indicator function of the set A.

algebra) like $\mathbb{R}^n$. A *random measure* μ on $\mathbb{R}_+ \times E$ is a family $(\mu(\omega, ds \times dz) : \omega \in \Omega)$ of non-negative measures on $(\mathbb{R}_+ \times E, \mathcal{B}_+ \otimes \mathcal{E})$ with $\mu(\omega, \{0\} \times E) = 0$ for all $\omega \in \Omega$.

Let us define $\tilde{\Omega} = \Omega \times \mathbb{R}_+ \times E$, $\tilde{\mathcal{O}} = \mathcal{O} \otimes \mathcal{E}$ and $\tilde{\mathcal{P}} = \mathcal{P} \otimes \mathcal{E}$. A real-valued function on $\tilde{\Omega}$ is called a *stochastic function.* A stochastic function is called optional (predictable) if it is measurable with respect to $\tilde{\mathcal{O}}$ ($\tilde{\mathcal{P}}$, respectively).

Let V be an optional stochastic function. Then we can define a process $V * \mu$ by

$$(V * \mu)(\omega, t) = \int_{[0,t]\times E} V(\omega, s, z)\mu(\omega, ds \times dz)$$

if $\int_{[0,t]\times E} |V(\omega, s, z)|\mu(\omega, ds \times dz)$ is finite, and $(V * \mu)(\omega, t) = \infty$ otherwise. (The condition that $\mu(\omega, \{0\} \times E) = 0$ for all $\omega \in \Omega$ is imposed so that such a process is null at zero.) As usual we omit the argument ω, and use the notation

$$\int_0^t \int_E V(s, z)\mu(ds \times dz) = \int_{[0,t]\times E} V(s, z)\mu(ds \times dz).$$

A random measure μ is said to be optional (predictable) if, for any optional (predictable) stochastic function V, the process $V * \mu$ is optional (predictable). The natural filtration $\mathbb{F}^\mu$ of a random measure μ is the smallest filtration with respect to which μ is optional.

Now let μ be an optional random measure. It is said to be integrable if $\mu(\mathbb{R}_+ \times E)$ is integrable. It is said to be $\tilde{\mathcal{P}}$-σ-finite if there is a partition $(A_n)_{n\in\mathbb{N}}$ of $\tilde{\Omega}$ with each A_n in $\tilde{\mathcal{P}}$ such that $(I_{A_n} * \mu)(\infty)$ is integrable for each $n \in \mathbb{N}$; this is equivalent to the existence of a strictly positive predictable stochastic function U such that $(U * \mu)(\infty)$ is integrable.

A *transition kernel* α of one measurable space $(C, \mathcal{C})$ into another $(D, \mathcal{D})$ is a family $(\alpha(c, \cdot) : c \in C)$ of non-negative measures on $(D, \mathcal{D})$ such that $c \mapsto \alpha(c, B)$ is $\mathcal{C}$-measurable for each $B \in \mathcal{D}$.

We will call a kernel K of $(\Omega \times \mathbb{R}_+, \mathcal{F} \otimes \mathcal{B}_+)$ into $(E, \mathcal{E})$ *predictable* if $(\omega, t) \mapsto K(\omega, t, Z)$ is $\mathcal{P}$-measurable for each $Z \in \mathcal{E}$, i.e. if K is actually a kernel of $(\Omega \times \mathbb{R}_+, \mathcal{P})$ into $(E, \mathcal{E})$.

Theorem 2.1. *Let μ be an optional $\tilde{\mathcal{P}}$-σ-finite random measure. Then there exists a predictable random measure ν satisfying the following equivalent properties:*

1. *for every non-negative predictable stochastic function V,*

$$E[(V * \mu)(\infty)] = E[(V * \nu)(\infty)],$$

or in other notation

$$E\left[\int_0^\infty \int_E V(s, z)\mu(ds \times dz)\right] = E\left[\int_0^\infty \int_E V(s, z)\nu(ds \times dz)\right].$$

2. if V is a predictable stochastic function such that $|V| * \mu \in \mathcal{A}^+_{\text{loc}}$, then $|V| * \nu \in \mathcal{A}^+_{\text{loc}}$ and $V * \mu - V * \nu$ is a local martingale.

This random measure ν is unique up to a P-null set, and is called the compensator, predictable compensator *or* dual predictable projection *of μ.*

Also, there exists a predictable $A \in \mathcal{A}^+$ and a predictable kernel K from $\Omega \times \mathbb{R}_+$ into E such that

$$\nu(\omega, ds \times dz) = A(\omega, dt)K(\omega, t, dz) \text{ a.s.}$$

For predictable stochastic functions V such that $|V| * \mu \in \mathcal{A}^+_{\text{loc}}$ we can define an integral process with respect to the compensated measure $(\mu - \nu)$ by

$$V * (\mu - \nu) = V * \mu - V * \nu.$$

This integral process can also be defined for a larger class of predictable stochastic functions; see Jacod & Shiryaev [24] or Liptser & Shiryaev [34]. It is always a local martingale.

A random measure μ is said to be *integer-valued* if (1) it is optional, (2) it is $\tilde{\mathcal{P}}$-σ-finite, (3) for each $t \in \mathbb{R}_+$, $\mu(\{t\} \times E) \in \{0, 1\}$ a.s., and (4) for each $A \in \mathcal{B}_+ \otimes \mathcal{E}$ the random variable $\mu(A)$ takes values in $\{0, 1, 2, \ldots, \infty\}$.

A *marked point process* with mark space E is an integer-valued random measure μ on $\mathbb{R}_+ \times E$ such that $\mu([0, t] \times E) < \infty$ for each $t \geq 0$ a.s. A *point process* is an integer-valued random measure μ on $\mathbb{R}_+$ such that $\mu([0, t]) < \infty$ for each $t \geq 0$ a.s. – essentially a marked point process with mark space $\{1\}$.[5]

Let μ be a marked point process with mark space E. Define, for each $n \in \mathbb{N}$, the random time $T_n = \inf\{t \in \mathbb{R}_+ : \mu([0, t] \times E) \geq n\}$. Then each T_n is a stopping time, $T_n < T_{n+1}$ on $\{T_n < \infty\}$ $(n \in \mathbb{N})$, and $T_n \uparrow \infty$ a.s. as $n \to \infty$. Also there exists a sequence $(Z_n)_{n \in \mathbb{N}}$ of E-valued random variables, Z_n measurable with respect to $\mathcal{F}_{T_n}$ for each n, such that

$$\mu(A) = \sum_{n=1}^{\infty} \varepsilon_{(T_n, Z_n)}(A) \text{ for all } A \in \mathcal{B}_+ \otimes \mathcal{E} \tag{2.1}$$

where

$$\varepsilon_{(T_n, Z_n)}(\omega, A) = \begin{cases} 1 & \text{if } (T_n(\omega), Z_n(\omega)) \in A \\ 0 & \text{otherwise} \end{cases}$$

for each $n \in \mathbb{N}$, $\omega \in \Omega$ and $A \in \mathcal{B}_+ \otimes \mathcal{E}$. The T_n's represent the times when our events occur, and the Z_n's represent the associated marks.

[5] Many definitions of point processes allow multiple events to occur at the same time. This definition does not.

2.5 Girsanov's Theorem

We will use the following version of Girsanov's Theorem, adapted from Schönbucher [40]. We work here with a finite time horizon $T^* > 0$; all processes and filtrations are indexed by $[0, T^*]$. Note the particular form that we assume for the compensator of the marked point process.

Theorem 2.2. *Let $(\Omega, \mathcal{F}, \mathbb{F}, P)$ be a filtered probability space that satisfies the usual conditions and supports an n-dimensional standard Brownian motion W $(n \in \mathbb{N})$ and a marked point process μ with Blackwell mark space $(E, \mathcal{E})$. Suppose that the compensator ν of μ takes the form*

$$\nu(ds \times dz) = K(s, dz)\lambda(s)ds$$

where K is a predictable kernel from $\Omega \times \mathbb{R}_+$ to E that is almost everywhere a probability measure, and λ is a non-negative, progressive process.

Let η be an $\mathbb{R}^n$-valued predictable process and ϕ a non-negative predictable stochastic function such that

$$\int_0^{T^*} ||\eta(s)||^2 ds < \infty \ a.s.$$

$$\int_0^{T^*} \int_E |\phi(s, z)| \nu(ds \times dz) < \infty \ a.s.$$

Define a process L by $L(0) = 1$ and

$$\frac{dL(t)}{L(t-)} = \eta(t) \cdot dW(t) + \int_E (\phi(s, z) - 1)(\mu - \nu)(ds \times dz).$$

If L is a martingale, then we can define a probability measure Q on $(\Omega, \mathcal{F})$ by

$$Q(A) = \int_A L(T) dP \text{ for all } A \in \mathcal{F},$$

and under this measure we have

1. *The process $\widetilde{W}$, defined by*

$$\widetilde{W}(t) = W(t) - \int_0^t \eta(s) ds,$$

is a Q-Brownian motion.

2. *The compensator of μ under Q is $\tilde{\nu}$ where*

$$\begin{aligned}\tilde{\nu}(ds \times dz) &= \phi(s, z)\nu(ds \times dz) \\ &= \tilde{K}(s, dz)\tilde{\lambda}(s)ds\end{aligned}$$

with $\tilde{K}$ a predictable kernel that is almost everywhere a probability measure, defined by

$$\tilde{K}(s, dz) = \frac{\phi(s, z)K(s, dz)}{\int_E \phi(s, x)K(s, dx)},$$

and $\tilde{\lambda}$ is a non-negative, progressive process defined by

$$\tilde{\lambda}(t) = \lambda(t) \int_E \phi(t, x)K(t, dx).$$

If L is positive, then also P and Q are equivalent.

Chapter 3

Ehlers' Modelling Framework

This chapter presents the modelling framework used in the rest of this dissertation.

Sections 3.1-3.5 recapitulate a general reduced-form model for a market with multiple currencies and a single credit-risky participant. This model comes from the doctoral thesis of Philippe Ehlers [14] and is presented with only minor adjustments. This model allows dependence between the exchange rate and the credit risk through correlation between the exchange rate and the default intensity and through a jump in the exchange rate at the time of default. Section 3.6 gives expressions for the prices of several standard instruments.

A more basic version of this model is presented in Chapter 4 and used in most of the rest of the dissertation.

3.1 Probability Space and Default

We work on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, P_d)$ satisfying the usual conditions. The filtration $\mathbb{F}$ represents the information available to the market. We will sometimes make use of other filtrations; when we do not specify a filtration (for example, if we say that a process is predictable) we mean $\mathbb{F}$. The model will be constructed so that P_d is a domestic risk-neutral measure. We give the model a finite time horizon T^*; all processes and filtrations are indexed by $[0, T^*]$.

This probability space is assumed to support both a standard n-dimensional Brownian motion W (for some positive integer n) and a marked point process μ with marks in $\mathcal{Z} = [0,1]^k$ for some positive integer k. We suppose that μ is almost surely $\{0,1\}$-valued, so the point process has at most one 'jump'.

We assume that only one agent in the market may default, and that this default occurs at the time of the jump in μ.[1] The default indicator process N is defined by

$$N(t) = \int_0^t \int_{\mathcal{Z}} \mu(ds \times dz).$$

Thus N jumps from zero to one at the time of default

$$\tau = \inf\{t \in [0, T^*] : N(t) = 1\},$$

and so $N(t) = I_{\{t \geq \tau\}}$, by which we mean that

$$N(\omega, t) = \begin{cases} 1 & \text{if } t \geq \tau(\omega) \\ 0 & \text{otherwise.} \end{cases}$$

Like many authors, we use 'default' as a generic term for any credit event. Where credit default swaps are concerned, it is especially important that the credit event is recognised as such by the International Swaps and Derivatives Association, who determine whether or not these swaps trigger. Recognised credit events differ from contract to contract, and may include failure to pay, bankruptcy (for corporates), debt restructuring, obligation acceleration, and repudiation or moratorium (for sovereigns).

We also define a process J by

$$J(t) = \int_0^t \int_{\mathcal{Z}} z\mu(ds \times dz).$$

J is zero (in $\mathbb{R}^k$) until the time of default when it jumps to $J(\tau)$, a $\mathcal{Z}$-valued random variable which is called the *severity of default.* After τ, J is constant. The severity of default does not usually represent anything directly. The appreciation or devaluation of the foreign currency at the time of default, and the recovery rates at default of various assets, will be defined as functions of $J(\tau)$. While these functions will be stochastic, they will also be predictable; thus $J(\tau)$ represents the randomness of the default event. Since the exchange rate jump and the recovery rates are all functions of the severity of default, dependence between them is easily allowed.

We denote the predictable compensator of μ by ν, and assume that ν takes the form

$$\nu(ds \times dz) = K(s, dz)\lambda^*(s)ds$$

[1] Ehlers does not restrict μ to be $\{0, 1\}$-valued – he allows μ to have multiple jumps. Since the default time is always taken to be the time of the first jump, this seems to be an unnecessary complication. Ehlers does indicate how multiple jumps could be incorporated into the model with the default time not necessarily being the time of the first jump, but he does not pursue this, and neither do we.

where K is a predictable kernel that for any fixed $(\omega, t) \in \Omega \times [0, T^*]$ is a probability measure on $(\mathcal{Z}, \mathcal{B}_{\mathcal{Z}})$; $\lambda^*(t) = I_{\{t \leq \tau\}} \lambda(t)$ for all $t \in [0, T^*]$; and λ is a non-negative process given by a deterministic initial condition $\lambda(0)$ and

$$d\lambda(t) = \alpha(t)dt + \phi(t) \cdot dW(t)$$

for some real-valued process α and some $\mathbb{R}^n$-valued process ϕ, both $\mathbb{F}^W$-adapted. Then λ is continuous and $\mathbb{F}^W$-adapted, while λ^* is $\mathbb{F}$-predictable, but not adapted to $\mathbb{F}^W$.

The probability measure $K(\omega, t, dz)$ is the conditional distribution of the severity of default given that we are on the path ω and that default occurs at time t. Note that

$$\int_{\mathcal{Z}} K(t, dz) = 1 \text{ for all } (\omega, t) \in \Omega \times [0, T^*].$$

The compensator of N is

$$\begin{aligned} A(t) &= \int_0^t \int_{\mathcal{Z}} \nu(ds \times dz) = \int_0^t \int_{\mathcal{Z}} K(s, dz) \lambda^*(s) ds \\ &= \int_0^t \lambda^*(s) ds = \int_0^{t \wedge \tau} \lambda(s) ds. \end{aligned}$$

So, roughly, $\lambda^*(s)ds$ is the probability at time s of default in the next time instant ds. Since N is constant after τ, A must be too, and we indeed have $A(t \wedge \tau) = A(t)$ for all $t \in [0, T^*]$. Since A is continuous, τ is totally inaccessible.

In several places in this model we will assume that quantities are not just $\mathbb{F}$-predictable or $\mathbb{F}$-adapted (properties necessary for the model to make sense) but are in fact $\mathbb{F}^W$-adapted, and hence $\mathbb{F}^W$-predictable. The first three of these assumptions were included in the paragraph on the form of ν – we could instead have used the weaker conditions that K be $\mathbb{F}$-predictable and α and ϕ be $\mathbb{F}$-adapted. These assumptions allow more tractable pricing expressions and will be noted whenever they are made. In each case, the assumption can be interpreted as forcing the quantity involved not to depend upon μ or any other information (besides the evolution of W) given in $\mathbb{F}$. This is stronger than not allowing the quantity to jump at the default time: for example, the process $(t \wedge \tau)_{t \in [0, T^*]}$ does not jump at default, but it does depend upon the point process.

Definition 3.1. Let τ be a stopping time on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, P)$. A progressively measurable process γ is called an intensity process of τ if

$$I_{\{t \geq \tau\}} - \int_0^{t \wedge \tau} \gamma(s) ds$$

is a local martingale. If γ is adapted to a subfiltration $\mathbb{G}$ of $\mathbb{F}$ then γ is also called a $\mathbb{G}$-intensity process.

A filtration $\mathbb{G}$ is said to be a subfiltration of $\mathbb{F}$ if $\mathcal{G}_t \subseteq \mathcal{F}_t$ for all t.

Both λ and λ^* are intensities of the default time, and λ is an $\mathbb{F}^W$-intensity. Clearly the default intensity is not unique, but since we work with a fixed λ we will call it (and λ^*) *the* default intensity.

Note that an intensity is defined with respect to a particular probability measure, which in our case is the domestic risk-neutral measure. We will later define a foreign risk-neutral probability measure, under which τ will have a different intensity. Thus we refer to λ as the domestic default intensity.

3.2 Interest Rates and the Exchange Rate

In our market there are two currencies, which we call domestic and foreign. We assume that riskless borrowing and investment are possible in each currency, and denote the domestic and foreign default-free short rate processes by r_d and r_f respectively. These processes are assumed to be $\mathbb{F}^W$-adapted and to satisfy

$$\int_0^{T^*} |r_d(s)|ds + \int_0^{T^*} |r_f(s)|ds < \infty \text{ a.s.}$$

The domestic money-market account is created by investing at time zero a unit of domestic currency at the domestic short rate. The domestic currency value of this account at any time t is

$$M_d(t) = \exp\left\{\int_0^t r_d(s)ds\right\}.$$

The foreign currency price process of the analogous foreign money-market account is

$$M_f(t) = \exp\left\{\int_0^t r_f(s)ds\right\}.$$

These processes are continuous and $\mathbb{F}^W$-adapted.

Note that we could instead have supposed that the interest rates were $\mathbb{F}$-adapted, rather than $\mathbb{F}^W$-adapted. We have made the stronger assumption in order to get more tractable pricing expressions later on.

The exchange rate is the price of a unit of foreign currency expressed in units of domestic currency. We suppose that the exchange rate process Q satisfies

$$\frac{dQ(t)}{Q(t-)} = (r_d(t) - r_f(t))dt + \eta(t) \cdot dW(t) + \int_{\mathcal{Z}} \delta(t,z)(\mu - \nu)(dt \times dz).$$

Here η is an $\mathbb{F}^W$-adapted $\mathbb{R}^n$-valued process, and δ is an $\mathbb{F}^W$-adapted stochastic function taking values in $(-1, \infty)$ and such that

$$\int_{\mathcal{Z}} |\delta(t,z)| K(t,dz) < \infty \text{ for all } t \in [0,T^*] \text{ a.s.}$$

Note that for tractability we have again imposed stronger assumptions than necessary on η and δ – we could have had η $\mathbb{F}$-adapted and δ $\mathbb{F}$-predictable.

The solution of this stochastic differential equation is

$$\begin{aligned} Q(t) &= Q(0)\left(1+\int_0^t\int_{\mathcal{Z}}\delta(s,z)\mu(ds\times dz)\right)\times \\ &\qquad \exp\left\{\int_0^t\left[r_d(s)-r_f(s)-\hat{\delta}(s)\lambda^*(s)-\frac{1}{2}||\eta(s)||^2\right]ds+\int_0^t\eta(s)\cdot dW(s)\right\} \\ &= Q(0)\left(1+I_{\{t\geq\tau\}}\delta(\tau,J(\tau))\right)e^{\int_0^t[r_d(s)-r_f(s)-\hat{\delta}(s)\lambda^*(s)-\frac{1}{2}||\eta(s)||^2]ds+\int_0^t\eta(s)\cdot dW(s)}\end{aligned}$$

where the process $\hat{\delta}$ is defined by

$$\hat{\delta}(t)=\int_{\mathcal{Z}}\delta(t,z)K(t,dz).$$

If $\delta = 0$ identically then Q is a geometric Brownian motion with stochastic volatility $||\eta(t)||$. When δ is not identically zero, the exchange rate jumps at the default time from $Q(\tau-)$ to

$$Q(\tau)=Q(\tau-)\big(1+\delta(\tau,J(\tau))\big).$$

This means that the exchange rate jumps by a fraction $\delta(\tau,J(\tau))$ of its value when default occurs. This represents an appreciation or devaluation of the foreign currency relative to the domestic currency. Since we restrict the devaluation to be less than 100%, Q is positive up to an evanescent set.[2]

The process $\hat{\delta}$ is $\mathbb{F}^W$-adapted, so for fixed t, $\hat{\delta}(t)$ is $\mathcal{F}_t^W$-measurable (and hence $\mathcal{F}_{t-}$-measurable) and represents the expected proportional appreciation of the foreign currency relative to the domestic currency if the default occurs at time t. It is called the *locally expected appreciation fraction* (with respect to P_d) and can depend upon the evolution of the Brownian motion as well as on time, but not upon μ.

The exchange rate drift is chosen so that the domestic currency value of the foreign money-market account, relative to the domestic money-market account, is a local martingale:

$$\begin{aligned} d\left(\frac{Q(t)M_f(t)}{M_d(t)}\right) &= \eta(t)\left(\frac{Q(t-)M_f(t)}{M_d(t)}\right)\cdot dW(t) \\ &\qquad +\int_{\mathcal{Z}}\delta(t,z)\left(\frac{Q(t-)M_f(t)}{M_d(t)}\right)(\mu-\nu)(dt\times dz).\end{aligned}$$

The only two assets in our model so far are the domestic and foreign money-market accounts. The domestic account, relative to itself, is clearly a P_d-martingale,

[2]Ehlers assumes that the foreign currency cannot gain value upon default; hence he subtracts the integral with respect to $(\mu-\nu)$ in the dynamics of Q, and assumes that δ is $[0,1]$-valued.

and we have just seen that the domestic value of the foreign account, relative to the domestic account, is a local martingale. Thus the value, relative to the domestic account, of any self-financing portfolio consisting of locally bounded predictable amounts of these two assets is a local martingale. This justifies calling P_d the domestic risk-neutral measure.

We will define the domestic currency price of any new asset as the conditional expectation under P_d, given the market information at the valuation time, of the asset's future payoffs in domestic currency discounted at the domestic default-free rate. This ensures that the domestic currency prices of all assets, relative to the domestic money-market account, are local martingales, and hence that the domestic currency value of any self-financing, locally bounded, predictable portfolio of them, relative to the domestic money-market account, is a local martingale.

This model, like many credit and interest rate models, is constructed under a risk-neutral measure (we have assumed dynamics for the quantities of interest under a risk-neutral, rather than real-world, measure). We do not assume that this risk-neutral measure is unique, but we do assume that the particular risk-neutral measure P_d above is used for all pricing.

3.3 The Foreign Risk-Neutral Measure

3.3.1 Definition

We define a process L^Q by

$$L^Q(t) = \frac{Q(t)M_f(t)}{Q(0)M_d(t)}.$$

Then L^Q is just a rescaling of the domestic value of the foreign money-market account, and is a positive local martingale:

$$dL^Q(t) = \eta(t)L^Q(t-)\cdot dW(t) + \int_{\mathcal{Z}} \delta(t,z)L^Q(t-)(\mu-\nu)(dt\times dz).$$

We assume the stronger property that L^Q is a true martingale. Then $E_d[L^Q(T^*)] = L^Q(0) = 1$ and so we can define a probability measure P_f on $(\Omega, \mathcal{F})$ by

$$P_f(A) = \int_A L^Q(T^*)dP_d \text{ for all } A \in \mathcal{F}. \tag{3.1}$$

This probability measure P_f is equivalent to P_d, and is called the foreign risk-neutral measure.

3.3.2 Consequences

From Girsanov's Theorem (Theorem 2.2, page 13) we know that

1. The $\mathbb{R}^n$-valued process $\widetilde{W}$, defined by

$$\widetilde{W}(t) = W(t) - \int_0^t \eta(s)ds,$$

is a P_f-Brownian motion.

2. The compensator measure $\tilde{\nu}$ of μ under P_f is given by

$$\tilde{\nu}(dt \times dz) = (1 + \delta(t,z))\nu(dt \times dz) = \tilde{K}(t,dz)\tilde{\lambda}^*(t)dt$$

where

(a) $\tilde{K}$ is a predictable kernel that is almost everywhere a probability measure, given by

$$\tilde{K}(t,dz) = \frac{(1+\delta(t,z))K(t,dz)}{\int_{\mathcal{Z}}(1+\delta(t,x))K(t,dx)}$$

(b) the process $\tilde{\lambda}$ is defined by

$$\tilde{\lambda}(t) = \lambda(t)\int_{\mathcal{Z}}(1+\delta(t,x))K(t,dx) = \left(1+\hat{\delta}(t)\right)\lambda(t)$$

(c) the process $\tilde{\lambda}^*$ is given by $\tilde{\lambda}^*(t) = I_{\{t\le\tau\}}\tilde{\lambda}(t)$.

Now $\tilde{\lambda} = (1+\hat{\delta})\lambda$ is $\mathbb{F}^W$-predictable and the compensator of N under P_f is

$$\begin{aligned}\int_0^t\int_{\mathcal{Z}} \tilde{\nu}(dt \times dz) &= \int_0^t\int_{\mathcal{Z}} \tilde{K}(t,dz)\tilde{\lambda}^*(t)dt \\ &= \int_0^t \tilde{\lambda}^*(t)dt \\ &= \int_0^{t\wedge\tau} \tilde{\lambda}(t)dt.\end{aligned}$$

Thus under P_f the default time τ has $\mathbb{F}^W$-intensity $\tilde{\lambda} = (1+\hat{\delta})\lambda$ and is still totally inaccessible. This means that the default intensity under the foreign measure is the domestic intensity scaled by the P_d-expected value of Q after default (if default occurs now) relative to its value a second ago.

Since η is $\mathbb{F}^W$-adapted, $\widetilde{W}$ is adapted to $\mathbb{F}$ and in fact $\mathbb{F}^W = \mathbb{F}^{\widetilde{W}}$ (i.e. $\mathcal{F}_t^W = \mathcal{F}_t^{\widetilde{W}}$ for all $t \in [0,T^*]$).

3.3.3 The Foreign Pricing Formula

Note that L^Q is the density process:

$$\left.\frac{dP_f}{dP_d}\right|_{\mathcal{F}_t} = L^Q(t) \text{ for all } t \in [0,T^*].$$

The point of using the foreign risk-neutral measure is the following. Let $T \in [0, T^*]$, and let $X(T)$ be an $\mathcal{F}_T$-measurable random variable. Then the foreign currency price $X(t)$ at time $t \in [0, T]$ of a claim that pays X units of foreign currency at time T is given by

$$Q(t)X(t) = E_d \left[e^{-\int_t^T r_d(s)ds} Q(T)X(T) \Big| \mathcal{F}_t \right]$$

and so

$$\begin{aligned} X(t) &= \frac{1}{Q(t)} E_d \left[e^{-\int_t^T r_d(s)ds} Q(T)X(T) \Big| \mathcal{F}_t \right] \\ &= \frac{L^Q(t)}{Q(t)} E_f \left[\frac{1}{L^Q(T)} e^{-\int_t^T r_d(s)ds} Q(T)X(T) \Big| \mathcal{F}_t \right] \\ &= \frac{M_f(t)}{Q(0)M_d(t)} E_f \left[\frac{Q(0)M_d(T)}{Q(T)M_f(T)} e^{-\int_t^T r_d(s)ds} Q(T)X(T) \Big| \mathcal{F}_t \right] \\ &= E_f \left[e^{-\int_t^T r_f(s)ds} X(T) \Big| \mathcal{F}_t \right]. \end{aligned}$$

The *foreign currency* price of a claim denominated in *foreign currency* is given by the usual formula with P_d and r_d replaced by P_f and r_f respectively.

Note that any domestic risk-neutral measure defines a corresponding foreign risk-neutral measure via (3.1). Since there might be many domestic risk-neutral measures, there might also be many foreign risk-neutral measures. However, we pick a particular domestic risk-neutral measure P_d for all our pricing, and the change from P_d to P_f is simply a technique to simplify P_d-expectations that involve the exchange rate. Thus we need not discuss the uniqueness of the foreign measure, which can be seen as a technical device for computing prices with no economic content of its own.

3.4 Martingale Invariance Property

A process is said to be square-integrable if $\sup E[X^2(t)] < \infty$, where the supremum is taken over the values of t for which the process is defined. Let M be a square-integrable $(P_d, \mathbb{F}^W)$-martingale. Then by the predictable representation property of Brownian motion (see Karatzas & Shreve [30]) there exists an $\mathbb{F}^W$-progressively measurable, $\mathbb{R}^n$-valued process θ such that

$$M(t) = M(0) + \int_0^t \theta(s) \cdot dW(s) \text{ for all } t \in [0, T^*] \tag{3.2}$$

and

$$\int_0^{T^*} ||\theta(s)||^2 ds < \infty.$$

Since W is also a Brownian motion with respect to $\mathbb{F}$, and θ is also progressively measurable with respect to $\mathbb{F}$, the representation (3.2) means that M is also a square-integrable martingale with respect to $\mathbb{F}$. Thus, the initial assumption that W is a Brownian motion with respect to $\mathbb{F}$ (and not just with respect to $\mathbb{F}^W$) implies that any square-integrable $(P_d, \mathbb{F}^W)$-martingale is also a square-integrable $(P_d, \mathbb{F})$-martingale. A similar argument implies that any square-integrable $(P_f, \mathbb{F}^W)$-martingale is also a square-integrable $(P_f, \mathbb{F})$-martingale.

Definition 3.2. Let $\mathbb{F}$ and $\mathbb{G}$ be two filtrations on a probability space $(\Omega, \mathcal{F}, P)$ with $\mathbb{G}$ a subfiltration of $\mathbb{F}$. $\mathbb{G}$ is said to have the martingale invariance property with respect to $\mathbb{F}$ under P if every $(P, \mathbb{G})$-martingale is also a $(P, \mathbb{F})$-martingale.

In this model we have shown that $\mathbb{F}^W$ has (a variation of) the martingale invariance property with respect to $\mathbb{F}$ under both the domestic and foreign risk-neutral measures. It seems no great stretch to assume the full martingale invariance property, but we will not do this.

In many reduced-form credit risk models, one starts with a reference filtration with respect to which all the default-free assets (relative to the money-market account) are martingales. One then constructs or assumes the existence of a default time τ which is not a stopping time in this reference filtration. The market filtration, which one uses for pricing, is then taken to be the smallest filtration that makes τ a stopping time and contains the reference filtration. Almost invariably the reference filtration satisfies the martingale invariance property (also called the $(\mathcal{H})$-hypothesis) with respect to the market filtration, by either construction or assumption. In such a model, the martingale invariance property means that the discounted price processes of all the default-free assets are still martingales with respect to the market filtration, and so the property is interpreted as 'if the default-free market is arbitrage free, then so is the market including default' – see, for example, Section 7.5 in Jeanblanc, Yor & Chesney [28].

Ehlers [14] gives the same interpretation in our model: 'The financial interpretation of [our variation on the martingale invariance property] is that if the (default-free) $\mathbb{F}^W$-market is arbitrage-free, then default does not introduce arbitrage in this market.' This property proves very useful in the analytical tractability of the general model.

Our weak version of the martingale invariance property implies that for any $\mathcal{F}^W_{T^*}$-measurable, P_d-square-integrable random variable Y, the conditional expectations $E_d[Y|\mathcal{F}_t]$ and $E_d[Y|\mathcal{F}^W_t]$ coincide for any $t \in [0, T^*]$.

Suppose that we are given such a Y. We define a process M by $M(t) = E_d[Y|\mathcal{F}^W_t]$. Then M is $\mathbb{F}^W$-adapted and $E_d[M(t)|\mathcal{F}^W_s] = M(s)$ for any $s < t$. Also, for any $t \in [0, T^*]$ we have

$$E_d\left[M^2(t)\right] = E_d\left[\left(E_d\left[Y|\mathcal{F}^W_t\right]\right)^2\right] \leq E_d\left[E_d\left[Y^2|\mathcal{F}^W_t\right]\right] = E_d\left[Y^2\right]$$

using Jensen's inequality, and so

$$\sup_{t\in[0,T^*]} E_d\left[M^2(t)\right] \leq E_d\left[Y^2\right] < \infty.$$

In other words, M is a square-integrable $\mathbb{F}^W$-martingale, and hence a square-integrable $\mathbb{F}$-martingale. Thus for any $t \in [0, T^*]$ we have

$$M(t) = E_d[M(T^*)|\mathcal{F}_t]$$

which we can rewrite as

$$E_d\left[Y|\mathcal{F}_t^W\right] = E_d\left[E_d\left[Y|\mathcal{F}_{T^*}^W\right]|\mathcal{F}_t\right] = E_d\left[Y|\mathcal{F}_t\right].$$

A similar argument shows that the corresponding result holds under P_f. If we also assume the full martingale invariance property, then the corresponding results also hold for Y integrable (and not necessarily square-integrable).

In the rest of the dissertation we will usually leave conditional expectations with respect to $\mathcal{F}_t^W$ as such, not rewriting them as conditional expectations with respect to $\mathcal{F}_t$ if we can. The purpose of the discussion above was to point out when this substitution is possible so that we do not attach special meaning to the fact that a conditional expectation is with respect to $\mathcal{F}_t^W$ rather than $\mathcal{F}_t$ when, as is often the case, no such meaning exists.

3.5 Two Conditional Expectation Results

The following results are taken directly from the doctoral thesis of Philippe Ehlers [14], to which we refer the reader for proofs.

The following result will give us an expression for the price of a defaultable claim that pays a possibly random amount upon survival to a future time.

Theorem 3.3 (Lemma 60, Ehlers [14]). *Let $0 \leq t \leq T \leq T^*$ and let Y be a P_d-integrable, $\mathcal{F}_T^W$-measurable random variable. Then*

$$E_d\left[I_{\{T<\tau\}}Y|\mathcal{F}_t\right] = I_{\{t<\tau\}}E_d\left[e^{-\int_t^T \lambda(s)ds}Y\middle|\mathcal{F}_t^W\right].$$

The following corollary will give us the price of a payment at the default time. This payment is a function (depending on W and time) of the severity of default.

Corollary 3.4 (Corollary 61, Ehlers [14]). *Fix t and T with $0 \leq t \leq T \leq T^*$. Let G be an $\mathbb{F}^W$-predictable stochastic function such that*

$$E_d\left[\int_0^{T^*}\int_{\mathcal{Z}} |G(t,z)|K(t,dz)\lambda^*(t)dt\right] < \infty. \tag{3.3}$$

Define the process $\hat{G}$ *by* $\hat{G}(t) = \int_{\mathcal{Z}} G(t,z)K(t,dz)$*. Then*

$$E_d\left[I_{\{\tau\in(t,T]\}}G(\tau,J(\tau))|\mathcal{F}_t\right] = I_{\{t<\tau\}}\int_t^T E_d\left[\lambda(s)e^{-\int_t^s\lambda(u)du}\hat{G}(s)\Big|\mathcal{F}_t^W\right]ds.$$

Ehlers also gives a result for the price of a defaultable continuous fee stream. Since fee streams are discrete in reality, we omit this and refer the interested reader to [14].

The results under the foreign measure P_f corresponding to Theorem 3.3 and Corollary 3.4 are given below. Note that we can still use conditional expectations with respect to elements of $\mathbb{F}^W$ because $\mathbb{F}^W$ and $\mathbb{F}^{\widetilde{W}}$ coincide.

Theorem 3.5. *Let* $0 \le t \le T \le T^*$ *and let* Y *be a* P_f*-integrable,* $\mathcal{F}_T^W$*-measurable random variable. Then*

$$E_f\left[I_{\{T<\tau\}}Y|\mathcal{F}_t\right] = I_{\{t<\tau\}}E_f\left[e^{-\int_t^T\tilde{\lambda}(s)ds}Y\Big|\mathcal{F}_t^W\right].$$

Corollary 3.6. *Fix* t *and* T *with* $0 \le t \le T \le T^*$*. Let* G *be an* $\mathbb{F}^W$*-predictable stochastic function such that*

$$E_f\left[\int_0^{T^*}\int_{\mathcal{Z}}|G(t,z)|\tilde{K}(t,dz)\tilde{\lambda}^*(t)dt\right] < \infty. \tag{3.4}$$

Define $\hat{G}$ *by* $\hat{G}(t) = \int_{\mathcal{Z}} G(t,z)\tilde{K}(t,dz)$*. Then*

$$E_f\left[I_{\{\tau\in(t,T]\}}G(\tau,J(\tau))|\mathcal{F}_t\right] = I_{\{t<\tau\}}\int_t^T E_f\left[\tilde{\lambda}(s)e^{-\int_t^s\tilde{\lambda}(u)du}\hat{G}(s)\Big|\mathcal{F}_t^W\right]ds.$$

3.6 Pricing Expressions

In this section we introduce notation for the prices of default-free zero-coupon bonds and defaultable zero-coupon bonds with zero recovery. We then briefly discuss recovery schemes before giving the prices of defaultable coupon-bearing bonds with positive recovery and credit default swaps (denominated in either domestic or foreign currency) in terms of these zero-coupon bonds and more complicated expressions for amounts paid at the time of default. These prices are also described by Ehlers [14].

3.6.1 Default-Free Zero-Coupon Bonds

Let $0 \le t \le T \le T^*$. We denote the domestic currency value at time t of a zero-coupon bond paying a unit of domestic currency at time T by $B_d(t,T)$. Then we have the definition

$$B_d(t,T) = E_d\left[e^{-\int_t^T r_d(s)ds}\Big|\mathcal{F}_t\right].$$

Similarly we denote the foreign currency value at time t of a zero-coupon bond paying a unit of foreign currency at time T by $B_f(t,T)$. Then we have

$$B_f(t,T) = E_f \left[e^{-\int_t^T r_f(s)ds} \middle| \mathcal{F}_t \right].$$

3.6.2 Defaultable Zero-Coupon Bonds with Zero Recovery

A unit domestic defaultable zero-coupon bond with zero recovery (DDZCB) is a security that pays a unit of domestic currency at its maturity if and only if default has not occurred by then. We denote the domestic currency price at time t of a DDZCB with maturity T $(0 \leq t \leq T \leq T^*)$ by $B_d^*(t,T)$. Then we have

$$\begin{aligned} B_d^*(t,T) &= E_d \left[e^{-\int_t^T r_d(s)ds} I_{\{T<\tau\}} \middle| \mathcal{F}_t \right] \\ &= I_{\{t<\tau\}} E_d \left[e^{-\int_t^T [r_d(s)+\lambda(s)]ds} \middle| \mathcal{F}_t^W \right]. \end{aligned}$$

Similarly, a security that pays a unit of foreign currency at its maturity if and only if default has not occurred by then is called a unit foreign defaultable zero-coupon bond with zero recovery (FDZCB). We denote the foreign currency price at time t of an FDZCB with maturity T $(0 \leq t \leq T \leq T^*)$ by $B_f^*(t,T)$ and we have

$$\begin{aligned} B_f^*(t,T) &= E_f \left[e^{-\int_t^T r_f(s)ds} I_{\{T<\tau\}} \middle| \mathcal{F}_t \right] \\ &= I_{\{t<\tau\}} E_f \left[e^{-\int_t^T [r_f(s)+\tilde{\lambda}(s)]ds} \middle| \mathcal{F}_t^W \right]. \end{aligned}$$

3.6.3 Including Positive Recovery

In the event of a default, payments promised by the defaulting party lose a large amount of their value. The recovery rate for a particular asset is its value after the default expressed as a proportion of some reference value for that asset. Various recovery schemes (choices of reference value) appear in the literature. The most common schemes are briefly explained below. We suppose that we are given a particular defaultable security issued by the credit-risky agent (usually a bond, though it may also be an over-the-counter derivatives position or some other asset).

In the *recovery of treasury* scheme, the reference value is a security that is identical to the one considered, except that it is default-free. In the *recovery of market value* scheme, the reference value is the value of a security that is identical to the one considered, except that default has not yet occurred. In the *recovery of par* scheme, the reference value is the nominal amount of the asset (for assets other than bonds we replace the nominal amount by the asset holder's legal claim in the event of default).

Of course, any actual recovery can be expressed in any of these forms. However, we usually assume that all assets issued by a specific agent within a specific seniority class have a common (and possibly constant) recovery rate. So the difference between these schemes is in what we assume to be common. Usually by a ‘recovery scheme’ we mean not just an expression for the recovery rate, but also the assumption that this recovery rate will be the same for all assets in some class.

We will most often express recovery rates as percentages of the asset’s nominal amount (recovery of par) which is a reasonable method for most purposes (see Schönbucher [40], to which we refer the interested reader for a thorough discussion of recovery assumptions). This scheme also makes the most sense when credit default swaps are being valued – the nominal amount of a vanilla CDS is clearly defined, while it makes little sense to speak of an equivalent undefaulted or default-free security. The pricing expressions for bonds using other recovery schemes are easily derived.

Note that in a multiple-currency situation, while we might assume that a recovery rate is constant for all bonds in a specific seniority class in one currency, the recovery rates for assets with the same seniority in different currencies are not necessarily equal, since the two recoveries may be affected by different legal procedures – see Davydenko & Franks [12].

The recovery rate on a credit default swap (CDS) is usually determined by an auction of the reference entity’s debt some time after the default – the protection seller then pays to the protection buyer $N(1-R)$ where N is the nominal amount of the CDS and R is the recovery rate. Thus credit default swaps denominated in two different currencies will have the same recovery rate if the same debt is used to determine the recovery rate for each swap. This situation is common but not universal.

(A fall-back method of settlement is for the protection buyer to deliver to the protection seller debt issued by the reference entity with the same nominal amount as the CDS in return for a payment of this nominal amount. This physical settlement gives the protection buyer a delivery option – he may present the protection seller with the cheapest portfolio with the correct nominal amount. A disadvantage of physical settlement is that the nominal amounts of CDSs traded on a particular entity may exceed the nominal amount of debt issued by that entity, leading to demand for these bonds exceeding supply in the event of default. If settlement is physical, then CDSs denominated in different currencies will have the same recovery rate if the deliverable obligations are the same, with nominal amounts converted at the spot exchange rate.)

3.6.4 Domestic Defaultable Coupon-Bearing Bonds

We consider a bond that pays, for $i = 1, 2, \ldots, n$, an amount c_i of domestic currency at time T_i if default has not occurred by then. The payment times (T_i) are an increasing sequence in $[0, T^*]$ and the amounts (c_i) are constants. Typically c_i, $i = 1, 2, \ldots, n-1$, represent coupon payments and c_n represents the final coupon and redemption of the bond's nominal amount. If default occurs before maturity T_n, a payment of $R(\tau, J(\tau))$ times the nominal amount of the bond is made to the bondholder at the time of default, where R is an $\mathbb{F}^W$-adapted, $[0,1]$-valued stochastic function.

Define $\beta(t) = \inf\{i : T_i > t\}$. Then the domestic currency value at time t of the payments in survival is

$$\begin{aligned}
&E_d\left[\sum_{i=\beta(t)}^{n} c_i e^{-\int_t^{T_i} r_d(s)ds} I_{\{T_i<\tau\}}\middle|\mathcal{F}_t\right] \\
&= \sum_{i=\beta(t)}^{n} c_i E_d\left[e^{-\int_t^{T_i} r_d(s)ds} I_{\{T_i<\tau\}}\middle|\mathcal{F}_t\right] \\
&= I_{\{t<\tau\}} \sum_{i=\beta(t)}^{n} c_i E_d\left[e^{-\int_t^{T_i}[r_d(s)+\lambda(s)]ds}\middle|\mathcal{F}_t^W\right] \\
&= \sum_{i=\beta(t)}^{n} c_i B_d^*(t, T_i).
\end{aligned}$$

For the domestic currency value of the recovery payment, first define the locally P_d-expected recovery rate $\hat{R}$ (an $\mathbb{F}^W$-adapted stochastic process) by

$$\hat{R}(t) = \int_{\mathcal{Z}} R(t,z) K(t, dz).$$

Denote the nominal amount of the bond (in domestic currency) by N. Then the value of the possible recovery payment is

$$\begin{aligned}
&E_d\left[I_{\{\tau\in(t,T_n]\}} e^{-\int_t^{\tau} r_d(u)du} N R(\tau, J(\tau))\middle|\mathcal{F}_t\right] \\
&= N E_d\left[I_{\{\tau\in(t,T_n]\}} R^1(\tau, J(\tau))\middle|\mathcal{F}_t\right]
\end{aligned}$$

where we define $R^1(s,z) = e^{-\int_t^s r_d(u)du} R(s,z)$ for $s \in [t, T_n]$. Then R^1 is a bounded, $\mathbb{F}^W$-predictable stochastic function with locally P_d-expected value

$$\hat{R}^1(s) = \int_{\mathcal{Z}} e^{-\int_t^s r_d(u)du} R(s,z) K(s,dz) = e^{-\int_t^s r_d(u)du} \hat{R}(s).$$

Now using Theorem 3.4 we conclude that the domestic value of the recovery payment is given by

$$\begin{aligned} &I_{\{t<\tau\}}N\int_t^{T_n} E_d\left[\lambda(s)e^{-\int_t^s \lambda(u)du}\hat{R}^1(s)\middle|\mathcal{F}_t^W\right]ds \\ &= I_{\{t<\tau\}}N\int_t^{T_n} E_d\left[\lambda(s)e^{-\int_t^s [r_d(u)+\lambda(u)]du}\hat{R}(s)\middle|\mathcal{F}_t^W\right]ds. \end{aligned}$$

Thus the value of the bond in domestic currency at time t is

$$\sum_{i=\beta(t)}^{n} c_i B_d^*(t,T_i) + I_{\{t<\tau\}}N\int_t^{T_n} E_d\left[\lambda(s)e^{-\int_t^s [r_d(u)+\lambda(u)du]}\hat{R}(s)\middle|\mathcal{F}_t^W\right]ds.$$

3.6.5 Foreign Defaultable Coupon-Bearing Bonds

Here we consider the foreign currency version of the bond above. It pays, for $i = 1, 2, \ldots, n$, an amount c_i of foreign currency at time T_i if default has not occurred by then, and $R(\tau, J(\tau))$ times the nominal amount of the bond at the time of default, where R is an $\mathbb{F}^W$-predictable, $[0, 1]$-valued stochastic function. Note that this recovery function (like the coupons and coupon payments dates) may be different from the corresponding quantity used to value any other domestic or foreign bond.

The foreign currency value at time t ($0 \le t \le T_n$) of the payments in survival is

$$E_f\left[\sum_{i=\beta(t)}^{n} c_i e^{-\int_t^{T_i} r_f(s)ds} I_{\{T_i<\tau\}}\middle|\mathcal{F}_t\right] = \sum_{i=\beta(t)}^{n} c_i B_f^*(t,T_i). \tag{3.5}$$

Using a procedure similar to the one above, we see that the foreign currency value of the recovery payment is

$$\begin{aligned} &E_f\left[I_{\{\tau\in(t,T_n]\}}e^{-\int_t^\tau r_f(u)du}NR(\tau,J(\tau))\middle|\mathcal{F}_t\right] \\ &= I_{\{t<\tau\}}N\int_t^{T_n} E_f\left[\tilde{\lambda}(s)e^{-\int_t^s [r_f(u)+\tilde{\lambda}(u)]du}\hat{R}(s)\middle|\mathcal{F}_t^W\right]ds \end{aligned} \tag{3.6}$$

where N is the nominal amount of the bond in foreign currency, and $\hat{R}$ is the locally P_f-expected recovery rate process

$$\hat{R}(t) = \int_{\mathcal{Z}} R(t,z)\tilde{K}(t,dz).$$

(Note the integration with respect to $\tilde{K}$ rather than K.) The foreign currency value of the bond at time t is then the sum of (3.5) and (3.6).

3.6.6 Credit Default Swaps

The International Swaps and Derivatives Association's website reports that at the end of 2010, the CDS market had a gross notional amount of USD 25.5 trillion and a net notional amount (excluding transactions that are back-to-back for one or the other party) of USD 2.3 trillion.

We consider a domestic credit default swap. Like all swaps, a CDS consists of two legs. The protection or floating leg consists of a domestic currency payment of $N(1 - R(\tau, J(\tau)))$ at the time of default, where N is the nominal amount of the CDS and R is an $\mathbb{F}^W$-predictable stochastic function representing the recovery rate (as a proportion of par) as a function of time and the default severity. This payment is only made if default occurs before a specified maturity $T > 0$. At time t $(0 \leq t \leq T \leq T^*)$ this leg has domestic value

$$\begin{aligned} &E_d \left[I_{\{\tau \in (t,T]\}} e^{-\int_t^\tau r_d(u)du} N R(\tau, J(\tau)) \middle| \mathcal{F}_t \right] \\ &= I_{\{t<\tau\}} N \int_t^T E_d \left[\lambda(s) e^{-\int_t^s [r_d(u)+\lambda(u)]du} (1 - \hat{R}(s)) \middle| \mathcal{F}_t^W \right] ds, \end{aligned} \tag{3.7}$$

with $\hat{R}$ the locally P_d-expected recovery rate process.

The premium or fixed leg consists of a sequence of domestic currency payments $(c_i)_{i=1,2,\ldots,n}$ made at the times in the increasing sequence $(T_i)_{i=1,2,\ldots,n}$. Typically this leg is specified by a rate, and the protection buyer pays

$$\text{Nominal} \times \text{CDS rate} \times \text{Year fraction since last payment}$$

quarterly to the protection seller. In the past, this rate was usually chosen so that the initial value of the swap was zero, though it is becoming common for the rate to be set at some standard value (e.g. 100bp for large corporates) and for the resulting non-zero value of the CDS at inception to be settled in cash. This has the advantage that when a market participant is party to two CDSs on the same reference entity, one as a protection seller and the other as a protection buyer, all future cashflows net and the participant's profit or loss on the transactions is realised immediately.

At the time of default, the protection buyer pays to the seller an accrual payment (the portion of the next fee payment that relates to protection already provided) and the rest of the premium leg is cancelled. The value of the premium leg at time t excluding the accrual payment is just

$$E_d \left[\sum_{i=\beta(t)}^{n} c_i e^{-\int_t^{T_i} r_d(s)ds} I_{\{T_i<\tau\}} \middle| \mathcal{F}_t \right] = \sum_{i=\beta(t)}^{n} c_i B_d^*(t, T_i). \tag{3.8}$$

The value of the accrual payment is

$$E_d\left[I_{\{\tau\in(t,T_n]\}}e^{-\int_t^\tau r_d(u)du}c_{\beta(\tau)}(\tau-T_{\beta(\tau)-1})\Big|\mathcal{F}_t\right]$$
$$=E_d\left[I_{\{\tau\in(t,T_n]\}}C(t,z)\Big|\mathcal{F}_t\right]$$

where C is a predictable stochastic function defined by

$$C(s,z)=e^{-\int_t^s r_d(u)du}c_{\beta(s)}(s-T_{\beta(s)-1})$$

for $s\in[t,T_n]$. Since this does not depend upon z, its locally P_d-expected value process is essential itself:

$$\hat{C}(s)=\int_{\mathcal{Z}}C(s,z)K(s,dz)=e^{-\int_t^s r_d(u)du}c_{\beta(s)}(s-T_{\beta(s)-1}).$$

We find that the value of the accrual payment is

$$I_{\{t<\tau\}}\int_t^{T_n}E_d\left[\lambda(s)e^{-\int_t^s\lambda(u)du}\hat{C}(s)\Big|\mathcal{F}_t^W\right]ds$$
$$=I_{\{t<\tau\}}\int_t^{T_n}E_d\left[\lambda(s)e^{-\int_t^s[r_d(u)+\lambda(u)]du}c_{\beta(s)}(s-T_{\beta(s)-1})\Big|\mathcal{F}_t^W\right]ds. \qquad (3.9)$$

Finally, the value of the CDS to the protection buyer is (3.7) less (3.8) and (3.9).

For a corresponding foreign credit default swap, the foreign currency value of the protection leg is

$$E_f\left[I_{\{\tau\in(t,T]\}}e^{-\int_t^\tau r_f(u)du}N(1-R(\tau,J(\tau)))\Big|\mathcal{F}_t\right]$$
$$=I_{\{t<\tau\}}N\int_t^T E_f\left[\tilde{\lambda}(s)e^{-\int_t^s[r_f(u)+\tilde{\lambda}(u)]du}(1-\hat{R}(s))\Big|\mathcal{F}_t^W\right]ds$$

(with $\hat{R}$ the locally P_f-expected recovery rate), the foreign currency value of the premium leg excluding the accrual payment is

$$E_f\left[\sum_{i=\beta(t)}^{n}c_ie^{-\int_t^{T_i}r_f(s)ds}I_{\{T_i<\tau\}}\Bigg|\mathcal{F}_t\right]=\sum_{i=\beta(t)}^{n}c_iB_f^*(t,T_i)$$

and the foreign currency value of the accrual payment is

$$E_f\left[I_{\{\tau\in(t,T_n]\}}e^{-\int_t^\tau r_f(u)du}c_{\beta(\tau)}(\tau-T_{\beta(\tau)-1})\Big|\mathcal{F}_t\right]$$
$$=I_{\{t<\tau\}}\int_t^{T_n}E_f\left[\tilde{\lambda}(s)e^{-\int_t^s[r_f(u)+\tilde{\lambda}(u)]du}c_{\beta(s)}(s-T_{\beta(s)-1})\Big|\mathcal{F}_t^W\right]ds.$$

The value of the CDS to the protection buyer is the first of these less the second and third.

3.6.7 Other Products

A quanto CDS is a credit default swap where the premium leg is paid in one currency and the protection leg is paid in another. Such a product is easily valued leg-by-leg as above.

Vulnerable foreign exchange forwards and fixed-for-fixed cross-currency swaps are also easily valued as (series of) defaultable payments in each currency. We can also price various forms of extinguishable cross-currency swaps – swaps that terminate on the default of a third party, or where the credit-risky party has a reduced claim in the event that he defaults while the swap is an asset to him.

Chapter 4

The Basic Model

This chapter details a simple, transparent version of Ehlers' general model. The models examined in Chapters 5, 6 and 7 will be particular specifications of this basic model, and the model of Chapter 8 is only slightly different.

First, we specify the model and consider its advantages. We will show how we can approximate payments at the default time with defaultable payments, which will allow us to focus only on the pricing of such defaultable payments. We also demonstrate a change of measure, simpler than the change from P_d to P_f, that gives our usual price for an FDZCB more directly.

4.1 Defining the Basic Model

4.1.1 Specifying the Default Intensity and Exchange Rate

Assumption 1. *We suppose that the Brownian motion W is two-dimensional.*

These two dimensions will drive the default intensity and the exchange rate. We separate the effects of the two components by the following assumption. (We use X_i to denote the ith component of a vector or vector-valued process X.)

Assumption 2. *The default intensity is driven only by the first component of the Brownian motion: $\phi_2 = 0$ identically, so*

$$d\lambda(t) = \alpha(t)dt + \phi_1(t)dW_1(t).$$

We will usually only consider diffusion models for λ, where the coefficients α and ϕ_1 are functions of λ and time.

We do not specify the process for the default intensity any further, but we do assume particular forms for the components of the exchange rate.

Assumption 3. *We assume that the random probability measure K in the compensator of μ is deterministic and time-homogeneous:*

$$K(\omega, t, Z) = K(Z) \text{ for all } (\omega, t, Z) \in \Omega \times [0, T^*] \times \mathcal{B}_{\mathcal{Z}}$$

in an abuse of notation.

We also assume that the function that converts the severity of default into the devaluation fraction is deterministic and time-homogeneous:

$$\delta(\omega, t, z) = \delta(z) \text{ for all } (\omega, t, z) \in \Omega \times [0, T^*] \times \mathcal{Z}.$$

This means that the appreciation or depreciation of the foreign exchange rate at the time of default, while still random, has a distribution that is non-random and time-homogeneous. Thus the locally P_d-expected appreciation fraction

$$\hat{\delta}(\omega, t) = \int_{\mathcal{Z}} \delta(\omega, t, z) K(\omega, t, dz) = \int_{\mathcal{Z}} \delta(z) K(dz)$$

is a constant $\hat{\delta} \in (-1, \infty)$ (in another abuse of notation).

Assumption 4. *We assume that the exchange rate volatility η is given by*

$$\eta(t) = \begin{pmatrix} \rho\sigma_Q \\ \sqrt{1-\rho^2}\sigma_Q \end{pmatrix} \text{ for all } t \in [0, T^*]$$

where $\rho \in (-1, 1)$ and $\sigma_Q > 0$ are constants.

To get a more explicit expression for Q, we define a process W_Q by

$$W_Q(t) = \rho W_1(t) + \sqrt{1-\rho^2} W_2(t).$$

Then W_Q is a Brownian motion that has instantaneous correlation ρ with W_1. Also

$$\begin{aligned} \eta(t) \cdot dW(t) &= \rho\sigma_Q dW_1(t) + \sqrt{1-\rho^2}\sigma_Q dW_2(t) \\ &= \sigma_Q dW_Q(t). \end{aligned}$$

This allows us to write the exchange rate Q as

$$\begin{aligned} \frac{dQ(t)}{Q(t-)} &= (r_d(t) - r_f(t))dt + \sigma_Q dW_Q(t) + \int_{\mathcal{Z}} \delta(t, z)(\mu - \nu)(dt \times dz) \\ &= (r_d(t) - r_f(t))dt + \sigma_Q dW_Q(t) + \int_{\mathcal{Z}} \delta(t, z)\mu(dt \times dz) - \hat{\delta}\lambda^*(t)dt. \end{aligned}$$

This stochastic differential equation has the solution

$$Q(t) = Q(0)\left(1 + I_{\{t \geq \tau\}}\delta(\tau, J(\tau))\right) e^{\int_0^t [r_d(s) - r_f(s) - \hat{\delta}\lambda^*(s)]ds - \frac{1}{2}\sigma_Q^2 t + \sigma_Q W_Q(t)}. \quad (4.1)$$

4.1.2 Intensity Dynamics under the Foreign Measure

The foreign default intensity is again $\tilde{\lambda} = (1 + \hat{\delta})\lambda$, though now $\hat{\delta}$ is constant, so $\tilde{\lambda}$ is a simple scaling of λ.

The P_f-Brownian motion $\widetilde{W}$ is given by

$$\widetilde{W}(t) = W(t) - \int_0^t \eta(s)ds;$$

in our case the two components are given by

$$\begin{aligned}\widetilde{W}_1(t) &= W_1(t) - \rho\sigma_Q t \\ \widetilde{W}_2(t) &= W_2(t) - \sqrt{1-\rho^2}\sigma_Q t.\end{aligned}$$

Thus we can rewrite the dynamics of λ in terms of $\widetilde{W}_1$:

$$\begin{aligned}d\lambda(t) &= \alpha(t)dt + \phi_1(t)dW_1(t) \\ &= \alpha(t)dt + \phi_1(t)d\left(\widetilde{W}_1(t) + \rho\sigma_Q t\right) \\ &= [\alpha(t) + \rho\sigma_Q\phi_1(t)]dt + \phi_1(t)d\widetilde{W}_1(t).\end{aligned}$$

The change of measure increases the drift of λ by $\rho\sigma_Q\phi_1$ while leaving the diffusion coefficient unchanged. The differential of the foreign default intensity is

$$d\tilde{\lambda}(t) = (1+\hat{\delta})[\alpha(t) + \rho\sigma_Q\phi_1(t)]dt + (1+\hat{\delta})\phi_1(t)d\widetilde{W}_1(t).$$

4.1.3 Interest Rates

Assumption 5. *The interest rates r_d and r_f are deterministic functions of time.*

Then the money-market accounts and default-free bond prices are also deterministic, and we have

$$B_d(t,T) = \frac{M_d(t)}{M_d(T)}$$

and similarly for $B_f(t,T)$.

The most useful consequence of this assumption is that the prices of defaultable zero-coupon bonds (with zero recovery) can be decomposed into discount factors and survival probabilities. In particular, we have

$$\begin{aligned}B_d^*(t,T) &= E_d\left[e^{-\int_t^T r_d(s)ds} I_{\{T<\tau\}} \middle| \mathcal{F}_t\right] \\ &= e^{-\int_t^T r_d(s)ds} E_d\left[I_{\{T<\tau\}} | \mathcal{F}_t\right] \\ &= I_{\{t<\tau\}} e^{-\int_t^T r_d(s)ds} E_d\left[e^{-\int_t^T \lambda(s)ds} \middle| \mathcal{F}_t^W\right] \\ &= I_{\{t<\tau\}} B_d(t,T) S_d(t,T)\end{aligned}$$

with the obvious definition. We refer to $S_d(t,T)$ as the domestic probability of survival to time T observed at time t ('on $\{t<\tau\}$' is implied).

Similarly, we have

$$\begin{aligned} B_f^*(t,T) &= I_{\{t<\tau\}} e^{-\int_t^T r_f(s)ds} E_f\left[e^{-\int_t^T \tilde{\lambda}(s)ds} \middle| \mathcal{F}_t^W \right] \\ &= I_{\{t<\tau\}} B_f(t,T) S_f(t,T) \end{aligned}$$

with $S_f(t,T)$ the foreign probability of survival to time T observed at time t.

The domestic instantaneous hazard rate for time T observed at time t is defined to be

$$h_d(t,T) = -\frac{\partial}{\partial T} \log P_d(T>\tau|\mathcal{F}_t) = -\frac{\partial}{\partial T} \log S_d(t,T)$$

on $\{t<\tau\}$, should this limit exist. Note that

$$\begin{aligned} h_d(t,T) &= -\frac{\partial}{\partial T} \log P_d(T>\tau|\mathcal{F}_t) \\ &= -\frac{1}{P_d(T>\tau|\mathcal{F}_t)} \frac{\partial}{\partial T} P_d(T>\tau|\mathcal{F}_t) \\ &= \frac{1}{P_d(T>\tau|\mathcal{F}_t)} \frac{\partial}{\partial T} P_d(\tau \le T|\mathcal{F}_t). \end{aligned}$$

This means that $h_d(t,T)$ is the density of the default time at T conditional on default not having occurred by time T and given the information in $\mathcal{F}_t$. Also, in our models we have $h_d(t,t) = \lambda(t)$.

The domestic average hazard rate to time T, observed at time t, is defined as

$$H_d(t,T) = \frac{1}{T-t}\int_t^T h_d(t,u)du = -\frac{1}{T-t}\log P_d(T>\tau|\mathcal{F}_t) = -\frac{1}{T-t}\log S_d(t,T).$$

(Note that our term 'average hazard rate' is not standard.)

As usual we have an analogy with interest rates: survival probabilities correspond to bond prices, instantaneous hazard rates to instantaneous forward rates, and average hazard rates to zero rates. In particular, we can interpret the average hazard rate as the credit spread at time t on a zero-recovery zero-coupon bond maturing at time T:

$$B_d^*(t,T) = I_{\{t<\tau\}} e^{-[Y(t,T)+H(t,T)](T-t)}$$

where $Y(t,T)$ is default-free zero rate at time t for maturity T.

The foreign instantaneous and average hazard rates h_f and H_f are defined analogously.

We could, of course, increase the dimension of W and relax the assumption of deterministic interest rates. We persist with the more restrictive condition as it simplifies the exposition in the rest of the dissertation, while stochastic rates would cloud the currency-credit interactions that we are studying.

On the other hand, in some cases the extension to stochastic rates is not difficult. If the rates are independent of the other random quantities in the model then we still have the factorization above, and most of the results in the rest of the dissertation still hold. And in any case we can force the factorization by a change of measure: if rates are stochastic we still have

$$B_d^*(t,T) = I_{\{t<\tau\}} B_d(t,T) E_d^{T\text{-forward}} \left[e^{-\int_t^T \lambda(s)ds} \middle| \mathcal{F}_t^W \right]$$

where the expectation is now taken under the domestic T-forward measure. The analogous result holds for the FDZCB price. Depending on the dynamics of the interest rates and the default intensity, the appropriate expectations may or may not be easy to evaluate.

4.2 Pricing Payments at Default in the Basic Model

Based on the factorization above, we see that in order to price defaultable payments it is sufficient to be able to calculate the survival probabilities $S_d(t,T)$ and $S_f(t,T)$. This section shows how we can approximate a payment at the default time using such probabilities; since most of our standard instruments consist of combinations of defaultable payments and payments at default, this means that we need only be able to calculate $S_d(t,T)$ and $S_f(t,T)$ for most pricing purposes.

Let us consider the price of the recovery payment on a domestic defaultable bond. We suppose that the recovery rate (as a proportion of par) R is a constant. Then the price of that recovery payment is

$$I_{\{t<\tau\}} NR \int_t^{\overline{T}} E_d \left[\lambda(u) e^{-\int_t^u [r_d(s)+\lambda(s)]ds} \middle| \mathcal{F}_t^W \right] du$$

(see page 29) where N is the bond's notional amount and $\overline{T}$ is its maturity. This simplifies in our case to

$$I_{\{t<\tau\}} NR \int_t^{\overline{T}} B_d(t,u) E_d \left[\lambda(u) e^{-\int_t^u \lambda(s)ds} \middle| \mathcal{F}_t^W \right] du, \tag{4.2}$$

which is NR times the value of a payment of 1 at the default time.

Now $E_d \left[\lambda(u) e^{-\int_t^u \lambda(s)ds} \middle| \mathcal{F}_t^W \right]$ as a function of u is the conditional density of the default time given $\mathcal{F}_t$ on $\{t < \tau\}$, while $S_d(t,\cdot)$ is the conditional survival function, so we have

$$E_d \left[\lambda(u) e^{-\int_t^u \lambda(s)ds} \middle| \mathcal{F}_t^W \right] du = \frac{d}{du}(1 - S_d(t,u))du = -dS_d(t,u)$$

and we can rewrite this payment as

$$-I_{\{t<\tau\}}NR\int_t^{\overline{T}} B_d(t,u)dS_d(t,u).$$

A simple trapezoidal approximation with grid points

$$t = T_0 < T_1 < \ldots < T_n = \overline{T}$$

gives the value of the payment at default as

$$\begin{aligned}
&- I_{\{t<\tau\}}NR\sum_{i=1}^{n}\frac{1}{2}[B_d(t,T_{i-1}) + B_d(t,T_i)][S_d(t,T_i) - S_d(t,T_{i-1})] \\
&= I_{\{t<\tau\}}\frac{NR}{2}\Bigg\{\sum_{i=0}^{n-1}[B_d(t,T_i) + B_d(t,T_{i+1})]S_d(t,T_i) \\
&\qquad\qquad - \sum_{i=1}^{n}[B_d(t,T_{i-1}) + B_d(t,T_i)]S_d(t,T_i)\Bigg\} \\
&= I_{\{t<\tau\}}\frac{NR}{2}\Bigg\{[1 + B_d(t,T_1)] + \sum_{i=1}^{n-1}[B_d(t,T_{i+1}) - B_d(t,T_{i-1})]S_d(t,T_i) \\
&\qquad\qquad - [B_d(t,T_{n-1}) + B_d(t,T_n)]S_d(t,T_n)\Bigg\}.
\end{aligned}$$

Of course in some cases it is perfectly simple to calculate (4.2) or its foreign equivalent directly – particular examples of this include the Hull-White model of Chapter 5 (both currencies), the Cox-Ingersoll-Ross model of Chapter 6 (domestic currency) and the Alternative CIR model of Chapter 8 (both currencies). The approximation above can be used when such direct calculations are not possible.

Note that by assuming R to be constant, we eliminate any possible dependence between the recovery rate and the appreciation fraction. This dependence may, however, be important: as noted by Ehlers, negative correlation between the recovery rate in the foreign currency and the appreciation fraction is beneficial to the protection buyer in a credit default swap denominated in foreign currency – when recovery is low and he must receive a large protection payment, it will likely have an increased value in local currency, while if he suffers from a less favourable exchange rate, he will likely do so only on a small receipt of foreign currency.

4.3 Construction and a Shortcut Change of Measure

This basic model is easily set up by

1. defining the default intensity λ

2. constructing the default time τ in the canonical way (see, e.g. Bielecki & Rutkowski [2])
3. defining Z as a $\mathcal{Z}$-valued random variable, independent of τ and the Brownian motions, with distribution K
4. defining μ by

$$\mu(A) = \begin{cases} 1 & \text{if } (\tau, Z) \in A \\ 0 & \text{otherwise} \end{cases}$$

for $A \in \mathcal{B}_+ \otimes \mathcal{B}_{\mathcal{Z}}$.
5. defining the exchange rate via (4.1).

Ehlers' framework takes the martingale approach to modelling a default time, in which the intensity λ is considered as a component in the compensator $\int_0^{t\wedge\tau} \lambda(s)ds$ of the default indicator process. The more common hazard rate approach considers the intensity as a component in the conditional probability of survival until a particular time given incomplete market information until that time:

$$P_d(t < \tau|\mathcal{G}_t) = \exp\left\{\int_0^t \lambda(s)ds\right\} \text{ for all } t$$

where $\mathbb{G}$ is some filtration (in our case $\mathbb{F}^W$) relative to which τ is not a stopping time. Bielecki & Rutkowski [2] discuss both approaches in detail. Due to our use of marked point processes and a change of measure that depends upon default, the martingale approach is more fruitful in the general model. However, hazard rate techniques are usually simpler and more elegant than martingale techniques, and the construction above lends itself naturally to using them.

The standard result that for any integrable, $\mathcal{F}_T^W$-measurable random variable Y and $t \in [0, T]$ we have

$$E_d[I_{\{T<\tau\}}Y|\mathcal{F}_t] = I_{\{t<\tau\}}E_d\left[e^{-\int_t^T \lambda(s)ds}Y\middle|\mathcal{F}_t^W\right]$$

follows as usual in the hazard rate approach (though the proof is a slight extension of usual results, since $\mathbb{F}$ contains information about Z as well as the usual information about τ).

We can then use the following shortcut change of measure, which uses only the simplest, Brownian motion-related version of Girsanov's Theorem, to arrive at our usual expression for the FDZCB price. This avoids martingale-based pricing completely and appears to be the most direct route to the FDZCB price, though relative to Ehlers' approach it lacks both finesse and generality.

First, we recognise that for any $t \in [0, T^*]$

$$\begin{aligned} I_{\{t<\tau\}} Q(t) &= I_{\{t<\tau\}} Q(0) e^{\int_0^t [r_d(s) - r_f(s) - \hat{\delta}\lambda(s)]ds - \frac{1}{2}\sigma_Q^2 t + \sigma_Q W_Q(t)} \\ &= I_{\{t<\tau\}} R(t) e^{-\hat{\delta} \int_0^t \lambda(s) ds} \end{aligned}$$

where we define R to be the standard geometric Brownian motion exchange rate

$$R(t) = R(0) e^{\int_0^t [r_d(s) - r_f(s)]ds - \frac{1}{2}\sigma_Q^2 t + \sigma_Q W_Q(t)}$$

with $R(0) = Q(0)$.[1]

Then the process L^F, given by

$$L^F(t) = \frac{R(t) M_f(t)}{R(0) M_d(t)} = e^{-\frac{1}{2}\sigma_Q^2 t + \sigma_Q W_Q(t)}$$

is a local martingale and in fact a true martingale with mean one, since Novikov's condition is satisfied:

$$E_d \left[e^{\frac{1}{2}[\sigma_Q W_Q](T^*)} \right] = e^{\frac{1}{2}\sigma_Q^2 T^*} < \infty.$$

Thus we can define a new probability measure P_F on $(\Omega, \mathcal{F})$ by

$$P_F(A) = \int_A L^F(T^*) dP_d \text{ for } A \in \mathcal{F}.$$

P_F is equivalent to P_d.

Note that $\widetilde{W}_1$, again given by $\widetilde{W}_1(t) = W_1(t) - \rho\sigma_Q t$, is also a P_F-Brownian motion, and so the domestic default intensity again has the differential

$$d\lambda(t) = [\alpha(t) + \rho\sigma_Q \phi_1(t)]dt + \phi_1(t) d\widetilde{W}_1(t).$$

Now the foreign currency price at time t of a unit foreign defaultable zero-coupon bond with zero recovery and maturity T $(0 \le t \le T \le T^*)$ is

$$\begin{aligned} B_f^*(t,T) &= \frac{1}{Q(t)} E_d \left[e^{-\int_t^T r_d(s)ds} I_{\{T<\tau\}} Q(T) \middle| \mathcal{F}_t \right] \\ &= \frac{B_d(t,T)}{Q(t)} E_d \left[I_{\{T<\tau\}} R(T) e^{-\hat{\delta}\int_0^T \lambda(s)ds} \middle| \mathcal{F}_t \right] \\ &= I_{\{t<\tau\}} \frac{B_d(t,T)}{Q(t)} E_d \left[R(T) e^{-\int_t^T \lambda(s)ds - \hat{\delta}\int_0^T \lambda(s)ds} \middle| \mathcal{F}_t^W \right]. \end{aligned}$$

[1]Note that we can and do use λ in place of λ^* in these expressions.

Then, changing measure, we get

$$\begin{aligned}
B_f^*(t,T) &= I_{\{t<\tau\}} \frac{B_d(t,T)}{Q(t)} L^F(t) E_F\left[\frac{R(T)}{L^F(T)} e^{-\int_t^T \lambda(s)ds - \hat{\delta}\int_0^T \lambda(s)ds} \middle| \mathcal{F}_t^W\right] \\
&= I_{\{t<\tau\}} \frac{B_d(t,T)R(t)M_f(t)}{Q(t)R(0)M_d(t)} E_F\left[\frac{R(0)M_d(T)}{M_f(T)} e^{-\int_t^T \lambda(s)ds - \hat{\delta}\int_0^T \lambda(s)ds} \middle| \mathcal{F}_t^W\right] \\
&= I_{\{t<\tau\}} B_f(t,T) \frac{R(t)}{Q(t)} E_F\left[e^{-\int_t^T \lambda(s)ds - \hat{\delta}\int_0^T \lambda(s)ds} \middle| \mathcal{F}_t^W\right] \\
&= I_{\{t<\tau\}} B_f(t,T) E_F\left[e^{-(1+\hat{\delta})\int_t^T \lambda(s)ds} \middle| \mathcal{F}_t^W\right] \\
&= I_{\{t<\tau\}} B_f(t,T) E_F\left[e^{-\int_t^T \tilde{\lambda}(s)ds} \middle| \mathcal{F}_t^W\right].
\end{aligned}$$

Note that while this change of measure gives us very directly the FDZCB price, it lacks interpretability. P_F is not the foreign risk-neutral measure – the foreign currency price of a foreign currency claim is not given by its discounted expectation under P_F because Q and R differ. Also, we can see from Girsanov's Theorem that because L^F does not depend upon μ, the compensator of μ under P_F is still ν and so the intensity of default under P_F is still λ, not $\tilde{\lambda}$.

Chapter 5

Hull-White Model

We can price defaultable bonds by calculating conditional expectations of the form

$$E_d\left[e^{-\int_t^T \lambda(s)ds}\middle|\mathcal{F}_t^W\right] \quad \text{or} \quad E_f\left[e^{-\int_t^T \tilde{\lambda}(s)ds}\middle|\mathcal{F}_t^W\right].$$

These expressions are of the same form as zero-coupon bond prices in short rate models, so it makes sense to reuse standard interest rate models as models of the default intensity.

Our first model uses a single-factor mean-reverting Gaussian process for the default intensity, based on the Hull & White [20] short rate process. This model proves to be extremely tractable, though it allows the default intensity to become negative.

This Hull-White model falls within the class of affine models considered by Ehlers [14] though he does not consider this model in particular. In an unpublished paper Li [33] briefly considers this model and obtains results that agree with our own independent work. Our primary reference for all information about interest rate models is Brigo & Mercurio [6].

5.1 Model Specification

We suppose that the domestic default intensity follows

$$d\lambda(t) = [\theta(t) - a\lambda(t)]dt + \sigma_\lambda dW_1(t)$$

where $\lambda(0)$, a and σ_λ are positive constants and θ is an integrable deterministic function from $[0, T^*]$ to $\mathbb{R}$. Then the default intensity is mean reverting with speed a and time-varying mean reversion level $\frac{\theta(t)}{a}$. The function θ will be chosen to fit the domestic survival probabilities observed in the market, though for the moment we take θ as some given function.

It is possible to have time-dependent a and σ_λ, which allows us to reproduce observed term structures of credit spread volatility. This has the disadvantage of allowing these term structures to become non-stationary, possibly in an unrealistic way, so we avoid it.

By integrating $d(e^{at}\lambda(t))$ one can easily show that for $0 \leq s \leq t$

$$\lambda(t) = \lambda(s)e^{-a(t-s)} + \int_s^t \theta(u)e^{-a(t-u)}du + \sigma_\lambda \int_s^t e^{-a(t-u)}dW_1(u).$$

Thus the distribution of $\lambda(t)$ given $\lambda(s)$, for any $0 \leq s \leq t$, is normal with mean

$$\lambda(s)e^{-a(t-s)} + \int_s^t \theta(u)e^{-a(t-u)}du$$

and variance

$$\frac{\sigma_\lambda^2}{2a}\left[1 - e^{-2a(t-s)}\right].$$

Since λ has normal marginal distributions, this model allows negative default intensities. When we use a similar process for the short rate r in interest rate modelling the same problem arises, but usually short rates have mean values high enough – and volatilities low enough – that the probabilities of negative rates occurring are negligible. In credit risk modelling, we may have low default intensities with high volatilities, leading to significant probabilities of a negative default intensity.[1]

The normal distribution is even more unreasonable in credit risk than in interest rates. While negative interest rates are economically implausible, they are at least possible – if one were quoted an interest rate of -1% one would know how to calculate bond prices. Negative default intensities, on the other hand, are completely nonsensical: it is meaningless to say that a company has a -1% chance of default in the next year.

Since an intensity is necessarily non-negative, our derivations of the prices of defaultable bonds are meaningless in this case and the model does not actually make sense. We still wish to investigate some version of this model, so we use the equations

$$B_d^*(t,T) = I_{\{t<\tau\}}B_d(t,T)E_d\left[e^{-\int_t^T \lambda(s)ds}\middle|\mathcal{F}_t^W\right]$$
$$B_f^*(t,T) = I_{\{t<\tau\}}B_f(t,T)E_f\left[e^{-\int_t^T \tilde{\lambda}(s)ds}\middle|\mathcal{F}_t^W\right]$$

as the definitions of the domestic and foreign defaultable bond prices.

[1] Though we do not focus on credit value adjustments (CVAs) in this dissertation, note that the CVA on a trade may depend strongly on the left (or right) tail of the distribution of the default intensity due to wrong- or right-way risk. If this is the case then using a normally distributed default intensity may result in unreasonable prices.

Of course this introduces arbitrage. Suppose that at some time t we have $\lambda(t) < 0$ and $t < \tau$. Then we can find a maturity $T > t$ such that $S_d(t,T) > 1$, and hence $B_d^*(t,T) > B_d(t,T)$. Then by selling the defaultable bond and investing in the default-free one, we make a risk-free profit.

5.2 Pricing Defaultable Bonds

5.2.1 Domestic Defaultable Bonds

Define the following function:

$$\mathrm{HW}(x,t,T,\theta,a,\sigma_\lambda) = \exp\{-A(t,T) - xC(t,T)\}$$

where

$$\begin{aligned} A(t,T) &= \int_t^T \theta(s)C(s,T)ds \\ &\qquad - \frac{\sigma_\lambda^2}{2a^2}\left\{T - t - \frac{3}{2a} + \frac{2}{a}e^{-a(T-t)} - \frac{1}{2a}e^{-2a(T-t)}\right\} \\ C(t,T) &= \frac{1}{a}\left(1 - e^{-a(T-t)}\right). \end{aligned}$$

Then for $0 \le t \le T$

$$S_d(t,T) = E_d\left[e^{-\int_t^T \lambda(s)ds}\middle|\mathcal{F}_t^W\right] = \mathrm{HW}(\lambda(t),t,T,\theta,a,\sigma_\lambda).$$

Thus the domestic currency price at time t of a domestic defaultable ZCB maturing at time T $(0 \le t \le T \le T^*)$ is

$$B_d^*(t,T) = I_{\{t<\tau\}}B_d(t,T)\mathrm{HW}(\lambda(t),t,T,\theta,a,\sigma_\lambda).$$

5.2.2 Foreign Defaultable Bonds

In terms of the P_f-Brownian motion $\widetilde{W}_1$, the domestic default intensity λ obeys

$$\begin{aligned} d\lambda(t) &= [\theta(t) - a\lambda(t)]dt + \sigma_\lambda d\left[\widetilde{W}_1(t) + \rho\sigma_Q t\right] \\ &= [\theta(t) - a\lambda(t) + \rho\sigma_\lambda\sigma_Q]dt + \sigma_\lambda d\widetilde{W}_1(t) \end{aligned}$$

and so the foreign default intensity $\tilde{\lambda} = (1+\hat{\delta})\lambda$ obeys

$$\begin{aligned} d\tilde{\lambda}(t) &= \left[(1+\hat{\delta})(\theta(t) + \rho\sigma_\lambda\sigma_Q) - a\tilde{\lambda}(t)\right]dt + (1+\hat{\delta})\sigma_\lambda d\widetilde{W}_1(t) \\ &= \left[\tilde{\theta}(t) - a\tilde{\lambda}(t)\right]dt + \tilde{\sigma}_\lambda d\widetilde{W}_1(t) \end{aligned}$$

where we define $\tilde{\theta} = (1+\hat{\delta})(\theta + \rho\sigma_\lambda\sigma_Q)$ and $\tilde{\sigma}_\lambda = (1+\hat{\delta})\sigma_\lambda$.

Note that the form of the default intensity is invariant to the change of measure: the foreign default intensity, under the foreign measure, is still a Hull-White process. The speed of mean reversion a is unchanged. The mean reversion level is shifted by $\rho\sigma_\lambda\sigma_Q$, and then the initial value, the mean reversion level and the absolute volatility are all scaled by $(1+\hat{\delta})$.

Figure 5.1 shows a simulated path of λ, using its P_d-differential, and a simulated path of $\tilde{\lambda}$, using its P_f-differential and the same random numbers. The expected appreciation fraction is set to zero, and the parameters are deliberately exaggerated to show the effect of the the correlation between the exchange rate and the default intensity. A non-zero mean devaluation would simply rescale the path of $\tilde{\lambda}$, as if the two paths were on different axes.

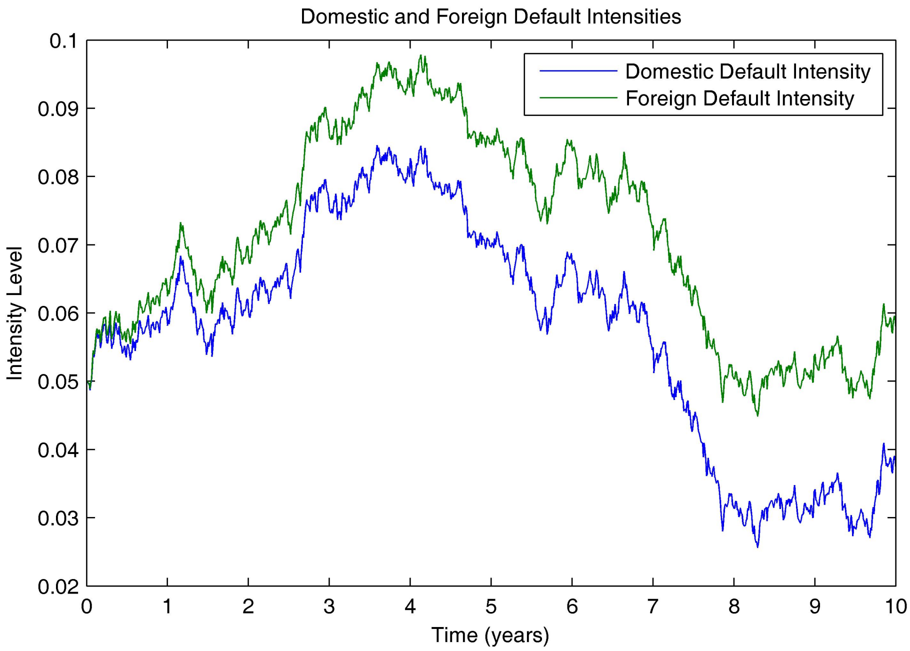

Figure 5.1: Simulated paths of the default intensity λ under the domestic and foreign measures. The mean devaluation is set to zero, so that $\lambda = \tilde{\lambda}$. The same random numbers were used for the two paths. Parameters: $\lambda(0) = 0.05$, $a = 0.2$, $\theta = 0.01$ (constant), $\sigma_\lambda = 0.001$, $\sigma_Q = 5$ and $\rho = 0.95$. Thus the mean reversion levels are $\theta/a = 0.05$ and $\tilde{\theta}/a = 0.07375$.

The foreign survival probabilities are given by

$$S_f(t,T) = E_f\left[e^{-\int_t^T \tilde{\lambda}(s)ds}\middle|\mathcal{F}_t^W\right] = \mathrm{HW}\left(\tilde{\lambda}(t), t, T, \tilde{\theta}, a, \tilde{\sigma}_\lambda\right)$$

and so the foreign currency price at time t of an FDZCB maturing at time T is

$$B_f^*(t,T) = I_{\{t<\tau\}} B_f(t,T) \mathrm{HW}\left(\tilde{\lambda}(t), t, T, \tilde{\theta}, a, \tilde{\sigma}_\lambda\right).$$

5.3 Foreign Survival Probabilities in Terms of Domestic Survival Probabilities

Here we give expressions for the foreign survival probabilities and average hazard rates in terms of their domestic counterparts.

Fix t and T with $0 \le t \le T \le T^*$. As stated above

$$S_d(t,T) = \exp\{-A(t,T) - \lambda(t)C(t,T)\} \tag{5.1}$$

where

$$A(t,T) = \int_t^T \theta(s)C(s,T)ds - \sigma_\lambda^2 J(t,T) \tag{5.2}$$

$$C(t,T) = \frac{1}{a}\left(1 - e^{-a(T-t)}\right) \tag{5.3}$$

$$J(t,T) = \frac{1}{2a^2}\left(T - t - \frac{3}{2a} + \frac{2}{a}e^{-a(T-t)} - \frac{1}{2a}e^{-2a(T-t)}\right). \tag{5.4}$$

We know that

$$S_f(t,T) = \exp\{-A_f(t,T) - \tilde{\lambda}(t)C(t,T)\} \tag{5.5}$$

where

$$A_f(t,T) = \int_t^T \tilde{\theta}(s)C(s,T)ds - \tilde{\sigma}_\lambda^2 J(t,T) \tag{5.6}$$

and C and J are as above (since $\tilde{\lambda}$ has the same speed of mean reversion under P_f as λ has under P_d).

We aim to express $S_f(t,T)$ in terms of $S_d(t,T)$ using (5.1)-(5.6). We calculate that

$$\begin{aligned}
A_f(t,T) &= \int_t^T \tilde{\theta}(s)C(s,T)ds - \tilde{\sigma}_\lambda^2 J(t,T) \\
&= \int_t^T (1+\hat{\delta})(\theta(s) + \rho\sigma_\lambda\sigma_Q)C(s,T)ds - (1+\hat{\delta})^2\sigma_\lambda^2 J(t,T) \\
&= (1+\hat{\delta})\int_t^T \theta(s)C(s,T)ds + (1+\hat{\delta})\rho\sigma_\lambda\sigma_Q\int_t^T C(s,T)ds \\
&\qquad - (1+\hat{\delta})\left[\sigma_\lambda^2 J(t,T) + \hat{\delta}\sigma_\lambda^2 J(t,T)\right] \\
&= (1+\hat{\delta})A(t,T) + (1+\hat{\delta})\frac{\rho\sigma_\lambda\sigma_Q}{a}[T - t - C(t,T)] \\
&\qquad - \hat{\delta}(1+\hat{\delta})\sigma_\lambda^2 J(t,T) \\
&= (1+\hat{\delta})[A(t,T) + G(t,T)]
\end{aligned}$$

where we define

$$G(t,T) = \frac{\rho\sigma_\lambda\sigma_Q}{a}[T - t - C(t,T)] - \hat{\delta}\sigma_\lambda^2 J(t,T).$$

Then

$$\begin{aligned} S_f(t,T) &= \exp\{-A_f(t,T) - \tilde{\lambda}(t)C(t,T)\} \\ &= \exp\{-(1+\hat{\delta})[A(t,T) + G(t,T)] - (1+\hat{\delta})\lambda(t)C(t,T)\} \\ &= \left(S_d(t,T)e^{-G(t,T)}\right)^{1+\hat{\delta}}. \end{aligned} \tag{5.7}$$

From (5.7) it is readily shown that the foreign average hazard rate $H_f(t,T)$ at time t for maturity T is

$$\begin{aligned} H_f(t,T) &= -\frac{1}{T-t}\log S_f(t,T) \\ &= (1+\hat{\delta})\left[-\frac{1}{T-t}\log S_d(t,T) + \frac{1}{T-t}G(t,T)\right] \\ &= (1+\hat{\delta})\left[H_d(t,T) + \frac{1}{T-t}G(t,T)\right] && (5.8) \\ &= (1+\hat{\delta})\left[H_d(t,T) + \rho\sigma_\lambda\sigma_Q\frac{T-t}{2} + o(T-t)\right] && (5.9) \end{aligned}$$

where $H_d(t,T)$ denotes the corresponding domestic average hazard rate and by $o(x)$ we mean a function g such that $\lim_{x\to 0}\frac{g(x)}{x} = 0$.

Expression (5.8) expresses the foreign average hazard rate in terms of the domestic average hazard rate and the model parameters, while expression (5.9) gives us a more intuitive approximation for the short-term foreign average hazard rates.

Figure 5.2 shows average hazard rate curves obtained using the Hull-White model. A non-zero value of $\hat{\delta}$ (roughly) shifts the foreign average hazard curve upwards or downwards relative to the domestic curve. A non-zero correlation results in an increased or decreased slope of the foreign hazard curve – the difference in slopes decreases as time passes. Note that while it appears in the figure that the foreign credit spread for $\hat{\delta} \neq 0$ is just $(1+\hat{\delta})$ times the foreign credit spread with $\hat{\delta} = 0$, this is not strictly true (except in the limit as $T \downarrow t$) because G depends upon $\hat{\delta}$.

5.4 Fitting Market-Implied Domestic Survival Probabilities

Thus far we have allowed θ to be any reasonable function. In practice, we would calibrate the model to the market, choosing θ to match observed time zero domestic

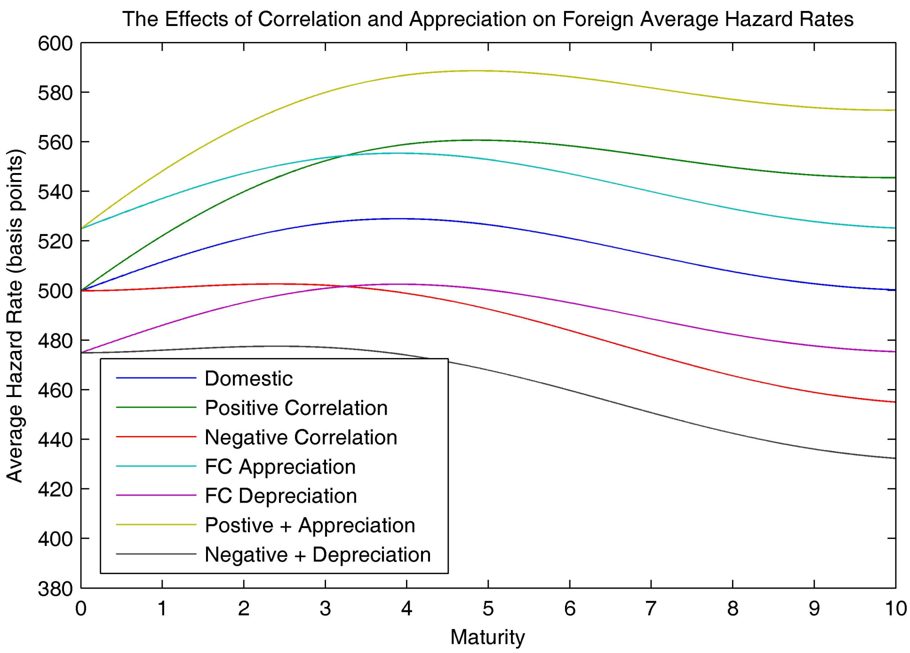

Figure 5.2: Foreign average hazard rate curves for various values of ρ and $\hat{\delta}$. The domestic average hazard rate curve is given by an instantaneous hazard rate of $0.05 + 0.004 \sin(0.6t)$ at time t; we also put $a = 0.4$, $\sigma_\lambda = 0.015$ and $\sigma_Q = 0.2$. The average hazard rate curves are (1) domestic, (2) foreign with $\rho = 0.8$ and $\hat{\delta} = 0$, (3) foreign with $\rho = -0.8$ and $\hat{\delta} = 0$, (3) foreign with $\rho = 0$ and $\hat{\delta} = 0.05$, (4) foreign with $\rho = 0$ and $\hat{\delta} = -0.05$, (5) foreign with $\rho = 0.8$ and $\hat{\delta} = 0.05$, and (6) foreign with $\rho = -0.8$ and $\hat{\delta} = -0.05$. 'FC' stands for 'foreign currency'.

survival probabilities. It is well known that, when a Hull-White process is used in interest rate modelling, pricing expressions can most often be written in terms of the initial term structure without reference to the mean-reversion level. Similarly, we need not extract θ from market-implied domestic survival probabilities – this section shows how to calculate foreign survival probabilities (at any time t) directly from the initial domestic probabilities and implied hazard rates.

We use the following: $S_d^M(T)$ is the market-implied domestic survival probability from time zero to time T, for any $T \in [0, T^*]$. We also define the market-implied domestic instantaneous hazard rate, as seen at time zero for time T, by

$$h_d^M(T) = -\frac{d \log S_d^M(T)}{dT}.$$

To match the initial term structure of DDZCB prices, we must have $S_d(0,T) = S_d^M(T)$ for each $T \in [0, T^*]$. Our model prices will satisfy this condition if and only if

$$\theta(T) = \frac{dh_d^M(T)}{dT} + ah_d^M(T) + \frac{\sigma_\lambda^2}{2a}\left(1 - e^{-2at}\right)$$

for each $T \in [0, T^*]$ (see Hull & White [23]). In this case, the function A in the expression $S_d(t,T) = \exp\{-A(t,T) - \lambda(t)C(t,T)\}$ for the domestic survival probabilities can be rewritten as

$$A(t,T) = \log\frac{S_d^M(t)}{S_d^M(T)} - C(t,T)h_d^M(t) + \frac{\sigma_\lambda^2}{4a}\left(1 - e^{-2at}\right)C^2(t,T).$$

Combining this with the results of the Section 5.3 (in particular equation (5.7)) we have that for $0 \le t \le T \le T^*$

$$\begin{aligned} S_f(t,T) &= \left(S_d(t,T)e^{-G(t,T)}\right)^{1+\hat{\delta}} \\ &= \exp\left\{-(1+\hat{\delta})[A(t,T) + G(t,T) + \lambda(t)C(t,T)]\right\} \end{aligned}$$

where

$$\begin{aligned} C(t,T) &= \frac{1}{a}\left(1 - e^{-a(T-t)}\right) \\ A(t,T) &= \log\frac{S_d^M(t)}{S_d^M(T)} - C(t,T)h_d^M(t) + \frac{\sigma_\lambda^2}{4a}\left(1 - e^{-2at}\right)C^2(t,T) \\ G(t,T) &= \frac{\rho\sigma_\lambda\sigma_Q}{a}[T - t - C(t,T)] - \hat{\delta}\sigma_\lambda^2 J(t,T) \\ J(t,T) &= \frac{1}{2a^2}\left(T - t - \frac{3}{2a} + \frac{2}{a}e^{-a(T-t)} - \frac{1}{2a}e^{-2a(T-t)}\right). \end{aligned}$$

Note that when $t > 0$, the foreign survival probabilities $S_f(t,T)$ depend upon $h_d^M(t)$, so for these probabilities to evolve smoothly with t we must construct the curve $S_d^M(t)$ such that it produces a smooth (or at least continuous) instantaneous hazard rate curve. Such a survival probability curve will also ensure that θ is smooth. This is a similar problem to that of bootstrapping a zero curve that produces a smooth instantaneous forward curve – see Hagan & West [17] for a discussion of the relevant complications and some possible solutions.

Even if we fit the model to arbitrage-free domestic survival probabilities (domestic survival probabilities that do not increase with term: $T \mapsto S_d^M(T)$ is non-increasing) there is no guarantee that the resulting initial foreign survival probabilities will be arbitrage-free for arbitrary ρ and $\hat{\delta}$. (Figure 5.3 gives an example of a curve of foreign survival probabilities that is increasing in term.) In fact, even if we choose ρ and $\hat{\delta}$ to match a given arbitrage-free curve of foreign survival probabilities, it is possible that the resulting model-produced foreign survival probabilities will not be arbitrage-free for long maturities. However, it seems difficult to produce initial foreign survival probability curves that do admit arbitrage with reasonable parameter values, and it is always simple to check whether or not a curve produced by the model is arbitrage-free over a given time period. So while the Hull-White model is flawed in principle, it may produce reasonable results in some applications.

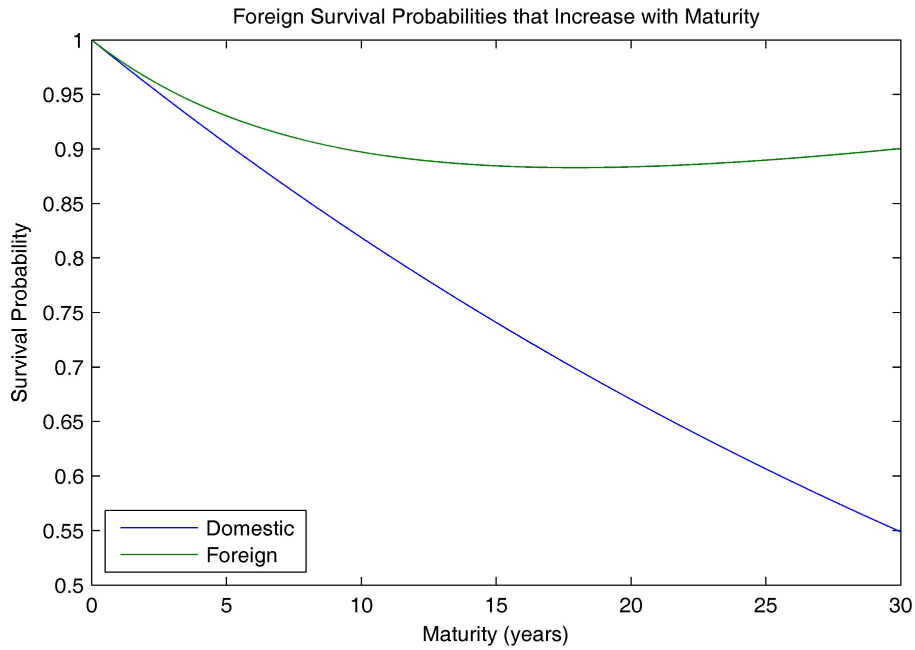

Figure 5.3: Domestic and foreign curves of survival probabilities given by a domestic average hazard rate curve flat at 2%, $a = 10\%$, $\sigma_\lambda = 1\%$, $\sigma_Q = 30\%$, $\rho = -0.8\%$ and $\hat{\delta} = -3\%$. Note how the foreign survival probabilities increase with term.

Chapter 6

Cox-Ingersoll-Ross Model

In this chapter we present a specification of our basic model where the default intensity (possibly shifted by a deterministic function of time) is a square-root process, as in the classical interest rate model of Cox, Ingersoll & Ross [10].

6.1 Model Specification

We suppose that the default intensity λ is given by $\lambda = \varphi + X$ where $\varphi : [0, T^*] \to \mathbb{R}_+$ is a deterministic, integrable function of time, and X is the square root process defined by

$$dX(t) = a[\theta - X(t)]dt + \sigma_\lambda \sqrt{X(t)} dW_1(t),$$

with a, θ and σ_λ positive constants. When $\varphi = 0$ identically, λ is the classical time-homogeneous square-root or Cox-Ingersoll-Ross (CIR) process. We may introduce a non-zero displacement φ in order to fit exactly an observed term structure of DDZCB prices – see Brigo & Mercurio [6].

The process X is non-negative up to an evanescent set, and if the Feller condition $2a\theta > \sigma_\lambda^2$ is satisfied then X is positive up to an evanescent set. Thus the condition $\varphi \geq 0$ is necessary and sufficient to have $\lambda \geq 0$, and having either $\varphi > 0$ or $2a\theta > \sigma_\lambda^2$ guarantees $\lambda > 0$.

The square root process X has non-central χ^2 marginal distributions, and so λ has shifted non-central χ^2 marginal distributions. In particular, for $t \geq 0$, $\lambda(t)$ has density (in y)

$$c f_{v,\delta}\big(c(y - \varphi(t))\big)$$

where

$$\begin{aligned} c &= \frac{4a}{\sigma_\lambda^2(1-e^{-at})} \\ v &= \frac{4a\theta}{\sigma_\lambda^2} \\ \delta &= c(\lambda(0)-\varphi(0))e^{-at} \end{aligned}$$

and $f_{v,\delta}$ is the density of the non-central χ^2 distribution with v degrees of freedom and non-centrality parameter δ (see Brigo & Mercurio [6] or Jeanblanc, Yor & Chesney [28]).

6.2 Pricing Domestic Defaultable Bonds

Define the following function:

$$\mathrm{CIR}(x,t,T,a,\theta,\sigma_\lambda) = A(t,T)e^{-xC(t,T)}$$

where

$$\begin{aligned} A(t,T) &= \left[\frac{2h\exp\{(a+h)(T-t)/2\}}{2h+(a+h)(\exp\{h(T-t)\}-1)}\right]^{\frac{2a\theta}{\sigma_\lambda^2}} \\ C(t,T) &= \frac{2(\exp\{h(T-t)\}-1)}{2h+(a+h)(\exp\{h(T-t)\}-1)} \\ h &= \sqrt{a^2+2\sigma_\lambda^2}. \end{aligned}$$

Fix t and T with $0 \le t \le T \le T^*$. Then

$$E_d\left[e^{-\int_t^T X(s)ds}\middle|\mathcal{F}_t^W\right] = \mathrm{CIR}(X(t),t,T,a,\theta,\sigma_\lambda)$$

and so the domestic survival probability from time t to time T is

$$\begin{aligned} S_d(t,T) &= E_d\left[e^{-\int_t^T \lambda(s)ds}\middle|\mathcal{F}_t^W\right] \\ &= e^{-\int_t^T \varphi(s)ds}E_d\left[e^{-\int_t^T X(s)ds}\middle|\mathcal{F}_t^W\right] \\ &= e^{-\int_t^T \varphi(s)ds}\mathrm{CIR}(X(t),t,T,a,\theta,\sigma_\lambda). \end{aligned}$$

6.3 The Foreign Default Intensity

We can rewrite the dynamics of X in terms of the P_f-Brownian motion $\widetilde{W}_1$ as

$$\begin{aligned} dX(t) &= a[\theta - X(t)]dt + \sigma_\lambda\sqrt{X(t)}d\left[\widetilde{W}_1(t)+\rho\sigma_Q t\right] \\ &= \left[a\theta+\rho\sigma_\lambda\sigma_Q\sqrt{X(t)}-aX(t)\right]dt + \sigma_\lambda\sqrt{X(t)}d\widetilde{W}_1(t). \end{aligned}$$

Then the foreign default intensity is

$$\tilde{\lambda} = (1+\hat{\delta})(\varphi + X) = \tilde{\varphi} + \tilde{X}$$

where $\tilde{\varphi} = (1+\hat{\delta})\varphi$ and the process $\tilde{X} = (1+\hat{\delta})X$ has the dynamics

$$\begin{aligned} d\tilde{X}(t) &= (1+\hat{\delta})\left[a\theta + \rho\sigma_\lambda\sigma_Q\sqrt{X(t)} - aX(t)\right]dt + (1+\hat{\delta})\sigma_\lambda\sqrt{X(t)}d\widetilde{W}_1(t) \\ &= \left[a\tilde{\theta} + \rho\tilde{\sigma}_\lambda\sigma_Q\sqrt{\tilde{X}(t)} - a\tilde{X}(t)\right]dt + \tilde{\sigma}_\lambda\sqrt{\tilde{X}(t)}d\widetilde{W}_1(t) \end{aligned}$$

where we define $\tilde{\theta} = (1+\hat{\delta})\theta$ and $\tilde{\sigma}_\lambda = \sqrt{1+\hat{\delta}}\sigma_\lambda$.

Now the foreign probability of survival from time t to time T $(0 \le t \le T \le T^*)$ is

$$\begin{aligned} S_f(t,T) &= E_f\left[e^{-\int_t^T \tilde{\lambda}(s)ds}\middle|\mathcal{F}_t^W\right] \\ &= e^{-\int_t^T \tilde{\varphi}(s)ds}E_f\left[e^{-\int_t^T \tilde{X}(s)ds}\middle|\mathcal{F}_t^W\right]. \end{aligned}$$

If we can calculate the expectation in the last line we will be able to price defaultable bonds denominated in the foreign currency.

6.4 Square Root Drift Process

The default intensity under the foreign risk-neutral measure is $\tilde{\lambda} = \tilde{X} + \tilde{\varphi}$ where

$$d\tilde{X}(t) = \left[a\tilde{\theta} + \rho\tilde{\sigma}_\lambda\sigma_Q\sqrt{\tilde{X}(t)} - a\tilde{X}(t)\right]dt + \tilde{\sigma}_\lambda\sqrt{\tilde{X}(t)}d\widetilde{W}_1(t).$$

Clearly we need only examine the stochastic part $\tilde{X}$ of $\tilde{\lambda}$: all the properties of $\tilde{\lambda}$ can easily be deduced from those of $\tilde{X}$.

Note firstly that the form of the default intensity is not invariant to the change of measure: $\tilde{X}$ is not a square root process under P_f. (If $\rho = 0$, then $\tilde{X}$ is a standard square root process. The analysis for this case is simple and we omit it.) It is easily shown that a process satisfying this stochastic differential equation does exist.

Since $1+\hat{\delta} > 0$, $\tilde{X}$ and X must have the same sign; thus $\tilde{X}$ is non-negative (or positive if $2a\theta > \sigma_\lambda^2$) up to an evanescent set. (Any P_d-evanescent set is also P_f-evanescent, and vice versa, because the measures are equivalent.)

Consider the drift of $\tilde{X}$ under P_f as a function of the level of the process:

$$f(x) = a\tilde{\theta} + \rho\tilde{\sigma}_\lambda\sigma_Q\sqrt{x} - ax.$$

Then we have

$$\begin{aligned} f'(x) &= \frac{1}{2\sqrt{x}}\rho\tilde{\sigma}_\lambda\sigma_Q - a \\ f''(x) &= -\frac{1}{4x^{3/2}}\rho\tilde{\sigma}_\lambda\sigma_Q. \end{aligned}$$

At $x = 0$, the drift is $a\tilde{\theta} > 0$. The shape of the drift depends upon the sign of the correlation.

For positive ρ, f'' is negative. As $x \downarrow 0$ the slope approaches positive infinity; the drift is increasing for

$$x < x_1 = \left(\frac{\rho\tilde{\sigma}_\lambda\sigma_Q}{2a}\right)^2$$

and decreasing for $x > x_1$. Note that for parameter values that we expect in practice, x_1 should be small. The drift has a zero at

$$x_0 = \tilde{\theta} + \frac{1}{2a^2}\left[[\rho\tilde{\sigma}_\lambda\sigma_Q]^2 + \rho\tilde{\sigma}_\lambda\sigma_Q\sqrt{[\rho\tilde{\sigma}_\lambda\sigma_Q]^2 + 4a^2\tilde{\theta}}\right]. \qquad (6.1)$$

For negative ρ, the drift is convex. As $x \downarrow 0$ the slope approaches negative infinity. The drift is always decreasing, and has a zero at x_0.

Note that $\tilde{X}$ is mean reverting under P_f: its drift is always towards x_0. For $\rho > 0$ the foreign mean reversion level x_0 is above the rescaled domestic mean reversion level $\tilde{\theta}$, and if the process drops below x_1 the mean reversion weakens as the process nears zero. For $\rho < 0$ the foreign mean reversion level is less than $\tilde{\theta}$, and the mean reversion speed increases strongly as the process nears zero.

In Figure 6.1 we plot the drift as a function of the process level with both $\rho > 0$ and $\rho < 0$. We choose $\varphi = 0$ identically and $\hat{\delta} = 0$ so that $\tilde{X} = \tilde{\lambda} = \lambda = X$. Since $\hat{\delta} = 0$, the domestic and foreign drifts coincide at zero. Note that for these parameters the curvature of the foreign drift is not particularly noticeable. When the correlation is positive the drift is increasing near the origin, but the turning point is so close to zero ($x_1 \approx 0.0008$) that this is not obvious.

It seems that reasonably extreme parameter values are required for the curvature in the foreign drift to become noticeable, or for the weakening mean reversion around zero to become significant. We exploit this fact to get approximate closed-form survival probabilities in Section 6.6.

The author has been unable to establish any other facts about this square root drift process or time integrals of it.

6.5 Numerical Evaluation of Foreign Survival Probabilities

In this section and the next we consider approximations to evaluate

$$E_f\left[e^{-\int_t^T \tilde{X}(s)ds}\middle|\mathcal{F}_t^W\right], \qquad (6.2)$$

from which we can calculate foreign survival probabilities. For brevity we will call (6.2) a survival probability; multiplication by $\exp\{-\int_t^T \tilde{\varphi}(s)ds\}$ gives the true probability.

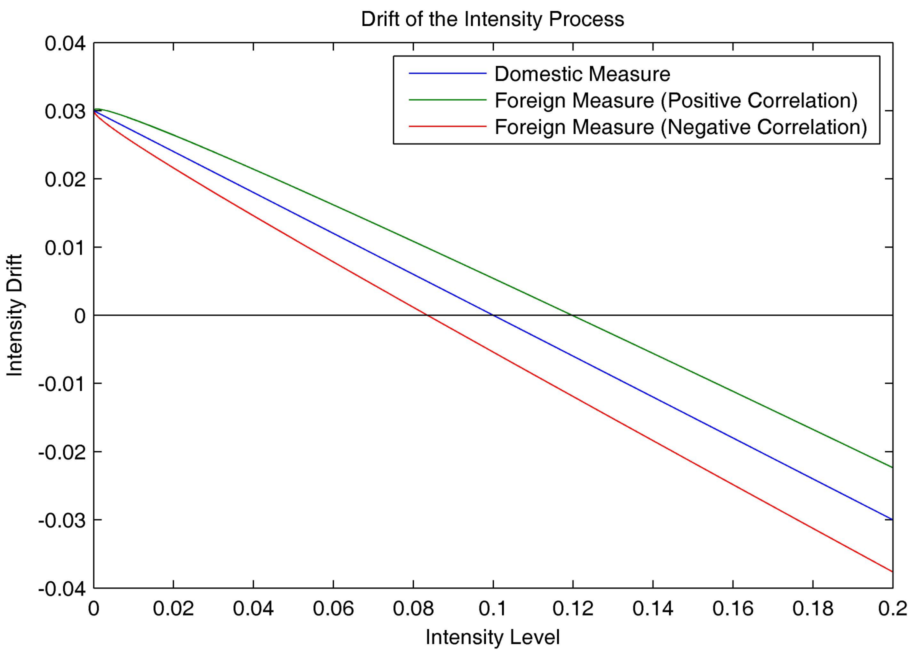

Figure 6.1: The drifts of the intensity process λ under the domestic and foreign risk-neutral measures P_d and P_f, with positive and negative correlation. Parameters: $\varphi = 0$, $\hat{\delta} = 0$, $a = 0.3$, $\theta = 0.1$, $\sigma_\lambda = 0.095$ (corresponding to a log-normal volatility of just over 30% at the domestic mean reversion level θ), $\sigma_Q = 0.3$ and $\rho = \pm 0.6$.

The simplest and most robust method for solving such a one-dimensional problem seems to be using Hull-White trinomial trees [21, 22, 23]. These trees are a form of explicit finite-difference scheme for solving the partial differential equation obeyed by the survival probability. Since the problem is time-homogenous, a faster implicit scheme could be used; we avoid such schemes because to implement them we must introduce artificial (and necessarily incorrect) boundary conditions.

To build a tree, we consider a transformation of $\tilde{X}$ with level-independent volatility. Let

$$Y = \sqrt{\tilde{X}}.$$

Then Y obeys

$$dY(t) = \left[\left(\frac{1}{2} a\tilde{\theta} - \frac{1}{8}\tilde{\sigma}_\lambda^2 \right) \frac{1}{Y(t)} + \frac{1}{2}\rho\tilde{\sigma}_\lambda\sigma_Q - \frac{1}{2}aY(t) \right] dt + \frac{1}{2}\tilde{\sigma}_\lambda d\widetilde{W}_1(t). \tag{6.3}$$

Using a discretisation of this stochastic differential equation, one can build the trinomial tree for Y.

One technical difficulty is the validity of (6.3), which was derived using the Itô formula with the function $x \mapsto \sqrt{x}$ which is only defined on $[0, \infty)$ and has a continuous second derivative only on $(0, \infty)$. If the Feller condition $2a\theta > \sigma_\lambda^2$ is satisfied, then $\tilde{X}$ only takes values in $(0, \infty)$ and so (6.3) appears to be valid.

Both $\tilde{X}$ and Y are always non-negative. Heuristically, $\tilde{X}$ remains non-negative because, as it approaches zero, its volatility approaches zero and its positive drift drags it upwards. The volatility of Y, on the other hand, remains constant as Y approaches zero; Y must remain non-negative because its drift becomes arbitrarily large as Y approaches zero. This will only be the case if

$$2a\theta > \frac{1}{2}\sigma_\lambda^2,$$

which is strictly weaker than the Feller condition. This shows that (6.3) is not valid for arbitrary parameter values. The author is unaware of whether or not (6.3) is valid for

$$\frac{1}{2}\sigma_\lambda^2 < 2a\theta \leq \sigma_\lambda^2.$$

We will assume that we have $2a\theta > \sigma_\lambda^2$.

The remaining difficulty is to ensure that the trinomial tree for Y does not become negative. One possibility is to truncate the tree at zero. We propose another (seemingly new) method.

The usual construction of such a tree (see Brigo & Mercurio [6] or Hull & White [23]) uses nodes at each time t_i with values $y_{i,j} = j\Delta y_i$ where j is allowed to be any integer and Δy_i is some positive quantity determined by the volatility of Y over the previous interval. From a particular node $y_{i,j}$ there is a positive probability of moving to three nodes at the next time step: $y_{i+1,k-1}$, $y_{i+1,k}$ and $y_{i+1,k+1}$ where k is chosen so that $y_{i+1,k}$ is as close as possible to $E_d[Y(t_{i+1})|Y(t_i) = y_{i,j}]$.

We propose that the nodes at an arbitrary time t_i be

$$y_{i,j} = \epsilon_i + j\Delta y_i$$

where the constants (Δy_i) are as before and (ϵ_i) are small positive quantities chosen so that

$$E_d[Y(t_{i+1})|Y(t_i) = y_{i,0} = \epsilon_i] = \epsilon_{i+1} + \Delta y_{i+1}. \tag{6.4}$$

This means that if the lowest possible node at time t_i is ϵ_i then the lowest possible node at time t_{i+1} is ϵ_{i+1}, so the tree nodes are always positive. Equation (6.4) can be solved analytically (we are finding the positive root of a quadratic). An example of the resulting tree structure is given in Figure 6.2.

This method avoids truncating the tree, and avoids having the intensity jump very sharply upwards (skipping multiple nodes) when it gets near zero. This scheme, applied to a Cox-Ingersoll-Ross process, was tested against the closed-form solution; we found that it gave good convergence and no systematic bias. We did not test

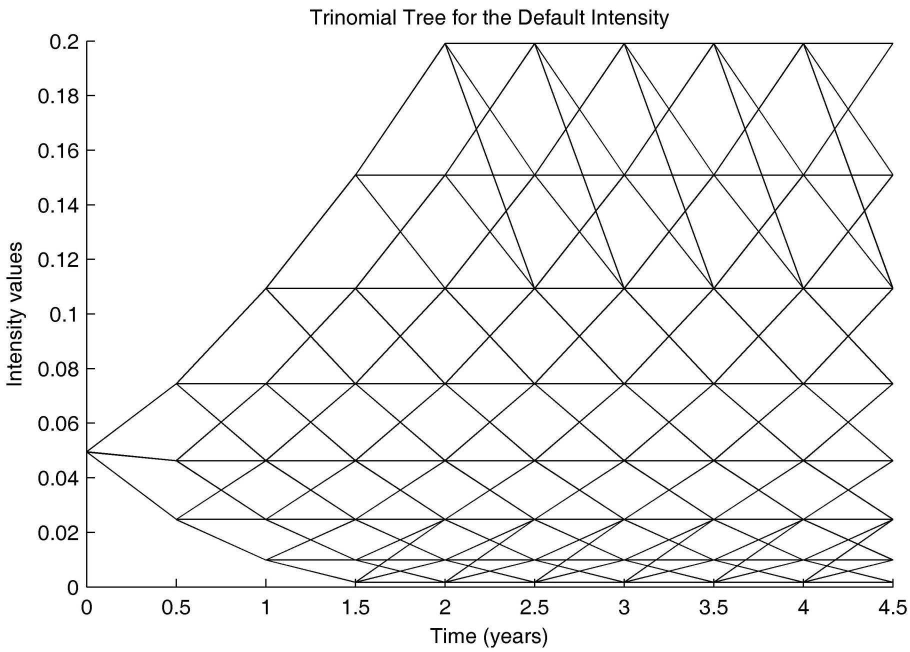

Figure 6.2: Trinomial tree for the foreign default intensity. Parameters: $\lambda(0) = 0.05$, $a = 0.3$, $\theta = 0.03$, $\sigma_\lambda = 0.0949$, $\sigma_Q = 0.2$, $\rho = -0.5$, $\hat{\hat{\delta}} = -0.01$ and time points half a year apart from zero to four and a half years.

this scheme against others proposed in the literature (see Nelson & Ramaswamy [39] and Nawalkha & Beliaeva [38]).

Implementing the trinomial tree is otherwise straightforward.

6.6 Nearest CIR Approximation

In Figure 6.1 we saw that the drift of $\tilde{X}$ sometimes exhibits very little curvature. This suggests that we could replace the true drift

$$f(x) = a\tilde{\theta} + \rho\tilde{\sigma}_\lambda\sigma_Q\sqrt{x} - ax$$

with a linear approximation

$$g(x) = p(q - x).$$

This would give us an approximate Cox-Ingersoll-Ross process for $\tilde{X}$, and hence a closed-form expression for the foreign survival probabilities. Even if this approximation is not precise enough to be used for pricing, it may be useful in calibration. For example, one might have to find $X(0)$, φ, a, θ, σ_λ, ρ and $\hat{\delta}$ so as to fit market-implied survival probabilities (both domestic and foreign) and maybe some CDS option prices. One could find approximate values for these parameters using the closed-form price, and use these values as the starting point for a calibration using the more precise trinomial tree pricing. This should be far less computationally demanding than a calibration using only the trees.

We consider the Taylor expansion of the true drift f about an arbitrary point κ. Then we have, for any x near κ,

$$\begin{aligned}
f(x) &= \sum_{n=0}^{\infty} \frac{1}{n!} f^{(n)}(\kappa)(x-\kappa)^n \\
&\approx f(\kappa) + f'(\kappa)(x-\kappa) \\
&= a\tilde{\theta} + \rho\tilde{\sigma}_\lambda\sigma_Q\sqrt{\kappa} - a\kappa + \left(\frac{1}{2\sqrt{\kappa}}\rho\tilde{\sigma}_\lambda\sigma_Q - a\right)(x-\kappa) \\
&= a\tilde{\theta} + \frac{1}{2}\rho\tilde{\sigma}_\lambda\sigma_Q\sqrt{\kappa} - \left(a - \frac{1}{2\sqrt{\kappa}}\rho\tilde{\sigma}_\lambda\sigma_Q\right)x \\
&= \left(a - \frac{1}{2\sqrt{\kappa}}\rho\tilde{\sigma}_\lambda\sigma_Q\right)\left[\frac{a\tilde{\theta} + \frac{1}{2}\rho\tilde{\sigma}_\lambda\sigma_Q\sqrt{\kappa}}{a - \frac{1}{2\sqrt{\kappa}}\rho\tilde{\sigma}_\lambda\sigma_Q} - x\right] \\
&= p(q-x)
\end{aligned}$$

with the obvious definitions of p and q.

One could use an expansion around the current level $\tilde{X}(t)$ or the foreign mean reversion level x_0 (see (6.1)). For a little more precision, an obvious goal is to use the expansion around the average (over the term of the survival probability considered) of the mean value of $\tilde{X}$ under P_f. One can first expand about, say, $(\tilde{X}(t)+x_0)/2$ to get an approximate CIR speed of mean reversion p_1 and mean reversion level q_1. Using these dynamics, the mean value of $\tilde{X}$ at time s, given $\lambda(t)$ (for $t \leq s$), is

$$E_f\left[\tilde{X}(s)|\mathcal{F}_t^W\right] = \tilde{X}(t)e^{-p_1(s-t)} + q_1\left(1 - e^{-p_1(s-t)}\right).$$

Thus the average over $[t,T]$ of the mean value of $\tilde{X}$ is

$$\frac{1}{T-t}\int_t^T E_f[\tilde{X}(s)|\mathcal{F}_t^W]ds = q_1 + \frac{\tilde{X}(t)-q_1}{p_1(T-t)}\left(1 - e^{-p_1(T-t)}\right).$$

We can then expand about this point, obtaining a new speed of mean reversion p_2 and mean reversion level q_2, and hence a new average mean value. We repeat the procedure until p_n and q_n are constant. Usually three or four iterations are enough.

In any case, once we have calculated the speed of mean reversion p and mean reversion level q, the approximate foreign survival probabilities are given by

$$S_f(t,T) \approx e^{-\int_t^T \tilde{\varphi}(s)ds}\mathrm{CIR}(\tilde{X}(t), t, T, p, q, \tilde{\sigma}_\lambda).$$

This 'Nearest CIR' approximation seems always to overestimate the absolute value of the adjustment – if the foreign average hazard rates are higher (lower) than the domestic rates, then the Nearest CIR foreign rates are higher (lower) than the true foreign rates. This is due to the fact that for positive (negative) ρ, the drift is concave (convex) and so the Taylor expansion overestimates (underestimates) the drift at all levels. The author could find no robust way of correcting this error.[1]

Figure 6.3 plots the foreign five-year average hazard rates obtained using a trinomial tree against the number of steps per year in the tree. The parameters used were: $\lambda(0) = 0.1$, $\varphi = 0$ identically, $a = 0.3$, $\theta = 0.05$, $\sigma_\lambda = 0.09487$ (30% lognormal volatility at initial level), $\sigma_Q = 0.2$, $\rho = -0.5$ and $\hat{\delta} = -0.01$. Also plotted are the domestic average hazard rate, and the Nearest CIR (NCIR) approximate average hazard rates using expansion about $\tilde{\lambda}(0)$ (Current), using expansion about the foreign mean reversion level x_0 (FMR) and using the iterative procedure (Iter). The iterative procedure gives the best performance (an error of less than half a basis point) which seems to be the norm.

[1]The most successful fix the author found was: use the CIR dynamics (from the iterative procedure) to give the 15% and 85% quantiles x_1 and x_2 of the distribution of $\tilde{\lambda}(T/2)$; calculate the values d_1 and d_2 of the true foreign drift at x_1 and x_2; then use the straight line connecting (x_1, d_1) and (x_2, d_2) as the drift in new approximate CIR dynamics. This was not tested extensively, and is sensitive to the percentiles used.

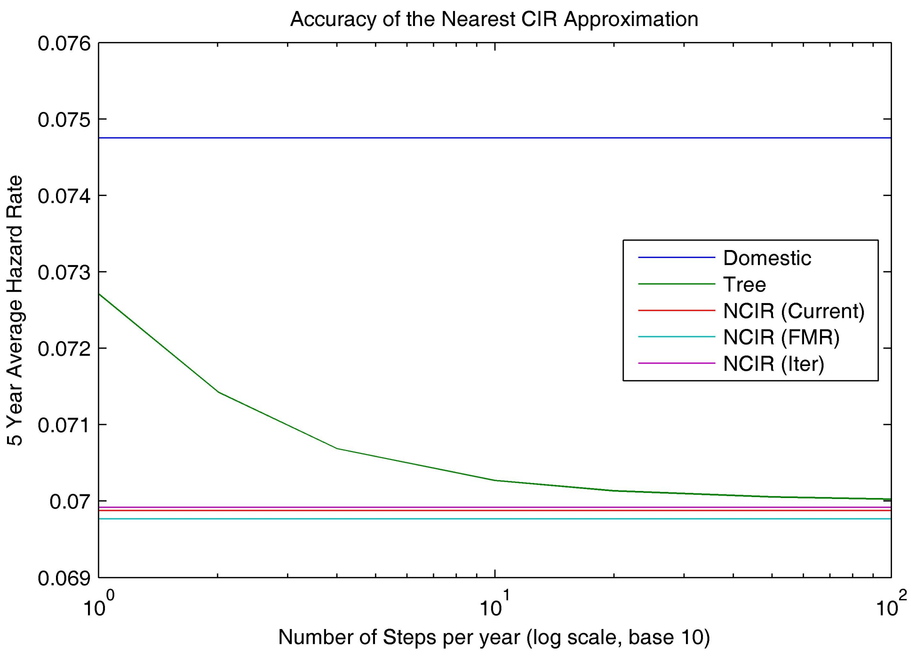

Figure 6.3: Foreign average hazard rates calculated using trinomial trees with various steps lengths, and using various Nearest CIR approximations.

6.7 Affine Models with a Geometric Brownian Motion Exchange Rate

A model for the default intensity λ is said to be affine under a particular measure P if we have, for $0 \leq t \leq T$,

$$E\left[e^{-\int_t^T \lambda(s)ds}\middle|\mathcal{F}_t\right] = e^{A(t,T)+B(t,T)\lambda(t)}$$

where A and B are deterministic functions and E denotes expectation with respect to P. Such models are extremely tractable.

The Cox-Ingersoll-Ross model of this chapter has a positive intensity, and is affine under P_d, but is not affine under P_f. The Hull-White model of Chapter 5, on the other hand, is affine under both P_f and P_d; however it has the flaw of allowing negative default intensities. So it is natural to ask if there are any other models that are affine under both P_f and P_d, in the hope that there may be such a model that also has a positive default intensity. We show that the answer to this question, for time-homogeneous models, is negative.

First, let us suppose that the coefficients of the stochastic differential equation defining λ are time-homogeneous:

$$d\lambda(t) = \alpha(\lambda(t))dt + \phi_1(\lambda(t))dW_1(t).$$

It can be shown (see Björk [3]) that this model is affine if and only if the drift coefficient and the square of the diffusion coefficient are affine (linear) functions:

$$\begin{aligned} \alpha(x) &= a_1 + b_1 x \\ \phi_1^2(x) &= a_2 + b_2 x \end{aligned}$$

for some constants a_1, b_1, a_2 and b_2. In this case the coefficients A and B above are given by a set of ordinary differential equations.

The differential of λ in terms of the P_f-Brownian motion $\widetilde{W}_1$ is

$$\begin{aligned} d\lambda(t) &= [a_1 + b_1\lambda(t)]dt + \sqrt{a_2 + b_2\lambda(t)}d\left[\widetilde{W}_1(t) + \rho\sigma_Q t\right] \\ &= \left[a_1 + b_1\lambda(t) + \rho\sigma_Q\sqrt{a_2 + b_2\lambda(t)}\right]dt + \sqrt{a_2 + b_2\lambda(t)}d\widetilde{W}_1(t) \\ &= \beta(\lambda(t))dt + \psi(\lambda(t))d\widetilde{W}_1(t) \end{aligned}$$

where $\beta(x) = a_1 + b_1 x + \rho\sigma_Q\sqrt{a_2 + b_2 x}$ and $\psi^2(x) = a_2 + b_2 x$. Thus the model is also affine under P_f if and only if $b_2 = 0$, i.e. if the stochastic differential equation for λ is

$$d\lambda(t) = [a_1 + b_1\lambda(t)]dt + \sqrt{a_2}dW_1(t),$$

which is the time-homogenous Hull-White model (the Vasicek model). So the only time-homogeneous version of our basic model that is affine under both measures is time-homogeneous Hull-White.

Now suppose that the coefficients α and ϕ_1 are allowed to depend upon time. If we have

$$\alpha(t, x) = a_1(t) + b_1(t)x \tag{6.5}$$

$$\phi^2(t, x) = a_2(t) + b_2(t)x \tag{6.6}$$

for some deterministic functions of time a_1, b_1, a_2 and b_2, then the model is affine (and the functions A and B are again given by the solutions to a set of ODEs). The converse, however, is not true – there are some time-inhomogeneous affine models where the coefficients do not satisfy (6.5) and (6.6). Thus we cannot repeat our argument above and cannot conclude that the only model which is affine under both measures is the general Hull-White model (with time-varying a and σ_λ).

On the other hand, we expect that any reasonable time-inhomogeneous model will contain as a special case a time-homogeneous version of itself – we should not be forced to include any dependence on time. Thus any model that is affine under both measures, and contains the time-homogeneous version of itself, must have as its time-homogeneous version the Hull-White model.

Note that these results depend upon our choice of geometric Brownian motion as a model for the exchange rate. Chapter 8 presents a model, due to Ehlers [14], where the use of a non-standard process for the exchange rate means that the Cox-Ingersoll-Ross form of the default intensity is unaffected by the change of measure from P_d to P_f, giving a model with a positive default intensity that is affine under both measures.

Chapter 7

Black-Karasinski Model

In this chapter we present a version of our basic model with a lognormally distributed default intensity, using the dynamics suggested for the short rate by Black & Karasinski [4].

Lognormal models are seldom considered for credit risk because the pricing of defaultable bonds requires a numerical procedure; Cox-Ingersoll-Ross-type models are preferred for their tractability. On the other hand, in the multiple-currency setting, the Cox-Ingersoll-Ross model (at least when combined with a geometric Brownian motion exchange rate) also requires a numerical procedure when pricing defaultable bonds denominated in foreign currency. So for our purposes there is no tractability advantage to using the Cox-Ingersoll-Ross model – in fact this Black-Karasinski model requires less computational effort and is more straightforward.

7.1 Model Specification

We assume that the logarithm of the domestic default intensity is a Hull-White process:

$$d \log \lambda(t) = [\theta(t) - a \log \lambda(t)]\, dt + \sigma_\lambda dW_1(t)$$

with $\lambda(0)$, a and σ_λ positive constants and θ a deterministic, integrable function from $[0, T^*]$ to $\mathbb{R}$.[1] Then the log-intensity is Gaussian and mean reverting with speed a and mean reversion level $\theta(t)/a$ at time t. The domestic default intensity itself is positive and obeys

$$d\lambda(t) = \left[\theta(t) + \frac{1}{2}\sigma_\lambda^2 - a \log \lambda(t)\right] \lambda(t)dt + \sigma_\lambda \lambda(t) dW_1(t).$$

[1] We can define a process X by $dX(t) = [\theta(t) - aX(t)]\, dt + \sigma_\lambda dW_1(t)$; such a process has a unique solution that is continuous and adapted to the natural filtration of W_1. We then define $\lambda = \exp X$.

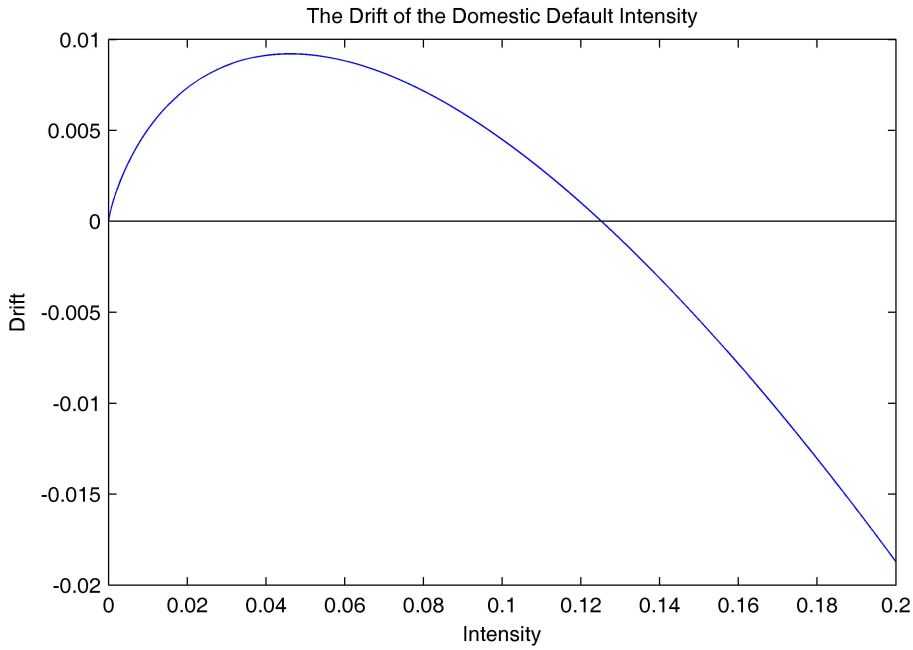

Figure 7.1: The drift of the domestic intensity as a function of the intensity level. Parameters: $a = 0.2$, $\sigma_\lambda = 0.3$ and $\theta = a\log(0.1)$ (so that the exponential of the mean reversion level of $\log\lambda$ is 0.1).

We consider, for some fixed t, the drift of λ as a function $d : (0, \infty) \to \mathbb{R}$ of the process level:

$$d(x) = \left[\theta(t) + \frac{1}{2}\sigma_\lambda^2 - a\log x\right] x.$$

This function is plotted, for particular parameters values, in Figure 7.1. The only zero of d is the mean reversion level of λ, $x_0 := \exp\{(\theta(t) + \sigma_\lambda^2/2)/a\}$. Due to the curvature of the log function, this is different to the exponential of the mean reversion level of $\log\lambda$, $\exp\{\theta(t)/a\}$. The drift is increasing in x for

$$\begin{aligned} x < x_1 &:= \exp\left(\frac{\theta(t) + \frac{1}{2}\sigma_\lambda^2}{a} - 1\right) \\ &= x_0 e^{-1} \end{aligned}$$

and decreasing for $x > x_1$. Its limit as the process level tends to zero is zero. Thus the mean reversion, while never negative, decreases as the process moves below x_1 and can become arbitrarily weak as the process nears zero.

We know that for any s and t with $0 \le s \le t \le T^*$,

$$\log\lambda(t) = \log\lambda(s)e^{-a(t-s)} + \int_s^t \theta(u)e^{-a(t-u)}du + \sigma_\lambda \int_s^t e^{-a(t-u)}dW_1(u).$$

Thus $\log \lambda(t)$, given $\lambda(s)$, has the normal distribution

$$N\left(\log \lambda(s) e^{-a(t-s)} + \int_s^t \theta(u) e^{-a(t-u)} du, \quad \frac{\sigma_\lambda^2}{2a}\left[1 - e^{-2a(t-s)}\right]\right).$$

In particular, λ has lognormal marginal distributions.

A theoretical problem with lognormal models for the short rate is that the expected value of the money-market account at any positive time is infinity. This is usually not a problem in practice since most calculations will use trinomial trees to approximate the dynamics of the short rate, and these trees give finite expectations. In credit risk modelling, when we use a lognormally distributed default intensity, the problem becomes that the price of an account that pays the default intensity has an infinite expected value. Such an account seems unlikely to exist – the defaultable entity is hardly likely to continuously change the interest rate it pays on deposits according to its own changing credit spreads.

If one could trade defaultable bonds with a continuum of maturities then by continuously rolling over an investment in short-dated bonds one could synthetically create this defaultable money-market account, but it seems reasonable to suggest that bonds are available with only finitely many maturities. Thus we should be able to ignore this problem safely.

Domestic and foreign survival probabilities are not known in closed form for the Black-Karasinski model. More pertinently, it seems impossible to give the foreign survival probabilities in terms of the domestic survival probabilities and the model parameters. We will build a trinomial tree to approximate the evolution of λ under P_d, then adjust the tree to give us foreign survival probabilities. The decomposition of λ given in the next section (taken from Brigo & Mercurio [6]) will enable us to build this tree so that it exactly reproduces a given curve of domestic survival probabilities.

7.2 Decomposing the Intensity

Define x as the Ornstein-Uhlenbeck process

$$x(t) = \sigma_\lambda \int_0^t e^{-a(t-u)} dW_1(u).$$

Then $x(0) = 0$ and $dx(t) = -ax(t)dt + \sigma_\lambda dW_1(t)$. We also then have

$$\log \lambda(t) = \alpha(t) + x(t)$$

where we define the function $\alpha : [0, T^*] \to \mathbb{R}$ by

$$\alpha(t) = e^{-at} \log \lambda(0) + \int_0^t \theta(u) e^{-a(t-u)} du.$$

Note that the functions θ and α cannot be calculated analytically from the curve of domestic survival probabilities that they are chosen to fit.

The importance of this decomposition is that we can split $\log\lambda$ into the simple stochastic process x and a deterministic function of time. This means that we can build a trinomial tree for x, then displace the nodes to reproduce a given curve of domestic survival probabilities. Alternatively, we could have tried somehow to calculate θ, then built a tree to approximate $\log\lambda$. Using this decomposition, we circumvent the need to calculate θ. Also, we do the fitting (to the given survival probabilities) after the discretisation of the continuous process to the tree, so that the discretisation does not disturb our fit; of course other errors remain, but our tree will at least reproduce the domestic survival probabilities exactly.

One can easily build a tree for x as explained by Brigo & Mercurio [6]. Let the times at which the nodes occur be $0 = t_0 < t_1 < \ldots < t_n$. We displace each node at each time t_i by a quantity β_i so that the correct survival probability at time t_{i+1} is obtained. (We use β to emphasise that the displacements are discrete approximations to the function α above.)

In particular, we can employ the following procedure (we follow Brigo & Mercurio [6] exactly). Denote the node spacing at time t_i by Δx_i; at the node at time t_i with index j, x has value $j\Delta x_i$ and $\log\lambda$ has value $\beta_i + j\Delta x_i$. Denote by $Q_{i,j}$ the value at time zero of one unit of currency paid at node (i,j) taking into account the probability of survival (but not discounting), and by $S_d^M(T)$ the time zero market-implied domestic probability of survival to time T.

We clearly have $Q_{0,0} = 1$. Since $(0,0)$ is the only node at time t_0, we must have

$$S_d^M(t_1) = \exp\{-\exp\{\beta_0\}(t_1 - t_0)\}$$

so we set

$$\beta_0 = \log\left(-\frac{1}{t_1 - t_0}\log S_d^M(t_1)\right).$$

Now for each $i = 1, 2, \ldots, n-1$, we can calculate $Q_{i,j}$ (for each appropriate j) from β_{i-1} as

$$Q_{i,j} = \sum Q_{i-1,h}\exp\{-\exp\{\beta_{i-1} + h\Delta x_{i-1}\}(t_i - t_{i-1})\}p_i(h,j)$$

where $p_i(h,j)$ is the probability of moving from node $(i-1,h)$ to node (i,j), and the sum is over all h such that this probability is positive. Then the value of β_i is numerically calculated to solve

$$S_d^M(t_{i+1}) = \sum_{\text{all } j} Q_{i,j}\exp\{-\exp\{\beta_i + j\Delta x_i\}(t_{i+1} - t_i)\}.$$

Proceeding in this way, we can calculate β_i for each $i = 0, 1, \dots, n$. (To calculate β_n we can include another time point t_{n+1} – our tree will then reproduce the domestic survival probabilities for times $t_0, t_1, \dots, t_n, t_{n+1}$.)

7.3 Foreign Measure Dynamics

The foreign default intensity is given by $\tilde{\lambda} = (1+\hat{\delta})\lambda$, and so $\log \tilde{\lambda} = \log(1+\hat{\delta}) + \log \lambda$. We can rewrite the dynamics of $\log \lambda$ in terms of the P_f-Brownian motion $\widetilde{W}_1$ as

$$\begin{aligned} d \log \lambda(t) &= [\theta(t) - a \log \lambda(t)]\, dt + \sigma_\lambda d \left[\widetilde{W}_1(t) + \rho \sigma_Q t\right] \\ &= [\theta(t) + \rho \sigma_\lambda \sigma_Q - a \log \lambda(t)]\, dt + \sigma_\lambda d\widetilde{W}_1(t). \end{aligned}$$

Now we define another Ornstein-Uhlenbeck process x_f by

$$x_f(t) = \sigma_\lambda \int_0^t e^{-a(t-u)} d\widetilde{W}_1(u).$$

Then we can write $\log \tilde{\lambda}(t) = \alpha_f(t) + x_f(t)$ where we define $\alpha_f : [0, T^*] \to \mathbb{R}$ by

$$\begin{aligned} \alpha_f(t) &= e^{-at} \log \lambda(0) + \int_0^t (\theta(u) + \rho \sigma_\lambda \sigma_Q)\, e^{-a(t-u)} du + \log(1+\hat{\delta}) \\ &= e^{-at} \log \lambda(0) + \int_0^t \theta(u) e^{-a(t-u)} du + \rho \sigma_\lambda \sigma_Q \int_0^t e^{-a(t-u)} du + \log(1+\hat{\delta}) \\ &= \alpha(t) + \frac{\rho \sigma_\lambda \sigma_Q}{a} \left(1 - e^{-at}\right) + \log(1+\hat{\delta}). \end{aligned}$$

7.4 Attempts at Closed-Form Foreign Survival Probabilities

For any time points t and T with $0 \le t \le T \le T^*$ we have

$$\begin{aligned} S_f(t,T) &= E_f \left[e^{-\int_t^T \tilde{\lambda}(s) ds} \middle| \mathcal{F}_t^W \right] \\ &= E_f \left[\exp\left\{ -(1+\hat{\delta}) \int_t^T e^{\alpha(s) + \frac{\rho \sigma_\lambda \sigma_Q}{a}(1-e^{-as}) + x_f(s)} ds \right\} \middle| \mathcal{F}_t^W \right] \end{aligned} \tag{7.1}$$

Despite the fact that we have

$$E_f \left[\exp\left\{ -\int_t^T e^{x_f(s)} ds \right\} \middle| \mathcal{F}_t^W \right] = E_d \left[\exp\left\{ -\int_t^T e^{x(s)} ds \right\} \middle| \mathcal{F}_t^W \right],$$

the author has been unable to reduce the foreign survival probability to an analytical expression in terms of the domestic survival probabilities and the model parameters. This is essentially due to the fact that the conditional expectation operator and the outer exponential in (7.1) do not commute.

We can use the approximation $e^x \approx 1 + x$ on the outer exponential, and integration by parts, to obtain an expression for the foreign survival probabilities in terms of the domestic survival probabilities. Unfortunately, due to the nature of the approximation, this works well only for high quality debt. It is impossible to add more terms to the approximation, as we will then need to commute the conditional expectation and the squaring operation.

Due to the failure of these two attempts, we resort to the trinomial tree procedure that we have been developing throughout this chapter.

7.5 Adjusting the Trinomial Tree

We now try to use a trinomial tree to approximate the evolution of $\log \tilde{\lambda}$ under P_f. One way to go about this would be to extract θ from the displacements (β_i), insert this function into the dynamics

$$d \log \lambda(t) = [\theta(t) + \rho\sigma_\lambda\sigma_Q - a \log \lambda(t)]\, dt + \sigma_\lambda d\widetilde{W}_1(t)$$

and approximate the evolution of $\log \lambda$ by a trinomial tree.

We use a more direct approach: replacing the domestic displacements (β_i) with foreign displacements (γ_i). The two displacements are related by

$$\alpha_f(t) = \alpha(t) + \frac{\rho\sigma_\lambda\sigma_Q}{a}\left(1 - e^{-at}\right) + \log(1 + \hat{\delta})$$

since (β_i) and (γ_i) are the trinomial tree equivalents of $\alpha(t)$ and $\alpha_f(t)$ respectively.

First, we should construct a tree for x_f with the same times $t_0, \ldots, t_n$. Clearly we will construct exactly the same tree as we constructed for x, so we can simply reuse it. The foreign intensity value at node (i, j) will be $\gamma_i + j\Delta x_i$.

The obvious foreign displacement is

$$\gamma_i = \beta_i + \frac{\rho\sigma_\lambda\sigma_Q}{a}\left(1 - e^{-at_i}\right) + \log(1 + \hat{\delta}) \tag{7.2}$$

for $i = 0, 1, \ldots, n$. However, we find that if we use this adjustment, the trinomial tree underestimates the quanto: as we decrease the step size in the tree, the absolute value of the difference between the foreign and domestic credit spreads increases. This is due to the fact that the absolute correlation adjustment $|\rho\sigma_\lambda\sigma_Q(1 - e^{-at})/a|$ is increasing in t. So while the adjustment in (7.2) is correct at time t_i, it is too small (in absolute value) at any other time in the interval (t_i, t_{i+1}) over which it applies.

A fix for this is to use the average adjustment over (t_i, t_{i+1}):

$$\begin{aligned}\gamma_i &= \beta_i + \frac{\rho\sigma_\lambda\sigma_Q}{a(t_{i+1}-t_i)}\int_{t_i}^{t_{i+1}} \left(1-e^{-as}\right) ds + \log(1+\hat{\delta}) \\ &= \beta_i + \frac{\rho\sigma_\lambda\sigma_Q}{a}\left[1 + \frac{1}{a(t_{i+1}-t_i)}\left(e^{-at_{i+1}} - e^{-at_i}\right)\right] + \log(1+\hat{\delta}). \qquad (7.3)\end{aligned}$$

This ad hoc formula performs quite well. Figures 7.2 and 7.3 show how the basic formula for γ_i underestimates the adjustment, and how the averaged formula corrects this.

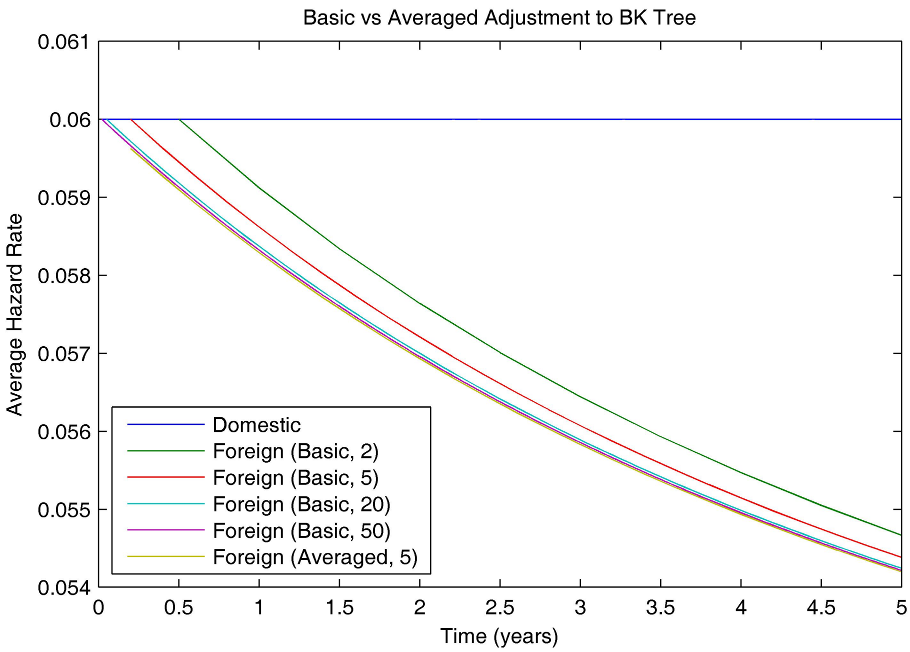

Figure 7.2: Comparing the foreign average hazard rate curves obtained using the basic formula for γ_i (7.2) and the averaged formula (7.3). The constant function is the domestic average hazard rate curve; the others are approximate foreign curves. For each curve the formula used and the number of time steps per year in the trinomial tree are indicated in parentheses. As the number of time steps increases, the basic formula should converge to the true average hazard rate curve. The averaged formula produces similar accuracy with fewer time steps. Parameters: $a = 0.3$, $\sigma_\lambda = 0.4$, $\sigma_Q = 0.2$, $\rho = -0.8$ and $\hat{\delta} = 0$. The domestic hazard curve is flat at 6%.

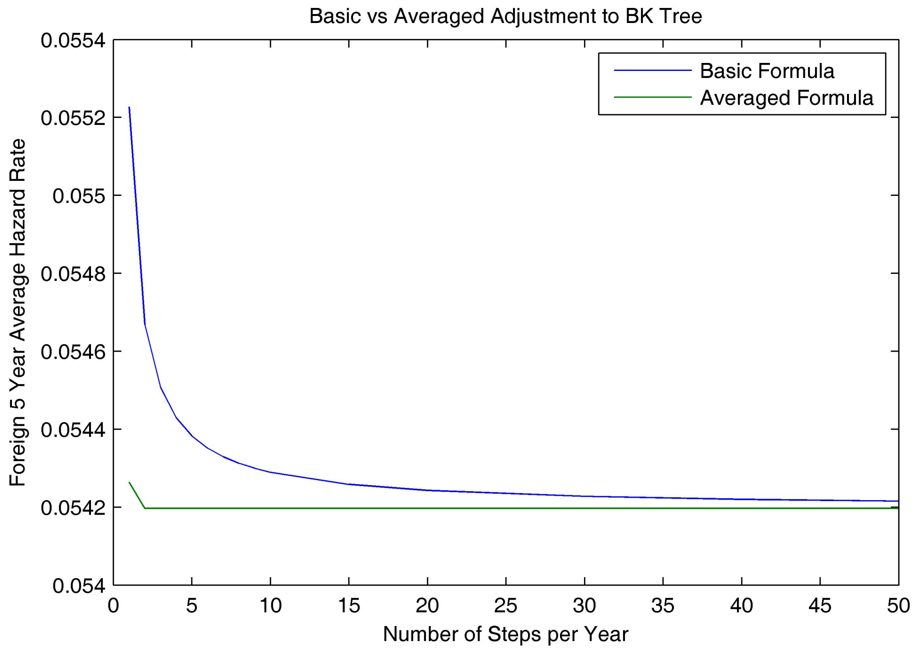

Figure 7.3: Convergence of the foreign five-year average hazard rates (using the basic and averaged formulae) as the number of steps per year in the trinomial tree increases from one to fifty. The same parameters as in Figure 7.2 are used. The 'averaged formula' curve is not actually flat, though no significant increase in accuracy is obtained by using more than a few steps per year in the tree. The difference between the average hazard rates obtained using the basic and averaged formulae decreases to less than two hundredths of a basis point if we increase the number of steps per year to five hundred.

Chapter 8

An Alternative CIR Model

In the chapter we examine a model considered by Ehlers in his doctoral thesis [14]. This model is a special case of his affine jump diffusion model (which we recapitulate in Chapter 9); we consider this particular case here because of its similarities with our basic models in Chapters 5, 6 and 7. The model of this chapter has the advantages of a positive default intensity and closed-form survival probabilities under both the domestic and foreign measures. The price we pay for these advantages is a non-standard model for the exchange rate – one in which the exchange rate volatility depends upon the level of the default intensity.

The model specification in Section 8.1 is taken from [14], though we add the deterministic displacement. Like us, Li [33] writes out the foreign survival probabilities explicitly, though his results appears to suffer from several typographical errors. The discussion of the alternative exchange rate in Section 8.2 is original.

8.1 Model Specification and Pricing

We again suppose that the default intensity λ is given as $\lambda(t) = \varphi(t) + X(t)$ where $\varphi : [0, T^*] \to \mathbb{R}$ is a deterministic, integrable function of time and X is the square root process

$$dX(t) = a[\theta - X(t)]dt + \sigma_\lambda \sqrt{X(t)} dW_1(t)$$

with $X(0)$, a, θ and σ_λ positive constants. Then the domestic survival probabilities are given by

$$S_d(t,T) = e^{-\int_t^T \varphi(s)ds} \mathrm{CIR}(X(t), t, T, a, \theta, \sigma_\lambda)$$

where the function CIR is as defined on page 52.

In contrast to our other basic models, we suppose now that the exchange rate is given by

$$\begin{aligned}\frac{dQ(t)}{Q(t-)} &= (r_d(t) - r_f(t))dt + \gamma_1\sqrt{X(t)}dW_1(t) + \gamma_2 dW_2(t) \\ &\qquad + \int_{\mathcal{Z}} \delta(t,z)(\mu-\nu)(dt\times dz) \\ &= (r_d(t) - r_f(t))dt + \gamma_1\sqrt{X(t)}dW_1(t) + \gamma_2 dW_2(t) \\ &\qquad + \int_{\mathcal{Z}} \delta(t,z)\mu(dt\times dz) - \hat{\delta}\lambda^*(t)dt.\end{aligned}$$

where γ_1 and $\gamma_2 \geq 0$ are constants. Notice that the two Brownian motions appearing in this equation are independent.

Defining the foreign risk-neutral measure P_f as before, we find that the process $\widetilde{W}_1$, defined by

$$\widetilde{W}_1(t) = W_1(t) - \gamma_1 \int_0^t \sqrt{X(s)}ds,$$

is a P_f-Brownian motion. We can rewrite the dynamics of X in terms of $\widetilde{W}_1$:

$$\begin{aligned}dX(t) &= a[\theta - X(t)]dt + \sigma_\lambda\sqrt{X(t)}d\left[\widetilde{W}_1(t) + \gamma_1\int_0^t \sqrt{X(s)}ds\right] \\ &= a[\theta - X(t)]dt + \sigma_\lambda\sqrt{X(t)}d\widetilde{W}_1(t) + \gamma_1\sigma_\lambda X(t)dt \\ &= [a\theta - (a-\gamma_1\sigma_\lambda)X(t)]dt + \sigma_\lambda\sqrt{X(t)}d\widetilde{W}_1(t) \\ &= (a-\gamma_1\sigma_\lambda)\left[\frac{a\theta}{a-\gamma_1\sigma_\lambda} - X(t)\right]dt + \sigma_\lambda\sqrt{X(t)}d\widetilde{W}_1(t).\end{aligned}$$

The foreign default intensity is again

$$\tilde{\lambda} = (1+\hat{\delta})\lambda = (1+\hat{\delta})\varphi + (1+\hat{\delta})X = \tilde{\varphi} + \tilde{X}$$

with the obvious definitions; the dynamics of $\tilde{X}$ is

$$\begin{aligned}d\tilde{X}(t) &= (1+\hat{\delta})(a-\gamma_1\sigma_\lambda)\left[\frac{a\theta}{a-\gamma_1\sigma_\lambda} - X(t)\right]dt + (1+\hat{\delta})\sigma_\lambda\sqrt{X(t)}d\widetilde{W}_1(t) \\ &= (a-\gamma_1\sigma_\lambda)\left[\frac{(1+\hat{\delta})a\theta}{a-\gamma_1\sigma_\lambda} - \tilde{X}(t)\right]dt + \sqrt{(1+\hat{\delta})}\sigma_\lambda\sqrt{\tilde{X}(t)}d\widetilde{W}_1(t) \\ &= \tilde{a}\left[\tilde{\theta} - \tilde{X}(t)\right]dt + \tilde{\sigma}_\lambda\sqrt{\tilde{X}(t)}d\widetilde{W}_1(t)\end{aligned}$$

with the definitions

$$\begin{aligned}\tilde{a} &= a - \gamma_1\sigma_\lambda \\ \tilde{\theta} &= \frac{(1+\hat{\delta})a\theta}{a-\gamma_1\sigma_\lambda} \\ \tilde{\sigma}_\lambda &= \sqrt{(1+\hat{\delta})}\sigma_\lambda.\end{aligned}$$

This means that the foreign survival probabilities are given in closed form by

$$S_f(t,T) = e^{-\int_t^T \tilde{\varphi}(s)ds}\mathrm{CIR}(\tilde{X}(t), t, T, \tilde{a}, \tilde{\theta}, \tilde{\sigma}_\lambda).$$

In many ways this model is ideal: the default intensity is always positive, and we have both domestic and foreign survival probabilities in closed form. The only cause for concern is the exchange rate process.

The next section will compare this alternative exchange rate and the standard geometric Brownian motion. Interestingly, if we consider the former as an approximation of the latter, then we have two 'ways around' the fact that a GBM exchange rate and a CIR domestic default intensity do not afford us closed-form foreign survival probabilities: the first is to approximate the foreign default intensity with a CIR process, as in the Nearest CIR approximation (Section 6.6); the second is to approximate the GBM exchange rate with the alternative exchange rate, resulting in this Alternative CIR model.

8.2 The Alternative Exchange Rate

Geometric Brownian motion, as a model for an exchange rate, has many well-known faults, as documented by many authors including Campa, Chang & Reider [7] and illustrated by volatility skews and smiles in foreign exchange option markets around the world. Nonetheless, it remains a standard model, a first approximation to reality. In this section we compare the alternative exchange rate model and geometric Brownian motion.

We consider devaluation to be impossible, $\delta = 0$ identically. This makes no difference except to simplify the discussion (the jump affects the two exchange rates in exactly the same way). Similarly we suppose that $\varphi = 0$ identically, so that $X = \lambda$. So we are considering the exchange rate model

$$\frac{dQ(t)}{Q(t-)} = (r_d(t) - r_f(t))dt + \gamma_1\sqrt{\lambda(t)}dW_1(t) + \gamma_2 dW_2(t).$$

Clearly the departure from the standard model will be small if γ_1 is small.

First, consider the percentage quanto (the foreign average hazard rate divided by the domestic one) that can be achieved by the model. We use the following parameters, from an example in Brigo & Mercurio [6]: $\lambda(0) = 0.035$, $a = 0.35$, $\theta = 0.045$ and $\sigma_\lambda = 0.15$ (which is equivalent to a lognormal volatility of 71% at the mean reversion level, though the effect of this is diminished by the mean reversion). While this volatility may seem excessive, very high volatilities are observed in CDS options markets and give this model greater flexibility. Lastly, we constrain the 'volatility' of the exchange rate to be 20%, where by 'volatility' we mean the square root of the quadratic variation of $\log Q$ when the intensity is at its domestic mean reversion level:

$$\sqrt{\gamma_1^2\theta + \gamma_2^2} = 20\%.$$

This means that the correlation-style parameter γ_1 is constrained to have an absolute value of at most $0.2/\sqrt{\theta} \approx 0.9428$. In Figure 8.1 we plot the percentage quanto (at the 5-year point) against the possible values of γ_1.

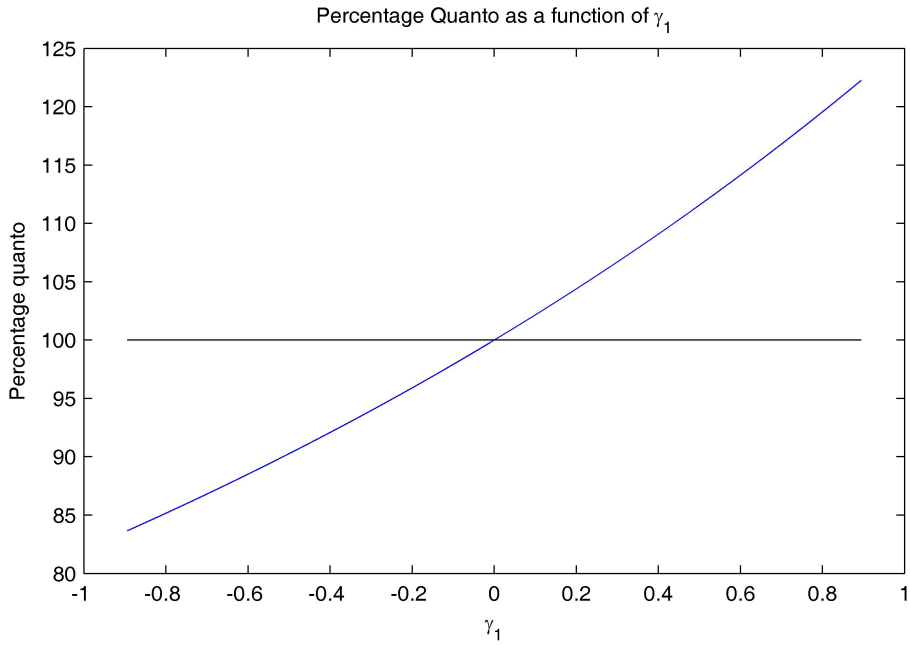

Figure 8.1: Percentage quanto as a function of γ_1.

Clearly we can generate reasonable-sized quantos with this model (at least with such high intensity volatility) even in the absence of a devaluation. The positive adjustments (foreign average hazard rates above domestic ones) are larger than the corresponding negative adjustments. This is due not to the speed of mean reversion increasing with the foreign mean reversion level (the reverse is true) but to the convexity of the foreign mean reversion level as a function of γ_1.

We now look at a particular point on this curve, the 90% quanto. This corresponds to $\gamma_1 \approx -0.5154$. In Figure 8.2 we plot the exchange rate 'volatility' as a function of the intensity level, where by 'volatility' we mean the square root of the quadratic variation of $\log Q$, $\sqrt{\gamma_1^2\lambda(t) + \gamma_2^2}$. Here γ_2 is chosen to give a 20% volatility at the domestic mean reversion level $\theta = 0.045$.

The exchange rate volatility changes roughly linearly from less than 17% to more than 23% as the intensity varies between zero and 10%.

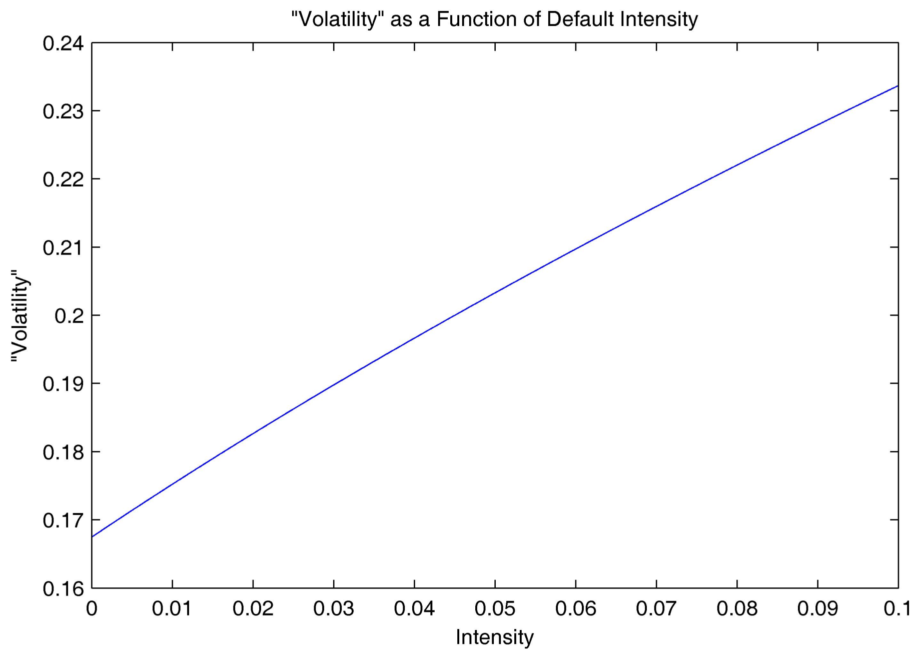

Figure 8.2: Exchange rate 'volatility' as a function of default intensity.

We also look at the departure from the standard model by comparing the alternative and normal probability density functions for the log exchange rate at one year. We simulated the exchange rate using a two-dimensional Euler scheme on the default intensity (with reflection at zero) and the log exchange rate. We created 10 000 samples with a step size of 0.001 years, starting at $Q(0) = 7$, and plotted the empirical density function of the log exchange rate (using Matlab's `ksdensity` function), with the normal density function with the same mean and variance, in Figure 8.3.

There seems to be no great departure from normality for these parameter values. One noticeable difference is that the log exchange rate has a thicker left tail and a thinner right tail in the alternative model than in the GBM model. We confirm this using the QQ-plot in Figure 8.4.

The reason for this is that if $W_1(t)$ has a large positive value, the exchange rate will tend be low (since γ_1 is negative) while the intensity and hence the exchange rate volatility will be high. This gives us the fat left tail. Conversely, if $W_1(t)$ becomes large negative, then the exchange rate will tend to be high, while the intensity and exchange rate volatility will be low, giving us a thin right tail. If this reasoning is correct, we would expect the opposite effect when γ_1 is positive. We confirm this with another QQ-plot in Figure 8.5, using the same parameters except that the sign of γ_1 is reversed.

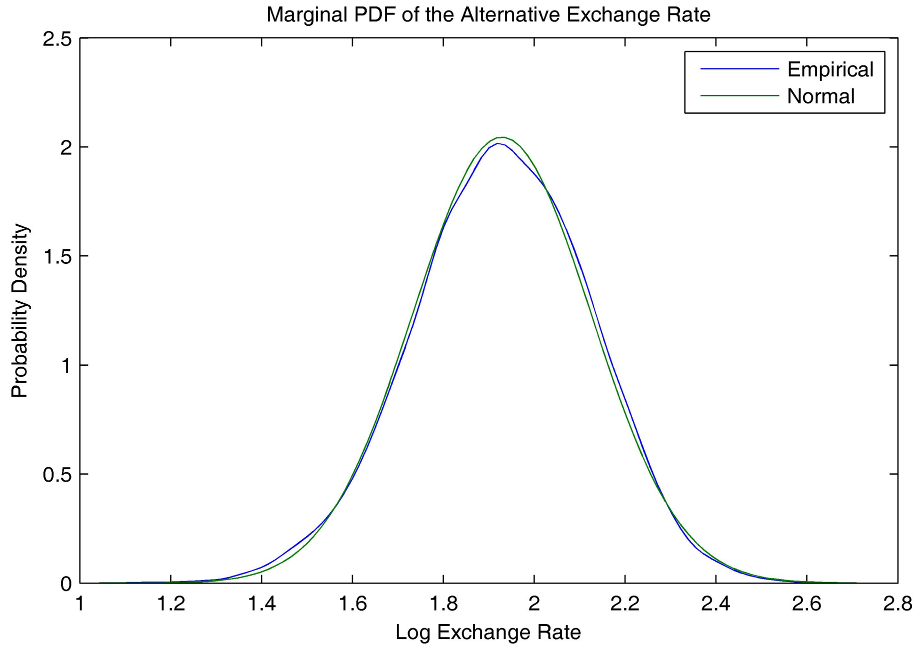

Figure 8.3: The probability density function of the logarithm of the alternative exchange rate and the probability density function of the normal distribution with the same mean and variance.

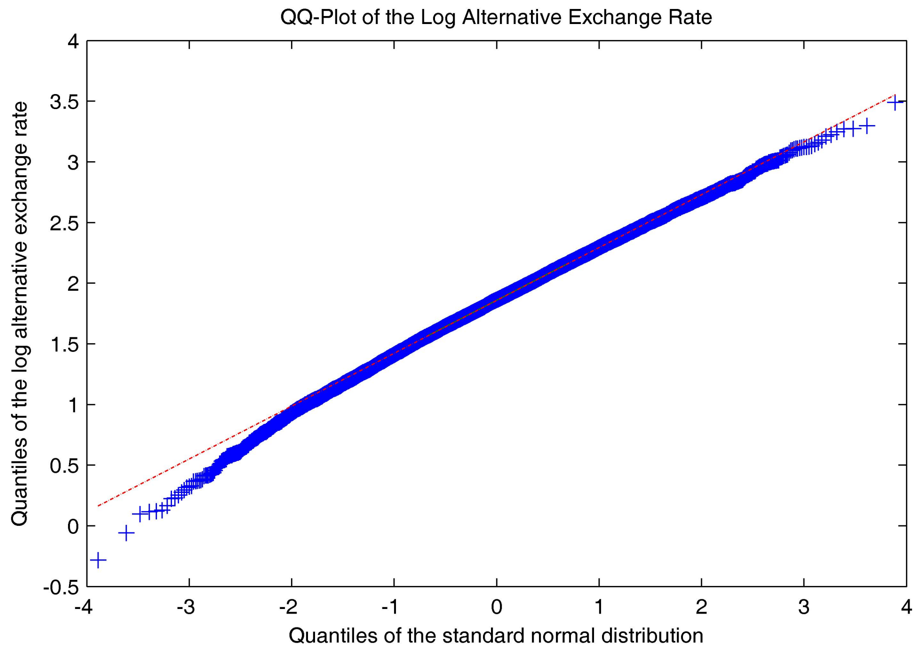

Figure 8.4: The QQ-plot of the log alternative exchange rate with negative γ_1.

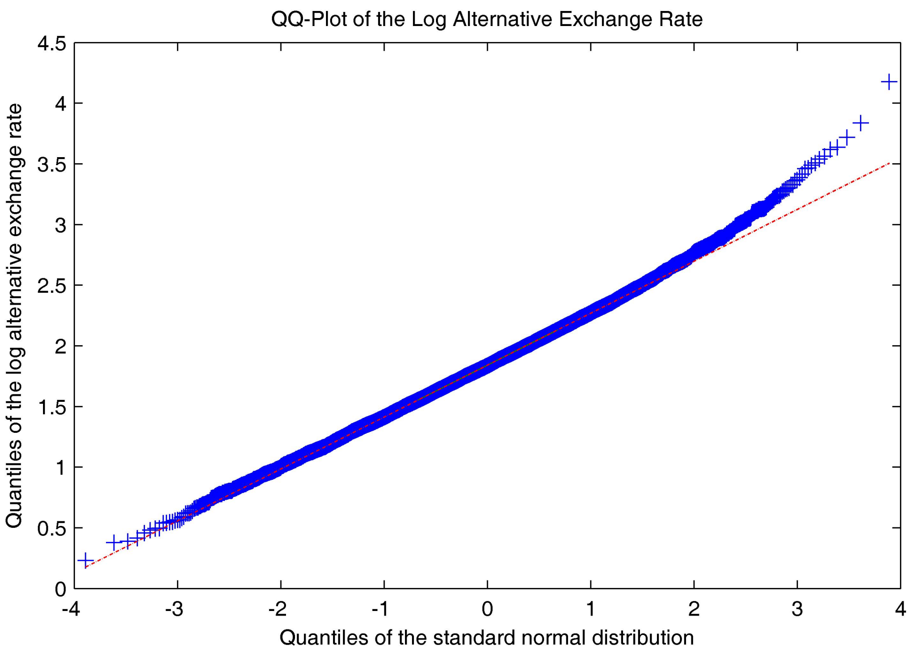

Figure 8.5: The QQ-plot of the log alternative exchange rate with positive γ_1.

Finally, Figure 8.6 shows the volatility smiles generated by the alternative exchange rate with the above parameters (with both positive and negative γ_1). Note that the slope of the smile is determined by the levels of foreign and domestic average hazard rates, and cannot be chosen separately. The volatility smiles are generated by a simple Monte Carlo scheme.

The alternative exchange rate may or may not represent a reasonable model for the exchange rate, depending on the volatility smile exhibited in the foreign exchange market and the quanto desired.

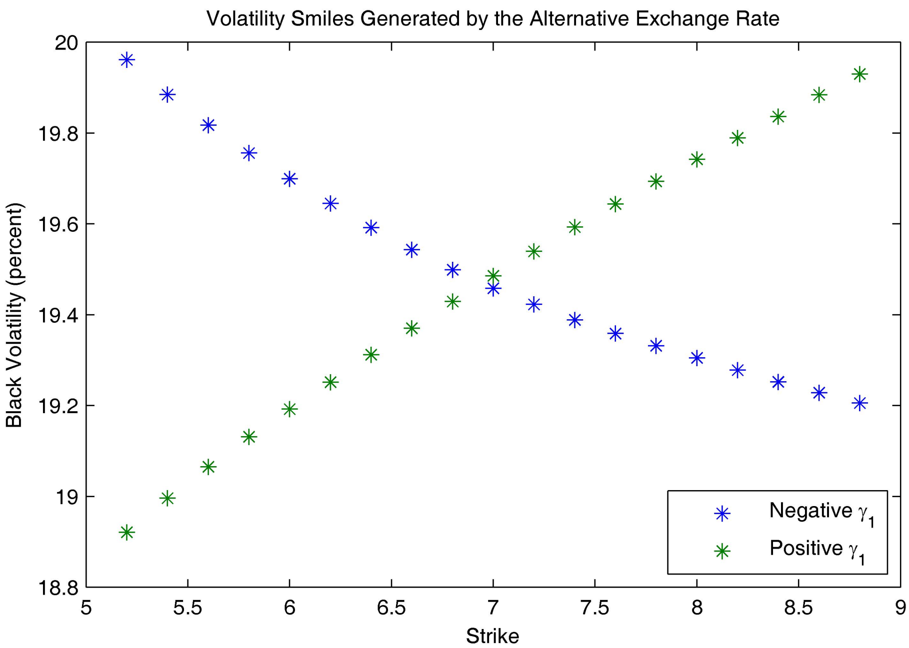

Figure 8.6: Volatility smiles generated by the alternative exchange rate.

Chapter 9

Application of Affine Diffusions

Ehlers [14] examines the use of affine diffusions in his general model. The great advantage of affine diffusions is that all the pricing problems that we will consider are reduced to systems of ordinary differential equations. Thus we can easily use high-dimensional models, including the effects of dependent stochastic interest rates and possibly multifactor models for r_d, r_f and λ. In this chapter we first discuss affine diffusions and the associated pricing advantages following Ehlers, before showing how the Alternative CIR model arises from the use of an affine diffusion, and constructing an affine model with stochastic rates.

9.1 Affine Diffusions

We denote the ith element of a vector x by x_i, and the transpose of a vector or matrix A by A'. A vector- or matrix-valued function g defined on $\mathbb{R}^n$ is said to be affine if, for every $x \in \mathbb{R}^n$, each element of $g(x)$ is given by $a_0 + \sum_{i=1}^{n} a_i x_i$ for some constants (a_i) depending on the element.

An affine diffusion is a continuous n-dimensional ($n \in \mathbb{N}$) solution to a stochastic differential equation

$$dY(t) = \chi(Y(t))dt + \sigma(Y(t))dW(t), \qquad Y(0) = Y_0,$$

where χ and σ are functions from $\mathbb{R}^n$ to $\mathbb{R}^n$ with the property that the drift function χ and the diffusion matrix function $\sigma\sigma' : \mathbb{R}^n \to \mathbb{R}^{n\times n}$ are affine.

Dai & Singleton [11] give a parameterisation of affine diffusions in which one can easily check sufficient conditions for the existence of a unique strong solution to the parameterised stochastic differential equation. This parameterisation also classifies an n-dimensional affine diffusion according to how many of its components are used to determine its diffusion coefficient. The parameterisation does not include all affine diffusions. We reproduce Ehlers' slight generalisation of Dai & Singleton's parameterisation.

An affine diffusion is called admissible if the stochastic differential equation defining it has a unique strong solution. For each $m \in \{0, 1, \ldots, n\}$, we denote by $\mathbb{A}_m(n)$ the set of admissible, n-dimensional affine diffusions where σ depends on m components of Y. Consider the stochastic differential equation

$$dY(t) = [\Theta - \mathcal{K}Y(t)]dt + \sqrt{S(Y(t))}dW(t), \qquad Y(0) = Y_0, \tag{9.1}$$

where Y_0 and Θ are n-dimensional vectors, $\mathcal{K}$ is an n-by-n matrix and S is a function from $\mathbb{R}^n$ to the space of diagonal n-by-n matrices, with each (i, i) component given by

$$S_{ii}(y) = a_i + \sum_{j=1}^{n} b_{ij} y_j$$

for some $a \in \mathbb{R}^n$ and $b \in \mathbb{R}^{n \times n}$. Now suppose that

1. each element of b, and the first m components of Y_0 and Θ, are non-negative.
2. the first m components of a are zero.
3. for i, $j \in \{1, 2, \ldots, m\}$ with $i \neq j$, $\mathcal{K}_{ij}$ is non-positive and b_{ij} is zero.
4. the first m diagonal entries of b are all ones
5. for $k > m$, Θ_k is zero and a_k is either zero or one
6. for $k > m$, the kth column of b is all zeros
7. for $k > m$ and $i \leq m$, $\mathcal{K}_{ik}$ is zero.

Then there is a unique strong solution Y to (9.1). This solution belongs to, and is called a canonical representative of, $\mathbb{A}_m(n)$. The first m components of Y are non-negative, and it is these m components that the diffusion coefficient depends upon. Every element Z of $\mathbb{A}_m(n)$ is given by

$$Z = \eta + \theta Y$$

for some $\eta \in \mathbb{R}^n$ and invertible $\theta \in \mathbb{R}^{n \times n}$.

If we also have

8. $\mathcal{K}$ has positive diagonal elements
9. the first m components of Y_0 are positive
10. the first m components of Θ are greater than $\frac{1}{2}$

11. for each $i \in \{0, 1, \ldots, n\}$, $a_i + \sum_{j=1}^{n} b_{ij}$ is positive,

then the first m components of Y remain positive, Y does not explode on $[0, T^*]$ and the filtration generated by Y coincides with the filtration generated by W. We will suppose that these conditions are satisfied.

9.2 Pricing Results

Let Y be the n-dimensional affine diffusion of the previous section. We define the functions v and w on $\mathbb{R}^n$ by

$$\begin{aligned} v(\beta) &= a'\text{diag}(\beta)\beta = \sum_{i=1}^{n} a_i \beta_i^2 \\ w(\beta) &= b'\text{diag}(\beta)\beta. \end{aligned}$$

Then, for any $\beta \in \mathbb{R}^n$, the quadratic variation of $\beta' Y$ is

$$[\beta' Y](t) = \int_0^t (v(\beta) + w(\beta)' Y(s)) ds.$$

Theorem 9.1 (Lemma 64, Ehlers [14])**.** *Let $\alpha \in \mathbb{R}$ and $\beta \in \mathbb{R}^n$. Fix t and T with $0 \leq t \leq T \leq T^*$. Let $\mathcal{A} : \mathbb{R} \to \mathbb{R}$ and $\mathcal{B} : \mathbb{R} \to \mathbb{R}^n$ be the solutions of the Ricatti ordinary differential equations*

$$\begin{aligned} \frac{d}{dx}\mathcal{A}(x) &= \alpha + \Theta' \mathcal{B}(x) + \frac{1}{2} v(\mathcal{B}(x)) \\ \frac{d}{dx}\mathcal{B}(x) &= \beta - \mathcal{K}' \mathcal{B}(x) + \frac{1}{2} w(\mathcal{B}(x)) \end{aligned}$$

with initial conditions $\mathcal{A}(0) = \mathcal{B}(0) = 0$. Suppose that there exists $\mathcal{B}^* \in \mathbb{R}^n$ such that $||\mathcal{B}(x)|| \leq ||\mathcal{B}^*||$ for all $x \in [0, T]$ and

$$E_d \left[e^{\frac{1}{2} \int_0^T w(\mathcal{B}^*)' Y(t) dt} \right] < \infty.$$

Then we have

$$E_d \left[e^{\int_t^T (\alpha + \beta' Y(s)) ds} \middle| \mathcal{F}_t^W \right] = e^{\mathcal{A}(T-t) + \mathcal{B}(T-t)' Y(t)}.$$

Theorem 9.2 (Lemma 65, Ehlers [14])**.** *Suppose that the conditions of Theorem 9.1 are satisfied. Suppose also that $\zeta \geq 0$ and $\xi \in \mathbb{R}_+^m \times \{0\}^{n-m}$. Let $A : \mathbb{R} \to \mathbb{R}$ and $B : \mathbb{R} \to \mathbb{R}^n$ be the solutions to the ODEs*

$$\begin{aligned} \frac{d}{dx} A(x) &= \Theta' B(x) + \mathcal{B}(x)' \text{diag}(a) B(x) \\ \frac{d}{dx} B(x) &= -\mathcal{K}' B(x) + b' \text{diag}(\mathcal{B}(x)) B(x) \end{aligned}$$

with $A(0) = \zeta$ *and* $B(0) = \xi$. *Suppose that*

$$E_d\left[\exp\left\{\frac{1}{2}\int_0^T w\left(\mathcal{B}^* + \frac{B_*}{A_* + B_*'Y(t)}\right)' Y(t)dt\right\}\right] < \infty$$

where $A_* = \min_{t\in[0,T]} A(t)$ *and* $B_* = \min_{t\in[0,T]} B(t)$. *Then we have*

$$\begin{aligned}&E_d\left[(\zeta + \xi' Y(T))e^{\int_t^T(\alpha+\beta' Y(s))ds}\middle|\mathcal{F}_t^W\right]\\&= (A(T-t) + B(T-t)'Y(t))e^{\mathcal{A}(T-t)+\mathcal{B}(T-t)'Y(t)}.\end{aligned}$$

Note that under the given technical conditions, if $\zeta = \alpha$ and $\xi = \beta$ then $A = d\mathcal{A}/dx$ and $B = d\mathcal{B}/dx$.

To exploit these properties of the affine diffusion Y, we suppose that the domestic and foreign interest rates, and the domestic default intensity, are affine functions of Y:

$$\begin{aligned}r_d &= \alpha_d + \beta_d' Y\\ r_f &= \alpha_f + \beta_f' Y\\ \lambda &= \alpha_\lambda + \beta_\lambda' Y\end{aligned}$$

for some α_d, α_f, $\alpha_\lambda \in \mathbb{R}$ and β_d, β_f, $\beta_\lambda \in \mathbb{R}^n$. Note that we can force any of r_d, r_f or λ to be positive by choosing the corresponding α to be non-negative and the corresponding β to be in $\mathbb{R}_+^m \times \{0\}^{n-m}$.

We also suppose that K and δ are actually deterministic. Then the locally expected devaluation fraction $\hat{\delta}$ is a constant.

Lastly, we suppose that the diffusion coefficient of the exchange rate is given by

$$\eta(t) = \sqrt{S(Y(t))}\overline{\gamma}$$

where $\overline{\gamma} \in \mathbb{R}^n$ and $\sqrt{S(Y(t))}$ is the diffusion coefficient of Y.

Note that $\overline{\gamma}$ determines both the dependence of Q on Y and the instantaneous correlation between Q and λ. The instantaneous covariation between Q and λ is given by

$$d[Q,\lambda](t) = \overline{\gamma}' S(Y(t))\beta_\lambda Q(t-)dt.$$

This specification of σ_Q is chosen so that we have the following result.

Theorem 9.3. *In terms of the* P_f*-Brownian motion* $\widetilde{W}$ *we have*

$$dY(t) = (\tilde{\Theta} - \tilde{\mathcal{K}}Y(t))dt + \sqrt{S(Y(t))}d\widetilde{W}(t) \tag{9.2}$$

where

$$\begin{aligned}\tilde{\Theta} &= \Theta + \mathrm{diag}(\overline{\gamma})a\\ \tilde{\mathcal{K}} &= \mathcal{K} - \mathrm{diag}(\overline{\gamma})b.\end{aligned}$$

In particular, Y is still an affine diffusion under the foreign risk-neutral measure P_f.

Ehlers notes that while a unique strong solution to (9.2) exists and belongs to $\mathbb{A}_m(n)$, Y is not necessarily a canonical representative of $\mathbb{A}_m(n)$. Also, the diagonal elements of $\tilde{\mathcal{K}}$ are not necessarily positive, so we might not have that the first m components of Y are positive under P_f. Hence we need to recheck the parameter restrictions.

Note also that for Theorem 9.3 to be true, it is not required that r_d and r_f be affine in Y. We assumed that they were so that we can apply Theorems 9.1 and 9.2 (reducing expectations to the solution of ODEs) in our pricing problems. If we assume that r_d and r_f are independent of the rest of the model then we need not assume that they are affine in Y.

When we constructed the foreign measure P_f, we assumed that

$$L(t) = \frac{Q(t)M_f(t)}{Q(0)M_d(t)}$$

was a true martingale, rather than just a local martingale. In our affine jump-diffusion framework, this will be the case if the expectation

$$E_d\left[\exp\left\{\frac{1}{2}\int_0^{T^*} w(\overline{\gamma})'Y(t)dt + \int_0^{T^*}\int_{\mathcal{Z}} [(1+\delta(z))\log(1+\delta(z)) - 1]K(dz)\lambda(t)dt\right\}\right]$$

is finite. If the foreign currency can only lose value at default (i.e. if $\delta \leq 0$) then the Novikov condition

$$E_d\left[\exp\left\{\frac{1}{2}\int_0^{T^*} w(\overline{\gamma})'Y(t)dt\right\}\right] < \infty$$

is also sufficient.

9.3 The Alternative CIR Model

In this section we show that a very simple version of this affine diffusion model is the Alternative CIR model of Chapter 8.

Firstly, let the domestic and foreign interest rates r_d and r_f be deterministic functions of time. We let W and Y be two-dimensional, since we will have two random factors – the default intensity and the exchange rate. We would like the default intensity to remain positive, while having as few other restrictions as possible, so we let Y be the canonical representative of $\mathbb{A}_1(2)$ (with the extra restrictions to keep Y_1 positive). Thus we let

$$dY_1(t) = [\Theta_1 - \mathcal{K}_{11}Y_1(t)]dt + \sqrt{Y_1(t)}dW_1(t)$$
$$dY_2(t) = [-\mathcal{K}_{21}Y_1(t) - \mathcal{K}_{22}Y_2(t)]dt + \sqrt{a_2 + b_{21}Y_1(t)}dW_2(t)$$

with $b_{21} \geq 0$, $Y_1(0) > 0$, $\Theta_1 > \frac{1}{2}$, $a_2 \in \{0, 1\}$, $\mathcal{K}_{11}$ and $\mathcal{K}_{22}$ positive, and $a_2 + b_{21} > 0$.

We now define $\lambda(t) = \alpha_\lambda + \beta_\lambda' Y(t)$ for some $\alpha_\lambda \in \mathbb{R}$ and $\beta_\lambda \in \mathbb{R}^2$. Since we want λ to be positive, the second component of β_λ must be zero. For the same reason, we must have $\alpha_\lambda \geq 0$; if we choose also not to bound λ away from zero, we must choose $\alpha_\lambda = 0$. Then the first component of β_λ scales the diffusion coefficient: we pick an arbitrary positive constant σ_λ and let the first component of β_λ be σ_λ^2. Then we have $\lambda = \sigma_\lambda^2 Y_1$ and so

$$\begin{aligned}
d\lambda(t) &= \sigma_\lambda^2[\Theta_1 - \mathcal{K}_{11}Y_1(t)]dt + \sigma_\lambda^2\sqrt{Y_1(t)}dW_1(t) \\
&= \sigma_\lambda^2\left[\Theta_1 - \frac{\mathcal{K}_{11}}{\sigma_\lambda^2}\lambda(t)\right]dt + \sigma_\lambda^2\sqrt{\frac{\lambda(t)}{\sigma_\lambda^2}}dW_1(t) \\
&= [\sigma_\lambda^2\Theta_1 - \mathcal{K}_{11}\lambda(t)]dt + \sigma_\lambda\sqrt{\lambda(t)}dW_1(t) \\
&= \mathcal{K}_{11}\left[\frac{\sigma_\lambda^2\Theta_1}{\mathcal{K}_{11}} - \lambda(t)\right]dt + \sigma_\lambda\sqrt{\lambda(t)}dW_1(t) \\
&= \breve{a}\left[\theta - \lambda(t)\right]dt + \sigma_\lambda\sqrt{\lambda(t)}dW_1(t)
\end{aligned}$$

with the obvious definitions. (We use $\breve{a}$ instead of a to distinguish the intensity's speed of mean reversion from the constants in the diffusion coefficient of Y.) Note that we can pick $\breve{a}$ and σ_λ with absolute freedom (except for the obvious requirement that they be positive). Since Θ_1 is free, we can also choose θ to be any value we please by setting

$$\Theta_1 = \frac{\breve{a}\theta}{\sigma_\lambda^2}$$

provided only that this value is greater than $\frac{1}{2}$, which is exactly the Feller condition $2\breve{a}\theta > \sigma_\lambda^2$.

Now the diffusion coefficient of $\log Q$ is, for some $\overline{\gamma} \in \mathbb{R}^2$,

$$\begin{aligned}
\eta(t) &= \sqrt{S(Y(t))}\overline{\gamma} \\
&= \begin{pmatrix} \sqrt{Y_1(t)} & 0 \\ 0 & \sqrt{a_2 + b_{21}Y_1(t)} \end{pmatrix}\begin{pmatrix} \overline{\gamma}_1 \\ \overline{\gamma}_2 \end{pmatrix} \\
&= \begin{pmatrix} \frac{\overline{\gamma}_1}{\sigma_\lambda}\sqrt{\lambda(t)} \\ \overline{\gamma}_2\sqrt{a_2 + \frac{b_{21}}{\sigma_\lambda^2}\lambda(t)} \end{pmatrix}.
\end{aligned}$$

Since we would like the exchange rate volatility to depend upon the default intensity as little as possible, we choose $b_{21} = 0$ and hence must have $a_2 = 1$. Thus

$$\eta(t) = \begin{pmatrix} \frac{\overline{\gamma}_1}{\sigma_\lambda}\sqrt{\lambda(t)} \\ \overline{\gamma}_2 \end{pmatrix}.$$

Finally, setting $\gamma_1 = \frac{\overline{\gamma}_1}{\sigma_\lambda}$ and $\gamma_2 = \overline{\gamma}_2$ we have retrieved the Alternative CIR model of Chapter 8 (with the deterministic displacement φ equal to zero identically).

9.4 A Model with Stochastic Interest Rates

In this section we look at an extension of the above model, where we have W and Y four-dimensional – the extra two components will be used to drive the domestic and foreign short rates.

For flexibility we choose Y to be the canonical representative of $\mathbb{A}_1(4)$ (with the extra restrictions) – this gives us the minimal number of parameter restrictions while still allowing us to force $\lambda \geq 0$. This unfortunately means that r_d and r_f will not be forced to remain positive, unless we choose them to be linear functions of λ – hopefully, when we use parameter values chosen to match market data, the probabilities of r_d and r_f becoming negative will be negligible.

We choose to have volatilities that depend upon the components of Y as little as possible – we set every element of b that we can to zero. In fact, the only element of b that cannot be set to zero is b_{11}. Then, since we require that $a_i + \sum_{j=1}^n b_{ij} > 0$ for each i, we must have $a_2 = a_3 = a_4 = 1$.

We will also want the drift of each of λ, r_d and r_f to depend only upon that process – the value of λ must not affect the drift of r_d and so on. We will construct λ as a function of Y_1, r_d as a function of Y_1 and Y_2, and r_f as a function of Y_1, Y_2 and Y_3. Thus we can choose to have $\mathcal{K}$ lower-triangular. This forces the drift of Y_1 to depend only upon Y_1, the drift of Y_2 to depend only upon Y_1 and Y_2, and so on.

This leads to the dynamics

$$\begin{aligned}
dY_1(t) &= [\Theta_1 - \mathcal{K}_{11}Y_1(t)]dt + \sqrt{Y_1(t)}dW_1(t) \\
dY_2(t) &= [-\mathcal{K}_{21}Y_1(t) - \mathcal{K}_{22}Y_2(t)]dt + dW_2(t) \\
dY_3(t) &= [-\mathcal{K}_{31}Y_1(t) - \mathcal{K}_{32}Y_2(t) - \mathcal{K}_{33}Y_3(t)]dt + dW_3(t) \\
dY_4(t) &= [-\mathcal{K}_{41}Y_1(t) - \mathcal{K}_{42}Y_2(t) - \mathcal{K}_{43}Y_3(t) - \mathcal{K}_{44}Y_4(t)]dt + dW_4(t)
\end{aligned}$$

with the diagonal elements of $\mathcal{K}$ positive, $Y_1(0) > 0$ and $\Theta_1 > \frac{1}{2}$.

The Default Intensity

Then, similarly to the work above, we choose $\alpha_\lambda = 0$ and $\beta_\lambda = \sigma_\lambda^2(1,0,0,0)'$, giving

$$\begin{aligned} d\lambda(t) &= \mathcal{K}_{11}\left[\frac{\sigma_\lambda^2\Theta_1}{\mathcal{K}_{11}} - \lambda(t)\right]dt + \sigma_\lambda\sqrt{\lambda(t)}dW_1(t) \\ &= \check{a}\left[\theta - \lambda(t)\right]dt + \sigma_\lambda\sqrt{\lambda(t)}dW_1(t). \end{aligned}$$

The Domestic Short Rate

Now we define r_d by

$$r_d = \psi_0 + \psi_1 Y_1 + \psi_2 Y_2$$

for arbitrary constants ψ_0, ψ_1 and ψ_2 (i.e. we let $\alpha_d = \psi_0$ and $\beta_d = (\psi_1, \psi_2)'$). Then

$$\begin{aligned} dr_d(t) &= \psi_1\left\{[\Theta_1 - \mathcal{K}_{11}Y_1(t)]dt + \sqrt{Y_1(t)}dW_1(t)\right\} \\ &\quad + \psi_2\left\{[-\mathcal{K}_{21}Y_1(t) - \mathcal{K}_{22}Y_2(t)]dt + dW_2(t)\right\} \\ &= \left[\psi_1\Theta_1 - (\psi_1\mathcal{K}_{11} + \psi_2\mathcal{K}_{21})Y_1(t) - \psi_2\mathcal{K}_{22}Y_2(t)\right]dt \\ &\quad + \psi_1\sqrt{Y_1(t)}dW_1(t) + \psi_2 dW_2(t) \\ &= \left[\psi_1\Theta_1 - \psi_1\left(\mathcal{K}_{11} + \frac{\psi_2\mathcal{K}_{21}}{\psi_1}\right)Y_1(t) - \psi_2\mathcal{K}_{22}Y_2(t)\right]dt \\ &\quad + \psi_1\sqrt{Y_1(t)}dW_1(t) + \psi_2 dW_2(t). \end{aligned}$$

In order for this drift to be a linear function of r_d, we must set $\mathcal{K}_{21}$ such that

$$\mathcal{K}_{11} + \frac{\psi_2\mathcal{K}_{21}}{\psi_1} = \mathcal{K}_{22}.$$

Also, we remember that $\lambda = \sigma_\lambda^2 Y_1$. Thus we have

$$\begin{aligned} dr_d(t) &= \left[\psi_1\Theta_1 - \mathcal{K}_{22}(r_d(t) - \psi_0)\right]dt + \frac{\psi_1}{\sigma_\lambda}\sqrt{\lambda(t)}dW_1(t) + \psi_2 dW_2(t) \\ &= \mathcal{K}_{22}\left[\left(\frac{\psi_1\Theta_1}{\mathcal{K}_{22}} + \psi_0\right) - r_d(t)\right]dt + \frac{\psi_1}{\sigma_\lambda}\sqrt{\lambda(t)}dW_1(t) + \psi_2 dW_2(t). \end{aligned}$$

Since ψ_0 and $\mathcal{K}_{22}$ are free, we can set the mean reversion level and speed of mean reversion to any values we please. Unfortunately we have only one parameter, ψ_1, to control both the correlation between λ and r_d and the dependence of the volatility of r_d on λ.

The Foreign Short Rate

We define the foreign short rate by

$$r_f = \phi_0 + \phi_1 Y_1 + \phi_2 Y_2 + \phi_3 Y_3$$

for some constants ϕ_0, ϕ_1, ϕ_2 and ϕ_3. Then we have

$$\begin{aligned} dr_f(t) &= \phi_1 \left\{[\Theta_1 - \mathcal{K}_{11}Y_1(t)]dt + \sqrt{Y_1(t)}dW_1(t)\right\} \\ &\quad + \phi_2 \{[-\mathcal{K}_{21}Y_1(t) - \mathcal{K}_{22}Y_2(t)]dt + dW_2(t)\} \\ &\quad + \phi_3 \{[-\mathcal{K}_{31}Y_1(t) - \mathcal{K}_{32}Y_2(t) - \mathcal{K}_{33}Y_3(t)]dt + dW_3(t)\} \\ &= \Big[\phi_1\Theta_1 - [\phi_1\mathcal{K}_{11} + \phi_2\mathcal{K}_{21} + \phi_3\mathcal{K}_{31}]Y_1(t) - [\phi_2\mathcal{K}_{22} + \phi_3\mathcal{K}_{32}]Y_2(t) \\ &\quad - \phi_3\mathcal{K}_{33}Y_3(t)\Big]dt + \phi_1\sqrt{Y_1(t)}dW_1(t) + \phi_2 dW_2(t) + \phi_3 dW_3(t). \end{aligned}$$

So if we choose the parameters $\mathcal{K}_{31}$ and $\mathcal{K}_{32}$ so that

$$\mathcal{K}_{33} = \mathcal{K}_{22} + \frac{\phi_3\mathcal{K}_{32}}{\phi_2} = \mathcal{K}_{11} + \frac{\phi_2\mathcal{K}_{21} + \phi_3\mathcal{K}_{31}}{\phi_1}$$

then we have

$$\begin{aligned} dr_f(t) &= \Big[\phi_1\Theta_1 - \mathcal{K}_{33}[\phi_1 Y_1(t) + \phi_2 Y_2(t) + \phi_3 Y_3(t)]\Big]dt + \phi_1\sqrt{Y_1(t)}dW_1(t) \\ &\qquad + \phi_2 dW_2(t) + \phi_3 dW_3(t) \\ &= \Big[\phi_1\Theta_1 - \mathcal{K}_{33}[r_f - \phi_0]\Big]dt + \phi_1\sqrt{Y_1(t)}dW_1(t) + \phi_2 dW_2(t) + \phi_3 dW_3(t) \\ &= \mathcal{K}_{33}\left[\left(\frac{\phi_1\Theta_1}{\mathcal{K}_{33}} + \phi_0\right) - r_f\right]dt + \frac{\phi_1}{\sigma_\lambda}\sqrt{\lambda(t)}dW_1(t) + \phi_2 dW_2(t) + \phi_3 dW_3(t) \end{aligned}$$

which is again mean-reverting. The speed of mean reversion $\mathcal{K}_{33}$ and the mean reversion level $\frac{\phi_1\Theta_1}{\mathcal{K}_{33}} + \phi_0$ are both arbitrary, since $\mathcal{K}_{33}$ and ϕ_0 are free. Again, allowing any dependence between r_f and λ (i.e. allowing $\phi_1 \neq 0$) results in the volatility of r_f depending upon the level of λ.

The Exchange Rate

The exchange rate volatility is given by

$$\begin{aligned} \eta(t) &= \sqrt{S(Y(t))}\overline{\gamma} \\ &= \begin{pmatrix} \overline{\gamma}_1\sqrt{Y_1(t)} \\ \overline{\gamma}_2 \\ \overline{\gamma}_3 \\ \overline{\gamma}_4 \end{pmatrix} \end{aligned}$$

and thus the exchange rate is given by

$$\frac{dQ(t)}{Q(t-)} = dt \text{ terms } + \frac{\overline{\gamma}_1}{\sigma_\lambda}\sqrt{\lambda(t)}dW_1(t) + \overline{\gamma}_2 dW_2(t) + \overline{\gamma}_3 dW_3(t) + \overline{\gamma}_4 dW_4(t).$$

Pricing

Pricing of our usual default-free and defaultable zero-coupon bonds is easily done using Theorems 9.1 and 9.2.

Comments

This model shares the main flaw of the related Alternative CIR model: the exchange rate volatility depends upon the level of the default intensity in a way that is fully determined by the domestic and foreign survival probabilities. Also the interest rate volatilities depend upon the default intensity if we wish to have correlation between λ, r_d and r_f. Thus in this model the volatilities of the exchange rate and the interest rates are not easily interpreted, and cannot be found separately from one another.

Chapter 10

Currency Options and Sophisticated Models for the Exchange Rate

This short chapter considers two issues. The first is the use of models for the exchange rate that are more sophisticated than geometric Brownian motion: in particular stochastic and local volatility models. The second is the pricing of vulnerable currency options – options on one or the other currency written by a default-risky agent. These two issues are linked by the usefulness in both cases of a change from the domestic risk-neutral measure to a domestic survival measure.

For this chapter we return from Ehlers' general model and its affine specifications, and work with a slight generalisation of the basic model of Chapter 4.

10.1 The Problems

10.1.1 Local and Stochastic Volatility

In most of our models we can easily find DDZCB prices – indeed they are a model input. The key to using a model for pricing the credit risk in the foreign currency is being able to calculate FDZCB prices:

$$B_f^*(t,T) = I_{\{t<\tau\}} \frac{B_d(t,T)}{Q(t)} E_d \left[e^{-\int_t^T \lambda(s)ds} Q(T) \Big| \mathcal{F}_t^W \right].$$

The expectation in this expression involves both the default intensity λ and the exchange rate Q.

By constructing the foreign risk-neutral measure P_f, we were able to reduce this expectation to

$$B_f^*(t,T) = I_{\{t<\tau\}} B_f(t,T) E_f \left[e^{-\int_t^T \tilde{\lambda}(s)ds} \middle| \mathcal{F}_t^W \right],$$

an expectation involving only $\tilde{\lambda}$. This reduction in dimensionality allows us to price FDZCBs, if not in closed form, then at least by use of a tree with only one space dimension.

We suppose in this section that the exchange rate is given by

$$\frac{dQ(t)}{Q(t-)} = (r_d(t) - r_f(t))dt + \sigma_Q(t)dW_Q(t) + \int_{\mathcal{Z}} \delta(t,z)(\mu - \nu)(dt \times dz)$$

where, in another abuse of notation, the exchange rate volatility σ_Q is some $\mathbb{F}^W$-adapted stochastic process. Examples include local volatility models, where the volatility is a function of time and the exchange rate itself, and stochastic volatility models such as those proposed by Heston [18] and Hagan et al. [16] (the SABR model).

Then, by a simple adjustment of our earlier work, the process

$$\widetilde{W}_1(t) = W_1(t) - \rho \int_0^t \sigma_Q(s)ds$$

is a P_f-Brownian motion. This means that if

$$d\lambda(t) = \alpha(t)dt + \phi_1(t)dW_1(t)$$

then in terms of $\widetilde{W}_1$ we have

$$\begin{aligned} d\lambda(t) &= \alpha(t)dt + \phi_1(t)d\left[\widetilde{W}_1(t) + \rho \int_0^t \sigma_Q(s)ds\right] \\ &= [\alpha(t) + \rho\phi_1(t)\sigma_Q(t)]dt + \phi_1(t)d\widetilde{W}_1(t). \end{aligned}$$

The dynamics of λ thus depend upon the exchange rate volatility, and we are still dealing with a multi-dimensional problem: a tree to approximate the evolution of λ under the foreign measure will have two space dimensions – one for λ itself and another for σ_Q.

10.1.2 Pricing Defaultable Options

We would like to price defaultable options. We will generally restrict ourselves to European, non-path-dependent options: securities that pay off some function of the exchange rate at their maturity if default has not yet occurred. For example, we would like to be able to value a defaultable European call option on the foreign currency with strike K, which pays off

$$I_{\{T<\tau\}}(Q(T) - K)^+$$

units of domestic currency at its maturity T. (A similar payoff is evaluated in calculating the credit value adjustment on a foreign exchange forward.) Note that we assume zero recovery in default.

The domestic currency price of such an option at time $t \in [0, T]$ is

$$\text{DC Call Price} = B_d(t,T)E_d\left[I_{\{T<\tau\}}(Q(T)-K)^+|\mathcal{F}_t\right].$$

Changing to the foreign risk-neutral measure we obtain

$$\begin{aligned}
&\text{DC Call Price}\\
&= B_d(t,T)\frac{Q(t)M_f(t)}{Q(0)M_d(t)}E_f\left[\frac{Q(0)M_d(T)}{Q(T)M_f(T)}I_{\{T<\tau\}}(Q(T)-K)^+\Big|\mathcal{F}_t\right]\\
&= B_f(t,T)Q(t)E_f\left[I_{\{T<\tau\}}\left(1-\frac{K}{Q(T)}\right)^+\Big|\mathcal{F}_t\right].
\end{aligned}$$

Clearly this change of measure has been of no use.

10.2 The Domestic Survival Measure

This section presents a measure that may be more useful than the foreign risk-neutral measure in the situations considered above, at least when combined with a Hull-White default intensity process: the domestic T-survival measure.

10.2.1 Definition and Use

We suppose that the exchange rate volatility σ_Q is stochastic, and consider the pricing of a general European, non-path-dependent, defaultable claim that pays $g(Q(T))$ at its maturity T if default has not occurred by then, where g is some deterministic function. We also suppose that δ is zero identically, so $g(Q(T))$ is $\mathcal{F}_t^W$-measurable.

Define a process L^T by

$$L^T(t) = \frac{E_d\left[e^{-\int_0^T \lambda(s)ds}\Big|\mathcal{F}_t^W\right]}{E_d\left[e^{-\int_0^T \lambda(s)ds}\right]} = \frac{e^{-\int_0^t \lambda(s)ds}E_d\left[e^{-\int_t^T \lambda(s)ds}\Big|\mathcal{F}_t^W\right]}{E_d\left[e^{-\int_0^T \lambda(s)ds}\right]}.$$

Then L^T is a martingale and $E_d[L^T(t)] = 1$ for all $t \in [0, T^*]$.

We define the domestic T-survival measure P_d^T by

$$P_d^T(A) = \int_A L^T(T^*)dP_d \text{ for all } A \in \mathcal{F}.$$

(This measure should not be confused with the T-forward measure, often denoted in the literature by P^T or a similar expression.) Note that L^T is the density process with respect to $\mathbb{F}^W$:

$$L^T(t) = \frac{dP_d^T}{dP_d}\bigg|_{\mathcal{F}_t^W}.$$

Then the domestic currency price at time $t \in [0, T]$ of our defaultable claim is

$$\begin{aligned}
&B_d(t,T)E_d\left[I_{\{T<\tau\}}g(Q(T))|\mathcal{F}_t\right] \\
&= I_{\{t<\tau\}}B_d(t,T)E_d\left[e^{-\int_t^T \lambda(s)ds}g(Q(T))\Big|\mathcal{F}_t^W\right] \\
&= I_{\{t<\tau\}}B_d(t,T)L^T(t)E_d^T\left[\frac{1}{L^T(T)}e^{-\int_t^T \lambda(s)ds}g(Q(T))\Big|\mathcal{F}_t^W\right] \\
&= I_{\{t<\tau\}}B_d(t,T)E_d\left[e^{-\int_t^T \lambda(s)ds}\Big|\mathcal{F}_t^W\right]E_d^T\left[g(Q(T))\Big|\mathcal{F}_t^W\right] \\
&= I_{\{t<\tau\}}B_d(t,T)S_d(t,T)E_d^T\left[g(Q(T))\Big|\mathcal{F}_t^W\right] \\
&= B_d^*(t,T)E_d^T\left[g(Q(T))\Big|\mathcal{F}_t^W\right]
\end{aligned}$$

with E_d^T denoting expectation under P_d^T. So if we can calculate the expected value of $g(Q(T))$ under P_d^T, we will be able to price our claim.

Note that the density process $L^T(t)$ depends upon λ, and so the differential of the exchange rate under P_d^T depends upon the process that we choose for λ.

10.2.2 Hull-White Model

In our Hull-White model we have

$$\lambda(s) = \lambda(0)e^{-as} + \int_0^s \theta(u)e^{-a(s-u)}du + \sigma_\lambda \int_0^s e^{-a(s-u)}dW_1(u)$$

and hence

$$\begin{aligned}
\int_0^T \lambda(s)ds &= \text{Constants} + \sigma_\lambda \int_0^T \int_0^s e^{-a(s-u)}dW_1(u)ds \\
&= \text{Constants} + \sigma_\lambda \int_0^T \int_u^T e^{-a(s-u)}dsdW_1(u) \\
&= \text{Constants} + \sigma_\lambda \int_0^T C(u,T)dW_1(u)
\end{aligned}$$

where

$$C(u,T) = \frac{1}{a}\left(1 - e^{-a(T-u)}\right).$$

Thus

$$
\begin{aligned}
L^T(T^*) &= \frac{e^{-\int_0^T \lambda(s)ds}}{E_d\left[e^{-\int_0^T \lambda(s)ds}\right]} \\
&= e^{\text{Other Constants}-\sigma_\lambda \int_0^T C(u,T)dW_1(u)} \\
&= e^{\text{Other Constants}-\int_0^{T^*} Y(u)dW_1(u)}
\end{aligned}
$$

where

$$Y(u) = I_{\{u\leq T\}}\sigma_\lambda C(u,T).$$

Thus the process W_3, defined by

$$W_3(t) = W_1(t) + \int_0^t Y(u)du = W_1(t) + \sigma_\lambda \int_0^{t\wedge T} C(u,T)du,$$

is a P_d^T-Brownian motion, as is W_2. Now, for any $t \in [0,T]$ we have

$$
\begin{aligned}
W_Q(t) &= \rho W_1(t) + \sqrt{1-\rho^2}W_2(t) \\
&= \rho\left[W_3(t) - \sigma_\lambda \int_0^t C(u,T)du\right] + \sqrt{1-\rho^2}W_2(t) \\
&= \rho W_3(t) + \sqrt{1-\rho^2}W_2(t) - \rho\sigma_\lambda \int_0^t C(u,T)du \\
&= W_4(t) - \rho\sigma_\lambda \int_0^t C(u,T)du
\end{aligned}
$$

where W_4, defined in the obvious way, is a P_d^T-Brownian motion.

Then, for $t \in [0,T]$, the dynamics of Q can be rewritten as

$$
\begin{aligned}
\frac{dQ(t)}{Q(t-)} &= [r_d(t) - r_f(t)]dt + \sigma_Q(t)dW_Q(t) \\
&= [r_d(t) - r_f(t)]dt + \sigma_Q(t)d\left[W_4(t) - \rho\sigma_\lambda \int_0^t C(u,T)du\right] \\
&= [r_d(t) - r_f(t) - \rho\sigma_\lambda C(t,T)\sigma_Q(t)]dt + \sigma_Q(t)dW_4(t).
\end{aligned}
$$

In particular, the dynamics of Q do not depend upon λ, which may enable us to evaluate the defaultable option price efficiently.

If we use a local volatility model for Q ($\sigma_Q(t) = v(t,Q(t))$ for some function $v : [0,T^*]\times\mathbb{R}_+ \to \mathbb{R}_+$) then the expectation $E_d^T[g(Q(T))|\mathcal{F}_t^W]$ is given by $f(t,Q(t))$ where f satisfies

$$f_t(t,x) + [r_d(t) - r_f(t) - \rho\sigma_\lambda C(t,T)v(t,x)]xf_x(t,x) + \frac{1}{2}v^2(t,x)x^2 f_{xx}(t,x) = 0,$$

where subscripts denote partial derivatives, and the terminal condition $f(T,x) = g(x)$.

If the exchange rate volatility $\sigma_Q(t)$ is in fact a constant σ_Q and the option considered is a vanilla call or put option,

$$g(x) = (\omega(x - K))^+$$

where $K \in \mathbb{R}_+$ and $\omega = 1$ (for a call) or -1 (for a put), then

$$E_d^T[g(Q(T))|\mathcal{F}_t^W] = \text{Black}(\omega, Q^T(t), K, T - t, \sigma_Q)$$

where

$$\begin{aligned} Q^T(t) &= Q(t)e^{\int_t^T [r_d(s) - r_f(s) - \rho\sigma_\lambda\sigma_Q C(s,T)]ds} \\ &= Q(t)\frac{B_f(t,T)}{B_d(t,T)}e^{-\frac{\rho\sigma_\lambda\sigma_Q}{a}[T-t-C(t,T)]} \end{aligned}$$

and Black$(\omega, F, K, \tau, \sigma)$ is the usual Black formula with put-call indicator ω, current forward F, strike K, time to expiry τ and volatility σ.

The same sort of reduction does not seem possible for our other models. In particular, for both the Cox-Ingersoll-Ross models and the Black-Karasinski model, we do not have explicit expressions for $\int_0^T \lambda(s)ds$ (and hence L^T) that can be used to infer the differential of Q in terms of P_d^T-Brownian motions. Even if we did, it seems likely that such expressions would involve integrals with respect to dW_1 of integrands involving λ, which would lead to the P_d^T-differential of Q depending on λ.

Chapter 11

Hedging in the Basic Model

This chapter considers the hedging of a position in foreign defaultable bonds using positions in domestic defaultable bonds and foreign default-free bonds. Here by 'bonds' we mean zero-coupon bonds that have zero recovery if they are default-risky. We work within the basic model of Chapter 4.

There are two possible objectives in hedging default-sensitive instruments (either the simple bonds above or real bonds and credit default swaps). The first is to be hedged against the default event: we refer to this as being default-neutral. Alternatively, one can aim to be hedged against movements in the default intensity (and hence credit spreads), which we refer to as being spread-neutral. Spread-neutrality is often assigned more importance when marking a trading book to market – the risk of loss due to a (hopefully rare) default is sometimes ignored.

The obvious instruments to use for hedging a short position in a T-maturity FDZCB are the T-maturity DDZCB and the T-maturity foreign default-free bond. These two instruments are used to hedge against movements in three random factors: the exchange rate, the default intensity and the default time. The only way that this will be possible is if default- and spread-neutrality are equivalent, in the sense that the position in the T-maturity DDZCB required to maintain default-neutrality is the same as the position required to maintain spread-neutrality. If default- and spread-neutrality are not equivalent, then a satisfactory hedge will have to involve DDZCBs with two different maturities.

This chapter shows firstly that, in the Hull-White model with no currency devaluation upon default, spread- and default-neutrality are equivalent. We then show that spread- and default-neutrality are not in general equivalent (even with zero currency devaluation) using the Alternative CIR and Black-Karasinski models as examples. The topic of proper hedging in the general case is left for further research.

11.1 Spread- and Default-Neutrality in Models that are Affine under Both Measures

We first consider hedging in a model that is affine under both measures, in the sense that, for any t and T,

$$S_d(t,T) = \exp\{-A_d(t,T) - C_d(t,T)\lambda(t)\}$$
$$S_f(t,T) = \exp\left\{-A_f(t,T) - C_f(t,T)\tilde{\lambda}(t)\right\}$$

for some deterministic functions A and C. Two special cases of this are the Hull-White and Alternative CIR models.

11.1.1 General Result

We fix a maturity T and consider the hedging of a unit short position (or equivalently the replication of a unit long position) in a foreign defaultable zero-coupon bond with zero recovery and maturity T.

The domestic currency value of the FDZCB is

$$Q(t)B_f^*(t,T) = (1 - N(t))Q(t)B_f(t,T)S_f(t,T)$$

or, omitting the arguments t and T,

$$QB_f^* = (1 - N)QB_fS_f.$$

Thus the differential of the domestic currency value of the FDZCB is

$$\begin{aligned} d(QB_f^*) &= dt \text{ terms} + (QB_f^*)_- \frac{dQ}{Q_-} - (QB_f^*)_- dN - C_f(QB_f^*)_- d\tilde{\lambda} \\ &= dt \text{ terms} + (QB_f^*)_- \frac{dQ}{Q_-} - (QB_f^*)_- dN - (1+\hat{\delta})C_f(QB_f^*)_- d\lambda. \end{aligned}$$

The replication/hedging portfolio that we consider is the natural one: a position Δ_1 in the foreign default-free bond and a position Δ_2 in the domestic defaultable bond, funded by domestic default-free borrowing. We denote the domestic currency value process of this portfolio by V. The differentials of the domestic currency values of the two hedge assets are

$$d(QB_f) = r_f(QB_f)_- dt + (QB_f)_- \frac{dQ}{Q_-}$$
$$dB_d^* = r_d(B_d^*)_- dt - (B_d^*)_- dN - C_d(B_d^*)_- d\lambda$$

The portfolio's evolution is thus given by

$$\begin{aligned} dV(t) &= \Delta_1 d(QB_f) + \Delta_2 dB_d^* + r_d(V(t) - \Delta_1(QB_f)_- - \Delta_2(B_d^*)_-)dt \\ &= dt \text{ terms} + \Delta_1(QB_f)_- \frac{dQ}{Q_-} - \Delta_2(B_d^*)_- dN - \Delta_2 C_d(B_d^*)_- d\lambda. \end{aligned}$$

Putting the two equations next to each other, we have

$$d(QB_f^*) = dt \text{ terms} + (QB_f^*)_- \frac{dQ}{Q_-} - (QB_f^*)_- dN - (1+\hat{\delta})C_f(QB_f^*)_- d\lambda$$
$$dV(t) = dt \text{ terms} + \Delta_1(QB_f)_- \frac{dQ}{Q_-} - \Delta_2(B_d^*)_- dN - \Delta_2 C_d(B_d^*)_- d\lambda.$$

To match the coefficients of dQ/Q_- we set

$$\Delta_1 = \frac{(B_f^*)_-}{B_f}.$$

In other words, we hold an amount of the foreign default-free bond with the same value as the FDZCB – this protects us against any exchange rate movements.

Now we have only one free quantity (Δ_2) to use in matching the coefficients of both dN and $d\lambda$. Let us hedge against default by matching the dN terms:

$$\Delta_2 = \frac{(QB_f^*)_-}{(B_d^*)_-}.$$

This means that we hold an amount of the DDZCB with the same domestic currency value as the FDZCB, which protects us against default. Then the coefficient of $d\lambda$ in the evolution of the portfolio value is

$$\begin{aligned} -\Delta_2 C_d(B_d^*)_- &= -\frac{(QB_f^*)_-}{(B_d^*)_-} C_d(B_d^*)_- \\ &= -C_d(QB_f^*)_-. \end{aligned}$$

On the other hand, the coefficient of $d\lambda$ in the evolution of the domestic currency FDZCB value is $-(1+\hat{\delta})C_f(QB_f^*)_-$. Thus, default- and spread-neutrality will be equal if and only if

$$C_d(t,T) = (1+\hat{\delta})C_f(t,T)$$

identically.

11.1.2 Special Case: Hull-White Model

In the Hull-White model we have, with a the speed of mean reversion of λ,

$$C_d(t,T) = C_f(t,T) = \frac{1}{a}\left(1 - e^{-a(T-t)}\right)$$

so that default- and spread-neutrality are equivalent if and only if $\hat{\delta} = 0$.

11.1.3 Special Case: The Alternative CIR Model

In the Alternative CIR model (see Chapter 8) we have

$$C_d(t,T) = \frac{2(\exp\{h(T-t)\}-1)}{2h+(a+h)(\exp\{h(T-t)\}-1)}$$

with $h = \sqrt{a^2 + 2\sigma_\lambda^2}$, while on the other hand

$$C_f(t,T) = \frac{2\left(\exp\left\{\tilde{h}(T-t)\right\}-1\right)}{2\tilde{h}+\left(\tilde{a}+\tilde{h}\right)\left(\exp\left\{\tilde{h}(T-t)\right\}-1\right)}$$

where $\tilde{h} = \sqrt{\tilde{a}^2 + 2\tilde{\sigma}_\lambda^2}$ with $\tilde{a} = a - \gamma_1\sigma_\lambda$ and $\tilde{\sigma}_\lambda = \sqrt{(1+\hat{\delta})}\sigma_\lambda$. Thus, even with $\hat{\delta} = 0$, we do not have $C_d(t,T) = (1+\hat{\delta})C_f(t,T)$ identically, and so default- and spread-neutrality are not equivalent.

This shows that we cannot in general expect to find that default- and spread-neutrality are equivalent. The default-delta (the holding in the DDZCB that will ensure default-neutrality) and the spread-delta (the holding that will ensure spread-neutrality) will in general be different. The following example illustrates possible default- and spread-deltas generated by various models.

11.2 Example

11.2.1 Example Composition

We now find the default- and spread-deltas in an example situation using the Hull-White (HW), Alternative CIR (ACIR) and Black-Karasinski (BK) models. For simplicity, we generate the domestic average hazard rate curve using the ACIR model – the HW and BK models can easily be fitted to this curve, while fitting the ACIR model to an exogenously given curve is somewhat more involved.

The example situation is:

Exchange rate: $Q(0) = 1$.

Interest rates: equal in domestic and foreign currencies.

These two points are chosen for simplicity, so that the default-deltas are simply the ratios of foreign to domestic survival probabilities.

Exchange rate volatility: 20%.

The exchange rate volatility is only used in the HW and BK models, though we check that the 'volatility' in the ACIR model (see page 73) is below 22% at the starting value and mean reversion level of λ.

Domestic average hazard rate curve: the average hazard rate curve generated by a CIR default intensity process with initial value 4%, mean reversion level 5%, speed of mean reversion 10% and square-root volatility 9%.

We will choose the volatility parameters a and σ_λ in the HW and BK models to match the variance of this CIR default intensity process at the 5- and 10-year points.

Mean foreign currency devaluation at default: $\hat{\delta} = 0$.

The volatility parameters that we use for the various models (found by the variance-matching described above) are:

	HW	CIR	BK
Speed of mean reversion:	8.2583%	10%	13.5349%
Volatility:	1.7869%	9%	40.6144%

We found the correlations ρ (in the HW and BK models) and exchange rate volatility parameters γ_1 (in the ACIR model) that lead to various levels of the five-year "quanto", by which we here mean the difference between the domestic and foreign five-year average hazard rates as a percentage of the domestic one:

$$\text{Quanto} = \frac{\text{Foreign average hazard rate}}{\text{Domestic average hazard rate}} - 1.$$

We plot these parameters in Figure 11.1. (By chance the values of γ_1 fell in the same region as the correlations.)

Figure 11.2 shows the domestic average hazard rate curve and the foreign average hazard rate curves generated by the various models, using the parameters that induce a 10% quanto at five years.

11.2.2 Example Deltas

The default- and spread-deltas in the Hull-White model are equal. We plot them in Figure 11.3 against the five-year quanto and the maturity of the FDZCB to be hedged (all maturities are in years). This figure exhibits the expected behaviour: if foreign hazard rates are higher than domestic ones, then FDZCBs are cheaper than DDZCBs and so to hedge a certain FDZCB exposure we need to hold DDZCBs in a smaller amount, i.e. $\Delta < 1$ – and vice versa. The effect of the correlation on the average hazard rates takes time to accumulate, so the deltas are near one for short maturities and increase/decrease with increasing maturity.

In Figure 11.4 we plot the default- and spread-deltas in the Alternative CIR model. The left-hand surface is the default-deltas, which exhibit the same pattern

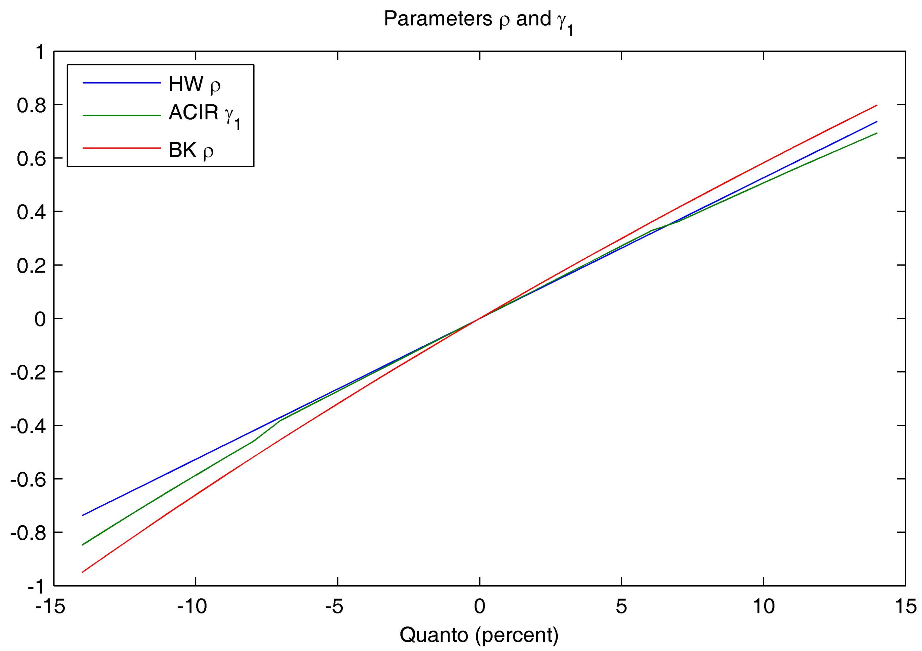

Figure 11.1: Values of ρ and γ_1 plotted against the five-year quanto that they induce.

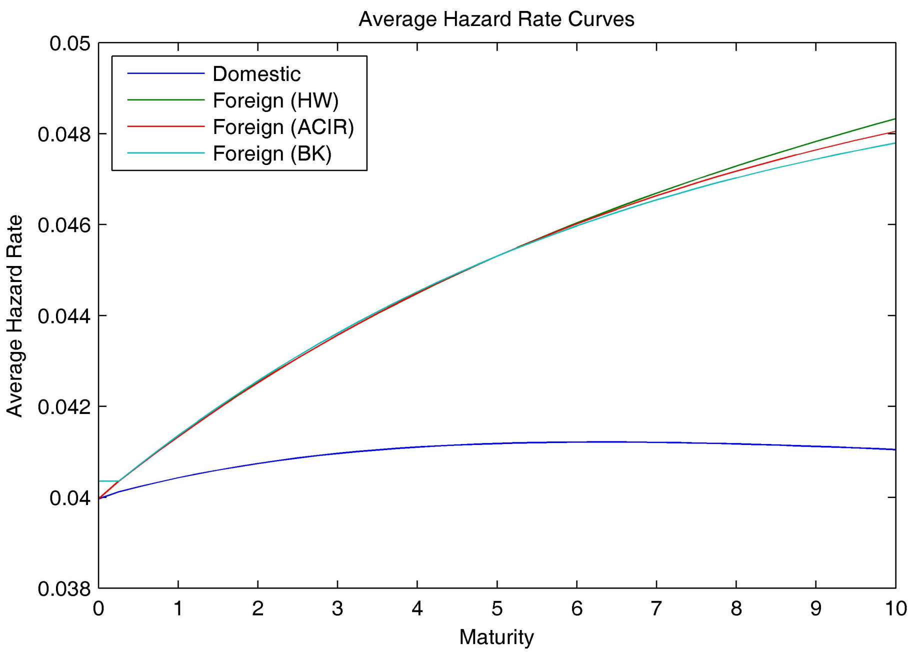

Figure 11.2: The domestic average hazard rate curve, and the foreign average hazard rate curves generated by the Hull-White, Alternative CIR and Black-Karasinski models, with parameters chosen to give a foreign five-year average hazard rate that is 10% higher than its domestic equivalent.

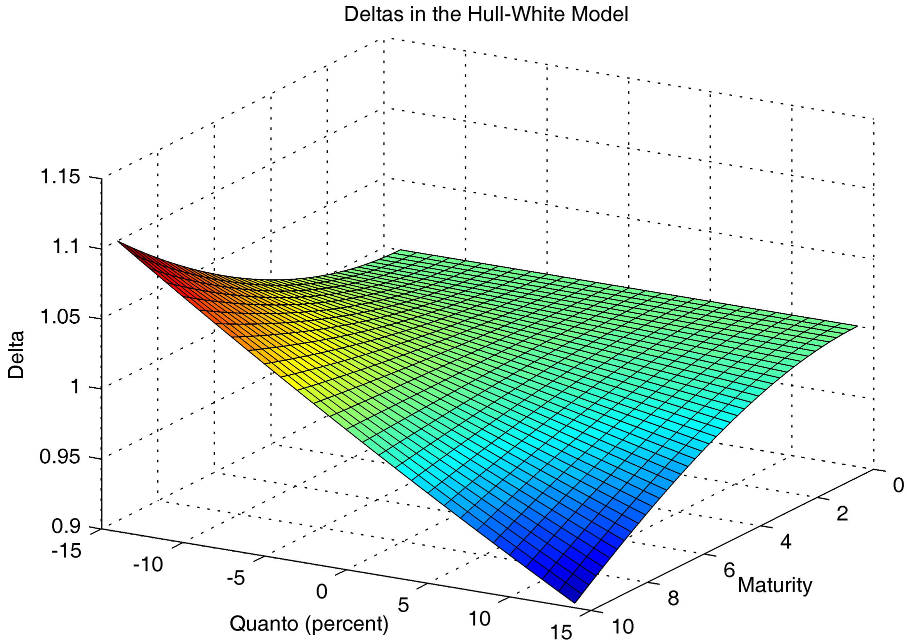

Figure 11.3: The delta (position in DDZCBs required to hedge a unit short position in FDZCBs) in the Hull-White model against the maturity of the defaultable bonds and the five-year quanto.

as the Hull-White deltas: a downward slope in the quanto that becomes more pronounced as maturity increases. The right-hand surface is the spread-deltas, which exhibit an upward slope in the quanto that first becomes steeper with increasing maturity and then flattens. In fact, with further increasing maturity this upward slope is reversed and the spread-deltas also become decreasing in the quanto; see Figure 11.5.

Figures 11.6 and 11.7 correspond to Figures 11.4 and 11.5, except that here the deltas are calculated using the Black-Karasinski model. The same pattern is exhibited.

11.3 Conclusions

It is clear that default- and spread-deltas may differ significantly in these basic models, and that a book of credit exposures that is immune to small movements in credit spreads will not in general also be immune to defaults. In all the considered cases the default-deltas behave as expected. The spread-deltas, on the other hand, move with increasing maturity from flat to increasing to decreasing in the quanto. The author has not been able to find a suitable intuitive explanation for this phenomenon.

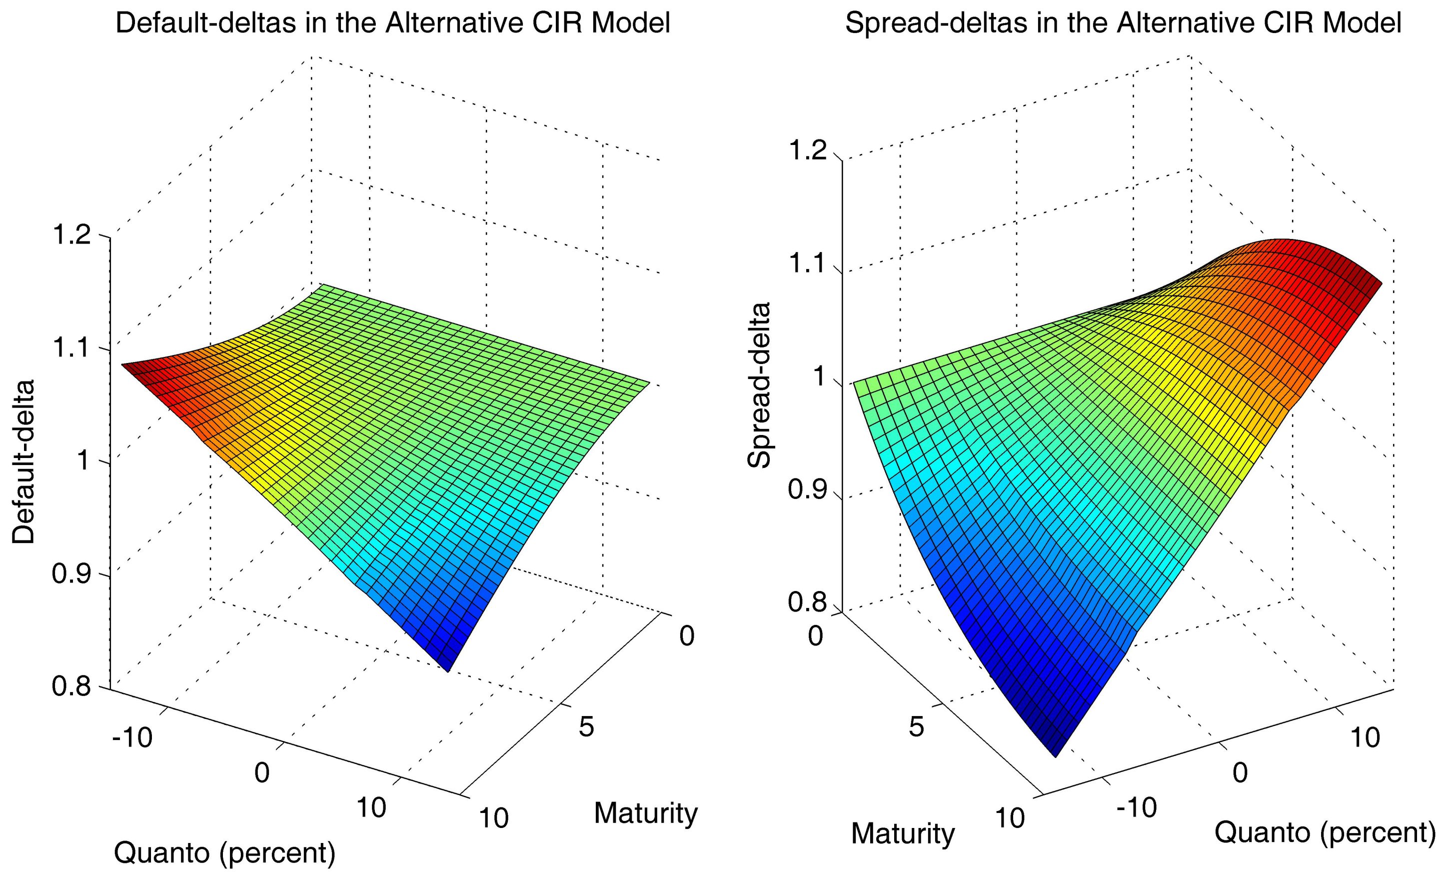

Figure 11.4: Default- and spread-deltas in the Alternative CIR model against the maturity of the defaultable bonds and the five-year quanto induced. Note the different orientations of the two axes.

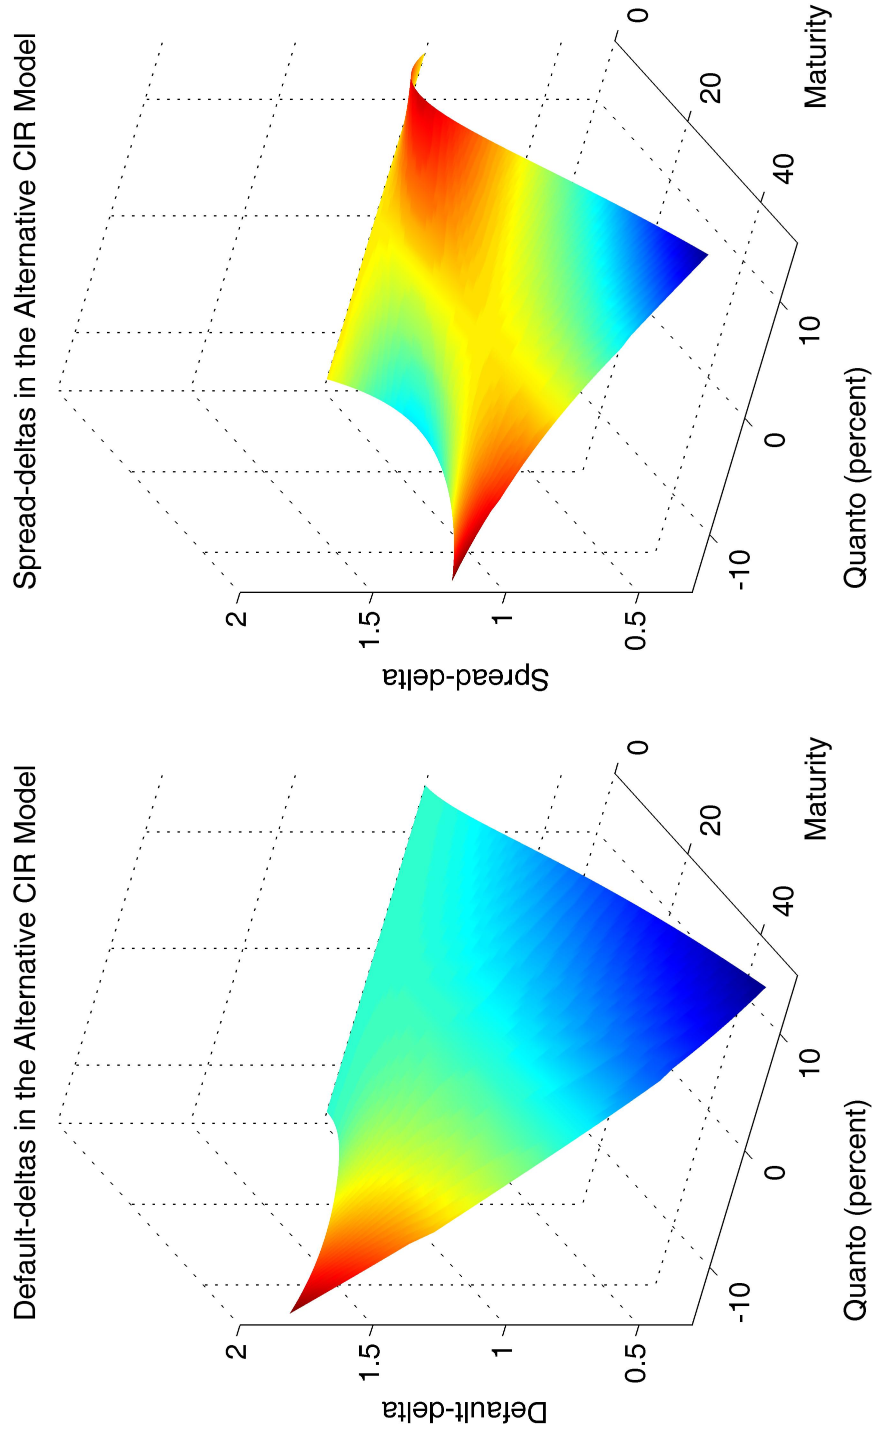

Figure 11.5: Default- and spread-deltas in the Alternative CIR model out to 50 years.

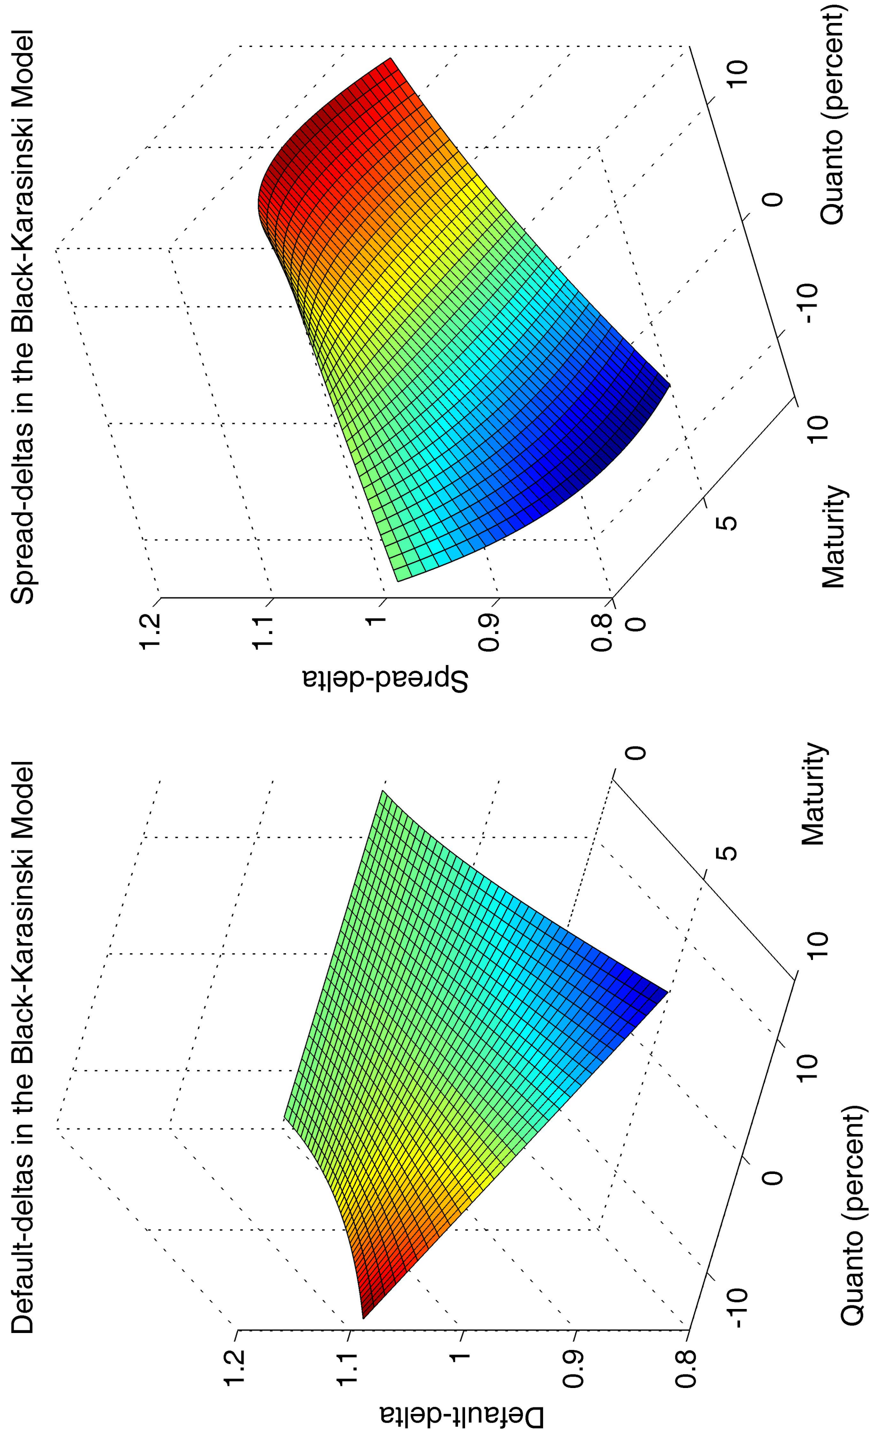

Figure 11.6: Default- and spread-deltas in the Black-Karasinski model against the maturity of the defaultable bonds and the five-year quanto induced. Note the different orientations of the two axes.

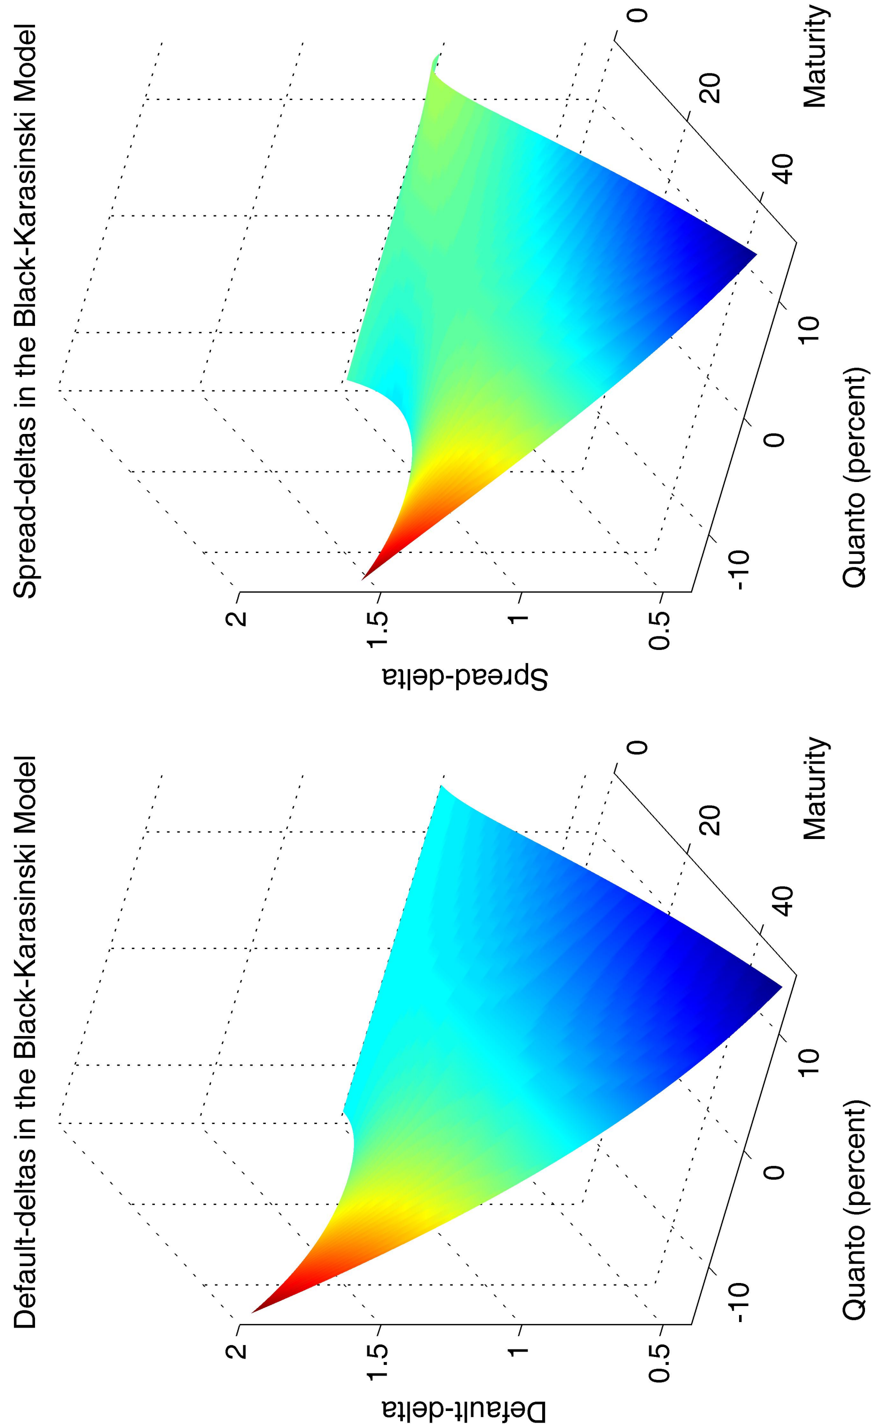

Figure 11.7: Default- and spread-deltas in the Black-Karasinski model out to 50 years.

Chapter 12

Conclusions

This dissertation has analysed the problem of valuing default-contingent claims in a market with multiple currencies and one default-risky participant. This valuation was conducted in the modelling framework proposed by Ehlers [14], which we recounted in Chapter 3. Most of the analysis was performed in the simpler model of Chapter 4, where we assumed that interest rates were deterministic and that the distribution of the random appreciation or devaluation of the foreign currency at the default time was deterministic and time-homogeneous. This allowed us to express the prices of coupon-bearing bonds and credit default swaps in terms of default-free discount factors and the domestic and foreign survival probabilities. We then proceeded to examine several possible specifications of this basic model.

The first was the Hull-White model, which is extremely tractable – we are able to express the foreign survival probabilities directly in terms of the domestic survival probabilities and the model parameters. This tractability comes at the cost of allowing the default intensity to become negative, which can result in survival probabilities increasing with maturity.

The Cox-Ingersoll-Ross model does not allow the default intensity to become negative. Unfortunately, the default intensity obeys an unusual stochastic differential equation under the foreign measure, and we were unable to find the foreign survival probabilities exactly in closed form. We proposed two approximate methods to calculate the foreign survival probabilities. The first is a standard trinomial tree (though we take a novel approach to the problematic behaviour of the tree near zero). The second relies on approximating the foreign default intensity process with a standard Cox-Ingersoll-Ross process, and gives us closed-form approximate foreign survival probabilities with reasonable accuracy. We also showed that if the exchange rate is modelled as a geometric Brownian motion, then there is no time-homogeneous model that is affine under both the domestic and foreign measures other than time-homogeneous Hull-White.

Our third model used lognormal Black-Karasinski dynamics for the default intensity, which means that we must calculate (or calibrate to) domestic survival probabilities using a trinomial tree. This disadvantage is overcome, though, by the fact that we can approximate the foreign survival probabilities using the same tree with just a simple adjustment, and that this approximation is excellent even with long tree steps (greater than three months).

The last of our specifications was what we called the Alternative CIR model. This model is due to Ehlers [14], and combines a Cox-Ingersoll-Ross process for the default intensity with a non-standard process for the exchange rate, where the exchange rate volatility depends upon the level of the default intensity. This exchange rate has the advantage that the default intensity remains a CIR process when we change to the foreign measure, allowing us closed-form survival probabilities in both currencies. The original part of that chapter was the comparison of this alternative exchange rate with standard geometric Brownian motion. Though no strong conclusions were drawn about the suitability of this model for practical use, the behaviour of the alternative exchange rate was clarified.

Chapter 9 recalled Ehlers' use of affine diffusions in his general model. We showed how the Alternative CIR model arises from the use of an affine diffusion, and illustrated a more general model where the domestic and foreign interest rates are allowed to be stochastic. This model shares the main fault of the Alternative CIR model: if we require the default intensity to remain positive, then the volatilities of the exchange rate and the interest rates depend upon the level of the default intensity.

Chapter 10 briefly considered the use of local and stochastic volatility models for the exchange rate, and the valuation of vulnerable options. We found that when using such models, or valuing such products, the change to the foreign measure was of little use, while a change to the domestic survival measure was of some use, at least if the default intensity is a Hull-White process.

Lastly, we considered the problem of hedging a short position in a foreign defaultable zero-coupon bond with zero recovery using a similar domestic bond (with the same maturity) and a foreign default-free bond. We found that in the Hull-White model, the amounts of the domestic defaultable bond required to hedge against default and against spread movements were identical, but that this does not hold in general: default- and spread-neutrality are not equivalent. The problem of proper hedging against both default and spread movements is left for further research.

While most of this dissertation consisted of reviews and applications of previous work, we have contributed the Nearest CIR approximtion in the Cox-Ingersoll-Ross model and clarified the characteristics of the Alternative CIR model. More significantly, the Black-Karasinski model was not previously considered in the literature and has been shown to be extremely useful and efficient. We have also contributed to the understanding of hedging credit risk in multiple currencies: the relationship between spread- and default-neutrality does not appear to have been considered before, at least in this context, and we have shown that the two are not equivalent.

Bibliography

[1] T.M. Barnhill, Jr. and W.F. Maxwell, *Modeling correlated market and credit risk in fixed income portfolios*, Journal of Banking and Finance **26** (2002), no. 2-3, 347–374.

[2] T.R. Bielecki and M. Rutkowski, *Credit risk: Modeling, valuation and hedging*, Springer Finance, Springer-Verlag, 2002.

[3] T. Björk, *Arbitrage theory in continuous time*, Oxford University Press, 1998.

[4] F. Black and P. Karasinski, *Bond and option pricing when short rates are log-normal*, Financial Analysts Journal **47** (1991), no. 4, 52–59.

[5] F. Black and M. Scholes, *The pricing of options and corporate liabilities*, Journal of Political Economy **81** (1973), no. 3, 637–654.

[6] D. Brigo and F. Mercurio, *Interest rate models – theory and practice: with smile, inflation and credit*, second ed., Springer Finance, Springer-Verlag, 2006.

[7] J.M. Campa, P.H.K. Chang, and R.L. Reider, *Implied exchange rate distributions: evidence from OTC option markets*, Journal of International Money and Finance **17** (1998), no. 1, 117–160.

[8] J.A. Chan-Lau, *FX-adjusted local currency spreads*, Working paper (available at *ssrn.com*), May 2008.

[9] J.A. Chan-Lau and A.O. Santos, Sr., *Currency mismatches and corporate default risk: Modeling, measurement, and surveillance applications*, IMF Working paper 06/269, Monetary and Capital Markets Department (available at *ssrn.com*), December 2006.

[10] J.C. Cox, J.E. Ingersoll, Jr., and S.A. Ross, *A theory of the term structure of interest rates*, Econometrica **53** (1985), no. 2, 385–407.

[11] Q. Dai and K.J. Singleton, *Specification analysis of affine term structure models*, The Journal of Finance **LV** (2000), no. 5, 1943–1978.

[12] S.A. Davydenko and J.R. Franks, *Do bankruptcy codes matter? A study of defaults in France, Germany, and the U.K.*, The Journal of Finance **63** (2008), no. 2, 565–608.

[13] D. Duffie and K.J. Singleton, *Modeling term structures of defaultable bonds*, The Review of Financial Studies **12** (1999), no. 4, 687–720.

[14] P. Ehlers, *Pricing credit derivatives*, Ph.D. thesis, Swiss Federal Institute of Technology, Zurich, 2007.

[15] P. Ehlers and P. Schönbucher, *The influence of FX risk on credit spreads*, Working paper (January 2006 version), March 2004.

[16] P.S. Hagan, D. Kumar, A.S. Lesniewski, and D.E. Woodward, *Managing smile risk*, Wilmott Magazine (2002), 84–108.

[17] P.S. Hagan and G. West, *Interpolation methods for curve construction*, Applied Mathematical Finance **13** (2006), no. 2, 89–129.

[18] S.L. Heston, *A closed-form solution for options with stochastic volatility with applications to bond and currency options*, The Review of Financial Studies **6** (1993), no. 2, 327–343.

[19] P. Houweling, J. Hoek, and F. Kleibergen, *The joint estimation of term structures and credit spreads*, Journal of Empirical Finance **8** (2001), no. 3, 297–323.

[20] J. Hull and A. White, *Pricing interest rate derivative securities*, The Review of Financial Studies **3** (1990), no. 4, 573–592.

[21] ———, *Valuing derivative securities using the explicit finite difference method*, Journal of Financial and Quantitative Analysis **25** (1990), no. 1, 87–100.

[22] ———, *Single-factor interest rate models and the valuation of interest rate derivative securities*, Journal of Financial and Quantitative Analysis **28** (1993), no. 2, 235–254.

[23] ———, *Numerical Procedures for Implementing Term Structure Models I: Single-Factor Models*, The Journal of Derivatives **2** (1994), no. 1, 7–16.

[24] J. Jacod and A.N. Shiryaev, *Limit theorems for stochastic processes*, Springer-Verlag, 1987.

[25] R. Jankowitsch and S. Pichler, *Currency dependence of corporate credit spreads*, Journal of Risk **8** (2005), no. 1.

[26] R.A. Jarrow and P. Protter, *Structural versus reduced form models: A new information based perspective*, Journal of Investment Management **2** (2004), no. 2, 1–10.

[27] R.A. Jarrow and S.M. Turnbull, *Pricing derivatives on financial securities subject to credit risk*, Journal of Finance **50** (1995), no. 1, 53–85.

[28] M. Jeanblanc, M. Yor, and M. Chesney, *Mathematical methods for financial markets*, Springer Finance, Springer-Verlag, 2009.

[29] M. Kafetzaki-Boulamatsis and D. Tasche, *Combined market and credit risk stress testing based on the Merton model*, Working paper, June 2001.

[30] I. Karatzas and S.E. Shreve, *Brownian motion and stochastic calculus*, second ed., Springer-Verlag, 1991.

[31] F.C. Klebaner, *Introduction to stochastic calculus with applications*, second ed., Imperial College Press, 2005.

[32] D. Lando, *On Cox processes and credit risky securities*, Review of Derivatives Research **2** (1998), no. 2-3, 99–120.

[33] A. Li, *Valuation of credit-contingent options with applications to quanto CDS*, Working paper (August 2008 version, available at *ssrn.com*), June 2006.

[34] Robert S. Liptser and Albert N. Shiryaev, *Stochastic calculus on filtered probability spaces*, Probability Theory III, Encyclopaedia of Mathematical Sciences, vol. 45, Springer-Verlag, 1998, pp. 111–157.

[35] ______, *Statistics of Random Processes II: Applications*, second ed., Applications of Mathematics: Stochastic Modelling and Applied Probability, Springer-Verlag, 2001.

[36] D.B. Madan and H. Unal, *Pricing the risks of default*, Review of Derivatives Research **2** (1998), no. 2-3, 121–160.

[37] R.C. Merton, *On the pricing of corporate debt: the risk structure of interest rates*, Journal of Finance **29** (1974), no. 2, 449–470.

[38] S.K. Nawalkha and N.A. Beliaeva, *Efficient trees for CIR and CEV short rate models*, The Journal of Alternative Investments **10** (2007), no. 1, 71–90.

[39] D.B. Nelson and K. Ramaswamy, *Simple binomial processes as diffusion approximations in financial models*, The Review of Financial Studies **3** (1990), no. 3, 393–430.

[40] P.J. Schönbucher, *Credit derivatives pricing models*, Wiley Finance, Wiley, 2003.

[41] E. Sener and T. Kenc, *Empirical investigation of currency dependence of credit spreads*, Working paper, April 2008.

[42] S.E. Shreve, *Stochastic Calculus for Finance II: Continuous-Time Models*, Springer Finance, Springer-Verlag, 2004.

[43] D. Tasche, *Incorporating exchange rate risk into PDs and asset correlations*, Working paper, December 2007.

[44] I. Warnes and G. Acosta, *Valuation of international corporate debt*, Working paper, 2002.

[45] A.B. Yigitbasioglu, *Pricing convertible bonds with interest rate, equity, credit, FX and volatility risk*, Working paper (August 2002 version), November 2001.