1
An Overview

The natural place to start the exposition is with the Black and Scholes (1973) and Merton (1974) milestones. The development of option-pricing techniques and the application to the study of corporate liabilities is where the modeling of credit risk has its foundations. While there was of course research out before this, the option-pricing literature, which views the bonds and stocks issued by a firm as contingent claims on the assets of the firm, is the first to give us a strong link between a statistical model describing default and an economic-pricing model. Obtaining such a link is a key problem of credit risk modeling. We make models describing the distribution of the default events and we try to deduce prices from these models. With pricing models in place we can then reverse the question and ask, given the market prices, what is the market's perception of the default probabilities. To answer this we must understand the full description of the variables governing default and we must understand risk premiums. All of this is possible, at least theoretically, in the option-pricing framework.

Chapter 2 starts by introducing the Merton model and discusses its implications for the *risk structure* of interest rates—an object which is not to be mistaken for a term structure of interest rates in the sense of the word known from modeling government bonds. We present an immediate application of the Merton model to bonds with different seniority. There are several natural ways of generalizing this, and to begin with we focus on extensions which allow for closed-form solutions. One direction is to work with different asset dynamics, and we present both a case with stochastic interest rates and one with jumps in asset value. A second direction is to introduce a default boundary which exists at all time points, representing some sort of safety covenant or perhaps liquidity shortfall. The Black–Cox model is the classic model in this respect. As we will see, its derivation has been greatly facilitated by the development of option-pricing techniques. Moreover, for a clever choice of default boundary, the model can be generalized to a case with stochastic interest rates. A third direction is to include coupons, and we discuss the extension both to discrete-time, lumpy dividends and to continuous flows of dividends and continuous coupon payments. Explicit solutions are only available if the time horizon is made infinite.

Having the closed-form expressions in place, we look at a numerical scheme which works for any hitting time of a continuous boundary provided that we know the transition densities of the asset-value process. With a sense of what can be done with closed-form models, we take a look at some more practical issues.

Coupon payments really distinguish corporate bond pricing from ordinary option pricing in the sense that the asset-sale assumptions play a critical role. The liquidity of assets would have no direct link to the value of options issued by third parties on the firm's assets, but for corporate debt it is critical. We illustrate this by looking at the term-structure implications of different asset-sale assumptions.

Another practical limitation of the models mentioned above is that they are all static, in the sense that no new debt issues are allowed. In practice, firms roll over debt and our models should try to capture that. A simple model is presented which takes a stationary leverage target as given and the consequences are felt at the long end of the term structure. This anticipates the models of Chapter 3, in which the choice of leverage is endogenized.

One of the most practical uses of the option-based machinery is to derive implied asset values and implied asset volatilities from equity market data given knowledge of the debt structure. We discuss the maximum-likelihood approach to finding these implied values in the simple Merton model. We also discuss the philosophy behind the application of implied asset value and implied asset volatility as variables for quantifying the probability of default, as done, for example (in a more complicated and proprietary model), by Moody's KMV.

The models in Chapter 2 are all incapable of answering questions related to the optimal capital structure of firms. They all take the asset-value process and its division between different claims as given, and the challenge is to price the different claims given the setup. In essence we are pricing a given securitization of the firm's assets.

Chapter 3 looks at the standard approach to obtaining an optimal capital structure within an option-based model. This involves looking at a trade-off between having a tax shield advantage from issuing debt and having the disadvantage of bankruptcy costs, which are more likely to be incurred as debt is increased. We go through a model of Leland, which despite (perhaps because of) its simple setup gives a rich playing field for economic interpretation. It does have some conceptual problems and these are also dealt with in this chapter. Turning to models in which the underlying state variable process is the EBIT (earnings before interest and taxes) of a firm instead of firm value can overcome these difficulties. These models can also capture the important phenomenon that equity holders can use the threat of bankruptcy to renegotiate, in times of low cash flow, the terms of the debt, forcing the debt holders to agree to a lower coupon payment. This so-called strategic debt service is more easily explained in a binomial setting and this is how we conclude this chapter.

At this point we leave the option-pricing literature. Chapter 4 briefly reviews different statistical techniques for analyzing defaults. First, classical discriminant analysis is reviewed. While this model had great computational advantages before statistical computing became so powerful, it does not seem to be a natural statistical model for default prediction. Both logistic regression and hazard regressions have a more natural structure. They give parameters with natural interpretations and handle issues of censoring that we meet in practical data analysis all the time. Hazard regressions also provide natural nonparametric tools which are useful for exploring the data and for selecting parametric models. And very importantly, they give an extremely natural connection to pricing models. We start by reviewing the discrete-time hazard regression, since this gives a very easy way of understanding the occurrence/exposure ratios, which are the critical objects in estimation—both parametrically and nonparametrically.

While on the topic of default probability estimation it is natural to discuss some techniques for analyzing rating transitions, using the so-called generator of a Markov chain, which are useful in practical risk management. Thinking about rating migration in continuous time offers conceptual and in some respects computational improvements over the discrete-time story. For example, we obtain better estimates of probabilities of rare events. We illustrate this using rating transition data from Moody's. We also discuss the role of the Markov assumption when estimating transition matrices from generator matrices.

The natural link to pricing models brought by the continuous-time survival analysis techniques is explained in Chapter 5, which introduces the intensity setting in what is the most natural way to look at it, namely as a Cox process or doubly stochastic Poisson process. This captures the idea that certain background variables influence the rate of default for individual firms but with no feedback effects. The actual default of a firm does not influence the state variables. While there are important extensions of this framework, some of which we will review briefly, it is by far the most practical framework for credit risk modeling using intensities. The fact that it allows us to use many of the tools from default-free term-structure modeling, especially with the affine and quadratic term-structure models, is an enormous bonus. Particularly elegant is the notion of recovery of market value, which we spend some time considering. We also outline how intensity models are estimated through the extended Kalman filter—a very useful technique for obtaining estimates of these heavily parametrized models.

For the intensity model framework to be completely satisfactory, we should understand the link between estimated default intensities and credit spreads. Is there a way in which, at least in theory, estimated default intensities can be used for pricing? There is, and it is related to diversifiability but not to risk neutrality, as one might have expected. This requires a thorough understanding of the risk premiums, and an important part of this chapter is the description of what the sources of excess expected return are in an intensity model. An important moral of this chapter is that even if intensity models look like ordinary term-structure models, the structure of risk premiums is richer.

How do default intensities arise? If one is a firm believer in the Merton setting, then the only way to get something resembling default intensities is to introduce jumps in asset value. However, this is not a very tractable approach from the point of view of either estimation or pricing credit derivatives. If we do not simply want to assume that intensities exist, can we still justify their existence? It turns out that we can by introducing incomplete information. It is shown that in a diffusion-based model, imperfect observation of a firm's assets can lead to the existence of a default intensity for outsiders to the firm.

Chapter 6 is about rating-based pricing models. This is a natural place to look at those, as we have the Markov formalism in place. The simplest illustration of intensity models with a nondeterministic intensity is a model in which the intensity is "modulated" by a finite-state-space Markov chain. We interpret this Markov chain as a rating, but the machinery we develop could be put to use for processes needing a more-fine-grained assessment of credit quality than that provided by the rating system.

An important practical reason for looking at ratings is that there are a number of financial instruments that contain provisions linked to the issuer rating. Typical examples are the step-up clauses of bond issues used, for example, to a large extent in the telecommunication sector in Europe. But step-up provisions also figure prominently in many types of loans offered by banks to companies.

Furthermore, rating is a natural first candidate for grouping bond issues from different firms into a common category. When modeling the spreads for a given rating, it is desirable to model the joint evolution of the different term structures, recognizing that members of each category will have a probability of migrating to a different class. In this chapter we will see how such a joint modeling can be done. We consider a calibration technique which modifies empirical transition matrices in such a way that the transition matrix used for pricing obtains a fit of the observed term structures for different credit classes. We also present a model with stochastically varying spreads for different rating classes, which will become useful later in the chapter on interest-rate swaps. The problem with implementing these models in practice are not trivial. We look briefly at an alternative method using thresholds and affine process technology which has become possible (but is still problematic) due to recent progress using transform methods. The last three chapters contain applications of our machinery to some important areas in which credit risk analysis plays a role

The analysis of interest-rate swap spreads has matured greatly due to the advances in credit risk modeling. The goal of this chapter is to get to the point at which the literature currently stands: counterparty credit risk on the swap contract is not a key factor in explaining interest-rate swap spreads. The key focus for understanding the joint evolution of swap curves, corporate curves, and treasury curves is the fact that the floating leg of the swap contract is tied to LIBOR rates.

But before we can get there, we review the foundations for reaching that point. A starting point has been to analyze the apparent arbitrage which one can set up using swap markets to exchange fixed-rate payments for floating-rate payments. While there may very well be institutional features (such as differences in tax treatments) which permit such advantages to exist, we focus in Chapter 7 on the fact that the comparative-advantage story can be set up as a sort of puzzle even in an arbitrage-free model. This puzzle is completely resolved. But the interest in understanding the role of two-sided default risk in swaps remains. We look at this with a strong focus on the intensity-based models. The theory ends up pretty much confirming the intuitive result: that swap counterparties with symmetric credit risk have very little reason to worry about counterparty default risk. The asymmetries that exist between their payments—since one is floating and therefore not bounded in principle, whereas the other is fixed—only cause very small effects in the pricing. With netting agreements in place, the effect is negligible. This finally clears the way for analyzing swap spreads and their relationship to corporate bonds, focusing on the important problem mentioned above, namely that the floating payment in a swap is linked to a LIBOR rate, which is bigger than that of a short treasury rate. Viewing the LIBOR spread as coming from credit risk (something which is probably not completely true) we set up a model which determines the fixed leg of the swap assuming that LIBOR and AA are the same rate. The difference between the swap rate and the corporate AA curve is highlighted in this case. The difference is further illustrated by showing that theoretically there is no problem in having the AAA rate be above the swap rate—at least for long maturities.

The result that counterparty risk is not an important factor in determining credit risk also means that swap curves do not contain much information on the credit quality of its counterparties. Hence swaps between risky counterparties do not really help us with additional information for building term structures for corporate debt. To get such important information we need to look at default swaps and asset swaps. In idealized settings we explain in Chapter 8 the interpretation of both the asset-swap spread and the default swap spread. We also look at more complicated structures involving baskets of issuers in the so-called first-to-default swaps and first m-of-n-to-default swaps. These derivatives are intimately connected with so-called collateralized debt obligations (CDOs), which we also define in this chapter.

Pricing of CDOs and analysis of portfolios of loans and credit-risky securities lead to the question of modeling dependence of defaults, which is the topic of the whole of Chapter 9. This chapter contains many very simplified models which are developed for easy computation but which are less successful in preserving a realistic model structure. The curse is that techniques which offer elegant computation of default losses assume a lot of homogeneity among issuers. Factor structures can mitigate but not solve this problem. We discuss correlation of rating movements derived from asset-value correlations and look at correlation in intensity models. For intensity models we discuss the problem of obtaining correlation in affine specifications of the CIR type, the drastic covariation of intensities needed to generate strong default correlation and show with a stylized example how the updating of a latent variable can lead to default correlation.

Recently, a lot of attention has been given to the notion of copulas, which are really just a way of generating multivariate distributions with a set of given marginals. We do not devote a lot of time to the topic here since it is, in the author's view, a technique which still relies on parametrizations in which the parameters are hard to interpret. Instead, we choose to spend some time on default dependence in financial networks. Here we have a framework for understanding at a more fundamental level how the financial ties between firms cause dependence of default events. The interesting part is the clearing algorithm for defining settlement payments after a default of some members of a financial network in which firms have claims on each other.

After this the rest is technical appendices. A small appendix reviews arbitrage-free pricing in a discrete-time setting and hints at how a discrete-time implementation of an intensity model can be carried out. Two appendices collect material on Brownian motion and Markov chains that is convenient to have readily accessible. They also contains a section on processes with jumps, including Itô's formula and, just as important, finding the martingale part and the drift part of the contribution coming from the jumps. Finally, they look at some abstract results about (special) semimartingales which I have found very useful. The main goal is to explain the structure of risk premiums in a structure general enough to include all models included in this book. Part of this involves looking at excess returns of assets specified as special semimartingales. Another part involves getting a better grip on the quadratic variation processes.

Finally, there is an appendix containing a workhorse for term-structure modeling. I am sure that many readers have had use of the explicit forms of the Vasicek and the Cox–Ingersoll–Ross (CIR) bond-price models. The appendix provides closed-form solutions for different functionals and the characteristic function of a one-dimensional affine jump-diffusion with exponentially distributed jumps. These closed-form solutions cover all the pricing formulas that we need for the affine models considered in this book.