The critical parameters

Frank Skinner and Antonio Díaz examine how different binomial stochastic processes alter the pricing of credit derivatives, and define which credit risk parameters are of critical interest

In recent years, we have learned much about modelling term structures subject to credit risk. However, there is little empirical work regarding such basic issues as the relative importance of different parameters, describing credit risk and whether correlation between credit risk and pure (default-free) interest rates really matters. Indeed, we do not know if we should not parameterise credit risk at all and instead apply pure interest rate modelling methods directly to interest rates subject to credit risk.

We address these questions by implementing binomial stochastic processes for pure rates of interest and credit risk in an arbitrage-free framework. The resulting models yield replicating portfolios of state prices (synthetic corporate zeros) that can be used to price credit derivatives. Since the whole point of modelling credit risk is to obtain accurate prices, we examine how different stochastic processes change the distribution of state prices and therefore credit derivatives prices.

Specifically, the idea is as follows. For each state security price, there exists a corresponding corporate interest rate. Credit derivative prices can be found as the present value of promised payoutsusing the distribution of corporate interest rates. Precisely the same price can be found as the sum of promised payments multiplied by the state security prices. Since all models force the structure of state prices at each point in time to replicate corporate zero prices that underlie the corporate term structure, then different parameterisation schemes force different distributions of state prices. Since changes in the distribution of state prices imply different prices for derivatives, then we can suggest which parameters seem critical. If by including, say, correlation between credit risk and pure rates of interest, the distribution of state prices remains the same, then we can suggest that correlation does not matter. On the other hand we may observe that adding negative correlation changes the distribution of state prices such that at low corporate interest rate states, state prices are higher, but at high corporate interest rates states, they are lower. This means that by neglecting negative correlation, we would underprice credit-risky call options that payout in low corporate interest rate states and overprice credit-risky put options that payout in high corporate interest rate states.

The binomial model

Under the risk-neutral probability measure Q conditional upon information available up to date t, Duffie & Singleton (1999) show that the price of a one-period defaultable zero is written as:

$$V_t = E_t^Q\left[h_t e^{-r_t}\omega_{t+1} + \left(1-h_t\right)e^{-r_t}V_{t+1}\right] \qquad (1)$$

Note that h_t is the conditional (upon no prior default) hazard probability and r_t is the pure interest rate at time t. Meanwhile, ω_{t+1} is the recovery rate and V_{t+1} is the promised payout of \$1 at maturity $t+1$. In other words, a defaultable zero promises to pay V_{t+1} at maturity $t+1$, but the promise may be broken at hazard rate h_t. If default occurs with hazard rate h_t at time t, an amount ω_{t+1} is paid at time $t+1$, conditional upon no prior default. Then, under the risk-neutral probability measure Q, these future expected cashflows are discounted by the pure rate of interest.

The above is a general expression for the value of a one-period defaultable zero. To highlight the challenges confronted when modelling credit risk, we rewrite (1) in state contingent format in the case of a two-period corporate zero:

$$\begin{aligned} V_t = \Big\{ & h(t,j)\left(\omega_{t+1}\right)\left[e^{0.5\left[r(t+1,i+1)+r(t+1,i)\right]}\right] \\ & + 0.5\left[h(t+1,j+1)+h(t+1,j)\right]\left[1-h(t,j)\right]\omega_{t+2} \\ & + \left[1-0.5\left\{h(t+1,j+1)+h(t+1,j)\right\}\right] \\ & \left[1-h(t,j)\right]V_{t+2}\Big\} e^{-\left\{r(t,i)+0.5\left[r(t+1,i+1)+r(t+1,i)\right]\right\}} \end{aligned} \qquad (2)$$

The above expression is a binomial implementation of (1). It says that a corporate zero may default during the first period with hazard rate $h(t, j)$ and recover ω_{t+1} at the end of the first period. The amount is reinvested in a Treasury security to earn interest at a high rate (interest state $i+1$) or a low rate (interest state i) next period where either state may occur with equal likelihood under the risk-neutral probability measure. If the corporate zero survives the first period with probability $[1 - h(t, j)]$, it may default at maturity in a high credit risk (high hazard rate) state with hazard rate $h(t+1, j+1)$, or it may default at maturity in a low credit risk (low hazard rate) state with hazard rate $h(t+1, j)$, where either credit risk state may occur with equal likelihood under the risk-neutral probability measure. If the corporate zero defaults during the second period, investors recover ω_{t+2} at maturity, which may be different from ω_{t+1}. The corporate zero pays the promised \$1 (V_{t+2}) at maturity conditional upon survival for both periods. All potential cashflows, both the terminal payout and recovery amounts, are discounted back to the present using binomial stochastic pure interest rates.

The above expression is an empty mathematical shell, as two critical challenges evident in (2) remain unresolved. First, what is the relationship between hazard probabilities $h(t, j)$ that evolve in credit risk state j and pure rates of interest $r(t, i)$ that evolve in interest rate state i? Second, hazard probabilities are conditional probabilities in that, to default at t_2, the bond must survive t_1. This means that in all possible credit risky states j and interest rate states i, one must measure expected conditional payouts in the event of default under the equivalent martingale measure for not only the current period, but also all possible prior periods.

It is tempting to solve these challenges by "brute force", that is, by calculating all possible hazard and pure interest rates for all time periods.

1. Non-parameterised corporate interest rate model compared to zero correlation, constant recovery model

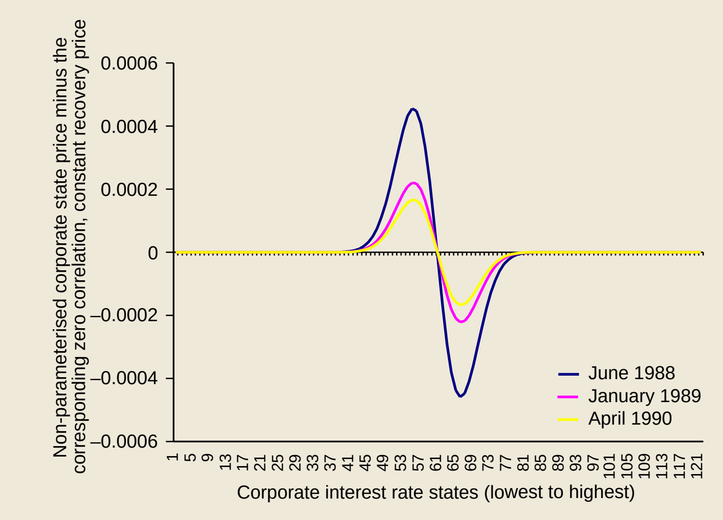

2. Bias in corporate state prices and correlation with Treasury interest rates (upward sloping term structure of June 30, 1988)

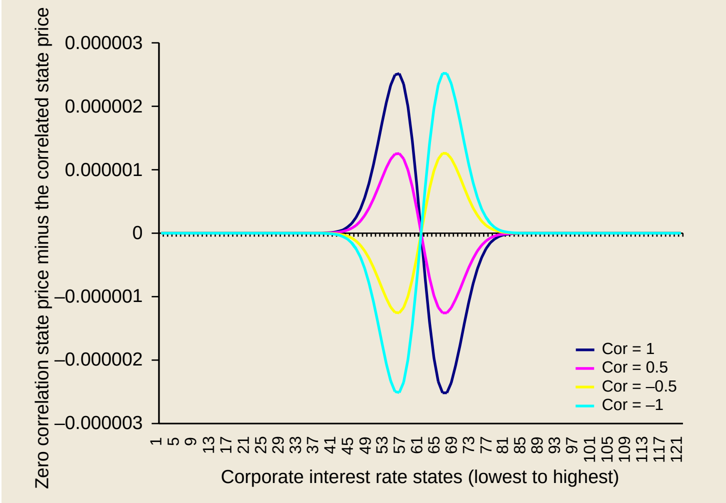

However, this would be computationally expensive. The number of pure interest rate states will equal $t + 1$, and the number of hazard states will be $t + 1$ and all possible combinations will be $(t + 1)^2$. Consequently, we impose distributional assumptions regarding the relationship between hazard and pure interest rates.

To obtain a binomial credit risk model from (2), we suggest the following binomial stochastic process for the pure rate of interest and the hazard rate:

$$r(t,i) = r(0,0)e^{\left(u_t \Delta T + (2i - t)\sigma_r \sqrt{\Delta T}\right)} \qquad (3)$$

$$h(t,j) = h(t,i) = h(0,0)e^{\left(v_t \Delta T + \rho_{h,r} \frac{\sigma_h}{\sigma_r} r(t,i)\Delta T\right)} \qquad (4)$$

The first binomial stochastic process is the familiar constant volatility version of Black, Derman & Toy (1990), where $r(t, i)$ refers to the pure interest rate that evolves in state i and time t, $r(0, 0)$ is the current observable short-term pure interest rate, u_t is a time-dependent parameter that calibrates the interest rate tree by forward induction through use of state prices to the Treasury yield curve, ΔT is the time step and σ_r is the constant pure interest volatility parameter. Note that when $t = 0$, $r(t, i)$ is defined to be $r(0, 0)$. This process was chosen since it prevents negative pure interest rates, and it is of simple form.

The second binomial stochastic process describes the evolution of the one-period hazard rate and the joint probability distribution between $r(t, i)$ and $h(t, j)$. Through covariance between the pure rate of interest and the hazard rate, correlation $\rho_{h,r}$ between these parameters is included. This covariance is scaled by the pure interest rate variance, leading to a multiplicative term that models the volatility of hazard rates as the sensitivity of hazard rates to the current pure rate of interest. Of course, this means (4) generates a recombining hazard rate process, since (3) is a recombining process. In (4), the time-dependent parameter v_t calibrates the hazard rate tree by forward induction through use of corporate state prices to the corporate yield curve and σ_h is the constant hazard rate volatility parameter. Note that when $t = 0$, then $h(t, j)$ is defined to be $h(0, 0)$.

Together, the binomial processes (3) and (4) form a model similar to Das & Tufano (1996) in that we assume a linear scaling of cashflows. By applying the law of iterated expectations, the two binomial trees (3) and (4) are combined to calculate corporate state prices, which forms a single binomial tree. Procedurally, we first calibrate the pure interest rate process at today's date $t = 0$ to the sovereign yield curve by adjusting the calibration factor u_t for all future dates. This obtains the pure interest binomial tree, the values $[r(t, i)]$ of which are included in the hazard rate process (4). We then calibrate the hazard rate process, which is correlated with the pure interest rate process generated by the first calibration, to a corporate yield curve by adjusting the calibration factor v_t. Simultaneously, this calibration adjusts the structure of corporate state securities until the yield implied by this replicating portfolio of corporate zeros agrees with our estimate of the corporate yield curve.

Substituting (3) and (4) into (2) and rewriting slightly, we obtain our binomial credit risk model. The result is as follows:

$$V_0 = e^{-\sum_{t=0}^{k=1}\left[r(0,0)e^{\left(u_t\Delta T + 0.5\sum_{i=0}^{1}(2i-t)\sigma_r\sqrt{\Delta T}\right)}\right]}$$
$$x\left[\left\{1 - h(0,0)e^{\left(\sum_{t=0}^{k=1}\left[v_t\Delta T + \rho_{h,r}\frac{\sigma_h}{\sigma_r}0.5\sum_{i=0}^{1} r(t,i)\Delta T\right]\right)}\right\} V_2 + \delta_2\right] \qquad (5)$$

where:

$$\delta_2 = h(0,0)e^{v_1\Delta T + \rho_{h,r}\frac{\sigma_h}{\sigma_r}0.5\sum_{i=0}^{1} r(1,i)\Delta T}$$
$$\omega_2\left(1 - h(0,0)\right) + h(0,0)\omega_1 e^{r(0,0)e^{\left(u_1\Delta T + 0.5\sum_{i=0}^{1}(2i-1)\sigma_r\sqrt{T}\right)}} \qquad (5a)$$

Equation (5) is our binomial model. This model allows for correlation of any type and for time-varying recovery rates. Equation (5a) is the recovery assumption where, conditional upon no prior default in any prior time and a pure interest rate state, an (possibly time-varying) amount ω is paid at the end of the current period. Should default occur at any prior time period, the recovery amount is reinvested in a Treasury security until promised maturity. The sum of these recovery amounts is then included in (5) and therefore these recovery amounts are expressed as a fraction of the value of a two-period Treasury zero.

Empirical procedures

We select all US Treasuries and double-A rated financial industry bonds that were quoted rather than matrix priced on June 30, 1988, from the University of Houston's fixed-income database. We choose bonds that have no call or put features. We select Treasury and banker acceptance interest rates as our shorter-term interest rates. Finally, we choose June 30, 1988, because we know from Standard & Poor's that the average recovery rate during the subsequent recession in 1990 was approximately 32.5%. We use this figure as our estimate of the recovery fraction.

3. State price bias and increasing recovery (upward sloping term structure of June 30, 1988)

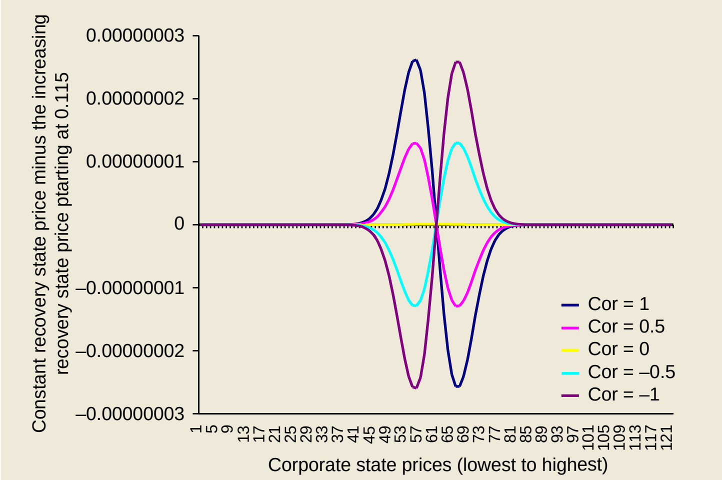

4. State price bias and decreasing recovery (upward sloping term structure of June 30, 1988)

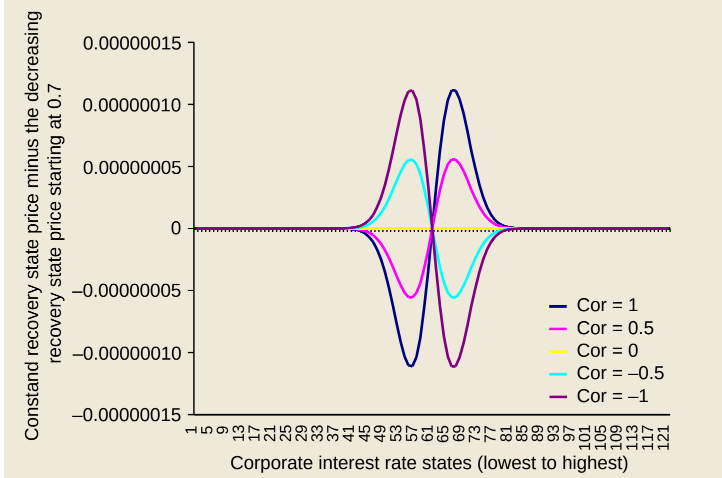

We then apply Vasicek & Fong (1982) to estimate the Treasury and double-A financial yield curves up to 10 years' maturity. To generate information about how changes in the credit spread may affect the influence of credit risk parameters, we estimate three sets of Treasury and double-A financial yield curves, one each on June 30, 1988, January 31, 1989, and April 30, 1990. On June 30, 1988, the credit spread is quite wide, 138 basis points on average, and the credit spread is fairly constant. In contrast, the January 31, 1989, credit spread widens with maturity but the average spread is now much smaller, at 81bp. Finally, the April 30, 1990, credit spread narrows to 42bp, and the credit spread is again fairly constant.

We then calibrate (5) to these yield curves using the procedures described in the previous section. To correspond to our monthly data observations, we use a monthly time step for a total of 120 steps for each 10-year yield curve. The base case uses June 30, 1988, yield curves, a constant recovery rate of 32.5%, a constant pure interest rate volatility of 10% and a constant hazard volatility of 1%.

First, we estimate (5) by setting the correlation between pure interest and hazard rates to zero. We then re-estimate (5) four times, assuming a constant recovery rate of 32.5%, but correlation of +1, +0.5, –0.5 and –1. We then re-estimate this set of correlated models along with the zero correlation case using two sets of time-varying recovery rates. Rather than keeping the recovery rate constant at 32.5%, the first time-varying recovery rate set smoothly increases the recovery rate from 11.5% to 70% by the 120th month, and the second set smoothly decreases the recovery rate from 70% to 11.5% by the 120th month.

Finally, we calibrate corporate interest rates to the corporate yield curve using (3) as our corporate interest rate process. This obtains estimates of corporate state prices in the same way we currently obtain pure interest state prices and so we make no attempt to parameterise the credit risk process.

Even the simplest parameterised model obtains large differences in state prices, when compared with a non-parameterised model

Empirical results

Now we plot differences in the corporate state prices generated by the various models at month 120. These state prices are today's value of hypothetical securities that promise to pay $1 only if the corresponding corporate interest rate state occurs at month 120. Hence, the y-axis reports the differences in today's cash price per dollar of promised future payout for this hypothetical security. Adding up all corporate state prices at month 120 replicates the value of a corporate zero that promises to pay $1 at maturity, no matter which corporate interest rate state occurs. Hence, areas under the curves represent differences in today's cash price of a corporate zero that promises to pay $1 in 10 years' time. These corporate state prices are used to price credit derivatives, so changes in the distribution of these corporate state prices caused by, say, high correlation imply that correlation may be important in modelling credit risk because adding correlation obtains different credit derivative prices.

To aid comprehension of the following figures, we always plot corporate state price differences as the "Strawman" model (the more simple model) minus the more complex model. This means that the distance from the x-axis represents the price bias created by using the Strawman model rather than the more complex model, assuming of course that the more complex model is "better". Furthermore, since call options pay out in low corporate interest rate states, then graphs plotting above the x-axis at low corporate interest rate states mean that calls are overpriced if we use the Strawman model and it is incorrect. Similarly, graphs plotting above the x-axis at high interest rate states mean that puts are overpriced.

Figure 1 shows that the non-parameterised corporate interest rate model generates corporate state prices that are substantially different from the corresponding zero correlation state prices. This suggests that even under the restrictive assumptions of zero correlation between credit risk and pure rates of interest and constant recovery rates, this parameterised credit risk model can achieve substantially more accurate derivative prices than those derived by a non-parameterised model. We observe that at low corporate interest rate states, the non-parameterised model leads to higher state prices. The opposite occurs at high interest rate states. In other words, if we believe that the zero correlation model is "more correct" than the non-parameterised model, then use of the non-parameterised model leads to overvalued call options and undervalued put options.

Notice that the bias in figure 1 decreases as we move from January 31, 1988, when the credit spread averaged 138bp, to April 30, 1990, when the credit spread averaged only 42bp. This suggests that the bias is related to the size of the credit spread rather than the shape of the credit spread. Furthermore, this suggests that if the credit spread is very narrow, there may be little point in parameterising credit risk since parameterised and non-parameterised credit risk models may yield similar prices.

Figure 2 compares zero correlation as the Strawman with a version that has non-zero correlation but constant recovery rates. The bias is related to the sign and size of correlation. For positive correlation, the Strawman

5. State price bias and zero correlation, time-varying recovery (upward sloping term structure of June 30, 1988)

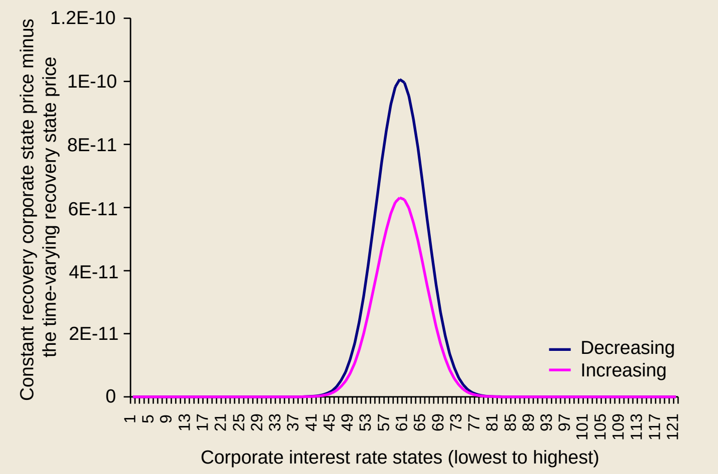

generates higher corporate state prices at lower corporate interest rate states and lower corporate state prices at higher corporate interest rate states. In contrast, negative correlation will lead to the oppose result.

Note that the size of the bias in figure 2 is less than the size of the bias reported in figure 1. However, the differences in figure 2 are still large. Imagine we are pricing a vulnerable European-style put option when credit risk and pure rates of interest rate have a negative 0.5 correlation.[1] Since the put will pay out at high corporate interest rate states, the maximum size of the price difference will be equal to the sum of price differences from corporate interest rate state 61 to state 120. This will be $0.0011 per dollar of future payout or $0.11 per hundred. For perfect negative correlation, the difference will be roughly double. Now imagine we are pricing a 10-year vulnerable interest rate cap, which pays off twice a year. As the size of the total difference will grow with more monthly time steps, we can roughly approximate the total price discrepancy using a linear approximation. This suggests that the cap may be overpriced by $1.10 per hundred. For investment banks that hold large positions, the absolute numbers will have a substantial impact.

Figures 3 and 4 compare a credit risk model with correlation but constant recovery rates as the Strawman, with a model that has correlation and time-varying recovery rates. Figure 3 uses a recovery rate that increases with time and figure 4 uses a recovery rate that decreases with time. For the upward sloping term structure of recovery rates, we observe the same pattern to pricing bias that we observed in figure 2. That is, for positive correlation, call options are overpriced and put options are underpriced, and for negative correlation, calls are underpriced and puts are overpriced. Figure 4 reports that for decreasing recovery rates, the opposite occurs. Figures 3 and 4 report biases that are roughly one-tenth the size of those in figure 2. This suggests that the potential improvement in pricing accuracy obtained by using a time-varying recovery rate is more modest than the potential improvement obtained by adding correlation and a constant recovery rate to the zero correlation model.

Figure 5 compares the zero correlation case as the Strawman with a binomial version that similarly assumes zero correlation, but with time-varying recovery rates. The state price differences thereby obtained are trivial. This finding supports the view that time-varying recovery rates are of secondary importance.[2]

Nevertheless, we note that this conclusion is based on changes in the distribution of state prices. For some securities, such as credit default swaps, a portion of the state price related to payments in the event of default forms a disproportionate part of the value of the security. The remaining (and usually much larger) portion of the state price related to survival contingent values have no direct influence on the value of default contingent payouts. Furthermore, this survival contingent value is more influenced by the correlation with pure rates of interest than the default contingent value simply because it forms a larger portion of the full state price. Hence, it is possible that for default contingent credit derivatives, such as credit default swaps, the recovery assumption may prove to be important, yet correlation with pure rates of interest is less important. This is precisely the opposite conclusion we reach when examining the distribution of the full state security price. This leads us to suggest that whether or not the correlation or recovery fraction is important for modelling credit derivatives depends upon the task at hand.[3]

Conclusions

By examining the behaviour of state prices obtained from a binomial credit risk model, we are able to suggest which credit risk parameters are of critical interest. It appears worthwhile to parameterise credit risk, since even the simplest parameterised model obtains large differences in state prices when compared with a non-parameterised model. While correlation between pure rates of interest and credit risk and time-varying recovery rates both appear influential in determining state prices, correlation appears more influential than time-varying recovery rates.

The latter conclusion is valid for all derivatives whose price is dependent upon both the survival and default contingent portions of the state price. However, unlike vulnerable options for example, credit default swap values depend more upon the default contingent portion of the state price, so we may reach precisely the opposite conclusion, namely that correlation is less important than time-varying recovery rates. This suggests that which of these parameters are the most important depends upon the task at hand.

Finally, apparent differences in state prices obtained as we vary recovery assumptions and parameter estimates appear related to the size rather than the shape of the credit spread. This suggests that if there is little credit risk when, say, examining US government agency yield curves, there is little point to parameterising credit risk. But if credit risk is large when, say, examining emerging market sovereign yield curves, how we parameterise credit risk becomes a critical issue. ■

Frank Skinner is an associate professor in the ISMA Centre at the University of Reading and Antonio Díaz is an associate professor in the Departamento de Economía y Empresa at the Universidad de Castilla – La Mancha.
They would like to thank Philip Schönbucher for important comments

Comments on this article can be posted on the technical discussion forum on the Risk website at http://www.risk.net

REFERENCES

Black F, E Derman and W Toy, 1990
A one-factor model of interest rates and its application to Treasury bond options
Financial Analysts Journal 46, pages 33–39

Das R and P Tufano, 1996
Pricing credit sensitive debt when interest rates, credit ratings and credit spreads are stochastic
Journal of Financial Engineering 5, pages 161–198

Duffie D and K Singleton, 1999
Modeling term structures of defaultable bonds
Review of Financial Studies 12, pages 687–720

Vasicek O and H Fong, 1982
Term structure modeling using exponential splines
Journal of Finance 37, pages 339–348

[1] *By vulnerable, we mean that the underlying asset is not subject to credit risk, but the writer is*
[2] *Using the non-constant credit spreads of January 31, 1989, and the narrow credit spread of April 30, 1990, we obtained basically the same result*
[3] *We replicated figures 3 and 4, which isolate the impact of time-varying recovery rates but include non-zero correlation, for the widening credit spread of January 31, 1989, and the narrow constant credit spread of April 30, 1990. We obtain the same pattern of price biases as reported in figures 3 and 4, except that the size of these biases decrease with the size rather than the shape of the credit spread*