ARTICLES

Connecting the dots: a yield curve perspective on New Zealand's interest rates

Leo Krippner

This article explains the concept of the yield curve, the pattern of interest rates by their time to maturity, and how that concept helps to provide some perspective on the multitude of different interest rates that exist in modern economies. The dominance of macroeconomic influences on the government yield curve are discussed, and then the additional factors of default risk and liquidity risk are introduced in the context of the bank and mortgage yield curves. The article also indicates how careful analysis of the yield curve may be used to back out information of interest to the Reserve Bank, such as the market expectations about the path of the Official Cash Rate (OCR) and future economic growth, and how interest rate influence the exchange rate.

1 Introduction

Interest rates play a fundamental role in the economy and financial markets. A glance at the business news illustrates how interest rates range from a financial investment in their own right to the tool the Reserve Bank uses to operate monetary policy. More broadly, interest rates, as the price of credit, affect people's decisions on consumption, investment, savings and borrowing. Given their pervasive influence, it is not surprising then that records of interest rates have existed over at least the last five millennia.[1]

When it comes to the modern global economy, there are certainly many different interest rates to contend with. For example, the average New Zealander will be familiar with the mortgage rates at which banks lend. But mortgage rates can vary considerably over time and by their term, or time to maturity. Figure 1 shows that, since 1998, the six-month fixed mortgage rate has varied between 5 and 10 percent. By time to maturity, as at June 2009, mortgage rates ranged from 5.39 percent for six months to 7.9 percent for five years. Such substantial differences in interest rates impact materially on borrowers' wallets, so it's little wonder that choosing when to fix, and the mortgage term, is an inevitable source of angst.

Besides mortgage rates, there are retail interest rates on bank deposits, non-bank deposits, personal loans, business loans, etc. At the wholesale level, there are interest rates on

Figure 1

Mortgage rates by time and term

(time to maturity)

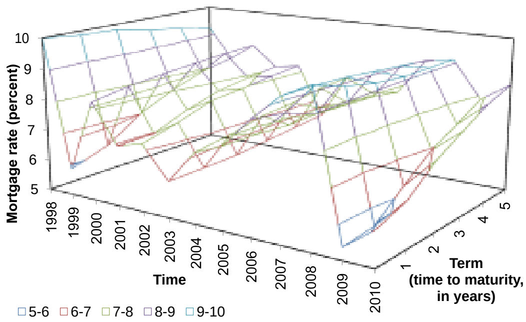

government bonds, bank bills, corporate bonds, etc. These interest rates all vary over time and by term, and can be found in almost all economies around the world (although this article will concentrate on New Zealand interest rates).

To help put the multitude of different interest rates in perspective, economists often construct so-called yield curves. Put simply, a yield curve is a plot of interest rates or yields with different maturities (but otherwise almost identical characteristics) observed at a single point in time.

As examples, figure 2 shows three key yield curves for the New Zealand economy as at 31 May 2010. The mortgage yield curve connects fixed mortgage rates with maturities of six months and one year to five years. The government and bank yield curve are two all-important wholesale benchmarks for interest rates throughout the economy.

[1] James and Webber (2000) chapter 2 gives a summarised history, and Homer and Sylla (2004) provides a very detailed history of interest rates.

Figure 2
The government, bank, and mortgage yield curves on 31 May 2010

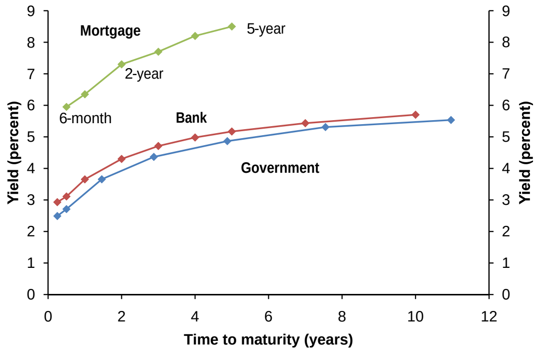

Once different classes of interest rates are grouped together within their appropriate yield curve as in figure 2, they each become easier to monitor over time and the relationships between different classes of interest rates becomes clearer. For example, one could simply say that the mortgage yield curve is roughly three percentage points above the bank yield curve, rather than specifying the gap for each individual maturity. As we shall also see, yield curves are also amenable to careful analysis that can reveal information about the economy and financial markets, and other aspects of interest to the Reserve Bank.

To follow through these ideas, this article proceeds as follows. Section 2 begins by discussing in more detail the government, bank, and mortgage yield curves from figure 2. Using the common themes that show up as influences on these yield curves, we then briefly introduce several other yield curves and their relationship to the government, bank, and mortgage yield curves. Section 3 describes the consistent government and bank yield curve data being produced and used by the Reserve Bank for yield-curve-related research. Section 4 introduces two topics from that research agenda. Section 5 concludes.

2 Three benchmarks: the government, bank and mortgage yield curves

2.1 The government yield curve

The government yield curve is an all-important benchmark because it represents the interest rates at which the state, almost always the largest and most creditworthy entity in any economy, can borrow at for different maturities. Other entities and individuals, being smaller and less credit worthy, should typically expect to borrow at higher rates. In this sense, the government yield curve may be seen as the base component in a hierarchy of yield curves within the economy, and so we spend some time discussing it.

Like any household, the government needs to borrow if it spends and invests more than the revenue it receives plus proceeds from any asset sales. Borrowing for the New Zealand government is undertaken by the New Zealand Debt Management Office (NZDMO) who issue (that is, sell) debt securities to investors. For the investors, the securities are assets (so they can be traded as discussed further below) and the cash payment investors make to buy the securities is government borrowing.

The two main securities issued are Treasury bills and government bonds and together these may be used to define the government yield curve.[2] Treasury bills are currently issued each week with maturities of three, six and 12 months. The interest return, or the yield to maturity, on Treasury bills results from the purchase price being lower than the face value that is promised to be returned on maturity, so the difference is effectively the interest payment.

Government bonds are issued for longer maturities. For example, the longest maturity bond currently on issue (15 May 2021) had a time to maturity of nearly 11 years when it was first issued.[3] The buyer of a government bond is promised fixed interest payments, so-called coupons, paid at

[2] Full details about these securities and the issuance process are available from the NZDMO website http://www.nzdmo.govt.nz. The site also provides information on Kiwi Bonds and inflation indexed bonds, the latter which we discuss briefly at the end of this section.

[3] Major international markets, such as the US, Japan, Germany, and the UK issue bonds with maturities of up to 30 years.

six-monthly intervals. The final interest payment and return of face value occurs on maturity. A simple calculation based on the purchase price, the coupon payments and principal repayment gives the yield to maturity of the bond. This is like an average return, or the internal rate of return, over the lifetime of the bond.

The advantage in issuing debt securities is that they may be traded in so-called secondary markets. That is, holders of the securities that have since decided they would rather have cash can sell them to investors that would rather have securities. The NZDMO has no further involvement, apart from registering the change of ownership so the promised future payments associated with the security are directed to the new security holder.

However, the secondary-market buyer of a security will only be prepared to purchase it at a yield that they think is appropriate in the current market, which will not necessarily be the yield for the original investor. For example, if interest rates have generally risen in the economy (we discuss possible reasons for this soon), then buyers in the secondary market will seek to pay a lower price for the security (thereby raising its yield) than the original buyer. Conversely, the security price will rise (lowering its yield) if interest rates have fallen.[4]

Through such market pricing and trading, the prices of the government securities change over time and therefore yields always reflect up-to-date market views. And grouping those yields together at any point in time always gives an up-to-date government yield curve.

Figure 3 illustrates the government yield curve from figure 2 along with two distinctly different yield curves from other dates. Using market terminology, the May 2010 yield curve is ‘positively sloped’, with yields rising as time to maturity increases. Conversely, the June 2007 curve is ‘negatively sloped’ (or inverted), with yields declining as time to maturity

Figure 3

Examples of the government yield curve

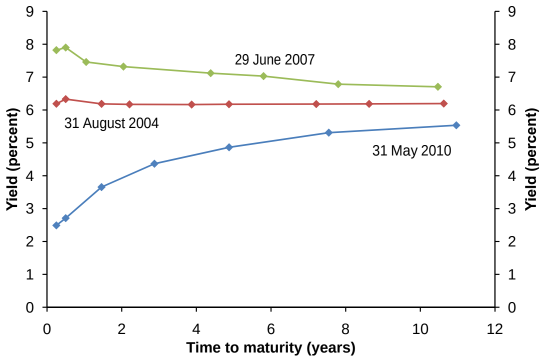

increases. The August 2004 yield curve is “flat”, with similar yields at all maturities.

So what causes the yield curve to change over time? The influences are many, but the most dominant are macroeconomic factors. As two key examples, we first discuss inflation and economic growth. Monetary policy and fiscal policy are then introduced as additional macroeconomic influences on the yield curve, although all of these factors are often intertwined. Finally, we note how similar global factors can also influence the New Zealand government yield curve.

Inflation has always been the enemy of bond investors, and for good reason. If one is promised a fixed payment in the future, as with a coupon or the return of face value from bonds, then the future purchasing power of that fixed payment (that is, the amount of goods and services it can buy in the future) becomes lower as inflation becomes higher. In other words, the real (that is, inflation-adjusted) return from bonds falls as inflation rises. The real return can even become negative if inflation rises sharply and unexpectedly, leaving an investor with less purchasing power than originally invested.

The yields of bonds therefore tend to include an element to compensate for inflation, or more precisely to compensate for expected inflation over the lifetime of the bond. So if inflation is expected to rise, market yields will typically rise to offset the anticipated impact on real yields. In general, if an investor is seeking a stable real return, then an expected

4 The precise relationships between security prices and yields are given by the Treasury bill and government bond price formulae available on the NZDMO website http://www.nzdmo.govt.nz. The concept is easiest to illustrate for Treasury bills: given the return of a fixed face value at maturity represents the return of the purchase price plus interest, the purchase price must fall for the interest return to increase.

increase in inflation should be compensated for by a one-for-one increase in yields.[5]

Real economic growth is also an important influence on the government yield curve. To illustrate this, it is useful to introduce the idea of opportunity cost. The opportunity cost of investing in a bond is the bond yield versus the return on alternative investments suitably adjusted for risk. So, for example, if prospective returns on shares rise, then the opportunity cost of holding an investment in bonds rises, and investors will have an incentive to sell their bonds to invest in shares. That would lead to lower bond prices, therefore raising bond yields to a level where the appropriate balance of risk and return between bonds and shares was attained.

More broadly, real Gross Domestic Product (GDP) measures the inflation-adjusted value of all final goods and services produced in the entire economy. GDP may be viewed as a measure of returns to all of the underlying factors of production in the economy (for example, labour, physical capital, entrepreneurship, etc.). Therefore, if GDP growth is expected to rise, the prospective returns from investing in at least some of the factors of production of the economy are rising, and so bond yields will tend to rise in response.

The Reserve Bank uses the OCR as its tool for the operation of monetary policy. The OCR determines the levels at which banks can borrow (against collateral) or lend overnight cash to the Reserve Bank. The setting of the OCR is an anchor for short-term government yields: banks won't buy or sell short-maturity government securities at yields materially different from the OCR while they know they have an alternative to lend to or borrow from the Reserve Bank via the OCR. Similarly, yields on longer-term government securities will be influenced to a large degree by expectations of where the OCR will average over the lifetime of the security. Those expectations are in turn influenced by the objectives of monetary policy.

The Reserve Bank sets the OCR with the intention of keeping CPI inflation between 1 and 3 percent on average in the medium term, while avoiding unnecessary instability in economic growth, interest rates and the exchange rate.[6] If market participants believe that the prevailing OCR might lead inflation to become inconsistent with the Reserve Bank's objectives (either by their own forecasts or relying on published Reserve Bank projections), they should naturally anticipate appropriate changes to the OCR to dampen or stimulate economic growth and inflation. That in turn changes the expected average of the OCR for all times to maturity, which therefore leads to changes in the shape of the yield curve.

Of course, forecasts of inflation and the appropriate decision on the OCR are made using a wide range of relevant macroeconomic data, including current inflation and economic growth (and other information such as wages, employment, retail sales, confidence surveys, exchange rates, etc.). So macroeconomic data, monetary policy and the yield curve may be viewed as being intertwined within a macroeconomic system. As new information comes to hand that affects expectations about the evolution of the macroeconomy and likely monetary policy responses, the yield curve will change to reflect that new information.

There is also another influence that monetary policy with an inflation-targeting objective has on the yield curve. The more strongly the market believes the Reserve Bank will adhere to its inflation objectives going forward, the less compensation the market will require in bond yields for the risk of rising inflation. That helps to keep the bond yields of all maturities lower in general, but particularly for longer maturities, where higher inflation would have a longer span to make its negative impact.

One final macroeconomic influence is fiscal policy (that is, government revenue and spending decisions), which also

[5] Indeed, because investors are taxed on their total return (inflation compensation plus the real yield), yields may rise even more to ensure the after-tax real return does not deteriorate unduly. For example, with a 10 percent bond yield and 2 percent inflation, the real annual return is (1+0.1)/(1+0.02)-1=7.8 percent. If that real return was taxed at 33 percent, the result would be 5.3 percent. However, the real after-tax return is actually (1+0.1x[1-0.33])/(1+0.02)=4.6 percent because tax is charged on the total return. The distortion gets worse with higher inflation.

[6] There are several caveats. The Reserve Bank website www.rbnz.govt.nz contains more detail on New Zealand's price stability framework, including the Reserve Bank Act 1989, the Policy Targets Agreement and related background information.

feeds into the macroeconomy/yield curve mix as discussed above. But fiscal policy also makes a contribution to net government debt issuance or repayment. That can directly affect the government yield curve via market perceptions about the government's future debt levels, and therefore the reliability of government promises to make interest and face value payments on its securities. That possibility is known as default risk, which we will discuss further in the next section.

One example of the impact of fiscal policy on the yield curve was during the period 2005 to 2008, when government budget surpluses led to a shrinking volume of government securities on issue. That resulted in lower yields across all bond maturities than would otherwise be expected purely from macroeconomic factors prevailing at the time.

Finally, it is worth noting that government yield curves around the world are also influenced by the same factors discussed above, as relevant to their economies. But global macroeconomic factors are also a major influence on New Zealand's macroeconomy, given the increasing interdependence between global economies (for trade and investment, etc.). Hence, global macroeconomic factors indirectly have a heavy influence on the New Zealand yield curve. Or more directly, from day to day, the movements in global yield curves are typically translated into similar movements in the New Zealand yield curve.

The bank yield curve

The banking system is essentially an intermediary between end borrowers and end lenders in the New Zealand economy and around the globe. That is, banks borrow from those with excess cash and then lend to those requiring funds, subject to making an appropriate assessment of the risk of that lending.

Of course, the amounts and maturities of the loans individual banks receive never precisely match the amount and maturities of loans they make. Hence, the banks use wholesale markets to borrow and lend from each other at different maturities to balance their own cash excesses or shortfalls and to better balance any maturity mismatches between borrowing and lending.

The main securities used in this process are bank bills and interest rate swaps. Together, these securities may be used to define a yield curve for the banking system. We will refer to this as the bank yield curve from this point but, for reasons explained below, it is important to bear in mind that the bank yield curve is distinct from the yield curve for any individual bank.

Bank bills are identical to Treasury bills, apart from their issuer being a bank rather than the government. A bank that requires funds will sell bank bills to other banks or institutional fund managers. That creates a liability for the seller/borrower and an asset for the purchaser/lender, which can again be traded in secondary markets.[7]

Interest rate swaps are a derivative agreement between two parties where fixed interest payments agreed at the time of the trade are exchanged for floating payments based on the realisation of future three-month bank bill rates.[8] Using market terminology, the 'swap receiver' is the party receiving fixed interest payments (and therefore making floating interest payments) and the 'swap payer' is the party making fixed interest payments (and therefore receiving floating interest payments).

As shown in figure 4, receiving fixed cashflows from an interest rate swap could be combined with a rolling investment in three-month bank bills to offset the floating payments of the interest rate swap. The end result would be like investing in a bond with a time to maturity equal to the interest rate swap tenure, which can be up to 15 years in New Zealand.[9] In reverse, paying fixed cashflows is like selling a bond.

7 There are also other avenues for banks to obtain short-term wholesale funds (for example, the currency forwards market), but the effective interest rates from those avenues are usually very similar to bank bills of the same maturity.

8 The terminology 'derivative' denotes that the security derives its value by reference to an underlying security or securities. In this case, the reference is to rates on three-month bank bills. Hawkesby (1999) provides an overview of derivatives commonly used in New Zealand markets.

9 Major foreign markets have interest rate swap tenures up to 30 years.

Figure 4
Stylised cashflow diagram for an interest rate swap and a bond

"Down" arrows are cash payments, "up" arrows are cash receipts, and longer arrows denote larger cashflows

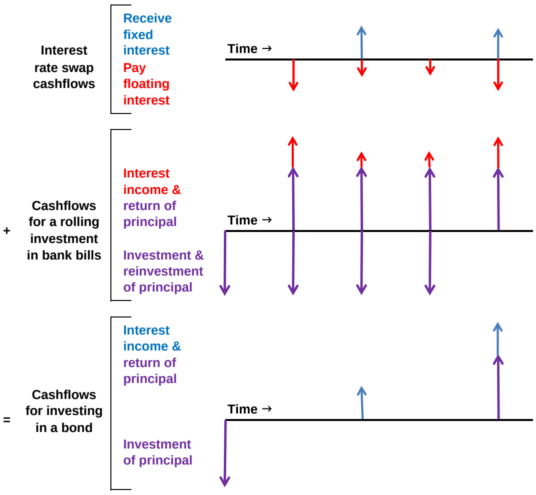

Note: principal amounts are typically much larger than interest amounts.

Therefore, the interest rate swap rate for a given maturity may be treated as the yield to maturity for a bond. However, such a "synthetic bank bond" is distinctly different from a bond issued by any individual bank where the default risk is to that bank for the term of the bond. Conversely, the rolling bank bill investment within the synthetic bank bond only ever has default risk for three-month periods, and the investment would be divided among the different banks of the New Zealand banking system for diversification (and, at any rollover state, less credit-worthy banks could be avoided altogether).

The influences on the bank yield curve are many. However, the macroeconomic influences of inflation, economic growth, monetary policy and fiscal policy already discussed for the government yield curve apply equally to the bank yield curve.

One major difference between the government and bank yield curves comes down to relative default risk: how reliable is the promise to have funds lent to the banking system returned relative to the government's promise of its future interest and face value payments?[10]

Government securities are usually regarded as having a lower risk of default than the securities that compose the bank yield curve. This is because governments generally have more flexibility and incentives to meet their obligations.[11] Conversely, banks are limited liability entities and so even a diversified rolling investment in bank bills, as for the synthetic bank bonds introduced above, carries an exposure to bankruptcy that would potentially result in some loss of value from funds lent.[12]

Another aspect that can differ between the government and bank yield curves is liquidity risk: essentially, how easy is it to trade securities for a cash amount close to their prevailing market value? Government and bank securities are generally regarded as being highly liquid in normal times, which is evidenced by the volumes traded and the presence of price makers in the interbank market with narrow bid-ask spreads. However, during periods of financial stress, such as the global financial crisis, government bonds are generally perceived to be much more liquid than securities in the banking system.

The combination of both relative default risk and relative liquidity risk means that the bank yield curve has usually remained above the government yield curve. That is, at each maturity, some extra yield is required to compensate investors for bearing the higher risk of default and the risk

10 Default risk is more precisely a combination of the risk of default (for example, missing an interest payment or the return of principal, or a company bankruptcy) and the expected value of the debt security given a default event (for example, the new market value of the security after a debt/capital restructuring, or the recovery of some cash as a creditor in the event of a company liquidation).

11 For example, raising taxes, cutting spending or, in some cases, printing money (albeit at the risk of generating higher inflation). Government securities are sometimes regarded as default-free, although it is more accurate to describe them as having low default risk. Reinhart and Rogoff (2009) contains a detailed history of explicit or implicit government defaults. Regarding views of the future, there are derivative securities known as credit default swaps available on the government debt of most countries, including New Zealand. They pay out in the event of a default, and so their prices provide a gauge of the market's perception of government default risk.

12 Retail deposits in banks are currently protected by the retail bank guarantee for up to $1 million.

of illiquidity. A notable exception occurred from late 2008 when the government introduced the wholesale and retail guarantees for banks. We shall revisit this episode in the discussion of figure 11 from section 3.

Because the government yield curve already captures macroeconomic factors, it makes sense to plot the difference, or spread, between the government and bank yield curves rather than plotting the bank risk curve outright. Examples of this government/bank spread curve at different points in time are presented in figure 5. Note that the maturities of the interest rate swaps don't match exactly to the maturities of the government bonds. The spreads beyond six months are calculated as the rate on the interest rate swap less a government yield curve for the same maturity obtained by linearly interpolating the actual government bond yields available.

Figure 5
Examples of the government/bank spread curve

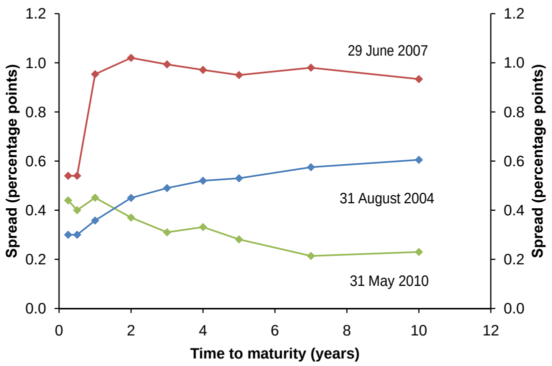

Just like the government yield curve itself, the government/bank spread curve can assume a variety of shapes over time. Hence, the August 2004 curve is positively sloped, May 2010 negatively sloped, and June 2007 fairly flat beyond the spread between Treasury bill and bank bill rates.

The mortgage yield curve

Like it or not, each individual bank is a business and seeks to make a profit for its shareholders to compensate them for the risk of being in the lending business. That profit mainly comes from the difference between the bank's revenue from lending less the cost of its funding and potential losses on lending.

On the revenue side, one avenue of bank lending is via mortgages to property owners, with the property involved as collateral. The mortgage conditions and different terms from major banks are usually standardised enough to consider the mortgage rates together within what we will call a mortgage yield curve. The Reserve Bank collects data for the mortgage yield curve as at the end of each month, which was the data used to create figure 1.

The cost of funding for an individual bank can be thought of as coming from a weighted average of the bank's unique yield curves for its different funding sources. For wholesale funding, the yield curve at short maturities will be around bank bill rates, but we have seen that the yields for longer-maturity bonds issued by individual banks will usually be above the bank yield curve.

To ensure they make a profit from mortgage lending, an individual bank will seek to ensure there is always a positive gap between its cost of funding at any point on its yield curve and the mortgage rates it has on offer. As the bank yield curve moves, the yield curves for the funding of individual banks will also move, and that will then flow through to the range of mortgage rates on offer and so a change in the mortgage yield curve.

For example, the prevailing two-year fixed mortgage rate will generally be related to the prevailing two-year interest rate swap yield, the three-year fixed mortgage rate to the three-year interest rate swap yield, etc. out to five years in New Zealand.[13] Floating rate mortgages are a bit different, being more related to the rolling average cost of funding from three-month bank bills rather than the three-month bank bill rate at a given point in time.

The connection between the bank yield curve, bank funding rates, and mortgage rates above means that movements in the mortgage yield curve will predominantly be influenced by the factors driving the bank yield curve. As we have seen, these factors include the macroeconomic influences on the government yield curve, including global factors, plus the additional influences of the differences in default risk and liquidity risk between the government and bank yield curves.

13 Fixed mortgage rates for longer terms are available internationally. For example, 30 years is a standard term in the US.

The additional influences for the mortgage yield curve relative to the bank yield curve should reflect shifts in the funding costs for individual banks, which we discuss further below, and the relative default and liquidity risk between mortgage borrowers and banks. Default risk is essentially the probability that borrowers might be unable to meet their payments for any reason. Liquidity risk reflects that, while mortgages can in principle be bought and sold via an appropriate legal process, it cannot be done as readily as for wholesale securities.

Analogous to the bank yield curve previously, it is more useful to plot the difference between the bank yield curve and the mortgage yield curve rather than plotting both yield curves outright. Examples at different points in time are given in figure 6. The May 2010 bank/mortgage spread curve is the difference between the bank and mortgage yield curves shown in figure 2.

Figure 6
Examples of the bank/mortgage spread curve

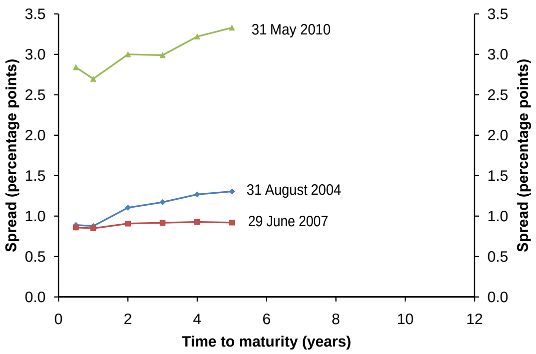

The bank/mortgage spread curve is noticeably wider as at May 2010 than past dates. This is mainly due to the effect of the global financial crisis, via several interwoven channels, on the cost of funding for individual banks. First, following the crisis, banks sought to reduce their susceptibility to disorder in short-term wholesale money markets, an arrangement that has recently been formalised by the Reserve Bank Core Funding Ratio requirements.[14] Specifically, banks have reduced their proportion of short-maturity wholesale funding at around bank bill rates, and have increased the proportion of retail deposits and longer-term bank bond issuance that are further above the bank yield curve. Second, the global financial crisis directly widened the spreads between bank-issued bonds and the bank yield curve, as markets now allow for higher bank default risk than in the past. Third, short-term retail deposit rates relative to bank bill rates have risen because the Core Funding Ratio treats deposits as substitutes for bank-issued bonds (given retail deposits are less prone to flight than short-maturity wholesale funding).

In summary then, individual banks now have a higher proportion of funding at more expensive rates (relative to the bank yield curve) than previously, but with the offsetting benefit of being less susceptible to potential adverse market events.

Ever more yield curves

We have now discussed three yield curves in detail, and have broadly categorised the influences on those yield curves in terms of macroeconomic factors, default risk and liquidity risk. From that perspective, we can consider other classes of interest rates within their own yield curves and even conjecture how they should usually sit relative to the government, bank and mortgage yield curves.

For example, the overnight indexed swap (OIS) yield curve is based almost purely on expectations about the OCR, which we have seen are dominated by macroeconomic factors. Not surprisingly then, the OIS yield curve is usually around par with the government yield curve.

As another example, personal loans and business loans will typically have higher default risk (and less security in the event of default) than mortgages. Therefore, the yield curves for personal loans and business loans will usually sit above the mortgage yield curve.

Finally, yield curves for individual corporates usually sit above the bank yield curve. That is especially so if they have higher default risk than the banks (which is often proxied by their relative credit ratings). However, the lower liquidity of most corporate-issued bonds in New Zealand means that even highly-rated securities can still sit above the bank yield curve.

14 The Reserve Bank website www.rbnz.govt.nz contains more detail on the requirements and related background information.

3 Yield curve analysis at the Reserve Bank

The discussion so far has indicated how different factors can influence the yield curve. Inverting that interpretation gives a basis for inferring information from the yield curve (or the spread curves). That is, if the yield curve is shaped in a particular way, what does that imply for expectations that the market holds about inflation, economic growth, default risk and liquidity?

Before addressing questions of that nature in section 4, we need to ensure that complete and consistent data sets for the relevant yield curves are available. The first part of this section describes the steps used by the Reserve Bank to create a comprehensive government yield curve dataset,[15] and the second section discusses the creation of a comprehensive bank yield curve dataset.

Note that the Reserve Bank has made representative yield curve data available in its *Bulletin* and on its website. However, from the perspective of undertaking research work in economics and finance, the representative yield curve data set has several unsatisfactory aspects. First, it mixes bank bill rates and government bond yields. Second, the bond yields are indicative only, given they use the yields for the government bonds closest to the maturities of one, two, five and ten years. Third, the data do not have the standardisation, to be introduced below, that is usually required for economic and financial research.

Government yield curve data

Securities

The data used to define the government yield curve at each point in time are the three-month and six-month Treasury bill rates and the yields of all major government bonds on issue at the time along with their exact maturities. The data are collected by the Reserve Bank and are originally sourced from an interbank broker. That source ensures up-to-date and consistent data at each point in time because the interbank broker provides market participants with indicative rates and yields for all securities even if trading is not occurring.

There are two technical issues with the data as it stands. The first issue is that the compounding period is different for the Treasury bill and bond data. Specifically, the compounding period for Treasury bills is effectively their time to maturity, while government bonds are compounded six-monthly (matching the spacing of their interest payments).

The second issue is that a bond yield is a 'fuzzy' measure of the interest rate for its maturity because the yield is effectively a weighted average of the interest rates applying to the different cashflows prior to and on maturity. Conversely, a Treasury bill yield is a precise measure of the interest rate for its maturity because there are no coupon payments. For that reason, securities like Treasury bills are often called zero-coupon securities.

A very convenient and useful standardisation for data on interest-bearing securities is to transform all rates and yields into zero-coupon continuously compounding interest rates.[16] Readers interested in an overview of how the Reserve Bank does this using the Nelson and Siegel (1987, hereafter NS) model are invited to read the box. The intuition is simply that the NS model represents almost any yield curve using just three components: a Level component that captures the general level of long-maturity interest rates; a Slope component that captures the extent to which short-maturity interest rates are below (or above) long-maturity interest rates; and a Bow component that captures a 'hump shape' where mid-term interest rates might lie above (or below) both short-maturity and long-maturity interest rates. The interest rate data that can be generated from the NS model is automatically zero coupon and continuously compounding.

15 **Krippner and Thorsrud (2009) contains more details on the process and the dataset itself.**

16 **While continuous compounding may sound strange, it can be seen simply as the natural limit as the compounding period gets shorter. To illustrate, assume a 10 percent annual interest rate, which would give a return of (1+0.1)=1.1. With semi-annual compounding (1+0.1/2)2=1.1025, quarterly (1+0.1/4)4=1.10381, monthly (1+0.1/12)12=1.10471, weekly (1+0.1/52)52=1.10506, daily (1+0.1/365)365=1.10516. As the compounding frequency N gets larger, the limit is (1+0.1/N) N=exp(0.1)=1.10517. That means a 10 percent continuously compounding rate would give the same return as a 10.517 percent annual rate. Or a 10 percent annual rate would give the same return as a 9.531 percent continuously compounding rate, because exp(0.9531)=1.1.**

Box

Fitting government yield curve data with the Nelson-Siegel model

The Reserve Bank uses the Nelson and Siegel (1987, hereafter NS) yield curve model to fit yield curve data.[17] The NS model uses Level, Slope, and Bow components to represent the yield curve, and the resulting fitted interest rate curve is a continuous, smooth and stable (that is, not unduly varying) function of time to maturity.

The NS interest rate curve has the following functional form:

$$R(t,m) = \overbrace{L(t)}^{\text{Level component}} \overbrace{+S(t)\left[\frac{\exp(-\phi m)-1}{\phi m}\right]}^{\text{Slope component}} \overbrace{+B(t)\left[\frac{1-\exp(-\phi m)}{\phi m}-\exp(-\phi m)\right]}^{\text{Bow component}}$$

where R(t,m) is the zero-coupon continuously compounding interest rate curve at time t as a function of time to maturity m. L(t), S(t), and B(t) are respectively the Level, Slope and Bow coefficients multiplied into the given functions of time to maturity illustrated in figure 7. The fixed parameter ϕ regulates the decay of the Slope function and the hump/decay of the Bow function. Estimating the NS coefficients from yield curve data observed at time t proceeds by finding the best fit to the prices of the securities used to define the yield curve.

Figure 7
Nelson-Siegel yield curve shapes

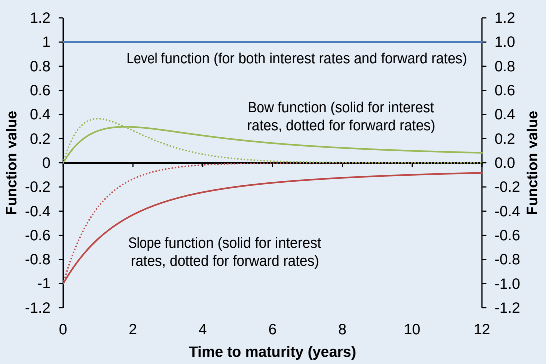

[17] The NS model and its variants have become very popular for fitting the yield curve, mainly from the perspective of their practical benefits and empirical successes. Bank for International Settlements (2005) provides an overview of routine central bank use of such models. Krippner (2006a) and Krippner (2009) provide further references on applications of NS models and also strengthen the theoretical foundation for their use.

A useful quantity that can be derived directly from the NS zero-coupon continuously compounding interest rate curve is the NS instantaneous forward rate curve. The NS forward rate curve turns out to be another simple function of time to maturity, and the components of the forward rate curve have intuitive shapes analogous to their counterparts in the interest rate curve. Specifically, the corresponding forward rate curve is:

$$f(t,m) = \overbrace{L(t)}^{\text{Level component}} \overbrace{+S(t)[-\exp(-\phi m)]}^{\text{Level component}} \overbrace{+B(t)[\phi m\exp(-\phi m)]}^{\text{Bow component}}$$

where f(t,m) is the forward rate at time t as a function of time to maturity m (or future time t+m). Figure 7 includes the functions of maturity for the NS forward rate curve.

Figure 8 shows an example of fitting the government yield curve data from May 2010 with the NS model. Together with our previous discussion on yield curve influences, this gives us a comprehensive perspective for discussing the yield curve shape at 31 May 2010. That is, we've already noted that the yield curve is 'positively sloped', with observed yields increasing by time to maturity. The NS interest rate curve 'fills the gaps' between the observed yields, clearly showing the rising trend of interest rates. Furthermore, the NS forward rate curve (derived from the NS interest rate curve) provides a quantitative indication of the expected path of the OCR, showing it rising to around 6 percent over roughly four years.[18] The expected rise in the OCR at that time was in turn consistent with market expectations of a gradual economic recovery following the 2008-09 global financial crisis.

[18] The NS forward rate curve may be interpreted as an indication of the market's expectation of the path of the OCR from the point in time of the yield curve observation, but without allowing for the effect of term premia. Term premia would need to be estimated and subtracted to produce an expected path of the OCR. That can be challenging, because term premia vary over time and by time to maturity. The OIS yield curve provides a more direct gauge of OCR expectations, but it only goes out to 12 months.

Figure 8

The government yield curve on 31 May 2010

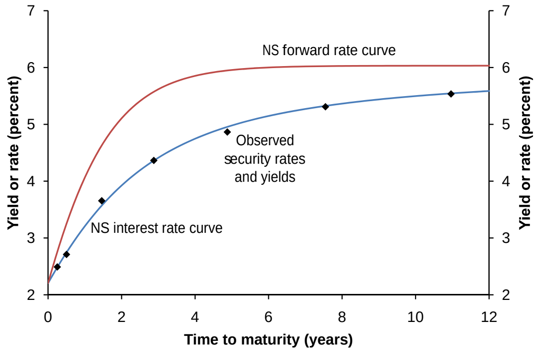

Repeating the curve fitting over the entire set of government yield curve data gives the estimated NS Level, Slope, and Bow components for the government yield curve over time. The results are plotted in figure 9.

Figure 9

Level, Slope and Bow coefficients for the government yield curve

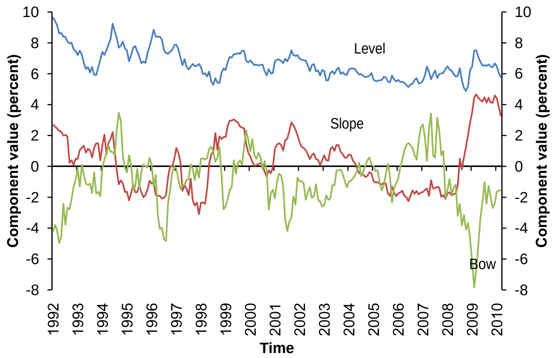

The positive slope of the yield curve for May 2010 is reflected in figure 9 by the positive value for the Slope component at that point in time. The government yield curve examples in figure 3 also included an observation from June 2007. That had a negative slope, which corresponds to the negative NS Slope component at that time in figure 9. The negative slope is consistent with an expectation of a declining OCR (from an historically high level of 8.25 percent), in line with expectations for a slowing economy. The yield curve example from August 2004 was flat, corresponding to the approximately zero Slope coefficient in figure 9. That is consistent with an expectation for the OCR to remain steady at around 6.00 percent. Note that combining the NS Bow component and Slope components for the forward rate curves provides more finesse about the path of OCR expectations at any given point in time than just using the NS Slope component itself.

Another point of interest in figure 9 is the mild downward trend in the NS Level component over the early part of the series. The NS Level component reflects the dynamics in the long-maturity part of the yield curve, which is the most sensitive to inflation expectations and inflation risk. Hence, the noted trend suggests growing market confidence in the Reserve Bank's operation of monetary policy to maintain low and stable inflation as time progressed.

The interpretations above show how the time series of NS components may be used as data to summarise the yield curve at any point in time, and its evolution over time. However, if zero-coupon continuously compounding interest rates for particular maturities are required as data, they can be evaluated from the NS interest rate curve at each point in time. The Reserve Bank calculates interest rates and forward rates for the times to maturity of three, six and nine months, and then each year from one to ten years.

Bank yield curve date

Securities used to define the bank yield curve

The securities used to define the bank yield curve are the three-month and six-month bank bill rates and the interest rate swap rates for the maturities of one, two, three, four, five seven and ten years. The data are all collected by the Reserve Bank, being originally sourced from an interbank broker. The data begin on 29 September 1997, when the ten-year swap was first collected as a standard quote.[19]

Fitting the bank yield curve with the NS model

As for the government yield curve, the Reserve Bank uses the NS model to fit the bank yield curve data observed at each point in time by finding the best fit to the prices of the securities. Figure 10 shows the results of fitting the bank yield curve data shown earlier from 31 May 2010, and then subtracting the fitted government yield curve at the same date.

[19] Data excluding the ten-year interest rate swap goes back further, and a standard quote for the 15-year swap also became available several years ago. We intend to incorporate that additional data into the bank yield curve data set, historically and going forward.

Figure 10

The government/bank spread curve on 31 May 2010

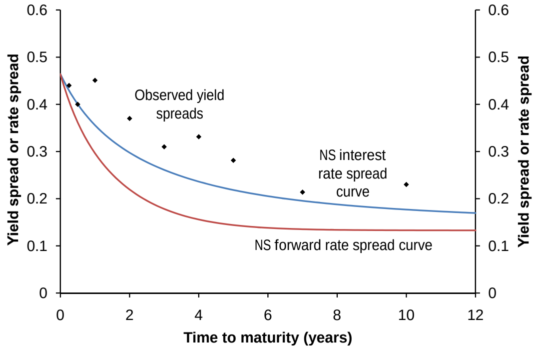

Note: The NS interest rate spreads are below the observed yield spreads beyond six months because the former are continuously compounding zero-coupon rates and the latter are yields to maturity.

Applying the NS model to each day of the entire data set gives the time series of the estimated NS Level, Slope, and Bow coefficients for the bank yield curve over the full sample. Figure 11 plots the spread between the NS Level component of the government yield curve and the NS Level component for the bank yield curve.

Figure 11

The government NS Level coefficient less the bank NS Level coefficient

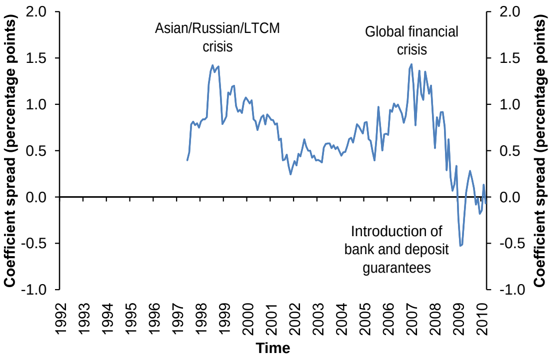

As we discussed in section 2, the spread in figure 11 should reflect changes in relative default and liquidity risk between the securities that compose the bank yield curve and government securities. That interpretation appears to hold. For example, the spread between the government and bank yield curves widened during the Asian/Russian/Long Term Capital Management financial crisis of 1998, and again during the early period of the global financial crisis. At the end of 2008, the spread narrowed markedly when the government guarantees on retail bank deposits and wholesale bank debt issuance were introduced.

Indeed, the spread became negative on occasions. On the face of it, that would suggest markets perceived the short-term default risk associated with synthetic bank bonds as lower than the default risk on long-maturity government bonds. That, in turn, would be consistent with the market anticipating a rising stock of government debt (and hence default risk) as the government's fiscal position turned from surplus to deficits. However, fluctuations in relative liquidity perceptions and the short-term effects of increased government bond issuance may also have made an impact.

4 Research projects

Given consistent sets of yield curve data, we can now turn to some examples of how that data may be used to infer information of use to the Reserve Bank. The first part of this section describes a research project that tests the slope of the yield curve as an indicator for future GDP growth. The second section describes an investigation into the relationship between interest rates and the exchange rate.

The relationship between the yield curve and economic growth

In section 2, we discussed how expectations of future economic growth should influence the shape of the yield curve. If those expectations prove correct on average, then the shape of the yield curve now should indicate how economic growth is likely to evolve over time.

We can capture the essence of that idea by plotting annual GDP growth one year ahead along with a measure of the yield curve slope now (in this case, the ten-year interest rate less the three-month interest rate from our government yield curve dataset). It is apparent that future economic growth tends to be higher, on average, when the yield curve is steeper, although the relationship is certainly not perfect (see figure 12).[20]

[20] The correlation between these two series is 0.47. The results from estimating a regression equation gave a constant of 2.95 and a regression coefficient of 0.86. The standard deviations are 0.25 and 0.20 respectively, so both estimated parameters are statistically significant to the 5 percent level. The results are typical: relationships between the yield curve slope and future GDP growth have been investigated and established as statistically significant for many international economies.

Figure 12

The slope of the government yield curve and GDP growth relative to average

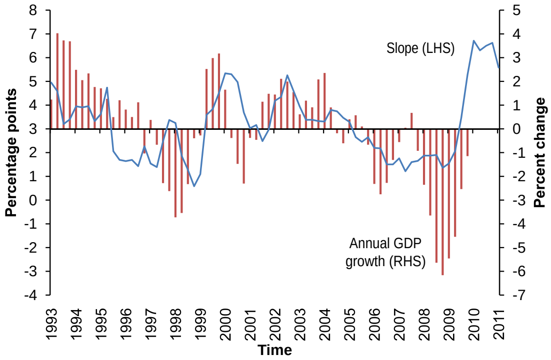

The slope of the yield curve could be expressed as the difference between any long-maturity and short-maturity yield. [21] The results of regressions using an extensive range of slope measures are detailed in Krippner and Thorsrud (2009), along with tests that use the data as it would have been available at each point in time. The combination of those regression results is now used within the Reserve Bank's suite of statistical models as one of the indicators of GDP growth.

The relationship between interest rates and the exchange rate

The uncovered interest parity hypothesis (UIPH) is the usual theoretical starting point for considering the relationship between interest rates and the exchange rate, and is based on the well-used adage "there are no free lunches". So the UIPH essentially maintains that the gain a foreign investor expects to make from investing at the higher interest rates in New Zealand should be offset by an equal expected depreciation in the New Zealand dollar against the investor's home currency.

Unfortunately, to the chagrin of economists the world over, the UIPH is often rejected at short horizons; that is, using short-term changes in the exchange rate and short-maturity interest rates. Indeed, a common finding is counter to the UIPH prediction (but consistent with conventional wisdom): higher interest rates tend to precede a currency appreciation. However, there is empirical support for the UIPH at longer horizons, and theoretical models show how cyclical factors in the economy and financial markets might lead tests of the UIPH to fail for short horizons even if the UIPH is explicitly incorporated within the model.[22]

Yield curve analysis offers a perspective for trying to link the different pieces of evidence outlined above. Firstly, the NS model components allow us to decompose interest rate differentials into a Level component (capturing long-maturity interest rates) and a non-Level component (that is, the sum of the Slope and Bow components). Both of those components may be tested separately within the UIPH regression equation. Secondly, Krippner (2008) establishes that the Level component of the NS model corresponds with the fundamentals of the economy, while the non-Level components together correspond to the cyclical components of the economy.

The results from Krippner (2006b) using Canadian and US data indicate that the non-Level/cyclical component is behind the rejection of the UIPH, while the UIPH would hold if just the Level/fundamental component were used. The results from Graham and Krippner (2009) using the three-month data for New Zealand and Australia show similar results.

Specifically, figure 13 shows that the exchange rate depreciates, on average when the Level/fundamental component of the three-month interest rate differential is positive (that is, the left side of the brown line). Similarly, the exchange rate appreciates on average when the Level/fundamental component of the three-month interest rate differential is negative (that is, the right side of the brown line).

Conversely, the non-Level/cyclical components of the three-month interest rate differential show the opposite relationship. As the interest rate differential moves from less positive to more positive (that is, following the green line from left to right), the exchange rate changes move from large depreciations to large appreciations. The non-zero intercept of the line also suggests a positive term premium,

21 **Alternatively, the measure of the yield curve slope could be the NS Slope coefficient, which effectively captures the average slope across the entire yield curve.**

22 **Krippner (2006) and Graham and Krippner (2009) contains further discussion and references.**

on average, of around 0.20 percentage points between New Zealand and Australia.

Figure 13
Data and results for testing the uncovered interest parity hypothesis between New Zealand and Australia

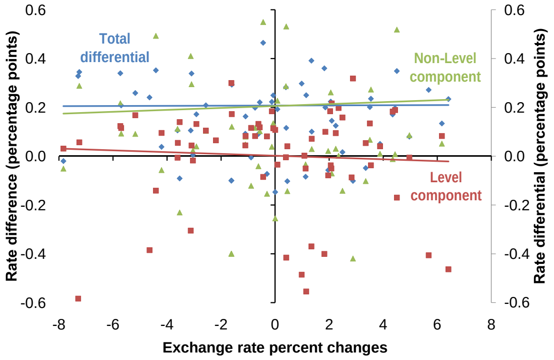

5 Conclusion

The yield curve is a concept that helps provide some perspective to the multitude of different interest rates that exist in modern economies. Specifically, we have seen how the interest rates of government securities with similar characteristics apart from their maturities may be grouped into a single government yield curve. Other groups of similar securities, such as bank securities or mortgage rates, may also be grouped into their own yield curves. These may be viewed as a hierarchy relative to the government yield curve.

Yield curves are predominantly influenced by macroeconomic factors, which implies that market expectations of those factors may be inferred via yield curve analysis. For example, we have seen how the slope of the yield curve may be used as an indicator of future GDP growth. Going forward, we would like to extend that simple relationship to simultaneously model inflation and economic growth in conjunction with the NS Level, Slope and Bow components of the yield curve, as undertaken by Diebold, Rudebusch and Aruoba (2006) for the US. It would also be interesting to add the exchange rate data to that mix, particularly given that our yield curve perspective offers an interesting avenue for research into the relationship between interest rates and the exchange rate. One aspect that might be addressed in such analysis is how (or if) exchange rate dynamics influence the persistent interest rate differentials that New Zealand usually experiences relative to Australia and to major global economies such as the US.

To conclude, interest rates in their own right have been interesting enough to have a recorded history of at least five millennia. And interest rates are even more interesting when viewed within the context of the yield curve. So the yield curve perspective on interest rates is likely to keep the interests of followers of economics and finance for a long time yet.

References

Bank for International Settlements (2005) "Zero-coupon yield curves: technical documentation", BIS Papers 25.

Diebold, F, G Rudebusch, and S Aruoba (2006) "The macroeconomy and the yield curve: a dynamic latent factor approach", *Journal of Econometrics*, 131, pp. 309-338.

Graham, J and L Krippner (2010) "Uncovered interest parity between Australia and New Zealand", forthcoming.

Hawkesby, C (1999) "A primer on derivatives markets", Reserve Bank of New Zealand *Bulletin*, Vol. 62. no. 2. June 1999, pp. 24-43.

Homer S and R Sylla (2005) *A history of interest rates* (fourth edition), John Wiley and Sons

James, J and N. Webber (2000) *Interest Rate Modelling*, John Wiley and Sons.

Krippner, L (2006a) "A theoretically-consistent version of the Nelson and Siegel class of yield curve models", *Applied Mathematical Finance*, 13(1), pp. 39-59.

Krippner, L (2006b) "A yield curve perspective on the uncovered interest parity hypothesis", Working Paper, University of Waikato.

Krippner, L (2008) "A macroeconomic foundation for the Nelson and Siegel class of yield curve models", Working Paper, University of Technology Sydney.

Krippner, L (2009) "A theoretical foundation for the Nelson and Siegel class of yield curve models", Reserve Bank of New Zealand Discussion Paper, DP2009/10.

Krippner, L and L Thorsrud (2009) "Forecasting New Zealand's economic growth using yield curve information", Reserve Bank of New Zealand Discussion Paper, DP2009/18

Nelson, C and Siegel, A (1987) "Parsimonious modelling of yield curves", *Journal of Business*, October, pp. 473-489.

Reinhart, C and K Rogoff (2009), T*his time is different: Eight centuries of financial folly*, Princeton University Press.