Markovian Annuities and Insurances

Thomas Struppeck, FCAS, MAAA, ASA

Markovian Annuities and Insurances

Abstract

Traditionally, property and casualty products have been thought of as "short duration contracts", while life insurance products have been thought of as "long duration contracts". Many modern property and casualty products have risk profiles and cash flow characteristics that are more akin to life insurance than to traditional property and casualty lines. In this paper, using bond insurance as a primary example, we show how such products can be priced and reserved using techniques from the capital markets and from life insurance.

The "life reserves" held by life companies are essentially premium deficiency reserves in that they are required not to pay losses that have occurred, but rather to make up the shortfall in future premium collections. Since bond insurance is so similar to life insurance, it is no surprise that the appropriate reserves for bond insurers are also premium deficiency reserves.

Introduction

Many insurance pricing and reserving problems can be phrased as questions about the value of a contingent annuity. This annuity might represent an anticipated stream of premium payments or a stream of loss payments. Typically, the stream of payments will terminate when a certain event occurs. This paper describes how to price and reserve for what we will call "Markovian annuities" and insurance products associated with them. Our main example will be bond insurance.

Markovian annuities are in some sense generalizations of level premium life insurance and, also for example, catastrophe reinsurance from the property and casualty side. As we will see, traditional life insurance pricing and reserving techniques suggest methods for valuing certain property and casualty reserves. These generalized methods in turn may be useful to life actuaries evaluating business priced with select and ultimate tables.

The paper is broken into thirteen sections. The first is this introduction, followed by a section describing what we will call the "risk-neutral world". Then Markov processes are discussed and a simple example is given. We then digress a little bit to discuss rating agencies. We then return to the topic of perpetuities and tie the first part of the paper together by introducing the notion of bond insurance.

We begin the second half of the paper by seeing how insurance can be used to turn risky assets into risky liabilities on an insurer's balance sheet. Valuing these liabilities is one of the central topics of this paper. To accomplish this, we first review the notion of a replicating portfolio, an idea that has its origins in the capital markets. Having built this machinery, we are finally ready to analyze bonds. The next two sections contain some remarks on accounting considerations and a detailed example. Finally, we make some concluding remarks and have a short bibliography.

The author would like to thank the Committee on Reserves for sponsoring the call, and to thank in particular the colleagues who read early drafts of this paper for their many helpful comments.

Perpetuities and the Risk-Neutral World

For ease of exposition, we will make several simplifying assumptions. None of these is necessary for what follows, but relaxing them introduces unnecessary complications that might mask what is really going on. Here and throughout the paper we will assume:

1) A flat, constant yield curve with an interest rate of 8%.
2) An unlimited supply of risk neutral investors willing to purchase or sell any stream of future cash flows, contingent or certain, at its expected present value.

3) No reporting lag.

4) Losses are paid at the end of the year.

5) Finally, assume that all losses occur at the end of the year.

Initially at least, we will examine perpetuities and contingent perpetuities. By a **contingent perpetuity** we mean a stream of payments of \$1.00 at the end of each year that terminates when a certain event occurs. The occurrence of this event we will call a **default**. A contingent perpetuity that cannot default we will call a **risk-free perpetuity**.

Contingent perpetuities are quite general; for example a life annuity payable to a 40-year old could be considered as a contingent perpetuity, the terminating event in this case being the annuitant's death.

As a first example, let's compute the market price in our risk-neutral world of a risk-free perpetuity. Denote by $\mathbf{a}_{rf}$ the market price of this perpetuity in our risk neutral world and let $v = 1/(1+i) = 1/1.08 = .926$ be the discount rate. We have:

$$\mathbf{a}_{rf} = v\,(1+ \mathbf{a}_{rf}).$$

That is, an investor is ambivalent between having the perpetuity today and having the present value of a portfolio consisting of the dollar that the perpetuity will pay in one year and another perpetuity one year from today. Equivalently, in the language of interest theory, an investor is ambivalent between a perpetuity-immediate and the present value of a perpetuity-due.

Solving, we obtain the familiar:

$\mathbf{a_{rf}} = v/(1 - v) = 1/i = 1/0.08 = 12.5$

Remark: If we had been evaluating an annuity that had a fixed number of payments, the annuity that we have after one year would not be identical to our initial annuity (it would have one less year remaining). In essence, perpetuities do not age, and this fact makes them easier to handle. This is an example where evaluating an infinite sum is easier than evaluating the corresponding finite sum.

Next we will evaluate a contingent perpetuity with a terminating event, but first we need a definition.

Markov Processes

A **(discrete) Markov process** is a stochastic process where the state at time t+1 depends only on the state at time t. Formally, it is a triple ($\mathbf{S}$, $\mathbf{p}$, $\mathbf{s_0}$) where:

$\mathbf{S}$ is the set of "states"

$\mathbf{p}$ is a function that given an element of $\mathbf{S}$ returns a probability measure on $\mathbf{S}$ and

$\mathbf{s_0}$ is an element of $\mathbf{S}$ called the **initial state**.

For our purposes, the set of states will be finite with, say, n elements. In this case, the mapping $\mathbf{p}$ can be expressed as an nxn matrix, called the **transition matrix**. The entries in the matrix will be real numbers between 0 and 1, inclusive. Also, each row of the matrix will sum to 1; such a matrix is called a **stochastic matrix**[1].

An Example

Suppose that every year there is a 10% chance of an earthquake of a certain magnitude. Our set of states will consist of two states: "no quake yet" (or NQY) and "had quake" (or HQ). Our transition matrix is 2x2 and looks like this:

[1] There is a vast literature on Markov processes; a good introduction is [R].

	NQY	HQ
NQY	0.90	0.10
HQ	0.00	1.00

The first row says, if we haven't had a quake yet, then there is a 90% chance that we won't have one this year and a 10% chance that we will. The second row simply says, if we have already had a quake, then we have already had a quake! Also, we will suppose that the initial state is NQY.

We now have the three ingredients needed to have a Markov process, namely, the set of states, the transition matrix, and the initial state. We will return to this example after a final definition.

Suppose that we have a Markov process, $(\mathbf{S}, \mathbf{p}, \mathbf{s_0})$. From the set of possible states, $\mathbf{S}$, we select a subset $\mathbf{T}$ and call these **terminating states**. Consider now a contingent perpetuity that pays \$1.00 at the end of each period until the Markov process enters one of the states in $\mathbf{T}$ at which point it permanently stops paying and becomes worthless. Such a contingent perpetuity we will call a **Markovian annuity**.

As an example, consider a life annuity on a 40 year-old. Let the set of states be his possible ages ("40", "41", "42", ...) along with a special state, "Dead". And let the transition probabilities be given by the life table (i.e., for each age N, state "N" goes to state "N+1" with probability $\mathbf{p_N}$ and to state "Dead" with probability $\mathbf{q_N}$). If we define "Dead" to be the terminating state, then this life annuity is a Markovian annuity.

Casualty actuaries reserving for certain worker's compensation claims, such as "permanent totals" and "permanent partials" already use similar techniques. In fact, in some jurisdictions, these are the only reserves that insurers can discount. This is the so-called "tabular discount" in statutory accounting.

Returning to our earthquake example from above, if we let the state HQ ("had quake") be the terminating state we can value the Markovian annuity that pays $1.00 at the end of each year until there is a quake. Denote this perpetuity by $\mathbf{a_{eq}}$. We have:

$$\mathbf{a_{eq}} = \nu\,(1+ \mathbf{a_{eq}})(.90)$$

This says that in our risk neutral world an investor is ambivalent between owning this annuity today and having the discounted value of a portfolio consisting of $1.00 and the annuity, a year from now, if he gets it. The difference between this formula and the formula for a risk-free perpetuity is the final factor of .90, which is the annual probability that the perpetuity does not default. Using the fact that the interest rate is 8% and solving, we obtain:

$$\mathbf{a_{eq}} = 5.$$

Observe that this is only 40% of the value of the risk-free perpetuity, $\mathbf{a}_{rf}$, which we earlier showed has value 12.5, even though the only difference between the two is a 10% annual default probability.

Suppose that an investor has $5,000 to invest. He could buy 400 risk-free perpetuities ("the risk-free portfolio") or 1,000 of these earthquake perpetuities ("the risky portfolio"). Assume for the moment that the default events are all independent. After one year, with the risk-free portfolio he will have on average the 400 perpetuities that he started with (no defaults) and $400 in cash. The market value of this portfolio is $5,400. With the risky portfolio at the end of one year, he will have (on average) 900 non-defaulted perpetuities and each of them will have paid him $1.00, so he will have $900 in cash. The market value of this portfolio is 900*5 + 900 = 5,400 --- the same as the risk-free portfolio.

The (expected) return of the risk-free portfolio consisted of interest of 400 (the cash) and capital gains of 0 (no defaults). The (expected) return of the risky portfolio consisted of interest of 900 (the cash) and capital gains of –500 (the value of the 100 defaulted perpetuities which are now worthless). This must be so, because in the risk neutral world all investments have the same expected returns (8%).

This example had only two states, defaulted and non-defaulted. In the next section we will consider an example that has four states and is considerably more interesting. To motivate it, we will briefly discuss rating agencies.

Rating Agencies

In our risk-neutral world securities are priced at their expected present values. In order to compute these expectations, investors need to know what the probabilities are that various cash flows will actually occur. In our earthquake example, all investors knew that the annual probability of an earthquake (default) was 10%. How do they obtain this information?

In our simplified risk neutral world (and in the real world) there are entities called rating agencies. Rating agencies evaluate investments and estimate the probabilities that various payments will be made. In our simple world, the rating agencies classify all risky perpetuities into one of four classes named A, B, C, and D.

Securities rated B by the rating agency are considered more risky (likely to default) than those rated A; those rated C are even more risky than those rated B; those rated D have already defaulted and are now worthless[2]. Each year the rating agency reevaluates each security and reclassifies it. Movements between the various non-defaulted classes are described as follows: if a security is now less risky than it was before (i.e. its rating has gone from B to A, C to B, or C to A) we say that the security has been **upgraded**; securities that are now riskier than before (A to B, B to C, or A to C) are said to have

[2] Real world defaulted securities may not be worthless. Estimating the amount of recovery available from a defaulted security is generally a difficult problem on which much research has been done. For simplicity, we will assume that the recovery is zero.

been **downgraded** and finally, securities that are left at their previous risk levels are said to have had their ratings **reaffirmed**.

The movements between rating classes in our simple world is given by the following transition matrix:

	A	B	C	D
A	0.90	0.05	0.04	0.01
B	0.09	0.81	0.05	0.05
C	0.01	0.14	0.75	0.10
D	0.00	0.00	0.00	1.00

Suppose that we wish to determine the price of an A-rated perpetuity, $\mathbf{a_A}$. Under the transition matrix, we have a Markovian annuity. To price this, we proceed as before:

$$\mathbf{a_A} = v\,(.90\,\mathbf{a_A} + .05\,\mathbf{a_B} + .04\,\mathbf{a_C} + (1-.01)\,(1))$$

where $\mathbf{a_B}$ and $\mathbf{a_C}$ are B-rated and C-rated perpetuities, respectively.

This comes directly from the first row of the transition matrix. An investor is ambivalent between an A-rated perpetuity today and the present value of a portfolio which contains an A-rated perpetuity 90% of the time, a B-rated perpetuity 5% of the time, a C-rated perpetuity 4% of the time, and \$1.00 that is paid unless the original perpetuity has defaulted (non-default = 99%).

Before, we had one equation in one unknown. Now it appears that we have one equation in three unknowns. Fortunately, there are more rows of the transition matrix and these supply us with more equations, namely:

$\mathbf{a_B} = v\,(.09\,\mathbf{a_A} + .81\,\mathbf{a_B} + .05\,\mathbf{a_C} + (1 - .05)\,(1))$ and

$\mathbf{a_C} = v\,(.01\,\mathbf{a_A} + .14\,\mathbf{a_B} + .75\,\mathbf{a_C} + (1 - .10)\,(1))$

Now we have three linear equations in three unknowns. Solving we obtain:

$\mathbf{a_A} = 9.027$

$\mathbf{a_B} = 7.687$ and

$\mathbf{a_C} = 6.262$

These are the market prices for risky perpetuities in our risk-neutral world; we will use these prices in the following sections.

Real world rating agencies such as Standard & Poor's (S&P) and Moody's Investors Service (Moody's) have much more refined class plans than we have shown here. Not only are there generally more rating classes, but also rating agencies will sometimes indicate that a rating is "on watch". This frequently means that a rating change is being considered or that new news is expected. Rating agencies serve an important role in financial markets by reducing information asymmetries between issuers and investors. Rating agencies are discussed more fully in [F], [M], and [W].

A final comment on transition matrices, the transition matrix describes the migration over time among the various rating classes. A portfolio initially consisting only of A-rated securities will, over time, become more risky as some of the securities get downgraded. On the other hand, a portfolio that consists of only C-rated securities will, over time, get less risky as securities get upgraded. Here we are only looking at the surviving (non-defaulted) securities. Is there a portfolio that maintains its riskyness over time?

It turns out that the answer is yes. This "eigenportfolio" for lack of a better name, is related to the dominant (left) eigenvector of a certain submatrix of the transition matrix. The corresponding eigenvalue turns out to be one minus the average default rate for the "eigenportfolio". As the reader may check, for the transition matrix given earlier, a portfolio consisting of 50.32% A-rated securities, 32.49% B-rated securities, and 17.19% C-rated securities will (in expectation) maintain its proportions over time, the eigenvalue in this case being 0.96153 and the average default rate being 0.03847.

Transition matrices appear in many fields of study. For example, they are used to study population dynamics in mathematical ecology where they are called "Leslie matrices".

Leslie matrices are named after P.H. Leslie who introduced them into biology in the mid-forties. See [L].

Perpetuities

Suppose that a company wishes to raise funds in our risk neutral world. The company wants to borrow \$1,000. In exchange for \$1,000 today the company will pay annual interest until it defaults. Further suppose that our company is rated "B" by our rating agency. Recall from our previous calculations that $\mathbf{a}_B$ a B-rated perpetuity paying \$1.00 each year has a value of \$7.687. We wish to find the amount of the coupon, K, that must be paid so that the market price of the security will be exactly \$1,000.00. In symbols:

$$1{,}000 = K\ \mathbf{a}_B$$

That is, an investor is ambivalent between keeping his \$1,000 today and getting the present value of a perpetual stream of payments of \$K annually until default. Replacing $\mathbf{a}_B$ with its value, \$7.687, and dividing we obtain:

$$K = \$130.09$$

Recall that in the risk neutral world, all investments are expected to yield 8%. The investor has only invested $1,000.00, so his expected yield must be 8% of this, namely $80.00. The "extra" $50.09 (=130.09 – 80.00) is compensation for the expected change in the market price of the perpetuity (a capital gain or loss). There are four possible outcomes. It is possible that the perpetuity had defaulted; in this case the investor gets no coupon payment and owns a worthless security. The other possibilities are that the perpetuity has been downgraded, upgraded, or has had its rating affirmed.

Notice that the coupon amount is fixed when the security is issued, and that subsequent upgrades or downgrades do not change the amount of the coupon. Suppose that the perpetuity has been downgraded, so it now is rated "C". The investor will still receive $130.09 per year until default, but now default is expected sooner. We previously computed the value of a stream of $1.00 payments from a C-rated security when we learned that $\mathbf{a}_C$ had a value of $6.262. Using this fact, we can find the market value of the downgraded security. It pays $130.09 per year, so its market value must be:

$$130.09\ \mathbf{a}_C = \$814.62$$

On a mark-to-market basis, the investor has suffered a loss, even though no cash payment has been late or missed. It is generally believed that investors like to get their principal back (although in the risk neutral world they really don't care provided that the coupon is adequate). Real world bonds have maturity dates when the principal is paid back.

Modeling this adds no real obstacles, and adds some interesting twists. To appreciate these subtleties, we will first examine perpetuities in more detail.

Bond Insurance

Suppose that our investor wants to purchase an insurance policy that will pay him $1 when his B-rated perpetuity defaults. Assuming that the insurance company cannot itself default[3], what is a fair premium for this insurance?

Denote by $\mathbf{A_B}$ the one-time premium that the insurer would charge for this insurance. In the risk neutral world, this premium is the expected present value of the benefit, so there will be no ambiguity in denoting the benefit by this same symbol. We have the tools to price this at our fingertips.

$$1.00 = 0.08\ \mathbf{a_B} + 1.08\ \mathbf{A_B}$$

What this says is: an investor is ambivalent between having $1.00 today and receiving the interest on the $1.00 ($0.08) every year until a default occurs. When the default occurs, he gets back his dollar and the final year's interest.

[3] One of the most contentious issues addressed by the white paper on fair value liabilities was related to how the fair value of a liability should depend on the creditworthiness of the parties. See [T], in particular item 15 of the Executive Summary.

This identity should look very familiar to students of life contingencies; it is the fundamental identity relating annuity values and insurance prices. The more traditional version involves annuities-due and discount rates (instead of annuities-immediate and interest rates), because life insurance premiums are paid in advance while bond interest is received in arrears. In this example we can solve and learn that the market price of this insurance is 0.3565 (recall that we computed that $\mathbf{a_B} = 7.687$ in an earlier section).

Suppose that an investor has $1,000 to invest. He elects to purchase a B rated perpetuity that will pay him $80/year (at a cost of 80 * 7.687 = 614.97) and he uses the rest to purchase an insurance policy that will pay him $1,080 when this perpetuity defaults (at a cost of 1,080 * .3565 = 385.03). He has now spent his $1,000 and he has created a synthetic risk-free bond. This bond will pay him $80/year until a default occurs at which point the insurance pays at the end of the year the final interest payment and the principal.

Suppose that a second investor purchases for $1,000 a B-rated perpetuity (which we learned earlier pays annual coupons of 130.09). If he now insures the perpetuity (for his principal plus the risk-free interest on it, i.e. $1,080), but arranges to pay premiums annually in arrears while the perpetuity has not defaulted, what will his annual premium be?

Well, he too has, in effect, turned his risky perpetuity into a risk-free perpetuity. His investment is $1,000, so he is entitled to exactly $80 per year (8%). The difference between the promised coupon, $130.09, and the risk-free coupon, $80.00, must be the insurance premium charged (if not an arbitrage would result)[4]. Bond traders call this difference the **spread**.

There is an interesting relationship between the spread and the default rate. To see it, consider a one year bond which will either default and be worthless (probability = 20%) or will mature and will pay $1,350 in one year (probability = 80%). What would an investor in the risk-neutral world pay for this bond?

The expected present value (at 8%) of this investment is $1,000. So the spread is 27%[5] The default probability is only 20%. The extra 7% is needed because only non-defaulted bonds pay the coupon. The 27% can be thought of as an assessment on the surviving bonds (80%) to pay the principal (100%) and the risk free interest on it (8%) for the defaulting ones (20%). We have:

Spread = 1/(1 – default) * (1 + risk-free) * (default)

0.27 = 1/(1 – 0.20) * (1 + 0.08) * (.20)

[4] Arbitrage opportunities are discussed in a subsequent section.
[5] Spreads are normally quoted in hundredths of a percent, called **basis points**; so, a 27% spread would be said to be a 2,700 basis point spread.

It is interesting to note that the above formula suggests that spreads should widen with increases in the risk-free rate, and that this effect should be more pronounced for worse credits.

Turning Assets Into Liabilities

By using bond insurance as described in the previous section, an investor can take a risky asset portfolio and turn it into a risk-free portfolio. The risk gets transferred to an insurance company where it resides on the liability side of the balance sheet. How should an insurance company account for contracts of this type? What constitutes a loss? How should reserves be valued?

Suppose that an entity purchases a risky perpetuity for $1,000 and insures it. We have seen that the premium paid will be the spread above the risk-free rate and that the insured amount will be $1,080, which is the $1,000 face amount plus the risk-free return (8%).

How does this look from the insurer's point of view? The insurer expects to receive the spread income until the year of the default. At the end of that year, the insurer will pay the $1,080 claim. A moment's thought reveals that the premium stream that the insurer expects to receive is, in fact, a Markovian annuity.

Suppose that we were insuring a B-rated perpetuity. At the end of the year, there are four possible states:

1) It has been upgraded (now rated A).
2) It has had its rating reaffirmed (now rated B).
3) It has been downgraded (now rated C).
4) It had defaulted (now rated D).

In the fourth case, we have paid the loss and there is no reserve. In the second case (rating has been affirmed), we will be receiving as premium the spread on a B-rated bond for insuring a B-rated bond. This premium is, of course, exactly adequate.

If the bond has been downgraded, however, the future spread income is no longer adequate. The expected future premium after the downgrade is $\mathbf{S_B}\ \mathbf{a_C}$, where $\mathbf{S_B}$ denotes the spread on a B-rated perpetuity. The required future premium becomes $\mathbf{S_C}\ \mathbf{a_C}$, where $\mathbf{S_C}$ denotes the spread on a C-rated perpetuity. The shortfall is $(\mathbf{S_C} - \mathbf{S_B})\ \mathbf{a_C}$.

Notice that increase in the bond's mortality contributes in two distinct ways to the shortfall. Not only has the expected future premium income decreased by $\mathbf{S_B}\ (\mathbf{a_C} - \mathbf{a_B})$, but also the required premium has increased from $\mathbf{S_B}\ \mathbf{a_B}$ to $\mathbf{S_C}\ \mathbf{a_C}$. Effectively, fewer premium payments are expected and, additionally, the expected loss payment has been accelerated.

The total shortfall in future premium should be recognized on the balance sheet (and the income statement) as an increase in the premium deficiency reserve. The appropriate accounting treatment of such changes in value is discussed briefly in the Accounting Considerations section.

The fourth possibility is an upgrade. In this case, the future premium income is excessive and under fair-value accounting this too would be reflected in the reserve for unexpired risks. Under codification, it appears that the negative premium deficiency could be used to offset premium deficiencies from other insured perpetuities (ones that had been downgraded), provided that management groups these together for internal reporting. Again, this will be discussed in more detail in the later section.

Remark: There is an important principle here. Memoryless ➔ No Reserve.

This is the case for constant mortality in whole life insurance and it is true here as well. Recall that for a whole life policy the reserve is really a premium deficiency reserve. Typically, premiums are level, but at most ages human mortality is increasing, so early on the premium is more than is needed for current mortality (the difference going into the reserve). Later on the premium is inadequate for the current mortality (but the reserve is there to fund the shortfall). In the constant mortality case, the level (constant) premium exactly matches the current (constant) mortality at all ages, hence there is no need for a reserve. In the same way for perpetuities, if at the end of the year there has been no change in rating (i.e. mortality has stayed constant), then there will be no change in the reserve[6].

While this holds for perpetuities, it does not in general hold for bonds. The difference is that over time bonds approach maturity, when the principal becomes due. A (non-defaulted) maturing bond pays its principal payment regardless of its rating. A risky bond one year from maturity and a risky bond two years from maturity may have very different prices. The life insurance analog of this phenomenon is that an endowment policy even with constant mortality still will build up a reserve (to pay the endowment amount at maturity). We will see how the prices of risky bonds change over time in a following section, but first we will examine a technique from the capital markets used for pricing risky cash flows.

Replicating Portfolios

Reserving frequently involves estimating the value of a collection of future cash flows. A very elegant technique for valuing such flows comes from modern finance theory. The crux of the idea is extremely simple: if two collections of cash flows are identical, then they must have identical prices.

---

[6] These are premium deficiency reserves and, as such, should be carried at discounted value. The annual unwind in the reserve is exactly enough to make up for the annual deficiency in premium.

Suppose that we have a collection of (contingent) cash flows that we wish to value. We try and find a second collection of securities that taken together have cash flows identical with our collection in all states of the world. For example, if the first one pays a dollar when there is a particular earthquake, the second one must also pay a dollar for the same earthquake. Such a collection is called a **replicating portfolio** for the first collection. Generally, it will be difficult to find such a portfolio because it must match exactly in all cases. However, if you are lucky enough to find one and the securities have market prices, then you have found the market value of your set of cash flows.

Let's look at some simple examples. Suppose that available in the market are three securities, all newly issued, risk-free annuities-immediate with terms of 3, 5, and 10 years, respectively. The market prices in our risk-neutral world for these annuities are given in the next table. (What is especially nice about this approach is that if you have real-world prices for these securities, you get the real-world price of your liability!)

$\mathbf{a}_3 = 2.577$

$\mathbf{a}_5 = 3.993$ and

$\mathbf{a}_{12} = 7.536$

Suppose that we wish to reserve for a stream of payments of \$8 for three years followed by \$2 for nine more years. A moment's thought reveals that this stream of payments can be obtained by buying 6 of $\mathbf{a}_3$ and 2 of $\mathbf{a}_{12}$. (Both of these types of annuities pay during the first three years yielding eight dollars per year; for the last nine years only the second type pays, yielding the required two dollars per year.) The cost of this portfolio is \$30.535 (= 6(2.577) + 2(7.536)) and, since it matches our payment stream exactly, is the market price of our liability.

As a second example, consider an obligation to pay \$1 per year for seven years starting in five years. We would like to reserve for this stream of payments by finding the market value of this liability. This is a 5-year deferred, seven-year annuity. It can be replicated as follows: purchase an $\mathbf{a}_{12}$ and sell an $\mathbf{a}_5$. You may wonder how we can sell something that we don't own, but for the moment, assume that this transaction can be done. What are the cash flows from the resulting portfolio? Well, in years one through five, we receive a dollar from the $\mathbf{a}_{12}$. The investor that purchased the $\mathbf{a}_5$ from us expects to receive a dollar. We take the dollar that we get from the $\mathbf{a}_{12}$ and give it to the purchaser of the $\mathbf{a}_5$. The investor is happy because he does not care which dollar he gets, he just wants a dollar to be paid to him at the end of each of five years. At the end of year five, the $\mathbf{a}_5$ makes its last payment and expires worthless. In years six through twelve we receive one dollar from the original $\mathbf{a}_{12}$. This exactly matches the payments that we will make on the deferred annuity, so this is a replicating portfolio. How much does this portfolio cost? Well, we know that we can buy an $\mathbf{a}_{12}$ for 7.536, since that is its market price. We can also sell an $\mathbf{a}_5$ for 3.993, since that is its market price, so the net cost of the portfolio is 3.543 (= 7.536 − 3.993). This is the market price for our liability.

Something interesting has happened; we have been able to compute the exact market price for this liability even though no market for it (directly) exists.

In the last example, we bought one annuity and sold another; practitioners would describe this as a **long position** in the $\mathbf{a}_{12}$ and a **short position** in the $\mathbf{a}_5$. We will use this terminology in what follows. We need to define one more term.

A portfolio with some positive cash flows, no negative cash flows, and zero net cost is called a **risk-free arbitrage opportunity**. Such a portfolio would also be a tremendous bargain! So much so, that there would be unlimited demand for it. This demand would be so great that it would cause market prices to shift to eliminate the opportunity. There are no risk-free arbitrage opportunities in the risk-neutral world, and it is generally believed that there are none in the real world either.

Suppose that two portfolios have identical cash flows, then they must have identical prices. Here is why. Suppose that the prices were different, then we would short the more expensive one (sell it) and go long the cheaper one (buy it). The resulting portfolio would have a positive cash flow at time zero (the difference in the prices), have no net cost, and would have no negative cash flows, so it would be a risk-free arbitrage opportunity. There would be unlimited selling pressure on the more expensive one, pushing its price down, while there would be unlimited buying pressure on the cheaper one, driving its price up. This process would continue until the two prices were equal.

The reader may have come across replicating portfolios before in studying the Black-Scholes solution to the call option-pricing problem. See for instance, [B].

Bonds

As previously noted, in the real world investors like to get their principal returned to them. A newly issued bond may have a maturity of thirty years. Such a bond will pay annual interest at the end of each of the first twenty-nine years and then will pay back the principal amount and the final year's interest at the end of year thirty. Of course, along the way, the bond may default.

Issuers tend to set the coupon so that their bonds will sell "at par". That is, they generally adjust the spread that they offer to pay, so that a bond with $1,000 in principal will sell for $1,000 at issue. Table 1, below, shows the annuity values and required coupon amounts for newly issued C-rated bonds to trade at par.

A comment on how the annuity values are computed is in order. The annuity values are computed recursively from the transition matrix. One year from maturity, the bond either defaults (probability = 10%) or it matures (probability = 90%). With i = 8%, we find the value of a one year C-rated annuity to be 0.9(1/1.08) = 0.83333. The values of A-rated and B-rated one-year annuities are found similarly. Once these values are in hand, we can value two year annuities using the transition matrix as we did above for perpetuities, then recursively we can compute the values for longer term annuities. The results are shown in Table 1.

TABLE 1

Newly issued C-rated bonds

Years to Maturity (N)	Actuarial PV of Principal	Annuity Value	"Required" Coupon	Actuarial PV of Coupons
1	833.33	0.83333	200.00	166.67
2	701.22	1.53455	194.70	298.78
3	595.12	2.12967	190.12	404.88
4	508.84	2.63851	186.15	491.16
5	437.88	3.07639	182.72	562.12
6	378.90	3.45529	179.75	621.10
7	329.39	3.78468	177.19	670.61
8	287.48	4.07216	174.97	712.52
9	**251.74**	**4.32389**	173.05	748.26
10	221.05	4.54494	**171.39**	778.95

Suppose that a firm issues for $1,000 a 10-year C-rated bond and that one year later the bond is still C-rated. What is the market price of the bond now?

The bond when issued was a 10-year bond and one year has passed, so it is now a 9-year bond. It is still rated "C", so Table 1 contains all of the information that we need. From column 2 we learn that the principal amount has an actuarial present value of $251.74.

From column 3 we see that each dollar of coupon has an actuarial present value of \$4.32389. Now the coupon gets set when the bond is issued, so it is still 171.39 (from column 4, row 10). Combining all of the information we see that the market price is

$251.74 + 171.39(4.32389) = 992.80$

The bond was worth \$1,000.00 at issue, but now it is worth only \$992.80. There has been no default nor has there been a downgrade, but the owner of the bond still lost \$7.20 (in market value). Figure 1 shows the annual change in the market price of this bond assuming that its rating never changes over its life.

Figure 1

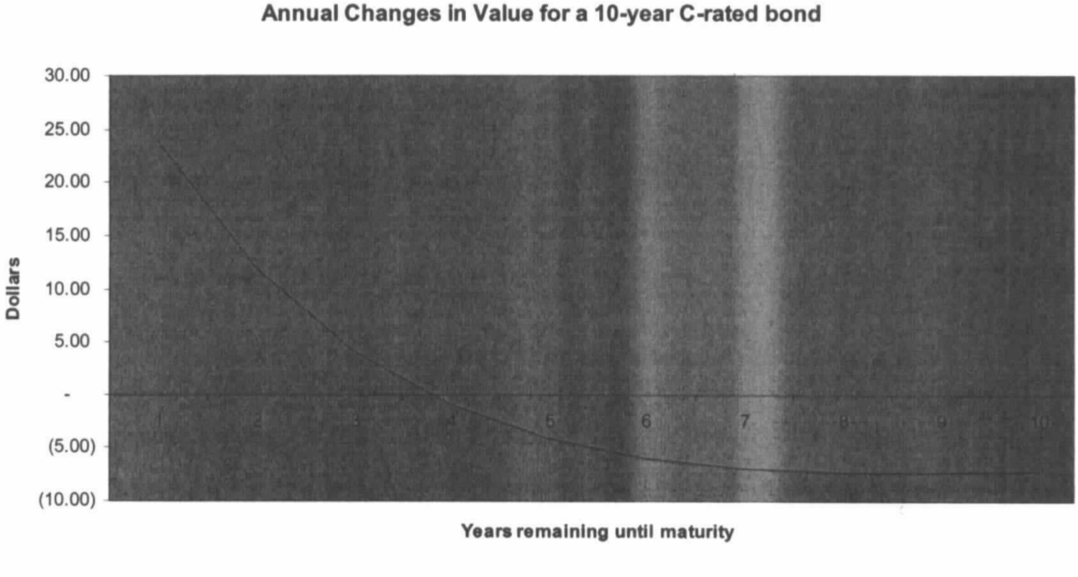

There are two competing forces affecting the bond price. Reviewing Table 1, we see that the required coupon increases as the time to maturity decreases, since the coupon is fixed at the 10-year value as the bond approaches maturity the coupons become more and more inadequate, pushing the price down. On the other hand, the actuarial present value of the principal payment rapidly increases as maturity nears. The combined effect is shown in Figure 1, where we can see that the coupon effect dominates when there are many years left to maturity, but when the bond is close to maturity the value of the principal starts to dominate.

Suppose that you are an insurer and that you have insured a 10-year C-rated bond against default. If one year has passed and the bond is still rated "C", you should put up a reserve. In particular, you should carry a premium deficiency reserve sufficient to allow you to reinsure your risk[7]. A loss reserve is not appropriate, because the covered event is default and default has not occurred. On the other hand, even in the risk-neutral world a reinsurer would require compensation in order to take over your current position. The amount that the reinsurer would require is exactly the difference between the current market price of the bond and the principal amount.

To see this we will create a replicating portfolio that exactly duplicates the cash flows of that the insurer will have to pay out. The cost of this perfect reinsurance will be the cost of the replicating portfolio. The required portfolio is a short position in the risky bond

[7] In the risk-neutral world, reinsurers will assume risks for the difference between their expected future discounted premiums and their expected future discounted losses.

(principal amount = $1,000) and a long position in a risk-free security (principal amount = $1,000). We will check the cash flows in each possible scenario.

During years when the bond does not mature and does not default, we receive premium equal to the spread, and investment income from the risk-free bond. The sum of these is exactly the coupon payment that we need to make on our short position, so we have no net flows. In the year that the bond matures if there is no default, things are exactly as in the previous case except that we need to pay the principal on our short position, we do this with the principal from the risk-free security. The short position is now closed, and the insurance has expired without a claim: no net cash flow, no outstanding liabilities (nor assets) remain. Finally, if there is a default, we sell the risk-free security (for $1,080); this is exactly the insured amount of the bond (recall that the policyholder insures the bond for principal and risk-free interest). In all three cases there are no net cash flows. That is, the portfolio exactly hedges the insurance policy and the cost of the portfolio is exactly what a reinsurer would charge (in the risk-neutral world) to take this risk from your books.

The cost of this portfolio is the difference between the cost of the risk-free bond ($1,000.00) and the market price of the risky bond which we earlier calculated to be $992.80 (the value of a 9-year C-rated bond, paying a 10-year C-rated coupon).

One might wonder why a premium deficiency arises in this case. We started with a C-rated bond and one year later we still had a C-rated bond --- no default, yet it appears that we have a loss. The reason is that in some sense you have had bad luck. While nothing explicitly bad has happened (a default), nothing good has happened either (an upgrade). The market had already priced the possibility of an upgrade into the required coupon. When the upgrade did not occur, the market price reflected the lack of good news.

Accounting Considerations

The NAIC's statutory accounting codification project now requires an estimation of the premium deficiency reserve for all property/casualty companies. Because of our simplifying assumptions (no reporting lag, losses and payments occurring only at the end of a year) the types of insurance products described here do not generate loss reserves, but they will generate premium deficiency reserves.

Accounting practice seems to be to earn spread income as it is received. Assuming that the spread income is treated as written when received, the insurer will carry no unearned premium reserve for these products. We have seen that earning the spread as received is exactly correct for perpetuities because of their memoryless feature. However for bonds, a premium deficiency could arise.

Should contracts such as bond insurance be treated as insurance at all? Guidance on this point under International Accounting Standards (IAS) rules can be found in [S]. Subissue 1-G states the Steering Committee's view that a contract is to be treated as insurance (and would come under IAS 37) if the triggering event is a failure "to make payment when due". However, if the triggering event were a downgrade, it would be treated as a financial instrument (and would come under IAS 39).

Under US GAAP, the line of demarcation seems less clear. FAS 133 covers derivatives and FAS 60 covers insurance. FAS 133 explicitly excludes "insurance" from its scope. I would presume then that bond insurance would be insurance, however it is not clear to me how a policy that protected against a rating agency downgrade would be treated under US GAAP. Anecdotally, I have heard that in the past "downgrade insurance" has been treated as insurance by some auditors, but I do not know if this is standard practice.

Assuming that these contracts are appropriately accounted for as insurance, they will generate premium deficiency reserves. Some contracts will generate positive premium deficiencies and others may generate negative premium deficiencies. Under codification, to the extent that management groups these contracts together for internal reporting they should be offset against one another for statutory accounting purposes, with only a net premium deficiency, if any, being reported.

Reserving the World Series

In this final example we will illustrate how an arbitrage argument can be used to evaluate the value of a wager on the outcome of a series when only partial information is available.

Suppose that you have wagered \$100 that team A will beat Team B in a best 4 out of 7 series. You believe that the probability that either team will win any given game is 50%. Your team (Team A) loses the first game. What is the value of your wager, given the first game result? In other words, what reserve should you be holding against the potential \$100 loss?

In the risk free world, answering this question is equivalent to determining what an investor would pay you (or demand that you pay him) to take over your position. This last question we can answer through an arbitrage argument. Let R(a,b) be the amount that the investor would be willing to pay you (or that he would demand) when Team A has won "a" games, and Team B has won "b" games. The possible states of the series are pairs (x,y) where "x" and "y" are each between zero and four (but they cannot both be four). Transitions between states occur based on the outcome of the next game, state (x,y) being equally likely to go to state (x+1,y) or state (x,y+1). The initial state was (0,0). We have a Markov process.

Since the series ends when either team has won 4 games, we have:

$R(0,4) = R(1,4) = R(2,4) = R(3,4) = -100$ and

$R(4,0) = R(4,1) = R(4,2) = R(4,3) = 100$

From this we conclude that $R(3,3) = .5(-100) + .5(100) = 0$. This follows because when you have a 3-3 tie the final game is decisive.

As we continue to back-solve we learn that:

$R(2,3) = .5\ R(3,3) + .5\ R(2,4) = 0 - 50 = -50$

$R(3,2) = .5\ R(4,2) + .5\ R(3,3) = 50 + 0 = 50$

$R(2,2) = .5\ R(3,2) + .5\ R(2,3) = 50 - 50 = 0$

$R(3,1) = .5\ R(4,1) + .5\ R(3,2) = 50 + 25 = 75$

$R(1,3) = .5\ R(2,3) + .5\ R(1,4) = -25 - 50 = -75$

$R(2,1) = .5\ R(3,1) + .5\ R(2,2) = 37.5 - 0 = 37.5$

$R(1,2) = .5\ R(2,2) + .5\ R(1,3) = 0 - 37.5 = -37.5$

$R(3,0) = .5\ R(4,0) + .5\ R(3,1) = 50 + 37.5 = 87.5$

$R(0,3) = .5\ R(1,3) + .5\ R(0,4) = -37.5 - 50 = -87.5$

$R(1,1) = .5\ R(2,1) + .5\ R(1,2) = 37.5 - 37.5 = 0$

$R(2,0) = .5\ R(3,0) + .5\ R(2,1) = 43.75 + 18.75 = 62.5$

$R(0,2) = .5\ R(1,2) + .5\ R(0,3) = -18.75 + -43.75 = -62.5$

$R(1,0) = .5\ R(2,0) + .5\ R(1,1) = 31.25 + 0 = 31.25$

$R(0,1) = .5\ R(1,1) + .5\ R(0,2) = 0\ \ -31.25 = -31.25$

So, the investor would take over your position for a payment of \$31.25. This is the reserve that you should carry for this wager. Note that it is a premium deficiency reserve, since the wager isn't lost yet, but your odds of winning have diminished.

It is interesting to note that the above calculation gives an explicit **defeasance strategy** for the wager from any point in time. A defeasance strategy is a set of explicit instructions on what bets to place and for how much to insure that the net cash flows from all of the bets exactly match the cash flows of the liability. In effect, we have explicitly exhibited a replicating portfolio of single game, even money bets that have a cumulative payoff of precisely \$100 if Team A wins the series and -\$100 if Team B wins the series.

This example is not as artificial as it might appear. A reinsurer negotiating a commutation of an inforce treaty could easily find itself in a comparable position. Determining the value of a reinsurance treaty midterm is generally a difficult problem, but if a replicating portfolio with market prices can be found, then the problem is solved.

Conclusion

Reserving actuaries need to opine on the adequacy of the unearned premium reserve for certain lines of business. Determining the existence of a premium deficiency or estimating its size can be difficult. For certain types of risks we have shown how it is possible to estimate the required premium deficiency reserve by using market prices and an arbitrage argument.

Spread income is traditionally earned as received. This is exactly correct for perpetuities that have not had their ratings changed. For bonds though, a premium deficiency can arise even if there is no change in rating.

In order to compute the premium deficiency future premium flows need to be estimated. Viewing these as Markovian annuities can facilitate this estimation. Life contingency techniques and notation, turn out to be quite convenient for this.

Life contingency texts have many formulas and identities that life reserves satisfy. Most of these have analogs for Markovian annuities and insurances. This is not surprising since such annuities are generalizations of level premium life insurance.

Bibliography

[B] Brealey, Richard A. and Myers, Stewart C., Principles of Corporate Finance, McGraw-Hill

[F] Fabozzi, Frank J. (editor), The Handbook of Fixed Income Securities, Irwin

[L] Leslie, P.H., "On the use of matrices in certain population mathematics" *Biometrika* (XXXIII), 1945, pp. 183-212

[M] Megginson, William L., Corporate Finance Theory, Addison-Wesley

[R] Ross, Sheldon M., Introduction to Probability Models, Academic Press.

[S] Steering Committee on Insurance, Issues Paper November 1999, Insurance, International Accounting Standards Committee, 1999

[T] CAS Task Force on Fair Value Liabilities, White Paper on Fair Valuing Property/Casualty Insurance Liabilities, CAS

[W] Wilson, Richard S. and Fabozzi, Frank J., Corporate Bonds, Fabozzi Associates