2015 Enterprise Risk Management Symposium

June 11–12, 2015, National Harbor, Maryland

A Cost of Capital Approach to Credit and Liquidity Spreads

By B. John Manistre

A Cost of Capital Approach to Credit and Liquidity Spreads

B. John Manistre[1]

Version 3.4, Dec. 12, 2015

Abstract[2]

The market cost of capital approach has emerged as the standard for estimating risk margins for insurers' fair value balance sheets. This paper takes some of the ideas developed for valuing life insurance liabilities and applies them to the problem of valuing credit-risky bonds. The basic idea is that credit spreads should cover the cost a) best-estimate defaults plus the cost of holding capital for b) contagion risk (e.g., a credit crunch) and c) parameter risk, the risk that the best estimate is wrong and must be revised. We argue that the margins required for parameter risk can capture liquidity issues. In addition, the models developed here allow the cost of capital rate itself to be a random variable that allows credit spreads to open and close stochastically.

Finally, the paper argues that it is reasonable to include something like AA best-estimate default rates and liquidity spreads when valuing insurance liabilities. The main rationale for doing so is the idea that there are elements of the total credit risk issue which can be hedged between the assets and liabilities.

Introduction

It is now 20 years since the Society of Actuaries held its first research conference dedicated to the topic of how one should determine the fair value of life insurance liabilities. One particularly vexing issue has always been the question of what yield curve one should use to discount some appropriately risk-adjusted sets of liability cash flows.

Some have argued that if a life insurer wants to be considered an AA risk, it should use an AA yield curve to discount its liabilities. While this makes some intuitive sense, it can lead to the conclusion that the lower a life insurer's credit rating is, the lower its liabilities. This is controversial.

If we try to fix this problem by ignoring all credit-spread issues in the liability valuation, we create another problem that many companies trying to develop market-consistent reporting models experienced during the financial crisis of 2008. At the height of the crisis, the flight to quality decreased yields on risk-free instruments (raising liability values) while credit spreads increased enough to more than offset the drop in risk-free yields. The result, for many companies, was that the market value of their assets dropped while the fair value of their liabilities increased. Many people, including this author, consider that result to be somewhat inappropriate, especially if assets and liabilities were reasonably well "matched" going into the crisis.

[1] The author is a research actuary at GGY AXIS based in Baltimore, Md.

[2] The views and opinions expressed in this paper are those of the author and not GGY AXIS.

The European Union's Solvency II Quantitative Impact Study (QIS) 5 specification tried to address the issues outlined above by allowing insurers to use interbank swap rates plus a spread designed to remove credit risk and capture the idea that many insurance obligations are fairly illiquid.[3] To the extent the allowed liquidity spread varies with market conditions, there is an element of credit spread hedging going on between the assets and liabilities. This goes some way to resolving issues raised above.

The models developed in this paper ultimately lead to a liquidity adjustment for liability valuation as in the Solvency II QIS 5 approach. However, the path we take to get there and the resulting pattern of liquidity adjustments are quite different.

Following this introduction, the paper is divided into three main sections.

1. We develop the key ideas in a simple two-state model where a bond is either in good standing or in default. This is where the analogy to life insurance is clearest. We are able to formulate a reasonably tractable affine model where the cost of capital rate is a random variable. For simplicity, we assume the risk-free interest rate is deterministic.

   The main conclusion is that the forward default rates for a credit-risky bond can be decomposed into the sum of

   i. A best-estimate default rate (i.e., one that assumes the law of large numbers for a portfolio of bond applies).
   ii. A spread for contagion risk, the risk that current experience differs materially from best estimate.
   iii. A dynamic spread for parameter risk, the risk that the best estimate is wrong and must be revised. We argue this spread also captures liquidity issues.
   iv. A final spread that arises if the cost of capital rate itself is stochastic.

   Each of items ii–iv is engineered to provide for holding a specific amount of economic capital.

   After this step, we have enough theory to develop the author's views on how credit risk issues should affect insurance liability valuation. The main points are

   a) On the asset side, the insurer should hold capital for each of risks ii–iv in the list above.
   b) For long liabilities, we can take a capital offset for risks iii–iv, because these risks are being effectively hedged. The contagion risk issue cannot be hedged.
   c) As a result of (b) it makes sense to build spreads iii–iv into the liability valuation. We also argue that best-estimate defaults should be included.

   One consequence of taking this point of view is that the net credit risk capital requirement for an insurer whose assets and liabilities are "matched" in the sense defined here is reduced to the contagion risk requirement. On the other hand, a financial

[3] CFO/CRO Forum, "QIS 5 Technical Specification: Risk-Free Interest Rates," 2010, http://ec.europa.eu/internal_market/insurance/docs/solvency/qis5/cfo-forum-cro-forum-paper-risk-free-rates_en.pdf.

institution, such as a bank, that does not have long liabilities would presumably have to hold more credit risk capital due to the larger mismatch between assets and liabilities.

2. The next step is to generalize the simple model to a multistate world where the best-estimate model is represented by a typical credit transition matrix. While there are many different ways in which the two-state model can be generalized, we choose an approach to minimize the technical details. The end result is a model in the same general family as a model published in 1997 by Jarrow, Lando and Turnbull.[4]

   We don't develop this model in a lot of detail. Instead, the point is to show that the additional complexity can be handled in a reasonable way.

3. The final section is a statement of conclusions.

There are two issues that are deliberately out of scope for this paper. These are

1. A calibration of the model to observed market data although we do briefly indicate how such a calibration could be performed, and
2. A detailed comparison with the current Solvency II approach to liability spreads.

These are both significant issues that deserve discussion. Unfortunately, an appropriate treatment of these issues could easily triple the length of the current paper. The author hopes that other interested risk professionals will rise to the occasion and engage in this important discussion.

The Two-State Model

The starting point for our model is the actuary's cost of capital approach. This is usually a three-step process where we start with a best-estimate model of the default process and then adjust that model for two kinds of risk: a) the risk that current experience differs materially from the best estimate and b) the risk that the best estimate is wrong and must be revised.

A key difference between the mortality model and the credit risk model developed here is that we also take market sentiment into account by allowing the cost of capital rate to be stochastic rather than a constant. This is a fourth step.

The author's life insurance version of this idea[5] was developed in detail in a paper presented at the ERM symposium held in Chicago in October 2014. A very short summary of that paper's main conclusion is that we can capture the essential elements by taking our best-estimate forward rates of default and adding a static margin for current experience risk and a dynamic margin for parameter risk.

[4] Robert Jarrow, David Lando, and Stuart M. Turnbull, "A Markov Model for the Term Structure of Credit Risk Spreads," *The Review of Financial Studies* 10, no. 2 (1997): 481–523.

[5] B. John Manistre, "Down but Not Out: A Cost of Capital Approach to Fair Value Risk Margins." This paper can be found on the Society of Actuaries' (SOA) website. A short summary of the paper appeared in the fall 2014 issue of the SOA's section newsletter *Risk Management*.

While this model was originally developed to value non-hedgeable insurance risk, there are insights that can be gained by applying the model to valuing credit risk. We will use those insights to justify our approach to applying credit spreads to the valuation of insurance liabilities.

The Best-Estimate Model

Let $V_0 = V_0(t, T)$ be the value at time t of a credit-risky bond that is scheduled to mature at time T for \$1 if it has not already defaulted. No other cash flows are assumed.

In a simple two-state model, the first key credit risk parameter is the best-estimate force of default $\mu_0(t)$ and the second key assumption is the residual value RV of the credit-risky bond once it has actually gone into default. A valuation equation that captures these two assumptions is

$$\frac{dV_0}{dt} + \mu_0(RV_0 - V_0) = rV_0, \;\; V_0(T, T) = 1.$$

What this equation says is that the total rate of change of the bond's value due to both the passage of time and the occurrence of defaults is equal to the risk-free rate r, which we assume is constant for now.

The resulting value V_0 is then calculated by discounting the maturity value with interest r and a recovery adjusted force of default $\mu_0(1 - R)$, i.e.

$$V_0(t, T) = \exp[-\int_t^T \{r + \mu_0(1 - R)\} ds].$$

This looks a lot like traditional actuarial discounting and could apply to a portfolio of bonds that was structured so the law of large numbers could be applied to average out the experience.

Adverse Current Experience: *Static Margins*

The only thing we know for sure about our best-estimate model is that it is wrong. No matter how much effort we put into developing our best-estimate assumptions, we cannot predict default costs precisely. For our two-state example, we will assume that, even if our model is correct in the long run, it is still possible to experience n years of best-estimate credit risk losses in a single year.

To protect its solvency, a risk enterprise that owns a portfolio of risky bonds should hold enough economic capital that it could withstand the adverse event if it occurred. The kind of event we have in mind is one where the law of large numbers would not help. In the mortality application, this might be the onset of a pandemic. In the bond application, this could be an economic crisis where many bonds are suddenly downgraded.

If $V(t, T)$ is the new risk-adjusted value of the bond, then the risk enterprise needs to hold economic capital in the amount $n\mu_0(1 - R)V$. The new valuation equation, which reflects the cost of holding this amount of risk capital, is now

$$\frac{dV}{dt} + \mu_0(RV - V) = rV + \pi n \mu_0 (1 - R)V, \qquad V(T,T) = 1.$$

Here π is a deterministic cost of capital rate that we will discuss in more detail later. To the extent that the economic capital was invested at the risk-free rate r, the total expected return to an investor putting up the risk capital is $r + \pi$.

A little algebra shows that the valuation equation above is equivalent to

$$\frac{dV}{dt} + \mu_0(1 + n\pi)(RV - V) = rV, \qquad V(T,T) = 1.$$

The important observation is that taking short-term risk issues into account is equivalent to adding a simple loading $\pi c = \pi n \mu_0$ to the best-estimate default rate. The resulting risk-adjusted value can then be calculated as

$$V(t,T) = \exp[-\int_t^T \{r + (\mu_0 + \pi c)(1 - R)\} ds].$$

Assumption Risk and Liquidity: *Dynamic Margins*

The previous analysis assumed our model was basically right over the long term but was vulnerable to adverse short-term fluctuations. We now relax the long-term assumption and ask what happens if we decide our best-estimate default cost $\mu = \mu_0(1 - R)$ is wrong and must be revised to a new value $\hat{\mu} = (\mu_0 + \Delta\mu)(1 - R)$. The economic loss that would occur if this happens is the difference between two bond values $V - \hat{V}$, each based on its own default assumption. The idea now is to work out what it means to build in the cost of holding this amount of capital into the bond value.

In a principles-based insurance model, assumptions are being examined and revised all the time as new information becomes available. To the extent that the insurance models are well understood, and based on credible experience, the potential assumption shock $\Delta\mu$ should be small. However, if the business is not that well understood, the assumption shock should be larger. At a high level, the size of the assumption shock therefore reflects the core elements of the liquidity issue.

We now examine the mathematical consequences of building in risk margins for holding this kind of economic capital. Once we have done this, the case for taking the assumption shock approach to modeling liquidity risk gets even stronger.

If we continue on the same path, we would write down a valuation equation of the form

$$\frac{dV}{dt} + \mu_0(RV - V) = rV + \pi c(1 - R)V + \pi(V - \hat{V}), \qquad V(T,T) = 1.$$

Unfortunately, this approach raises the thorny issue of how to calculate the shocked value $\hat{V}(t,T)$. The obvious next step is to write down a valuation equation for $\hat{V}$ of the form

$$\frac{d\hat{V}}{dt} + (\mu_0 + \Delta\mu)\big(R\hat{V} - \hat{V}\big) = r\hat{V} + \pi c(1-R)\hat{V} + \pi(\hat{V} - \hat{\hat{V}}), \qquad \hat{V}(T,T) = 1.$$

This equation makes the reasonable assumption that the shocked default rate is $\mu_0 + \Delta\mu$ but has introduced a second shocked value $\hat{\hat{V}}$ which presumably depends on some secondary shocked default rate like $\mu_0 + \Delta\mu + \Delta\mu$. Writing down a valuation equation for $\hat{\hat{V}}$ simply leads to the same problem. This is called the circularity problem in the insurance literature.

A number of practical ways to resolve the circularity problem were discussed in the author's paper presented at the 2014 ERM symposium. The approach taken here was called the explicit margin method in that paper. We assume there is a margin variable β that allows the risk-loaded default rate to be written for times $s > t$ as

$$\mu(s) = \mu_0(s) + \ \pi c(s) + \beta(t,s)\Delta\mu(s).$$

The margin variable β is assumed to be 0 in the real world so $\beta(t,t) = 0$, but for $s > t$, it is assumed to evolve according to the dynamical rule

$$d\beta = [\pi - \beta\Delta\mu(1-R)]ds.$$

When we come to do a new valuation at some later point in time $t' > t$, the margin variable resets to zero at that time. Chart 1 illustrates the idea just described.[6]

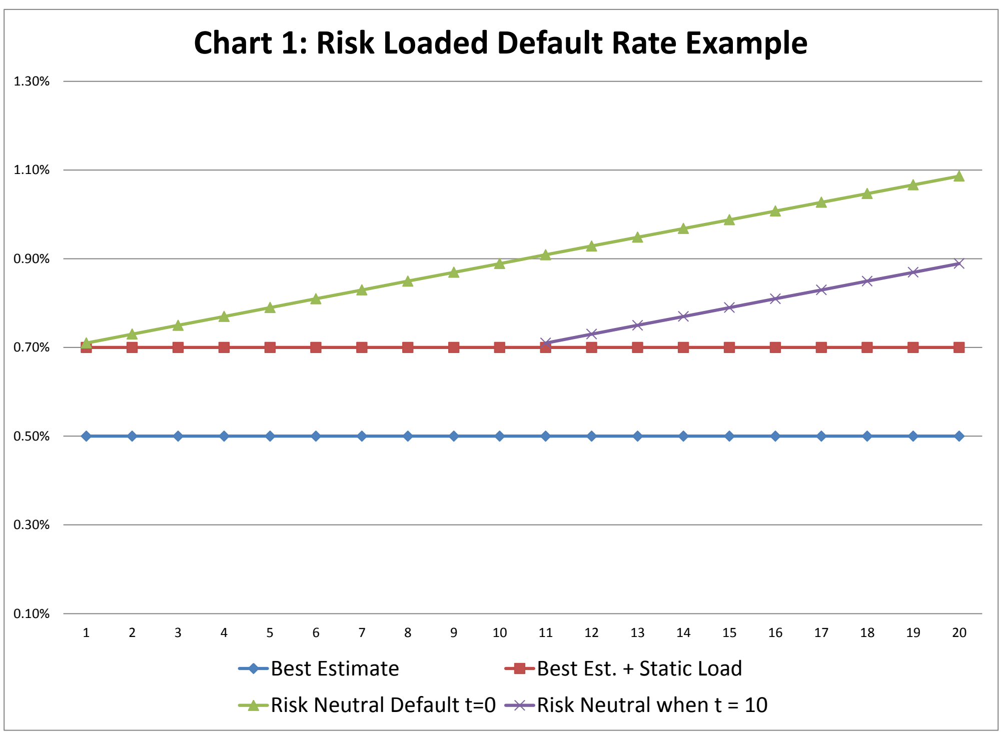

Chart 1: Risk Loaded Default Rate Example

[6] Under the stated assumptions, we can calculate the margin variable in closed form as $\beta(t,s) = \pi\frac{1-e^{-\Delta\mu(1-R)(s-t)}}{\Delta\mu(1-R)}$ which is approximately $\pi(s-t)$ when $\Delta\mu(1-R)$ is small. Later in this paper it will be convenient to approximate the dynamics of β by $d\beta = \pi ds$.

In the example above, the best-estimate forward default rate is $\mu_0 = .50\%$. The static load assumes a short-term shock equal to four years' worth of best-estimate defaults and a cost of capital rate equal to $\pi = 10\%$. The resulting load is then $.10 \times 4 \times .005 = .20\%$.

The dynamic load calculation assumes the same 10 percent cost of capital rate but then uses a parameter shock of $\Delta\mu = .20\%$ and a recovery rate of $R = 50\%$. Under these assumptions, the margin variable β grades from 0.00 to roughly 1.98 over the 20-year projection. When we come to do a new valuation 10 years later, the static load has not changed but the dynamic load has been pushed out.

Deriving the dynamic load model from first principles is beyond the scope of this paper but it is fairly easy to show that the resulting structure has the desired properties. To do this, we think of the value of the risky bond, in the valuation measure, as a function $V = V(s, \beta, T)$ that for $s \geq t$ it satisfies the following evolution equation,

$$\frac{\partial V}{\partial s} + [\pi - \beta\Delta\mu(1-R)]\frac{\partial V}{\partial \beta} = [r + (\mu_0 + \pi c + \beta(t,s)\Delta\mu)(1-R)]V, \qquad V(T,\beta,T) = 1.$$

What this equation says is that the total expected rate of change of the value, in the valuation measure, is equal to the risk-free interest rate plus the appropriate risk-loaded default cost. On the valuation date, when $s = t$ and $\beta(t,t) = 0$, the real-world expected rate of change is

$$\frac{\partial V}{\partial s}|_{s=t} + \mu_0(t)(R-1)V(t) = [r(t) + \pi c(t)(1-R)]V(t) - \pi\frac{\partial V}{\partial \beta}|_{s=t}.$$

This result makes sense if the greek $\delta = \frac{\partial V}{\partial \beta}|_{s=t}$ can be interpreted as a reasonable negative amount of economic capital to hold for parameter risk, i.e., $= \hat{V} - V$. That this is in fact the case can be shown by differentiating the valuation equation above with respect to the margin variable β to get

$$\frac{\partial \delta}{\partial s} + [\pi - \beta\Delta\mu(1-R)]\frac{\partial \delta}{\partial \beta} = [r + (\mu_0 + \Delta\mu + \pi c + \beta(t,s)\Delta\mu)(1-R)]\delta + \Delta\mu(1-R)V.$$

The solution to this equation at time t is

$$\delta(t) = -\int_t^T e^{-\int_t^s [r + (\mu_0 + \Delta\mu + \pi c + \beta(t,v)\Delta\mu)(1-R)]dv}\,\Delta\mu(1-R)V(s)ds.$$

This result tells us that the greek δ, at the valuation date, is minus the present value of losses that would occur if we valued the risky bond assuming the best-estimate default rate was μ_0 but the actual experience turned out to be $\mu_0 + \Delta\mu$. This means $\delta = \hat{V} - V$ where $\hat{V}$ is calculated using shocked best-estimate default rates but with the same static and dynamic default rate loadings deemed appropriate for the base case. Variations on this theme are possible.

We conclude that the dynamic margin mechanism is releasing margin into income in a way that is consistent with the cost of capital concept provided the greek method outlined above is used to estimate the economic capital for parameter risk.

Chart 2 below is the same as Chart 1 except that we have added the “capital” forward default rate scenario used to compute $\widehat{V}$ at the valuation date.

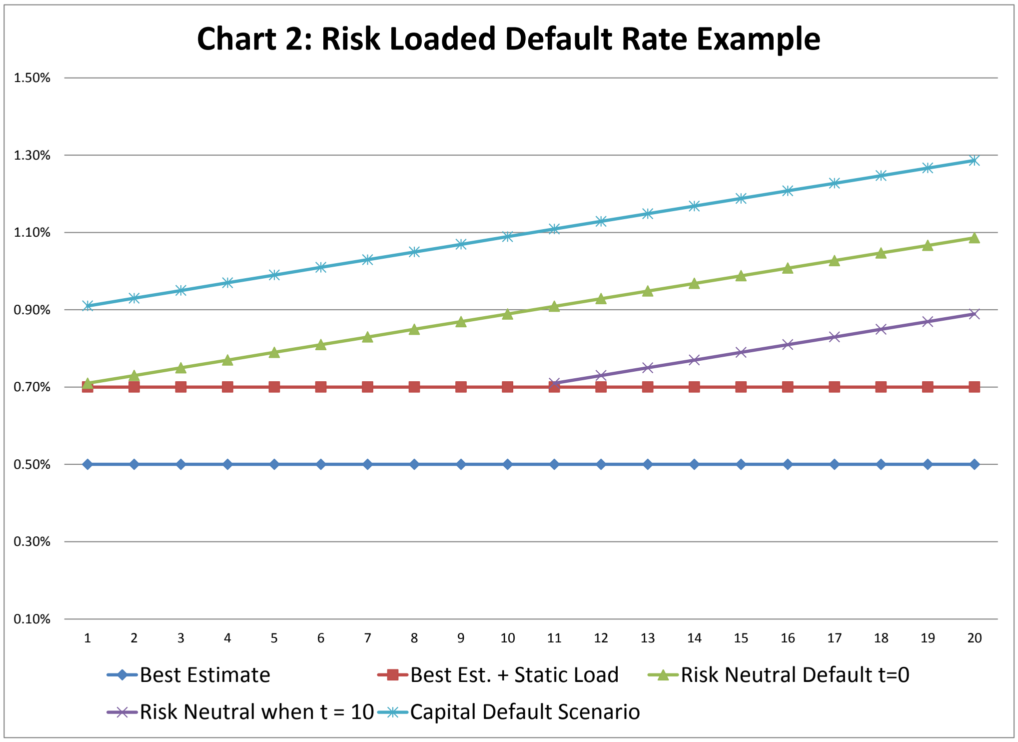

We can now develop the argument that we can capture bond liquidity issues by varying the size of the parameter shock. Liquid bonds would have small parameter shocks while illiquid bonds should have larger parameter shocks, even if they have the same short-term best-estimate default rates. This suggests a bond valuation model in which a risky bond’s value is driven by both its rating (a proxy for current best-estimate default rate) and a liquidity-rating factor embedded in the shock $\Delta\mu$.[7]

A short summary of the argument so far follows.

The valuation model developed has three components:

1. A best-estimate default assumption.
2. A capital requirement and associated static risk loading to deal with the risk that the best estimate could be wrong in the short run, i.e., a credit crunch.
3. A capital requirement and associated dynamic risk loading to deal with the risk that the best estimate itself could be wrong and require revision.

[7] For most bonds, it would be logical to assume $\Delta\mu > 0$ but for some applications this might not be the case. For example, people often take the swap curve to define the risk-free rate. If we do this, we might have to use a negative parameter shock to model U.S. Treasury bonds. Note this leads to the conclusion that the economic capital required for holding Treasury bonds could be negative. This makes sense in the context of the current model. It makes even more sense once we consider flight-to-quality issues in the next section since a credit crisis often raises the value of high quality sovereign debt.

We argued that we can capture liquidity issues with this last element.

One thing this model cannot explain yet is the way credit spreads open and close in a seemingly random fashion. A classic example of such a phenomenon would be a flight to quality. Note this is different from a contagion event because it does not necessarily imply any ratings downgrades.

Market Sentiment and Stochastic Cost of Capital Rates

When the model described above is applied to value life insurance underwriting risk, one normally assumes a constant cost of capital rate such as $\pi = 6.00\%$. For the credit risk example, it makes more sense to allow the cost of capital rate itself to be a random quantity in order to allow the credit risk spreads to open and close with changing market sentiment.

We therefore extend the model by allowing the cost of capital rate π to follow a Cox-Ingersoll-Ross process of the form

$$d\pi = \kappa(\pi_\infty - \pi)dt + \xi\sqrt{\pi}dz.$$

This will lead to a tractable affine model with a nonnegative cost of capital rate. For the moment, we continue to assume a deterministic risk-free interest rate r and we will also be deliberately ambiguous as to whether these are P measure (real-world) or Q measure (risk-adjusted) parameters.

If the cost of capital rate is stochastic, then we have to think of the risky bond's value V as a function of time and two-state variables β, π, i.e., $V = V(s, \beta, \pi, T)$. The fundamental valuation equation then generalizes to

$$\begin{aligned}&\frac{\partial V}{\partial s} + [\pi - \beta\Delta\mu(1-R)]\frac{\partial V}{\partial \beta} + \kappa(\pi_\infty - \pi)\frac{\partial V}{\partial \pi} + \frac{\xi^2\pi}{2}\frac{\partial^2 V}{\partial \pi^2}\\ &= [r + (\mu_0 + \ \pi c + \beta(t,s)\Delta\mu)(1-R)]V, \qquad V(T,\beta,\pi,T) = 1.\end{aligned}$$

This equation can be solved by using the well-known financial engineering trick of assuming an affine solution of the form

$$V(s,\beta,\pi,T) = \exp[A(s,T) + \beta B(s,T) + \pi P(s,T)].$$

We can then derive the following system of ordinary differential equations for the quantities A, B, P, as a function of s,

$$\begin{aligned}&\dot{A} + \mu_0(R-1) + \kappa\pi_\infty P = r, \ \ A(T,T) = 0,\\ &\dot{B} + \Delta\mu(R-1) = \Delta\mu(1-R)B, \ \ B(T,T) = 0,\\ &\dot{P} - \kappa P + \tfrac{\xi^2}{2}P^2 = c(1-R) - B, \ \ P(T,T) = 0.\end{aligned}$$

We will analyze this system of equations assuming that

- The parameters $R, \Delta\mu, c$ are constants as are κ, π_∞ and ξ.
- The quantities μ_0 and r are deterministic functions of time

These assumptions can be relaxed at the expense of greater analytical complexity. The main benefit of doing this is that the two functions $B(s,T), P(s,T)$ will both satisfy a relation of the form $F(s,T) = F(T-s)$.

Before solving these equations, we do some analysis to get a sense of how these quantities affect the model's forward discount rates. These rates are defined by

$$F(t,T) = -\frac{\partial}{\partial T}\ln(V(t,\beta,\pi,T),$$
$$= -\left[\frac{\partial A}{\partial T} + \beta\frac{\partial B}{\partial T} + \pi\frac{\partial P}{\partial T}\right].$$

On the valuation date t the margin variable β is zero so this simplifies to

$$F(t,T) = -\left[\frac{\partial A(t,T)}{\partial T} + \pi\frac{\partial P(t,T)}{\partial T}\right].$$

From the equation for dA/ds above we can write

$$A(T,T) - A(t,T) = \int_t^T \frac{dA(s,T)}{ds}ds,$$
$$= \int_t^T [r + \mu_0(1-R) - \kappa\pi_\infty P(s,T)]ds.$$

From which it follows that

$$-\frac{\partial A(t,T)}{\partial T} = r(T) + \mu_0(T)(1-R) - \kappa\pi_\infty P(T,T) - \kappa\pi_\infty \int_t^T \frac{\partial P(s,T)}{\partial T}ds.$$

We'll soon see that, when $\Delta\mu$ is a constant, $P(s,T) = P(T-s)$ so $\frac{\partial P(s,T)}{\partial T} = -\frac{\partial P(s,T)}{\partial s}$. This allows the last term in the equation above to be integrated to get

$$-\frac{\partial A(t,T)}{\partial T} = r(T) + \mu_0(T)(1-R) - \kappa\pi_\infty P(t,T).$$

A similar analysis can be used to show that, since $B(s,T) = B(T-s)$,

$$-\frac{\partial P(t,T)}{\partial T} = c(1-R) + \kappa P(t,T) - B(t,T) - \frac{\xi^2}{2}P^2(t,T)$$

Putting the two results above together, we get the following expression for the model's forward discount rates

$$F(t,T) = r(T) + (\mu_0(T) + \pi c)(1-R) - \pi B(t,T) + \kappa(\pi - \pi_\infty)P(t,T) - \frac{\pi\xi^2}{2}P^2(t,T).$$

We now work through this expression in three steps. The first step assumes $\pi = \pi_\infty$ and $\xi = 0$, i.e., the cost of capital rate is just a constant. In this case, the discount rates are given by

$$F(t,T) = r(T) + (\mu_0(T) + \pi c)(1-R) - \pi B(t,T).$$

The middle equation for $B(t,T)$ does not depend on any of the other variables and is simple enough that it can be solved in closed form. The result is

$$B(t,T) = -(1 - \exp[-\Delta\mu(1-R)(T-t)]) \approx -\Delta\mu(1-R)(T-t),$$

so

$$\begin{aligned} F(t,T) &= r(T) + (\mu_0(T) + \pi c)(1-R) + \pi(1 - \exp[-\Delta\mu(1-R)(T-t)]), \\ &\approx r(T) + \big(\mu_0(T) + \pi(c + \Delta\mu(T-t)\big)(1-R). \end{aligned}$$

This is exactly the same result we got in the previous section of this paper. The formulae reconcile once you notice that $-\pi B(t,T) = \beta(t,T)\Delta\mu(1-R)$.

The discount rate now has four components:

1. The risk-free rate $r(T)$.
2. The best-estimate default cost $\mu_0(T)(1-R)$.
3. A static risk loading $\pi c(1-R)$ for contagion risk.
4. A dynamic risk loading for parameter/liquidity risk that is about $\pi\Delta\mu(T-t)(1-R)$.

The fact that the credit spread for liquidity risk is zero at the valuation date and then grows with time is a reasonable property of the model. There is no reason for a cash flow due almost immediately to be illiquid. The longer the maturity of the cash flow, the more risk there is in any default assumption we develop, so the growing risk load makes sense.

We also saw in the previous section that the economic capital required for parameter/liquidity risk is

$$-\frac{\partial V}{\partial \beta} = -B(t,T)V(t,\beta,\pi,T) \approx \Delta\mu(1-R)(T-t)V(t,\beta,\pi,T).$$

This is a reasonable formula for the assumed liquidity capital that varies in a simple way by maturity. The longer the bond, the riskier it appears to be. This makes intuitive sense and differs from the contagion risk capital factor $c(1-R)$, which does not depend on the maturity of the risky bond.

For a more complex bond with coupon payments, the formula for liquidity risk capital clearly generalizes to

$$EC \approx \Delta\mu\,(1-R)DV,$$

where D is the usual (modified) duration of the bond.

Now that we know $B(t,T)$, we can substitute it into the third equation for $P(t,T)$ to get

$$\dot{P} - \kappa P + \frac{\xi^2}{2} P^2 = c(1-R) + (1 - \exp[-\Delta\mu(1-R)(T-t)]).$$

The case $\xi = 0$

When $\xi = 0$, the equation for P above can also be solved in closed form to get a function we will call $P_0(t,T)$. The result is

$$P_0(t,T) = \frac{e^{-\Delta\mu(1-R)(T-t)} - e^{-\kappa(T-t)}}{\kappa - \Delta\mu(1-R)} - (c(1-R)+1)\frac{1 - e^{-\kappa(T-t)}}{\kappa}.$$

When the term $\kappa(\pi - \pi_\infty)P_0(t,T)$ is added to the expression for the forward discount rates, we find

$$\begin{aligned} F(t,T) = r(T) + \big(\mu_0(T) + c\{\pi_\infty + (\pi - \pi_\infty)(1 - e^{-\kappa(T-t)}\}\big)(1-R) \\ + \pi_\infty\big(1 - e^{-\Delta\mu(1-R)(T-t)}\big) \\ +(\pi - \pi_\infty)\tfrac{\Delta\mu(1-R)}{\Delta\mu(1-R)-\kappa}\big(e^{-\kappa(T-t)} - e^{-\Delta\mu(1-R)(T-t)}\big). \end{aligned}$$

This is the result we should expect if the cost of capital rate is undergoing a deterministic mean reversion from its current value of $\pi(t) = \pi$ to its long-run target of π_∞, i.e.,

$$\pi(t,s) = \pi_\infty + (\pi - \pi_\infty)e^{-\kappa(T-t)}.$$

The contagion spread term is just $\pi(t,s)c(1-R)$ and the liquidity spread has the form $\beta(t,s)$ $\Delta\,\mu(1-R)$ where $d\beta = [\pi(t,s) - \beta\,\Delta\,\mu(1-R)]ds$.

A quantity of some interest is the sensitivity $-\frac{\partial \ln(V)}{\partial \pi} = -P$, which is another duration-like quantity that measures the impact on value if the cost of capital rate itself were to change. We will therefor call $-P$ the *capital duration*. In the simple case where $\pi = \pi_\infty$, we can see that the capital duration is just the total economic capital required for contagion and liquidity risk.

Asset/liability management in a world driven by this kind of model would want to keep the difference between the capital and liquidity durations of assets and liabilities under close scrutiny.

The case $\xi \neq 0$

We now discuss the case where the volatility ξ of the cost of capital rate is nonzero. In this situation, the differential equation for $P(t,T)$ does not have a simple closed-form solution although it can be simplified by introducing a new unknown function $U(t,T)$ such that $P(t,T) = \frac{2}{\xi^2}\frac{\dot{U}}{U}$. This will give us the right behavior as long as U satisfies the linear second-order differential equation[8]

$$\ddot{U} - \kappa\dot{U} = \frac{\xi^2}{2}\{c(1-R) + (1 - \exp[-\Delta\mu(1-R)(T-t)]\}U.$$

[8] This is a standard applied math trick for solving differential equations of the Ricatti type.

This equation can be attacked numerically, or, if $\Delta\mu$ is constant, by assuming a power series solution of the form

$$U(t,T) = \sum_{j=0}^{\infty} a_j [T-t]^j,\ a_0 = 1,\ a_1 = 0.$$

A recurrence relation for the constant coefficients a_j is

$$a_{j+2} = -\kappa \frac{a_{j+1}}{j+2} + \frac{\xi^2}{2} \frac{1}{(j+2)(J+1)} [c(1-R)a_j - \sum_{k=1}^{j} \frac{(-\Delta\mu(1-R))^k}{k!} a_{j-k}].$$

This series converges quickly and is fairly easy to implement in a spreadsheet environment.

It is not hard to show that, if $\Delta\mu > 0$, then $0 \geq P(t,T) \geq P_0(t,T)$ and the new forward discount rates are then given by

$$\begin{aligned} F(t,T) &= r + \mu_0(1-R) + \pi c(1-R) \\ &\quad + \pi\left(1 - e^{-\Delta\mu(1-R)(T-t)}\right) + (\pi - \pi_\infty)\kappa P - \frac{\xi^2}{2} P^2, \\ &= F_0(t,T) + \kappa(\pi - \pi_\infty)(P - P_0) - \pi \frac{\xi^2}{2} P^2,\ \ (P > P_0). \end{aligned}$$

Interestingly, there is no simple conclusion which states that allowing $\xi > 0$ makes the forward default rates go consistently up or down. The last term, $\pi \frac{\xi^2}{2} P^2$, clearly reduces the forward default rates but that may, or may not, be offset by the term $\kappa(\pi - \pi_\infty)(P - P_0)$. Only when $\pi < \pi_\infty$ is it clear that using a nonzero volatility ξ reduces the forward default rates.

One definite conclusion we can draw is that assuming a stochastic cost of capital makes risky bond values less sensitive to a change in the cost of capital rate.

Chart 3 below shows the impact of using a nonzero cost of capital rate volatility under the assumptions that $\pi = \pi_\infty = 10\%,\ \kappa = 15\%$ and $\xi = 50\%$.

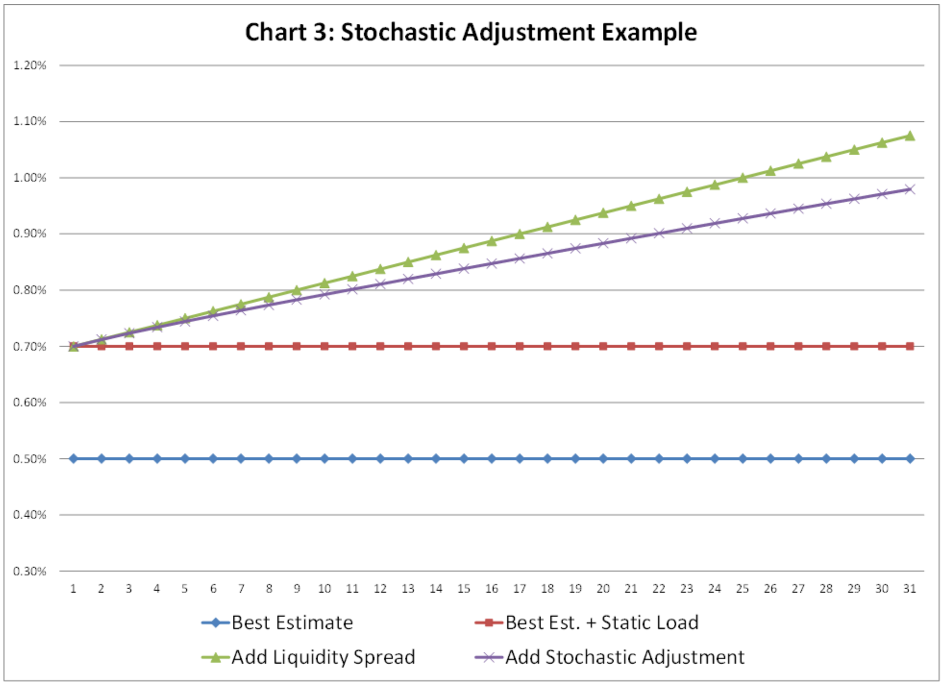

Chart 4 is presented to show what happens if the cost of capital rate at the valuation date is not equal to the long-term mean. This particular example assumes $\pi = 15\%$ at the valuation date with all other parameters unchanged from Chart 3.

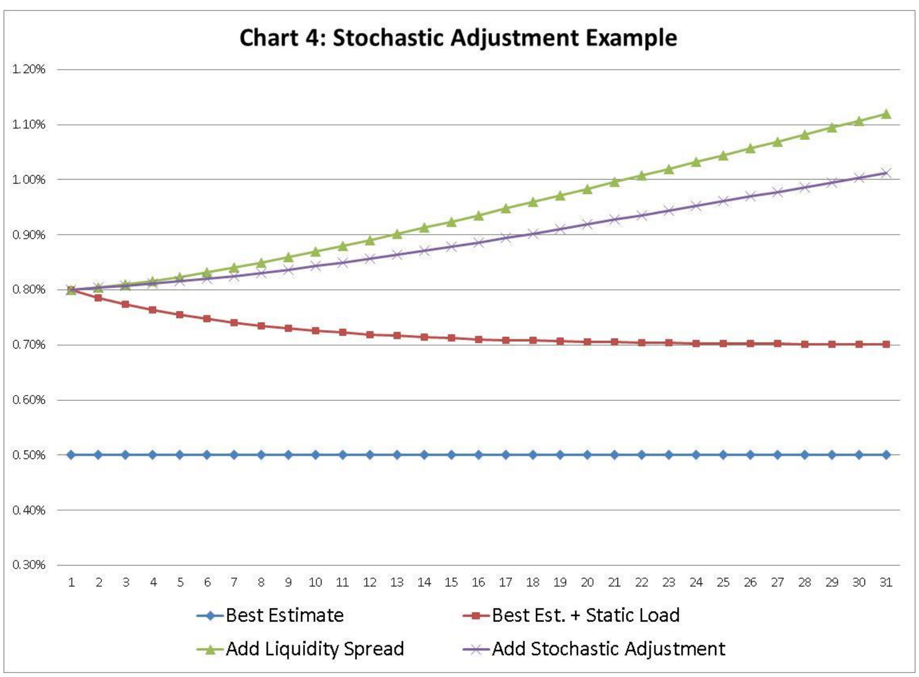

As expected, the long-term default spreads have not changed.

While cost of capital volatility may not have that much impact on model values, it does have significant risk management implications. Imagine, for simplicity, that the two-state model we have been developing is actually good enough to describe the real world of risky bonds. Assume also we have found a reasonable way to calibrate the model so we know the key parameters and state variables.

An investor holding a credit-risky bond is then subject to a number of risks:

1. Best-estimate credit default experience, portfolio risk diversified away by the law of large numbers.
2. Short-term credit crunch (correlated ratings downgrades in a more sophisticated model).
3. A change in the bond's perceived liquidity.
4. Fluctuations in the risk-free yield curve.
5. Fluctuations in market sentiment.

To the extent that the risky bond is being used to back a long-term insurance liability, we can ask which, if any, of these risks can be naturally hedged between the asset and liability. If an insurer is holding capital and risk margins for all of risks 2–5, and the bond's value reflects that, then we are entitled to do two things for risk which can be hedged:

- Take a capital offset for any risk that can be reduced by taking on a matching liability.
- Take credit for the cost of capital savings when putting a fair value on the liability.

The author's point of view is that it is appropriate to take credit for items 1, 3, 4 and 5 in the list above when valuing life insurance liabilities. We discuss each item in turn.

1. As noted near the beginning of this paper, the idea of best-estimate default experience is controversial and has been debated for many years. The author's point of view is that insurance company customers are taking some credit risk when buying a life insurance product and are entitled to some form of premium for taking that risk. The best-estimate default probability makes sense in this case.

2. To the extent a credit crunch occurs, and n years' worth of defaults and credit downgrades happened over night, this is not the policyholder's problem. The contagion spread should not be used when valuing an insurance liability.

3. If the market suddenly changes its point of view about bond liquidity (at a portfolio level), it makes sense for this risk to be passed through to the liability side by introducing a similar adjustment to liability values. A liquidity spread should therefore be included when valuing insurance liabilities. At a high level, this is consistent with Solvency II in Europe. As mentioned earlier, the details of the liquidity model developed here are different from the details of the current Solvency II model.[9]

4. Most insurers already assume that fluctuations in the risk-free yield curve can be hedged between assets and liabilities. The usual way to handle this issue is to hold capital for the net mismatch between assets and liabilities. When pricing liabilities, some companies assume a mismatch budget to take account of the fact that matching can never be perfect.

[9] The current Solvency II model for this issue takes observed credit default swap (CDS) spreads as an input. This is reasonable in that it takes observable market data into account. However, the resulting pattern of liquidity adjustments is flat for a certain period and then drops to zero after a fixed time horizon. This is not consistent with the risk insights that come out of the model described in this paper.

In principle, the idea of a mismatch budget could be expanded to cover the broader sense of "match" discussed in this paper.

5. The last issue requires more discussion because we have not stated, yet, what the risk associated with a change in market sentiment really is. Our point of view though is that this risk can be hedged between assets and liabilities.

At the beginning of this section, we stated that our assumption for the dynamics of the cost of capital rate was

$$d\pi = \kappa(\pi_\infty - \pi)dt + \xi\sqrt{\pi}dz.$$

We will now argue that the development above makes sense if this is the risk-neutral process. To see this, assume we start with a real-world process of the form

$$d\pi = \kappa'(\pi'_\infty - \pi)dt + \xi'\sqrt{\pi}dz.$$

A sudden change in market sentiment $\pi \to \pi + \Delta\pi$ causes the risky bond's value to change by

$$V(t,\beta,\pi+\Delta\pi,T) - V(t,\beta,\pi,T) \approx \Delta\pi\frac{\partial V}{\partial \pi} + \frac{(\Delta\pi)^2}{2}\frac{\partial^2 V}{\partial \pi^2}.$$

Now assume that the bond's owner is holding economic capital to cover this potential loss. The cost of capital concept then says the fundamental valuation equation should be

$$\begin{aligned}&\frac{\partial V}{\partial s} + [\pi - \beta\Delta\mu(1-R)]\frac{\partial V}{\partial \beta} + \kappa'(\pi'_\infty - \pi)\frac{\partial V}{\partial \pi} + \frac{\xi'^2\pi}{2}\frac{\partial^2 V}{\partial \pi^2} + \mu_0(R-1)V \\ &= rV + [(\pi c + \beta(t,s)\Delta\mu)(1-R)]V - \pi\left[\Delta\pi\frac{\partial V}{\partial \pi} + \frac{(\Delta\pi)^2}{2}\frac{\partial^2 V}{\partial \pi^2}\right].\end{aligned}$$

On rearranging, this becomes the risk-neutral equation studied earlier

$$\begin{aligned}&\frac{\partial V}{\partial s} + [\pi - \beta\Delta\mu(1-R)]\frac{\partial V}{\partial \beta} + \kappa(\pi_\infty - \pi)\frac{\partial V}{\partial \pi} + \frac{\xi^2\pi}{2}\frac{\partial^2 V}{\partial \pi^2} \\ &= [r + (\mu_0 + \pi c + \beta(t,s)\Delta\mu)(1-R)]V, \qquad V(T,\beta,\pi,T) = 1,\end{aligned}$$

provided the following relationships hold between risk-neutral and real-world parameters

$$\begin{aligned}\kappa &= \kappa' - \Delta\pi, \\ \pi_\infty &= \pi'_\infty\kappa' / (\kappa' - \Delta\pi), \\ \xi^2 &= \xi'^2 + \Delta\pi^2.\end{aligned}$$

These are all reasonable results. Risk adjustment reduces the rate of mean reversion, increases the long-term cost of capital assumption and also increases the assumed volatility.

Using the risk-neutral parameters to value insurance liabilities is then equivalent to assuming we take economic capital credit for the market sentiment risk on the liability side of the balance sheet. That this is reasonable is the current paper's main argument.[10]

The Multistate Model and Other Enhancements

Once a modeling process has started, there is never an end to the enhancements that could be made. The main point of this section is to show that some important issues ignored so far do not really change any of the important risk management conclusions derived in the context of the two-state model.

1. Stochastic risk-free rates. The model developed here could easily be incorporated into any affine model of the risk-free rate. Examples of affine models are the Hull-White model and its higher dimensional cousins such as G2++. One would have to consider how the risk-free rate and cost of capital rate processes are correlated.

2. Additional parameter risk. The model developed here assumed the parameters governing the cost of capital process were known with certainty. Since this is certainly not the case, one could argue that an important parameter such as π_∞ should get the same kind of treatment we gave to μ_0. This is certainly possible and should be considered if the issue is material to the problem at hand.

3. Income tax effects. To the extent one thinks of an income tax system as a risk-sharing arrangement, it may be appropriate to tax effect both the required economic capital and the associated risk margins described in this document.[11] The models in this document ignore income tax issues.

4. Recovery rates. The models discussed here assume a constant recovery rate. In the real world, recovery rates have a stochastic element and can vary by both issuer and the particular rank of the bond in a credit hierarchy.

5. Multistate bond ratings. There are a number of commercial bond-rating services that publish their analysis of the credit worthiness of individual bond issues. In most situations, these ratings are intended to indicate a given bond's probability of default over some relatively short time frame such a year or less.

Many authors have studied the process of bond transitions where a rating agency changes a bond's rating as new information about the bond issuer's credit worthiness becomes available. The simplest version of a model that takes this multistate issue into account is an annual ratings transition matrix T which we assume most readers of this paper are already familiar with.

The two-state model building process described earlier can then be generalized as follows:

[10] Note that if $\Delta\mu < 0$, then a flight-to-quality event where $\Delta\pi > 0$ could cause the bond's value to rise. This makes sense as long as we allow very high quality bonds to have negative liquidity capital requirements.

[11] The author's views on this subject were laid out in the article "An ERM Approach to Income Tax Risk," which appeared in the spring 2009 edition of the SOA newsletter *Risk Management*.

1. Best-estimate defaults are modeled by a transition intensity matrix M such that the annual transition probabilities are given by a matrix $T = \exp[M]$. The value function V gets generalized to a vector of values $\boldsymbol{V} = (V^1, \dots, V^n)$. Here V^i is the value of the risky bond given that it is in rating class i. Realistic examples of such matrices are given in tables 1 and 2 below.[12]

Table 1: Best Estimate Transition Matrix T

	AAA	AA	A	BBB	BB	B	C	D
AAA	91.95%	7.24%	0.77%	0.00%	0.03%	0.00%	0.00%	0.00%
AA	1.10%	90.99%	7.56%	0.26%	0.07%	0.01%	0.00%	0.00%
A	0.05%	2.40%	91.91%	5.01%	0.48%	0.12%	0.01%	0.02%
BBB	0.05%	0.24%	5.27%	88.52%	4.82%	0.78%	0.16%	0.17%
BB	0.01%	0.04%	0.50%	5.69%	84.99%	6.98%	0.54%	1.25%
B	0.01%	0.03%	0.13%	0.43%	6.63%	83.15%	3.15%	6.47%
C	0.00%	0.00%	0.00%	0.57%	1.71%	4.35%	68.21%	25.16%
D	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%

Table 2: Transition Intensity Matrix M

	AAA	AA	A	BBB	BB	B	C	D
AAA	-8.4%	7.9%	0.5%	0.0%	0.0%	0.0%	0.0%	0.0%
AA	1.2%	-9.6%	8.3%	0.1%	0.1%	0.0%	0.0%	0.0%
A	0.0%	2.6%	-8.7%	5.5%	0.4%	0.1%	0.0%	0.0%
BBB	0.1%	0.2%	5.8%	-12.5%	5.5%	0.7%	0.2%	0.1%
BB	0.0%	0.0%	0.4%	6.6%	-16.8%	8.3%	0.5%	1.0%
B	0.0%	0.0%	0.1%	0.2%	7.9%	-18.9%	4.2%	6.5%
C	0.0%	0.0%	0.0%	0.6%	2.0%	5.7%	-38.4%	30.1%
D	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

The state $i = 'D'$ is usually assumed to mean actual default. A typical simplifying assumption is that $V^D(t,T) = Re^{-r(T-t)}$ where R is the recovery rate.

The best-estimate equation of value is then

$$\frac{d\mathbf{V}}{dt} + M\boldsymbol{V} = r\boldsymbol{V}.$$

Based on this equation, one can develop time-dependent best-estimate forward default rates that take into account a scenario where a risky bond gradually declines from a high credit rating to default.

Putting $r = 0$ for convenience, the solution to the equation above can be written as

$$V^i(t,t) = (1,..,1,R)'$$
$$V^i(t,T+1) = \sum_j T^i_j \, V^j(t,T).$$

[12] This particular matrix was adapted by the author from a study by Moody's covering the decade of the 1990s. All we claim here is that it is broadly representative of what a realistic transition matrix looks like.

This could be one way to develop a best-estimate default assumption $\mu_0 = \mu_0(t)$ that reflects the multistate environment. Table 3 below shows the forward default rates implied by the assumed best-estimate transition matrix. These rates were calculated using

$$d_T^j(t) = \ln[\frac{V^j(t,T-1)}{V^j(t,T)}].$$

Table 3: Best Estimate Forward Default Rates

Time	AAA	AA	A	BBB	BB	B	C
1	0.00%	0.00%	0.01%	0.08%	0.63%	3.29%	13.44%
2	0.00%	0.00%	0.02%	0.15%	0.84%	3.29%	10.53%
3	0.00%	0.00%	0.04%	0.21%	1.00%	3.18%	8.08%
4	0.00%	0.01%	0.05%	0.27%	1.12%	3.02%	6.13%
5	0.00%	0.01%	0.07%	0.32%	1.20%	2.82%	4.61%
10	0.02%	0.05%	0.17%	0.51%	1.26%	1.87%	1.24%
15	0.04%	0.11%	0.27%	0.59%	1.09%	1.23%	0.52%
20	0.08%	0.18%	0.34%	0.60%	0.90%	0.86%	0.31%
25	0.13%	0.23%	0.38%	0.58%	0.74%	0.63%	0.22%
30	0.18%	0.28%	0.41%	0.54%	0.61%	0.48%	0.16%

Note that if we want to assume insurance liabilities have a AA rating, we are talking about a fairly small best-estimate default spread. The behavior of the C bond's forward default rates can be explained by noting that such a bond will only survive to 30 years by migrating back to a higher rating class at some future point in time.

2. Contagion risk can still be modeled by assuming a capital requirement equal to, approximately, n years' worth of best-estimate credit transitions overnight. The capital requirement would be, if the eigen values of T are small enough,

$$\boldsymbol{EC} = (I - T^n)\boldsymbol{V} \approx -nM\boldsymbol{V}.$$

The new valuation equation becomes

$$\frac{d\mathbf{V}}{dt} + (1 + n\pi)M\boldsymbol{V} = r\boldsymbol{V}.$$ [13]

This equation can be solved for a set of risk-loaded forward default rates by using the same approach as was used for the best-estimate case except that we now use a risk-loaded transition matrix $\hat{T} = \exp[M(1 + n\pi)]$. Table 4 shows the impact on default rates

[13] At this point, we are close to the model of Jarrow, Lando and Turnbull referenced in footnote 4. Their model could be understood as a version of the current model where the cost of capital is a vector $\boldsymbol{\pi}(t)$, which varies by rating class and with time. They then use the time dependence of $\boldsymbol{\pi}$ to calibrate the model to observed yield curves by rating class. Their derivation does not use cost of capital concepts. The author's main critique of this approach is that the concept of, say, a AA yield curve, does not really exist since many similar bonds, with the same current rating, can have different prices due to liquidity considerations.

of assuming $n = 4$ and $\pi = .10$. Table 4 shows the difference between the forward default rates consistent with the loaded model and the forward default rate in Table 3.

Table 4: Contagion Loaded Forward Default Spreads

	AAA	AA	A	BBB	BB	B	C
1	0.00%	0.00%	0.01%	0.05%	0.31%	1.32%	4.53%
2	0.00%	0.00%	0.02%	0.11%	0.48%	1.23%	2.10%
3	0.00%	0.01%	0.04%	0.17%	0.56%	1.04%	0.52%
4	0.00%	0.01%	0.06%	0.21%	0.59%	0.82%	-0.34%
5	0.01%	0.02%	0.08%	0.24%	0.58%	0.61%	-0.70%
10	0.03%	0.08%	0.17%	0.30%	0.33%	0.03%	-0.39%
15	0.08%	0.15%	0.22%	0.25%	0.13%	-0.10%	-0.11%
20	0.14%	0.19%	0.22%	0.19%	0.03%	-0.10%	-0.05%
25	0.18%	0.22%	0.21%	0.14%	0.00%	-0.09%	-0.03%
30	0.21%	0.22%	0.19%	0.11%	-0.01%	-0.07%	-0.02%

According to the author's point of view, this is the component of the forward default rate that should not be used to discount insurance liabilities. It is very similar to the best-estimate forward default rate if $n = 4$.

The fact that some of the risk-adjusted forward default rates can be lower is not an error. This table simply shows that adding a contagion loading to the best-estimate default matrix simply exaggerates the survival issue we saw in Table 3.

We can also use these calculations to derive economic required capital factors for contagion risk. Table 5 below shows the factors that should apply by rating class and cash flow maturity

$$c_{T-t}^{i} = \frac{n \sum_j M_j^i V^j(t,T)}{V^i(t,T)}.$$

Table 5: Contagion Capital%

	AAA	AA	A	BBB	BB	B	C
1	0.0%	0.0%	0.1%	0.6%	3.3%	13.2%	43.1%
2	0.0%	0.0%	0.2%	0.9%	4.2%	12.6%	29.7%
3	0.0%	0.0%	0.3%	1.2%	4.7%	11.5%	20.0%
4	0.0%	0.1%	0.4%	1.5%	5.0%	10.4%	13.5%
5	0.0%	0.1%	0.5%	1.7%	5.1%	9.3%	9.2%
10	0.2%	0.4%	1.0%	2.3%	4.4%	5.1%	2.2%
15	0.4%	0.8%	1.4%	2.4%	3.4%	3.1%	1.1%
20	0.6%	1.1%	1.6%	2.2%	2.6%	2.1%	0.7%
25	0.9%	1.3%	1.7%	2.0%	2.0%	1.5%	0.5%
30	1.1%	1.5%	1.7%	1.8%	1.7%	1.2%	0.4%
OSFI	0.25%	0.5%	1.0%	2.0%	4.0%	8.0%	16.0%

The bottom row of this table shows the current Canadian[14] regulatory capital requirements by bond-rating class. They would appear to be reasonable for a mix of bond maturities if contagion risk were the only issue on the agenda.

3. The idea of parameter or liquidity risk can be incorporated by assuming a potential shock $\Delta M = \varphi M$. Assuming the shock is proportional to the best estimate is merely the simplest place to start. More complex models are possible.

   Given this particular approach to liquidity risk, the valuation equation becomes, for $\boldsymbol{V} = \boldsymbol{V}(t, \beta, T)$,

$$\frac{\partial \mathbf{V}}{\partial t} + \pi \frac{\partial \boldsymbol{V}}{\partial \beta} + (1 + n\pi + \beta\varphi) M \boldsymbol{V} = r\boldsymbol{V}.$$[15]

   This equation can be solved by assuming we know the left eigen vectors of the matrix M. This is a matrix L such that

$$LM = -DM,$$

   i.e., each row of L is an eigen row of M and D is a diagonal matrix of the form

$$D = \begin{pmatrix} \mu_1 & 0 & 0 \\ 0 & \dots & 0 \\ 0 & 0 & \mu_n \end{pmatrix}.$$

   For this particular example, the eight diagonal eigen values are, in increasing order

0.0%	1.0%	5.9%	9.0%	13.3%	18.0%	26.6%	39.6%

   and the eigen row matrix is given by

			Eigen Row Matrix *L*					
	AAA	AA	A	BBB	BB	B	C	D
AAA	0.125	0.125	0.125	0.125	0.125	0.125	0.125	0.125
AA	0.259	0.229	0.197	0.154	0.096	0.050	0.015	-
A	1.184	0.379	(0.001)	(0.170)	(0.210)	(0.141)	(0.041)	-
BBB	0.898	(0.051)	(0.135)	0.013	0.129	0.115	0.031	-
BB	3.118	(1.930)	0.405	0.593	(0.390)	(0.636)	(0.160)	-
B	1.273	(1.634)	1.480	(1.737)	0.023	1.290	0.305	-
C	(0.045)	0.131	(0.254)	0.873	(2.209)	1.907	0.597	-
D	0.000	(0.001)	0.003	(0.017)	0.076	(0.272)	1.211	-

[14] These factors are very similar to the ones used by U.S. regulators. OSFI stands for Office of the Super-Intendant of Financial Institutions, which is the Canadian federal government regulator for both banks and insurers.
[15] This particular model assumes $d\beta = \pi dt$ without any adjustment. This is a simplification that is usually immaterial in most practical examples. One alternative is to assume the eigen vectors of M are known with certainty and then apply the two-state model to each eigen value of M sepateately. The next level of complexity is to allow for uncertainty in the eigen vectors of M. Such considerations are beyond the scope of this paper.

If we now define a new vector by $\boldsymbol{W} = L\boldsymbol{V}$, then the fundamental valuation equation becomes

$$\frac{\partial \boldsymbol{W}}{\partial t} + \pi \frac{\partial \boldsymbol{W}}{\partial \beta} = \big(r + D(1 + n\pi + \beta\varphi)\big)\boldsymbol{W}.$$

This equation is subject to the boundary condition $\boldsymbol{W}(T,T) = L\boldsymbol{V}(T,T)$. Based on the particular matrix chosen[16] above, we find $\boldsymbol{W}(T,T) = (.9375, 1, 1, \ldots, 1)'$.

Because the matrix D is diagonal, we only need to consider the diagonal components, which are very similar to the two-state model described earlier. Each diagonal component is of the form

$$\frac{dW^{\boldsymbol{i}}}{dt} + \pi \frac{\partial W^{\boldsymbol{i}}}{\partial \beta} = \big(r + \mu_i(1 + n\pi + \beta\varphi)\big)W^{\boldsymbol{i}}.$$

If the cost of capital rate π is a constant, then $\beta(t,s) = \pi(s-t)$ and the closed form solution for $W^{\boldsymbol{i}}(t,T)$ is

$$W^{\boldsymbol{i}}(t,T) = W^{\boldsymbol{i}}(T,T)\exp[-\int_t^T \big(r + \mu_i(1 + \pi(n + (s-t)\varphi)\big)ds],$$

$$= W^{\boldsymbol{i}}(T,T)\exp[-\left[r(T-t) + \mu_i(T-t)\left(1 + \pi\left(n + \varphi\frac{(T-t)}{2}\right)\right)\right].$$

We can recover the original vector value from $\boldsymbol{V} = L^{-1}\boldsymbol{W}$ and then compute forward default spreads as before. Table 6 below shows the results that follow from the assumptions stated earlier along with $\varphi = .25$, i.e., $\Delta M = .25\,M$.

Table 6: Parameter/Liquidity Risk Spreads $\phi = 25\%$

	AAA	AA	A	BBB	BB	B	C
1	0.00%	0.00%	0.00%	0.00%	0.01%	0.04%	0.13%
2	0.00%	0.00%	0.00%	0.01%	0.04%	0.12%	0.23%
3	0.00%	0.00%	0.01%	0.02%	0.08%	0.17%	0.19%
4	0.00%	0.00%	0.01%	0.04%	0.12%	0.19%	0.10%
5	0.00%	0.01%	0.02%	0.06%	0.15%	0.20%	0.03%
10	0.02%	0.05%	0.09%	0.15%	0.20%	0.13%	0.00%
15	0.07%	0.12%	0.17%	0.20%	0.16%	0.08%	0.02%
20	0.16%	0.21%	0.23%	0.21%	0.13%	0.05%	0.02%
25	0.26%	0.28%	0.27%	0.22%	0.12%	0.05%	0.01%
30	0.34%	0.33%	0.29%	0.22%	0.12%	0.05%	0.02%

Interestingly, a 25 percent parameter shock can have more impact on high quality bond spreads than it does on low quality bond spreads as the bonds get longer.

[16] The rows of the matrix L can be permuted and scaled arbitrarily without affecting the end results.

4. Allowing the cost of capital rate to be stochastic still does not change very much.

   The model becomes

$$\frac{d\mathbf{V}}{dt} + \kappa(\pi_\infty - \pi)\frac{\partial \boldsymbol{V}}{\partial \pi} + \frac{\xi^2 \pi}{2}\frac{\partial^2 \boldsymbol{V}}{\partial \pi^2} + \pi\frac{\partial \boldsymbol{V}}{\partial \beta} + (1 + n\pi + \beta\varphi)M\boldsymbol{V} = r\boldsymbol{V}.$$

   As before, we let $\boldsymbol{W} = L\boldsymbol{V}$ and $\boldsymbol{W}(T,T) = (.9375,1,1,\dots,1)'$. The diagonal component equations are then

$$\frac{dW^{\boldsymbol{i}}}{dt} + \pi\frac{\partial W^{\boldsymbol{i}}}{\partial \beta} + \kappa(\pi_\infty - \pi)\frac{\partial W^{\boldsymbol{i}}}{\partial \pi} + \frac{\xi^2 \pi}{2}\frac{\partial^2 W^{\boldsymbol{i}}}{\partial \pi^2} = \big(r + \mu_i(1 + n\pi + \beta\varphi)\big)W^{\boldsymbol{i}}.$$

   This is slightly different from the equation studied in the two-state model section but the main technical conclusions do not change. We can still solve each equation by assuming a solution of the form

$$W^{\boldsymbol{i}}(t,T) = W^{i}(T,T)\exp[A^{i}(t,T) + \beta B^{i}(t,T) + \pi P^{i}(t,T)]$$

Model calibration

While we will not present an actual model calibration in this paper, we can indicate the steps one could go through to do this.

1. Choose a universe of credit-risky instruments that can be valued using a model of the kind described above. This could include both vanilla bonds and credit default swaps. Index these instruments with lower case roman indices i.

2. Obtain real market values for these instruments at a number of historical time points t_A. Let V_{iA} denote the matrix of raw data. For out-of-sample verification, collect the same data for a different set of time points $t_{A'}$. Call this data matrix $V_{iA'}$.

3. Now divide the model's input parameters into three groups:

   - Global parameters that should be constant for all instruments and for all points in time. These are $M, \kappa, \pi_\infty, \xi$.
   - Parameters that vary by time only. These are the risk-free yield curve (e.g., the swap curve) and the point-in-time cost of capital rate $\pi(t_A) = \pi_A$.
   - Parameters that vary by instrument only. This would include the bond's rating and its own unique liquidity parameter φ_i and recovery rate R_i.

4. Now go through a traditional fitting exercise where we solve for the parameters that minimize a reasonable measure of fit. One place to start is to consider a fit measure of the form

$$F = \sum_{A,i}[V_{iA} - T(i, t_A, M, \kappa, \pi_\infty, \xi, \pi_A, \varphi_i, R_i)]^2/E_i.$$

Here $T(i, t_A, M, \kappa, \pi_\infty, \xi, \pi_A, \varphi_i, R_i)$ is the theoretical value of the instrument given by the model and E_i is a factor designed to give an appropriate weight to the particular instrument. One choice for E_i might be the total outstanding face amount.

The final step of the fitting process is to see how well the model does on the out-of-sample data. In this test, we fix the global and bond-specific parameters but allow the time dependent variables to move. If it does well, we are done. If not, we have to go back and revise the model until it is deemed to be working. This is clearly a subjective step in the process.

One test of the model is to see whether the "observed" cost of capital rates $\pi_A, \pi_{A'}$ are consistent with a process of the form $d\pi = \kappa'(\pi'_\infty - \pi)dt + \xi'\sqrt{\pi}dz$.

5. Once we have a set of credible parameters, we still have to decide how the model should, be applied to value insurance liabilities. This amounts to assigning both a rating (e.g., AA) and a liquidity factor φ to the insurance enterprise. A reasonable starting point might be to look at an average liquidity factor for all instruments in the target rating class.

Conclusion

The main risk management conclusion of this paper is that there are several credit-related risks which can be naturally hedged between the asset and liability side of a life insurer's balance sheet. In particular, we have argued that the risks associated with credit loss parameter risk are essentially the same as liquidity risk and that the issues associated with market sentiment can also be passed through to the liability side of the balance sheet. If the reader accepts this point of view, then it makes sense to include those components of a credit-risky bond yield curve into the valuation of life insurance liabilities. Taking this approach to liability valuation would go a long way toward resolving the issues raised in the introduction to this paper.

To the extent a life insurer matches its long assets with its long liabilities, the only remaining risk it needs to hold capital for is the contagion risk element (credit crunch). At a high level, this is already where most regulatory capital models are.[17]

A second conclusion is that institutions that do not have long liabilities (such as banks) should be holding more regulatory capital for risky bonds because there is no capital offset with their liabilities.

[17] To be fair, most of these regulatory models were developed for a book-value-accounting world where the risks associated with short-term market fluctuations are swept under the rug by accounting tricks such as book-value accounting. If you take the risk management point of view, as we do here, that this is inappropriate, then the capital models need to be adjusted. This may include allowing very high quality bonds to have negative capital requirements.