Modelling Corporate Bonds
Considerations for Stochastic Modelling

by Jon Exley & Andrew Smith
June 2002

Abstract:

Many of the newly issued corporate bonds are finding their way into insurance and pension fund asset portfolios. Modelling these bonds is tricky. Historic return distributions are well behaved and the observed credit spreads exceed all estimates of historic default probabilities. As a result, corporate bonds may appear anomalously attractive. We show that some of the more widely used models of corporate bonds have some major weaknesses from a theoretical perspective and in particular fail to explain this apparent anomaly convincingly. However, our analysis of the economics of corporate debt early in the paper suggests that credit spreads are option premiums and thus have option-like characteristics. This leads to the Merton Model as a means of capturing these characteristics. The model reveals that the apparent anomaly in credit spreads is another facet of well known apparent anomalies in option prices under the standard Black Scholes Model (such as the existence of skews and smiles in implied volatility). The latter anomalies are not generally regarded as exploitable in asset allocation studies. Armed with the insights from this model, suitable for heavy duty application in pricing contingent claims based on default, we return to the simpler problem of modelling corporate bond portfolios. We suggest that complexity in these models inevitable, given that a rebalancing investment *process* must be modelled and that there is a fundamental arbitrariness in the relevant grade categories on which the process is based. Having considered these issues we suggest that, despite its sledge-hammer appearance the Merton Model provides a solid foundation for alternative approaches to modelling such portfolios that ensures consistency and reduces the risk of anomalies. In particular we reconcile well-known approximations in terms of special cases of this model.

1. INTRODUCTION

1.1 The last few years have seen record levels of corporate debt issuance in the UK and elsewhere in Europe. Some of these have found their way into insurance and pension funds, where Actuaries have to figure out how to model them.

1.2 There is as yet no perfect or generally accepted technique for modelling corporate bonds. In this paper we seek to outline the chief strengths and weaknesses of existing methods. It is certain, given the pace of development in this area, that our paper will quickly become out of date.

1.3 The remainder of this paper falls into four main sections. In section 2, we investigate the economics behind the existence of corporate bonds. Why are companies not financed entirely by equity? How can we use corporate capital structures to explain some relevant features of corporate debt?

1.4 The third section looks at some properties of the corporate bond market. These include rating categories, default probabilities and yield spreads. We also examine the kind of models, which might explain the relationships we see between these quantities.

1.5 Our fourth section is taken up entirely with consideration of Merton's famous corporate bond model, and its various generalisations.

1.6 In our fifth section, we offer some practical tips on the modelling corporate bond portfolios for asset-liability modelling purposes.

1.7 We wrap up with our conclusions in the sixth and final section.

2. WHY DO CORPORATE BONDS EXIST?

2.1 A recent SIAS paper by Paul Sweeting, entitled "The Role of High Yield Corporate Debt in Pension Schemes" gives a breakdown of market size, growth and other statistics. We do not reproduce this material here, but instead refer the reader to Paul's paper at:

http://www.sias.org.uk/papers/pensiondebt.pdf

2.2 Corporate financial theory and what it has to say about debt versus equity finance is the subject of many textbooks (references eg Brealey & Myers, Copeland & Weston). We do not intend to repeat all of the theory here. However, an introduction to some of the basic issues seems like a good starting point and is important for an understanding of the model descried later. In particular, there are a number of common fallacies about debt finance, its "cost" and why companies issue debt in the first place that can get in the way of the construction of a good model.

2.3 Why Do Companies Issue Debt?

Perhaps the most basic question is why would a company issue debt in the first place? After all, on the face of it there is no reason why a company could not be 100% equity financed.

2.4 Let us start with the state of understanding pre 1950. Economists had started to observe that equities tend to return more than bonds over long time horizons. So if equities return more than bonds over such horizons, perhaps that would make equities an *expensive* way of financing the company, relative to bonds? Would companies create value by financing to the maximum extent by "cheap" debt?

2.5 The riddle of whether equity or debt was the cheaper way to finance a firm is still an area of debate. Indeed, when applied "the other way round" to pension funds, the mirror image of the same riddle is still just as contentious. In 1958, Modigliani and Miller made the first significant breakthrough in this area. Their first proposition is that, ignoring tax, bankruptcy and other second order effects (which we will discuss below), the *value* of a firm is independent of the way it is financed.

2.6 We can argue their proposition using a thought experiment. Consider a firm that is 100% equity financed, and suppose an investor owns all of the equity so he is entitled to all future profit cash flows from the business. Now suppose that the firm substitutes equity for debt, and the investor holds all of the equity and debt. If the underlying business is unaffected by this financing (discussed later) then clearly the same profit cash flows are simply divided between equity dividends and debt coupons – but in aggregate they are the same cash flows. Ignoring tax, we can argue by arbitrage that therefore the value of the equity plus debt under the new structure must equal to the value of the equity under the old structure. Q.E.D. (It is easily seen that the same result holds if the investor owns a given fraction of equity originally, and equity is swapped for debt.)

2.7 So although it is common to hear debt finance justified in terms of its effect on "cost of capital" or "earnings per share", it should always be borne in mind that at a basic level, the *value* of the company should be unaffected by capital restructuring. Likewise the (aggregate) "cost of capital" is also unaffected. For example, consider a company that is 100% equity financed, with a beta of 1. If we substitute half of the equity with debt, and ignore tax and bankruptcy for the time being, then all of the firm's risk is borne by the remaining equity, which acquires a beta of 2. (We invoke the Capital Asset Pricing Model to explain the issue, not because it is relied upon in the proof of the Modigliani & Miller proposition). The aggregate cost of capital is then as follows, if i is the risk free interest rate and e is the equity risk premium:

	Equity Beta	Aggregate Cost of Capital
100% Equity Financed	1.00	$i + e$
50% Equity Financed	2.00	$0.5i + 0.5(i + 2e) = i + e$

2.8 So this tells us nothing about why firms gear, and we are no nearer to understanding why firms issue debt. However, so far we have ignored the possibility of bankruptcy, perhaps if we explore this in more detail we can understand more about debt issue.

2.9 Understanding Bankruptcy from Equity Holder Perspective

The interesting aspects of corporate bonds arise because of the possibility of bond default. The difficult modelling aspects relate to yield spreads and to default probabilities. We will see that the precise sequence of events prior to default has a profound impact on the value of the debt, and on the reason for raising debt. Initially, we will focus on a simple model in which dividends and interest payments are ignored. We develop the possible sequence events in the form of a traditional good news – bad news joke, from a shareholder's perspective.

2.10 **Bad News**. Borrowings need to be repaid, and these must come out of shareholder wealth. Consider a firm with current value (equity and debt) of S_0 and future value S_1. The debt is K. The existence of the debt means that the equity holder's stake is not S_0 but S_0-K.

2.11 **Good News**. The shareholder has limited liability. He has an option to default. If the total assets S_1 are less than the debt K then the shareholder can walk away, leaving the assets in the hand of the bondholder. Thus, the shareholder receives $\max\{S_1-K, 0\}$ and the bondholder $\min\{S_1, K\}$. This is essentially the Merton model of bonds; it uses the Black-Scholes model to value the equity as net assets (S-K) plus a put option on the firm. Likewise, the bond is worth the debt K minus the same put option; the bond yield spread is the put premium which shareholders must pay to bondholders for the option to default.

2.12 **Bad News**. Bond holders can take control of the company at the moment of insolvency, at the instant when S falls to K. However, if this is the case then bondholders lose nothing at this point – the firm value is sufficient to repay debt in full. In other words, provided the company value moves continuously and does not jump, bondholders can avoid any loss by realising assets when S first touches K. At this point, there are still sufficient assets to pay bondholders. That implies the shareholder option to default is worthless. On the other hand, the fact that debt spreads in the market are positive suggests that debt holders do expect to lose on default – which begs two questions: How could bondholders lose, and who gains?

2.13 **Good News**. There may still be opportunities to pass losses onto bondholders. For example covenants may be based on accounting data that takes time to become available, and the process of gaining control may be lengthy. Furthermore, there may well be unreported losses ("black holes") in the firm accounts. Specifically, suppose the accounting assets are A, which is greater than the true value S. In the event that $A>K>S$, shareholders can continue to operate the firm, in the hope that S rises. Shareholders need only hand over control when A falls below K, because the test of bankruptcy is on an accounting basis and not an economic basis. Then the asymmetry is as follows: If the value of the firm S rises back above K then the equity holders keep control of the firm and the bond holders just sit and watch. On the other hand, if the value of the firm continues its downward path then the bondholders acquire the firm when A hits K, at which point S may be much lower. Even if the delay in taking control is short, this could generate quite a substantial loss, especially if we consider, realistically, a model where the volatility of the firm rises as its value falls, or where the firm's value can jump instantaneously downwards.

2.14 **Bad News**. While some companies run black holes, most firms trade at a premium to accounting assets, that is, $S>A$. This is because of intangible aspects of value, which are not recognised in financial statements. However, bondholder control is triggered when $A<K$. In this case, bondholders will be interested in maximising their entitlement; as they have no upside participation beyond their debt, bondholders may be reluctant to let a firm continue trading in the hope of future profits. In this way, default can destroy the intangible value. Notice that this is the first point in our argument where overall value destruction or creation enters the fray. Up until now, we have looked only at wealth transfer between shareholders and bondholders. Allowing for a loss of intangibles is an overall loss of aggregate shareholder and bondholder value.

2.15 **Good News**. Shareholders have an option to subscribe more capital. It will be optimal to do this in the event where accounting assets are insufficient to pay the debt but total value is more than sufficient. Then, assuming the company is penniless and has exhausted its credit lines, shareholders are faced with a choice of subscribing new capital of K to meet the debt repayment or losing E (ie passing control of the firm to the bond holders). Clearly, they should subscribe new equity capital as long as $E > 0$ (assuming only that they prefer more to less). Thus in simple terms, shareholders will only allow a firm to go bankrupt if it is worth nothing to them. By injecting more capital, shareholders can pay off bondholders in full yet still safeguard a portion of the intangible assets by continuing to support the firm as an ongoing concern. Equally well, shareholders are not obliged to subscribe more capital, and will not do so if the additional capital merely fills a black hole or perpetuates a value-destroying management strategy. Thus shareholders now have a three-way option to obtain the greater of S-K, A-K or 0. This aspect of gearing is sometimes called the value of monitoring – a more highly geared company will need to return more frequently to shareholders to endorse its continuing strategy, which controls management's ability to sponsor value-destroying investments.

2.16 **Bad News**. There are costs associated with subscribing new capital, in the form of investment bank fees, market impact and moral hazard (from asymmetric information; management have an incentive to misrepresent the investment opportunity). For these reasons, shareholders may be unable to capture intangible value in an accounting default – because the funds injected must be shared with bondholders, and the cost of refinancing exceeds the benefit.

2.17 Market Pricing

We should recognise that in this story, the good news is good news only to the extent that it is not recognised in the initial price of the debt. When a firm raises debt, at the same time various default options are granted to shareholders. Bondholders deduct the price of these options from the amount they are prepared to subscribe for the debt issue. Therefore, the yield spread is effectively an option premium.

2.18 This illustrates two points. Firstly, bankruptcy is not in itself a cost. Indeed, it is usually the shareholders, or managers acting on their behalf, who file for bankruptcy in order to gain protection from other creditors. Bankruptcy is a potentially valuable option, which is an asset to shareholders.

2.19 Secondly, the interest spread on risky debt is not a cost to shareholders, either. An institution whose cost of debt service has increased does not necessarily leave shareholders worse off. Although accounting standard may require a firm to expense the interest cost immediately, it is economically a premium, which the shareholders pay in return for an option. As with any other option, the premium is not a deadweight cost, merely the price of an investment.

2.20 The Advantages of Debt Issue

If we accept the argument that bond spreads accurately reflect shareholder premiums, then shareholders will not generally gain advantage from a high or low credit rating. Changes in interest cost merely reflect changes in the value of the corresponding options. Instead of looking for ways in which shareholders could short-change bond holders, or vice versa, we can explain the phenomenon of debt issuance only by considering possible third party effects which could advantage or disadvantage a firm's aggregate capital providers.

2.21 The most widely quoted valid reason for firms issuing debt is to exploit the tax shelter of profits used to pay debt interest. In broad terms, debt interest is drawn from pre tax profits whilst dividend payments are drawn from post tax profits.

2.22 Miller (1977) has pointed out that in reality the equation is more complex than this, since investors are interested in returns after personal tax. Although debt is privileged in terms of corporate taxes, it is heavily taxed in personal terms, whilst the reverse is true of equities. On the other hand, if the debt is held by pension funds, for example, the Miller argument is less clear, especially (in a UK context) after the abolition of ACT reclaim, which previously levelled the playing field partly. These issues have been discussed in detail in previous papers to the Finance and Investment Conference (see for example Bezooyen et al (1998)).

2.23 Once tax is dealt with, we are left with some more subtle, but still valuable benefits from gearing. These relate to the actual running of the business, with the argument that a certain level of debt leads to optimal management of a company. It is not our intention to describe each of these in detail, although we are keen not to appear to understate their importance as in a Miller (1977) tax indifference framework these are the *only* reasons for gearing a firm. However, in summary they are as follows:

- Agency Monitoring: It is easier to monitor management in a more highly geared firm; they have less access to shareholder reserves and must seek shareholder approval for new ventures, rather than using internal funds to pursue “pet” projects. It is also easier to give managers an equity stake in the business and thus incentivise value building activities(see Jensen & Meckling, 1976)

- Clientele Effects: Different investors have different asset preferences. By carving up profits into different risk characteristics, these preferences can be better catered for. For example, an insurer or pension fund may be able to tolerate the risk of bond default but less able to tolerate the swings in equity market values. Tax clienteles can be particularly important.

- In some cases, a firm may take advantage of third party guarantees, such as banking deposit protection schemes or insurance insolvency pools. To the extent that contributions to these pools are not priced according to risks, a firm can be a net beneficiary of the pool by being more than average risk. Gearing is one way to achieve this.

- The cost of capital includes the transaction costs of raining the capital in the first place. Investment banks typically charge a pro rata fee for capital raised. To the extent that the fees are lower on debt, this provides a reason for companies to gear.

2.24 Disadvantages of Debt Issue and Optimal Gearing Ratio

Given the above benefits of gearing, including (to a debatable extent) the tax benefit, we now find ourselves in the position of asking why firms are not more highly geared and why particular levels of gearing are chosen. For example, why aren’t firms 99% geared?

2.25 Of course, as the debt levels rose, the rating of the debt would fall and the credit spread rise. This leads us neatly to our final misconception about debt issuance – namely the belief that companies with higher rated debt are necessarily superior to those with lower rated debt. We need to distinguish between two distinct costs of issuing debt here. Firstly, there are default “costs” which are borne by bondholders, but for which they derive a premium in the form of a credit spread on the bond and accounted for as a higher “cost of debt capital”. As discussed above, this premium (i.e. the credit spread) is a benefit paid to bond holders and a cost to shareholders, which in theory should be equal to the value to shareholders of the option to default and its cost to bondholders. It must be stressed that in these terms there is nothing particularly desirable to shareholders about low credit spreads – the spread is a cost, but the option to default is a benefit. The two are on the face of it, equal and opposite.

2.26 In theory therefore, although the cost to bondholders of default increases the "cost of debt capital", it reduces the "cost of equity capital" by an equal amount. Thus the aggregate "cost of capital" is unaffected by these costs of default representing transfers between equity and debt holders. Once again we need to look to secondary costs to determine not only the optimal level of debt, but also to address how the degree of flexibility or rigidity in bond covenants might be chosen to maximise aggregate value (rather than simply affect the cost and value of the various options held by shareholders).

2.27 In order to probe these secondary costs, consider first some of the parties who might gain in the event of a default. Possible candidates might be accountants and lawyers (particularly fees earned whilst administering a Company), employees (redundancy costs) and the Inland Revenue (if the tax shelter of profits is lost by capital restructuring or capital gains are realised for example). However, empirical studies (see for example Altman & Kishore, 1996) show that the average losses to bond holders are between 45% (senior debt) and 68% (subordinated) of the face value of debt. Although the leakage to third parties may be large, it appears implausible to attribute all of the loss to these sources.

2.28 In summary, increasing the level of debt, or, in some cases, making debt covenants more onerous has at least four adverse effects on shareholder wealth:

- It restricts management flexibility – which according to Agency Monitoring theory is good up to a point. This can however destroy value if the rigidity reaches such a level that, for example, value building business opportunities are missed.
- To the extent that it increases the likelihood of bankruptcy, it increases the cost of leakage to third parties (lawyers, accountants, employees, Inland Revenue) associated with this event.
- Even without bankruptcy, a perception that a firm's balance sheet is less than robust can lead to a deterioration in business opportunities, especially for financial firms such as insurers and banks. The complexity of refinancing a marginally solvent company, and the fact that newly injected equity also benefits bondholders and other creditors, can lead to high negotiation costs.
- In addition, the firm value may comprise many intangible items (brands, customer relationships) which are not recognised on the balance sheet, but which are forfeit in the event of bankruptcy.

2.29 The optimal level of debt issuance will be a trade off between these costs and the benefits described above. In theory the financial structure, including covenants and debt grades, should be targeted to maximise firm value. Although company management often seek to maintain "prestige" credit ratings , in reality the firm might benefit from a lower credit rating and greater covenant flexibility or greater debt issuance.

2.30 Asset Substitution and Debt Term Structure

We conclude this section with a comment on some of the more subtle aspects of debt issuance. The options and counter-options we have considered so far are broadly understood and accepted by all parties. There are also some sneaky options which the parties seek to control, with some success, by statute and adherence to externally monitored standards.

2.31 The cunning reader may already be hatching a wheeze for shareholders here, which is known in the literature as "asset substitution". Essentially, a trick might be to issue a lot of debt as a staid, conservative firm and then change colours into a high-risk venture. In this way the equity holders pay a low credit premium for the option to default and then reorganise the firm to maximise the value of this option (in simple terms, by increasing volatility).

2.32 Indeed actuaries have a good example of asset substitution close to home – when a pension scheme is immature the scope for the sponsor to gamble with the risk of default is small. This is because the granting of new accrual is the same as issuing more pension promises to employees, which the employees will discount deeply if the company plays fast and loose with funding or investment policy. It is only when the need to issue more promises is removed (a very mature scheme) that the company is able to play a version of this asset substitution game with the pension fund.

2.33 Asset substitution is one of the aspects that covenants attempt to address, by identifying particular assets as security and triggering repayment of the debt in certain circumstances. The problem largely disappears with short dated debt, since the firm will find that subsequent attempts to tap the credit markets will encounter sharply increased spreads to compensate for the possibility of repeating the trick. Dealing with the risk of asset substitution remains a challenge for any model of credit risk. For example, models such as Leland (1994) have attempted to explain term structures of debt in terms of the insurance sought by bondholders against changes in firm volatility, with varying degrees of success.

2.34 Another sneaky aspect is the possibly sub-optimal behaviour of management in refinancing. Even where there is a positive franchise value to protect, poor risk management or a lack of contingency planning may mean that equity holders cannot inject new capital in time to fight off bondholder control.

2.35 The last wheeze we consider is manipulation of accounting statements. There is some evidence that corporate managers may seek generous interpretations of their balance sheet for financial statements. The most commonly alleged motive for this is to create a false market in which the firms share price is overstated, which may trigger management compensation payments. However, there is little evidence that markets are actually fooled by firms' manipulation of accounting earnings. Is it more convincing to argue that a marginally solvent firm might overstate its balance sheet in order to keep bondholders at bay? Sometimes this process works in reverse as accounting standards are tightened. For example, the introduction of FRS17 for pension cost accounting, largely thwarts management efforts to conceal pension deficits, thereby strengthening the hand of bondholders and other creditors (including scheme members).

3. ASPECTS OF CORPORATE BOND BEHAVIOUR

3.1 Returns

The chart below shows percentiles of daily log returns for UK government bonds, corporate bonds of various grades (7-10 year maturity) and for equities. These data go back to mid 1988, before which we were unable to source daily corporate bond indices – so we show around 1000 data points for each asset class. We have plotted these against the corresponding percentiles of a standard normal distribution. A straight line implies a normal distribution. A curve, which is convex on the downside, (or convex on the upside) implies tails fatter than normal distributions.

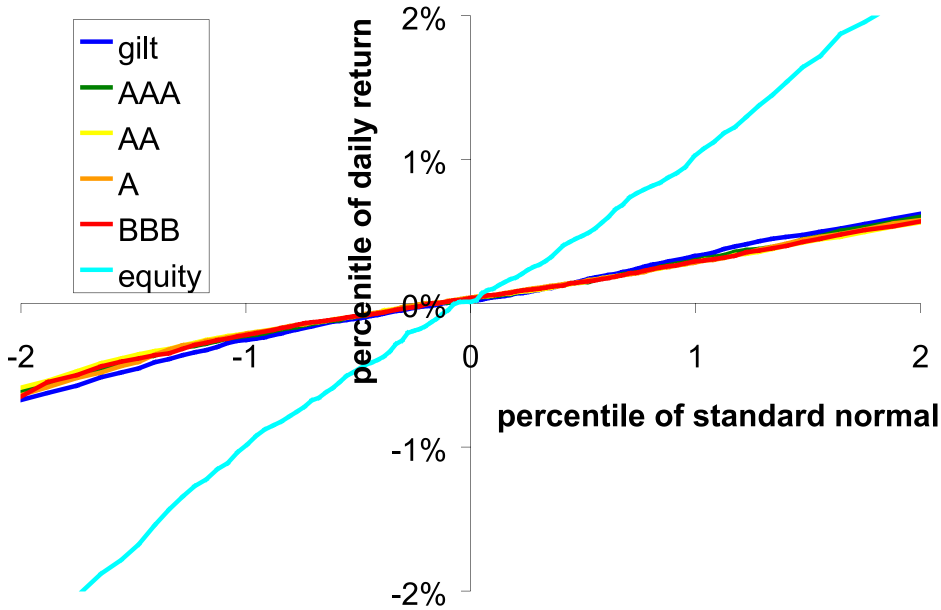

The most remarkable aspect of these data is the similarity in return distributions between government and various grades of corporate bonds. The famous skewness we might expect to see in corporate bonds, is there in the concave curve on the left hand size, but is note marked and is no more conspicuous for corporate bonds than for gilts or equities.

3.2 This is not to say that non normality can be ignored, it is just a warning to modellers that this feature is not empirically unique to corporate bonds. However, a significant part of the risk premium could still be attributable to very low frequency, but extreme, downside events and we will see that if we delve into the more subtle aspects of modelling in the following section we will find a "smoking gun" in terms of evidence of more complex processes afoot. We will consider the implications of this for simple modelling applications in Section 5.

3.3 Default Statistics

Credit rating agencies regularly publish statistics of bond defaults, split by credit grade. This table, taken from Standard and Poor’s 2001 rating analysis is typical, showing estimated probabilities of default, split by initial credit grade:

NR-Adjusted Static Pools Cumulative Average Default Rates, by Year from Rating

Rating	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
AAA	0.00%	0.00%	0.00%	0.03%	0.07%	0.11%	0.21%	0.31%	0.49%	0.55%	0.63%	0.63%	0.63%	0.63%	0.63%	0.63%
AA	0.00%	0.01%	0.03%	0.09%	0.16%	0.26%	0.39%	0.56%	0.71%	0.82%	0.99%	1.14%	1.32%	1.43%	1.56%	1.72%
A	0.00%	0.05%	0.14%	0.26%	0.43%	0.64%	0.85%	1.08%	1.33%	1.62%	1.90%	2.11%	2.30%	2.49%	2.66%	2.98%
BBB	0.00%	0.27%	0.62%	0.99%	1.63%	2.26%	3.00%	3.67%	4.27%	4.76%	5.35%	6.03%	6.53%	7.06%	7.71%	8.37%
BB	0.00%	1.29%	3.62%	6.57%	9.35%	11.90%	14.76%	16.79%	18.60%	20.45%	21.86%	23.20%	24.20%	25.11%	25.43%	25.62%
B	0.00%	6.71%	14.08%	20.59%	25.54%	29.12%	31.96%	34.58%	36.75%	38.22%	39.63%	40.56%	41.31%	41.70%	42.39%	43.06%
CCC	0.00%	28.76%	37.98%	43.98%	48.76%	53.96%	55.93%	56.66%	57.06%	58.80%	60.25%	60.80%	60.80%	60.80%	61.87%	61.87%

From this, we can see for example that the probability of a AA bond defaulting within the next 10 years is estimated at 0.99%.

3.4 For later comparison with yield spreads, it is helpful to express default probabilities as annual rates, that is the uniform annual rates of decrement which would reproduce the stated cumulative default probabilities. For later comparisons we record these annualised default probabilities (in basis points)

	Horizon (years)			
Grade	3	5	7	10
AAA	1	2	4	6
AA	3	5	8	10
A	9	13	16	19
BBB	33	46	53	55

3.5 Rating agencies also publish transition matrices, showing the frequency with which bonds are upgraded or downgraded. We show below the corresponding Standard and Poors table:

	AAA	AA	A	BBB	BB	B	CCC	D
AAA	93.28%	6.16%	0.44%	0.09%	0.03%	0.00%	0.00%	0.00%
AA	0.62%	91.64%	7.04%	0.53%	0.06%	0.09%	0.02%	0.01%
A	0.06%	2.20%	91.77%	5.26%	0.44%	0.18%	0.04%	0.05%
BBB	0.04%	0.25%	4.64%	89.45%	4.35%	0.75%	0.25%	0.27%
BB	0.03%	0.07%	0.45%	6.31%	83.07%	7.58%	1.20%	1.29%
B	0.00%	0.09%	0.31%	0.41%	5.38%	82.79%	4.31%	6.71%
CCC	0.13%	0.00%	0.27%	0.81%	1.75%	10.08%	58.20%	28.76%
D	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%

It is no coincidence that the last column of the transition matrix is equal to column 1 of the default matrix – going into state “D” is the polite expression for defaulting.

3.6 Switching between States

The transition matrix gives us a simple way to model corporate bonds – as a Markov process. The process is widely used by investment analysts.

3.7 The Markov variable is the credit grade. We can take repeated products of the transition matrix to estimate the probabilities of multi-year transitions. In particular, we can estimate the probability of default over many years using the Markov model. This does not exactly replicate the S&P multi-year observed default rates, but is remarkably close.

Cumulative default rates: Derived from Markov Model

Rating	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
AAA	0.00%	0.00%	0.00%	0.01%	0.01%	0.03%	0.05%	0.08%	0.11%	0.16%	0.22%	0.29%	0.38%	0.48%	0.59%	0.72%
AA	0.00%	0.01%	0.04%	0.08%	0.14%	0.22%	0.33%	0.45%	0.60%	0.77%	0.98%	1.20%	1.46%	1.74%	2.06%	2.40%
A	0.00%	0.05%	0.14%	0.27%	0.45%	0.68%	0.95%	1.28%	1.65%	2.08%	2.55%	3.06%	3.62%	4.21%	4.84%	5.51%
BBB	0.00%	0.27%	0.69%	1.25%	1.94%	2.74%	3.64%	4.63%	5.69%	6.81%	7.97%	9.17%	10.41%	11.66%	12.92%	14.18%
BB	0.00%	1.29%	3.23%	5.60%	8.23%	11.01%	13.83%	16.64%	19.39%	22.07%	24.64%	27.09%	29.44%	31.66%	33.76%	35.75%
B	0.00%	6.71%	13.58%	20.12%	26.13%	31.56%	36.41%	40.73%	44.57%	47.99%	51.04%	53.76%	56.20%	58.40%	60.39%	62.18%
CCC	0.00%	28.76%	46.20%	57.08%	64.12%	68.87%	72.24%	74.75%	76.70%	78.28%	79.61%	80.74%	81.73%	82.60%	83.38%	84.08%

3.8 However, although this is a neat application of Markov processes, the general methodology has a number of drawbacks. One rather philosophical potential criticism of the whole approach, to which we will return later, is the primary role of credit ratings in the process. The construction assumes that ratings (and their associated spreads) are immutable, but what if the rating agencies grades for a given risk vary over time - are we sure that a AAA rating in 2002 is the same as a AAA rating in 2001? For example, in broad terms the spread on a AA bond in 2002 is the same as the spread on a AAA bond was in 2001. Now, should we regard the spread as an indicator of transition probabilities or the grade?

3.9 At a more basic level, the chief difficulty in a portfolio context is in dealing with more than one bond. For example there is a 6.16% chance that a AAA bond is downgraded to AA, but if we have two AAA bonds, is the probability of them both being downgraded still 6.16% or $(6.16\%)^2$ or, more likely, somewhere in between? It depends whether the correlation between them is 100%, 0% or somewhere in between (or indeed they could even be negatively correlated, which would increase the range of possibilities even further).

3.10 A number of popular implementations assume, in effect, that bond transitions are independent of each other. Suppose for example that we ignore correlations between individual bonds, as well as ignoring over arching correlations with other asset classes. In other words, suppose bond transitions happen independently from one bond to the next, and independently of everything else in the universe. Then, by holding a large enough bond portfolio, we ought to obtain essentially risk free cash flows, equal to the expected payments under the Markov process. A corollary of this approach would be – contrary to empirical observations – that credit losses can be removed to a negligible level merely by diversification. Allowing positive correlations, both between different bonds and also between bonds and equities, would in principal remove some of this diversification effect. However, it is difficult to construct large correlations between continuous variables such as equity returns on the one hand and discrete state transitions on the other.

3.11 Even if we can build an appropriate correlation model, the methodology solves *half* of the problem – given a vector of credit spreads for each grade we can calculate a (discrete) distribution of total returns. We are no nearer to explaining where the spreads come from.

3.12 The Credit Spread Paradox

3.13 Some historic credit spreads are shown below (based on 7-10 year Merrill Lynch indices). Decent historic data is often difficult to obtain – for example, at many points in the past, the yield on the AA index has exceeded that on the A index- an obscure anomaly which is presumably caused by some artefact of either the rating process or the index construction.

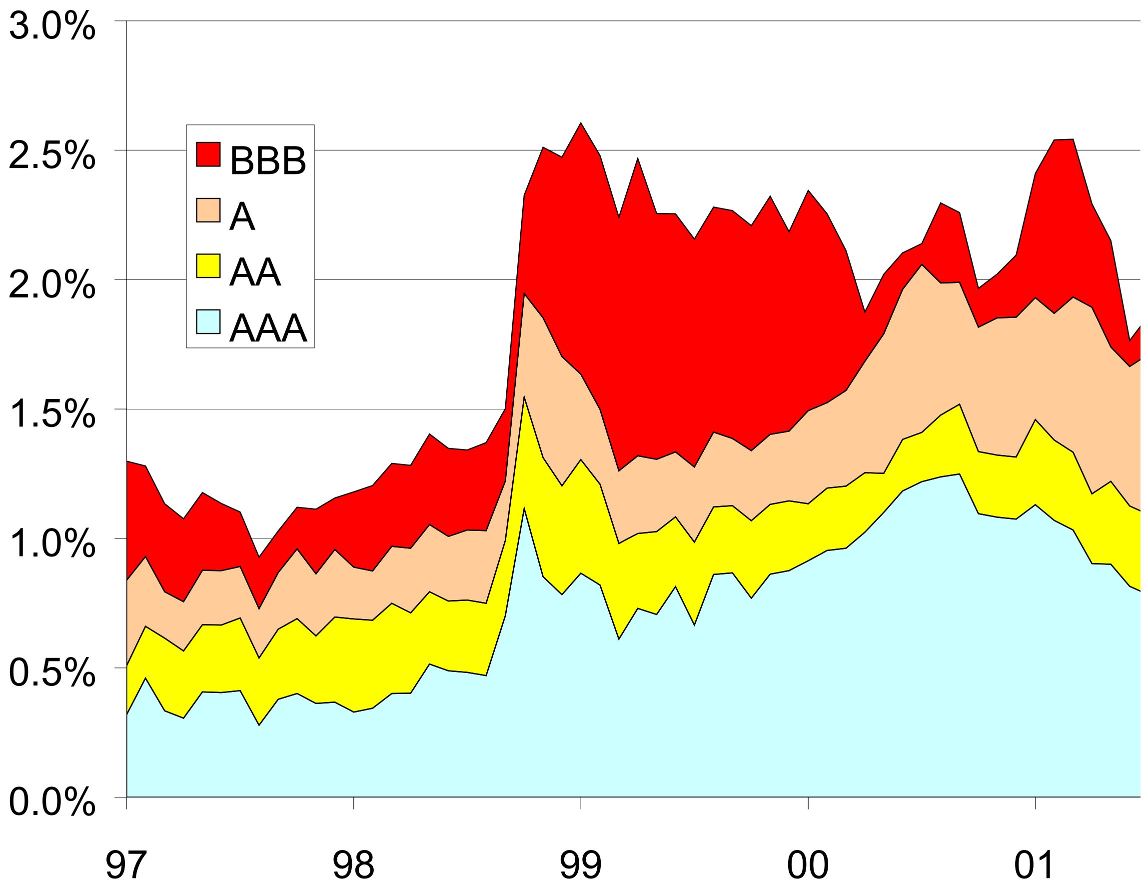

3.14 Figures are also available at other terms. Figures at 1 February 2002 are as follows:

	Horizon (years)			
Grade	1-3	3-5	5-7	7-10
AAA	44	46	53	71
AA	51	71	74	90
A	95	108	116	134
BBB	122	175	225	202

3.15 A comparison with the previously tabulated default probabilities reveals that yield spreads exceed all historically plausible estimates of default rates. In other words, even after allowing for expected defaults, bonds provide a risk premium in excess of government bonds. This risk premium is often very substantial.

3.16 A traditional approach might argue that the yield spread in excess of the annualised default probability is a form of free lunch. Why then are these free lunches not exploited until they disappear? We could perhaps argue that such arbitrages are difficult, and possibly capital intensive to execute. Although traders might hope to make arbitrage profits on bonds held to maturity, in the short term the market value of these bonds can be volatile. However, this seems to stretch credibility too far. Not only does the effect apply to short term as well as long-term credit risks, but, if the effect were robust there would seem to be many ways of implementing the basic "arbitrage".

3.17 Perhaps unsurprisingly, there is therefore another view, which implies very different models and rejects the possibility of arbitrage by simple diversification. The efficient market perspective is that yield spreads are predictors of changing default expectations. If the yield spread on a bond widens, this must be because the conditional expectation of downgrade or default has increased**.** Furthermore, precisely because bond defaults are not independent, there is an element of systematic risk to default which means that these expectations contain a risk premium representing compensation for systematic (rather than diversifiable non systematic) risk and not a free lunch.

3.18 There is a familiar actuarial analogy here. Suppose it was established that equity dividends had historically grown at 2%pa in real terms. Suppose we have a portfolio of equities – we could argue that over a large number of equities this growth rate becomes a certainty and therefore that equities yielding 2.5%pa were anomalously cheap relative to inflation linked bonds also yielding 2.5%. Although this view was widely held by actuaries in the last century it is clear that there are two related flaws in this argument that are precisely analogous to the bond risk premium "puzzle". Firstly, just because dividends have grown at 2%pa in real terms over the last x years does not mean that my best estimate of future dividend growth, conditional on what is known today, is also a constant 2%pa. In the same way, if we are in the middle of a recession (say) our best estimate of default rates over the next decade is not the same as the unconditional average level of defaults in the past. Secondly, if default rates of individual companies are, like dividend growth rates, correlated, then you don't diversify the systematic element of risk in these rates. Thus just as the "implied" dividend growth rate (gilt yield less equity yield) may contain a risk premium, so the "implied" default rate (corporate bond yield less gilt yield) may contain an embedded compensation for risk (again, to repeat, not a free lunch).

3.19 To take an extreme example of the first of these phenomena, suppose one hundred years of history showed one year with defaults of 10% and nil defaults in other years – giving an average default rate of 0.1%pa. Suppose that the market is currently anticipating one such event over the next decade – the "breakeven" spread would then be 1%pa, not 0.1%pa. The fact that this risk of loss contains systematic elements means that the actual observed spread could be even higher than this.

3.20 Under the efficient market hypothesis, the Markov transition model becomes untenable. Constant transition probabilities should imply constant yield spreads for each grade – a prediction that is clearly not borne out in practice. Instead, we must model upgrade and downgrade frequencies stochastically, in a way that is correlated with varying default rates. Furthermore, yield spreads show significant correlations to other markets. The most noticeable effect is that falling equity markets often accompany rising yield spreads on particular bonds. The correlation effect with yields spreads for a particular credit grade is less marked because of re-grading effects. All of this echoes our more philosophical objection to the transition matrix approach – namely the reliance on grade as a predictor of default or transition risk rather than working directly with spreads.

3.21 The difficulty with the efficient market route is that the resulting models are more complex than for inefficient models. The advantage of efficient models is a pricing framework consistent with risk. The same framework can then be applied to other cash flows coming out of the model, in order to derive valuations consistent with credit and other markets.

3.22 The exception here is where we are not expecting to do anything particularly useful with the model in the first place. For example if we are interested only in a model of corporate bonds for a portfolio selection study (e.g. an "efficient frontier" or "value at risk" analysis) then we may be able to get away without modelling total default. For example we might limit the analysis to a model for the change in credit spreads for an investment grade portfolio.

3.23 We will return to these sorts of model in our penultimate section on modelling bond portfolios.

3.24 However, such models are of little use for pricing contingent claims based on default events, except in limited circumstances. For the hard core modelling applications we really need a model that captures both the movements in credit spreads and the process of default, all in a framework that can be calibrated to observed market prices. This is where Merton's model comes in. We will therefore consider these more sophisticated approaches in the next section before returning, armed with insight, to simpler models.

4. MERTON'S MODEL

4.1 The basic model due to Merton (1974) treats default as a put option on the value of the firm. We have already introduced this concept in section 2. The equity holders have a call option on the value of the firm struck at the face value of the debt, the bond holders write a put option on the value of the firm struck at this face value. If the switch between puts and calls is confusing, remember put-call parity. The equity holder holds a call on the value S of the firm with strike K, but $call = S - K + put$, which ties together with the debt holders position.

4.2 We have already commented on the reasons why debt holders might allow the value of the firm to fall below this debt value without intervention. It is vital to the Merton model that bondholders cannot intervene prior to the maturity of their debt.

4.3 We now apply the simple Black Scholes put option formula, with a strike equal to the future debt value, to determine the cost of the option to default. We can similarly use Black and Scholes' call option formula to value the bond. Denoting the forward value of the firm by F and the debt redemption value by K, we can see that

$$\text{equity} = e^{-eT} F\Phi\left(\frac{\log(F/K)+\frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}\right) - e^{-rT}\Phi\left(\frac{\log(F/K)-\frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}\right)$$

$$\text{debt} = e^{-rT}\Phi\left(\frac{\log(F/K)-\frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}\right) + e^{-rT} F\Phi\left(\frac{\log(K/F)-\frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}\right)$$

4.4 If we assume a firm risk premium of μ, we can also identify a real world default probability, given by

$$\begin{aligned}
\pi &= \mathbf{Prob}\left\{F\exp\left(\mu T + \sigma Z\sqrt{T} - \tfrac{1}{2}\sigma^2 T\right) < K\right\} \\
&= \mathbf{Prob}\left\{\mu T + \sigma Z\sqrt{T} - \tfrac{1}{2}\sigma^2 T < \log(K/F)\right\} \\
&= \mathbf{Prob}\left\{\sigma Z\sqrt{T} < \log(K/F) - \left(\mu - \tfrac{1}{2}\sigma^2\right)T\right\} \\
&= \mathbf{Prob}\left\{Z < \frac{\log(K/F) - \left(\mu - \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}\right\} \\
&= \Phi\left(\frac{\log(K/F) - \left(\mu - \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}\right)
\end{aligned}$$

4.5 The immediate question is whether we can use this model to reconcile yield spreads to default probabilities. For example, let us assume a debt of K=1 of 10 year term, with a 4% equity risk premium and 30% share price volatility. Let us also assume a 1% probability of default, equivalent to 10bp per annum, around the historic default rates for AA bonds.

4.6 Back-solving for the firm parameters, we find that the forward firm value is 6.40, with risk premium of 3.39%, volatility of 25.43%. Valuing the debt gives us a yield spread of 7bp above the risk free rate. This is disappointing, as it gets us nowhere near the 90bp we see in the market.

4.7 We can make the fit somewhat better by assuming a larger loss given default than provided in the Merton model. In an extreme case, let us suppose that in the event of default, bondholders get nothing. This dramatically reduces the bond price, now implying a credit spread of 29bp per annum. This is still less than half way to the market of 90bp, but least we have demonstrated that in a rationally priced market the credit spread can substantially exceed default probabilities.

4.8 What is actually going on here? The defaults happen in the situations of poor equity returns, where deflators are largest. The market therefore gives these defaults a higher weight than could be justified by empirical probabilities alone. The Merton model only gives a mean deflator of 3 (~ 29 / 10) in this outcome. To reproduce market prices we need a conditional mean deflator of 9 = 90/10. More complex distributional forms showing fatter tails – such as a mixture of two lognormal distributions, can achieve this.

4.9 We will note in passing subtlety here. Conventionally Black Scholes models for equity options regard the process for equity prices as lognormal – meaning in simple terms that equity prices cannot reach zero. This is clearly inconsistent with a model that allows default. In the above formula, the process for the firm is lognormal and the equity value is derived from the difference between this firm value and debt value. It would appear to be inconsistent to model corporate bond default in detail without also allowing equity values to fall to zero. The difference between the lognormal model for the equity price versus the firm value is one (of many) possible explanations for the well known "skew" seen in implied volatilities in equity options.

4.10 It is perhaps worth noting here an alternative use of the simple model. The KMV rating agency (recently acquired by Moody's) used the principles of the above model to establish a ratings process. Rather than working from observed spreads to implied volatilities, this model worked the other way around. Instead, market observed equity volatilities were *inputs* into the model and KMV circumvented the problem of reproducing observed spreads by basing their rating on a parameter known as "distance to default", which is the number of standard deviations required for the firm value to fall below the face value of debt. In other words, in very simple terms:

$$\text{Distance to Default} = \frac{\ln(K / S)}{\sigma}$$

There is a slight complexity in the calculation of firm volatility from equity volatility, which we also note here, although it applies to the previous description of the Merton Model also. If we observe the volatility of equity E, we are actually observing the volatility of a call option, not of the firm itself. Thus in theory we need to unravel the volatility of S out of a recursive equation, since in simple terms the volatility of E is suppressed by the rise in the value of the put option as the value of the firm falls, and vice versa.

4.11 Once the above calculation has been performed, a statistical process then takes over, which maps historic "Expected Default Frequency" to "Distance to Default" for various horizons. The basic principle of the approach, which echoes our comments about over reliance on traditional ratings as immutable constants, is that all useful information about the likelihood of default is embedded in information about the current equity levels and the volatility of equity. (In the context of the *raison d'etre* of a rating agency it would perhaps be a step too far to suggest that all useful rating information is also embedded in existing credit spreads, so at least for quoted debt the KMV model naturally goes only part of the way towards a fully market based approach).

4.12 We do not seek to enter into the previously heated debate between the traditional rating methods and the KMV method. Naturally, KMV claimed successes where traditional approaches have failed and the traditional agencies for their part strongly defend their own methods. As noted previously, we understand that Moody's have now purchased KMV, from which readers can draw their own conclusions.

4.13 Extensions to the Merton model

We now look at a number of extensions of the Merton model. Our first challenge is to bridge the gap between yield spreads predicted by the Merton model, and those seen in the market. As observed by Bohn (2000), some form of fat - tailed distribution is required to achieve this.

4.14 Other extensions to the Merton model are needed to reflect more realistically the various options we listed in section 2. One simple enhancement, to allow for third party leakage, or for modelling mezzanine debt, is to treat the debt as a call spread on the firm, with the lowest layer of value accruing to liquidators, tax authorities and other third parties.

4.15 We can further adjust the model by allowing for early default. One way of modelling this, following Black & Cox [1976], is to allow the call spread to be knocked out if the firms value touches a lower barrier. This approach uses the barrier options cited, for example, in Hull [1992].

4.16 A further relevant aspect of corporate debt in practice is that firms can periodically re-gear. In other words, if the firm value rises to a certain level, management may raise additional debt. In this case, the call spread defining the debt is struck relative to the maximum firm value so far reached. Such options can be priced using the look-back structures described in Goldman, Sosin and Gatto (1979)

4.17 Ideally, we would like to allow for both refinancing at an upper barrier, and early default at a lower barrier. Such structures are difficult to price, but asymptotic solutions exist using trigonometric functions. Bohn (2000) gives a more detailed survey of the various available contingent claims approaches to pricing credit risk.

4.18 As we seek more realistic models, so the models become more complex. For example, we have already demonstrated the need for fat-tailed distributions in order to replicate observed default frequency and credit spreads. Combining such distributions with barrier options is very challenging – although Avram (2001) has made some progress using the Wiener-Hopf factorisation for Levy processes. Needless to say, calibration of such models remains a formidable challenge.

Non – Merton Aspects of Bond Pricing

4.19 At this stage, we consider some aspects of pricing which lie outside the Merton modelling framework. The first often-quoted effect is a tax effect. This is particularly noticeable in the US, where corporate bonds are taxed more heavily than some government bonds – for as recent synthesis of results, see Elton, Gruber et al (2001). The corporate bonds must carry a higher yield if only to compensate investors for the tax. In the UK this effect does not seem to be so material.

4.20 A further often-claimed effect is liquidity. Indeed, many UK investors seem keen to argue that most of the corporate yield spread is a premium for liquidity – with the implication that an investor with little need to realise assets in the short term can collect the liquidity premium as a free lunch. The logical conclusion of this argument is a clientele effect where such investors should dominate the corporate bond market to such an extent that the liquidity premium reflects the likely future sales from those investors – which are likely to be small. In other words, any liquidity premium in the bond market at best offsets the contingent transaction costs in the event that the bond must be sold. This is consistent with empirical studies, which find that corporate bonds yields are explainable in terms of defaults, risk premiums and tax effects. For a recent summary of this literature, see Duffie and Singleton (1999) or Dullaway (2002).

5. MODELLING BOND PORTFOLIOS

5.1 It should be clear from the preceding sections that in a heavy duty pricing or contingent claims application for credit risk modelling, we need most of the machinery of the Merton Model and its extensions, to capture the important aspects. However, we suggested in Section 3 that if the actuary has more limited ambitions, such as simply modelling a corporate bonds in a portfolio of other assets, then we might get away with a simpler model. Hopefully, this might just mirror the way that conventional bonds or, in particular, equities are already modelled. After all, we might not be worried about the truncation of the distribution of equity returns (due to the option to default)– so if the non-normality of bond returns is just the opposite side of this truncation, maybe we can apply a similarly cavalier approach to corporate bonds.

5.2 However, this logic for ignoring the complexities of corporate bond returns is far from water tight. Primarily this is because, in contrast with equities, 100% (or thereabouts) of the excess return from corporate bonds is derived from the features under consideration. A better analogy, as we can now see from the previous section, would be whether it would be sensible to include in our asset allocation model the possibility of *writing* out of the money put options as part of our asset allocation. If we naively booked the income derived from writing such options using *observed* put option prices then this income might appear to be anomalously high (due to skews in implied volatilities) when compared with the returns derived from modelled equity exposures with similar apparent risk characteristics. Faced with such a situation, if we were determined to include options, we would have to adjust our model – say by using lower premium income from writing such options – if we want to avoid spurious asset allocation decisions (namely writing unlimited amounts of them). Sadly, it seems that corporate bond prices, to the extent that the credit spread is equivalent to writing a put option, exploit precisely this grey area. So modellers beware!

5.3 The simplest possible way of modelling individual corporate bonds is to treat them as government bonds, that is, to ignore both the default option and the yield spread. Equivalently, this would assume that future yield spreads are exactly offset by defaults. Given the closeness of the return distributions for gilts versus corporate bonds in the UK, this approach is not as silly as it sounds. It may be appropriate if corporate bonds are only a small part of a bond portfolio. On the plus side, it also has the advantage of avoiding foolish conclusions about free lunches. By analogy it is after all no different from excluding the writing of options as a separate asset class - something that few asset modellers would balk at.

5.4 A simple approach we have seen in practice is to allow for a fixed corporate yield spread, but to ignore defaults. On further probing, practitioners may refer to the spread as a "net spread, after allowance for defaults". We do not recommend this approach, because it commits the fundamental error of booking the reward without modelling the risk. Any approach with a positive net spread would declare corporate bonds to be universally preferable to risk-free gilts. Further criticism is probably unnecessary, but it might also be worth repeating that the other danger with this approach is that the model often combines a current credit spread with a historic risk of default to quantify this free lunch. To the extent that changes in observed spreads reflect changes in conditional default expectations the asset allocation signals given by such models over time may also be highly dubious.

5.5 **Modelling Individual Bonds as Mixtures of Equity and Risk-free Debt**

5.6 Typical aspects of an individual corporate bond would include:

- a mean return intermediate between equity returns and gilt returns
- volatility also intermediate between gilts and equities
- strong positive correlation with gilts of corresponding term
- some additional correlation with the equity market reflecting credit risk

5.7 We can capture these aspects by treating individual corporate bond returns as a weighted average of the return on risk free bonds and the return on an equity, with some suitably calibrated fixed proportion in equity. This captures the link with equity prices, as demanded by Merton's model. The proportion of equities in the fund could be determined by reference to the bond volatility or risk premium (but not both).

5.8 A further enhancement to this approach would be a dynamic replication algorithm, which increased the effective equity exposure in a falling equity market. This reflects the dynamic aspect of the put option which shareholders purchase from bondholders.

5.9 To reflect the less than perfect correlation between an equity and bond portfolio, additional noise could be added. These approaches can readily be accommodated in a conventional mean-variance framework. However, much more work is required on realistic assessments of correlation. For example, many investors may find themselves exposed simultaneously to equity and bonds in the same firm. Is there some way of using a contingency matrix of such double hits to improve correlation estimates between the particular equity and bond portfolios?

5.10 A potential objection to all of these approaches is that there is a maximum value for any corporate bond, that its spread over gilts should not become negative. In principle, our replication algorithm could violate this principle, if equity markets jump upwards. This limitation of the method is reduced by dynamic replication, which reduces equity exposure in rising markets. We can also argue that the limitation is either not a problem or less of a problem if we move to a bond portfolio model, where we do not keep track of individual bonds, as discussed below. The lack of any apparent cap in historic corporate bond return indices, suggests the approximation of corporate bond returns as a mixture of equities and gilts may be acceptable going forward, especially if we are considering a portfolio of corporate bonds. We will later encounter complications when these bonds are held to maturity, for example to match particular liabilities.

5.11 Constant Spread Portfolios: Modelling Total Returns

Once we move on to consider portfolios of bonds (as typically appear in asset allocation studies) rather than individual bonds, we encounter some surprisingly subtle issues. In particular, we are no longer modelling assets alone, but an investment *process*, whereby, for an investment grade portfolio, we sell bonds that fall below investment grade and rebalance. Once again the same is true of modelling a typical equity index, but the rebalancing of a corporate bond portfolio plays a more important role. In particular, if we consider a *continuous* process for the value of the firm, then it could be said that although AA bonds appear to have a probability of default over the time interval of a discrete Markov process, in reality this would actually be an artefact of the discrete time interval. If the process for the firm was continuous then the AA bond would become a A, BBB, BB etc bond in quick succession, but would not default as a AA bond.

5.12 So, if we model AA bonds (say) we actually need to model not defaults but the process of transitioning back from A grades to AA grades for bonds that suffer downgrades. In doing so we hit a number of snags. Not least the fact that in reality some firms have defaulted with apparent AA grades and the artificiality of the grade categories discussed earlier. Faced with a plethora of such issues when we delve into a process defined by grades, it is convenient to start with a convenient approximation that cuts through the complexities and reveals much about the underlying *process*.

5.13 Let us therefore consider an approximation to an investment grade (or other grade selected portfolio) represented by a notional *constant spread portfolio*. Basically, instead of deciding to hold AA bonds, say, we hold bonds with a spread in the region of 100bps pa. If spreads rise in the portfolio, we sell our existing holdings and migrate up the credit curve (to higher security), if spreads fall we reverse this. For good measure we can also enforce a constant duration on the portfolio, by continually reinvesting back into a fixed duration bond.

5.14 With these sweeping assumptions, modelling all becomes rather simple. If we suppose, as is common in portfolio selection models, that volatilities are constant, then the only parameter that could make spreads fall or rise would be lengthening or shortening of the distance to default. Furthermore, over short time periods, the movement in each bond price can, under the Merton model, be hedged by a combination of equities and a riskless bond. If we aggregate a portfolio of such bonds, all with the same or similar distance to default and term then the aggregate hedge portfolio is just a larger version of the individual bond hedge. (If we have a continuous process for the firm value then this construction also eliminates the issue described in 5.10 in relation to individual bonds).

5.15 Under these assumptions, none of the bonds in the portfolio ever defaults. However, that is not to say that the credit spread becomes a free lunch. In fact this conceptual model neatly links the transition matrix approach to the observed credit spread. Under the transition martix approach the one period return on a portfolio of bonds can be computed by calculating the prices of bonds in each state at the next time step, multiplying by the probability of each state and deriving an expected price of the total portfolio. In contiuous time with a continuous process for the firm, we consider an infinitessimal version of this. The infinitessimal change in bond price, comprises the interest on the risk free bond element of the bonds (in the Merton framework), plus the change in the value of the put option. Thus we see that the loss on default is replaced in the context of an investment *process* representative of a grade-based bond portfolio (such as AAA or investment grade), the credit spread can be reconciled with the time decay (or "theta") of the option to default and not the actual rate of default seen within the grade.

5.16 What this suggests, in a roundabout way, is that we can approximate a constant spread portfolio by representing corporate bonds as a simple linear combination of the returns on equities and risk free bonds, as we had first proposed. The coefficients could in principle be determined by building a Merton style model and deriving hedge ratios, but this seems like using a sledgehammer to crack a nut. More simply the coefficients can be derived empirically from historic total returns. In practical calibrations, it is rare that these two calculations are close – we nearly always get a higher hedge ratio from the historic data calculation. We prefer to use empirical correlations, and blame the known weaknesses in Merton's model for the lack of a consistent parameter choice. We should not beat ourselves up over this though, as explained earlier this is no different to making an adjustment to option prices (if included in a model) and blaming the observed skew in implied volatilities.

5.17 The reader probably did not require any of the theory in the preceding sections to arrive at this conclusion, but it is comforting that the theory is consistent in this regard. Furthermore, modelling bonds as a linear combination of two other assets is of little help to determined efficient portfolio seekers – an optimiser will be indifferent between corporate bonds and this combination of alternatives. However, it is a useful starting point and provides a reality check for alternative approaches; theory says that the linear combination should be a pretty good approximation and has the advantage of being mean-variance consistent.

5.18 Modelling Corporate Yield Spreads

In some applications of portfolio selection modelling, risk relative to accounting liabilities based on corporate bond discount rates may be of interest (for example when modelling risk relative to FRS17 liabilities in a UK pension fund context). Once we move to the idea of modelling spreads on portfolios rather than total returns, a number of pitfalls emerge. Accordingly it is worthwhile weighing up whether this degree of sophistication is really necessary in many instances. For example, if the portfolio risks are dominated by equity risk then it may be wasteful of resources to devote energy to precise modelling of the nuances of spreads on corporate bond portfolios. In many applications, calculating risk relative to a well-constructed hedge portfolio of assets will be the most sensible way to proceed. In these instances, it may once again be possible to work with total returns on corporate bonds rather than working with models for spreads.

5.19 Returns and yield spreads are not independent. If the yield on a corporate bond rises then the price of the bond will fall. However, unlike the situation for default-free bonds, the total return must include an allowance for bond defaults, or (as discussed above) transitions. We repeat that where an investor has a policy of investing in (say) AA bonds, the total return on a AA bond portfolio will not reflect only AA yields; the effect of upgrades and downgrades must taken into account. Even if, as a matter of policy, bonds are sold following a re-grade, the sale price will reflect the new credit grade rather than AA rates.

5.20 Once we consider spreads the natural place to look for inspiration is in the standard literature on term structure models. Although we may not require all of the arbitrage free machinery of these models, it is worthwhile borrowing some of the ideas to avoid unwanted "features" emerging in models, such as spurious profits from barbell strategies within the corporate bond market.

5.21 However, we need to be careful in applying these models and it is useful to think through the theory a little first. We will proceed by analogy. Consider inflation linked bonds and nominal bonds. The difference between the interest term structure on these two bonds gives a market implied inflation rate. Furthermore, and most importantly in this context, the process for actual inflation can be modelled as emerging from the difference between the short nominal yield and the short real yield .

5.22 By analogy, if we consider the difference between corporate bonds and nominal bonds, we might think that the difference between the respective short term corporate bond rate and the risk free interest rate gives us a process for the actual rate of default. Of course, regarding this difference in short dated yields as an actual default rate in a model with diffusion processes would be useless for a single bond – it is either in default or it is not. However, we might consider exploiting the assumption that we have a portfolio of bonds in the hope that it is reasonable to assume that in aggregate the default process looks a bit like a diffusion process. However, once again, if we are modelling an investment *process* where we keep rebalancing our portfolio back into a given grade range, we cannot interpret this spread as a default rate but as a loss from rebalancing (identified above as the time decay or "theta" of the option to default).

5.23 Given the availability of off the shelf term structure models it is natural to start with the classical "short rate" models, where we model the short rate of interest on corporate bonds, and project the rates along the rest of the term structure from this process. The simplest model for spreads, as ever, is to assume that changes in the short rate follow a random walk.

5.24 Such models will, however, generally be found to generate excessive volatility in longer dated corporate bonds, relative to short dated bonds This not a problem if we are only interested in one particular duration of bond, as we can calibrate the volatility accordingly, but in that case the machinery of these term structure models is probably unnecessary anyway. If we do indeed want to model several different bond durations simultaneously, then we need a more complex driving process. The easiest way to reduce the volatility of longer dated bonds relative to shorter dated bonds is to follow the natural progression to a mean reverting model.

5.25 There is a subtlety here though. The term structure machinery is most easily applied to the overall interest *rate*, rather than the *spread*. However, if we do apply the process to the overall interest rate, then unless the same mean reversion parameter is adopted for all interest rates, the process for the spread will be quite complex.

5.26 However, as with other applications of this model, the main drawback is that it admits negative spreads. This may not be a serious problem in the context of many modelling applications and is less serious when modelling credit as a single asset class, rather than modelling individual grades. However, it is a serious problem if we try to model individual grades as we get a multitude of messy possibilities – such as A grade spreads falling below AA grades etc.

5.27 One possible way to address this is to model separately the AAA spread, the spread between AAA and AA, AA and A, between A and BBB etc, using a model that only admits positive rates. The next candidates on the list being either the Cox-Ingersoll-Ross or Black-Karasinski models (see Baxter & Rennie, 1996). However, once again we must recognise that the term structure mathematics applies to the total rates and not these spreads alone and quite messy adjustments are required to work back and forth between these if we want to remove spurious effects from various barbell strategies between grades.

5.28 The complexity of these models denies us any form of analytical solution. Existing numerical methods are based on the numerical solution of (potentially multi factor) partial differential equations. In principle these may have to be re-solved at each time point in each simulation – which would multiply many fold the computational burden of generating the model. It would be necessary to control the number of stochastic factors, and path dependence, to make these computations practical. A possible alternative approach, which the authors have not implemented in practice, would be to solve the differential equations numerically before simulating, and store a lookup table of solutions, which could be interpolated as each simulation is generated.

5.29 Further work is required to determine whether in practice such models can account simultaneously for observed default losses compared to credit spreads in the market. The models may be so inflexible that initially observed credit yield spreads cannot be fitted – an eventuality made more likely by the apparently anomalous (or even negative) spreads which occasionally crop up in corporate bond index data. A further calibration complication arises from the use of continuous default rates, since in practice observed defaults are discrete. In a good year there may be only a few dozen defaults. This makes empirical correlations difficult to observe.

5.30 **Modelling Spread Term Structures**

An alternative way to model yield spreads is to start with bond prices, from compounding total returns. If bonds are modelled individually, it is then possible to work backwards from the price to an implied yield. In many cases, it would be acceptable to model zero coupon bonds of each term.

5.31 Candidate models for the price include a dynamically rebalanced mix of bond returns and equity returns. As we must compute yields, these returns must necessarily be split by term of the bond. Any bond we may in future buy must be modelled from time zero. For example, if 10years from the date of projection we plan to purchase a 15 year corporate bond, we are forced to model a 25 year bond from the start of the projection (and hope a default event does not wipe it out in the meantime).

5.32 In addition to the total return, we need to decompose this into the return on continuing bonds less the loss (if any) on defaults (or transitions) during the year. It is common, and easy, to ignore defaults prior to maturity. Including defaults would require a greater return on continuing bonds, and so for a given set of return assumptions would lead to lower future yield spreads.

5.33 A difficulty with this approach is that it is by no means guaranteed to provide feasible corporate bond prices, which do not exceed the price of a corresponding risk free bond. This difficulty is particularly acute prior to a bond's expiry. If the bond market value is coming in below face value, we can easily interpret the difference as a default loss. But what do we do if our projected bond price is above the relevant risk-free bond?

5.34 One way around this problem is to readjust the hedge portfolio to include a sold call option on the equity market, with a term of one year or whatever frequency is chosen for generating model results. The strike of the call option needs to be expressed relative to the risk free bond index. In a diffusion setting, Margrabe's formula would be an appropriate way to price the out-performance option. A minimum yield spread over gilts is required as an additional model input to determine the strike of the option.

5.35 As an alternative, we could instead model the corporate bonds using much longer-term options. This technique uses the Merton model formula for the corporate bond payoff. However, instead of modelling individual bonds, we now apply this to a market index. Once again, Margrabe's formula gives an initial approach for pricing the options, but here the option term runs from the valuation date to the maturity date in one leap, rather than in annual intervals.

5.36 In the end of section 4, we discussed some enhancements to Merton's model. Although these were in the context of individual bonds, the same approaches may well also improve models for bond portfolios. As yet, there do not seem to be efficient algorithms for calibrating these more exotic structures. As a result, we have a preference for a simple model, which we can calibrate transparently.

5.37 Further empirical work is needed to establish how well in practice the Margrabe formula works in the presence of failures in the underlying assumptions. For example, the model may show stochastic volatility or jumps in the underlying price process, either of which would theoretically invalidate the Margrabe formula. In this situation, we recommend the discipline of market calibration as a partial protection against meaningless output. Specifically, we recommend the generation of simulations as a two-step process. First – produce the simulated equity returns, risk free rates and deflators. Secondly, iterate to find the implied volatility whose insertion into Margrabe's formula produces bonds which are correctly priced back to time zero using the chosen deflators. This process is akin to the implied volatility construction, widely used by derivatives traders even when strict Black-Scholes assumptions are known to fail. Our recommended approach takes this further to incorporate not only volatilities today but also at future dates.

5.38 Some of the approaches above will, in an attempt to better reflect reality, depart from lognormal returns for corporate bonds. This may be better at one level, but could create problems when we try to conduct standard portfolio optimisations. Unless the optimisation process is correctly adjusted (to reflect skewness in returns, for example) the simpler approach may give more understandable results.

5.39 Faced with the problems of maintaining positive spreads and various other anomalies that need to be ironed out when modelling spreads, we have generally found that despite its "sledgehammer" appearances at first sight, the use of a model such as the Merton model or its extensions to keep track of credit spreads is often convenient. In other words, we would first create a shadow default process based on firm values and then calculate spreads at each step directly from this model. Providing that the shadow model is tractable (and does not require simulation at each step) this provides a surprisingly effective work around the hidden problems of more conventional models for credit spreads.

6. CONCLUSIONS

6.1 A naïve approach to modelling corporate bonds might start with an observation that the return from corporate bonds (as reflected in the current yield, less an allowance for defaults) appears anomalously high. It seems that little would be done to shake this by analysing the distribution of corporate bond returns – there is little evidence for non-normality. A model built around these observations would inevitably lead to high allocations to corporate bonds in an asset allocation model.

6.2 However, such a model would be of little use in heavy duty modelling applications aimed at pricing contingent claims based on credit events. Such applications would require a model that reconciles credit spreads with default events.

6.3 In order to build such a model we draw instead on our initial analysis of the economics of corporate bonds. This reveals credit risk premiums as being akin to option premiums and leads naturally to the Merton Model and its extensions.

6.4 The Merton Model explains the apparent anomalies between credit spreads and actual default frequencies in terms of the nuances of options pricing. In particular evidence from plain vanilla options suggests that the implied volatility in a conventional Black Scholes Model needs to be calibrated carefully rather than drawn from historic data. Equally, apparent anomalies between implied volatility and historic volatility in such options are generally regarded as a feature of the approximations in the simple Black Scholes Model, rather than exploitable investment opportunities.

6.5 Our contention is that apparent anomalies in credit spreads need to be viewed carefully in a similar light and can be reconciled with plausible model assumptions relating to departures from simple Black Scholes Model.

6.6 Armed with these insights from the Merton Model we re visit the apparently much simpler application of modelling corporate bonds in portfolio selection studies. We find a large number of confounding and subtle nuances in this application and question the validity of building models without some element of economic consistency.

6.7 We suggest further that simple approaches without the capacity to create apparent anomalies may be far more reliable in such applications than models going to great lengths to model superficial features (such as positive spreads) without getting to grips with the underlying process that is generating returns (namely the implicit writing of options). We contend that modelling corporate bond *portfolios* primarily relates to modelling an investment *process* of continuous rebalancing. The constant spread approximation of this rebalancing in continuous time gives a clean insight into this process in terms of the time decay of options. However, the over arching arbitrariness of grades and grade ranges on which this rebalancing is based makes any efforts at realistic and consistent modelling of transitions between typical portfolio grade bands a tall order, if not doomed from the outset.

6.8 Our conclusion is that although apparently a "sledge hammer" approach, and despite its unfamiliarity within a traditional portfolio selection context, approximations to the Merton Model represent the most promising way forward. The common approximation of treating a corporate bond as a combination of equities and risk free bonds is shown to be the simplest example of applying this principle to a constant spread portfolio, provided that risk premiums and correlations are chosen consistently. We argue that any approximations made in linking more complex versions of the Merton Model to familiar credit ratings is a consequence of the arbitrariness of these ratings, rather than a fundamental weakness of the Merton approach.

Acknowledgements

We are grateful to Shyam Mehta, Jeroen van Bezooyen, David Dullaway, Frances Southall, Gareth Collard and Deniz Sumengen for useful discussions relating to corporate bond modelling. All opinions expressed and errors contained are solely the responsibility of the authors alone. These views do not represent the views of our employers nor – heaven forbid – the Institute and Faculty of Actuaries!

Bibliography

Altman, E.I & Kishore, V.M (1996) "Almost Everything You Wanted to Know about Recoveries on Defaulted Bonds", Financial Analysts Journal November/December 57-64.

Avram, F. (2001). Pricing of barrier option under the Levy model. Working paper, Heriot Watt University.

Black, F & Cox, J.C (1976) "Valuing Corporate Securities: Some Effects of Bond indenture Provisions", journal of Finance, 31(2), 351-367.

Bohn & Crosby (1993) Modeling Default Risk, KMV (revised 2002)

Bohn (2000) A survey of Contingent-Claims approaches to Risky Debt Valuation. Journal of Risk Finance,, Spring 2000, pp 53-70.

Baxter & Rennie (1996) Financial Calculus. Cambridge University Press.

Bezooyen et al (1998) Taxation and Investment Joint Faculty and Institute of Actuaries Investment Conference (available on www.Gemstudy.com)

Brealey& Myers (2002) Principles of Corporate Finance, McGraw Hill (7th Edition)

Copeland & Weston (1988) Financial Theory and Corporate Policy Addison Wesley (3rd Edition)

Duffie, D and Singleton, K J (1999) Modeling term structires of defaultable bonds. Review of Finacial Economics 25, 23-49.

Dullaway, David (2002) Corporate Bonds and the Liquidity Premium, BAJ forthcoming

Elton, EJ, Gruber, M J, Agrawal, D and Mann, C (2001). Explaining the rate spread on Corporate Bonds", *Journal of Finance*, 56, 247-277.

Goldman, B, Sosin, H and Gatto, M (1979). Path dependent options- buy at the low, sell at the high. Journal of Finance 34: 1111-1127.

Hull, J.C (1992) Options Futures and Other Derivative Securities (2nd edition) Prentice-Hall International

Jensen & Meckling (1976) "Theory of the firm, managerial behaviour, agency costs and ownership structures", Journal of Financial economics October **3** No.4 pp 305-60

Leland, H.E. (1994a). Corporate Debt Value, Bond Convenants, and Optimal Capital Structure. *Journal of Finance*, 1213-1252.

Leland, H.E. (1994b). Bond Prices, Yield Spreads and Optimal Capital Structure with Default Risk. *Walter A. Haas School of Business, University of California, Berkeley*.

Merton, R.C. (1974). On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. *Journal of Finance*. 449-469.

Margrabe, W. (1978). The value of an option to exchange one asset for another. *Journal of Finance*, 33, 177-186.

Sweeting, Paul (2002). The role of high yield corporate debt in pension schemes. SIAS.