

Modelling Credit Risk

L.C.G. Rogers*
University of Cambridge

January 31, 2008

FOREWORD. The notes which follow were originally prepared for a course given in 1999, and in the intervening years the whole subject of credit has boomed; accordingly, though in parts the material of these notes is timeless, elsewhere it appears distinctly dated. Please read with indulgence; and in the Epilogue I will make some remarks on the subject more recently, and where it might look in the future.

1 What is the problem?

The first and most important thing to realise about modelling of credit risk is that we may be trying to answer questions of *two different types*, and the links between the two are somewhat tenuous.

To fix some notation, let's suppose that the value of a firm's assets at time t are denoted by V_t , and that these evolve as

$$dV_t = V_{t-}(\sigma_t dW_t + (\mu_t - c_t)dt - dJ_t), \quad (1)$$

where W is a standard Brownian motion, μ is some process, the *rate of return* process, c is the dividend process, and J is some jumping process. The volatility process σ can be quite general. The value at time t of the total equity of the firm is denoted by S_t , and so the value at time t of all debt must be simply $V_t - S_t$. Default happens at some time τ ; upon default, the control of the firm passes from the shareholders to the bondholders, and there may be restructuring

*Statistical Laboratory, University of Cambridge, Wilberforce Road, Cambridge CB3 0WB, UK;

losses incurred. Default may happen because of some failure of the firm to fulfil its obligations to its creditors, or it may happen because the shareholders decide to surrender control by declaring bankruptcy. For now, we will not be specific about the circumstances of default.

We might be interested in questions of *risk-management*; what is the probability that the firm will default in the next 5 years? what is the expected loss on default if that happens? what are our expected losses from all the defaults of our obligors over the next 5 years? To answer these, we would presumably need to know the dividend policy, and the statistics of the rate of return. We would also need to know under what conditions default happens; if we make the simplifying assumption that it happens when V falls to some level, then we are dealing with a first-passage problem for a continuous process, which we may or may not be able to solve. Our approach to the problem would involve us in estimating the dynamics of various processes, as well as the dependence of μ on them. This would require historical data, and perhaps some judgemental inputs.

Contrast this with the situation we face if we are trying to answer *pricing* questions. This time, we are working in the *risk-neutral* or *pricing* measure, and so (1) is modified to make the rate of return equal to the riskless spot-rate r . To estimate this model, we would be looking to market prices - prices of equity, and credit-sensitive instruments. In the extreme case of $J = 0$, we would simply have $\mu = r$, and all structural dependence of the rate of return on economic fundamentals washes out!

The two situations involve working in two different measures, and they are linked only rather indirectly; if we had built a good model for one class of questions, we would have to make some (usually quite arbitrary) assumptions on risk premia to translate to a model for the other class of questions. We shall study in some detail below an example which displays clearly the kinds of difficulty involved.

We posed the problem in the generality of (1) because this form embraces the two main types of approach in the literature: the *structural* approach, and the *hazard rate* or *reduced form* approach. Typically in the first, the term dJ is absent, the value of the firm is modelled as a continuous process, with default occurring when the value reaches some (possibly time-dependent) barrier. In the second, the emphasis is on the jump process dJ , and default will occur at the first jump time of J .

It seems that we really need to include both components in our analysis¹. The structural approach fails to match the observed evidence that corporate spreads

¹The paper [4] is a first (rather simple-minded) step in this direction

do not decrease to zero as maturity decreases to zero; even for short maturities, the market does not neglect the possibility that some disaster may strike. On the other hand, the reduced-form approach can struggle to capture the dependency between defaults of different firms; in principle, by making the hazard rate depend on a range of other processes we can incorporate this, but this is rather artificial - we still need to understand the other processes on which the hazard rate depends.

Data issues. Let's firstly consider the risk-management questions. In order to estimate the probability of default of the firm, we would like to know as much as possible about the rate of return process. This may be affected by:

- costs of labour and raw materials;
- interest rates in countries where the firm produces and sells;
- exchange rates between countries where the firm produces and sells;
- recent and projected sales of the corporation;
- other debt issues outstanding, their priority, maturity and other characteristics;
- failure of obligors;
- perceived credit-worthiness of the corporation; as well as by
- taxation levels in various countries;
- technological progress and new products of the firm and competitors;
- possible major falls in value (e.g., litigation);
- continuity and competence of management;
- major political and market changes.

If we were well informed about all these things (as we would be if the corporation had approached us for a \$5 bn loan, or for a credit rating), most of the uncertainty about default would be removed. On the other hand, much of this information about the firm would be very difficult for individual investors to discover, so most would be relying on coarser information, such as credit ratings, share prices and prices of the firm's bonds. We must be cautious in using *all* of these!

To begin with, credit ratings are obtained by some gross aggregation of many diverse

corporations in order to make some estimates of (unlikely) changes in credit class, or default, and then all corporations with the same credit rating are treated the same for the purpose of assessing default risk. Now this is clearly too simplified; Hersheys will be significantly affected by the price of cocoa, Ford will not. Also, the moves between credit classes are often modelled as a continuous-time Markov chain, which means that the times in ratings classes will be exponentially distributed, but more importantly, the probability of a downgrade given that a firm has just experienced one is higher than for a firm that has been in that class for some time. This is not supported by evidence. Credit ratings convey only very crude information about the riskiness of a firm's debt - it would be tempting to omit them entirely from any modelling effort, were it not for the fact that there are various credit-sensitive products whose payoffs depend on the credit class to which the firm is assigned!

As far as share and bond prices go, these are calculated using the pricing measure, so we can't expect them to tell us much of use for risk management, apart from information about volatility.

How about the pricing questions? This time the useful data is the data relating to the pricing measure, so the prices of shares and corporate bonds; empirical estimates of ratings class transitions will not tell us anything we can use directly here. One point to note is that for sovereign debt, we do not have any share prices, so the range of usable data is much less; we would like our models to work without share price data, therefore.

Is there no useable link between the pricing measure and the real-world measure? Not entirely; the dividend policy of the firm will presumably depend on various economic fundamentals as well as the value of the firm, and the share price is just the net present value of all future dividends, so there is a link here. However, we still have to understand the law of the fundamentals in the pricing probability, so the matter is not ended.

2 Hazard rate models

There are two broad classes of models, the *structural* models (characterised by an attempt to model default by modelling the dynamics of the assets of the firm) and the *hazard rate* models, where the idea is that the default comes 'by surprise' in some sense, and we merely try to model the infinitesimal likelihood of a default. Hazard rate models are also called *reduced form* models by some authors.

In hazard rate models, the fundamental modelling tool is the Poisson process, and we begin by recalling the definition and some properties.

Definition 2.1 A Poisson counting process $(N_t)_{t \geq 0}$ is a non-decreasing process with right-continuous paths and values in \mathbb{Z}^+ such that

(2.1.i) $N_0 = 0$;

(2.1.ii) for any $0 \leq s_1 \leq t_1 \leq s_2 \leq t_2 \leq \dots \leq s_n \leq t_n$, the random variables $X_i \equiv N(t_i) - N(s_i)$ are independent, and the distribution of each X_i depends only on the length $t_i - s_i$;

(2.1.iii) for all $t \geq 0$, $N_t - N_{t-}$ is either 0 or 1.

The definition of the Poisson process uniquely determines its distribution to within a single positive parameter λ . When $\lambda = 1$, we speak of a standard Poisson process. Here are other key properties, in which the positive parameter λ appears explicitly.

(2.2) the process $\tilde{N}_t \equiv N_t - \lambda t$ is a martingale;

(2.3) the inter-event times $T_n - T_{n-1}$ are independent with common exponential(λ) distribution:

$$P[T_n - T_{n-1} > t] = \exp(-\lambda t)$$

for all $t \geq 0$; (Here, $T_n \equiv \inf\{t \geq 0 \mid N_t = n\}$.)

(2.4) For any $s \leq t$, $N_t - N_s \sim \mathcal{P}(\lambda(t-s))$, the Poisson distribution with mean λ :

$$P[N_t - N_s = k] = e^{-\lambda(t-s)} \lambda^k (t-s)^k / k!$$

for $k \in \mathbb{Z}^+$;

This much is known from any introductory text on stochastic processes, where the Poisson process will be motivated by descriptions of the arrivals of radioactive particles at a Geiger counter, or customers at a post office counter. But suppose we were counting the radioactive particles arriving from some source at the Geiger counter, and after one minute we halved the distance from source to counter. Physics tells us that the intensity of counts would be multiplied by 4, but how would we model it? We could suppose that we have two independent Poisson processes N' and N'' with intensities λ and 4λ respectively, and set up the counting process

$$\tilde{N}_t = N'(t \wedge 1) + N''(t \vee 1) - N''(1),$$

but a neater way to do it is to suppose we have a standard Poisson process N and define the counting process

$$N_t^* \equiv N(H_t),$$

where

$$\begin{aligned} H_t &= \lambda(t + 3(t-1)^+) \\ &= \int_0^t h_s ds, \end{aligned} \tag{2}$$

where $h_s = \lambda(I_{\{s < 1\}} + 4I_{\{s \geq 1\}})$. The function h is the *intensity* or *hazard rate* function of the counting process N^* ; the bigger it is, the faster the events (the jumps of N^*) are coming. This way of looking at the problem is powerful, because it permits immediate generalization to intensity functions which are allowed to be stochastic, and this is really all that is going on in the hazard rate approach to credit risk modelling. In more detail, we model the time τ of default as the first time that N^* jumps, so we shall have

$$H(\tau) = T_1. \quad (3)$$

This is true for stochastic hazard rate processes as well, of course. In our modelling, we shall suppose that we have defined the hazard rate process in some way, and then take an *independent* standard Poisson process N and define τ by way of (2) and (3). From this, we have immediately the key relation

$$\begin{aligned} P[\tau > t] &= P[T_1 > H(t)] \\ &= E\left[\exp\left(-\int_0^t h_s ds\right)\right], \end{aligned} \quad (4)$$

utilising property (3) and the independence assumption². Differentiating (4) gives us an expression for the density of τ :

$$P[\tau \in dt] = E\left[h_t \exp\left(-\int_0^t h_s ds\right)\right] dt.$$

Once this is understood, deriving expressions for prices of various credit-sensitive instruments becomes a straightforward application of the arbitrage-pricing principle. For example, if we wish to find the time- t price $P_C(t, T)$ of a zero-coupon corporate bond with expiry T which delivers 1 at time T if there were no default before T and delivers δ_τ at time T if default occurred at time $\tau \leq T$, then we have simply

$$\begin{aligned} P_C(t, T) &= E_t\left[e^{-R_{tT}}(I_{\{\tau > T\}} + \delta_\tau I_{\{\tau \leq T\}})\right] \\ &= P(t, T) - E_t\left[e^{-R_{tT}}(1 - \delta_\tau)I_{\{\tau \leq T\}}\right], \end{aligned} \quad (5)$$

$$= P(t, T) - E_t\left[e^{-R_{tT}} \int_t^T (1 - \delta_s)h_s e^{-H_{ts}} ds\right], \quad (6)$$

where $P(t, T)$ is the time- t price of a riskless zero-coupon bond with expiry T , and $R_{st} \equiv \int_s^t r_u du$, $H_{st} \equiv \int_s^t h_u du$.

Expression (6) for the price of a risky zero-coupon bond appears in various places at various levels of generality; it appears in modified guises according to the assumptions

²It is worth emphasising that we *do* need independence here; to derive (4) we use the argument $P[T_1 > H(t)] = E[P[T_1 > H(t)|\mathcal{G}]] = E[\exp(-H(t))]$, where \mathcal{G} is a σ -field with respect to which h is measurable but which is independent of N . The independence assumption is key to a number of the expressions which are derived in the literature of the subject.

made about what happens on default (Is payment made immediately? Is the loss proportional to the value of the asset immediately prior to default?), and we shall shortly discuss some of the papers where it features. For the moment, though, notice that the key components of the pricing problem are to model the riskless interest rate, the timing of default, and the recovery process; and notice also that without some very strong assumptions about the dynamics of these processes, simple closed-form prices for corporate bonds are unlikely to arise.

Example 1. If the recovery process is identically zero, the price of the corporate bond becomes

$$P_C(t, T) = E_t[e^{-R_{tT}} I_{\{\tau > T\}}] = E_t[e^{-R_{tT} - H_{tT}}],$$

so what we see is like a riskless zero-coupon bond with spot rate $r + h$. We can now view the problem as similar to the problem of pricing index-linked bonds, or bonds denominated in a foreign currency, thus making an existing literature available. It seems however that we may get further by exploiting the additional structure of the credit interpretation.

Example 2. If we assume that the recovery process δ is constant, and that the hazard rate takes the constant value μ , then the price given by (6) for the corporate bond simplifies to

$$P_C(t, T) = P(t, T)(\delta + (1 - \delta)e^{-\mu(T-t)} I_{\{\tau > t\}}). \quad (7)$$

In this case, the credit spread (if $\tau > t$) is given simply by

$$(T - t)^{-1} \log(P(t, T)/P_C(t, T)) = \mu - \frac{1}{T - t} \log(1 + \delta(e^{\mu(T-t)} - 1)), \quad (8)$$

which is a decreasing function of $T - t$; if $\delta = 0$ then the spread is constant. The paper [11] presents a fairly general framework for credit modelling, and then specialises to this example with a Gaussian HJM interest rate model. The choice of the Gaussian HJM description obscures the simplicity and generality of their work, in my view.

Example 3. In a more recent work, [12] extend their earlier paper by considering a situation where the Vasicek interest rate process is used, and where the hazard function h_t is some linear function of r_t and Z_t , where Z is some Brownian motion which may be correlated with the interest-rate process. More specifically,

$$\begin{aligned} dr_t &= \sigma dW_t + \beta(r_\infty - r_t)dt, \\ dW_t dZ_t &= \rho dt, \\ h_t &= a_0(t) + a_1(t)r_t + a_2(t)Z_t. \end{aligned}$$

As with the Vasicek model itself, the use of a process which may take negative values for the intrinsically non-negative process h is questionable, but if we close our eyes to this problem, then simple formulae result for the price of the corporate bond:

$$P_C(0, T) = \delta P(0, T) + (1 - \delta) \exp(-\mu T + \frac{1}{2}v_T),$$

where μ_T and v_T are the mean and variance of $R_{0T} + H_T$ respectively:

$$\begin{aligned}\mu_T &= \int_0^T \{(1 + a_1(s))(r_\infty + e^{-\beta s}(r_0 - r_\infty) + a_0(s))\} ds \\ v_T &= 2E \int_0^T ds \int_s^T dv \left[(1 + a_1(s))(1 + a_1(v))e^{\beta(s-v)} f(2\beta, s) + \rho a_2(v)(1 + a_1(s))f(\beta, s) \right. \\ &\quad \left. + \rho a_2(s)(1 + a_1(v))e^{\beta(s-v)} f(\beta, s) + s a_2(s) a_2(v) \right],\end{aligned}$$

where $f(\lambda, t) = (1 - e^{-\lambda t})/t$. The freedom to choose the three functions a_i gives a great deal of flexibility in fitting the model, and the involvement of the spot rate and another Brownian motion (for which Jarrow and Turnbull offer the interpretation of the log of some index price) certainly incorporates a desirable dependence of credit risk on economic fundamentals.

Example 4. This example (see [10]) is again in the same spirit as the earlier Jarrow & Turnbull paper, with various structural assumptions on the hazard-rate process. The idea is to model moves between credit classes as a time-homogeneous Markov chain X in the real-world measure (which facilitates estimation from historical data). One then assumes that in the pricing measure the riskless rate and the transitions are *independent*, and additionally that the recovery rate is *constant*. This leads to a neat formula for the price of risky bonds:

$$\begin{aligned}P_C(t, T) &= E[e^{-R_{tT}}(\delta + (1 - \delta)\tilde{P}(\tau > T | X_t, \tau > t))] \\ &= P(t, T) - E[e^{-R_{tT}}(1 - \delta)]\tilde{P}(\tau > T | X_t, \tau > t) \quad (9)\end{aligned}$$

$$= P(t, T) - (1 - \delta)P(t, T)\tilde{P}(\tau > T | X_t, \tau > t). \quad (10)$$

The probability \tilde{P} is the law governing the Markov chain of credit class transitions under the pricing measure. The link between the law of X under the two measures is achieved by assuming that the intensity matrix $\tilde{Q}(t)$ in the pricing measure may be expressed as $\tilde{Q}(t) = U(t)Q$, where $U(t)$ is diagonal, and Q is the Q -matrix in the real-world measure.

It is worth following through in some detail the steps of the analysis, because among all reduced-form models, this one is making perhaps the most sophisticated use of the most readily-available information about the riskiness of a firm's debt, namely credit ratings. The difficulties we encounter along the way will arise in any similar model.

Our goal is to estimate the model for the default process under the pricing measure, and the expression (9) for the price of a risky zero-coupon bond is the starting point. We assume that we know the riskless bond prices $P(t, T)$; finding these from market data is a non-trivial but well-studied problem. Next we need to know the risky zero-coupon bond prices $P_C(t, T)$. These are harder to tease out of market data, because most bonds are coupon-bearing, and many have convertible features. The procedure advocated by Jarrow, Lando & Turnbull goes as follows:

- separate bonds into buckets by maturity and credit class;
- within each bucket, compute the market-value-weighted average (MVWA) coupon, and the MVWA yield-to-worst³;
- treat each bucket as if it were a single bond with the MVWA coupon and MVWA yield-to-worst, and recursively compute $P_C(t, T_i)$, $i = 1, \dots, n$, from these synthesised bond prices. The treatment of convertible bonds is rather crude, and Jarrow, Lando & Turnbull find that the procedure sometimes results in a lower-rated bond being worth more than a higher-rated one! They comment that this problem is accentuated when there are comparatively few bonds in a bucket (apart from A and BAA1 grades, few of the buckets contain more than 30 bonds, and in many cases the number is less than 10).

Having got this far, there remains only the estimation of δ between us and estimates of the default probabilities in the pricing measure. The method used in the paper is to take for δ the MVWA of recovery rates over all classes of debt in 1991. This takes the value 0.3265; the values for the five classes of debt are 0.6081, 0.4550, 0.3368, 0.1658, and 0.0363, which we see varies very considerably, so this is a significant simplification. We now are able to use (2) to give us estimates of $\tilde{P}(\tau > T_j | X_t = i, \tau > t)$ for a range of maturities T_j , and for each credit class i .

If we knew the jump intensities $Q \equiv (q_{ij})$ between credit classes, we could compute the matrix of transition probabilities over time Δt as $\exp(Q\Delta t)$; assuming that Δt is small enough that we can ignore the possibility of more than one jump, we then have an approximation for the Δt -transition probabilities given by

$$p_{ij}(\Delta t) = q_{ij}(1 - e^{-q_i\Delta t})/q_i, \quad (i \neq j) \quad (11)$$

where $q_i = \sum_{j \neq i} q_{ij}$. The Standard and Poor's Credit Review provides an estimate of the one-year transition probabilities, and using these for the left-hand side of (11) it is easy to deduce the values of q_{ij} corresponding. This then deals with the estimation of the transitions between credit classes in the real-world probability, and now it remains to estimate the transformation from real-world to pricing probability.

This last step must of course use information from prices, and we use the risky zero-coupon bond prices $P_C(t, T)$ for each of the credit classes, and each of the maturities T_j . Using (2), we transform this into $\tilde{P}(\tau > T_j | X_t = i, \tau > t)$ for each j , and each credit class $i = 1, \dots, K$, where credit class K is the default state. Now recall that we are going to write the jump-rate matrix $\tilde{Q}(s)$ in continuous time as $U(s)Q$ for some diagonal matrix $U(s)$, and so the transitions in the pricing probability will be approximately

$$\tilde{P}(X_{s+\Delta t} = j | X_s = i) \doteq \delta_{ij} + \Delta t U_{ii}(s) q_{ij}.$$

³For a non-convertible bond, this is the yield; for a convertible bond, it is the yield calculated under the assumption that the bond will be called at the earliest allowable date.

Supposing that we knew the diagonal matrices $U(T_j)$ for $j = 1, \dots, m - 1$, we would then know $\tilde{P}(X_{T_m} = j | X_t = i)$, and so we could use the identity

$$\tilde{P}(\tau \leq T_{m+1} | X_t = i, \tau > t) = \sum_k \tilde{P}(X_{T_m} = k | X_t = i) \{ \delta_{kK} + (T_{m+1} - T_m) U_{kk}(T_m) q_{kK} \}$$

to find the unknown $U_{kk}(T_m)$ - we have $K - 1$ linear equations in $K - 1$ unknowns. This way, we build up recursively the estimates of the transition rates between states in the pricing probability, and can in principle answer any credit-sensitive pricing question in this framework.

There are several features of this modelling approach which pose problems (most of them signalled by Jarrow, Lando & Turnbull in their paper):

- By inspection of (2), we see that the ratio $P_C(t, T)/P(t, T)$ of the price of risky to riskless zero-coupon bonds depends only on t, T , and the current credit class. This seems an improbable feature, and disappears in the extension of [5], who allow the recovery rate to be random and correlated with the assumed Vasicek term structure.
- It appears hard to deal realistically with convertible bonds. There are also problems related to estimation issues:
 - The estimation of risk premia described above actually leads to some extremely negative values of $U_{kk}(t)$, so Jarrow, Lando & Turnbull find that it is better to make a best-fit estimate subject to the constraint that all the $U_{kk}(t)$ are non-negative. This certainly cures the negative values problem, but we end up (of course!) with zero values for some of the $U_{kk}(t)$ - in fact, for quite a lot of them - which would have the unacceptable consequence that transitions out of some classes in some years would be impossible. In particular, a AA-rated firm would stay AA rated after the third year of their study going out 14 years, which seems difficult to accept.
 - Can we accurately estimate the transition intensities of the Markov chain? If we have a Poisson random variable, and we want to be 95% certain that we know the mean of that random variable to within 5%, we would need the mean to be of the order of 1500. In terms of transitions between credit classes, this is quite a large number, and in terms of defaults of investment-grade bonds it is a very large number! If we had observed 100 changes of credit class of a certain type, we would be 95% certain only that we knew the transition rate to within about 20%.

In addition to these features of the chosen modelling framework, that framework itself is open to question:

- Are transitions between credit classes really governed by a Markov chain? If so, then we would see that the times spent in different credit classes would have exponential

distributions independent of the jumps, and there would be no tendency for a company to continue to fall through credit classes, contrary to some empirical evidence.

- Can we justify the assumed independence of the ratings transitions and everything else in the pricing probabilities?

Despite these difficulties, the approach is a sensible attempt to make use of widely-available credit ratings to model the default of corporate bonds.

Example 5. [8] assume in contrast to the situation in [11] that at the moment τ that default occurs, the corporate bond loses a fraction L_τ of its value. Denoting the hazard rate for default by h_t , and the payment at the maturity T of the bond by X , they find that the value at time $t < \tau$ of the bond is given by

$$S_t = E_t \left[\exp \left(- \int_t^T (r_s + h_s L_s) ds \right) X \right]. \quad (12)$$

Duffie & Singleton present a proof of this result using Itô's formula for jumping processes, but this is unnecessarily complicated. Firstly, observe that if the fraction lost on default were 1, the expression for the bond price if $t < \tau$ is

$$E_t \left[\exp \left(- \int_t^T r_s ds \right) X I_{\{\tau > T\}} \right] = E_t \left[\exp \left(- \int_t^T (r_s + h_s) ds \right) X \right]. \quad (13)$$

This establishes the result (12) in the special case $L \equiv 1$. Now suppose that the default time happens exactly as before, at intensity h_t , but that now when default happens at time t , with probability L_t the bond becomes worthless, while with probability $1 - L_t$ the value of the bond is unchanged. It is clear that the pre-default value of the bond is not changed by this way of thinking; just prior to default, the expected value of the bond is $S_{\tau-}(1 - L_\tau)$ in either case. However, we now can think of two types of default, harmless (with intensity $h_t(1 - L_t)$), and lethal (with intensity $h_t L_t$). As far as valuing the bond prior to default is concerned, we may simply ignore the harmless defaults, and price using the intensity hL of the lethal defaults. This reduces the problem to the simple situation where the bond loses all value on default, which we solved at (13). As Duffie & Singleton observe, the model does not allow for the effects of h and L separately, only for the product hL ; estimation of the two terms would require other data.

Duffie & Singleton offer various forms for the 'adjusted' default-rate process $r + hL$ (which they also allow may include a spread for convenience yield.) In a subsequent paper (1997), they examine the situation for an affine diffusion model in some depth, using interest-rate swap data for credit-risky counterparties.

Note that it is *essential* that we have independence of the Poisson process governing default, and the intensity h and loss-on-default L ; if this were not the case, any processes h and L which agreed up to the default time could be used, and the value of (12) could be varied at will!

In this approach, the bond loses L_τ of its value on default, which contrasts with the assumption of Jarrow & Turnbull mentioned earlier, namely that on default the bond is replaced with $1 - L_\tau$ riskless bonds with the promised payout; under which assumption will the price of the bond be larger?

Summary of the reduced-form approach.

- The existence of convertible bonds really forces one to consider firm value - so maybe we should go for a structural approach anyway?
- Bucketing complicates the estimation procedure. If we allow default rates to depend on economic fundamentals and certain gross features of the firm, then we may well end up estimating fewer parameters - and in particular, making some structural assumptions valid for all firms, the estimates are based on the whole sample, which would be advantageous for AAA, where the credit event data is so scarce.
- Modelling the moves between credit classes as a fundamental process leads to issues of estimation and interpretation. Perhaps it would be better to regard the credit class as a noisy observation of some more informative underlying process describing the credit-worthiness of the firm, and then to use a filtering approach.

3 Structural models.

The hallmark of a structural model is some attempt to model the value of the assets of the firm, and deduce the value of corporate debt from this. The paper of [16] is the first and simplest approach of this kind which we shall discuss.

Example 1. The model of Merton assumes a fixed rate of interest $r > 0$, and that the value V_t of the firm's assets at time t may be described by

$$dV_t = V_t(\sigma dW_t + r dt). \quad (14)$$

It is assumed that the firm is financed in part by the issue of bonds, and the face value B of the bonds must be repaid in full at time T . The shareholders are not allowed to pay dividends nor issue debt of equal or higher rank in the meantime. At time T , the bondholders will receive $\min\{V_T, B\}$, so the value of the bonds at time $t < T$ will be simply

$$E_t[e^{-r(T-t)} \min\{V_T, B\}] = Be^{-r(T-t)} - P(t, V_t, B)$$

where $P(t, V_t, B)$ is the value at time t of a put option with strike B if the current value of the firm's assets is V_t . But this is just the familiar Black-Scholes formula:

$$Be^{-r(T-t)}\Phi(-d_2) - V_t\Phi(-d_1),$$

where Φ is the cumulative distribution function of the standard normal distribution, and

$$\begin{aligned} d_1 &= \frac{\log(V_t/B) + (r + \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}, \\ d_2 &= \frac{\log(V_t/B) + (r - \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}} = d_1 - \sigma\sqrt{T-t}. \end{aligned}$$

The spread on corporate debt is

$$-\frac{1}{T-t} \log\left[\Phi(-d_2) - \frac{1}{d}\Phi(-d_1)\right],$$

where we have written $d \equiv Be^{-r(T-t)}/V$ for the debt-equity ratio, expressed in terms of the current value of the debt. It is easy to see that in fact the spread depends only on d , the time and the volatility. Merton studies the comparative statics of this model, and shows among other things that the spread is a decreasing function of maturity if $d \geq 1$, but for $d < 1$ it is humped.

Example 2. In one of the most intellectually satisfying papers in the literature, [13] consider the impact of the maturity of debt on the optimal exercise of the default option by the shareholders. The assumptions of the model are:

- constant interest rate r ;
- the value V_t of the firm's assets at time t evolves as

$$dV_t = V_t(\sigma dW_t + (r - \delta)dt),$$

where δ is the constant rate of dividends paid to the shareholders, and σ is a positive constant;

- upon default, a fraction α of the value of the firm is lost through restructuring;
- there is a constant rolling debt structure, with total outstanding principal of P and maturity T , with new debt being issued (and old debt retired) at rate P/T , and coupons being paid continuously at rate C annually;
- tax benefits accrue at rate γC on the coupon payments;
- the shareholders declare bankruptcy when the value of the firm's assets falls to V_B .

The value of the firm is given by the expression

$$v(V, V_B) = V + \frac{\gamma C}{r} \left[1 - \left(\frac{V}{V_B} \right)^{-x} \right] - \alpha V_B \left(\frac{V}{V_B} \right)^{-x}, \quad (15)$$

where x is the larger root of $\sigma^2\theta^2/2 + (r - \delta - \sigma^2/2)\theta - r = 0$. Noticing that $E \exp(-r\tau) = (V/V_B)^{-x}$, we may interpret the three terms in (15) as the value of the firm's assets, the net present value of all future tax refunds, and the net present value of the loss on default. A bondholder who will receive a coupon at fixed rate c , and will be repaid p at time t provided this was before default, but who receives ρV_B at the default time (if this was earlier than t) has an asset worth

$$d(V, V_B, t) = \int_0^t c e^{-rs} [1 - F(s)] ds + e^{-rt} p [1 - F(t)] + \int_0^t e^{-rs} \rho V_B F(ds), \quad (16)$$

where F is the distribution function of the default time, which depends of course on the values of V and V_B (in fact, only through their ratio). The total value $D(V, V_B, T)$ of the firm's debt is obtained by integrating (16) from 0 to T , using $c = C/T$, $p = P/T$, and $\rho = (1 - \alpha)/T$. The value of the firm's equity is therefore the difference

$$\begin{aligned} eq(V, V_B, T) &= v(V, V_B) - D(V, V_B, T) \\ &= v(V, V_B) - \int_0^T d(V, V_B, t) dt. \end{aligned}$$

A closed-form expression is available for D in terms of the normal distribution function.

The level V_B is determined endogenously as the level which maximises the value of equity subject to $eq \geq 0$, and Leland & Toft obtain a closed-form expression for V_B . Assuming that the coupon on debt is chosen so that new debt is issued at par, they go on to examine various comparative statics of the optimal solution, and they find (among other things) that:

- the longer the maturity of the debt, the higher the value of the firm, and the greater the optimal leverage;
- bond values are humped for low to moderate leverage, but for high leverage the bond sells below par for long time to maturity and above for short time to maturity, the effect becoming more pronounced as T increases;
- the credit spreads are increasing with T for low leverage, but become humped for moderate to large leverages;
- credit spreads for values of T up to 2 years are negligible.

Example 3. [13] assume that the interest rate is constant, which is a reasonable assumption in order to get insight into the influence of various effects, but the assumption of constant interest rates is too restrictive for a working model. [15] embrace the possibility of stochastic interest rates, modelling the spot rate as a Vasicek process correlated with the log-Brownian share price process. They assume that there is some threshold value K such that if the value of the firm ever falls to that level, then restructuring takes place,

and the bond is replaced by $(1 - w)$ riskless bonds of the same maturity. Longstaff & Schwartz derive an expression for the price of the risky bond, but their derivation contains a flaw; they apply results of [3] concerning the first-passage distributions of one-dimensional diffusions to the log of the discounted firm value, but this process is not a diffusion. It appears, therefore, that the pricing of a corporate bond in this modelling framework remains an open question.

Example 4. [2] consider a variant of the problem dealt with by Merton; control of the firm passes to the bondholders not only if the value of the firm is below some value B at the maturity T of the debt but also if in the meantime the value of the firm falls below some value (which depends on time as $Ce^{-\gamma(T-t)}$). They derive a closed-form expression for the value of the corporate bond, under the assumption of zero restructuring costs on default. They also derive the values of two bond issues, the senior and junior bonds, by identifying the prices in terms of the solution to the first problem.

Example 5. The KMV method for pricing risky debt relies on a structural-type approach. The description that follows is vague, not least because the details of the methodology are proprietary. The value of the firm's assets are modelled by a log-Brownian motion, $V_t = V_0 \exp(\sigma W_t + (m - \sigma^2/2)t)$, and the probability of default at time T is the probability that the value of the firm does not cover the liabilities K of the firm at that time, namely,

$$\Phi(-d_2)$$

where

$$d_2 = (\log(V_0/K) + (m - \sigma^2/2)T) / \sigma\sqrt{T}$$

is the so-called *distance to default*. In common with other structural approaches, the estimation of the parameters is a difficult matter, and the identification of the equity as a call option on the value of the firm allows an estimate of the volatility to be made. The total liabilities and market value of equity need to be observed or estimated. It is not clear how m is determined. The distance to default is used in conjunction with empirical data on the relation of defaults to the distance-to-default to estimate the probability of default.

The use of widely-available equity price data is an appealing feature of this approach (though this would render it unsuitable for pricing sovereign debt). The assumption of constant interest rates is a limitation also.

Example 6. Another structural approach to credit risk is given by [14], who assume that the value of the firm's assets obeys the SDE

$$dV_t = V_t(\sigma dW_t + (\alpha - \gamma)dt)$$

for constants σ , α , and γ . They assume that the interest rate process is a Cox-Ingersoll-Ross model, and that the bond-holders must be paid coupons at constant rate

c. Bankruptcy is triggered when the cash flow γV_t from the firm is no longer sufficient to cover the coupon payments which have to be made, that is, when V drops to c/γ . The authors compute values of convertible and non-convertible bonds in this model, and assert that the spreads which result are consistent with market values.

4 Some nice ideas.

This short section gathers some neat ideas which do not fit obviously in any of the preceding sections.

The paper of [9] contains some simple but attractive ideas for dealing with credit risk. They present a general characterisation of the time- t price of some risky contingent claim paying off X at time T if there is no default before time T , and otherwise paying $\delta_\tau Y_\tau$ at default time τ , where Y_t is the time- t price of a riskless asset paying X at time T . Their expression is

$$E_t [e^{-R_{tT}} w_{tT} X], \quad (17)$$

where w_{tT} is the expectation of δ_τ conditional on the interest-rate process between t and T , and on the final contingent claim X . Of course, this is too general to be of much use as such; we could think of this expression as an alternative description of the price in a hazard-rate model, so until we have been much more specific about the hazard rate, we can go no further. Nevertheless, Hull & White use (17) quite effectively to bound the price of a credit-risky call option on a log-Brownian stock, assuming constant interest rate. In this situation, we shall have that w_T is a function of S_T , $w_T = u(S_T)$. The price of the risky option is

$$e^{-r(T-t)} E_t [u(S_T) (S_T - K)^+], \quad (18)$$

which has to be consistent with the market price of the credit-risky zero-coupon bond of the call writer,

$$P_C(t, T) = e^{-r(T-t)} E_t [u(S_T)]. \quad (19)$$

Since $0 \leq u(S) \leq 1$, we maximise the value of the risky option by taking $u(S) = I_{\{S > s_1\}}$ for a constant s_1 chosen to make (19) hold, and we minimise it similarly by taking $u(S) = I_{\{S < s_2\}}$ for suitable s_2 . Numerical examples using a call writer with log-Brownian asset value process and bankruptcy when the value falls to some trigger level show that these bounds are not very tight, but perhaps by incorporating the information from other market prices they could be improved.

Hull & White also remark that a credit-risky American option will be exercised no later than its credit-risk-free counterpart; the reason is easy to see on a moment's reflection.

For a very quick and dirty approach, Hull & White also discuss the situation where the default process is independent of the interest-rate and the payoff of the contingent claim

in the pricing probability. Then using the market prices of credit-risky and default-free bonds, it is immediate that

$$P_C(t, T)/P(t, T) = E(w_{tT}), \quad (20)$$

and so the price of the risky contingent claim would be simply

$$Y_t P_C(t, T)/P(t, T),$$

where Y_t is as before the time- t price of the default-free contingent claim. This approach can be extended to deal with swaps by using (20) for a range of values of T .

One neat idea, to be found in [1] and [19], is to try to hedge out all credit risk, using a single credit-sensitive instrument. The idea is very simple. If your portfolio is vulnerable to default of a counterparty, and if there is a liquid asset which is also sensitive to the default of the same counterparty, then you take up a dynamically-adjusted position in the liquid asset so that upon default, the loss to your portfolio is zero. Thus you choose the holding of the liquid asset to exactly cancel out the loss that the rest of your portfolio will make on default. This done, there are no jumps in the value of your combined portfolio and (under Brownian market assumptions) you may therefore hedge the combined portfolio perfectly.

As a parting remark, it may be of interest to note that formally the reduced-form approach may be thought to include the structural form approach, in that the default intensity becomes infinite at the moment that the asset price in the structural description reaches the default boundary. This does not (of course!) mean that we can throw away the structural approach.

5 Summary.

Each of the two main classes of approach has its strengths and weaknesses. For the *structural approach*, we have:

- a clear link between economic fundamentals and defaults. This helps to understand losses on default, and the correlation of defaults of different firms;
- reliance on economic fundamentals and the value of the firm's assets which may be hard to estimate with any accuracy;

On the other hand, features of the reduced-form approach are:

- a model which is sufficiently close to the data that it is always possible to fit *some* version of the model;

- the fitted model may not perform well ‘out of sample’;
- in the case of proportional losses, it is hard to distinguish the hazard rate and the percentage loss on default;
- pricing of convertible bonds does not fit well into this framework.

Where might the modelling of credit risk be going now? Within the reduced-form framework, it seems that there is little one may do except explore further parametric forms of the intensity and loss-on-default processes. In the structural approach, we need to incorporate jumps in the value of the firm in a reasonable way, and we need to develop a filtering approach to the estimation; realistically, we cannot assume that we know the value of the firm with precision, nor how its rate of return will depend on the economic fundamentals, so we have to confront that uncertainty honestly. Ultimately, the quality of what we can create will be constrained by the quality of the data to calibrate it, so we probably should not be trying to do anything too sophisticated!

6 Epilogue.

... and yet we have! Since 1999, we have seen the improbable rise and overdue demise of the ‘industry-standard’ Gaussian copula, we have seen the publication of excellent monographs on the subject, such as [18], and vigorous development of new modelling ideas. We have, for example, seen developments of filtering ideas, such as in [7], and in works of Monique Jeanblanc and her co-workers; and we have seen enormous effort expended in modelling and fitting CDOs and other derivatives which depend on the defaults of more than one name. Most attention has focused on the reduced-form approach, and while it is regrettable that we are not able to offer a structural story that really works in practice, it is inevitable that we will not be able to build something that could handle the complexities and heterogeneities of the practical world. Ease of calibration has been the over-riding consideration, and many quite strange models (in reality, fits) have been pressed into service because they were easy to align to market data, without possessing any other virtue, such as intertemporal consistency.

This is an area which I have stood back from since these notes were first prepared, with one exception, the paper [6]. The approach adopted there is I believe quite simple; it is self-consistent; it can handle corporate and government debt in one model; and it can in principle embrace FX also, and therefore offers a sensible approach to hybrid pricing. Indeed, the whole modelling approach can be thought of in terms of the potential approach [17].

In the markets, we have experienced the liquidity drought following from the sub-prime

fallout of the summer of 2007, exacerbated by the profusion of repackaged credit risks, and understanding these presents huge challenges for the industry and the academic profession. Credit risks are more similar to insurance risks than to market risks, as the range of useable hedging instruments is far more restricted, and in the end if everything goes sour at once, no portfolio of market instruments will save you. Ratings agencies have been criticised in the wake of the sub-prime disaster, and I feel that some more rational characterisation of corporate credit-worthiness is required. The very granular nature of credit ratings causes corporations a lot of difficulty if a downgrade should occur, and probably some index of credit-worthiness which is real-valued would make more sense. Even better would be some index of exposure of the firm to a number of major macroeconomic indicators; this would prevent the meaningless comparison of one company being ‘more risky’ than another.

References

- [1] J.G.B. Beumee, B. Hilberink, and M.H. Vellekoop. Pricing and hedging options on defaultable assets. 2001.
- [2] F. Black and J.C. Cox. Valuing corporate securities: some effects of bond indenture provisions. *Journal of Finance*, 31:351–367, 1976.
- [3] A. Buonocore, A.G. Nobile, and L.M Ricciardi. A new integral equation for the evaluation of first-passage-time probability densities. *Advances in Applied Probability*, 19:784–800, 1987.
- [4] L. Cathcart and L. El-Jahel. Pricing defaultable bonds: a middle way approach between structural and reduced form models. *Quantitative Finance*, 6:243–253, 2006.
- [5] S.R. Das and P. Tufano. Pricing credit-sensitive debt when interest rates, credit ratings, and credit spreads are stochastic. *Journal of Financial Engineering*, 5:161–198, 1996.
- [6] G. Di Graziano and L.C.G. Rogers. A dynamic approach to the modelling of correlation credit derivatives using markov chains. Technical report, University of Cambridge, 2005.
- [7] J.D. Duffie and D. Lando. Term structure of credit spreads with incomplete accounting information. *Econometrica*, 69:633–664, 2001.
- [8] J.D. Duffie and K.J. Singleton. Modelling term structures of defaultable bonds. *Review of Financial Studies*, 12:687–720, 1999.
- [9] J. Hull and A. White. The impact of default risk on the prices of options and other derivative securities. *Journal of Banking and Finance*, 19:299–322, 1995.

- [10] R.A. Jarrow, D. Lando, and S.M. Turnbull. A markov model for the term structure of credit risk spreads. *Review of Financial Studies*, 10:481–523, 1997.
- [11] R.A. Jarrow and S.M. Turnbull. Pricing derivatives on financial securities subject to credit risk. *Journal of Finance*, 50:53–85, 1995.
- [12] R.A. Jarrow and S.M. Turnbull. The intersection of market and credit risk. *Journal of Banking and Finance*, 24:271–299, 2000.
- [13] H.E. Leland and K. Toft. Optimal capital structure, endogenous bankruptcy, and the term structure of credit spreads. *Journal of Finance*, 51:987–1019, 1996.
- [14] I.J. Lim, K. Ramaswamy, and S. Sundaresan. Does default risk in coupons affect the valuation of corporate bonds?: a contingent claims model. *Financial Management*, 22:117–131, 1993.
- [15] F. A. Longstaff and E.S. Schwartz. A simple approach to valuing risky fixed and floating rate debt. *Journal of Finance*, 50, 1995.
- [16] R.C. Merton. On the pricing of corporate debt: the risk structure of interest rates. *Journal of Finance*, 29:449–470, 1974.
- [17] L. C. G. Rogers. The potential approach to the term structure of interest rates and foreign exchange rates. *Mathematical Finance*, 7:157–176, 1997.
- [18] P. Schoenbucher. *Credit Derivatives Pricing Models:Models, Pricing, Implementation*. Wiley Finance, 2003.
- [19] D. Wong. *A unifying credit model*. PhD thesis, Carnegie-Mellon University, 1998.