Extracting expectations of New Zealand's Official Cash Rate from the bank-risk yield curve

Leo Krippner

March 2002

Discussion Paper Series

DP2002/01

Extracting expectations of New Zealand's Official Cash Rate from the bank-risk yield curve[1]

Abstract

The hypothesis that a forward term-premium (FTP) exists between forward 1-day rates calculated from the New Zealand bank-risk yield curve and the corresponding ex-post Official Cash Rate (OCR) is tested by applying a single equation method for a cointegrated system to daily data from March 1999 to December 2001. The results indicate that the FTP is statistically significant for all forward horizons tested. The results also indicate that the estimates of the FTP appear to be an increasing function of the forward horizon, and the FTP may be tentatively represented as a simple monotonically-increasing analytical function. The model may be used in reverse to imply current ex-ante expectations of the OCR.

"Time is a train,
    makes the future the past,
        leaves you standing in the station,
            your face pressed up against the glass."

- U2 (1991), "Zoo Station",
      Achtung Baby, Island Records Ltd.

---

1 Introduction and background

The Official Cash Rate (OCR) has been used by the Reserve Bank of New Zealand (the Bank) since 17 March 1999 for adjusting the stance of monetary policy. Market expectations of the OCR are a useful element of information to the Bank when forming its OCR decision, and in the Bank's ongoing monitoring of markets. Further discussion on these matters is contained in Krippner and Gordon (2001). Many private sector researchers are also interested in gauging market expectations of official rates from market data, as a basis for trading relative to their own OCR expectations.

This article outlines a method for extracting OCR expectations from market-quoted bank-risk interest rates. The main contribution is to outline a relatively straightforward method for estimating the term premium component, a key unknown within the method, and to provide empirical estimates of this component for the New Zealand market.

As background, one approach to calculating interest rate expectations is to assume the pure expectations hypothesis of the yield curve (hereafter PEH); ie mechanically calculate forward rates from the current yield curve, and assume these correspond to unbiased estimates of expected future rates. The PEH has received much attention in the literature. For example, a selection of recent international work includes a survey of related US literature by Campbell (1995), an analysis on selected European countries by Gerlach and Smets (1997), a study for the United Kingdom by Cutherbertson (1996), and for Canada by Deaves (1996). The empirical results from that work generally rejects the PEH, although the expectations hypothesis allowing for a term premium (hereafter EHP) is sometimes accepted. For New Zealand, Krippner (1998) does not reject the PEH for a horizon of up to six months, based on quarterly bank-bill and bank-bill futures data. However, using weekly data and the horizon of the full yield curve, Guthrie, Wright and Yu (1999) rejects the PEH, but accepts the EHP.

The more specific topic of calculating official monetary policy rate expectations from the yield curve has been approached by several authors, typically associated with central banks. Essentially, this

generally involves an EHP model; ie calculating or observing forward rates from market data, and then subtracting an estimate of the forward term premium (FTP).

For example, work undertaken at the Bank of Canada suggests that an FTP exists between official rate expectations and Canadian forward rate agreements (FRAs) on bankers accepted paper. Empirical results noted in Paquette and Streliski (1998) indicate that the FTP increases with the maturity of the FRA (ie the horizon of expectations), and Gravelle (1998) discusses an estimate of an FTP that varies over time. In work undertaken at the Bank of England, Brooke and Cooper (2000) notes that United Kingdom interbank forward rates have an upward bias compared with actual policy rate expectations, and that the bias increases with maturity. A popular method used by many central banks to extract interest rate expectations, especially in Europe, is that based on the work of Nelson and Siegal (1987) and Svensson (1994). This essentially involves fitting a parametric function to represent the entire zero-coupon yield curve, and then using the implied forward rates from the related parametric forward rate function as a gauge for market expectations of interest rates. A review of these methods and the extent of their application is contained in Bank for International Settlements (1999), although the discussion does not mention whether an FTP is, or should be, allowed for in practice.

Published private sector work on official rate expectations is less common, perhaps because of its proprietary nature. One article is that by Porter (1999), which discusses the use of money-market rates to gauge official rate expectations for the United States, the European currency, and the United Kingdom. King (1999) discusses the time-varying nature of the difference between survey-based expectations of official rates and expectations implied by market prices using the PEH. The analysis is undertaken for the United States, Japan, Germany, and the United Kingdom.

The approach to estimating OCR expectations outlined in this article is similar in principle to the work noted above; that is forward rates are mechanically calculated from the yield curve, and an estimate of the FTP is subtracted from those forward rates to leave an estimate

of the official rate. This procedure and the appropriate background is discussed in section 2.[2]

Given that the FTP is a key requirement for an OCR expectations model, the major focus of the remainder of the article is concerned with the estimation of a plausible, and practically useful FTP function. Specifically, section 3 discusses the models proposed for estimation, section 4 discusses the data used in the estimation, and section 5 discusses the results. Section 6 concludes by discussing the main results, and noting potential areas for further investigation.

2   A model for OCR expectations

Forward interest rates may be calculated from current physical interest rates using the following relationship (contained, for example, in Hull (2000) and Svensson (1994)):

$$f(n, m-n) = \frac{mr(m) - nr(n)}{m - n} \tag{1}$$

where:

- $f(n,m-n)$ is the $(m-n)$-day rate, n days forward;

- $r(m)$ is the current physical interest rate for maturity m (note that all interest rates in this article are expressed on a continuously compounding basis, unless specifically stated otherwise);

- $r(n)$ is the interest rate for maturity n ($n<m$).

Substituting $m=n+1$ in equation 1 gives the following relationship:

$$f(n,1) = (n+1)r(n+1) - nr(n) \tag{2}$$

---

[2]   Note that there are alternative, perhaps more efficient, ways to extract OCR expectations when an estimated FTP function is already available. However, the outline in this paper is retained because it leads directly to the method of estimating the FTP function, as subsequently discussed.

where $f(n,1)$ is the 1-day rate, n days forward. Hence, a series of current interest rates with day-by-day maturities, $r(n)$ (ie the current yield curve), could be equivalently expressed as a series of forward 1-day rates, $f(n,1)$.

One further piece of notation is to add a time index to denote when the series of forward 1-day rates were calculated (ie when the yield curve was observed). Hence $f(t,n,1)$ is the 1-day rate, n days forward, measured at time t. In this notation, $f(t,0,1) = r(t,1)$, which is the 1-day interest rate at time t.

According to the EHP, $f(t,n,1)$ with an allowance for a constant should provide an unbiased expectation of $r(t,1)$ in n days time, or $r(t+n,1)$. This may be most generally represented as:

$$r(t+n,1) = \alpha(t,n) + \beta(t,n) \cdot f(t,n,1) + u(t+n,n) \qquad (3)$$

where:

- $\alpha(t,n)$ is the FTP parameter;

- $\beta(t,n)$ is the parameter relating the 1-day forward rate and the expected 1-day interest rate; and

- $u(t+n, n)$ is the model disturbance.

The notation (t,n) indicates that each parameter is associated with the given horizon n, and each could potentially vary over time in the most general representation.

If t is today, equation 3 provides the basis for calculating ex-ante expectations of $r(t+n,1)$ from today's yield curve, once certain assumptions about the parameters are made. For example, the PEH would suggest that $\alpha(t,n)=0$ and $\beta(t,n)=1$, and then $f(t,n,1)$ alone would provide an unbiased forecast of $r(t+n,1)$, (since $E[u(t+n,n)] = 0$). The EHP would suggest $\beta(t,n)=1$ and $\alpha(t,n) \neq 0$, so $f(t,n,1)$ would require an adjustment by $\alpha(t,n)$ to provide an unbiased forecast of $r(t+n,1)$. The analysis in this article assumes that $\alpha(t,n) = \alpha(n)$, a constant that varies with horizon

n, but not over time. Of course, a more flexible, and perhaps realistic, model would allow $\alpha(t,n)$ to potentially vary over time, and this is discussed later as an avenue for further investigation.

The practical issues behind using equation 3 as a basis for an OCR expectations model are then: defining the current yield curve using market-quoted data; calculating forward 1-day rates, $f(t,n,1)$, from that yield curve; and estimating $\alpha(n)$ for the required horizons that OCR expectations are required for. This enables a "genuine" expectation of future 1-day rates, $r(t+n,1)$, to be calculated, and then a link from $r(t+n,1)$ to the OCR is required to formally complete the OCR expectations model.

In this article, bank-risk interest rates are used to define the yield curve, so $r(t,1)$ is naturally the overnight interbank rate, and $\alpha(n)$ is the FTP between $f(t,n,1)$ and "genuine" expectations of the overnight interbank rate. In practice, as noted in Brookes and Hampton (2000), the overnight interbank rate has almost always been identical to the OCR since the OCR system was introduced. Hence, it is reasonable to assume that OCR(t) is identical to $r(t,1)$, or more importantly, that estimates of $r(t+n,1)$ are directly related to the expected OCR. This is a convenient but not critical assumption; any systematic difference between OCR(t) and $r(t,1)$ or their expectations could be captured in the $\alpha(n)$ if the OCR did not equal the overnight interbank rate on average.

Another element of practice is that the OCR has almost always remained constant during the period between pre-specified OCR announcement dates, as per the Bank's stated intention.[3] Hence, if $r(t+n,1)$ shows any variation between OCR announcement dates (as it often will), then it is more realistic to treat the average of $r(t+n,1)$ between OCR announcement dates as an estimate of the expected OCR for that period.

---

[3] The one exception to-date was the 50 basis point cut to the OCR on 19 September 2001, in the aftermath of the 11 September terrorist attacks on the United States.

Figure 1:
An illustration of the OCR expectations model
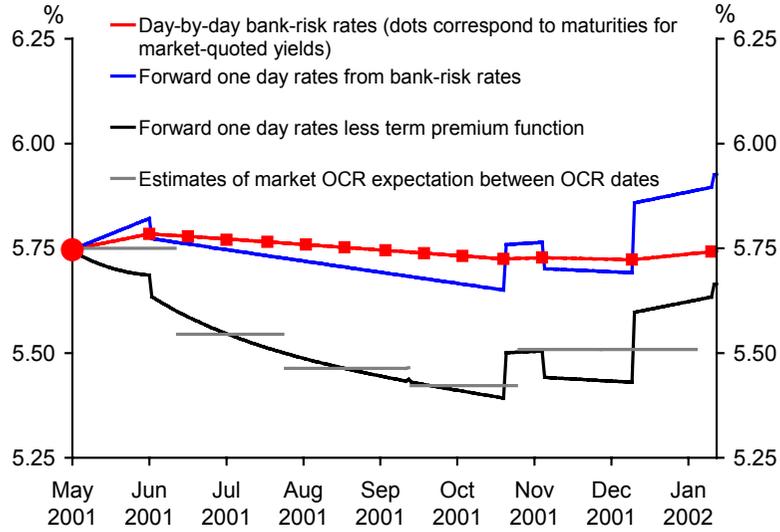
Anticipating the later discussion of calculating forward 1-day rates from market data (in section 4), and the estimates of the FTP function (in sections 3 and 5), figure 1 illustrates the concepts discussed above for the bank-risk yield curve observed on 24 May 2001:

- The first line shows bank-risk zero-coupon rates, with the dots indicating the actual observed rates (ie market-quoted rates transformed to a zero-coupon continuously compounding basis), and the connecting line showing linearly interpolated rates.

- The second line shows the forward 1-day rates calculated from the day-by-day term-rates.

- The third line shows the forward 1-day rates less the FTP function.

- The fourth (stepwise) line shows the average of the forward 1-day rates less the FTP for each OCR period. Hence, this line is

an estimate, as at 24 May 2001, of the market expectation for the OCR at each of the future OCR announcement dates.

The example above highlights that the estimate of the FTP function is the key consideration in the estimation of OCR expectations from the yield curve. An FTP generally exists to allow for factors such as the credit-risk of the issuer relative to a rolling overnight investment, and the liquidity preferences of the investor. However, the analysis discussed in this article only involves a statistical measurement of the FTP using historical data, so an underlying theory for the FTP is not necessarily required.

3   Estimating the FTP function

With the assumption that the parameters $\alpha(n)$ and $\beta(n)$ do not vary over time and $OCR(t)$ equals $r(t,1)$, equation 3 may be equivalently written using $OCR(t)$ and $f(t,n,1)$ measured n days ago, or $f(t-n,n,1)$:

$$OCR(t) = \alpha(n) + \beta(n) \cdot f(t-n,n,1) + u(t,n) \qquad (4)$$

This offers an approach to estimating $\alpha(n)$ and $\beta(n)$ using historical data. But although the estimation of equation 4 appears straightforward, there are actually several statistical issues to address.

Firstly, evidence indicates that the OCR and forward 1-day rates are cointegrated (ie both move as a "common random-walk" - the empirical results are presented and discussed in section 5). Hence, the statistical process may be represented as the following cointegrated system, which is adapted from Hamilton (1994):

$$OCR(t) = \alpha(n) + \beta(n) \cdot f(t-n,n,1) + u_1(t,n) \qquad (5a)$$

$$f(t,n,1) = f(t-1,n,1) + u_2(t) \qquad (5b)$$

$$u_1(t,n) = \Psi_1(L)\varepsilon(t), \; u_2(t) = \Psi_2(L)\varepsilon(t) \qquad (5c)$$

$$\varepsilon(t) \in N(0,\sigma^2) \qquad (5d)$$

The system of equations 5 may be termed a physical-forward system. The intuition underlying equations 5 is: $\varepsilon(t)$ represents unpredictable new information; that new information influences market expectations of the OCR, which therefore changes forward 1-day rates; and that same new information also causes forecast errors (the difference between ex-ante expectations of the OCR and the ex-post actual realisation of the OCR later in time).

Stock and Watson (1993) outline a method, asymptotically equivalent to full information maximum likelihood, for estimating the parameters in equation 5a from a single equation. This analysis follows the Stock and Watson (1993) approach as outlined in Hamilton (1994), pages 608 to 612:

$$\text{OCR}(t) = \alpha(n) + \beta(n) \cdot f(t-n,n,1) + \sum_{i=-p}^{+p} \gamma_i \Delta f(t-n-i,n,1) + v(t,n)$$
(6a)

$$v(t,n) = \Psi_3(L)v(t,n)$$
(6b)

The appropriate order of p in equation 6a is chosen so that $v(t,n)$ is uncorrelated with all leads and lags of $u_2(t)$, and then the parameters of equation 6 may then be estimated without bias using OLS. However, because the error $v(t,n)$ will not in general be i.i.d normal, then an adjustment to the standard error of the parameter estimates may be required before any statistical inference or hypothesis testing of the estimated parameters is undertaken. $\Psi_3(L)$ may be modelled as an $AR(Q)$ process, and used to adjust the t-statistics from the regression in equation 6a:

$$v(t,n) = \sum_{q=1}^{Q} \kappa_q \cdot v(t-q,n) + e(t,n)$$
(7a)

$$\left[\frac{\alpha(n)}{\sigma_\alpha}\right] = \left[\frac{\alpha(n)}{\sigma_\alpha^+}\right] \frac{\sigma_v}{\sigma_e}\left(1 - \sum_{q=1}^{Q} \kappa_q\right)$$
(7b)

where:

- $\alpha(n)/\sigma_\alpha$ is the adjusted t-statistic;

- $\alpha(n)/\sigma_\alpha^+$ is the original t-statistic from equation 6a;

- σ_v and σ_e are the standard deviations of $v(t,n)$ (from 6a) and $e(t,n)$, respectively; and

- κ_q are the estimated parameters of the $AR(Q)$ process.

One approach to choosing p is to calculate the empirical correlations of $\Delta f(t-n-i,n,1)$ (which is $u_2(t-n-i)$) with the residuals $u_1(t,n)$ obtained from an initial single-equation estimation of equation 5a, and then choose p to capture the significant correlations. However, the equations 5 underlying equations 6 suggest a more direct approach. Specifically, one can write a finite-order representation for $u_1(t,n)$ and $u_2(t)$ and then calculate the expected correlations by inspection:

$$u_1(t,n) = \sum_{j=0}^{n} \rho_{1j} \cdot \varepsilon(t-j)$$
(8a)

$$u_2(t) = \sum_{k=0}^{K} \rho_{2k} \cdot \varepsilon(t-k)$$
(8b)

$$\text{cor}(u_1(t,n),u_2(t-n-i)) = \frac{\text{cov}[u_1(t,n),u_2(t-n-i)]}{\sqrt{\text{var}[u_1(t,n)]} \cdot \sqrt{\text{var}[u_2(t)]}}$$

$$= \frac{\text{cov}[\sum_{j=0}^{J} \rho_{1j} \cdot \varepsilon(t-j), \sum_{k=0}^{K} \rho_{2k} \cdot \varepsilon(t-n-i-k)]}{\sqrt{\text{var}[u_1(t,n)]} \cdot \sqrt{\text{var}[u_2(t)]}}$$
(8c)

The indices for the ε terms in the first summation range from $t-n$ to t, and these will "overlap" the indices for the ε terms in the second summation (and hence yield finite covariance) whenever $t-n \leq t-n-i-k \leq t$, or $-n+k \leq i \leq k$. Hence, the minimum i for non-zero correlation is $-n$ (when $k=0$), and the maximum i for

non-zero correlation is $+K$ (when $k = K$). All other correlations will be zero. If it is assumed that the market is perfectly efficient, then it should be the case that $u_2(t,n) = \varepsilon(t)$, ie the new information should be incorporated into market expectations immediately, K would equal zero, and the coefficients on the lagged $\Delta f(t-n-i,n,1)$ terms in equation 6a (ie those terms associated with i ranging from $-p$ to 0) should be insignificantly different from zero. This is tested empirically in section 5.

Finally, following the advice of Hamilton (1994) regarding imposing the cointegrating vector when it is suggested by theoretical considerations (page 582), and also the approach of Gravelle (1998), the restriction $\beta(n) = 1$ is imposed in equation 6a to yield the actual equation to be estimated:[4]

$$\left[\mathrm{OCR}(t) - f(t-n,n,1)\right] = \alpha(n) + \sum_{i=-n}^{+n}\gamma_i\Delta f(t-n-i,n,1) + v(t,n) \qquad (9)$$

Equation 9 is estimated for each horizon n. The intuition behind the estimation is: "*after allowing for the unexpected events on the* n *days between the expectation being formed and the actual OCR being realised, what is the average systematic difference between the 1-day forward rate* n *days forward, and the OCR in* n *days time?*" Without this allowance for unexpected events (as proxied by $\Delta f(t-n-i,n,1)$), the estimate of $\alpha(n)$ would be dominated by the inherent random-walk in the physical/forward system, and may not be very representative of the systematic difference as sought. Even in cases where the original data is stationary (ie the OCR and forward rates are not a random-walk), a large autoregressive component in the residuals could lead to a mis-leading estimate of $\alpha(n)$ and/or its standard error in small samples.

For the horizon n, there are $T-n$ historical data points available, and a further n are "lost" due to the leads and lags required by the

approach. The number of independent variables used in the horizon n estimation is $2n+2$. Hence, the degrees of freedom for the estimation is $T-4n-2$. With just under three years (1021 days) of data at present, the degrees of freedom would be fully exhausted with $n = 255$, and $n = 244$ is the maximum horizon investigated in this analysis.[5]

Once the estimates of $\alpha(n)$ are available for each horizon, a smooth functional form by horizon may be estimated from the $\alpha(n)$ estimates:

$$\alpha(n) = \theta \cdot \mathrm{UFTP}(n) + \delta(n) \qquad (10)$$

where:

- θ is the magnitude of the FTP;

- $\mathrm{UFTP}(n) = 1 + \dfrac{365}{\phi}\left[\exp\left(-\dfrac{\phi(n+1)}{365}\right) - \exp\left(-\dfrac{\phi n}{365}\right)\right]$, ($\phi > 0$), which is the "unit shape" of the function; and

- $\mathrm{FTP}(n) = \theta \cdot \mathrm{UFTP}(n)$.

$\mathrm{UFTP}(n)$ is an adapted version of a related function for instantaneous forward rates discussed in Nelson and Siegal (1987) and Svensson (1994), ie $f(n) = \theta \cdot [1 - \exp(-\phi n)]$. Essentially, ϕ is a parameter that determines the rate at which the exponential term shrinks to zero with maturity, and hence how the function rises asymptotically to the constant θ. Apart from providing a "plausible" functional form for the FTP (ie in accordance with a prior that the FTP should be a smooth, monotonically-increasing function of maturity any point in time), using this basis may also enable a direct comparison to any subsequent analysis of the yield curve based on the Nelson and Siegal (1987) and Svensson (1994) models.

---

[4] With this restriction, it is possible to set up a simple moving-average representation for the estimation of $\alpha(n)$. This is more appealing in theory, but is not as straightforward to estimate in practice. Further background and discussion on this potential approach is contained in Appendix 1.

[5] Imposing exclusion restrictions on some of the γ coefficients, and/or using a low-order polynomial representation for the series of γ coefficients (if appropriate) would increase the degrees of freedom. The polynomial approach has not been investigated by the author.

Given that an estimate of standard error is associated with each estimate of $\alpha(n)$, then it is most efficient to estimate equation 10 allowing for heteroskedasticity ie using OLS with weights of $1/\text{stderr}[\alpha(n)]$. This is also intuitive; there is more data available for shorter horizon estimates of $\alpha(n)$, hence those estimates should be determined with more precision, and that precision should be reflected by placing more weight on those estimates when estimating the FTP function.

Note that the above specification suggests that the FTP function is constant over time, which is implicitly assumed to be the case in this analysis. However, if the FTP function was allowed to be fully time-varying, it could potentially change in both shape and magnitude over time (ie ϕ and θ may both vary over time)

4   The data

A series of forward 1-day rates are not typically quoted in New Zealand (or other countries for that matter). Hence, the first task is to generate forward 1-day rate data from the associated zero-coupon yield curve data constructed with market data. The data (all bank-risk interest rates) used to construct the zero-coupon yield curve are:

- the current OCR, quoted on a discount basis, as a proxy for the overnight interbank rate;

- bank-bill rates for half-monthly maturities (ie 1st to 15th, or 16th to end-of-month) from 1 to 6 months, quoted on a discount basis;

- forward rate agreements (FRAs) on 3-month bank-bills with monthly settlements (from the 4x7 FRA to the 8x11 FRA), quoted on a discount basis; and

- swaps rates, with annual maturities from 1 to 5 years, and also 7 and 10 year maturities, all quoted on a semi-annual basis.

The reasons for selecting these data is that bank-bills, FRAs, and swaps are "liquid" instruments in the New Zealand market (ie

frequently traded), quoted with "high density" on the yield curve (ie with relatively small periods of time between adjacent maturities or settlements), and they all have a clear dependence on the current and expected OCR (through the overnight interest rate). All of the data are sourced from the Bank, being originally collected from a New Zealand interbank broker. Note that the 1-year swaps rate is chosen to define the 1-year point rather than the 9x12 FRA, since the former provides a better link to longer-maturity swaps rates which will be used to define the zero-coupon curve for longer horizons as more data become available.

As an aside, there are other money-market instruments in the New Zealand market that could potentially be used to define the bank-risk yield curve, such as bank-bill futures and foreign exchange forwards, but their use is not explored in this article.[6] Also, using the government yield curve to extract OCR expectations is not practical in the New Zealand context, mainly because Treasury bills are relatively illiquid, and hence their quoted yields may not be very representative of market OCR expectations.

The zero-coupon rates for the maturities of the instruments noted above are calculated in the usual manner:

- the OCR and bank-bill rates are transformed directly into their equivalent continuously compounding form;

- the FRA rates are transformed to into their equivalent continuously compounding form, and then combined with the rate corresponding to the settlement of the FRA to create an equivalent zero-coupon rate to the maturity of the FRA; and

- the swaps instruments are treated as semi-annual par bonds and the implied zero-coupon rates corresponding to the maturity of each swap are "bootstrapped" from the market-quoted swaps rates.

---

[6] Bank bill futures do have greater liquidity than FRAs in the New Zealand market, but lack the "high density" of FRAs (since the futures are only quoted for quarterly settlements, while FRAs are quoted for monthly settlements).

Chapters 4 and 5 in Hull (2000), for example, provide useful background to these steps, but obviously New Zealand market conventions for settlement, pricing, and maturity are used. In particular on the latter, the maturity associated with the bank-bill rates in New Zealand tends to depend on the slope of the yield curve at the time of issuance. Bank-bill rates tend to apply to the back of the half-month tranche in an upward sloping curve, and to the front of the half-month tranche in a downward sloping yield curve, since this is most advantageous to the interbank issuer, which is the party that specifies the exact maturity. Specifically allowing for this "switching" of bank-bill maturities to the back or front of the tranche made an immaterial difference to the estimates of $\alpha(n)$ and the FTP function compared to simply using the middle of the tranche as the maturity date for bank-bills in all cases. Hence, only the results obtained from assuming bank-bill maturities to the middle of the tranche are contained in the article. Further details on this, and New Zealand market conventions are available from the author.

The zero-coupon rates for each day-by-day maturity are calculated from the points on the zero-coupon curve using linear interpolation, and these are used to calculate all of the forward 1-day rates corresponding to the zero-coupon yield curve at that time (using equation 2). This process is repeated using yield curve data for each trading day since the introduction of the OCR. The method outlined above requires data for every calendar day, so the data for the previous trading day is used to create the data for non-trading calendar days.

The final dataset is then a sequence of time-series $f(t,n,1)$ by horizon n, where time t spans from 17 March 1999 to the latest available data (31 December 2001 in this case). The horizon n can potentially span from 1 day to approximately 3650 days (10 years), although a maximum of 244 days is sufficient in this case, given the estimation method and data currently available, as already discussed.

5   The results

The valid estimation of equation 9 requires both the OCR and $f(t,n,1)$ to be cointegrated for each horizon n, which is generally

accepted by the data. Figure 2 shows the results of unit root tests for each horizon, indicating that the null hypothesis of a unit root is not rejected for any $f(t,n,1)$ series, to 5 percent significance. The relevant ADF statistic for the OCR(t) itself is -0.97, which also does not reject the null hypothesis of a unit root to 5 percent significance. Figure 3 indicates that the unit root hypothesis is decisively rejected for the first difference of each series, and the value of -31.9 for the first difference of OCR (t) itself also indicates a strong rejection.

Figure 4 shows that the hypothesis of no cointegration between the OCR(t) and each $f(t,n,1)$ series is usually rejected, although there are many exceptions for longer horizons. Tests for cointegration using data to early 2001 did not show this systematic pattern, which suggests that the sharp and unanticipated cuts to the OCR in the second half of 2001 may have influenced the subsequent results. In any case, the rejection of cointegration between the OCR(t) and $f(t,n,1)$ does not make economic sense, since it would imply that the OCR could potentially diverge arbitrarily from forward rate in the long-term. Hence, the analysis proceeds on the assumption that OCR(t) is cointegrated with $f(t,n,1)$ for all horizons.

Figure 5 illustrates the tests of restricting the coefficients γ_i in equation 9 to be all zero for $i=-n$ to 0. The results of the F-test (indicated as a percentage, allowing for the changing degrees of freedom) shows that this restriction is usually valid, but there are occasionally some significant exceptions. Hence, for consistency between estimates of equation 9 for each horizon, the subsequent analysis uses only the unrestricted results for each horizon.

Figure 6 illustrates the results of estimating equation 9, indicating that $\alpha(n)$ is positive for each horizon.[7] This is apparent from the 95 percent confidence intervals (which allow for the changing degrees of freedom), indicating that the hypothesis that $\alpha(n)=0$ may be rejected at the 5 percent level of significance for all horizons. Also,

---

[7]   Strictly, the estimates of $\alpha(n)$ from equation 9 are actually all negative, due to the way that equation is defined. However, a term-premium is more intuitive as a positive number, so the absolute value of $\alpha(n)$ is used.

the rising lower bound of the 95 percent confidence band suggests that there is a rising dependence of $\alpha(n)$ on maturity.

As an example of the estimation process undertaken for each horizon, figures 7a and 7b respectively illustrate the actual data used for the n=90 day horizon estimation, and the series of residuals from that estimation. Note firstly that the forward rate generally overstates the OCR realised in the future, which is readily evident from the typically positive "raw residuals", $f(t-90,90,1) - OCR(t)$. The "raw residuals' also display large variance and autocorrelation. The residuals from the Stock and Watson (1993) method show smaller variance but still material autocorrelation. The residuals from the AR(Q) process show negligible autocorrelation and small variance, except for the "spikes" when the OCR was changed by discrete amounts at particular points in time. These "spikes" (indicated by the square dots) lead the AR(Q) residuals to be non-normal, and it is not immediately obvious how this "discontinuity" can be dealt with, at least within the approach taken in this article.

Figure 8 contains the results for estimating the FTP function using the estimates of $\alpha(n)$ for datasets with varying ranges of n. The first FTP function estimate uses all $\alpha(n)$ estimates from 0 to 244 days. However, figure 9 shows that this FTP function has the undesirable property that it lies outside the upper 95 confidence bound of $\alpha(n)$ estimates for horizons from 92 to 215 days. The second FTP function estimate excludes the $\alpha(n)$ estimates for horizons greater than 223 days, on the grounds that the variance of the $\alpha(n)$ estimates suspiciously decreases steadily for longer horizons, and this would tend to dominate the weighted OLS estimates. This exclusion results in a far better fit of the estimated FTP function to the $\alpha(n)$ estimates and the associated confidence intervals, as illustrated in figure 9. The third FTP function estimate excludes the estimates of $\alpha(n)$ above 223 days and below 31 days, since the latter $\alpha(n)$ estimates may be sensitive to the simple assumption of linear interpolation between the OCR and the 1-month bank-bill rate. Figure 9 shows that this makes only a marginal difference compared to the second estimate of the FTP function.
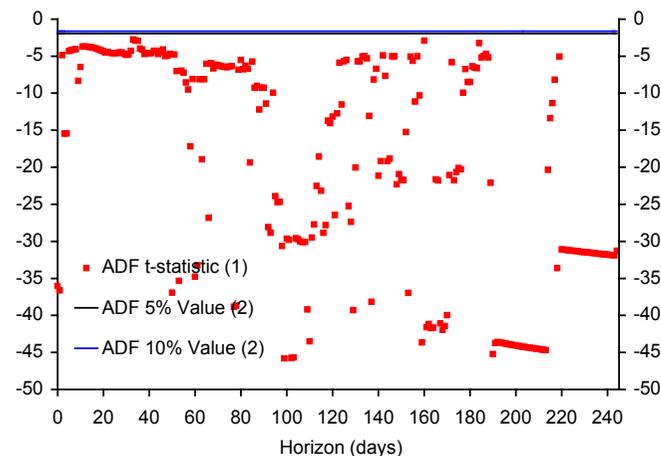
Figure 2:
ADF tests for $f(t,n,1)$, with critical values
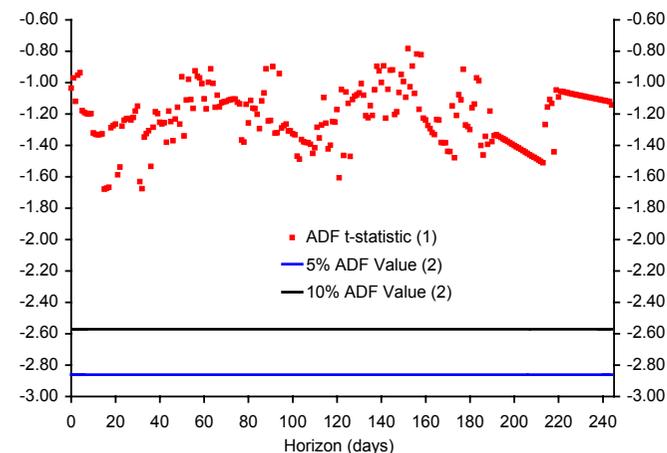
Notes:
(1) ADF regression includes constant. Lag selection for each horizon as in Hamilton (1994), page 530. (2) From Hamilton (1994).
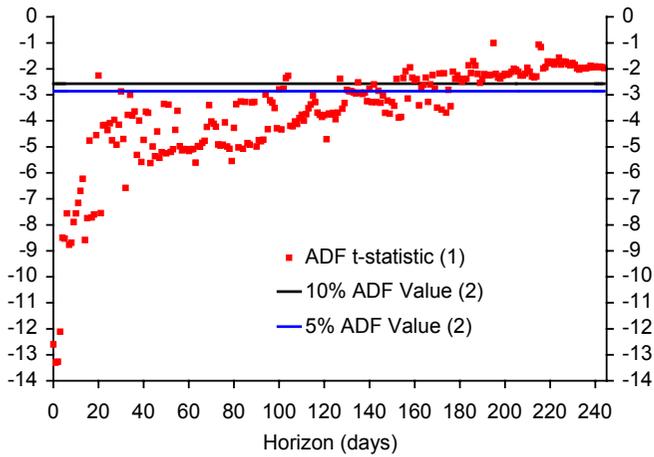
Figure 3:
ADF tests for $\Delta f(t,n,1)$, with critical values

Notes:
(1) No constant in ADF regression. Lag selection for each horizon as in Hamilton (1994), page 530. (2) From Hamilton (1994).

Figure 4:
ADF OCR(*t*)/*f*(*t*,*n*,1) cointegration tests, with critical values

Notes:
(1) No constant in ADF regression on residuals. Lag selection for each horizon as in Hamilton (1994), page 530. (2) From Hamilton (1994).

Figure 5:
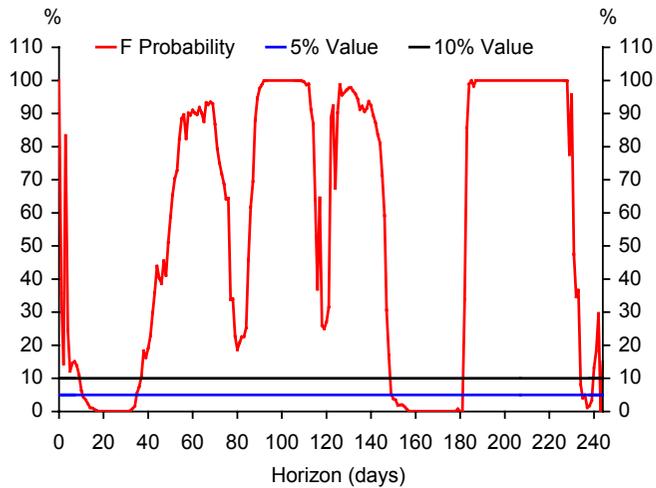**Test that the lagged coefficients *γ* in equation 9 are all zero**

Figure 6:
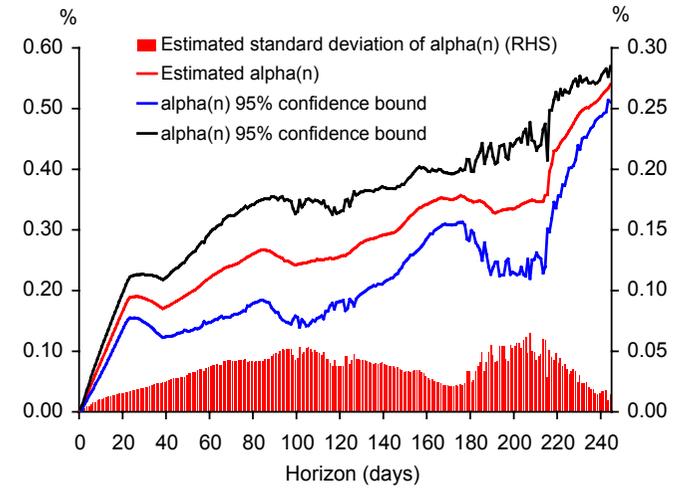**α(*n*) estimates with 95 percent confidence bounds**
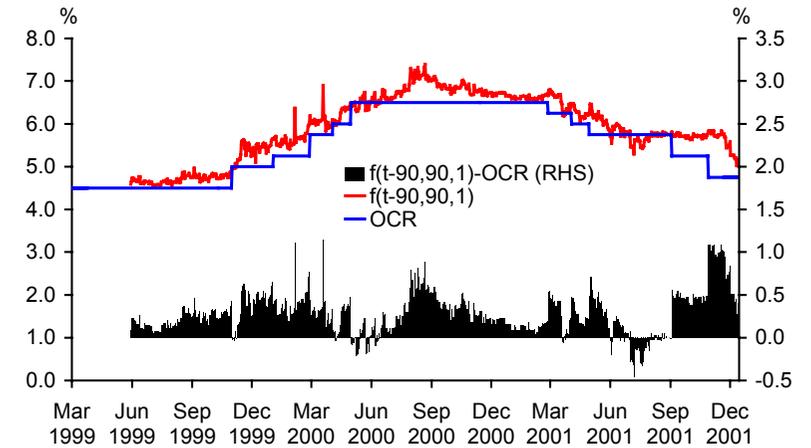

Figure 7a:
The data used for the *n*=90 day estimation
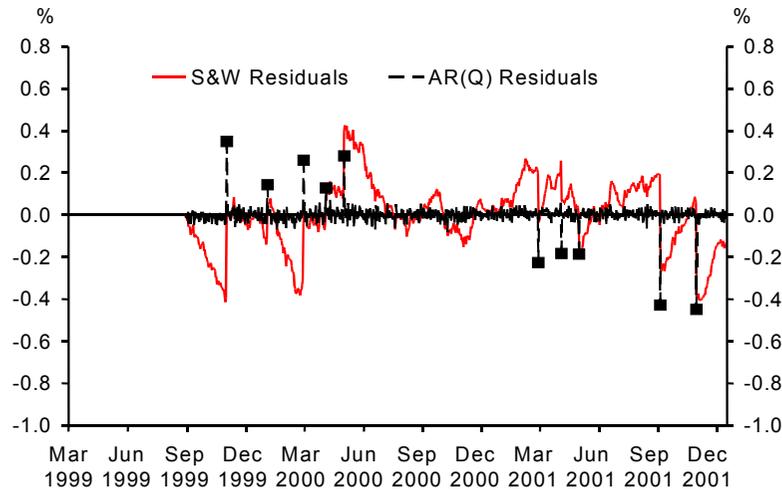
Figure 7b:
The residuals for the *n*=90 day estimation

Figure 8:
The FTP function parameter estimates

0-244	ϕ	θ
Estimate	4.82	0.510%
Std error	0.16	0.01%
t-statistic	30.6	62.9

0-223	ϕ	θ
Estimate	7.23	0.361%
Std error	0.25	0.01%
t-statistic	29.3	38.4

31-244	ϕ	θ
Estimate	5.88	0.367%
Std error	0.18	0.00%
t-statistic	32.8	100.7

Figure 9:
FTP function estimates and $\alpha(n)$ 95 percent confidence intervals
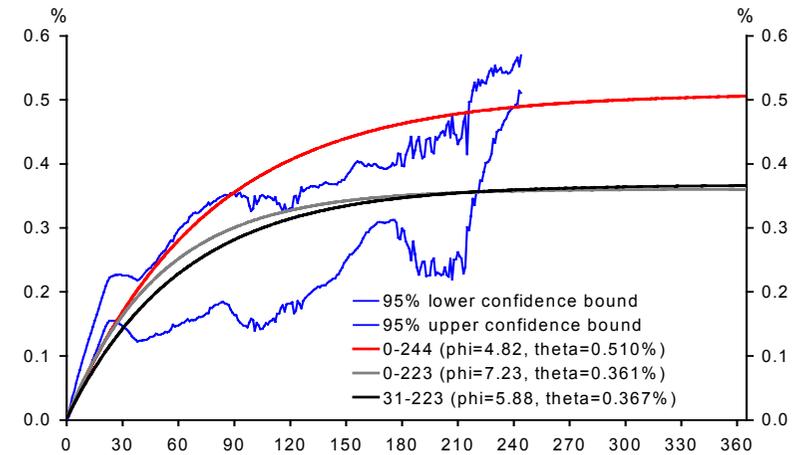
6   Conclusions and areas for further work

The strongest point to note in conclusion is that the evidence for a positive FTP in bank-risk interest rates relative to the OCR is very strong. This suggests that expectations of the OCR calculated on the basis of the PEH (ie assuming forward rates based on bank-risk instruments are an expectation of the OCR) are likely to be biased above actual expectations held by the market.

Secondly, the evidence that the OCR and forward one-day rates are unit root series and cointegrated is also very strong. This suggests that simply using the average of the difference between historical forward 1-day rates and the realised OCR may provide a misleading estimate of $\alpha(n)$. This estimate would tend to be dominated by the random-walk inherent in the physical/forward system, particularly in small samples (as with the New Zealand data currently available).

The third point is that the estimates of $\alpha(n)$ tend to show an increase with horizon, and are also quite variable. This suggests that modelling the FTP as the smooth, monotonically-increasing

parametric function noted in the text is a useful additional step, and each of the FTP functions estimated from the initial estimates of $\alpha(n)$ have very significant parameters. However, the FTP function estimates can vary widely depending on which $\alpha(n)$ estimates are included in this second estimation, and hence some pragmatism and/or experience-based judgement (see below) is useful at this point.

In answer to the natural question "Which FTP function should we use?", the author suggests the 31-223 day estimate rather than the 0-223 day estimate. Part of the reason for this is that the OCR to 1-month spread calculated using the 31-223 day estimate is closer to the "typically observed" spread in the market, whereas the alternative would be a little higher. Another reason relates to the general shape of the bank-risk yield curve. As background to this suggestion, the bank-risk yield curve shows a spread to the government yield curve that continues to increase (albeit slowly) out to a 10-year maturity, and one would naturally expect the government yield curve to be based on OCR expectations. In combination, these stylised facts suggest that the degree to which the bank-risk yield curve overstates OCR expectations should continue to increase with horizon. Hence, the curve that "flattens out" the slowest with horizon (ie the curve with the lowest ϕ) is likely to be the best approximation to the "true" FTP function. This also suggests that the estimate of ϕ may continue to decline as more data for the FTP function estimation becomes available, and this indeed has proved the case between subsequent updates of the analysis outlined in this article.

The final point in conclusion is that the assumption of constant $\alpha(n)$ coefficients and a constant FTP function over time may be too restrictive. The author intends to investigate the general topic of a time-varying FTP function in future work, and this may be facilitated by the moving-average form of analysis as outlined in Appendix 1.

In the meantime, the method outlined in this article is still likely to prove useful in extracting OCR expectations from the New Zealand bank-risk yield curve. Similar models may also be useful for

extracting official rate expectations from the yield curves of other countries, and this will also be investigated in future work.

Appendix: The moving-average estimation of $\alpha(n)$

The model represented by equation 9 may alternatively be expressed in a moving-average form. This form is more appealing in theory, since it is both fully parametric and parsimonious.

According to EHP, the evolution of the forward rate corresponding to a given date over time may be written as:

$$f(t-n+1, n-1, 1) = f(t-n, n, 1) + q(t-n+1) + \varepsilon(t-n+1)$$
(A.1)

where:
- $f(t-n, n, 1)$ is the 1-day rate, n days forward, measured at time $t-n$. This is the time $t-n$ predictor of $r(t, 1)$.
- $f(t-n+1, n-1, 1)$ is the 1-day rate, $n-1$ days forward, measured at time $t-n+1$. This is the time $t-n+1$ predictor of $r(t, 1)$.
- $q(t-n+1)$ is the marginal 1-day step, from n to $n-1$, in the forward term premium of horizon n; and
- $\varepsilon(t-n+1)$ is the innovation for time $t-n+1$, which is assumed for now to have constant variance σ_ε^2.

This process may be progressively iterated from day to day, until the given prediction date is reached, and the actual $r(t, 1)$ is realised. Iterating equation A.1 results in a collection of the marginal forward term premiums and innovation terms:

$$f(t-n+n, n-n, 1) = f(t-n, n, 1) + \sum_{i=1}^{n} q(t-n+i) + \sum_{i=1}^{n} \varepsilon(t-n+i)$$
(A.2)

and since $f(t-n+n, n-n, 1) = f(t, 0, 1) = r(t, 1)$ then:

$$r(t, 1) - f(t-n, n, 1) = \sum_{i=1}^{n} q(t-n+i) + \sum_{i=1}^{n} \varepsilon(t-n+i)$$
(A.3)

Equation A.3 is of a similar form to equation 9, which is easily seen by expressing the collections of innovation terms as single variables:

$$[r(t, 1) - f(t-n, n, 1)] = \alpha(n) + w(t, n)$$
(A.4)

where:
- $\alpha(n) = \sum_{i=1}^{n} q(t-n+i)$; and
- $w(t, n) = \sum_{i=1}^{n} \varepsilon(t-n+i)$.

However, the key difference is that equation A.4 contains a "known" and parsimonious moving-average (MA) structure in $w(t,n)$ (according to the model assumptions), compared to the corresponding non-parsimonious expression (the summation and residual terms) in equation 9. For example, according to the assumption of constant innovation variance, the overlap of innovations gives an obvious correlation structure for $w(t,n)$:

$$\text{cov}[w(t, n), w(t+s, n)] = \sigma_\varepsilon \begin{cases} n-s & \text{if } s < n \\ 0 & \text{if } s \geq n \end{cases}$$

where s is the number of days between different observations of $w(t,n)$ over time.

With the "known" correlation matrix for $w(t,n)$, estimates of both $\alpha(n)$ and the time-series of innovations $\varepsilon(t)$ may be obtained by direct application of Generalised Least Squares, ie factorising the correlation matrix, creating the transform of $[r(t, 1) - f(t-n, n, 1)]$ and the constant using that factor matrix, and then applying OLS to the transformed data (Hamilton (1994), for example, contains further details on this standard procedure). The set of residuals from the OLS regression is then the estimate of the time-series of innovations $\varepsilon(t)$, and these may then be used as the basis for investigating the potentially time-varying nature of $\alpha(n)$ and the FTP function.

One practical problem in this direct MA approach is that the correlation matrix for $w(t,n)$ can become very large using daily data (for the dataset used in this analysis it would be a 1021 x 1021 matrix), and so the factorisation process is time consuming (even when the process is made more efficient by exploiting the banded property of the symmetric correlation matrix).

Another problem is theoretical/empirical: preliminary investigation by the author using the direct MA approach indicates that the assumption of a constant innovation variance is not adequate. This inadequacy seems to arise from the fact that innovations can be much larger for OCR review days, compared to non-OCR review days (as one would intuitively expect, due to the stepwise nature of the OCR adjustments). Not specifically allowing for this difference in the assumptions for the innovation variance results in residuals that have a large amount of negative autocorrelation, making them unsuitable for the ongoing work into investigating an FTP that potentially varies over time. This suggests that any further investigation using the direct MA form would need to parametrically account for at least two different daily innovation variances; one for OCR review dates, and another for non-OCR review dates.

Note that the moving-average structure of the correlations in $w(t,n)$ makes equation A.4 a natural candidate for estimation using the Newey-West approach for calculating heteroskedastic and autocorrelation consistent standard errors for $\alpha(n)$. Indeed, this may be more straightforward than the method outlined in this article if only estimates of a constant FTP were required. However, the application of this method has not been investigated by the author, because it only offers a means to improve the estimate of the variance associated with $\alpha(n)$, rather than a means to estimate the innovations in form suitable for investigating an FTP that potentially varies over time.

References

Bank for International Settlements (1999), "Zero-coupon yield curves: technical documentation," Bank for International Settlements, Switzerland.

Brooke, M and N Cooper (2000), "Inferring market expectations from money market rates," *Bank of England Quarterly Bulletin*, November.

Brookes, A and T Hampton (2000), "The Official Cash Rate one year on," *Reserve Bank of New Zealand Bulletin*, 63, 2, pp 53-62.

Campbell, J Y (1995), "Some lessons from the yield curve," *Journal of Economic Perspectives*, 9, 3.

Cutherbertson, K, (1996), "The expectations hypothesis of the term structure: The UK interbank market," *The Economic Journal*, 106.

Deaves, R, (1996), "Forecasting Canadian short-term interest rates," *Canadian Journal of Economics*, 39.

Gerlach, S and F Smets (1997), "The term structure of euro-rates: Some evidence in support of the expectations hypothesis," *Journal of International Money and Finance*, 16.

Gravelle, T (1998), "Extracting market expectations at the Bank of Canada*," Centre for Central Banking Studies*.

Guthrie, G, J Wright and J Yu (1999), "Testing the expectations theory of the term structure for New Zealand," *New Zealand Economic Papers*, 33, 1.

Hamilton, J D (1994), *Time Series Analysis*, Princeton University Press.

Hull, J (2000), *Options, Futures and Other Derivatives* (4th edition), Prentice Hall International Inc.

King, M (1999) "Term premia: what's priced in?" *JP Morgan Securities Limited*, Fixed Income Research.

Krippner, L (1998), "Testing the predictive power of New Zealand bank bill future rates," *Reserve Bank of New Zealand Discussion Paper* G98/8.

Krippner, L and M Gordon (2001), "Market expectations of New Zealand's official rate," *Reserve Bank of New Zealand Bulletin*, June.

Nelson, C and A Siegal (1987), "Parsimonious modelling of yield curves," *Journal of Business*, October.

Paquette, J and D Streliski (1998), "The use of forward rate agreements in Canada," *Bank of Canada Review*.

Porter, W (1999), "Three month fixings and policy rates," *Credit Suisse First Boston*.

Stock, J and M Watson (1993), "A simple estimator of cointegrating vectors in higher order integrated systems," *Econometrica*, 61.

Svensson, L (1994), "Estimating and interpreting forward interest rates, Sweden 1992-4, *Centre for Economic Policy Research discussion paper* 1051.