

A Risk Management Tool for Long Liabilities: The Static Control Model

B. John Manistre¹ FSA, FCIA, MAAA, CERA

January 31, 2009 (Revised March 31, 2009)

Abstract²

This paper looks at the problem of valuing and managing the A/L M risks associated with insurance liabilities that are too long to be matched by available investments. Two very different approaches to the problem are explored. The first approach called Yield Curve Extension starts with a number of simple ideas for extrapolating a yield curve and analyzes them from a risk management perspective. The paper concludes that these methods lead to unnecessarily extreme A/L M strategies. The paper then describes a second approach called the Static Control Model which allows one to use a total return vehicle as part of the A/L M strategy. The model decomposes a long liability into fixed income and total return components in a market consistent way. The fixed income component is a static hedge for the liability in the sense that it matches the first order sensitivities of the model liability as observable market information changes. The paper concludes by arguing that the Static Control Model leads to more useful A/L M strategies for long liabilities.

Introduction and Summary of Results

The issue of managing long liabilities has been a practical problem for life insurers and pension plan sponsors for many years. In developed economies products such as Long Term Care insurance in the United States, or Term to 100 life insurance in Canada, have very low lapse rates and can create liabilities with very long durations. In less developed economies even traditional life insurance products can be difficult to match simply because the local debt markets are not well developed.

While the problem is far from new the advent of market consistent financial reporting requirements and the demands of evolving Enterprise Risk Management standards mean that the inherent difficulties should be addressed more comprehensively than they have in the past.

So what should a comprehensive approach to managing long liabilities look like? In this author's opinion a comprehensive approach should be able to do the following

1. Put a value on the liability that can be used in a market consistent balance sheet.
2. The value must roll forward in time in such a way that the resulting market consistent income statement makes sense.

¹ The author is Vice President for Risk Research in the Group Risk Dept. of AEGON NV. 1111 North Charles St. Baltimore, Md. USA.

² The views and opinions expressed in this paper are those of the author and not AEGON NV.

3. The method should identify which components of the risk can be hedged in the capital markets and which cannot. The hedgeable component of the risk must be valued in a market consistent way.
4. Risks that cannot be hedged should be valued using either the cost of capital method as described by the European Chief Risk Officers Forum or some other method acceptable to the actuarial profession.

This paper illustrates takes two very different approaches to meeting the above requirements.

The first approach, called a Yield Curve Extension Method, starts with some simple ideas for extrapolating a forward rate curve and then tries to fine tune the extrapolation so as to meet the requirements listed above. The end result of this process is a formal theory of yield curve extension similar to the first generation of stochastic interest rate models developed by Oldrich Vasicek and other researchers.

The paper ultimately concludes that this approach is a failure because the underlying assumed hedging strategies are unrealistic but the model building exercise is very instructive. The results of this approach are used as a benchmark for testing the next set of ideas.

The second approach starts with the idea of using available market instruments to match the hedgeable component of the risk while taking a total return approach to the unhedgeable risk. As before, we start with some simple ideas and then fine tune the process. The resulting model is called the Static Control Model here because it ends up solving an optimization problem over a space of static (buy and hold) investment strategies. The Static Control Model is technically more complex than any simple yield curve extension but it usually arrives at similar values while using hedging strategies that, in the author's opinion, are more realistic.

Despite the use of a total return asset class in the Static Control Model we are able to get an extended yield curve out of the model that we call the Marginal Cost Yield Curve (MCYC). This turns out to be a very important tool for constructing hedge portfolios and understanding the roll forward in time of the model's values.

One result of interest to come out of the Static Control Model is an answer to the question of what the long forward rate in the extended yield curve should be. If the total return vehicle is assumed to be a standard log normal process $\exp[\mu t + \sigma Z(t)]$ with parameters μ, σ calibrated such that the expected total return is 8.00%, i.e. $\exp[\mu + \sigma^2 / 2] = 1.08$ and the volatility is $\sigma = 16\%$ then the model's implied long forward rate turns out to be 5.27%. This is the rate that produces the expected discount factor $\exp[-\mu + \sigma^2 / 2] = 1/1.0527$.

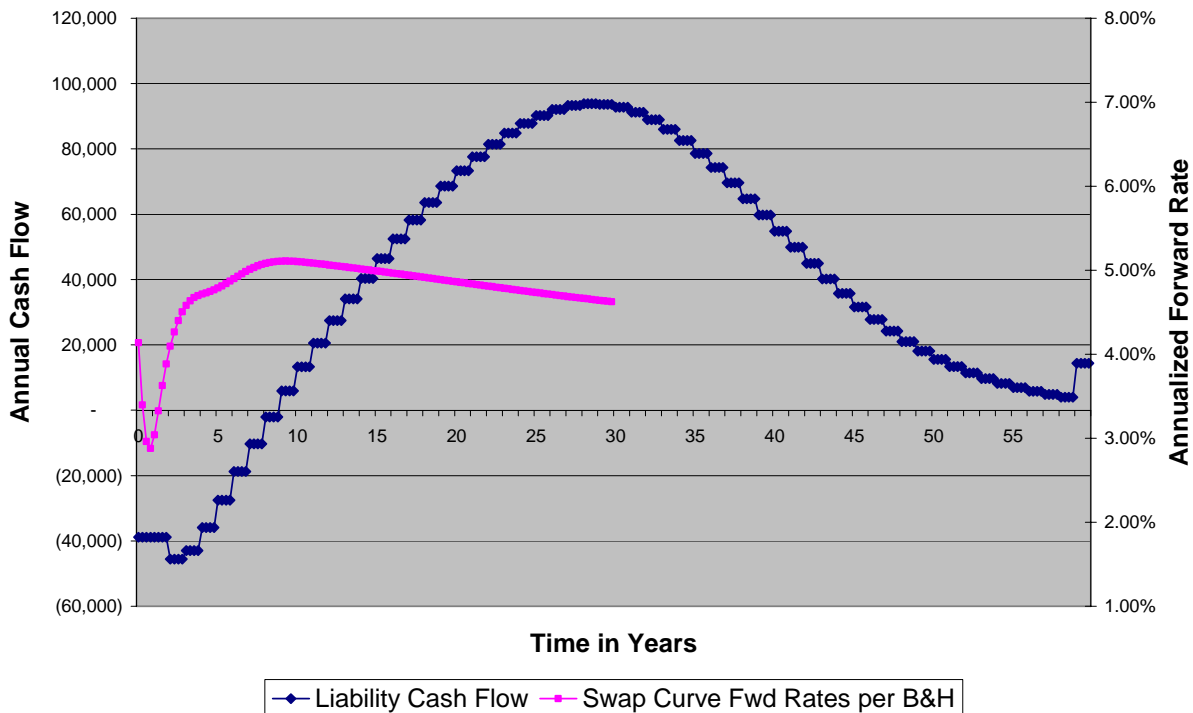
A property of the Static Control Model that some readers may not like is the fact that the implied extension of the forward rate curve is not necessarily continuous. However, we'll argue in the section on Yield Curve Extension Methods that demanding continuity of the forward rate curve is part of the problem and not part of the solution.

Following this introduction the paper is divided into three main sections followed by two appendices. The first section summarizes the ideas behind the Yield Curve Extension Method while the second section describes the Static Control Model. A short Conclusions section

summarizes the author’s arguments in favor of the Static Control Model. The more mathematical details required to support both methods are in Appendices 1 & 2 respectively.

The following data will be used to provide numerical examples of all the methods discussed in this paper. The chart below illustrates 60 years of quarterly cash flow adapted from one of the business units operated by the author’s employer. Like most long liability examples it has cash inflows during the early years which then turn into cash payouts that extend for many years. Cash flows beyond year 60 have been simplified by collapsing them into year 60.

Model Inputs: Long Liability at 9/08



The chart also shows 30 years worth of the \$US swap rates at Sept. 31, 2008 as determined by the consulting firm Barrie & Hibbert³. The numerical examples will assume that we can buy, or sell, zero coupon bonds at prices determined by these forward rates. The last quarterly forward rate was 4.63%.

The methods of this paper will then be used to add an additional 30 years to this observed yield curve.

³ While we treat the “observable” yield curve as a hard fact in this paper this is not really the case. There is an element of art in the process of estimating the yield curve.

The Yield Curve Extension Approach

The Simple Monopole

Perhaps the simplest idea one could start with is to assume a constant single forward rate for all durations greater than 30. To be consistent with later examples we'll assume this is an annual effective rate $f = 5.27\%$ which corresponds to a continuously compounded rate $\bar{f} = 5.14\%$. The value that we assign to cash flows beyond the yield curve horizon can then be written as $V = Z_{30}K$ where Z_{30} is the value of a 30 year 0 coupon bond as determined by the yield curve and K is the present value of cash flows in years 30+ discounted to year 30 using f . In symbols

$$\begin{aligned} V &= Z_{30}K, \\ &= Z_{30} \sum_{j>30} a_j (1+f)^{-(j-30)}, \end{aligned}$$

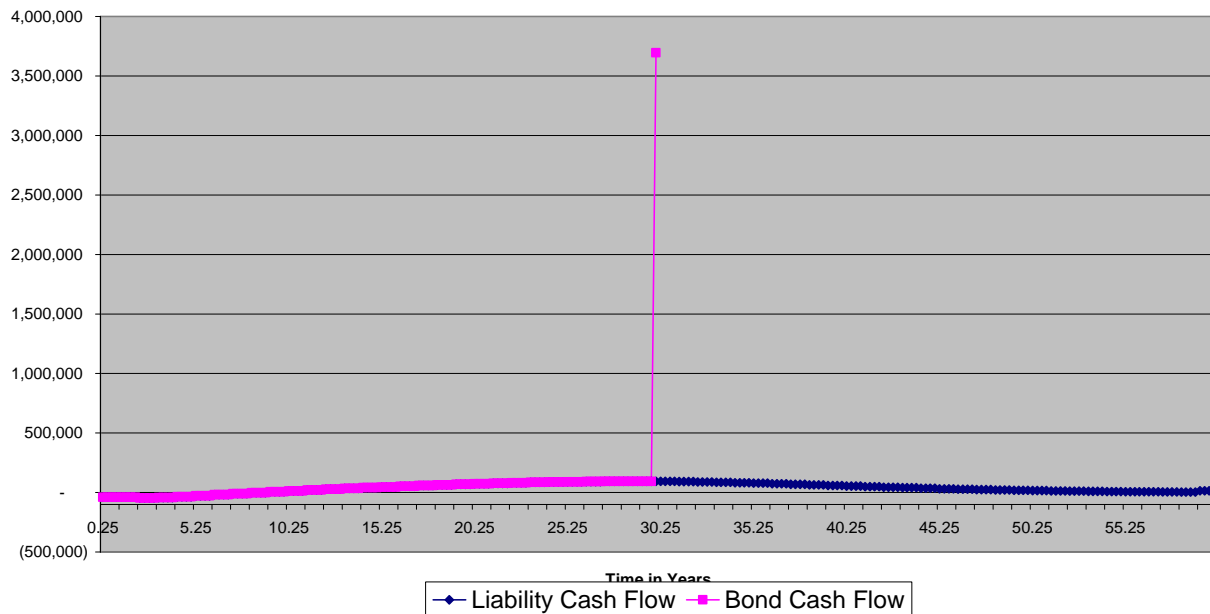
where the a_j , $j > 30$ are the cash flows beyond the yield curve horizon which we will assume are deterministic in this paper.

We ask two types of questions of this model:

1. How does the value react to an instantaneous shock in the observed yield curve? Can we hedge random movements in the observed yield curve? If so, we will call the resulting hedge portfolio the *static hedge*.
2. Is the static hedge portfolio, if it exists, risk free and self financing as time moves forward? If not we will say that the model has an *unhedged risk*, a *bias* or perhaps both.

The assumed model is simple enough that we can give a complete answer to both questions. Since the quantity K is independent of the observed yield curve we can construct a static hedge simply by investing the entire value $V = Z_{30}K$ in a 30 year 0 coupon bond. The chart below shows what the resulting static hedge portfolio looks like for the entire liability assuming we buy 0 coupon bonds to match all cash flows less than 30 years out.

Long Liability Monopole Static Hedge



For obvious reasons we will refer to this as the monopoly strategy. Let's ignore the issue of how impractical such a strategy might be and look at how the static hedge rolls forward in time. Table 1 below walks through the steps.

Table 1 Monopole Static Hedge

$K = \sum_{j>30} a_j (1+f)^{-(j-30)}$	Assets	Liabilities	Difference
$t = 0$	KZ_{30}	KZ_{30}	0
$t = 1$	$KZ'_{29} = (1+f'_{30})KZ'_{30}$	$(1+f)KZ'_{30}$	$(f'_{30} - f)KZ'_{30}$

The table assumes that at time $t = 0$ assets and liabilities are matched in the sense that we are invested in the static hedge portfolio.

One time period later the yield curve has changed and our 30 year 0 coupon bond has become a shorter bond, say 29 years, whose value is given by $KZ'_{29} = (1+f'_{30})KZ'_{30}$ where f'_{30} is the new forward rate at year 30. However, the liability K factor has rolled forward at the discount rate used to extend the yield curve. The net result is a gain/(loss) which is determined by the difference $(f'_{30} - f)$ between the new forward rate and the discount rate.

Is this an acceptable risk management strategy? If the 30 year forward rate were known to fluctuate randomly around a well defined mean \bar{f} in such a way that the expected gain/(loss) in any given reporting period was always zero then this might be acceptable. Economic Capital would be required to protect against an unexpected loss in any given period and providing for the cost of that capital would require a margin $\pi\Delta$ in the discount rate. Here π is the company's cost

of capital and Δ is the size of the spread loss at the company's VaR level. Plausible estimates for these parameters $\pi \approx .04$, $\Delta \approx .025$ suggest the required margin is on the order of 10bp so a final discount rate might be $5.17\% = 5.27\% - 0.10\%$. This valuation would be supplemented by holding economic capital equal to 2.5% of the long liability value. Problem solved.

What if the 30 year forward rate did not behave so nicely? The observed forward rate of 4.63% at Sept. 30, 2008 might take several years before it reverts back to its assumed mean of 5.27%. There is also the risk that it goes lower before it mean reverts. If that is the case then the simple model has a bias since we would expect to take losses for several years. The model therefore needs to be modified to remove the bias and account for the risk of further random movements. This is done in Appendix 1 where we show that correcting the bias creates almost as many problems as it solves with the result that much larger margins are required. However, correcting the bias and building in a cost of capital for the unhedged risk does not change the monopole nature of the static hedge strategy.

We'll explain the source of the new problems in the corrected model after discussing the dipole model.

Finally, we note that there is also an issue of parameter risk in this model. However we determined that the right expected forward rate should be 5.27% it is possible that new information could arrive that causes us to reevaluate this assumption. An appropriate response might be to shock this important parameter down to say 5.00% and hold enough economic capital to cover the shock. Covering the cost of that capital would require additional margins. This problem is solvable but beyond the scope of the current paper.

The Simple Dipole

If the simplest idea was a failure then why don't we try the next simplest idea of using the last observed forward rate $f_{30} = 4.63\%$ and using that to discount cash flows beyond year 30? This method has been widely used by many researchers and, at first glance, appears to solve the bias problem in the monopole model just discussed while eliminating the parameter risk issue.

As before we start by looking for a static hedge portfolio and then see what roll forward issues we may have.

Let $Z = Z_{30}$ be the value of the longest 0 coupon bond available and let \tilde{Z} be a shorter bond used to determine the forward rate f that will be used for discounting. Thus $\tilde{Z} = (1 + f)^\Delta Z$ where Δ is the time difference between Z and \tilde{Z} . In our numerical examples $\Delta = .25$ years.

In this notation the model value of the long liability is

$$\begin{aligned}
V &= ZK, \\
&= Z \sum_{j>30} a_j (1+f)^{-(j-30)}, \\
&= Z \sum_{j>30} a_j \left(\frac{\tilde{Z}}{Z} \right)^{\frac{-(j-30)}{\Delta}}.
\end{aligned}$$

A portfolio of assets $A = bZ + \tilde{b}\tilde{Z}$ which has the same value and instantaneous sensitivities to market movements as the liability is given by setting


$$\begin{aligned}
b &= \frac{\partial V}{\partial Z} = K \left(1 + \frac{D}{\Delta} \right), \\
\tilde{b} &= \frac{\partial V}{\partial \tilde{Z}} = -K \frac{D}{\Delta} (1+f)^{-\Delta}.
\end{aligned}$$

Here $D = \left[\sum_{j>30} a_j (j-30)(1+f)^{-(j-30)} \right] / K$ is the duration of the long liability as seen from the 30 year point.

The static hedge consists of investing the value K in a monopole together with a strategy that we will call a dipole in this paper. The dipole consists of taking a short position $KD(1+f)^{-\Delta} / \Delta$ in the \tilde{Z} bond and investing the proceeds in the Z bond. This strategy is necessary to get the additional duration needed to hedge the liability which is now much more sensitive to a change in the observed yield curve.

The chart below shows the size of the dipole components relative to the monopole component and the underlying liability cash flows in our example.

Long Liability Dipole Static Hedge



If the reader thought the monopoly strategy was unrealistic then this is even more so. Ignoring this issue again we examine the roll forward properties of this new strategy.

As was the case in the monopoly model the dipole model is not self financing. The details of the analysis are more complex than they were for the monopoly strategy so they are presented in Appendix 1. However, the end results are fairly intuitive. To first order in small quantities there are two sources of bias in the dipole model

$$Gain = gDVdt - (1/2)(C + D\Delta)Vdf^2.$$


The first term is driven by the slope g of the forward curve at the 30 year point. This is the cost or benefit of carry that arises from borrowing at one point on the curve and investing at a neighboring point. The second term is a convexity cost that arises because the dipole strategy matches the first order sensitivities of the liability but not the second order sensitivities. The liability is more convex than the assumed assets and the result is a convexity bias.

In addition to the bias terms described above there is an unhedged risk that arises from random movements in the quantity g .

Appendix 1 goes through a tuning exercise to remove this bias from the valuation and to estimate the capital required to deal with the g risk. . The conclusion reached is that the first bias term is immaterial while the second term creates significant costs that must be priced into an extended yield curve.

The chart below shows the end result of the fine tuning process for both the monopoly and dipole models.

Yield Curve Extension Models: Fwd Rates



The key simplifying assumptions used here are:

- The last observed continuously compounded forward rate f satisfies a mean reverting process $df = \alpha(\bar{f} - f)dt + \sigma dz$.
- The last 0 coupon bond Z behaves as it would in a Vasicek type model $dZ = -30Z[\mu dt + s dw]$ where s is a constant and $dwdz = \rho dt$. The constant ρ is essentially the correlation between the long forward rate and the long spot rate. This assumption is only used for the corrected monopole model.
- The slope of the forward curve $g = df / ds$ at the end point also satisfies a mean reverting process of the form $dg = \beta(\bar{g} - g)dt + v d\tilde{w}$. This assumption is only used for the corrected dipole model.

In terms of this notation the extrapolated forward rates for the corrected dipole model are approximately given by the formula

$$\delta(f, g, u) \approx f + g \frac{1 - e^{-\beta u}}{\beta} - \frac{1}{2} \Delta \sigma^2 u - \frac{1}{2} \sigma^2 u^2.$$

We can think of this result as the starting simple assumption $\delta(f, g, u) = f$ together with an adjustment to correct for the bias described earlier. Appendix 1 has all the details and an exact expression for $\delta(f, g, u)$ under the stated assumptions.

The end result for the corrected monopole model came as a surprise to the author. There is a large spread of just under 200bp between the expected future forward rate and the ultimate forward rate in the extrapolated yield curve. There are two reasons for this

1. Correcting the bias in the base monopole strategy basically requires a term that depends on the current forward rate f . The resulting bias correction then has an element of dipole risk in it. By sticking with a monopole investment strategy we are essentially going naked on this risk so we have to put up some economic capital. The cost of that economic capital explains roughly a third of the spread.
2. The majority of the spread margin is related to the correlation between movements in the long bond and the K factor. In the bias corrected model $V = ZK$ the quantity K must depend on the current forward rate f . The roll forward of the liability then satisfies a relation of the form $dV = KdZ + ZdK + dKdZ$. The monopole investment strategy hedges out the 'Z' risk in first term but not the third. This gives rise to a new bias that must be priced into the yield curve.

The final formula for the extrapolated forward rates in the corrected monopole model is

$$\delta(f, u) = fe^{-\alpha u} + (\bar{f} - \frac{30s\rho\sigma + \pi D\lambda\sigma}{\alpha})(1 - e^{-\alpha u}) - \frac{(1 + \pi D\lambda^2)\sigma^2}{2\alpha^2}(1 - e^{-\alpha u})^2.$$

This formula assumes the cost of capital is π and if $V(t, Z, f) = ZK(t, f)$ we hold economic capital in the amount $DZ[K(t, f - \lambda\sigma) - K(t, f)]$ for the unhedged dipole risk.

The table of values below shows what happens when we apply each of these approaches to value the example long liability. These are values for the entire liability but only the portion beyond 30 years changes when we change valuation methods. The value of the first 30 years is 981.

Table 2: Numerical Examples

	Value	Duration	Capital	IRR
Simple Monopole	1,796	34.1	20	4.84%
Simple Dipole	1,832	37.4	N/A	4.78%
Bias Corrected Models				
Corrected Monopole	1,873	34.0	74	4.72%
Corrected Dipole	1,847	37.5	0	4.76%

The table shows the total value after an arbitrary scaling, the duration in years of the static hedge portfolio, an economic capital number where available, and the interest rate that discounts all 60 years of cash flow to the model value.

A capital number has been estimated for the corrected dipole model but it is immaterial.

We note that, despite large differences in method, the total values are actually quite close together. This is not an artificial result due to a poor choice of example.

It should be clear that the models presented here are the tip of a very large iceberg. There are many possible variations on the theme introduced here. However the author believes that all such approaches will end up relying on what are unnecessarily extreme investment strategies. To actually justify using such a model an insurer would have to hold additional economic capital for the mismatch gains and losses that occur when the actual assets held fail to behave like the static hedge portfolios.

The Static Control Model

As before we start with a very simple idea and then evolve the process to deal with the issues we encounter. The initial idea is to use nothing but a simple total return vehicle as the only asset class available to back the liability.

The Pure Total Return Model

In this section we will assume that the only available total return vehicle can be modeled by a standard lognormal process $S(t) = S(0) \exp[\mu t + \sigma Z(t)]$ where $Z(t)$ is a standard Wiener process. As noted in the introduction the parameters have been chose so that $E[S(t)] = S(0)(1.08)^t$ and $\sigma = 16\%$. These are clearly important parameters subject to parameter risk considerations.

Given this assumption we generate a large number $N = 5,000$ of total return scenarios and then calculate N present values PV_α , $\alpha = 1, \dots, N$ where each present value is the amount of money that must be invested in the total return vehicle to satisfy the liability on that scenario.

$$PV_\alpha = \sum_{j>0} a_j \exp[-j\mu - \sigma Z_\alpha(j)]$$

In order to get a single number for the liability we apply a coherent risk measure, such the Conditional Tail Expectation or CTE($x\%$), to the distribution of outcomes. CTE($x\%$) is defined as the average of the worst (100- x)% of outcomes.

In practice one sorts the results to get the order statistics $PV_{(1)} \geq PV_{(2)} \geq \dots$ and then estimates the risk measure by

$$\begin{aligned} L(a_1, \dots, a_{60}) &= \frac{1}{(1-x\%)N} \sum_{\alpha=1}^{(1-x\%)N} PV_{(\alpha)}, \\ &= \frac{1}{(1-x\%)N} \sum_{\alpha=1}^{(1-x\%)N} \sum_{j>0} a_j \exp[-j\mu - \sigma Z_{(\alpha)}(j)]. \end{aligned}$$

The function $L(a_1, \dots, a_{60})$ is homogeneous of degree 1 in the sense that if $\lambda > 0$ then $L(\lambda a_1, \dots, \lambda a_{60}) = \lambda L(a_1, \dots, a_{60})$. Given any function with this property we can differentiate with respect to the parameter λ and then set $\lambda = 1$. The result is called the Euler allocation of the value back to the individual cash flows

$$\sum_{j=1}^{60} \frac{\partial L}{\partial a_j} a_j = L(a_1, \dots, a_{60}).$$

We conclude that the partial derivatives $\partial L / \partial a_j$ can be thought of as discount factors. If we knew what all of these discount factors were we could build a yield curve that we will call the Marginal Cost Yield Curve (MCYC). This yield curve has two useful properties

1. It discounts the cash flows to the desired risk measure.
2. It tells us how the risk measure would change if we perturbed the cash flows

$$a_j \rightarrow a_j + \Delta a_j \text{ in some way since } L(a_1 + \Delta a_1, \dots, a_{60} + \Delta a_{60}) \approx L(a_1, \dots, a_{60}) + \sum_j \frac{\partial L}{\partial a_j} \Delta a_j.$$

A theoretical result due to the German risk manager D. Tasche⁴ (2000) tells us that we can actually estimate the first partial derivatives from the simple expression

$$\frac{\partial L}{\partial a_j} = \frac{1}{(1 - x\%)^N} \sum_{\alpha=1}^{(1-x\%)^N} \exp[-j\mu - \sigma Z_{(\alpha)}(j)].$$

The resulting MCYC is a function of both the cash flows and the risk measure.

The table below shows the results of applying these ideas to the long liability cash flow for a range of CTE risk measures.

Table 3: Simple Total Return Approach

CTE	Static Hedge	Duration	Total Return	Total Return Hurdle	Total Liability	IRR	VaR Level	Std Err
0%			1,522	5.23%	1,522	5.30%	64%	28
20%			1,911	5.21%	1,911	4.66%	72%	34
40%			2,394	5.06%	2,394	4.01%	80%	43
60%			3,092	4.77%	3,092	3.25%	87%	60
80%			4,366	4.33%	4,366	2.21%	94%	102
90%			5,829	3.89%	5,829	1.33%	97%	173

Since the process used to estimate these values is statistical there is sampling error in each value. The right hand column reports an estimated standard error in the CTE level using the method of Manistre and Hancock⁵ (2005).

⁴ For a rigorous derivation of this result see Tasche, D., "Risk Contributions and Performance Measurement" Preprint (2000).

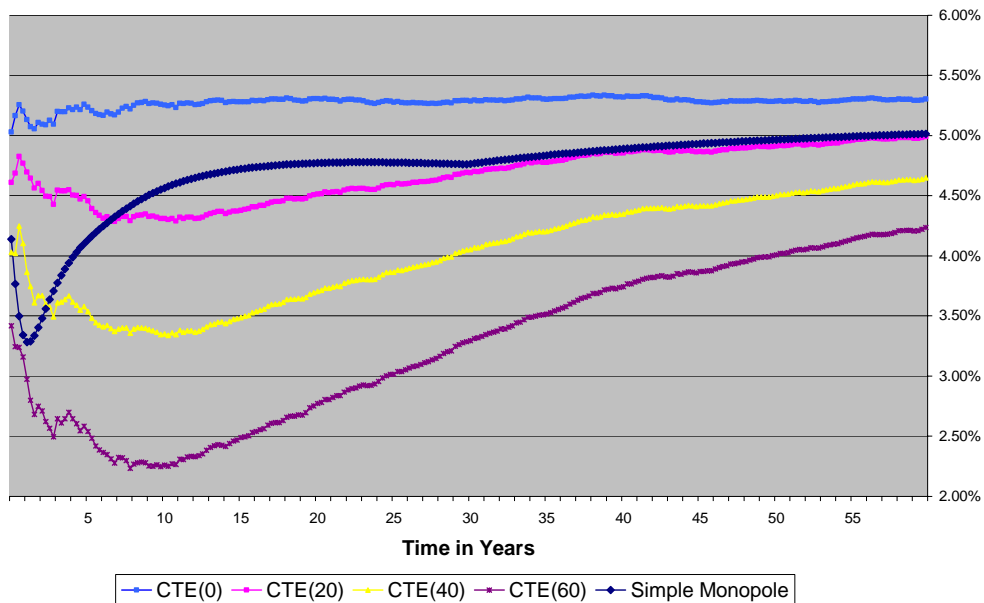
⁵ Manistre, B.J. & Hancock, G.H., *Variance of the CTE Estimator*, North American Actuarial Journal (May 2005).

The next column over reports the Value at Risk level indicating what percentage of the simulated results are lower than the estimated CTE level. The IRR is simply the rate that discounts all cash flows to CTE value.

The Total Return Hurdle is an estimate of the rate we would have to earn on any assets backing the liability in order to avoid a loss over the next reporting period. Except for the CTE(0%) result all Total Return Hurdles are higher than the corresponding IRR.

In order to explain the Total Return Hurdle we first need to look at the Marginal Cost Yield Curve for each line in the table above. The chart below does this for the entire 60 year period.

Simple Total Return: Annual MCYC Spot Yields



Were it not for sampling error, the MCYC for the CTE(0) case would be a flat 5.27%. This is because CTE(0) corresponds to taking the mean of the distribution and the expected discount factor is, for our choice of parameters

$$E[\exp(-\mu - \sigma z)] = \exp\left[-\mu + \frac{1}{2} \sigma^2\right],$$

$$= \frac{1}{1.0527}.$$

It is clear from the chart that all of the marginal cost yield curves are heading toward this value but the time scale required to get there depends on the risk measure. The model is clearly telling us that 5.27% is where the long end of the yield curve should be. Since the spot yield curves are almost parallel after 40 years it appears as if the MCYC forward rates have converged to 5.27% by about year 40.

The earlier parts of the MCYC are telling us what “bad” scenarios look like from a CTE perspective. The curves tell us that the worst that can happen is for bad returns to occur early on.

This is not entirely intuitive since the early cash flow in the example are negative so early bad return might indicate an opportunity to buy equity on the cheap.

We can now explain why the Total Return Hurdles H behave the way they do. Over a short time period we can model the passage of time as a change in cash flow $a_j \rightarrow a_{j+1} = a_j + (a_{j+1} - a_j)$. Assuming this change is small enough that we can use first order analysis we can estimate the change in value as $\Delta L = \sum_j \frac{\partial L}{\partial a_j} (a_{j+1} - a_j) = HL$. This is essentially equivalent to assuming that

the MCYC does not change, to first order, as time moves forward. The present value of each cash flow therefore accumulates at the corresponding MCYC forward rate. The aggregate total return hurdle is therefore a present value weighted average of all the MCYC forward rates. Since the early cash flows, with low forward rates, are negative this leverages up the average.

The chart also shows the spot yields implied by the simple monopole model introduced at the beginning of this paper. Since the MCYC spot yields do not agree with the observed market spot rates there is an opportunity to reduce the cost of meeting the liability by using a mixture of bonds and the total return strategy.

The Simple Bond Strategy

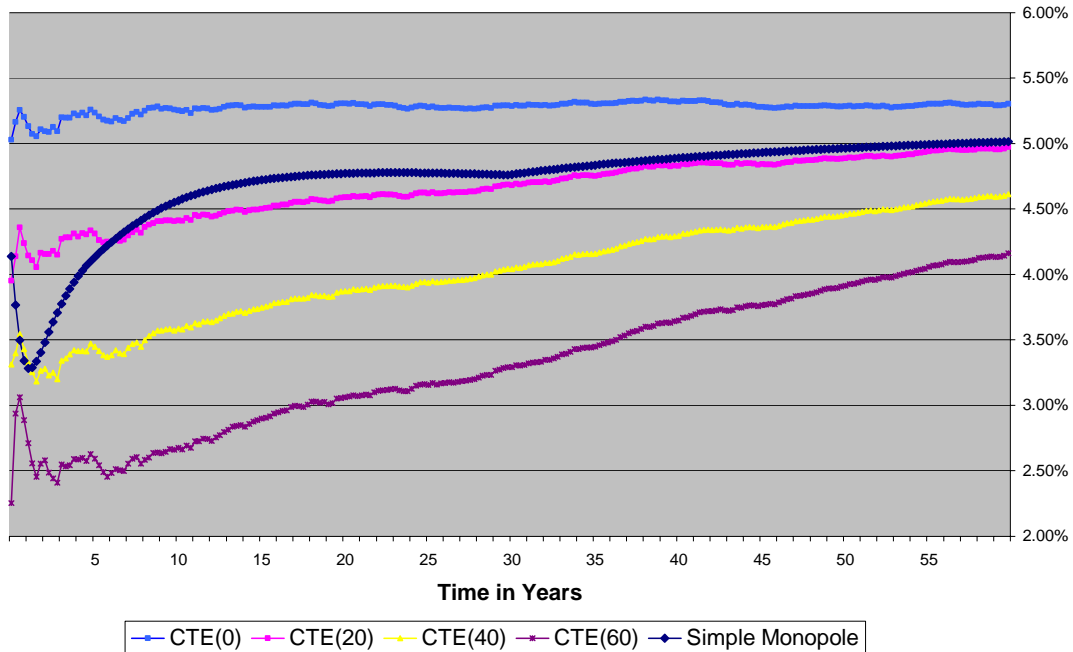
An intuitive next step is to use bonds to match the first 30 years of cash flow while using the total return approach only for the cash flow beyond the horizon. The table below shows those results together with the results from the previous section.

Table 4: Simple Bond Approach

Quarterly Time Step		Static Hedge	Dur'n	Total Return	Total Return Hurdle	Total Liability	IRR	VaR Level	Std Err
Strategy	No Bonds	0%		1,522	5.23%	1,522	5.30%	64%	28
		20%		1,911	5.21%	1,911	4.66%	72%	34
		40%		2,394	5.06%	2,394	4.01%	80%	43
		60%		3,092	4.77%	3,092	3.25%	87%	60
		80%		4,366	4.33%	4,366	2.21%	94%	102
		90%		5,829	3.89%	5,829	1.33%	97%	173
Simple Bonds (match first 30 years)	0%	981	37.6	696	5.21%	1,677	5.03%	67%	12
	20%	981	37.6	836	5.08%	1,817	4.81%	74%	15
	40%	981	37.6	1,020	4.90%	2,001	4.53%	81%	19
	60%	981	37.6	1,298	4.60%	2,279	4.15%	87%	27
	80%	981	37.6	1,833	4.12%	2,814	3.53%	94%	47
	90%	981	37.6	2,469	3.74%	3,450	2.92%	97%	83

There is a substantial reduction in values at all CTE levels, which was expected. Furthermore the total return hurdle rates have dropped a bit. The next chart shows how the MCYC has changed

Simple Bonds: Annual MCYC Spot Yields



The CTE(0) curve has not changed but all of the other curves reflect the new lower risk profile. The range of values has compressed considerably. The main point to be made here is that there is still room for improvement by modifying the bond strategy. For each CTE level (other than 0) we can still get a lower value by shorting bonds where the MCYC is below the swap curve or doing the reverse if the MCYC is above the swap curve.

The Optimal Bond Strategy

There are a number of technical challenges associated with estimating the optimal bond strategy but it can be done. Some additional detail is in Appendix 2. Here we will focus only on results. Table 5 below shows the numbers that come out of the optimization process.

Table 5: The Static Control Model

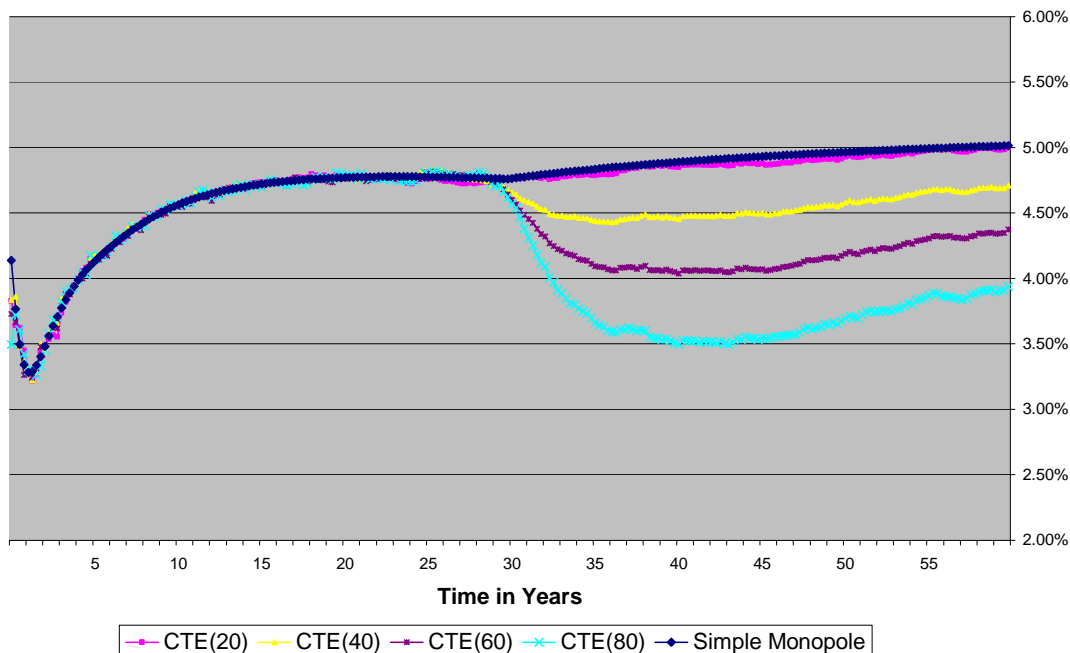
Quarterly Time Step									
Strategy	CTE	Static Hedge	Dur'n	Total Return	Total Return Hurdle	Total Liability	IRR	VaR Level	Std Err
Simple Bonds (match first 30 years)	0%	981	37.6	696	5.21%	1,677	5.03%	67%	12
	20%	981	37.6	836	5.08%	1,817	4.81%	74%	15
	40%	981	37.6	1,020	4.90%	2,001	4.53%	81%	19
	60%	981	37.6	1,298	4.60%	2,279	4.15%	87%	27
	80%	981	37.6	1,833	4.12%	2,814	3.53%	94%	47
	90%	981	37.6	2,469	3.74%	3,450	2.92%	97%	83
Simple Monopole	N/A	1,796	34.0		N/A	1,796	4.84%		
Simple Dipole	N/A	1,832	37.5		N/A	1,832	4.78%		
Corrected Monopole	N/A	1,873	34.1	-	N/A	1,873	4.72%		
Corrected Dipole	N/A	1,847	37.4	-	N/A	1,847	4.76%		
Static Control	20%	1,016	40.7	792	4.78%	1,807	4.82%	75%	13
	40%	1,509	35.3	389	2.35%	1,898	4.68%	82%	10
	60%	1,683	34.5	321	-2.17%	2,004	4.53%	88%	10
	80%	1,883	34.1	276	-10.90%	2,159	4.31%	94%	12
	90%	2,108	32.5	193	-27.62%	2,301	4.13%	97%	16

Not surprisingly, values have dropped again and the range of results has compressed even further. The total return hurdles have also come down. If one used the total return hurdle as a criteria for setting the CTE level one could easily justify a number in the CTE(20) to CTE(40) range.

We don't show a value for CTE(0) since this optimization problem is not well defined. However this numerical example, and logical reasoning, suggest that the CTE(0) limit value is equal to the simple monopole value. A limiting risk management strategy may not exist.

The next chart shows that the marginal cost yield curves now agree with the swap curve for the first 30 years. This is one of several ways in which the method is market consistent.

Static Control: Annual MCYC Spot Yields



In Appendix 2 we show that the method described above has three important theoretical properties.

1. Market Consistency – the model reproduces the market information that has been input. In this example the values of 0 coupon bonds for 30 years.
2. Existence of a Static Hedge Portfolio – the solution to the optimization problem is the static hedge portfolio from a market shock perspective.
3. Convexity Margin – liabilities which cannot be expressed as a static linear combination of the available instruments have a convexity margin which gets released as time moves forward.

In addition to these very mathematical properties we also have

4. Method is Sub-Additive – using a coherent risk measure on the present value distribution means that the method is sub additive rather than additive. However this only applies to the portion of the risk that cannot be hedged so market consistency is not violated. One implication is that the static control method must be applied at a product portfolio level. Once the total value is determined, tools like the MCYC can be used to allocate that value back to individual contracts.
5. The extrapolated forward rate in the optimal MCYC is not continuous. This is clear from the spot rate chart above. The extrapolated forward rate takes an immediate drop at time 30+0 and then starts grading back to 5.27%. While some may find this property to be undesirable it one of the ways in which we avoid dealing with the dipole issue that arose in the early models.

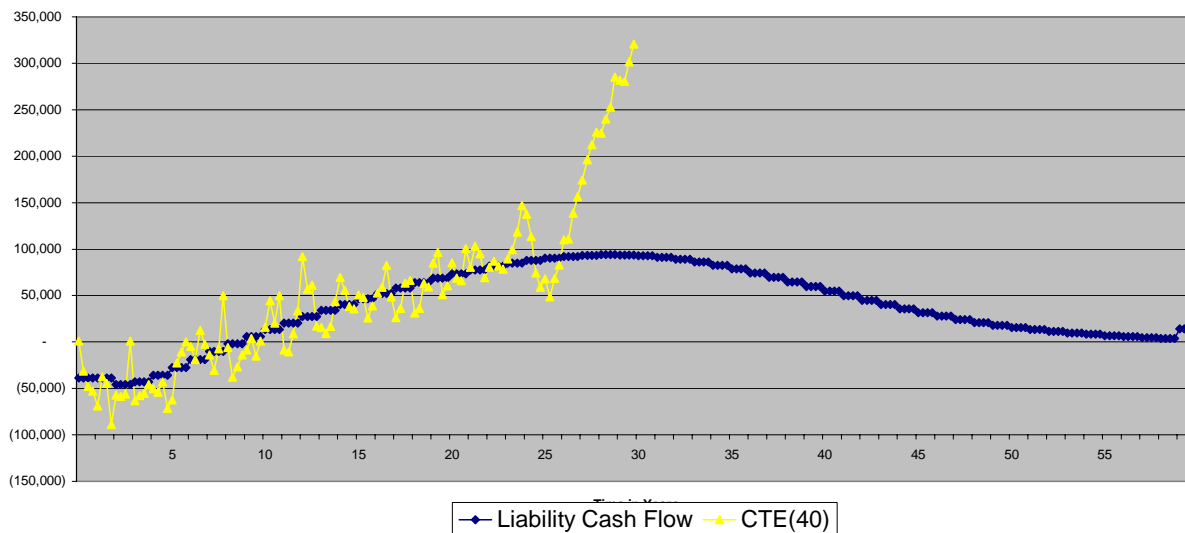
We discuss some of these points a bit more here.

Market Consistency: The fact that the model reproduces the input swap curve is an example of market consistency. One might argue however that we did not use market volatility in the total return model and that if we did the estimated liability values would go up. However, suppose there was an observable market for puts and calls on the total return vehicle. If we made those instruments available to the model two things would happen 1) the optimization process would mean that the model prices for the puts and calls would be equal to the observed market price and 2) the value of the long liability would go down because the universe of available investment strategies over which the optimization takes place got larger.

The implication is that “simple” models which use a relatively small number of market instruments have an element of conservatism that can be removed by making the model more sophisticated.

Static Hedge: The fact that the solution to the optimization problem is also a static hedge is an extremely useful result. The mathematics behind this statement are in Appendix 2. The chart below shows what the optimal bond portfolio looks like for the CTE(40) risk measure as compared to the input liability cash flows.

Static Control CTE(40) Optimal Bonds



The right way to think of this static hedge is that is the sum of the simple bond portfolio that matches the liability cash flows for the first 30 years plus a second portfolio whose purpose is to mitigate the total return risk. This second portfolio appears to have a monopole component and a number of seemingly random fluctuations during the first 30 years.

The monopole component is fairly intuitive. What is important here is that the size of the monopole is much smaller than it was in the early monopole model. This is a much more realistic static hedge portfolio.


From a theoretical view point the fluctuations are an important element of the total return risk mitigation process. A succession of over/under cash flow mismatches means that we are never, in the model, consistently buying or selling the total return vehicle.

From a more practical perspective the fluctuations can almost be ignored. Since the static hedge changes over time it should not be thought of as a cash flow matching target. Rather, the right way to use the static hedge portfolio is as a tool to estimate quantities like the sensitivity of the liability to principal component shocks in the yield curve.

The detailed structure of the fluctuations is also sensitive to things like the size of the time step and the particular sample of total return scenarios. The above example was calculated using a quarterly time step. If we increase the time step to a year the fluctuations get bigger but the more important information, such as principal component shock sensitivities and values, does not change materially.

The next chart shows how the static hedge changes as we vary the CTE level. As we become more risk averse the size of the monopole grows and the fluctuations become more pronounced. At the CTE(20) level the total return vehicle is so much more attractive than the available bonds that the long bond is being shorted. At the other extreme we see a dipole structure starting to emerge at the CTE(80) level.

Static Control Optimal Bonds



In the author’s opinion all of these static hedge portfolios are more useful than the simple monopole introduced earlier.

Table 6 below shows how the CTE(40) values change when the base yield curve is shocked up and down by 50bp.

Table 6 : Yield Curve Shock Analysis

Yield Curve	CTE	Static Hedge	Duration	Total Return	Total Return Hurdle	Total Liability	Level Val'n %	VaR Level	Std Err
+50 bp	40%	1,360	36.8	276	-0.72%	1,636	5.10%	82%	8
Base	40%	1,509	35.3	389	2.35%	1,898	4.68%	82%	10
- 50 bp	40%	1,477	34.7	697	4.47%	2,174	4.29%	82%	14
Static Hedge \$Dur'n				53,286	Simple \$Dur'n Calc		53,781		

There are two points we would like to make here. The first point is that the static hedge is working. As a simple example consider the dollar duration of the static hedge which is $35.3 \times 1,509 = 53,286$. This value is within 1% of the dollar duration estimated by looking at the shocked values $(2,174 - 1,636)/.01 = 53,781$.

The aggregate duration of the total liability is $(53,781/1,898) = 28.3$ years. This duration is shorter than the result produced by any of the other models considered here. The reason this happens is because the long end of the yield curve is essentially fixed at 5.27%. This is a major reason why the static control model makes the A/L M process more manageable.

The second point is that the mix of bonds and total return vehicle can change considerably as the yield curve moves. In particular, as the yield curve drops the bonds look less attractive so the model uses more of the total return instrument. Unfortunately, the total return hurdle is also going up. This is the element of interest rate risk that has not been hedged.

If a company using this model encountered a prolonged period of low interest rates (a “Japan” scenario) they would ultimately find the total return requirements of the model to be unmanageable. The only solution to this problem would be to adjust one or more of the key parameters ($\mu, \sigma, x\%$) so that the total return issue became manageable again. The shock from the resulting increase in liability value would be a risk for which an appropriate amount of economic capital should be held. For example, if the company decided to manage off the CTE(40) model then it might be appropriate to hold enough economic capital to cover a jump to the CTE(60) level.

Convexity Margin: As a partial offset to the cost of holding economic capital the inherent conservatism of the static control model gets released over time as a series of convexity gains. This conservatism arises because we are optimizing over a set of static (buy and hold) investment strategies rather than a set of dynamic strategies. This turns out to imply the mathematical fact, see appendix 2, that the static hedge portfolio is always more convex than the model liability. Quantifying the convexity margin issue is possible but we do not do so here due to time and space limitations.

For an alternative approach to long liabilities that does not build in a convexity margin see Platen⁶ who argues that long liabilities can be priced using a concept he calls the Growth Optimal Portfolio.

Discontinuous Forward Rates: As mentioned earlier the MCYC extrapolated forward rates are not continuous. If the risk measure is conservative enough it is possible to get a negative forward rate. If this occurred in the “tradable” part of the yield curve it would be a serious problem for the model. However, we can’t actually trade beyond the 30 year horizon so this is not a problem.

A reasonable way to think about this result is that the bulk of the unhedged interest rate risk is in the middle of the yield curve. Short cash flows (<30 years) can be matched with market instruments and the longest cash flows have the benefit of time diversification. The MCYC is telling us that the risk is in the middle.

Summary and Conclusion

This paper has looked at two very different approaches to issue of valuing and managing long liabilities. The first half of the paper started with some very simple yield curve extension ideas and then used risk management concepts to fine tune those ideas. A formal theory of yield curve extension was built that corrected for any bias in the original models. Unfortunately all of these models relied on static hedge strategies that, in this author’s opinion, are unrealistic.

The second half of the paper described a version of the static control model which was developed to allow the use of a total return vehicle in the risk management of a long liability. After choosing a coherent risk measure we were able to put a value on the liability and come up with a matching A/L M strategy that consists of a static hedge portfolio and a total return component.

A key advantage of the static control model is that the implied hedging strategies are more realistic.

A disadvantage of the static control model is that it is technically more complex.

The static control model does have bias. If one were actually invested in the model’s hedge portfolios there would be three sources of gain or loss as time moves forward.

1. There is a gain or loss on the total return component of the investment strategy equal to the difference between the actual returns and the total return hurdle on the assets backing the total return piece.
2. There is an interest rate convexity gain on the fixed income component.
3. Any changes to the key parameters would result in a gain or loss. Some amount of economic capital should be held for this parameter risk.

⁶ Platen, E. “Investments for the Short and Long Run” available at http://www.qfrc.uts.edu.au/research/research_papers/rp163.pdf

On balance, the author believes that if a company has a material exposure to long liability issues then an investment in the tools and ideas needed to make the static control model work are worth the effort.

Appendix 1 – Mathematics of the Yield Curve Extension Approach

Corrected Monopole Model

Assume we invest in a monopole A/L M strategy in the amount $A = KZ$ where Z is the value of the n period 0 coupon bond (e.g. $n = 30$ years) and K is the value of cash flows beyond the n period horizon. Let f be the instantaneous forward rate at time t , then the first order roll forward for the monopole asset is

$$dA = fAdt + KdZ .$$

One way to derive this is to note that at time $t + dt$ the n period 0 coupon bond has become shorter so its value at that time is

$$\exp[(f + df)dt](Z + dZ)$$

where $f + df$ is the new forward rate and $Z + dZ$ is the new n year 0 coupon bond. Expanding this out using the rules of stochastic calculus we get

$$\begin{aligned} \exp[(f + df)dt](Z + dZ) &= (1 + fdt)(Z + dZ) + \dots \\ &= Z + fZdt + dZ + \dots \end{aligned}$$

To analyze the liability we assume $K = K(t, f)$ is the present value of cash flows beyond the horizon discounted using an extrapolated yield curve that depends on f . Since K depends on f the roll forward of the liability is

$$dL = ZdK + KdZ + dZdK .$$

The net surplus now satisfies $dS = dA - dL = fAdt - ZdK - dZdK$. We see that the monopole strategy has hedged out the 'Z' risk but leaves us naked on any risk associated with random movements in the variable f since $K = K(t, f)$. This is the risk we must hold economic capital for.

In order to make further progress we need to make some assumptions about how the random variables f, Z move. The following assumptions were chosen to be simple enough that we can get closed form solutions for the yield curve extension problem. This simplicity does not affect the final conclusions.

For the forward rate we will assume a mean reverting process of the form

$$df = \alpha(\bar{f} - f)dt + \sigma dw$$

where α, \bar{f}, σ are constants and dw is standard Brownian motion.

For the zero coupon bond we will assume that we can write

$$dZ = -n\Delta Z[\mu dt + sdw']$$

where n is the number of periods of length Δ in the observed yield curve, s is a constant and $dwdw' = \rho dt$. We won't need to make any assumptions about μ because this is the risk that has been hedged out. The quantity ρ is effectively the correlation between the n period spot rate and n period forward rate. It will play an important role in what follows.

Using Ito's lemma we can write the roll forward for surplus as

$$\begin{aligned} dS &= dA - dL = fAdt - ZdK - dZdK, \\ &= fKZdt - Z\left\{\frac{\partial K}{\partial t} + [\alpha(\bar{f} - f) - n\Delta s\rho\sigma]\frac{\partial K}{\partial f} + \frac{\sigma^2}{2}\frac{\partial^2 K}{\partial f^2}\right\}dt - \sigma\frac{\partial K}{\partial f}dw. \end{aligned}$$

Note that the term $dZdK$ in the surplus roll forward ends up adjusting the effective expected future forward rate term.

Since we are going naked on the 'f' risk, we need hold enough capital that we can withstand an unexpected movement in the long forward rate. A reasonable amount of capital to hold is then

$$EC = DZ[K(t, f - \lambda\sigma) - K(t, f)] \approx DZ[-\lambda\sigma\frac{\partial K}{\partial f} + \frac{(\lambda\sigma)^2}{2}\frac{\partial^2 K}{\partial f^2}].$$

Here λ is the number of standard deviations needed to get the required VaR level and D is a factor which reduces the amount of required capital to recognize any appropriate diversification benefits.

Let π be the cost of capital. The pricing equation for this risk is then the statement that the expected return is equal to the cost of capital.

$$E[dS] = \pi ECdt,$$

or

$$fKZdt - Z\left\{\frac{\partial K}{\partial t} + [\alpha(\bar{f} - f) - n\Delta s\rho\sigma]\frac{\partial K}{\partial f} + \frac{\sigma^2}{2}\frac{\partial^2 K}{\partial f^2}\right\}dt = \pi DZ[-\lambda\sigma\frac{\partial K}{\partial f} + \frac{(\lambda\sigma)^2}{2}\frac{\partial^2 K}{\partial f^2}]dt.$$

Simplifying this expression we find that the Z factor drops out and we are left with a partial differential equation for $K = K(t, f)$ that is formally identical to that derived by Vasicek⁷ in 1974. The result is

$$\frac{\partial K}{\partial t} + [\alpha(\bar{f} - f) - n\Delta s\rho\sigma - \pi D\lambda\sigma]\frac{\partial K}{\partial f} + \frac{(1 + \lambda^2)\sigma^2}{2}\frac{\partial^2 K}{\partial f^2} = fK.$$

One more manipulation puts it in the form

⁷ There are now many good books on interest rate modeling. One source is Hull, J.C. "Options Futures and Other Derivatives" fifth edition. Prentice Hall 2003.

$$\frac{\partial K}{\partial t} + \alpha[\bar{f} - \frac{n\Delta s\rho\sigma + \pi D\lambda\sigma}{\alpha} - f]\frac{\partial K}{\partial f} + \frac{(1 + \pi D\lambda^2)\sigma^2}{2}\frac{\partial^2 K}{\partial f^2} = fK.$$

This form of the pricing equation tells that the yield curve extension is priced as if the forward rate f mean reverts to a risk adjusted target $\theta = \bar{f} - [n\Delta s\rho\sigma + \pi D\lambda\sigma]/\alpha$ with a risk-adjusted volatility $\tilde{\sigma} = \sqrt{(1 + \pi D\lambda^2)\sigma}$. The adjustments coming from the cost of capital are fairly intuitive but the correlation adjustment will be explained more fully later.

If we are valuing a cash flow at time $n\Delta + w$ the formal solution to the above equation can be written as

$$K(t, f, w) = \exp[-\int_0^w \delta(f, u) du],$$

where the extrapolated force of interest is given by

$$\delta(f, w) = fe^{-\alpha w} + \theta(1 - e^{-\alpha w}) - \frac{\tilde{\sigma}^2}{2\alpha^2}(1 - e^{-\alpha w})^2.$$

The first two terms represent a simple grading from the current forward rate f to the long term target θ while the last term builds in a convexity margin. As $w \rightarrow \infty$ we find the ultimate long forward rate is

$$\begin{aligned} \delta(\infty) &= \theta - \frac{\tilde{\sigma}^2}{2\alpha^2}, \\ &= \bar{f} - \frac{n\Delta s\rho\sigma + \pi D\lambda\sigma}{\alpha} - \frac{(1 + \pi D\lambda^2)\sigma^2}{2\alpha^2}, \\ &= \bar{f} - \frac{\sigma^2}{2\alpha^2} - \frac{n\Delta s\rho\sigma}{\alpha} - \pi D[\frac{\lambda\sigma}{\alpha}(1 + \frac{\lambda\sigma}{2\alpha})]. \end{aligned}$$

We will call the three adjustments to the expected forward rate the convexity spread, the correlation spread and the cost of capital spread respectively.

In order to get a sense of the relative importance of these issues we use some ball park estimates for the parameters to estimate each spread component,

For the forward rate process let's assume $\bar{f} = 5.15\%$, $\sigma = 1\%$ and $\alpha = 20\%$. This leads to a convexity spread of $\sigma^2/(2\alpha^2) = 12.5bp$ which is fairly modest.

Next we estimate the cost of capital spread assuming $\pi = 5\%$, $D = 80\%$, $\lambda = 2.8$. The resulting spread estimate is $\pi D[\frac{\lambda\sigma}{\alpha}(1 + \frac{\lambda\sigma}{2\alpha})] \approx 60bp$. This is large enough to get noticed but reasonable given the risk that is going unhedged.

Finally, we estimate the correlation spread assuming $n\Delta = 30 \text{ yrs}$, $\rho = 75\%$, $s = 1\%$. The estimated spread is $(n\Delta s \rho \sigma) / \alpha \approx 113 \text{ bp}$.

The final estimated long forward rate for the extrapolated yield curve is then about 3.30% which is significantly lower than what most people would find intuitive. The correlation spread is clearly an issue.

Since the correlation spread is almost twice the size of the other two issues put together it requires some discussion. Mathematically this spread arises from the $dZdK$ term in the liability roll forward. The problem is that the monopoly investment strategy assumes the liability is independent of the forward rate and so has no offsetting term. The resulting bias must then be priced into the liability. The only way to correct this issue is to take on some dipole exposure in the investment strategy.

We conclude that there is a hidden cost in the monopoly approach which makes this approach undesirable. We therefore explore the dipole model to see if there are any surprises there.

The Corrected Dipole Model

As before we assume cash flows beyond the yield curve horizon have been discounted to the horizon date to get a value K . We will assume this amount is fully invested in a monopoly as before but we will also assume a position in a dipole.

A dipole strategy consists of selling the $(n-1)\Delta$ period bond short and investing the proceeds in the $n\Delta$ period bond. The starting value of such a strategy is zero but the end result is not.

To analyze this situation let Z' denote the value of the $(n-1)\Delta$ bond. If f is the forward rate as before we assume $Z' = \exp[f\Delta]Z$. Also let $g = df / ds$ be the slope of the forward curve at the horizon. Using the same logic that was developed earlier we find that an investment $A = J'Z'$ in the $(n-1)\Delta$ bond rolls forward according to

$$\begin{aligned} dA &= (f - g\Delta)Adt + J'dZ', \\ &= (f - g\Delta)Adt + J'd(\exp[f\Delta]Z), \\ &= (f - g\Delta)Adt + J' \exp[f\Delta]dZ + Z\Delta \exp[f\Delta]J'\left\{df + \frac{1}{2}\Delta df^2 + df dZ / Z\right\}. \end{aligned}$$

If the same amount $A = JZ$ is invested in the $n\Delta$ bond we must have $J' = \exp[-f\Delta]J$. The net return for investing in a dipole is then, to first order of small quantities

$$\begin{aligned} dP &= J[fZdt + dZ] - (f - g\Delta)J'Z'dt - J'dZ', \\ &= g\Delta JZdt - JZ\Delta\left\{df + \frac{1}{2}\Delta df^2\right\} - J\Delta df dZ. \end{aligned}$$

Note that a correlation term $df dZ$ has now appeared in the asset roll forward.

Assuming the dipole can actually be put into practice we can go through the same modeling steps that were used for the monopoly model. The main difference is that we now assume that the quantity K is a function of the two state variables f, g .

We also have to make an assumption about the dynamics of g which we will write as

$$dg = \beta(\bar{g} - g)dt + v d\tilde{w}, \quad dwd\tilde{w} = \tilde{\rho}dt.$$

Plausible assumptions for these parameters are that \bar{g} is small but positive while $\tilde{\rho}$ is close to 0 and β is larger than α .

As before we find

$$\begin{aligned} dS = & fKZdt + g\Delta JZdt - JZ\Delta\{df + \frac{1}{2}\Delta df^2\} - J\Delta dfdZ \\ & - Z\left\{\frac{\partial K}{\partial t} + \alpha(\bar{f} - f)\frac{\partial K}{\partial f} + \beta(\bar{g} - g)\frac{\partial K}{\partial g} + \frac{\sigma^2}{2}\frac{\partial^2 K}{\partial f^2} + \frac{v^2}{2}\frac{\partial^2 K}{\partial g^2} + \sigma v\tilde{\rho}\frac{\partial^2 K}{\partial f\partial g}\right\}dt \\ & - \frac{\partial K}{\partial f}dfdZ - \sigma Z\frac{\partial K}{\partial f}dw - vZ\frac{\partial K}{\partial g}d\tilde{w}. \end{aligned}$$

If we set $J = \frac{-1}{\Delta}\frac{\partial K}{\partial f}$ then the ‘ f ’ risk is hedged, we are naked on the ‘ g ’ risk and both correlation terms $dfdZ$ drop out. The simplified result is

$$dS = fKZdt - Z\left\{\frac{\partial K}{\partial t} + \left[g - \frac{\Delta\sigma^2}{2}\right]\frac{\partial K}{\partial f} + \beta(\bar{g} - g)\frac{\partial K}{\partial g} + \frac{\sigma^2}{2}\frac{\partial^2 K}{\partial f^2} + \frac{v^2}{2}\frac{\partial^2 K}{\partial g^2} + \sigma v\tilde{\rho}\frac{\partial^2 K}{\partial f\partial g}\right\}dt - vZ\frac{\partial K}{\partial g}d\tilde{w}.$$

As before we set up economic capital for the unhedged ‘ g ’ risk

$$EC = DZ[K(t, f, g - \lambda v) - K(t, f, g)] \approx DZ\left[-\lambda v\frac{\partial K}{\partial g} + \frac{\lambda^2 v^2}{2}\frac{\partial^2 K}{\partial g^2}\right]$$

and then write the pricing equation as $E[dS] = \pi ECdt$. The final equation for K is

$$\frac{\partial K}{\partial t} + \left[g - \frac{\Delta\sigma^2}{2}\right]\frac{\partial K}{\partial f} + \beta\left(\bar{g} - \frac{\pi D\lambda v}{\beta} - g\right)\frac{\partial K}{\partial g} + \frac{\sigma^2}{2}\frac{\partial^2 K}{\partial f^2} + \frac{(1 + \pi D\lambda^2 v^2)}{2}\frac{\partial^2 K}{\partial g^2} + \sigma v\tilde{\rho}\frac{\partial^2 K}{\partial f\partial g} = fK.$$

We now have a two factor Vasicek model which has a closed form solution of the form

$K(t, f, g, s) = \exp\left[-\int_0^{s-t}\delta(f, g, u)du\right]$ where the extrapolated force of interest is linear in the state

variables f, g . The formal solution is

$$\begin{aligned} \delta(f, g, u) = & f + g\frac{1 - e^{-\beta u}}{\beta} - \frac{1}{2}\Delta\sigma^2 u + G(u)[\beta\bar{g} - \pi D\lambda v] \\ & - \frac{1}{2}\left[\sigma^2 u^2 + 2\tilde{\rho}\sigma v u G(u) + v^2(1 + \pi D\lambda^2)G^2(u)\right] \end{aligned}$$

where $G(u) = [e^{-\beta u} + \beta u - 1]/\beta^2 = \int_0^u \frac{1 - e^{-\beta u}}{\beta} du$.

The economics of this model are quite different from the monopole model. In the limit where $u \rightarrow 0$ we see that $G(u) \approx u^2 / 2$ and so

$$\delta(f, g, u) \approx f + (g - \frac{1}{2} \Delta \sigma^2) u + O(u^2).$$

The short end of the extrapolated yield curve is approximately constant at the current forward rate with a small adjustment to account for the bias in the dipole strategy. However, as $u \rightarrow \infty$ we find that $G(u) \rightarrow u / \beta^2$ and the extension is dominated by the convexity adjustment in square brackets.

What this result is telling us is that, although we have hedged out the 'f' risk, there is a large convexity mismatch between the liability and dipole asset strategy that has to be priced into the liability. Plausible values for the new parameters appearing in the extension formula suggest that terms involving the 'g' risk and cost of capital are negligible compared to the pure convexity terms. A simplified extension formula which captures the most significant issues would then be

$$\delta(f, g, u) \approx f + g \frac{1 - e^{-\beta u}}{\beta} - \frac{1}{2} \Delta \sigma^2 u - \frac{1}{2} \sigma^2 u^2.$$

The details of these two extensions depend on the specific simplifying assumptions that have been made here but the general features do not. No matter how we decide to model the details the correlation issue in the monopole model and the convexity cost in the dipole model will not go away. Any reasonable model will suggest that material margins need to be added to the simple intuitive models in order to correct for the risk management shortcomings of each approach.

Appendix 2 - Mathematics of the Static Control Model

The main purpose of this appendix is to derive the key theoretical results stated in the section on the Static Control Model. In addition to that goal we want to emphasize that the Static Control model can be considered to be a very general approach to valuing unhedgeable risks that extend over time. In this paper the concept is being applied to long liabilities but the tool has more potential than that.

In its most general form the Static Control has five basic ingredients

1. A large set of real world (P measure) economic scenarios indexed by $\alpha = 1, \dots, N$.
2. An instrument that we wish to value that produces cash flow a_t^α at time t on scenario α .
3. A universe of available market instruments whose current observed market values are $Z_i, i = 1, \dots, M$ and which produce cash flow CF_{it}^α on the modeled scenarios.
4. A coherent risk measure, such as CTE($x\%$), which assigns a weight w_α to the scenario with rank order α .
5. An asset class to serve as numeraire. We will denote the discount factor corresponding to time t on scenario α by v_t^α .

In the body of the paper, these assumptions were simplified by taking the numeraire to be a log normal equity process and the universe of available market instruments was restricted to 0 coupon bonds with maturities out to 30 years.

The method starts by considering an arbitrary position b_i in each of the available instruments which generates cash flow $\sum_i b_i CF_{it}^\alpha$ at a cost of $\sum_i b_i Z_i$. If this does not match the liability perfectly we apply the coherent risk measure to the present value distribution of the residual cash flow $a_t^\alpha - \sum_i b_i CF_{it}^\alpha$. The value we put on this structure is then written as

$$W(a, b, Z) = \sum_i b_i Z_i + \sum_\alpha w_\alpha \sum_t [v_t^{(\alpha)} (a_t^{(\alpha)} - \sum_i b_i CF_{it}^{(\alpha)})].$$

The round bracket around the index (α) indicates that sorting has been done on the present values $PV^\alpha = \sum_t v_t^\alpha (a_t^\alpha - \sum_i b_i CF_{it}^\alpha)$.

Given this set up the formal definition of the static control value is

$$\begin{aligned} V(a, Z) &= \min_b W(a, b, Z), \\ &= \min_b \left(\sum_i b_i Z_i + \sum_\alpha w_\alpha \sum_t [v_t^{(\alpha)} (a_t^{(\alpha)} - \sum_i b_i CF_{it}^{(\alpha)})] \right). \end{aligned}$$

The first theoretical result, **market consistency**, follows from Tasche's theorem which tells us how to differentiate once through a risk measure. The result is

$$\frac{\partial W}{\partial b_i} = \left(Z_i - \sum_\alpha w_\alpha \sum_t [v_t^{(\alpha)} CF_{it}^{(\alpha)}] \right).$$

The first order optimality condition $\frac{\partial W}{\partial b_i} = 0$ is therefore equivalent to the statement that the weighted scenarios price the available market instruments properly. This result is also important for estimating the optimal mix $b_i^* = b_i^*(a, Z)$ since having an estimate of the first derivative of the function W we are trying to optimize allows us to use the tools of non-linear optimization.

Another way of stating this result is that the optimization process has allowed us to assign weights to each P measure scenario in such a way that we now have a calibrated risk neutral set of scenarios. CTE risk measures are particularly useful in this regard since they assign a weight of 0 to many scenarios. This implies the resulting calibrated scenario set is smaller than the set we started with.

A technical caveat is appropriate at this point. There is no guarantee that, for a general set of inputs, there is a unique local minimum. If the problem is improperly specified or the scenario set is too small there can appear to be a statistical arbitrage. If that happens then the minimum is effectively $-\infty$ and the optimization step fails. An example is using a risk measure that is "too close" to CTE(0) for a given problem.

The author's practical experience is that a failed optimization step is usually the indicator of an error that, once corrected, is no longer a problem⁸.

Static Hedge: Having solved the optimization problem we can write the static control value as

$$V(a, Z) = \left(\sum_i b_i^* Z_i + \sum_\alpha w_\alpha [v_t^{(\alpha)}(a_t^{(\alpha)} - \sum_i b_i^* CF_{it}^{(\alpha)})] \right).$$

Using Tasche's theorem again we calculate

$$\frac{\partial V}{\partial Z_k} = b_k^* + \left(\sum_i \frac{\partial b_i^*}{\partial Z_k} (Z_i - \sum_\alpha w_\alpha \sum_i v_t^{(\alpha)} CF_{it}^{(\alpha)}) \right) = b_k^*.$$

The large round bracket vanishes by virtue of the first order optimality condition so we get the result stated earlier that the optimal asset mix also gives us the static hedge exposure.

We can also write the optimized value as

$$\begin{aligned} V(a, Z) &= \left(\sum_i b_i^* \sum_t (Z_i - \sum_\alpha w_\alpha v_t^{(\alpha)} CF_{it}^{(\alpha)}) + \sum_t \sum_\alpha w_\alpha v_t^{(\alpha)} a_t^{(\alpha)} \right) \\ &= \sum_t \sum_\alpha w_\alpha v_t^{(\alpha)} a_t^{(\alpha)} \end{aligned}$$

This simply says that the static control value of the liability is the risk neutral present value of the cash flows we are trying to value where the value is calculated using the calibrated scenario set. In particular this means that if the liability can be expressed as a static linear combination of the available market instruments then it will be priced as such.

Convexity Margin: If we start with the static hedge result just derived we find

$$\frac{\partial V}{\partial Z_k} = b_k^* \Rightarrow \frac{\partial^2 V}{\partial Z_j \partial Z_k} = \frac{\partial b_k^*}{\partial Z_j}.$$

We will show that $\frac{\partial b_k^*}{\partial Z_j} = - \left(\frac{\partial^2 W}{\partial b_j \partial b_k} \right)^{-1} \Big|_{b=b^*}$. If the optimization was successful the matrix

$\left(\frac{\partial^2 W}{\partial b_j \partial b_k} \right)$ is positive definite at the point $b = b^*$ implying that $\frac{\partial^2 V}{\partial Z_j \partial Z_k}$ is negative definite.

To do this it helps to introduce the notation $W(a, b, Z) = \sum_i b_i Z_i + L(a, b)$. From this we easily

⁸ If the risk measure is coherent the function $W(a, b, Z)$ is convex in the variables b . This means we have one of a) a unique local and global minimum, b) the minimum is at infinity or c) we have convex set of local minima with the same value. Option c) won't occur if the set of available market instruments is linearly independent and option b) usually means we have improperly specified the problem.

find $\frac{\partial W}{\partial b_i} = Z_i + \frac{\partial L}{\partial b_i}$, $\frac{\partial^2 W}{\partial b_i \partial b_j} = \frac{\partial^2 L}{\partial b_i \partial b_j}$. On the other hand when $b_i = b_i^*$ the first order

optimality condition is satisfied so $Z_i = -\frac{\partial L(a, b^*)}{\partial b_i}$. Differentiating this result with respect to

Z_k we find $I_{ik} = -\sum_j \frac{\partial L^2(a, b^*)}{\partial b_i \partial b_j} \frac{\partial b_j^*}{\partial Z_k}$ where I_{ik} is the identity matrix. We conclude that both

matrices in this equation are invertible and $\frac{\partial b_k^*}{\partial Z_j} = -\left(\frac{\partial^2 L}{\partial b_j \partial b_k}\right)^{-1}\bigg|_{b=b^*} = -\left(\frac{\partial^2 W}{\partial b_j \partial b_k}\right)^{-1}\bigg|_{b=b^*}$ as

required.