Kamakura Corporation
2222 Kalakaua Avenue
Suite 1400
Honolulu, Hawaii 96815, USA
phone 1 808 791 9888
facsimile 1 808 791 9898
www.kamakuraco.com

Advanced Financial Risk Management:

Tools and Techniques for Integrated Credit Risk, Market Risk, Liquidity Risk and Interest Rate Risk Management

April 2012

Donald R. van Deventer
Kenji Imai
Mark Mesler

Part 4: Risk Management Applications, Instrument by Instrument

Chapter 35:
Valuing insurance policies and pension obligations

Our final chapter in individual instrument valuation and simulation is focused on insurance policies and pension obligations, since insurance companies and pension funds make up a substantial proportion of the world's financial institutions. To reiterate, our objective in reviewing valuation and simulation methodologies is to incorporate the insurance liabilities and pension liabilities in the Jarrow-Merton put option on the financial institution's assets and liabilities. We want to use a consistent methodology to calculate this comprehensive measure of integrated interest rate risk, market risk, liquidity risk, foreign exchange risk, and credit risk. This is particularly interesting to discuss for insurance and pension obligations because the links with what we have already covered are so strong and yet so few members of the financial community would acknowledge that that is the case. The authors hope that this chapter assists in narrowing this perception gap.

We start first with life insurance policies and then move on to pension obligations and property and casualty insurance.

Life Insurance: Mortality Rates versus Default Probabilities

The analogy between the mortality rate on a life insurance policy and the default probability on a bond of a particular issuer is very strong. Figure 1 below shows annual mortality rates for a 50 year old male and 50 year old female taken from the U.S. Society of Actuaries RP-2000 mortality tables:

Figure 1 [Chapter35-Figure01.tif]

The mortality rate graph is surprisingly similar in shape and meaning to the evolution of the one year default probability for Tokyo Electric Power prior to its nationalization by the Japanese government on May 8, 2012:

Figure 2 [Chapter35-Figure02.tif]

Like default or bankruptcy, mortality on a life insurance policy is a zero/one kind of event. The fundamental approach to measuring mortality risk was proposed by D.R. Cox in 1972.[1] The 'Cox process' that he developed is in fact the same 'Cox process' used in reduced form modeling in the Jarrow [1999, 2001] and Duffie and Singleton [1999] credit models that we discussed in Chapter 16. The mathematics is identical. The term structure of mortality rates shown above is very directly related to the term structure of default probabilities that we can observe for Tokyo Electric Power Corporation on May 8, 2012 shown in Figure 3:

Figure 3 [Chapter35-Figure03.tif]

This parallel is so strong and so obvious to financial academics that Robert Jarrow described the relationship between credit default swaps, which we discussed in Chapter 20, and life insurance in this way in his paper "Problems with Using CDS to Infer Default Probabilities" (Cornell University and Kamakura Working Paper, January 25, 2012, page 1):

> "It is commonly believed that credit default swap (CDS) implied probabilities provide unbiased estimates of a corporation's or sovereign's actual default probabilities or credit risk. The purpose of this paper is to show why this common belief is false using an intuitive analogy between CDS and life insurance premiums. For life insurance, the difficulties of using insurance premiums to imply the mortality probability are crystal clear. The difficulties are unraveling the impact of risk premium, counterparty risk, market frictions, and strategic trading. These difficulties are well known in the academic literature (references are provided below). Of course, all is not lost because there is no need to use these implied probabilities. Readily available actuarial based mortality tables are used to determine mortality probabilities. And insurance premium implied mortality probabilities are not used for risk management nor in regulation."

The other strong parallel between life insurance and credit risk, as Professor Jarrow's quote implies, is the progressive use of logistic regression to determine the probability of the event (bankruptcy as in Chapter 16 or mortality in this Chapter) with more precision than we would get from a simple table (a transition matrix with default probability by ratings grade or a mortality table like the graph

[1] 'Regression Models and Life Tables', Journal of the Royal Statistical Society, B34, 187-220, 1972.

in Figure 1). A recent paper by Guizhou Hu[2] shows how mortality rate prediction can be improved by incorporating explanatory variables that are much more detailed than those in mortality tables (age, male or female, smoker or non-smoker); among the variables he mentions are the following:

- Systolic blood pressure;
- Total cholesterol;
- Body mass index;
- Smoker or non-smoker;
- Albumin;
- Physical activity — i.e. light, moderate or heavy;
- Income

All of these input variables are specific to the insured, just like the accounting and stock price inputs to the Jarrow-Chava model of Chapter 16 are specific to the firm being modeled.

Cyclicality in Default Probabilities and Mortality Rates

Logistic regression can also confirm the importance of cyclicality in mortality just as it does (as required by the Basel Capital Accords) for default. Any Japanese banker can tell you that the mortality rate of his co-workers is much different when the Nikkei stock index is below 10,000 yen than it was when it is nearly 39,000 yen in December 1989. Similarly, just as there is seasonality in the default of small businesses, there is seasonality in mortality that is logical and predictable (due to factors like weather, stresses of the school calendar year, low stress of summer vacations, and so on).

Valuing Life Insurance Policies

Life insurance policies are closely related to digital default swaps that we discussed in Chapter 20. The protection buyer on the digital default swap called a life insurance policy is (usually) the insured. The insured is also the reference name on the digital life insurance policy. The counterparty on the digital life insurance policy is the life insurance company. The amount paid on the digital life insurance policy in the event of “default” is the policy amount. The only significant difference is the maturity—term life insurance policies, which have an explicit maturity, are almost literally identical to digital default swaps. A ‘whole life’ policy runs until mortality or until cancelled. Ignoring the cancellation provisions, it’s like a infinite maturity digital default swap except that we know the term structure of mortality/default probabilities approaches 100% in time. In fact, the graph for age 120 in Figure 1 would have been 100%, the amount specified in the Society of Actuaries RP-2000 tables for that age.

What about cancellation privileges? This is a real option of the insured that can be exercised rationally or irrationally as discussed in the prepayment context in Chapters 27 to 29. It is rational to cancel when (a) the mortality risk of the insured becomes lower than that implied in the current policy, allowing the insured to reinsure at a lower rate, or (b) when the insured becomes so wealthy that they no longer need insurance (they save money by not paying the insurance company costs embedded in the policy and effectively ‘self insure’. ‘Irrational’ cancellation (like irrational

[2] ‘Mortality Assessment Technology: A New Tool for Life Insurance Underwriting,’ undated memorandum, BioSignia, Inc., Durham, North Carolina.

prepayment) is usually for a good reason unknown to the insurance company. The financial condition of the insured, for example, may have deteriorated so much that they can't afford the premiums, even though they need the insurance.

What about investment options on a whole life policy? We know from finance theory that an option to switch from one instrument at its market price to another instrument at its market price has zero value, since there is no benefit to the options holder (if we ignore transactions costs). Many of the investment options in a life insurance policy have this nature. To the extent there is a real benefit to the holder of the insurance policy from the investment options offered, we can value them using the techniques in other chapters in this book using the "default adjusted" approaches from Chapter 19 (corporate bonds) or Chapter 20 (default swaps or CDOs).

One of the key benefits of life insurance to the buyer is the ability to defer taxes on the investment income of a whole life insurance policy. This tax deferral is of course a real benefit to the buyer of the insurance policy, relative to making the identical investments in the name of the insured. The value of this benefit to the buyer of the insurance policy depends on the state of the tax code now and in the future, the source and variability of income now and in the future, and the default probability of the insurance company. Clearly, valuing these tax benefits with precision is a challenge, and one that we're happy to defer to another time.

Pension Obligations

Defined benefit pension obligations have a very interesting mix of terms that combine part of the life insurance problem with more traditional financial instruments. The U.S. Social Security System is exactly such an obligation. We encourage readers to explore the excellent paper on this topic by Dartmouth College Prof. Samir Soneji and Prof. Gary King of Harvard: "Statistical Security for Social Security," June 30, 2011, Harvard University Working Paper. One of the advantages of a pension fund in an economic sense is the on-going stream of current payments for future pension benefits from both employees and their employers. The present value of this cash inflow from any one employee, say Mr. Jones, depends on three factors: the mortality probability of Mr Jones (he could be unlucky enough to pass away before leaving employment and receiving pension benefits), the probability that Mr Jones will change jobs and end up contributing to a different pension fund, and the probability of default of Mr Jones' employer. The liabilities of a pension fund, the future payments to the beneficiaries of the fund, depend on the mortality rate of the pension beneficiary, just like a life insurance policy. The difference is that the payments are made in the opposite circumstance of a life insurance policy—they are paid while the beneficiary is still living, and a "default event" terminates this stream of cash flows just like a mortgage default terminates payments on a mortgage held by a bank. In this sense, the future streams of pension payments are each a 'digital default' swap of sorts again, except the key probability used in valuation is the probability of no mortality prior to the payment, not the probability of bankruptcy/mortality relevant to valuing a life insurance policy.

For defined benefit pension plans, which is our focus in this section, there are normally a complex set of rules or circumstances which determine when and how pension benefits are increased over time. Similarly, even with government pension insurance, there is the probability that the pension fund defaults on its obligations to the pension beneficiaries. Both of these complexities can be dealt with using a combination of the life insurance analysis above and the techniques of other chapters.

In particular, the increased use of the logistic regression technique and reduced form modeling approach allows for comprehensive credit-adjusted risk management of both the assets and liabilities of a pension fund.

One of the greatest fallacies in the world of legacy pension analytics is the assumption of a "portfolio return" on the assets of the pension fund. This assumed return has huge political implications, in the case of a government sponsored pension fund, if an analysis of the fund uses this assumed return in the calculation of whether or not the pension fund is "underfunded." A more rational analysis marks to market both the assets and the liabilities of the pension fund using the techniques of this book. Discounting the mortality adjusted cash flows of the liabilities at the risk-free rate is the liability risk-neutral value on the assumption that the liabilities must and will be paid. Discounting the liabilities at any higher discount rate implies it is economically and politically acceptable for the pension fund to default on its obligations to pensioners. Similarly, any form of valuation on the asset side which allows assets to be "marked" at a higher value than their market value (like the "assumed return" above) is a thinly disguised ploy to overstate the financial health of the pension fund. A recent study ("Going For Broke: Reforming California's Public Employee Pension Systems," April 2010) by a team from the Stanford Institute for Economic Policy Research made this particularly clear with respect to the pension funds overseen by the State of California.

Property and Casualty Insurance

Property and casualty insurance differs from life insurance in two dimensions: the term of the insurance contract is typically shorter and the events being insured are quite diverse, ranging from auto insurance to catastrophic weather-related events. Aside from events that are acts of nature, like weather and earthquakes, logistic regression again plays a powerful role in both pricing and risk management of liabilities on a property and casualty insurance contract. The price of an auto insurance policy depends on factors such as the following:

- The age of the driver;
- The sex of the driver (in some jurisdictions);
- The health of the driver (if the data is available);
- The income of the driver (again, if the data is available);
- The driving history of the driver;
- The amount of driving to be done (and the possibility that the true amount is not disclosed);
- The type of car being insured (i.e. a bus or a sports car);
- The value of the car being insured;
- The place in which the car will be stored (a deserted urban parking lot in a high crime area or a locked garage in a gated community);
- The place in which the car will be driven (i.e. downtown New York City or rural France).

In addition, it is an empirical question whether macroeconomic factors have an impact on both the probability of an insured event and the payout on the occurrence of the event. Clearly, we know from the statistics on new auto sales that new sales of cars decline in bad times so the average age of all cars being driven increases during recessions. This in turn would be expected to affect the probability of a claim and the payout on the claim. For this reason, depending on the nature of the insured event, it is often the case that macro factors drive the probability of an insurance payout just like they affect the probability of default in the reduced form credit models of Chapter 16. When this is the case, it is essential that these macro factors are incorporated in policy pricing and total balance sheet risk management or the risk of providing the policy can be dramatically underestimated. If these macro factors are relevant, they will increase the correlation in the occurrence of insured events and the fat tails of the loss distribution.

The Jarrow-Merton Put Option

In many insurance companies and pension funds, there is a sharp division between the investment or asset side of the organization and the insurance or liability side of the organization. The investment function is dominated by finance experts and the liability side is dominated by insurance experts and actuaries. This contrasts sharply with risk management at banks and securities firms where asset and liability management on a joint basis has been standard for 40 years.

Looking at the risk on only one side of the balance sheet at a time is dangerous because common macro factors drive the risk on both sides of the balance sheet. If you are looking at only half the picture, it is impossible for management, the board of directors and regulators to have an accurate view of total risk. How can we value the Jarrow-Merton put option on the assets and liabilities of the firm as a comprehensive integrated measure of interest rate and credit risk if half of the information is missing or based on a completely different and totally inconsistent ("We have assumed that the assets of the pension fund will return a compounded rate of 12% annually for the next 30 years") set of assumptions?

Perhaps the most tragic recent example of the damage caused by these Chinese Walls in insurance is the joint impact of the 2011 Sendai tsunami and Fukushima nuclear disaster in Japan. Japanese life insurance companies were very heavily invested in the bonds of Tokyo Electric Power Company at the same time mortality rates and property and casualty damages rates sky-rocketed in the aftermath of the tsunami. If the Japanese government had not rescued Tokyo Electric Power, the default probabilities for TEPCO shown above make it clear that a "double whammy" of wrong way risk on the asset and liability sides would have occurred.

Fortunately, as we have seen in this chapter, the mathematics of credit risk models is taken directly from insurance expert Cox and there is no conceptual barrier to integrated risk management in the insurance industry. There is literally no difference in the mathematical approach, so there are only emotional and political reasons for the long-standing divide between the actuaries and the investment side of the organization when it comes to risk management. The best practice insurance companies are solving this problem by moving actuaries to the investment side of the organization and investment managers to the insurance side of the organization. Integrated risk management teams that have both actuaries and finance experts are being established and a common view of integrated risk is widely recognized as necessary and desirable. Insurance firms are increasingly bringing in chief risk officers from the banking business who are determined to break down the legacy Chinese Walls between the actuaries and the rest of the organization. We expect that the traditional 'walls' down the middle of insurance companies and pension funds will break down rapidly in the years ahead. That, in the end, is good for all concerned.

We now step back and do exactly that, ourselves. We look at the entire integrated risk of the entire organization in the next five chapters, now that we have completed our transaction level survey of best practice valuation techniques.