Actuaries Summit

get involved, get ahead

20-21 May 2013
Hilton Sydney

Actuaries Institute

A topic of interest – how to extrapolate the yield curve

Prepared by Peter Mulquiney & Hugh Miller

Presented to the Actuaries Institute
Actuaries Summit
20-21 May 2013
Sydney

Institute of Actuaries of Australia
ABN 69 000 423 656
Level 7, 4 Martin Place, Sydney NSW Australia 2000
t +61 (0) 2 9233 3466 **f** +61 (0) 2 9233 3446
e actuaries@actuaries.asn.au **w** www.actuaries.asn.au

A topic of interest – how to extrapolate the yield curve

Peter Mulquiney Hugh Miller

This paper addresses key practical questions concerning the setting of interest rate assumptions in actuarial work. In particular, we look at the situation where the liabilities being valued extend beyond the term of available market instruments, and so the yield curve must be extrapolated. We find that, in Australia, the yield curve up to 2 years before the longest dated bond can be estimated reliably. We look at international evidence and conclude there is reasonable evidence of reversion to a flat long term forward rate. However, the evidence suggests that the rate of reversion is slow with a term of about 40 years being the minimum point to reversion, and 60 years the central estimate. We also propose hedging strategies for long-dated liabilities and show how these give further insight into appropriate rates.

1 Introduction

One of the most fundamental concepts in actuarial practice is the time value of money. For any work where future cash flows are allowed for, such as reserving or pricing, it is natural to discount to present values so that an appropriate amount of money can be set aside today, allowing for future investment returns.

It is widely accepted that for claims reserving, liabilities should be discounted using the prices of the "risk-free" assets available in the financial markets. This means that the present value of a liability cash flow should be set equal to the market price of a basket of risk-free assets that provides a matching cash flow.

Although the principle of using risk-free rates for discounting is widely accepted, for some time there has been considerable debate on some practical aspects of the principle. This debate gained intensity following the global financial crisis of 2007-08 which saw large increases in the price of risk-free assets and correspondingly, large decreases in risk-free interest rates. Issues of debate include:

- What are the best instruments to use to determine risk-free interest rates?
- Should the risk-free rate include an "illiquidity premium"?
- What should be done when the liabilities being valued extend beyond the term of available market instruments?

In the Australian context, where we have a deep and liquid market in AAA rated Commonwealth Government Bonds, the Australian Prudential Regulation Authority (APRA) has made it clear that it regards these bonds as the best instruments to use to determine the risk-free rate (APRA 2010;

APRA Draft GPS 320 - Discount rates). Further, for general insurance liabilities at least, APRA do not allow the inclusion of an illiquidity premium.

However, the issue of what should be done when the liabilities being valued extend beyond the term of available market instruments has, in our opinion, not been fully addressed in any current Australian regulations. Further, actuaries operating in the Australian market have adopted a wide range of approaches to this issue and this has led to inconsistent valuations of long-term liabilities across entities. This issue is particularly relevant for Australia where the term of the longest dated government bond is currently around 15 years - in markets such as the US, UK and Canada government bonds are available for terms of up to 30 years or more.

Common to all approaches for the valuation of long-term liabilities is the requirement that the yield curve be extrapolated to terms beyond those available in the financial markets. The aim of this paper is to review the issues relevant to yield curve extrapolation.

Of course, the issue of yield curve extrapolation is not new. In particular it has received considerable discussion in Europe with the move towards a market-based approach to valuation under Solvency II (e.g. CEIOPS, 2010; CRO Forum, 2010). In presenting this work we have been heavily influenced by these discussions. In particular, we have found the papers and discussion notes by Barrie & Hibbert (Barrie and Hibbert, 2008; Hibbert, 2008; Carlin 2010; Hibbert 2012) helpful and have chosen to organise this paper around the three questions they recognise as being central to the technical issue of extrapolation:

- What is the longest-maturity market forward interest rate we can estimate reliably?
- For the purposes of extrapolation, what is the ultimate very long-term “unconditional” forward interest rate?
- What path should be set between the longest market rate and the unconditional forward rate and how many years should it take to the final level?

These issues are summarised in Figure 1.

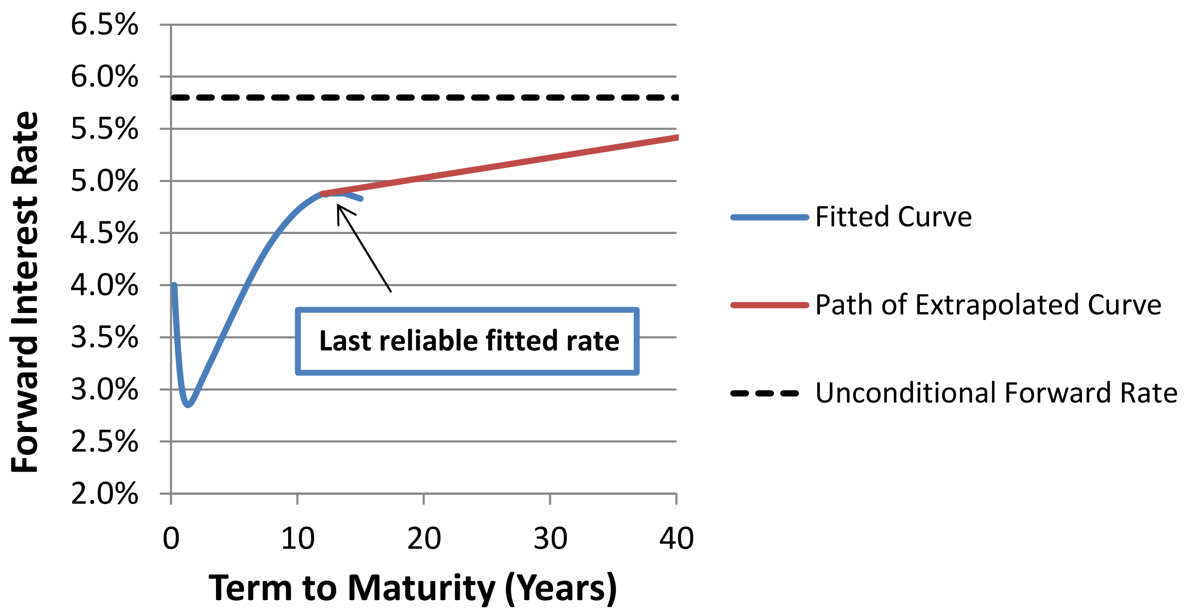

Figure 1 **Yield curve extrapolation**

Our main contribution to the debate relates to the third question. In particular, we present analysis which suggests that the yield curve may only move towards its ultimate level at a relatively slow rate. The implication of this is that for most long-tailed general insurance liabilities, the assumed long-term "unconditional" forward interest rate should have little impact on the valuation result. We start though, with a discussion of the first two technical questions as they give a fuller context to the results we have found.

Before starting on the technical issues, we give a brief overview of the philosophical and regulatory issues surrounding yield curve extrapolation - these issues will influence how the technical issues should be solved.

2 Philosophical and regulatory considerations for yield curve extrapolation

2.1 Philosophical considerations

Broadly speaking there are two philosophical approaches to yield curve extrapolation; one which emphasises market consistency at a point-in-time, and another which emphasises liability stability (Hibbert, 2012) across time.

With the market consistency approach, the aim of yield curve extrapolation is to estimate what the market price of longer-term assets would have been on a particular day if those assets actually existed and were traded in the financial markets. The emphasis here is on producing a liability value that is equivalent to the value that would be required to transfer that liability in current market conditions.

On the other hand, those adopting an approach emphasising liability stability believe that the volatility induced by requiring market consistent prices at a point-in-time is unhelpful for running a business over many years. Also they may claim that insurance liabilities are rarely traded and so the idea of identifying market value is less relevant. Some proponents of this view may be fully supportive of the idea of valuing liabilities for which hedging instruments are available with reference to the prices of those instruments. But where those instruments don't exist, or for some proponents, are not sufficiently liquid, they would rather emphasise liability stability over market consistency (e.g. see CRO Forum, 2010).

Each approach has its advantages and disadvantages. While an approach which emphasises liability stability may be superficially attractive, it may cause insurers to underestimate the true economic cost of writing long-term insurance contracts. It would not be in the firms, policyholders or shareholders future interests if long-term contracts were issued too cheaply. In addition an artificially stabilised liability valuation may also make it harder to hedge balance sheet risk over the longer-term as our analysis in Section 6 seems to indicate.

On the other hand it has been suggested as part of the Solvency II debate that emphasising market price consistency may play a role in magnifying economic or financial downturns, or in other words it may add to pro-cyclicality (CRO Forum, 2010).

2.2 Regulatory and professional considerations

So what approach do Australian accounting and prudential standards suggest is required?

The Australian Accounting Standard AASB 1023 states:

AASB 1023 6 – Discount Rates

6.1 The outstanding claims liability shall be discounted for the time value of money using risk-free discount rates that are based on current observable, objective rates that relate to the nature, structure and term of the future obligations.

6.1.2 Typically, government bond rates may be appropriate discount rates for the purposes of this Standard, or they may be an appropriate starting point in determining such discount rates.

And so is unfortunately unhelpful when observable, objective rates do not exist.

APRA's current prudential standard applicable to the valuation of insurance liabilities states:

APRA GPS 310 – Audit and Actuarial Reporting and Valuation. Appendix A

31. The value of an insurer's insurance liabilities is typically independent of the value of the insurer's underlying assets. A discount rate that is based on current observable, market-based and objective rates that directly relate to the nature, structure and term of the future obligations must therefore be used.

32. The rates to be used in discounting the expected future claims payments for a class of business are derived from the gross redemption yields, as at the calculation date, of a portfolio of sovereign risk securities in the currency of, and with a similar expected payment profile to, the insurance liabilities for that class (for example, Commonwealth Government securities for Australian liabilities). It is acceptable to use either the average rate or a series of discount rates taken from the relevant yield curve.

33. Where the expected payment profile of the insurance liabilities cannot be matched (for example, because the duration is too long), a discount rate regarded as consistent with the intention of paragraphs 31 and 32 of this Attachment must be used.

Although not entirely clear, clause 33 seems to imply that yield curve extrapolation should be performed on a point-in-time market consistent basis. APRA has also released the draft GPS 320 that may be applicable to general insurance valuations from 1 January 2013:

APRA Draft GPS 320 - Discount rates

35. The rates to be used in discounting the expected future claims payments of insurance liabilities denominated in Australian currency for a class of business are derived from yields of Commonwealth Government Securities (CGS), as at the calculation date, that relate to the term of the future insurance liability cash flows for that class.

36. Where the term of the insurance liabilities denominated in Australian currency exceeds the maximum available term of CGS, other instruments with longer terms and current observable, objective rates are to be used as a reference point for the purpose of extrapolation. If there are no other suitable instruments or the Appointed Actuary elects to use an instrument that does not meet this requirement, the Appointed Actuary must justify the reason for using that particular instrument in the insurer's ILVR. Adjustments must be made to remove any allowances for credit risk and illiquidity that are implicit in the yields of those instruments.

The suggestion to use other instruments where the term of the liabilities exceeds the maximum available term of Commonwealth Government Bonds suggests a preference for a point-in-time market consistent valuation and so this standard appears consistent with the earlier standard. Unfortunately, at present there does not appear to be a large and transparent market in alternative instruments for durations beyond 15 years. For example, the only publicly available information on bank bill swaps is for terms up to 15 years.

Finally, the Institute's PS 300, the standard applicable to general insurance valuations, mandates:

IAA PS300 – General Insurance Business

8.2.2 Legislative and/or regulatory requirements may prescribe whether Claim Payments are to be discounted. The Member must consider the purpose of the valuation and document whether the future Claim Payments are to be discounted. Discount rates used must be based on the redemption yields of a Replicating Portfolio as at the valuation date, or the most recent date before the valuation date for which such rates are available.

8.2.3 If the projected payment profile of the future Claim Payments cannot be replicated (for example, for Classes of Business with extended runoff periods), then discount rates consistent with the intention of Paragraph 8.2.2 must be used.

"Replicating Portfolio" means a notional portfolio of current, observable, market-based, fixed-interest investments of highest rating, which has the same payment profile (including currency and term) as the relevant claim liability being valued.

Here the wording of 8.2.3 is remarkably similar to that used in clause 33 of APRAs GPS 310 - Appendix A. This suggests to us that the professional standards require a market consistent approach to be used, also.

2.3 Approach taken in this paper

Although not entirely clear, it seems to us that the intention of both APRA's prudential standards and the Institute's PS300 is for yield curve extrapolation to be performed on a point-in-time market consistent basis. In this paper we will approach the technical issues from this viewpoint. However, when there is uncertainty about how the market may behave at longer terms we will favour an approach that contributes to liability stability rather than an approach that may add unwarranted volatility to an insurer's balance sheet.

3 What is the longest market forward interest rate we can estimate reliably?

The question we consider in this section is: at what point along the forward interest rate yield curve should we start to extrapolate? A common starting point for extrapolation is the term of the last available market instrument. In this section we consider reasons why this may not always be appropriate.

Before considering this question we will briefly discuss why we are working with the forward rate yield curve rather than one of the other possible representations of the yield curve.

3.1 Extrapolation of forward rates rather than spot rates

Yield curves can be represented in a variety of ways but those with which actuaries tend to work show either spot rates or forward rates at various terms to maturity. For this paper we have chosen to work with instantaneous (continuously compounding) forward interest rates. A reason for this is that if we were to use say, a linear extrapolation of the spot rate curve, the change in slope would lead to a jump in the instantaneous forward rate and this would create a yield curve with arbitrage opportunities. As a general principle is seems preferable that extrapolated interest rates be arbitrage free (CRO Forum, 2010).

3.2 Uncertainty in forward rate estimates

The starting point for forward rate estimation, in our case, will be the prices of Commonwealth Government bonds. When fitting these prices to estimate the forward rates, if we were to fit the prices exactly we would end up with a very irregular (bumpy) looking forward curve. The reason for this is that that bonds of adjacent maturities can trade at prices that are slightly off-trend as traded volumes across different maturities vary. Or in other words, liquidity effects can lead to noise in reported bond prices. Because of this noise, a better estimate of the true underlying forward curve is obtained by imposing some smoothness on the estimated forward rates. In fact, a number of studies have shown that smooth forward rate curves perform better than bumpy ones in predicting the prices of bonds deliberately left out of the forward rate estimation process. That is, the smoother curves give better out-of-sample performance (Waggoner, 1997; Bolder and Gusba, 2002).

Because of the noise present in bond prices, any forward rate curve will have some estimation error associated with it. To illustrate this we estimated 90% confidence intervals for forward rate curves from two different days over the last year (Figure 2).

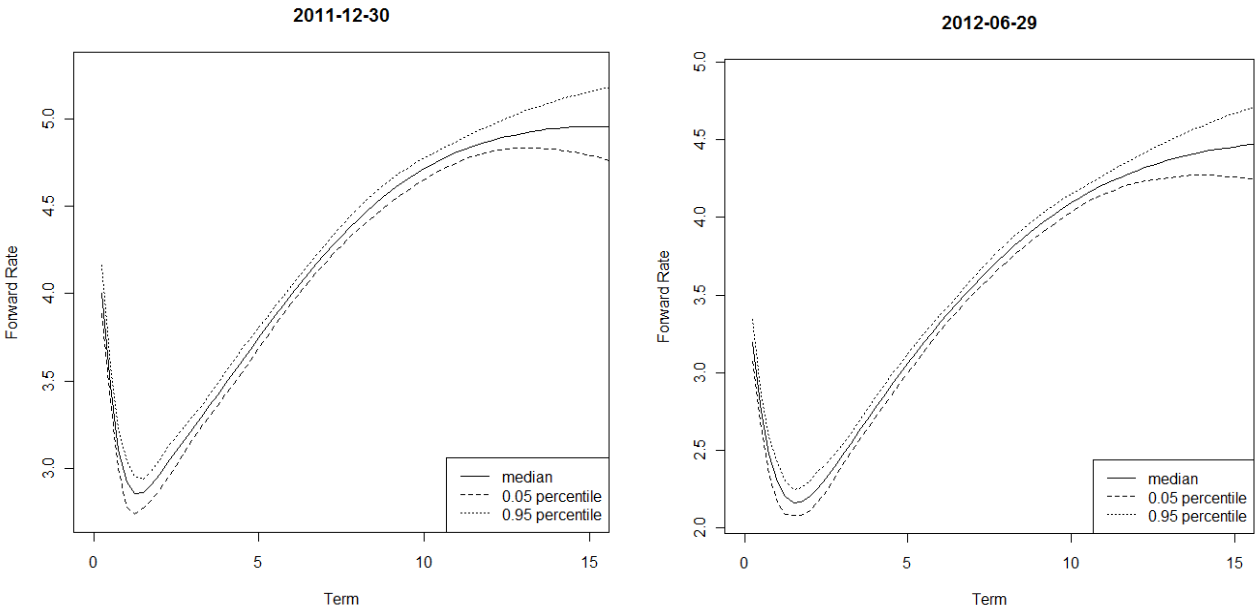

Figure 2 **Estimation error in the forward rate curve**

The forward curves were estimated from the quoted prices of Australian Government Bonds on two separate days using a modified version of the Merrill Lynch Exponential Spline Method as presented in Finlay and Chambers (2008). This is the method the Reserve Bank of Australia currently uses for forward rate estimation. Error estimates were made using the bootstrap method with 1000 bootstrap samples. The forward rates that we estimated minimised the weighted least squares error in bond prices. The weights used were the reciprocal of the modified duration for each bond. The effect of this was to place less weight on longer term bonds. The specific assumption is that the variance of the bonds pricing error is proportional to its duration. Or more simply, that we expect the observed prices for these bonds to be more variable.

Figure 2 shows that up until a term of 10 years, the 90% confidence interval is around 10-15 basis points in width. But by year 13 it has spread to 25 basis points, and reaches about 40 basis points by year 15. Part of the reason for this increase between years 13 and 15 is the sparseness in bonds in this region. For example, the last two bonds are at terms of roughly 12 years and 15 years - a 3 year separation.

The implication of this analysis is that the uncertainty in forward rate estimates will be far greater at the long end of the curve. Indeed, similar studies from periods when the longest maturity bonds were closer to 12 years, suggest that forward rates up until about 2 years prior to the longest maturity date are fairly reliable, with estimation errors increasing rapidly after that.

A further consideration that may affect the reliability of forward rate estimates at the long end is that, in most markets, the liquidity of instruments with long terms is generally lower. For example, in Australia the value on issue of 15 year Commonwealth Government bonds is only 2% of the total value on issue of all Commonwealth Government Bonds. It may be reasonable to expect that these bonds are subject to more price noise due to their reduced liquidity.

It is worth mentioning here that the estimation error measured in Figure 2 is specific to the model we have used to estimate the forward rates. The model we used permits a great deal of flexibility in long-term forward rates. We could, however, have used an alternative model that tried to minimise estimation error at the long end, for example, by using a model that assumed the forward curve levelled out at some duration prior to 15 years. While this would reduce estimation error it would leave us no wiser as to whether the curve had actually flattened out or not, and this would, in our view, be less informative for the problem of extrapolation.

3.3 Impact on extrapolation

The analysis presented in this section suggests that we should not be overly reliant on forward rate estimates made at the long end of the fitted forward curve, in particular the last 2 years of the observable range. At present it appears that terms of around 13 years would be an appropriate point to start an extrapolation in Australia.

Although we are recommending not placing too much reliance on the forward rates estimated at the long end, if in the process of extrapolation there was a significant deviation between the fitted curve and the extrapolated curve some care must be taken. For example, it would seem to us to be important to check that the prices of the longest-dated bonds were still estimated to within an acceptable tolerance level when the extrapolated curve was used.

With these considerations in mind some other questions are posed by the results of Figure 2. In particular, given the curves are still rising at a term of 13 years:

- At what level should the forward rates ultimately be extrapolated to?; and
- How quickly should they get there?

The first of these questions is answered in the next section.

4 For the purposes of extrapolation, what is the ultimate very long-term "unconditional" forward interest rate?

The rational expectations hypothesis of the term structure of interest rates suggests, in its more general form, that long-term forward rates are the sum of the expected future short-term interest rate plus a constant "term premium" that varies only by term and not over time. Taking this to be true, the process of determining the ultimate very long-term "unconditional" forward interest rate ("UFR") would involve making estimates of what these two components are. Unfortunately, things are not so simple.

The rational expectations hypothesis has been shown to be inconsistent with the empirical evidence. In particular, a number of studies have shown that long-term interest rates are significantly more volatile than would be expected if the rational expectations hypothesis held. In addition, studies have also shown that long-term rates respond to information that would reasonably be thought to influence short-term rates only. These inconsistencies result from time-varying term premia (Shiller, 1990).

It is now commonly accepted that time-varying term premia account for most of the variability in long-term forward rates (Kim and Orphanides, 2005; Finlay and Chambers, 2008). A fundamental question to be answered then is how far along the term structure do these term premia variations extend? (Hibbert, 2008).

In the following section we discuss this question and other important considerations needed to project expected future short-term interest rates and term premia out to very long terms.

4.1 The difference between term and time

Before we begin this section it is important to understand the distinction between term and time. On any given day there is a yield curve of spot rates or forward rates that varies by term to maturity, or equivalently, duration. The yield curve itself will also evolve over calendar time as the prices of bonds of different durations change over time. The problem of determining the UFR involves determining what the forward rate is at very long durations on a specified day. In this section we will discuss the idea that the UFR can be determined as the sum of the expected future short-term interest rate plus the term premium applicable at very long durations on the specified day. Here the expected future short-term interest rate is the markets expectation on the specified day of what short-term interest rates will be at a maturity date long into the future.

4.2 Expected short-term interest rate

When projecting expected future short-term interest rates to maturity dates long into the future it is reasonable to expect that our ultimate rate should only change in response to fundamental changes in the structure of the economy and should not be affected by short-term economic changes.

With this in mind, it is usual (e.g. Barrie and Hibbert, 2008; CRO Forum, 2010) to separate the problem of setting the expected short-term interest rate into one of determining:

- Expected future inflation; and
- Expected future real short-term interest rate.

In relation to expected future inflation, the Reserve Bank of Australia has been very successful in targeting inflation and entrenching low and stable inflation expectations for at least the last 15 years

(Finlay and Chambers, 2008). It seems reasonable then to adopt the mid-range of the bank's current CPI target of 2-3% as our future inflation expectation. The issue of forecasting inflation rates was dealt with in more detail in a recent paper by Miller (2010).

For expected real short-term interest rates a typical approach has been to look at historical averages for real cash returns across several countries (e.g. CEIOPS, 2010; Barrie & Hibbert, 2008). The underlying assumption is that, in the long-run, real interest rates should not differ substantially across economies which are at a similar stage of economic development. As part of Solvency II, the QIS 5 assumed the expected real rate of return to be 2.2% (CEIPS, 2010). This figure was determined by reference to a study by Dimson et al. (2000) looking at bond returns over the second half of the 20th century for 12 major economies, including Australia. A similar study by Barrie & Hibbert (2008) found a median estimate of 1.8% from 16 developed countries over the period 1970 – 2007. In light of these studies a figure of around 2% seems reasonable to us.

4.3 Term premia

The definition of term premia that is most useful in the current context is that they are the difference between the forward rate and the expectation of the future short-term interest rate. Term premia have a number of causes:

- Investors demand a term premium for locking into long-term investments. In this case the term premium acts as compensation for holding long term bonds whose value will fluctuate in the face of interest rate uncertainty, exposing the holder to mark to market losses. Term premia in this case are often referred to as risk premia and will be positive.

- Alternatively, demand for long-term government securities from large institutional investors such as insurance companies and pension funds can drive down long-term forward rates because these long-term bonds offer a closer match to liabilities and are less risky investments to these investors. Such forces can cause term premia to be negative. This phenomenon is often called "term preference".

- Additionally, convexity effects also cause term premiums to decrease at longer maturities. Fixed income securities have positive convexity. This means that the capital gains and losses from equal sized interest rate swings are not matched – the gains will be greater than the losses. This effect can theoretically cause very-long duration bonds to trade at higher prices (lower yields).

Term premia tend to increase quickly up to around 2 years duration and reach a relatively flat level by about 10 years. Term premia at around 10 years are on average of the order of 1% to 2% and at this duration appear to be dominated by risk premia effects. In many currencies, starting around 15 to 20 years, there is a decline in the forward rates which may be of the order of 1% by year 30 (see Figure 3). This decline is primarily due to the term preference of pension and life insurance companies, but also due the increased convexity of long-maturity bonds.

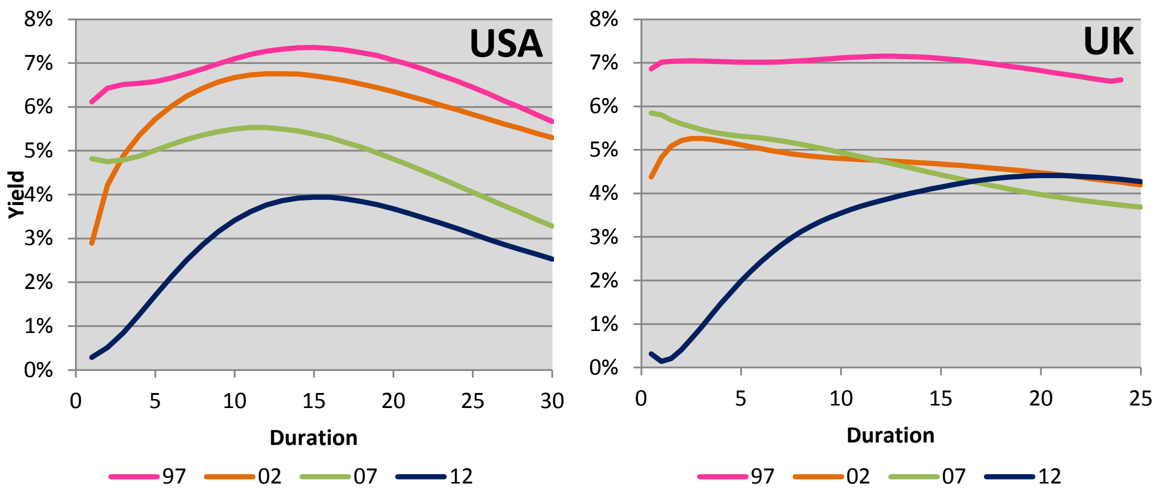

Figure 3 **Forward rate yield curve at the end of June in each year**

However, this is a stylised view of term premia and ignores their important time-varying behaviour. This behaviour is discussed in the following sub-section.

4.3.1 The time-varying behaviour of term premia

Model based techniques exist which allow the decomposition of forward rates into expected future short-term interest rates plus term premia. These models may incorporate survey data about analysts' forecasts of short-term interest rates. Those who use these models acknowledge that they are complex and difficult to calibrate and that the results should not be interpreted too precisely. However they do appear to provide important insight into the time-varying nature of term premia.

The following figure shows the results from one such analysis (Kim and Wright, 2005) on US Treasuries. The figure shows how forward rates at different terms can be broken down into the underlying expected short rate and term premium components. The results show the daily variation.

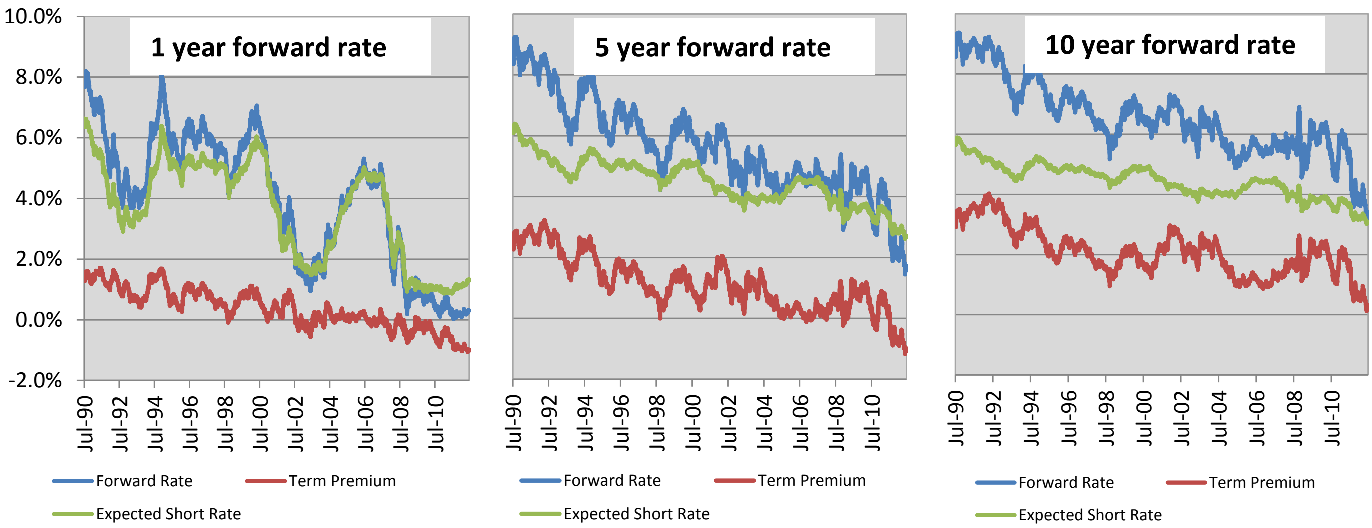

Figure 4 **Term premia estimated for US Treasuries**

The data in this figure was obtained from the US Federal Reserve website and was estimated using the method outlined in the paper by Kim and Wright (2005).

Some relevant observations from Figure 4 include:

- Term premia are particularly volatile. There can be large month to month swings and over the course of a year it is not uncommon for the term premia to change by 1% - 2%. The changes in term premia at year 5 are often mirrored in the changes at year 10.

- Term premia tend to increase when yields are rising and tend to decrease when yields are falling. These changes sometimes appear as over-reactions to tightening or loosening monetary policy.

- There are some noticeable trends in term premia. In particular term premia declined over the early 90s until around mid-1998. This decline corresponded to a general decline in US inflation expectations over this period.

- Towards the end of 2011 term premia have declined dramatically at all durations. For some durations the term premia are negative suggesting investors are willing to pay a premium for government bonds. This is a result of the “flight to quality” in response to the bad news from Europe.

For our purposes, the most important observation from Figure 4 relates to the variability of the term premia. We see that while most of the variability of the 1 year forward rate is explained by variability in the expected short rate, by year 10, most of the variability is explained by variability in the term premia. This is, in part, because short-term interest rates are expected to mean revert over time. Or in other words, expectations of future short-term interest rates are relatively fixed by durations of 10 years or more. On the other hand, the term premia appear to display little (or no? - it is not certain from the study) mean reversion with term.

Similar studies have been performed on Australian Commonwealth Government bonds with similar results. Indeed there is a very close correlation between Australian and US term premia suggesting that they are driven by global rather than country specific factors (Finlay and Chambers, 2008).

4.3.2 Implications for UFR determination

As discussed earlier, the term structure for a particular date depends only on the prices of bonds quoted on that date. These prices in turn depend on the economy wide supply and demand for risk free credit on that particular day. And as we have seen in Figure 4 supply and demand changes may result in large swings in term premia from day to day and month to month.

An unresolved question is how far down the term structure do these day to day and month to month swings extend? Unfortunately there are no empirical results which can help us with this question. It certainly seems possible that the volatility in term premia could extend out across the entire term structure if bonds of very long maturities were actually traded and particularly, if this volatility was driven by changes in the risk premia required for investing long-term. But this is mere speculation and in the words of Hibbert (2008) "how far is it reasonable to adjust the 50-year, 100-year and 1000-year forward rate?"

So the pragmatic response is to set the long-term term premium assumption to some assumed stable level. This may not be correct, but it steers away from introducing potentially unwarranted volatility onto an insurer's balance sheet.

4.4 Assumptions for components of the UFR

Pulling together the arguments of this section, at the present time a UFR of about 5.8% seems reasonable for Australia. The components of this estimate are detailed in the following table.

Table 1 **Components of the Unconditional Forward Rate for Australia in 2012**

Component	**Rate**
Expected future short-term interest rate	
Expected future inflation	2.5%
Expected future real interest rate	2.0%
	4.5%
Term Premium	
Risk Premium	1.5%
Convexity adjustment	-0.2%
	1.3%
Unconditional forward rate	**5.8%**

As discussed earlier the component of the UFR that is most uncertain is the term premium. A value of 1.3% was chosen using the following considerations:

- Using the data from Figure 4, we calculate that the average term premium on 10 year bonds in the US, since the stabilisation of US inflation expectations from around 2007, is 1.7%;
- Work by the Reserve Bank of Australia has shown that since 1997, term premia on 5 year forward rates have been around 0.5% to 1% lower than those in the US (Finlay and Chambers, 2008). However in that work, estimated short-term rates that were about 1.5% higher than those assumed here;
- While there are **term preference** effects in markets such as the US and UK beyond terms of around 15 years, we are unsure if it is appropriate to allow for them in the UFR for Australia. It has been suggested that at very long terms there are no liabilities to hedge anyway so the

hedging activities of institutional investors are unlikely to bid down term premia beyond say 100 years (Carlin, 2010);
- For the convexity adjustment we have relied on the work of the CRO Forum (2010) which showed that the convexity premium decreased about 0.2% from year 10 to year 30 (where it reached a minimum). Other estimates have been in the range of 0.4% (Barrie and Hibbert, 2008); and
- Given uncertainties in both the risk premium component and the convexity adjustment, a combined figure of around 1.3% seems reasonable.

As a reasonableness check on our adopted UFR of 5.8%, we note that the average 10 year forward rate in Australia has been 5.8% since 1998. Strictly, speaking, to compare our UFR to the recent average 10 year forward rate we need to remove the convexity adjustment - since by our definition this adjustment is zero at year 10 - giving a rate of 6.0%, which is slightly above the recent 10 year forward rate average.

It should be apparent from the discussion in this section that there is some uncertainty about what an appropriate UFR assumption should be – a figure anywhere in the range 5.4% - 6.2% could be reasonably justified at this point in time. For most long-tailed general insurance liabilities, this uncertainty will only have a material impact if the UFR is reached relatively quickly in an extrapolation. This issue of the speed of reversion to the UFR is discussed in the next section.

5 What path should be set between the longest market rate and the unconditional forward rate?

5.1 Introduction

We now move on to the final, and arguably most important, aspect of yield curve extrapolation: what path should be set between the longest duration market rate and the unconditional forward rate and how long should we take to reach it?

Note there is an implicit assumption in the previous question – that the extrapolated rate should approach to the UFR. There is a possibility that the extrapolated curve should be flat, or equivalently that it takes a very long time to approach the UFR. The international evidence offered below suggests that this is not the case.

The speed of return to the UFR is of practical importance; in the Australian context, if it returns quickly (by duration 20 years say), then long-tailed liabilities will be considerably more stable over time than if it returns slowly (e.g. by 100 years).

Throughout this section we will focus primarily on a linear extrapolation between observed bond rate and the UFR, illustrated schematically in Figure 5. This is partly out of a desire for simplicity, but is also consistent with much of current practice in Australia and New Zealand (including the approach mandated by the New Zealand Treasury for government accounting work). The conclusion of Section 5 will give some consideration to non-linear patterns of extrapolation.

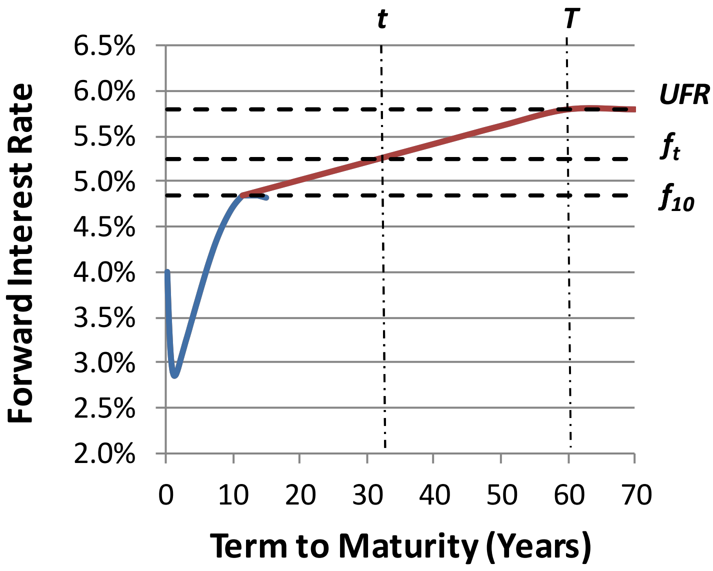

Figure 5 **Linear extrapolation model**

5.2 Insights from countries with longer dated government bonds

5.2.1 Background and data sources

In Australia our longest bond term is around 15 years. In countries such as the US, UK and Canada however government bonds are available at terms of 30 years or longer. So it seems to us forward rate curves from these countries would be a good place to start to understand how the Australian yield curve should be extrapolated, at least in the range 15-30 years.

The central banks from the US, UK and Canada make public their historical estimates of the forward rate curve for government bonds in their respective countries. We have used these forward rate curves for our analysis. Details on data used in the analysis, including some of the issues that surround the use of such data, are discussed in Appendix A. We have relied only on the forward rate curves available from 1 January 1998. We believe this is a good starting point for the current era of monetary policy characterised by low and stable inflation expectations.

5.2.2 A regression based analysis

Consider the simple linear regression equation:

$$f_{s+t} = \alpha_{s:t} + \beta_{s:t} f_s$$

Here the intercept $\alpha_{s:t}$ allows for the UFR as well as any (fixed) term premiums across the yield curve, and $\beta_{s:t}$ is the linear dependence of the forward rate at term $s+t$ on the forward rate at term s. If the schematic in Figure 5 is correct, then $\beta_{s:t}$ should be a good indicator of progress towards the UFR; if the slope is close to 1, this implies the forward rate at $s+t$ is still moving in sync with the rate at s, so no reversion to the UFR has taken place. Note too that if the slope was consistently close to 1 as t grows, this would be strong evidence against any reversion to the UFR. Conversely, if the slope is close to zero then the forward rate at $s+t$ is largely independent of the rate at s, suggesting that it has reverted to a constant level. Values in between can indicate the rate of progress, so a value of 0.5 indicates a place on the extrapolation halfway along the reversion.

To illustrate, we regress f_{20} against f_{10} using the UK data. The resulting fit is $f_{20} = 0.78\% + 0.74f_{10}$. The slope parameter suggests a heavy relationship between the two, but some evidence of mean reversion. We can then hold $s = 10$ fixed, and vary t between 0 and 20 to build up a more complete picture. The results are presented in Figure 6.

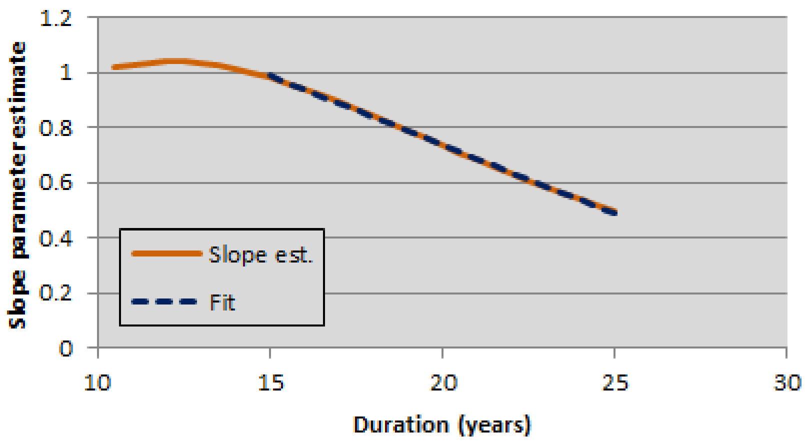

Figure 6 **Slope (β) coefficients for UK forward rates regressed against f_{10}.**

The results here are surprisingly clear. The coefficients are relatively stable until term 15, where they begin to decrease at a rate very close to linear. Extrapolating that rate of decrease would lead to the slope equalling zero (and thus reaching the UFR) at duration 35 years. We discuss the uncertainty around this estimate below.

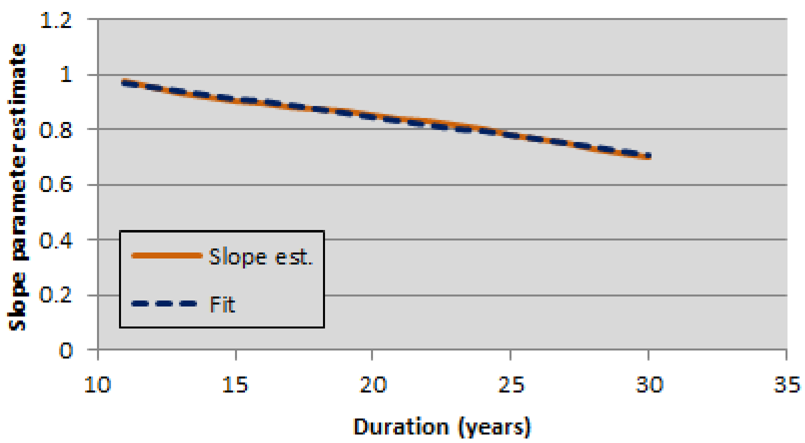

Figure 7 **Slope (β) coefficients for USA forward rates regressed against f_{10}.**

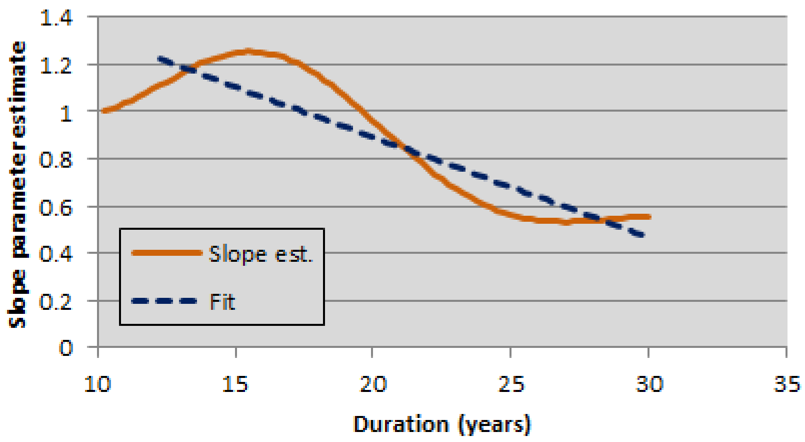

Figure 8 **Slope (β) coefficients for Canadian forward rates regressed against f_{10}.**

Figures Figure 7 and Figure 8 show the corresponding analyses for the USA and the Canadian forward rate data. The USA result again has a clear interpretation; a linear decay from duration 10, but at a significantly slower rate than the UK result; the point estimate for reaching the UFR is 83 years.

The Canadian results are the most difficult to interpret, due to the highly nonlinear shapes observed in the slope estimates. We believe this is partly due to the original methods used to construct the forward rate curves, and partly due to some unique patterns seen in the Canadian data at durations 15-20. Despite this, there is still reasonable evidence for a UFR reversion effect, with slopes generally decreasing with duration. For completeness we have attempted to impose a linear shape to estimate the term at which the UFR is reached, 41 years.

We can attempt to get a rough measure of uncertainty on the point estimates by bootstrapping the historical data series and refitting the linear trends. The table below summarises our results.

Table 2 **Regression results for the linear extrapolation model (Equation 1). Standard errors in brackets**

Country	*Duration decay starts*	Duration when reach UFR	*95% confidence interval*
US	10	82	(55, 168)
UK	15	34	(31, 40)
Canada	10	41	(35, 47)

The large degree of uncertainty, particularly in the slow decay seen in the USA data, is reflected in the relatively large confidence intervals. However, the results from this analysis appear clear enough to allow us to conclude that:

- There is reasonable international evidence for reversion to a UFR;
- The reversion starting at somewhere between duration 10 to 15 is plausible; and
- The reversion is quite slow – it is not complete before duration 30, and in some cases is considerably longer.

5.2.3 Principal components analysis of yield curves

Principal components analysis (PCA) is a means of determining what yield curve shapes account for the majority of the variation in the yield curve over time. PCA has been used in a variety of finance contexts, including yield curves (Litterman and Scheinkman, 1991), Foreign exchange (Avellaneda and Zhu, 1997) and equities (Gourrieroux et al, 1997 and Laloux et al, 1999). We have given a brief description of the PCA approach in Appendix B; the most important feature is that the shape of the components shows the typical movement of the yield curve around the mean. This is illustrated in the following schematic:

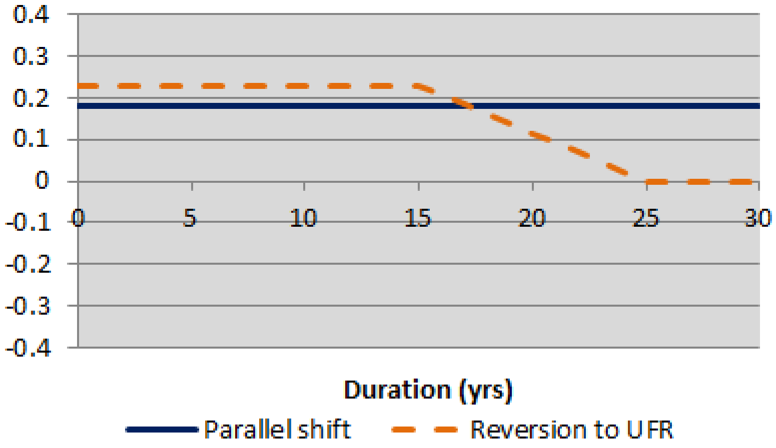

Figure 9 **Expected principle component curves under reversion to UFR and parallel shift hypotheses**

The two curves shown in Figure 9 are the expected shapes under two different hypotheses. The flat line is the component expected when most movement in the yield curve is attributable to parallel shift – the short and long durations move in equal amounts. The reversion shape is what would be expected when the forward rates return to a UFR type average in the long term (here after 25 years); the shape of the curve means that there is more movement in the short part of the curve and little at the high durations, where it remains closer to the UFR. By examining the shape of principal components in other countries we can gain significant insight as to the speed for which forward rates might return to the UFR.

Data for the analysis was the same as that used in the previous regression analysis.

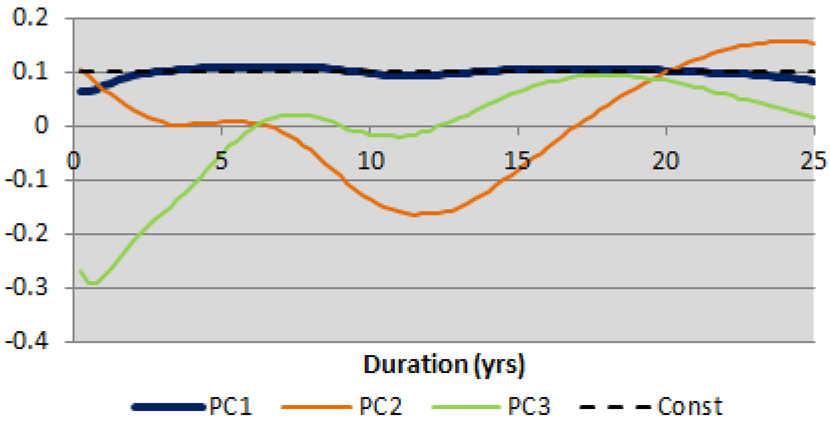

Figure 10 **Principal components analysis of Canadian yield curve**

Figure 10 shows the Canadian principal component analysis. The first heavy blue component dominates the results – it accounts for over 80% of the total variation. This curve is relatively flat, which we can interpret to represent shifts in the yield curve that are close to parallel. In addition to the primary principal component, there are other shapes that explain other types of movements in the yield curve. The second curve (8% of variation) is a flex shape for when medium durations move relative to short and long rates, and the third curve (6%) is a flex primarily at the short durations. The other 6% of total yield curve variation is accounted for by the remaining principal components, which are not shown here.

The shape and heavy importance of the first component (ten times as much variation explained compared to any other component) suggests that the **bulk of the yield curve movement is explained by parallel shift**. This fact is consistent with the observations of Section 4.3.1 about the behaviour of time-varying term-premia. That said, there is some evidence that the component is decaying from about duration 16 onwards; this can be seen as the slight decrease observable in the blue curve from this term. While this decay is slight, it is statistically significant. We examine these trends in further detail at the end of this section.

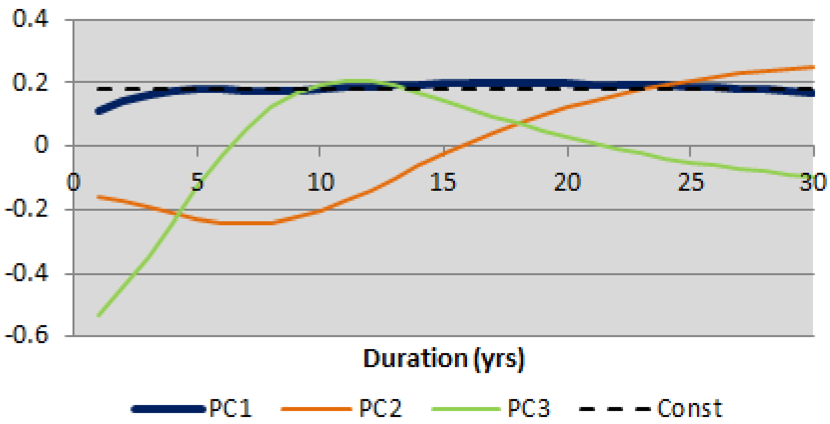

Figure 11 **Principal components analysis of USA yield curve**

The USA decomposition shows very similar results to Canada, with the leading component accounting for 80% of the total variation. The shape of the first component is again very flat, with a slight decay (decrease in the blue line towards zero) from duration 17.

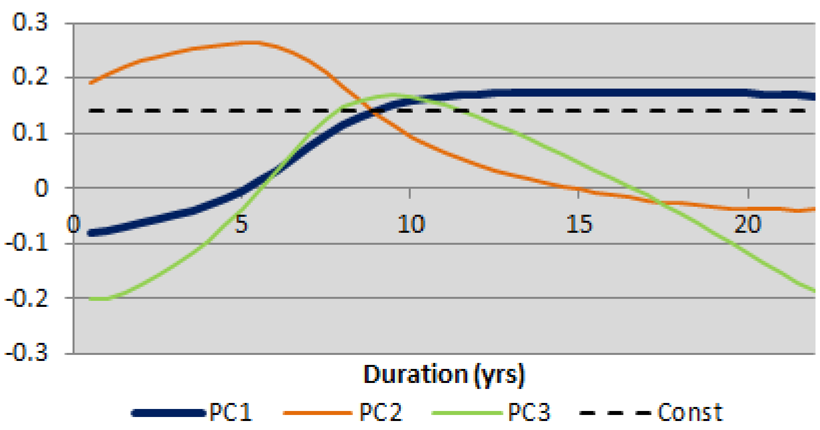

Figure 12 **Principal components analysis of UK yield curve**

The UK decomposition is the most unusual. Its shape is due to the very large drop in short duration forward rates compared to the longer duration ones observed over the last 15 years. Despite this, the pattern observed at longer durations is consistent with Canada and the USA; the shape is very flat, but with a slight decay at higher terms, here beginning from 15 years.

We observe that each of the leading components show some decay towards zero at high durations. Although slight, these trends are statistically significant (if we take bootstrap samples and refit, the trends persist), and the (assumed linear) slope of the downtrend can be used to estimate the point at which the curves reach zero, which is analogous to reaching the UFR. The zero point corresponds to reaching the UFR because this is the point at which there is no (or relatively little) movement in the forward rates. This provides a further check to the results in Section 5.2.2. The results are summarised in the table below – we have given estimates of the slope, the duration at which the curves reach zero using this slope, as well as 95% confidence intervals on these estimates (using bootstrap resampling).

Table 3 **Results for extrapolation of decay in leading principal component**

Country	Starting duration	Slope	95% CI	**Duration when reach UFR**	95% CI
Canada	16	-0.0023	(-0.0032,-0.0015)	64	(46, 101)
UK	15	-0.0026	(-0.0037, -0.0018)	84	(60, 122)
USA	17	-0.0021	(-0.0027, -0.0016)	110	(87, 147)

The results between countries are strikingly similar, despite the natural differences expected between countries and the sensitivity of the duration estimates to small changes in slope. In particular, all three confidence intervals include the duration range (87,100).

While there are issues with the analysis (such as the statistical validity of performing regression on a PCA), the results do point to a flat shape to the forward rate curve, with very slow decay to the UFR.

5.2.4 Reconciling the regression and PCA results and other methodological considerations

The PCA analysis gives a reversion period significantly longer than the regression analysis. While we do not completely understand the reasons for this, we favour the results of the regression analysis as is the most direct test of the problem of extrapolation: given the last reliable estimate of the forward rate what projected future forward rate will give me the least prediction error? Further the PCA analysis only analyses the behaviour of the first principal component and ignores other sources of variation. However, we like the insights provided by PCA – that the bulk of the yield curve movement is explained by parallel shift and we are pleased that it provides results broadly consistent with the regression analysis.

5.3 Discussion of results and other research

The concept of extrapolation to a long term rate over a long period is not unique to this paper. Two relevant examples are:

- QIS 5 undertaken for Solvency II suggested that convergence to an UFR should be reached between terms of 70 – 90 years; while
- The method presented by Barrie and Hibbert (2008) suggested convergence should occur by terms of around 100 years.

The approach used by Barrie and Hibbert to arrive at their 100 year convergence period is not in the public domain, but as far as we can tell it seems to involve making subjective judgements about what forward rate volatility should be at longer terms based on what it was in measured forward rates up to about terms of about 15 years. An approach obviously requiring a significant amount of judgement, but from what we have seen, it doesn't appear entirely unreasonable and it has also been endorsed by the CRO Forum (2010).

Both the QIS 5 and the Barrie and Hibbert approach involve extrapolations that are not linear – the paths chosen have a decaying rate – which would lead to a somewhat longer time to reach the UFR.

The latest indications are that the Solvency II guidelines will require a significant earlier convergence term than was initially suggested in QIS 5 – possibly somewhere between 10 to 40 years after the last reliable (in this case liquid) term – but this seems to be driven in part by some of the stability concerns outlined in Section 2.

The choice of non-linear paths is partly a tension between simplicity and aesthetics, given that identifying the true shape involves considerable uncertainty. We believe the linear path is reasonable, but recognise that there are more sophisticated techniques out there such as:

- The Smith-Wilson technique (described in CEIOPS, 2010); and
- The use of the Nelson-Siegal functional form (described in Barrie and Hibbert, 2008).

Discussions of the pros and cons of each of these methods can be found in the cited references. An advantage of a method such as Nelson-Siegal is that it allows one to match the slope of the fitted curve at the start of the extrapolation. This would be an advantage if the slope of the fitted curve was rising relatively rapidly or if it was turning downwards and away from the UFR at its final point.

5.4 Final thoughts on the extrapolation path

The analysis of yield curves in countries with longer dated securities indicates that it may revert to the UFR somewhere between duration 40 and 100 years, with about term 60 years a reasonable recommendation from our results. This is mostly based on the average of the regression results, but gives some small weight to the longer PCA estimates. While there may be instances where quicker convergence to the ultimate forward rate could be justified - for example where the fitted forward curve was continuing to rise rapidly near its final terms - the practice of just simply assuming quick convergence, say within a 5 or 10 year period, cannot be justified.

6 Implications for hedging balance sheet risk

It is possible to use a duration matching strategy to hedge risks that are beyond the longest term assets available, if you are allowed to take short positions in assets, and have an estimate for the risk free rate beyond the longest term asset. We have not seen this hedging concept in the actuarial literature (in which case it is overdue) but we would expect it to be common knowledge amongst traders that hedge long positions.

Given that this hedging approach requires an extrapolated yield curve, a natural question is whether hedging performance is better or worse if the extrapolated yield curve is assumed to revert quickly to the UFR?

We have attempted to answer this question by testing a duration matching hedging strategy under two alternative assumptions:

1) That the forward rates reach ultimate level very slowly, say by a term of 80 years (the point-in-time market consistent approach); or

2) That the forward rates converge to ultimate very quickly; say by a term of 20 years (the stability approach).

The tests have been performed using historical data from the Australian Government Bond market.

6.1 Details of the hedging strategy

The hedging strategy is as follows. Suppose we have a distant fixed liability L_T, payable at time T in the future. Suppose further that there are two shorter dated risk-free zero coupon bonds available in the market with durations s and t, with $s < t < T$. Then it is possible to match the present value and the (modified) duration of the liability by shorting A_s and going long A_t. In formulae, it is possible to choose loadings a and b, with $a < 0 < b$, such that:

$$aPV(A_s) + bPV(A_t) = PV(L_T)$$

And

$$Dur(aA_s + bA_t) = Dur(L_T)$$

Here Dur is the modified duration measuring the price sensitivity to changes in interest rate y,

$$Dur(L_T) = -\frac{1}{PV(L_T)}\frac{\partial PV(L_T)}{\partial y}.$$

Thus we have duration matched our long liability using shorter term assets. This portfolio can be rebalanced at regular intervals to ensure duration remains matched. We make the following comments:

- There is no reason why only two assets have to used, or that the assets are zero coupon bonds – these are simplifying assumptions for illustration and testing purposes. In fact, more complex duration matched portfolios are likely to exhibit superior properties (e.g. better convexity properties);

- The concept of duration matching for hedging is a type of immunization (Redington, 1952). In standard immunization the portfolio is chosen to ensure the convexity of the assets is greater than that of the liabilities, giving small (second order) profits in the presence of parallel shifts in the yield curve. In our situation the convexity of the assets is always less, exposing possible second order loss. This may be a secondary issue in practice, if:
  - There is enough non-parallel shift to swamp these second order effects; and
  - The existence of term premiums (usually higher yields for longer bonds) means that this strategy collects some of this long term premium, partly offsetting convexity losses;

- This approach results in large opposing long and short positions in government bonds. This obviously comes with a number of associated costs which we have not evaluated and so we are unsure how attractive this hedging strategy would actually be in practice; and

- The hidden assumption in the equations above is that $PV(L_T)$ is known. Of course, this requires extrapolation of the risk free rates between terms t and T. Thus the effectiveness of the hedging strategy depends (somewhat) on the appropriateness of the extrapolation assumptions. We discuss this further below.

6.2 Australian market examples

We have run a number of historical experiments testing this strategy. The general setup is:

- The distant liability of $100 falls due in 20 years;

- We invest $PV(L_T)$ in the market to cover this liability;

- We invest in only 4 and 10 year zero coupon bonds, returning the risk free market rate;

- We rebalance every quarter. This involves closing the positions in the 3.75 and 9.75 year bonds and opening a new position in 4 and 10 year bonds, still matching the duration of the liability (19.75 after the first quarter etc);

- After any quarter we can calculate the asset/liability ratio to assess how well we have hedged balance sheet movements.. The variability of the asset/liability ratio around 100% indicates how good the hedging is; and

- We close the position after 10 years of managing (as at that time we can perfectly match the liability by investing in a 10 year ZCB, crystallising any gain/loss).

Note that the assumption of quick reversion to the UFR has the effect of reducing the modified duration of the liability, because a portion of its value is stabilised by the path assumption. A linear reversion to the UFR between terms 10 and 20 gives a modified duration of the liability of 14.9 years. Thus this assumption requires a less extreme short-long position of bonds, compared to the slow reversion.

The figure below shows the inflated value of our assets if we use a flat forward rate (a reasonable proxy for any slow path of return) beyond 10 years, starting at time 30 June 1995. That is, we set the forwards between term 10 and 20 at the same level as the 10 year forward rate.

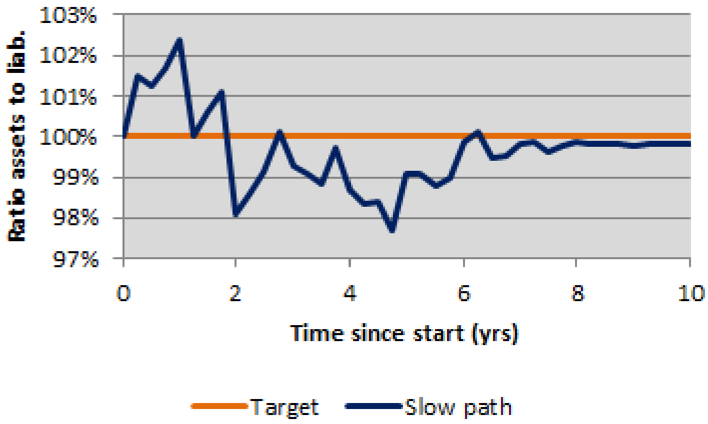

Figure 13 **Long liability hedging strategy performance, starting June 1995**

The performance of this strategy from 1995 is remarkably good despite large changes in interest rates over the date range. At the end of the management period a tiny loss is made (about 20c on the $100), but over the course of the ten years of management the implied value hugs the target 100% value fairly tightly. One way to quantify this is the average absolute departure from 100% (that is, the average distance between blue and orange lines), which in this case is 0.8%.

We can then compare the hedging performance if we use an alternative yield curve beyond ten years. Here we test a mean reversion assumption that the forward rate returns to 6.0% (consistent with that recommended in Section 4.4 without allowance for the convexity assumption) linearly over the duration years 10 through 20. The resulting hedging performance is shown in the next figure.

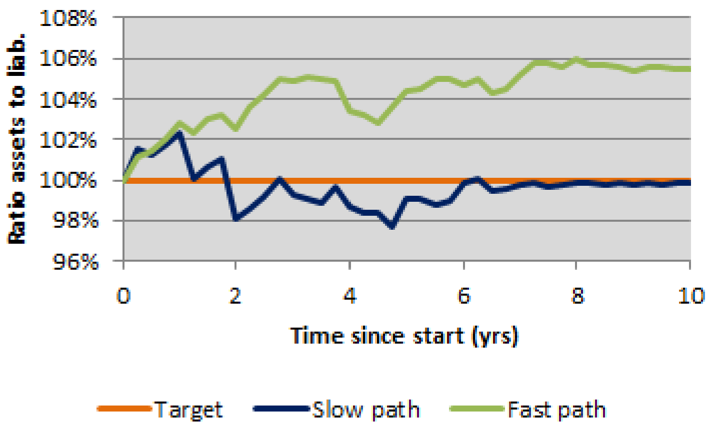

Figure 14 **Comparison of hedging performance, starting June 1995**

The performance is markedly worse, with an average absolute difference of 4.2%. In this case the eventual result is a substantial profit (due to the falling rates in the late 90s), but this is not the aim of a hedging program. The slow path seems far superior at reducing interest rate risk. We note however, that while we have assumed a constant UFR over the 10 year period, in practice one could imagine a higher UFR being adopted in the first few years from 1995, with some lower rate adopted towards the end of the 1990's when lower inflation expectations had been cemented. However, we tried a number of alternative scenarios anticipating such a UFR revision without changing our general conclusion - that the market consistent approach gave a better result.

We can run the experiment over different historical periods. In the figures and table below we summarise results for June 1998, 2000, and 2002 starting points.

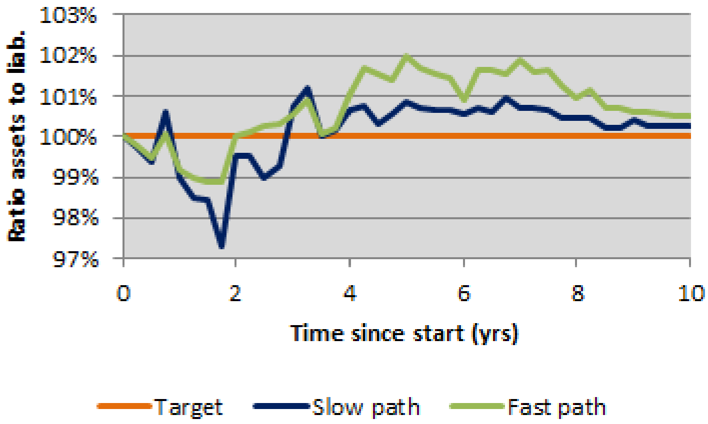

Figure 15 **Comparison of hedging performance, starting June 1998**

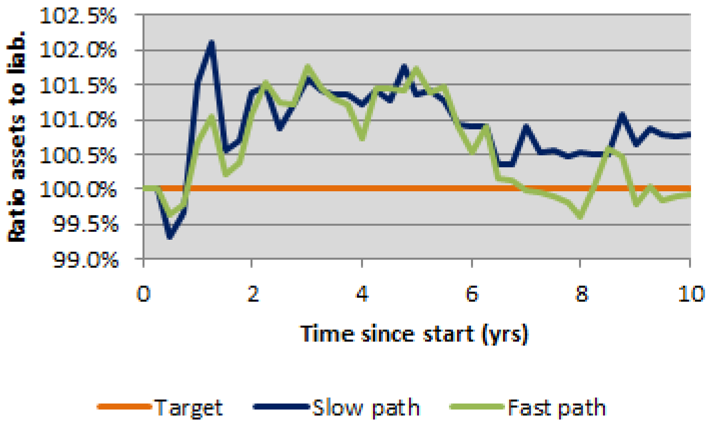

Figure 16 **Comparison of hedging performance, starting June 2000**

With regard to the June 2000 result, we note that the hedging performance is fairly similar over the first six years, where they then diverge. Even though the fast path result ends up closer to 100% after 10 years, we prefer the slow path performance due to its flatness from years 6 to 10. To express this another way: if a manager had "booked" the 0.9% profit after six years, there would be minimal further hedging loss over the subsequent four years using the slow path assumption.

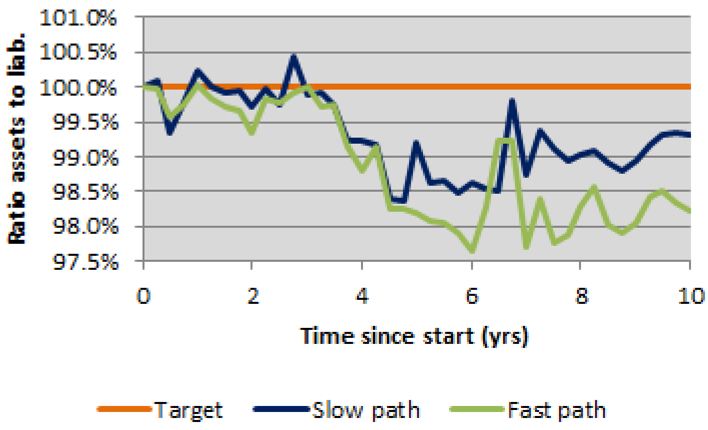

Figure 17 **Comparison of hedging performance, starting June 2002**

Table 4 shows the average distance between the each hedging performance line in the figures and the target 100% line. The slow path outperforms in all except the year 2000 (discussed above) according to this metric.

Table 4 **Comparison of hedging performance using average absolute difference metric**

Year (30-Jun)	Slow path	Fast path	Ratio
1995	0.8%	4.2%	5.64
1998	0.6%	0.9%	1.44
2000	0.9%	0.7%	0.73
2002	0.7%	1.1%	1.57

To summarise, the differences in the later experiments are less stark than the 1995 result, with the slow path clearly better in 1998 and 2002, but more ambiguous in 2000. The extreme difference in the first date range is not completely unexpected; the period 1995 to 1998 was unique due to the very strong decrease in yields from 9% to about 6%; this is the type of situation where a good hedging strategy is most critical.

In summary, the approach to hedging appears fairly legitimate from both a theoretical and historical data perspective. It gives hedging error of less than 1% in all experiments. The fast path is inferior, giving further evidence that a slower reversion to UFR is closer to "truth". Investing only in the 10 year bond (a common strategy in practice – "hold the longest term bond possible") by contrast would result in significant losses in the first and last experiments, where a slide in the 10 year forward rate is observed.

7 Conclusions and recommendations

We can summarise our findings relatively succinctly:

- **The yield curve up to 2 years before the longest dated bond can be estimated reliably.** For the last year or so, the noise and method of fit can cause significant (relative) error;

- **There is reasonable international market evidence for reversion to a flat long term forward rate**. This rate is reached via extrapolation from the end of the observable yield curve;

- **The rate of reversion is slow**. We believe term 40 is about the minimum point to reversion based on the bond markets examined, with a central estimate closer to term 60. This conclusion rests on the assumption that the unconditional forward rate has been stable over the period 1998 to 2012;

- **Linear path reversion is plausible, with other approaches possible.** Non-linear paths may have implications for term at which the long term rate can be reached.; and

- **Long term risk free hedging is possible, at least for moderate term extrapolations of the yield curve**. The lack of long dated risk free assets does not mean no hedging strategies exist for long dated liabilities. The methods proposed here appear plausible and give reasonable results on historical data for cash flows up to 20 years in the future. However, they do require a reasonable estimate of yields beyond the observable range.

We believe that these results make significant contributions to actuarial assumption setting.

Finally we make two comments. First, the current Australian accounting and prudential standards, along with the Actuaries Institute’s PS300, are not entirely clear, in our opinion, in stating the philosophy that should be used for yield curve extrapolation. Attempts to rectify this would improve consistency in the financial reporting of long-term liabilities.

Second, the conclusion that the rate of reversion is slow means there is an inevitable tension between market consistency and liability stability; while it may be attractive from a stability standpoint to have a yield curve that quickly reverts to an assumed long-term average, we believe the valuation would be incorrect if appropriate long term risk free assets existed. The tension deserves further consideration in the actuarial community.

8 References

1. Avellaneda, M. and Zhu, Y. (1997). An E-arch model for the term structure of implied volatility of FX options, Int. J. Theor. Appl. Finance.

2. Barrie and Hibbert (2008) A framework for estimating and extrapolating the term structure of interest rates. Exposure Draft.

3. Bolder, D J and Gusba, S (2002) Exponentials, Polynomials, and Fourier Series: More yield curve modelling at the Bank of Canada. Bank of Canada working paper.

4. Bolder, D J, Johnson, G., Metzler, A. (2004). An empirical analysis of the Canadian Term structure of zero-coupon interest rates. Bank of Canada working paper.

5. Carlin, S (2010) Thoughts on QIS5 yield curves. Barrie+Hibbert Insights.

6. CEIOPS (2010) QIS5 Risk free interest rates. Extrapolation Methods.

7. CRO Forum (2010) CRO Forum Best Practice Paper - Extrapolation of Market Data.

8. Dimson E, Marsh P, Staunton M (2000) Risk and return in the 20th and 21st. Business Strategy Review, Volume 11, Issue 2, 1-18.

9. Finlay, R and Chambers M (2008) A term structure decomposition of the Australian yield curve. Reserve Bank of Australia Research Discussion Paper.

10. Gourrieroux, Ch., Scaillet, O. and Szafarz, A., (1997). *Econometrie de la Finance*, Ed. Economica, Paris.

11. Gurkaynak, R. S., Sack, B., and Wright, J. K. (2007). *The U.S. Treasury Yield Curve: 1961 to the Present*. Finance and economic discussion series, Federal Reserve Board, Washington D.C.

12. Hibbert, J (2008) Market-consistent valuation of ultra long-term cash flows, Barrie+Hibbert Insights.

13. Hibbert, J (2012) Alternative views on extrapolated yield curves: A fundamental question remains unanswered. Barrie+Hibbert Insights.

14. Kim D H and Orphanides A (2005) Term structure estimation with survey data on interest rate forecasts. Federal Reserve Board Finance and Economics Discussion Series.

15. Kim D H and Wright J H (2005) An arbitrage-free three-factor term structure model and the recent behaviour of long-term yields and distant-horizon forward rates. Federal Reserve Board Finance and Economics Discussion Series.

16. Laloux, L., Cizeau, C., Bouchaud, J.-P. and Potters, M., (1999). Noise dressing of financial correlation matrices, Physical Review Letters, **83,** 1467.

17. Litterman, R. and Scheinkman J., (1991). Common factors affecting bond returns, J. Fixed Income 55-61.

18. Miller, H., (2010). Towards a better inflation forecast. Paper for the 2010 General Insurance Seminar, Institute of Actuaries Australia.

19. Redington. , F. M., (1952). Review of the Principles of Life Office Valuations. Journal of the Institute of Actuaries, **78**, 286-340.

20. Shiller R J (1990) The term structure of interest rates in *Handbook of Monetary Economics Volume 1*. Friedman B M and Hahn F H (eds).

21. Svensson, Lars E. O. (1994), Estimating and Interpreting Forward Rates: Sweden 1992-4, National Bureau of Economic Research Working Paper #4871.

22. Waggoner, D (1997), Spline methods for extracting interest rate curves from coupon bond prices, Federal Reserve Bank of Atlanta, Working Paper series, 97-10.

APPENDIX A DATA USED IN THE ANALYSIS

Historical series of forward rates were available for the USA, UK and Canada. We have from the following sources:

- **Canadian zero coupon bond data:** daily ZCB yields, downloadable from the Bank of Canada Website. The ZCB curve has been estimated by Bank of Canada employees from market prices of bonds, again using the exponential spline methodology (Waggoner, 1997). Curve is given at quarter intervals for durations 0 to 20 years. Further details can be found in Bolder et al., 2004. We extracted an end of month series. We calculated half-year forward rates from the ZCB series;

- **UK zero coupon bond data:** monthly ZCB yields and forward rates, downloadable from the Bank of England Website. The ZCB curve has been estimated by Bank of England employees from market prices of bonds, using Waggoner' s variable roughness penalty model . Curve is given at half intervals for durations 0 to 25 years, although shorter in periods where longer dated bonds are unavailable. We extracted an end of month series.

- **US Treasury Yield curve:** Data series release in conjunction with Federal Reserve discussion paper by Gurkaynak et al. Yield curve estimated by Svensson's (1994) approach, another exponential spline based model.

We make further comments regarding the data:

- We have relied on the yield curve fitting approach done by each of the central banks. One consequence of this is that any biases in the fitting approach will be inherited. In particular, the degree to which smoothness is imposed on the yield curve could materially affect results. We have attempted to mitigate this by using durations for which there is reasonable market evidence, but also note that:

  - The results obtained have been remarkably consistent across the three countries. All three central banks have used different techniques for estimating the forward rate curves. Further the methods used spanned the two main classes of term structure estimation methods, smoothing splines and function based methods (Bolder and Gusba, 2002);

  - The standard deviation of forward rates at longer terms did not display any excess volatility; and

  - The standard deviation of the forward rate curve slope at durations above 10 years did not indicate that any constraints (apart from smoothness) were placed on the curve fits.

  Given all this, we feel that the fitted yield curves are a reliable basis for analysis.

- The regression and PCA analyses make use of data from 1998. We believe this is a good starting point for the current era of monitory policy characterised by low and stable inflation expectations. For earlier time periods, significant complexity would be needed to distinguish between changes due to real interest rates and those due to inflation expectations. One consequence is that results appear closer to the no reversion scenario - Figure 18 shows a slope very close to 1 when regressing USA 25 forwards against 10 year forwards; and

- There are some time periods where the estimated yield curve appears highly unstable in the tail. We have used our judgement in removing these time periods from the analysis. For the regression and PCA analysis, the only excluded time periods were for Canada, where the yield curve at high durations was quite unstable between 1998 and March 2002. It is also subjective whether to leave in or omit points relating to the GFC (see those coloured in the figure below). Our analysis has managed to retain them, which slightly increases the noise in the analysis. We believe that estimates for the reversion term to UFR would be slightly longer if these points were removed.

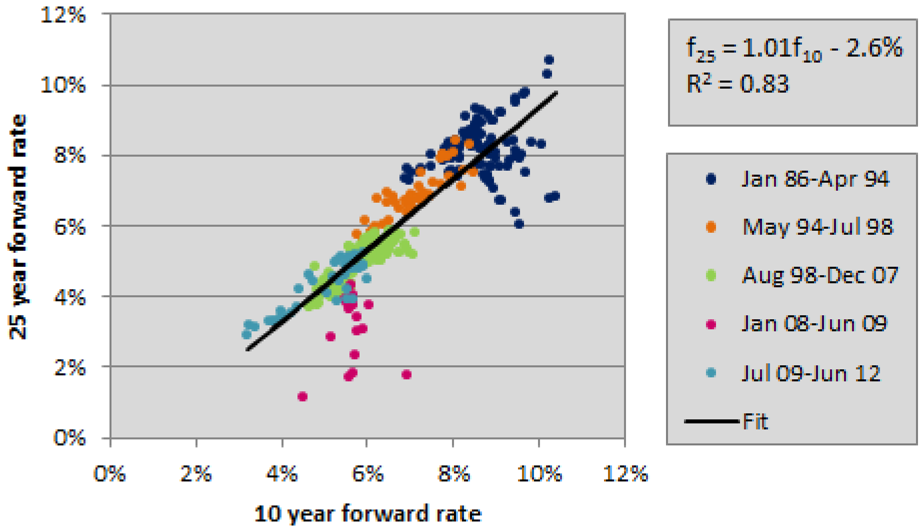

Figure 18 **Regression of USA forward rates, 1986 to 2012**

APPENDIX B BACKGROUND ON PCA

Principal component analysis (PCA) is a means of finding vectors that characterise maximal variation in a dataset. Suppose that we have n observations $X_1, X_2, \ldots, X_n$, where each observation is a p-vector $X_i = (X_{i1}, \ldots, X_{ip})$. The first principal component for this dataset is the p-vector $P_1 = (P_{11}, \ldots, P_{1p})$ such that:

- $\sqrt{\sum_{j=1}^{p} P_{1j}^2}$ = 1; and
- $Var\left(\sum_{j=1}^{p} P_{1j} X_{ij}\right)$ is maximised.

The first condition is a normalising condition (otherwise the maximisation would be satisfied by setting all the P_{1j} to be very large. The interpretation of the second condition is that we attempt to find the vector that captures as much variation possible. The shape of this vector, and the amount of variance explained, often yields important insights into the model.

Once the first component is found, subsequent principal components can also be found. These have the same properties as above, with the additional constraint that they are orthogonal to the components found so far. Because variances can vary between columns, it is common to scale variables beforehand. In the case of the yield curves scaling is not necessary.

As an illustrative example, consider the (simulated) dataset of students exam results presented on the next page. Results are for 100 students across five subjects (Maths, Science, English, History and Health). The first three principal components are:

	Math	Science	English	History	Health	% total var.
PC1	0.56	0.42	0.49	0.44	0.28	67%
PC2	-0.42	-0.51	0.50	0.54	-0.14	16%
PC3	0.24	0.17	0.13	0.00	-0.95	13%

We can interpret these results as follows:

- The first component is an overall “intelligence” measure, and explains two thirds of the variation. This accounts for students’ scores tending to rise and fall together. Note Maths and English have higher loadings, and health lower. This suggests that Maths and English are more important in determining overall intelligence, and Health less so;

- The second component is negative for Maths and Science, positive for English and History, and small for Health. This is a stereotypical difference between equation based and essay based courses;

- The third component is almost entirely loaded onto the Health subject. This suggests that there is a component of the Health result that is largely independent of the other subjects; and

- The first three components account for 96% of the variation. It is reasonably safe to conclude that the remainder is “noise”.

Thus the principal components have neatly summarised the data. In theory, we could use these vectors to report three numbers that neatly characterise each student – an overall intelligence score, an equation versus essay score and a health outperformance score. The first score in particular might be useful in ranking student performance relative to peers.

Table B **Data used in PCA example**

ID	Math	Sci	Eng	His	Heal	ID	Math	Sci	Eng	His	Heal	ID	Math	Sci	Eng	His	Heal
1	55	62	52	57	67	34	53	57	52	58	61	67	71	77	65	65	90
2	80	84	69	62	77	35	62	59	69	66	64	68	58	65	62	59	63
3	72	68	75	72	82	36	76	81	70	71	76	69	81	83	83	74	77
4	79	71	78	73	71	37	79	77	91	95	89	70	81	82	68	71	61
5	71	70	69	70	79	38	77	73	80	78	81	71	71	72	69	75	76
6	58	59	65	69	56	39	74	77	69	68	73	72	80	76	68	69	82
7	89	80	85	82	75	40	86	84	91	88	91	73	68	71	72	73	69
8	57	60	55	58	63	41	71	77	67	61	67	74	57	60	74	70	66
9	75	72	74	68	61	42	68	70	73	70	60	75	75	74	59	59	79
10	75	79	68	66	66	43	88	84	76	78	76	76	84	77	75	74	82
11	72	78	66	67	70	44	65	59	62	67	65	77	69	73	59	57	72
12	70	66	68	67	71	45	63	63	72	65	70	78	76	77	74	78	80
13	73	70	65	73	66	46	53	57	62	63	75	79	72	72	72	74	74
14	73	71	77	74	69	47	66	63	63	62	69	80	77	69	75	68	81
15	76	71	75	77	72	48	91	83	90	93	80	81	72	73	77	73	73
16	63	62	56	59	79	49	61	64	69	69	77	82	78	80	67	66	71
17	78	74	75	74	74	50	72	72	61	64	76	83	50	52	61	60	56
18	79	72	82	78	65	51	68	67	57	54	48	84	62	63	73	81	70
19	77	73	71	66	63	52	58	62	65	60	71	85	77	72	68	66	67
20	60	61	73	75	66	53	63	68	63	72	73	86	54	64	48	60	60
21	86	79	84	80	69	54	66	71	70	68	81	87	70	64	84	85	72
22	65	64	64	60	78	55	86	82	76	76	61	88	74	77	68	66	75
23	62	67	75	72	68	56	76	73	80	77	77	89	79	75	86	85	75
24	70	69	89	77	66	57	50	58	71	73	66	90	68	72	72	74	62
25	60	59	63	62	63	58	84	82	80	71	68	91	81	77	71	76	69
26	68	74	65	64	72	59	71	72	69	68	77	92	79	75	73	68	68
27	76	73	70	70	75	60	61	60	68	62	65	93	64	69	59	58	66
28	75	65	69	77	69	61	65	68	65	65	70	94	65	69	65	66	71
29	68	66	62	62	61	62	82	85	81	83	76	95	60	62	59	52	62
30	68	65	72	80	72	63	62	70	70	68	55	96	81	73	83	75	74
31	55	53	67	67	71	64	54	57	65	60	72	97	78	78	73	70	63
32	64	65	59	63	74	65	70	73	73	77	71	98	66	66	63	70	83
33	59	57	62	61	64	66	74	67	76	73	74	99	80	80	66	65	85
												100	72	75	77	80	77