Kalman filter estimates of time-varying term premia for New Zealand and Australia

Michael Gordon[1]

Financial Markets Department

Reserve Bank of New Zealand

Forward rates in the money market exhibit excess returns over spot rates, which can be attributed to an unobservable term premium and a forecast error. This paper uses a Kalman filter approach to characterise the evolution of the term premium over time. The approach used here accounts for the possibility of a unit root in the term premium, and adjusts for overlapping forecast errors. Using monthly data for New Zealand from 1995 and for Australia from 1993, the time-varying estimates are found to be preferable to those assuming a constant premium. There is strong evidence that the term premium was influenced by changes in the implementation of monetary policy in New Zealand.

INTRODUCTION

It is common practice for central banks to either manipulate or directly set a short-term interest rate as their primary tool for implementing monetary policy. Currently, both the Reserve Bank of New Zealand (RBNZ) and the Reserve Bank of Australia (RBA) set the rates at which they make settlement cash available overnight, and they review these rates at scheduled intervals. This creates a high degree of stability in the overnight cash market, and allows market participants to assess longer-term interest rates based on their forecasts of the policy rate.

There is a clear motive for deriving a measure of the market's expectations of the policy rate. Market participants are concerned with how the "average" market expectation compares to their own views, and will trade accordingly. For the central banks themselves, market expectations provide a reality check on their assumptions about the state of the economy, and help them to recognise when markets may respond to a change in monetary policy in an

[1] I would like to thank Kelly Eckhold, Wai Kin Choy, and Leo Krippner for comments and ongoing discussions on this topic. I am extremely grateful to James Morley and Toni Gravelle for providing me with the GAUSS code used in their paper. The views expressed in this paper are my own, and do not necessarily reflect the views of the Reserve Bank of New Zealand.

adverse way.

The theoretical basis for deriving market expectations is the Expectations Hypothesis of the term structure of interest rates, hereafter shortened to EH. This proposes that forward interest rates are unbiased predictors (up to constant term premia) of future spot interest rates. In other words, forward rates represent the 'average' market expectation of future spot rates, plus a constant forward term premium. Standard linear regression tests can be used to test the validity of this assumption, and to produce estimates of the term premia.

However, the vast body of literature on this subject has frequently failed to support the EH - that is, a number of studies have found that forward rates are not unbiased predictors of spot rates. (See Cook and Hahn (1990), for example, for a review of the post-war US literature.) The evidence for New Zealand, such as Guthrie, Wright and Yu (1999) and Krippner (2002) has supported the assumptions of the EH. In addition, an earlier study (Krippner, 1998) failed to find support for the existence of a positive term premium. The evidence for Australia from Lowe (1994), Karfakis and Moschos (1995), and Gordon (2002) has generally supported the EH as well.

The lack of empirical support for the EH in some regions could be due to the presence of a time-varying term premium (Mankiw and Miron, 1986). With this in mind, more recent studies have attempted to produce time-varying estimates of the term premium. Finance theory does not suggest an appropriate specification, so these studies have focused mainly on those factors that might plausibly alter the term premium, such as the volatility of financial and economic variables. For example, Engle, Lilien and Robins (1987) introduced the ARCH-M approach, where the term premium is influenced by the conditional variance of interest rates. Lee (1995) uses ARCH measures of the conditional variances of macroeconomic variables.

An alternative approach is to employ an adaptive filtering technique where the estimates of the parameters are allowed to vary stochastically over time. Previous studies have used the Kalman filter to extract the unobserved term premia from forward rates, under an identifying assumption of rational expectations. If the estimated term premia vary sufficiently over time, we can reject the EH. Iyer (1997) applies this approach to US short-term rates and rejects the EH. Bhar (1993) also rejects the EH for Australia, albeit using a very small sample.

The advantage of using the Kalman filter is that it is possible to estimate how the term premia have evolved over time without a prior specification of the factors that cause the premia to change. From a practical perspective, the term premium can be thought of as compensation for interest rate volatility, credit risk, market liquidity, political risk, and any other factors that could cause forward rates to systematically differ from the spot rate. These factors cannot be measured directly, are difficult to proxy with other observable variables, and are likely to vary in importance over time.

This paper uses Gravelle and Morley's (2002) Kalman filter approach to derive term premia for New Zealand and Australian short-term rates. The model specification accommodates overlapping market forecast errors in the excess forward returns used to estimate the term premia. This allows us to examine the behaviour of term premia on forward rates with settlement dates more than one period ahead, whereas previous Kalman filter studies have examined the one-period-ahead case only. As a result, we are able to estimate a sizeable portion of the term premium profile. As well as producing constant term premia estimates for the base case, we consider two alternative time series specifications for the unobserved premia; namely, a mean reverting specification, and a random walk specification that allows for the possibility of permanent shocks to the term premia.

The motivation for examining both New Zealand and Australian short-term rates is that the money markets are similar in most respects, and some of the 'shocks' that could have altered the term premia over the sample period would have affected both markets to some extent. However, the operation of monetary policy has differed significantly between the countries at times. The RBA used the overnight cash rate as its policy instrument during the whole sample period, whereas the RBNZ used three distinct approaches during this time, adopting an overnight cash rate regime in March 1999. This analysis should highlight whether changes in the implementation method influenced the term premia.

The rest of the paper is organised as follows: Section 2 presents the model used to estimate the term premium, and the rationale behind the model. Section 3 describes the data used. Section 4 presents the empirical results. Section 5 concludes and notes avenues for further work. Tables and figures follow the references.

MODEL

Background

Conceptually, a forward rate can always be decomposed as follows:

$$f_{t,j} = E_t[r_{t+j}] + a_{t,j} \quad (1)$$

Where $f_{t,j}$ is the forward rate at time t with settlement date $t+j$, $E_t[r_{t+j}]$ is the conditional expectation of the j-period-ahead spot rate, given the information available at time t, and $a_{t,j}$ is the systematic difference between the spot and j-period-ahead forward rates, which is known as the 'term premium' because it is associated with the term to settlement. Subtracting the realised spot rate from both sides of (1) gives:

$$efr_{t,j} = a_{t,j} + u_{t+j} \quad (2)$$

where $efr_{t,j} = f_{t,j} - r_{t+j}$ is the j-period excess forward return, and $u_t = E[r_{t+j}] - r_{t+j}$ is the j-period market forecast error. Since expectations will be realised on average, the forecast error in equation (2) will have a mean of zero. The Expectations Hypothesis is defined by the restrictions placed on this relationship: first, the term premium alpha is assumed to remain constant over time. Second, the market forecast error is assumed to be uncorrelated with information available at time t. This corresponds to the standard definition of rational market expectations.

Model specification

We test three different specifications for the unobserved term premium in equation (2). The first is the "constant" specification, as implied by the EH:

$$a_{t,j} = \bar{a}_j \quad (3)$$

The second is the "non-stationary" specification, where the term premium is assumed to follow a driftless random walk:

$$a_{t,j} = a_{t-1,j} + v_t \tag{4}$$

The final specification is "mean-reverting", where the term premium is assumed to follow a stationary first-order autoregressive process:

$$a_{t,j} = c + f\,a_{t-1,j} + v_t \tag{5}$$

where f is less than one in absolute value. Note that the first two specifications are restrictions of the third; for the constant scenario, $f = 0$ and the variance of v_t is zero, and for the non-stationary case, $f = 1$ and the constant term c is not required.

The question of whether the term premium is stationary or non-stationary is more of an empirical than a theoretical one. Some evidence from the US suggests that interest rates themselves are mean-reverting over very long periods, but given the relatively small period examined here, we are unlikely to find strong evidence either way. Of course, the relevant issue is the degree of mean reversion – that is, the relative importance of transitory and permanent shocks to the term premia. The non-stationary specification tends to attribute more of high-frequency movements in excess forward returns to changes in the term premium. Hence, it is likely to provide the strongest case against the EH. On the other hand, the mean-reverting specification is less restrictive and is likely to better reflect the 'true' term premia.

The market forecast error in equation (2) is specified as a $(j$-1)th-order moving average process. This follows from the fact that market forecasts more than one period ahead will have overlapping forecast errors from each additional period, which are themselves and independent and normally distributed:

$$u_{t+j} = e_{t+j} + q_1 e_{t+j-1} + q_2 e_{t=j-2} + \ldots + q_{j-1} e_{t-1} \tag{6}$$

$$e_{t+j} \sim N(0, s^2_{t+j}) \tag{7}$$

State-space specification and Kalman filter

The model can now be written in state-space form. The state-space specification requires two equations: the *observation* equation, where the observed excess forward returns are expressed as a function of the term premium; and the *state* equation, which determines how the term premium evolves over time. The model below is for the mean-reverting case; as noted before, the other two cases are restrictions of this specification.

Equations (2) and (6) imply the following observation equation:

$$efr_{t+j,j} = \begin{bmatrix} 1 & 1 & q_1 & \cdots & q_{j-1} \end{bmatrix} \begin{bmatrix} t_{t,j} \\ e_{t+j} \\ e_{t+j-1} \\ \vdots \\ e_{t+1} \end{bmatrix} \qquad (8)$$

Or more simply,

$$efr_{t+j,j} = Hb_t . \qquad (9)$$

Equations (5), (6) and (7) give the following state equation:

$$\begin{bmatrix} t_{t,j} \\ e_{t+j} \\ e_{t+j-1} \\ \vdots \\ e_{t+1} \end{bmatrix} = \begin{bmatrix} c \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} + \begin{bmatrix} f & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix} \begin{bmatrix} t_{t-1,j} \\ e_{t+j-1} \\ e_{t+j-2} \\ \vdots \\ e_t \end{bmatrix} + \begin{bmatrix} v_t \\ e_{t+j} \\ 0 \\ \vdots \\ 0 \end{bmatrix} \qquad (10)$$

Or more compactly,

$$b_t = \tilde{c} + Fb_{t-1} + \tilde{v}_t . \qquad (11)$$

v_t represents the shocks to the process that generates the term premium, and e_t is the measurement error, due to the fact that the term premium is not directly observed (only the excess forward returns). The errors are independently distributed and are uncorrelated with past errors:

$$\tilde{v}_t \sim N(0,Q) \tag{12}$$

$$Q = E[\tilde{v}_t \tilde{v}_t'] = \begin{bmatrix} s_v^2 & 0 & 0 & \cdots & 0 \\ 0 & s_{t+j}^2 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix} \tag{13}$$

The Kalman filter is started by entering a set of priors for the vectors *H*, $\tilde{c}$, *F* and *Q*. The filter then estimates the vector of unobservable variables b_t period by period, using a two-step procedure. The first step is to "predict" the vector b_t given the information set in the period *t*-1. The second step is to "update" the prediction of b_t using the one-step-ahead forecast error of the excess forward returns (generated by the filter at each execution).

Because the Kalman filter generates a time series of the one-step-ahead forecast errors made in estimating $efr_{t+j,j}$, it can be run in conjunction with an optimisation routine to produce the maximum likelihood estimates of the parameters. The parameters are fixed for each run through the model, but the likelihood function can be calculated for each run. For all of the models in this paper, maximum likelihood estimates were obtained using the OPTMUM procedure in the GAUSS programming language. The parameters f, s_v and s_{t+j} were constrained appropriately. In most cases, the results were robust to a range of initial parameter settings.

DATA

The sample periods examined here are those for which a full data set is available. The data set begins in November 1993 for Australia and March 1995 for New Zealand, and both sets end in October 2002. After allowing for the required number of lags in the estimation, the model produces full estimates of the term premium function up to November 2001.

We use interest rates from the bank risk curve, which will naturally include a term premium due to the credit risk of banks. This data is preferable to government risk rates (that is, Treasury bills), because the markets are more liquid and there is a greater range of maturities. The rates are drawn from:

- The market cash rate, which is published every business day by the respective central banks. The actual market rate can differ from the policy rate, but under an overnight cash rate regime these deviations have tended to be small and are not persistent. Figure 1 shows the time series of the cash rates for New Zealand and Australia.
- Bank bills with monthly maturities up to six months. The rates are taken from the fixings submitted by price makers every business day, which are compiled and published by the respective Financial Markets Associations (NZFMA and AFMA).
- Forward rate agreements (FRAs) on three-month bank bills, with monthly settlement dates from four to nine months ahead. The FRA rates are end-of-day rates, and are taken from Bloomberg.

All of the observations are taken from the last business day of each month. The data is used to build a yield curve of spot rates out to 12 months, by compounding the appropriate spot and forward rates:

$$r_{t,b} = [(1 + r_{t,a} * a/36500) * (1 + f_{t,b} * b/36500) - 1] * \frac{36500}{b-a} \tag{14}$$

where a and b are the number of days to maturity for the relevant rates, and $b>a$. Forward one-month rates, with settlement dates one to eleven months ahead, are then calculated by rearranging equation (14).

Strictly speaking, the spot and forward rates must have the same length of time from settlement to maturity. Since we are comparing the one-day spot rate with forward one-month rates, we account for this by treating the "spot" rate as the average daily cash rate over the relevant month.[2] For example, the "spot" rate at the end of January will actually be the average of the cash rate during February. As a result, the one-month-ahead forward rate (j=1) will actually be the one-month bank bill rate at the end of January, since this is the market forecast of the average cash rate during February. For the horizon j=2, the forward rate would apply from the end of February to the end of March.

[2] The alternative approach, used in Krippner (2001) and Gordon (2002), is to use daily observations and interpolate forward one-day rates from the monthly rates.

Excess forward returns are calculated by expressing the spot and forward rates as continuously compounding rates, then subtracting the spot rate from the forward rate. Figures 2 and 3 show samples of the excess forward returns at various horizons for New Zealand and Australia. It appears that the time series become more persistent as the forward horizon increases, which is consistent with the implications of the moving average error process.

RESULTS

Constant specification

Figures 4 and 5 show the profiles of the constant term premia estimates for New Zealand and Australia. The results fit with the traditional view of the term premium function, in that the estimates of the premia are positive at every forward horizon, and tend to increase with the horizon. The premium 12 months ahead is about 60 basis points in both countries. This is slightly larger than the estimates from Krippner (2002) and Gordon (2002), which used data from 1999 onward. Notably, the estimates have very wide confidence intervals at every horizon - in fact, except for the horizon j=1, the constant term premia estimates are not significantly different from zero.

Mean-reverting and non-stationary specifications

Space considerations prevent a complete list of the parameter estimates for each specification and time horizon. Tables 1 and 2 report the log-likelihood values for the three term premium specifications and the twelve forward horizons, and Table 3 details the parameter estimates for the horizon j=4.

The log-likelihood values in Tables 1 and 2 provide clear evidence against the EH. Using the non-stationary estimates for New Zealand as the alternative, the implied likelihood ratio statistics for the null hypothesis $H_0 : s_v^2 = 0$ range from 3.78 to 6.90, and are almost all significant at the 5% level. Similarly, for Australia the non-stationary estimates are all significant at the 1% level.

If the mean-reverting specification is used for the null hypothesis, the likelihood ratios for New Zealand are not significant at the 5% level. However, the estimates of f for the mean-reverting specification are in the order of 0.7 to 0.9, and are not significantly different from 1.

Nevertheless, we avoid drawing any strong conclusions on whether the term premia are mean-reverting or non-stationary.

Interestingly, the mean-reverting estimates for Australia converge to the constant specification, with f estimated as zero for every forward horizon. Closer inspection shows that the horizons up to seven months are robust to a wide range of initial parameter settings, but for the eight-month horizon and beyond, the results are sensitive to the initial settings. For these longer horizons, it is possible to produce estimates where f is in the order of 0.8, which is comparable to the estimates for New Zealand. For the sake of consistency, the mean reverting estimates detailed here are those with $f = 0$ at every horizon.

While the time-varying estimates are statistically significant, there is also the question of whether they are economically significant. Tables 4 and 5 show the variation of the term premia as a proportion of the total variation in forward premia. The contribution of the term premia in New Zealand ranges from 20% to 60% for short horizons, but falls off significantly for longer horizons. For Australia, the term premia accounts for a significant portion of the total variance at almost every horizon. This implies that the contributions of the time-varying term premia are economically significant.

Figures 6 through 9 provide one perspective on the output from the Kalman filter. These figures show the cross-sections of the term premia estimates – that is, each line represents the term premia estimates for a particular month. It is apparent from these figures that the term premia for New Zealand have been much more variable at horizons of six months or less, whereas the Australian term premia have followed a more 'normal' pattern, with changes in the term premium being larger for longer horizons.

One possible explanation for this comes from work by Gurkaynak, Sack and Swanson (2002) showing that a significant portion of the market impact of a monetary policy 'shock' is due to the timing of the change in policy, rather than the size. Over longer horizons, the market may accurately forecast the total amount of monetary easing or tightening, but the timing of the changes is a significant source of forecast error. However, since this effect is not apparent in the Australian estimates, it is more likely to reflect the impact of changes to the implementation of monetary policy in New Zealand, which deserves a closer examination.

Changes in monetary policy implementation

Since 1989, the RBNZ's sole policy goal has been price stability, with an explicit inflation rate target. However, it is free to choose the tools it uses to achieve this goal. Over the sample period, there have been three distinct methods of implementing monetary policy. Prior to 1997, the central bank would issue statements to signal when it was uncomfortable with the current levels of the exchange rate and short-term interest rates. From 1997, the RBNZ used the Monetary Conditions Index (MCI), made up of the trade-weighted exchange rate and the 90-day interest rate, to explicitly signal the desired level of monetary conditions. During these two phases, the RBNZ would signal the desired setting frequently and irregularly. In March 1999 the MCI was discontinued in favour of an official cash rate (OCR), which is reviewed at scheduled intervals and changed in increments of 25 basis points.[3]

The change in the implementation method is likely to have led to a change in the transmission mechanism between the cash rate and bank-bill rates. Prior to the OCR regime, the market cash rate was highly volatile, and it would take a significant period of time for the market to decide whether a fall in the cash rate was genuine or simply the result of 'noise'. Because of this ambiguity around the true level of the cash rate, it is plausible that the premium required on bank bills was significantly larger during this period.

If so, the constant and mean-reverting estimates of the mean will be biased upward by the data prior to March 1999. As confirmation of this, we estimated the term premia for New Zealand using data from March 1999 to October 2002. Figure 10 shows that the constant premia estimates were about six basis points lower than those using the full sample period. The mean-reverting estimates converged to a constant premium, although this may be a result of the very small sample size and should be treated with caution. Even so, it is consistent with the fact that excess forward returns were relatively stable over this period.

The non-stationary estimates do not rely on an estimate of the mean, so the results would be unchanged if the sample period is restricted. Figure 11 shows the cross-section of the non-stationary premia estimates from March 1999 to October 2002. Compared to Figure 7, it is apparent that most of the extreme values for the term premia estimates related to the period prior to the OCR regime.

---

[3] The official cash rate is not a target rate, but rather the mid-point of the rates at which the RBNZ is willing to borrow and lend cash overnight.

Single-factor representation

To examine how the profile of term premium evolves over time, it is useful to express it as a function of the time horizon. We would expect the shape of the premium function to resemble the 'normal' shape of the yield curve – that is, increasing with the time horizon but at a decreasing rate. As with yield curve analysis, there is a wide range of suitable specifications for 'fitting' a function to the data. Here we have used a one-parameter square-root function, which is not meant to fully explain the shape of the term premium function, but is useful for demonstration purposes.

$$\overline{a}_t = j\sqrt{k} \qquad (15)$$

where k is defined as the forward horizon in days, and is treated as the mid-point of the forward horizon, i.e. it is 15 days for j=1, 45 days for j=2, and so on. This is because the monthly term premium can be though of as the average of the one-day premia over the period, or approximately the one-day premium at the midpoint of the period. This produces a smooth function for daily horizons, which can then be used to calculate the premium for any forward horizon, such as weekly or quarterly horizons.

This function is fitted to the term premia estimates for each month, for the mean-reverting and non-stationary estimates. The time series of the slope coefficients j are shown in Figure 12 for New Zealand and Figure 13 for Australia. (Note that a 0.01 change in j is roughly equivalent to a 19 basis point change in the term premium at the one-year horizon.) The first thing that is notable is the lack of any correlation between the term premium estimates for New Zealand and Australia (in fact, the correlation for the non-stationary estimates is –0.35). While short-term rates themselves have often diverged, the term premia on these rates could be expected to be reasonably similar for New Zealand and Australia, if they respond similarly to external shocks. This suggests that variations in the term premia may be largely due to country-specific factors.

Second, it shows that New Zealand's term premium function was negative on several occasions between 1995 and 1997 (more so if we examine the non-stationary estimates). Gravelle and Morley (2002) also produced negative term premia for Canada, and suggested the presence of a "peso problem", as discussed in Bekaert, Hodrick and Marshall (1997).

Early on in an inflation-fighting regime, the market may be sceptical about the central bank's commitment to the policy, and may continue to anticipate a return to more expansive monetary policy. As a result, markets may under-predict short-term rates for extended periods, and still be considered rational given the information available at the time. While inflation was well established in New Zealand by 1995, it is still possible that the RBNZ kept short-term rates high for longer than the market believed was required.

CONCLUSIONS AND FURTHER WORK

This paper provides some evidence of a time-varying term premium in short-term rates in New Zealand and Australia, which contradicts the traditional view of the Expectations Hypothesis. The variation in the term premium is both statistically and economically significant. The non-stationary specification of the term premia has significantly greater explanatory power than the constant premia estimates. The mean-reverting specification does not appear to be significantly better than the constant premia estimates – this is somewhat surprising, given that this specification is the least restrictive. However, the autoregressive parameter f is generally close to one, which suggests that it is reasonable to restrict the term premium to a non-stationary specification.

A major caveat is that the results appear to be greatly influenced by the monetary policy implementation regime. The Australian central bank used the cash rate as its policy instrument over the whole period examined here, and New Zealand has used it since March 1999. Over these data sets, the term premia are relatively stable – so much that the time varying estimate of the premium often converges to a constant premium. However, we should be careful about drawing conclusions about New Zealand post-1999, due to the small sample size. These estimates will benefit from re-estimation in coming years as more data becomes available.

One possible use of these results is to examine the market reaction to past economic and policy shocks. For example, monetary policy 'surprises' depend on market expectations prior to the policy decision; understanding the degree of surprise requires us to know the size of the term premia at the time. Another avenue is to examine the impact of external shocks; for example, the estimates presented here end a couple of months after the terrorist attacks on 11 September 2001. It will be interesting to see how the term premia evolved after this event – again, this will be possible as more data becomes available.

Another task remaining is to relate the term premia estimates to observable variables, with the goal of finding some contemporaneous indicators of the size of the premia. We are able to describe how the premium evolved up to about a year ago; ideally we would also like to have a gauge of where the premium lies today. Market participants tend to use simple indicators to adjust their assumptions about the term premium, such as the slope of the money market yield curve. This in turn largely reflects the current stage in the monetary policy cycle.

Another unresolved issue is finding a parsimonious representation of the term premium profile. It is possible that estimating the term premia in a panel specification, rather than equation-by-equation, could shed some light on this issue. It may be necessary to use an estimation technique that filters the estimates cross-sectionally as well as over time.

REFERENCES

Bekaert, G, R Hodrick and D Marshall (1997), "'Peso problem' explanations for term structure anomalies", *Federal Reserve Bank of Chicago Working Paper* 97-7.

Bhar, R (1993), "Modelling Australian bank bill rates: a Kalman filter approach", *University of Western Sydney Discussion Paper* E9313.

Cook, T and T Hahn (1990), 'Interest rate expectations and the slope of the money market yield curve', *Federal Reserve Bank of Richmond Economic Review* 76, 3-26.

Engle, R, D Lilien, and R Robins (1987), "Estimating time varying risk premia in the term structure: the ARCH-M model", *Econometrica* 55(2), 391-407.

Gordon, M (2002), "Deriving expectations of the Australian cash rate target from money market rates", *Econometric Society Australasian Meeting*, July 2002.

Gravelle, T, and J Morley (2002), "A Kalman filter approach to characterizing the Canadian term structure of interest rates", *Journal of Money, Credit and Banking*, forthcoming.

Gurkaynak, R, B Sack and E Swanson (2002), "Market-based measures of monetary policy expectations", *Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series* 2002-40.

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.

Iyer, S (1997), "Time-varying term premia and the behaviour of forward interest rate prediction errors", *Journal of Financial Research* (20), 503-7.

Karfakias, C, and D Moschos (1995), "The information content of the yield curve in Australia", *Journal of Macroeconomics* 17(1), 93-109.

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 (2002), "Extracting expectations of New Zealand's Official Cash Rate from the bank-risk yield curve", *Reserve Bank of New Zealand Discussion Paper* 2002/01.

Krippner, L and M Gordon (2001), "Market expectations of the Official Cash Rate", *Reserve Bank of New Zealand Bulletin* 64(2), 25-39.

Lee, S-S (1995), 'Macroeconomic sources of time-varying risk premia in the term structure of interest rates', *Journal of Money, Credit and Banking* 27, 549-69.

Lowe, P (1994), "The link between the cash rate and market interest rates", *Reserve Bank of Australia Discussion Paper* 9504.

Mankiw, G and J Miron (1986), "The changing behaviour of the term structure of interest rates", *Quarterly Journal of Economics* 101, 211-28.

Table 1: Log-likelihood values for New Zealand.

	Term premia specifications		
Forward horizon	Constant	Mean-reverting	Non-stationary
1	65.64	66.38	67.53
2	80.90	82.14	83.96*
3	79.29	80.44	82.05*
4	81.65	83.44*	84.79*
5	80.55	82.78*	84.00**
6	77.61	79.34	80.80*
7	83.62	86.21*	86.84*
8	84.01	86.36*	87.12*
9	79.04	81.00*	82.15*
10	82.42	84.27	85.48*
11	77.35	79.84*	80.44*
12	80.00	81.66	83.04*

Table 2: Log-likelihood values for Australia.

	Term premia specifications		
Forward horizon	Constant	Mean-reverting	Non-stationary
1	26.43	26.43	36.03**
2	48.85	48.85	58.64**
3	62.05	62.05	66.74**
4	70.53	70.53	74.54**
5	74.84	75.20	81.10**
6	81.29	81.29	86.70**
7	86.17	86.17	90.13**
8	84.31	84.36	88.98**
9	85.16	85.38	88.69**
10	92.71	92.76	98.63**
11	84.54	85.03	88.78**
12	82.86	82.97	86.40**

Critical value for the $\chi^2(1)$ distribution is 3.84 at the 5% level and 6.63 at the 1% level. * (**) denotes that the estimates are more significant than the constant premium estimates at the 5% (1 %) level.

Table 3: Parameter estimates for New Zealand and Australia, horizon *j*=3.

	New Zealand			***Australia***		
Specification	*Constant*	*Mean-reverting*	*Non-stationary*	*Constant*	*Mean-reverting*	*Non-stationary*
c	0.282	0.290	-	0.142	0.142	-
	(0.244)	*(0.415)*		*(0.103)*	*(0.103)*	
f	-	0.852*	-	-	0.000	-
		(0.111)			*(0.002)*	
s_v	0.049	0.479*	0.334*	0.018	0.088	0.000
	(6.391)	*(0.214)*	*(0.159)*	*(0.025)*	*(3.421)*	*(0.011)*
s_{t+j}	0.697*	0.314*	0.486*	0.411*	0.153	0.276*
	(0.091)	*(0.151)*	*(0.206)*	*(0.164)*	*(0.462)*	*(0.048)*
q_1	1.015	1.302*	1.182*	0.629*	1.413	1.465*
	(1.292)	*(0.472)*	*(0.524)*	*(0.287)*	*(2.364)*	*(0.276)*
q_2	0.742	1.677	0.809*	0.518	1.700	0.498*
	(0.528)	*(1.053)*	*(0.325)*	*(0.268)*	*(4.182)*	*(0.179)*
q_3	0.458	1.375	0.628	0.359	2.593	0.829*
	(0.322)	*(0.967)*	*(0.365)*	*(0.257)*	*(5.557)*	*(0.244)*

Figures in bracket are the standard errors. * indicates that the parameter estimates are significant at the 5% level.

Table 4: Standard deviations of innovations to term premia for NZ.

Forward horizon	$s_{\Delta fp}$	s_v	*Contribution*
1	0.430	0.095	22.1%
2	0.390	0.224	57.4%
3	0.402	0.330	82.2%
4	0.394	0.334	84.7%
5	0.415	0.275	66.2%
6	0.407	0.263	64.6%
7	0.465	0.002	0.4%
8	0.463	0.003	0.6%
9	0.443	0.003	0.7%
10	0.474	0.008	1.7%
11	0.472	0.003	0.6%
12	0.458	0.001	0.2%

Table 5: Standard deviations of innovations to term premia for Australia.

Forward horizon	$s_{\Delta fp}$	s_v	*Contribution*
1	0.154	0.117	75.8%
2	0.208	0.195	93.7%
3	0.268	0.007	2.6%
4	0.272	0.121	44.5%
5	0.326	0.316	96.9%
6	0.362	0.274	75.6%
7	0.415	0.066	16.0%
8	0.378	0.295	78.1%
9	0.383	0.355	92.6%
10	0.474	0.434	91.7%
11	0.411	0.282	68.7%
12	0.399	0.130	32.5%

Forward premia statistics are calculated using the first differences of the constructed series $fp_{t,j} = f_{t,j} - r_t$. Term premia statistics are maximum likelihood estimates for the non-stationary specification.

Figure 1: Monthly market cash rates for New Zealand and Australia.

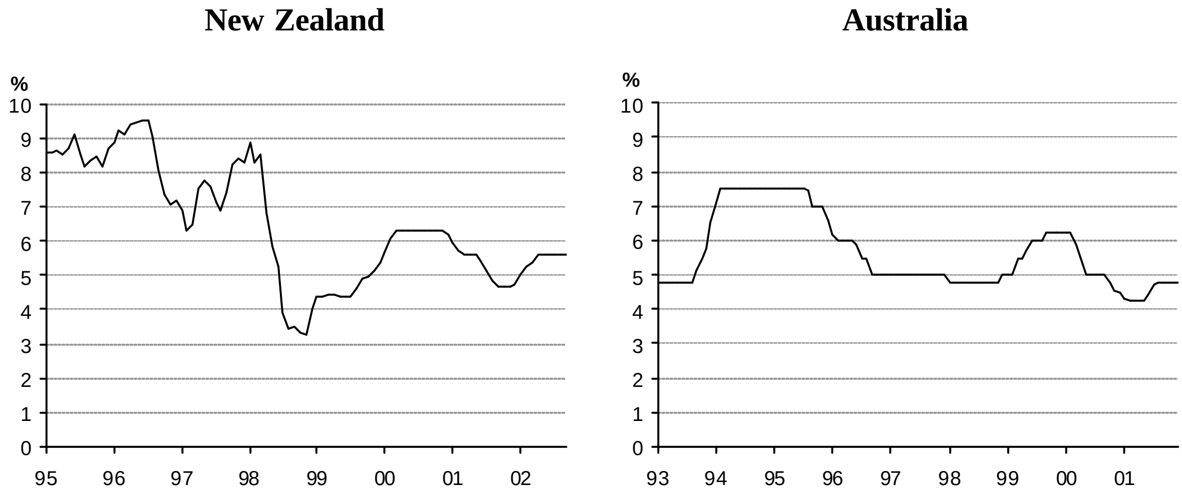

Figure 2: Sample of excess forward returns for New Zealand.

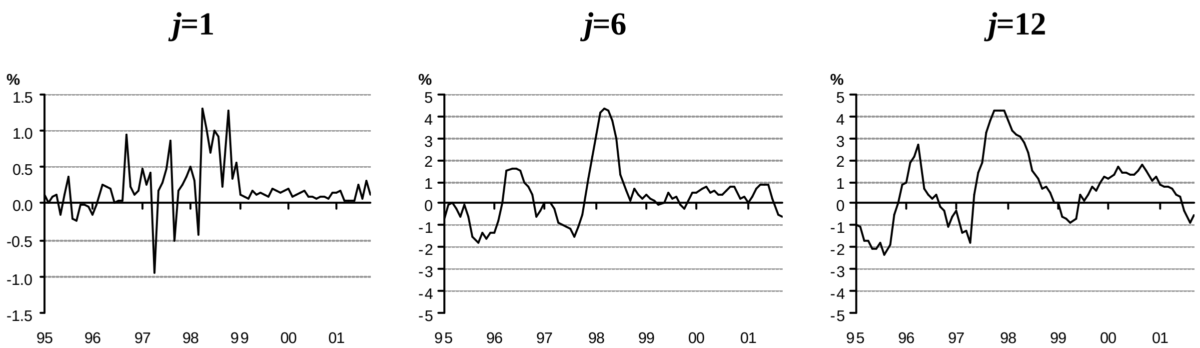

Figure 3: Sample of excess forward returns for Australia.

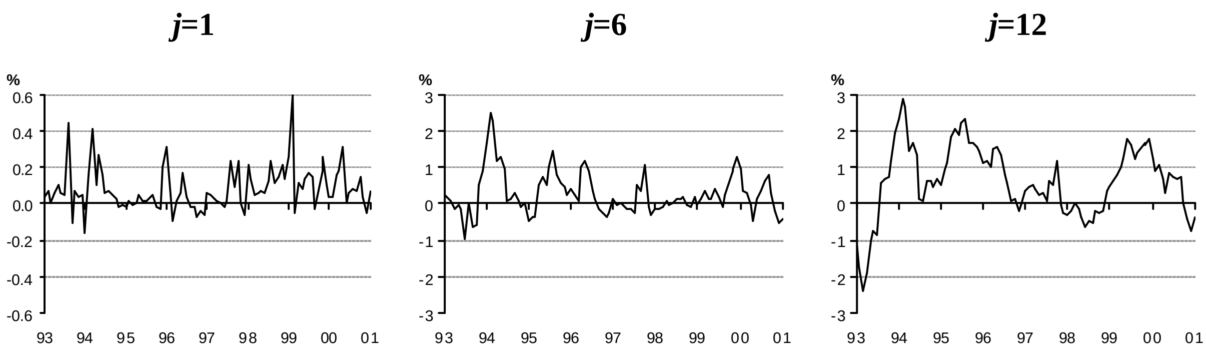

Figure 4: Constant term premium estimates for New Zealand.

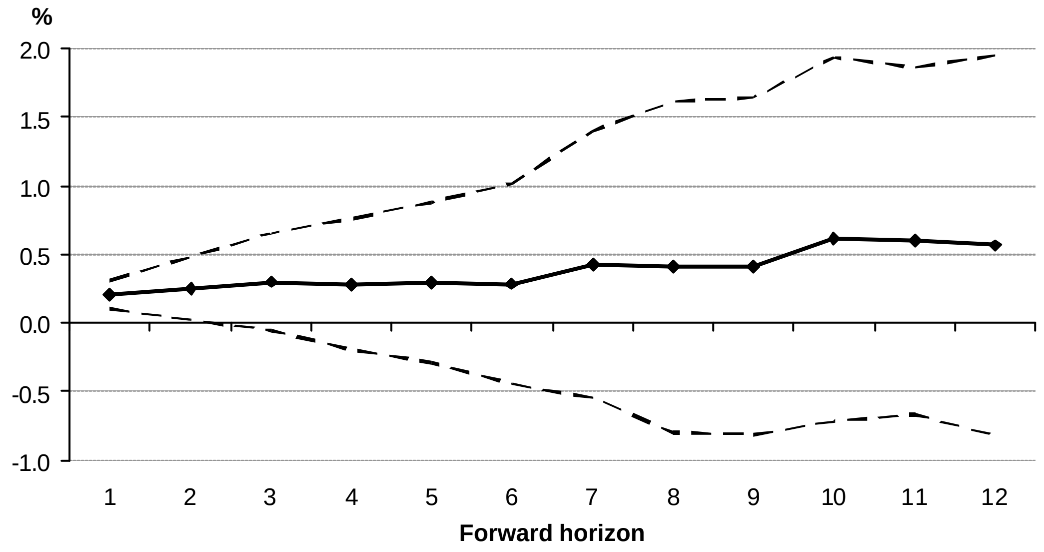

Figure 5: Constant term premium estimates for Australia.

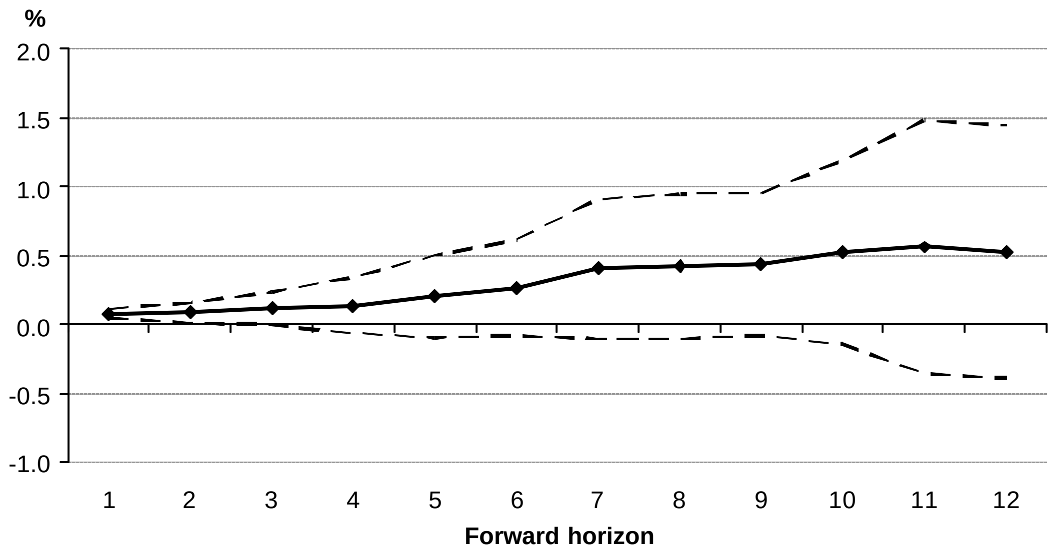

The dashed lines are the 95% confidence intervals of the estimates.

Figure 6: Profile of mean-reverting estimates of term premium function for NZ.

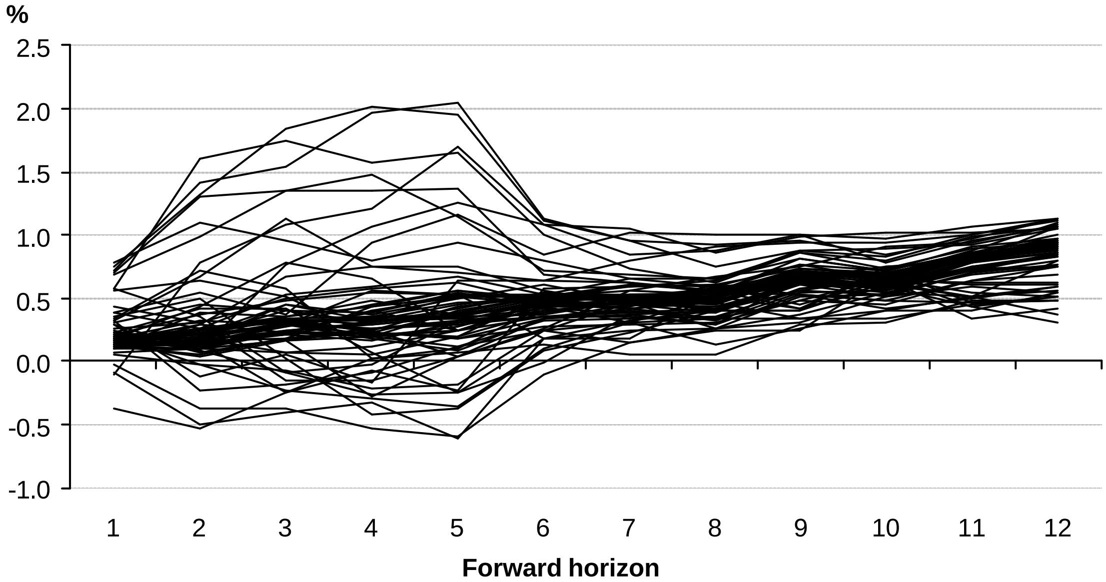

Figure 7: Profile of non-stationary estimates of term premium function for NZ.

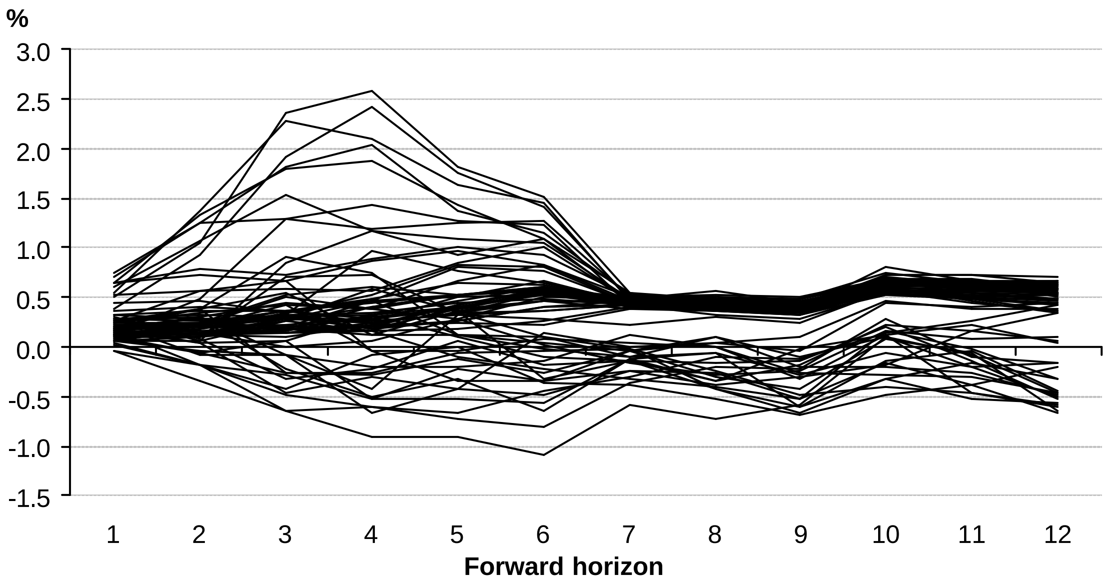

Figure 8: Profile of mean-reverting estimates of term premium function for Australia.

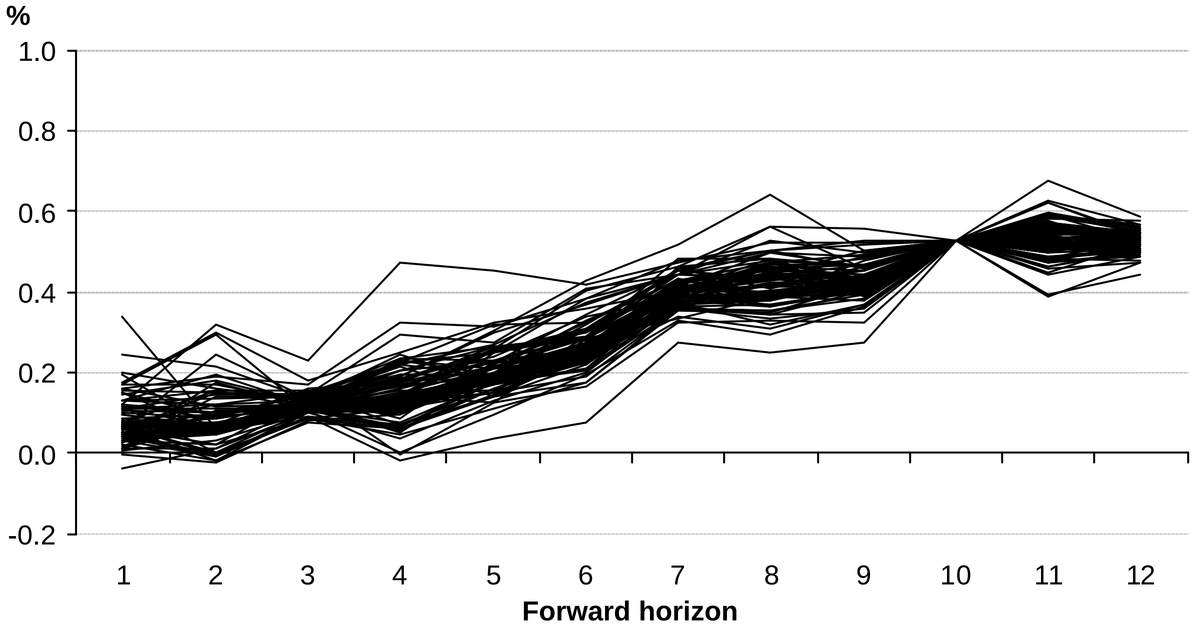

Figure 9: Profile of non-stationary estimates of term premium function for Australia.

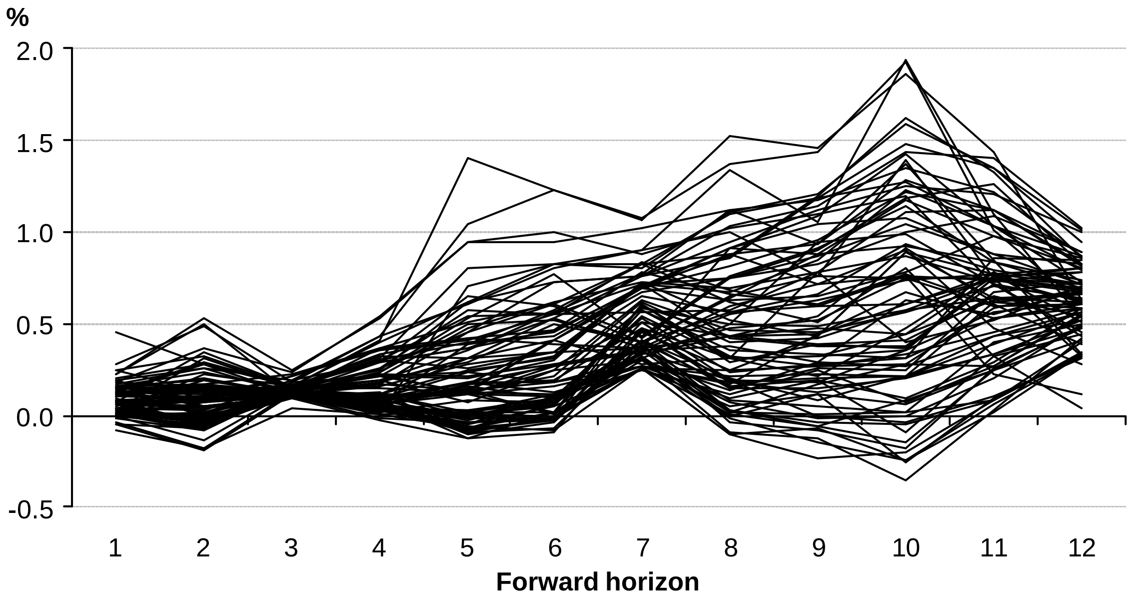

Figure 10: Term premia response to change in NZ monetary policy regime

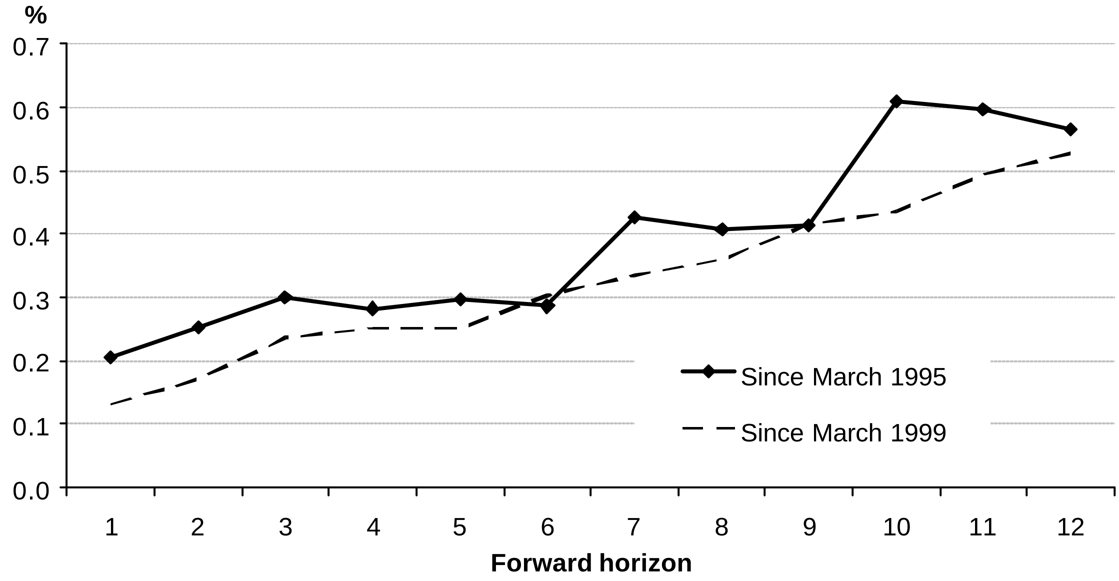

Figure 11: Profile of non-stationary estimates for NZ from March 1999.

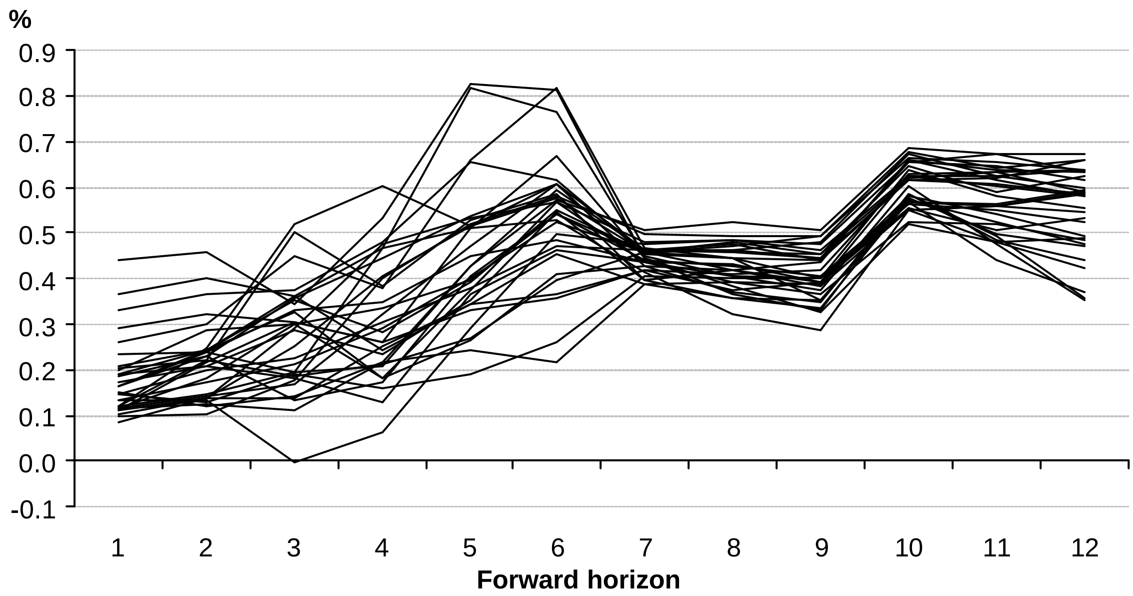

Figure 12: Term premium slope parameters for New Zealand.

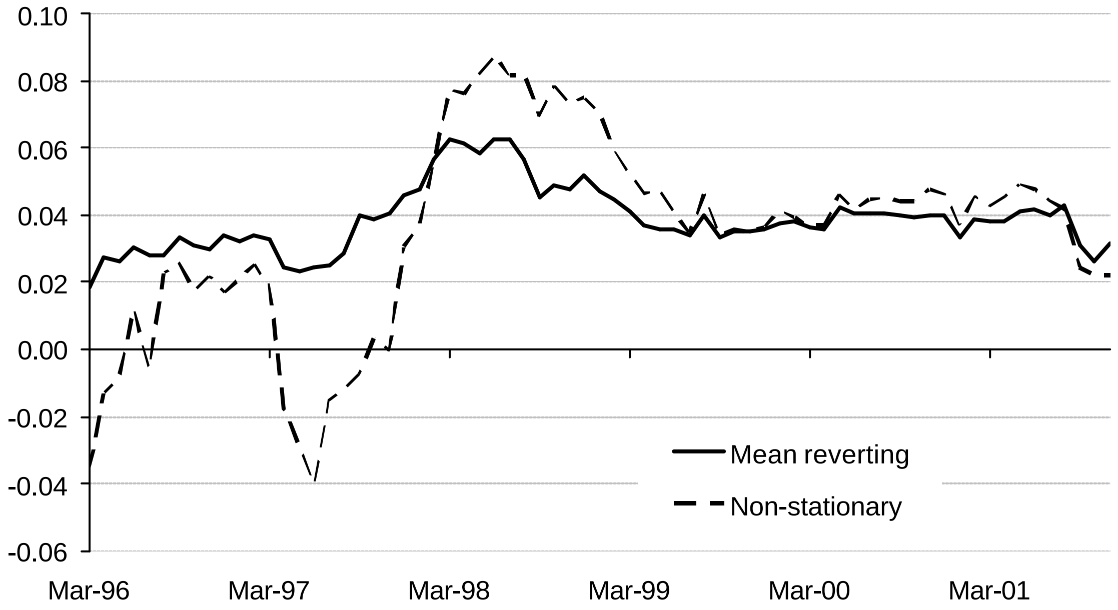

Figure 13: Term premium slope parameters for Australia.