THE NATIONAL BANK OF NEW ZEALAND LIMITED

Forecasting Volatility

DAVID RAE

ECONOMICS DIVISION

FINANCIAL RESEARCH PAPER NO. 9 - AUGUST 1997

FOR ALL ENQUIRIES PLEASE CONTACT:

The National Bank of New Zealand Limited
Economics Division
P O Box 540
WELLINGTON
Fax No: 0-4- 473 4928
Phone No: 0-4-802 2212
Internet: www.nbnz.co.nz

EXECUTIVE SUMMARY

Many derivative markets can be thought of as places in which people trade 'volatility'. And wherever there are traders, there is a demand for forecasters. This paper compares different methods for forecasting the volatility of NZ's interest rates and exchange rates. The methods are compared based on their ability to forecast out-of-sample at a variety of horizons: 1, 3, 6, and 12 months ahead. The results have significant implications for the pricing of options and other derivatives.

Eight forecasting methods are compared: four 'simple' methods, and four 'regression-based' methods. The simple methods consist of: the assumption of no change (i.e., random walk); the historical average; a three-month average of historical volatility; and a one-year average of history. The regression-based methods are: OLS (which is closely related to GARCH models of returns); Least Absolute Deviations (which places less weight than OLS on extreme observations); a Vector Autoregression or VAR (which allows for the correlation amongst volatilities); and a regression model in which the coefficient values can vary through time (using a Kalman filter).

The main conclusions are:

- The assumption of no change to volatility is a poor predictor over all forecast horizons. The average forecast error can be halved by choosing one of the other methods. This should not be surprising: there is a pattern to volatilities that we should be able to model and use for forecasting.
- The regression methods are the clear winners at all forecasting horizons out to one year ahead. This is especially true over the post-1989 period. This implies that GARCH models would also do well.
- There is no clear difference between the different regression models. Methods that allow for outlier observations, volatility correlations, or changing coefficients, give forecasts that are no better than simple OLS.
- Moving average models perform well, especially over periods that show a large structural shift in volatility. This is because in these cases they discard the now-irrelevant historical information.
- If you are going to use a moving average method, the one-year average always does better than the three-month average. This is because it captures the mean reversion better.

In summary, the preferred forecasting method is simple OLS. But if there is a clear structural break in the policy *framework* then more emphasis should be placed on moving average models because they place less weight on historical information.

CONTENTS

1. INTRODUCTION

There has been an explosion in markets that trade volatilities. Options markets are the best known, but several of the other derivatives markets can be thought of as markets in which traders buy and sell volatility. At the same time, there has been an explosion of research aimed at modelling movements in volatility, but there have been two comparative gaps in the research. The first is that little work has been published specific to New Zealand's financial markets. The second is that most studies have modelled volatility over a certain period but not asked whether those models have any significant forecasting power out of sample.

Simple within-sample comparisons of different volatility models are only moderately useful because the most flexible models – the ones with the greatest number of free parameters to be estimated – will always fit the data better. A more powerful test is to compare models based on how well they forecast volatility outside the sample on which they were estimated. We compare different methods based on their ability to forecast up to one year out of sample.

We first construct volatility measures of ten financial variables: four interest rates (30-day and 90-day bills; 3-year and 10-year bonds); and six exchange rates (the USD, AUD, GBP, JPY, DEM, and TWI, all measured relative to the New Zealand dollar).[1] Eight methods of forecasting – of various degrees of complexity – are compared based on 4, 12, 26, and 52 week ahead out-of-sample forecasts.

To get a feel for the data, Figure 1 shows returns on the NZD/AUD cross rate, defined as the weekly percentage change in the currency. Figure 2 graphs the volatility of this exchange rate from 1986 to 1997, plus a 9-week centered moving average of the volatility (the next section describes how volatility is calculated). The corresponding graphs for the other variables have features that are similar to the NZD/AUD rate.

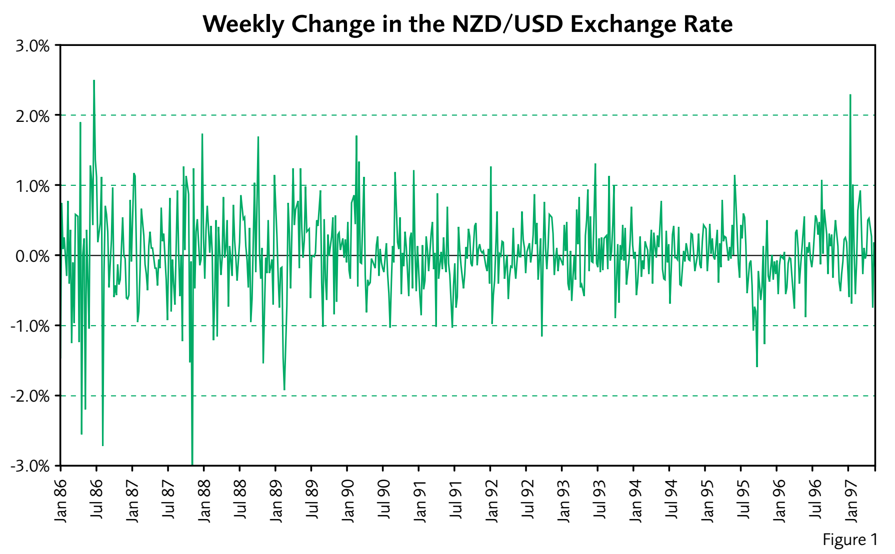

Figure 1

1 The TWI is the trade-weighted exchange rate index, consisting of the other five currencies.

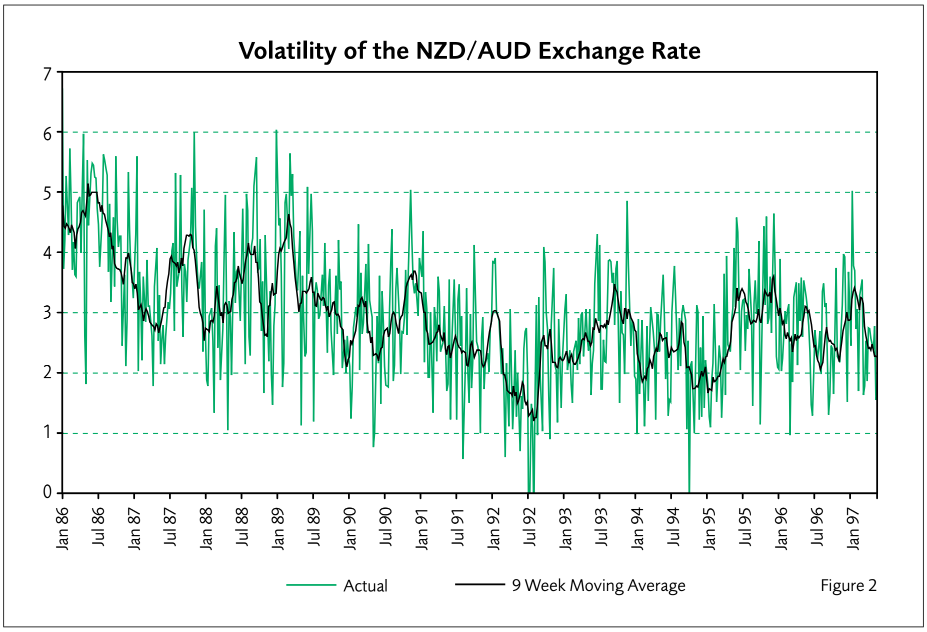

Figure 2

Several features are apparent from these graphs and from a more detailed examination of the data:

- *Volatility has fallen.* The exchange rate is noticeably less volatile from 1989 onwards. The Reserve Bank put more emphasis on the exchange rate when setting monetary policy over this period. Since 1989, volatility has shown no trend in either direction and looks distinctly mean-reverting. Unit root tests confirm that the volatility of all ten variables is stationary (mean-reverting) over this period.
- *Volatility comes in clusters.* There are clear periods of high volatility followed by calm patches. This finding is common overseas and led to the development of GARCH models of financial returns. It also suggests that volatility is not random and should therefore be predictable.
- *Long memory and persistence.* Although volatility appears to be mean-reverting, the process seems surprisingly slow. Volatility can remain unusually high for periods of a year or more.
- *Volatility is not isolated.* Volatilities of different variables are correlated. From 1989 onwards the correlation coefficients amongst exchange rate volatilities are all positive and vary between 0.2 and 0.7 (this suggests that much of the volatility reflects the volatility of the NZD against *all* currencies). Volatilities amongst the interest rates are also positively correlated while the correlation between interest rate and exchange rate volatility is approximately zero.

These features give some hints about what types of forecasting models may work well. First, the structural drop in volatility after 1989 suggests that simple historical averages may not work well. Second, the mean-reversion and the apparent pattern suggests that models that place more weight on recent data should do well. Third, these same two features also suggest that an assumption of no change (i.e., a random walk) will not do well. Section 4 puts these guesses to the test in a more formal way.

2. MEASURING VOLATILITY

The 'right' measure of volatility depends on the theory of asset pricing that we are testing. Encouraged by the pioneering work of Louis Bachelier (1900), option pricing models have worked in continuous time and hence some variant on the *instantaneous variance* of the asset pricing process has been used to measure volatility. However, there is no consensus on which of the many pricing models should be used. Should we work with the geometric Brownian motion assumed by Black and Scholes (1973) or broaden the process to include jump diffusion processes (Merton 1976), mean-reversion in volatility (an Ornstein Uhlenbeck process), variable (stochastic) discount rates, and so on?

The most common approach is to assume that the Black Scholes formula is correct and use it to back out the *implicit volatility.* The snag with this approach is that it is contradictory: the Black Scholes formula assumes that volatility is constant. This led to models of stochastic (changing) volatility (eg, Hull and White 1987) which make more theoretical sense but have the drawback that they depend on parameters that are not known. This makes their pricing formulas difficult to use in practice. A fuzzy compromise was reached whereby the Black Scholes implicit volatility was interpreted as an implied *average* volatility from a more complex stochastic volatility model.

These problems led to a second strand of research that used *ad hoc* estimates of volatility. These include a simple variance of returns, the average of the absolute value of returns, Schwert's measure (which is just a jazzed-up version of the average absolute return), and various one-sided measures of volatility which assume that, for example, investors only worry about volatility on the down-side and not on the up-side. However these one-sided measures are becoming less relevant because investors can go short in most markets in which case the up-side becomes the down-side and *vice versa.*

In the interests of simplicity – and because I am not testing any asset pricing models – I work with a volatility measure based on the average of squared daily returns. If r_t is the daily rate of return then a measure of weekly volatility is

$$\sigma_t^2 = \tfrac{1}{n}\sum_{i=1}^{n} r_t^2 \qquad (1)$$

where n is the number of working days in that week. It is common to σ_t^2 use as the measure of volatility (e.g. Brailsford and Faff 1996) but this will cause severe modelling problems. Obviously, σ_t^2 is always greater than zero and hence its distribution will be highly skewed upwards – making it almost impossible to get a sensible forecasting model based on OLS for example. The solution is to take logs of (1) which 'squares up' the distribution. I then forecast σ_t rather than σ_t^2.

3. FORECASTING MODELS

Each forecasting model is estimated from January 1986 to April 1997 using weekly data. This covers the period over which New Zealand's financial markets have been deregulated. The models can be grouped into 'simple' models and 'regression-based' models. They are described below:

Simple Models

1. Random Walk (no change)

The random walk model assumes that the best forecast of next week's volatility is this week's volatility, *ie* no change. That is,

$$\sigma_{t+1}^{RW} = \sigma_t \qquad (2)$$

The 4, 12, 26, and 52 week ahead forecasts have a similar form.

2. Historical Average

If volatility were mean-reverting, we would expect the historical average to provide a reasonable forecast. The forecast based on the historical average is

$$\sigma_{t+1}^{HA} = \frac{1}{t}\sum_{i=1}^{t} \sigma_t$$

3. Moving Average

The historical average places equal weight on all volatility observations since the start of the sample (1985). If there has been a fundamental shift in market volatility we would expect the historical average to forecast poorly. An alternative is to put more weight on more recent history. To do this we test two moving average models: a three-month and a one-year moving average.

Regression-Based Models

4. OLS

This model uses an OLS regression of volatility on recent observed volatilities. We use the three most recent weeks' data, so the OLS forecast becomes

$$\sigma_{t+1}^{OLS} = \beta_0 + \beta_1 \sigma_t + \beta_2 \sigma_{t-1} + \beta_3 \sigma_{t-2} \qquad (3)$$

where the β's are the estimated coefficients from a regression of σ_t on a constant and three lags of itself.

OLS forecasts of volatility are closely related to GARCH models. For example, let r_t be returns from some asset and assume that returns are a random walk. Then a GARCH (0,3) model of returns would be

$r_t = \gamma + \varepsilon_t$, where $\varepsilon_t \sim N(0, \sigma_t^2)$ and

$$\sigma_t^2 = \beta_0 + \beta_1 \sigma_{t-1}^2 + \beta_2 \sigma_{t-2}^2 + \beta_3 \sigma_{t-3}^2 \qquad (4)$$

The idea behind GARCH models and their variants is that more efficient estimates of the process driving returns can be estimated if the returns themselves and the volatility of the returns are modelled together. The hope is to catch the patterns in both the level and the variance of returns. But equations (3) and (4) are clearly similar in form; this shows that an OLS regression of volatilities is capturing basically the same patterns as a GARCH model but takes the simpler approach of not bothering to estimate a returns equation for r_t alongside it. If OLS models give good forecasts of σ_t, then this would imply that a GARCH model of r_t would also do well.

5. Least Absolute Deviations (LAD)

As is now well established, the distribution of returns tends to have fat tails: there are more outlier observations than would be expected if the underlying distribution was normal. This can cause problems for OLS estimation because OLS can place too much weight on the handful of extreme observations. To avoid this, we estimate the same regression equation (3) but estimate it using LAD rather than OLS.

6. Vector Autoregression (VAR)

Both OLS and LAD are single-equation methods – they ignore the possibility that volatility may be related across markets. For example, volatility of the USD and AUD exchange rates may be positively related while volatility of exchange rates and interest rates may be negatively related. To allow for this possibility, equations for all ten variables (the six exchange rates and four interest rates) are estimated together as a VAR and the VAR is then used to forecast volatility.

7. Time-varying coefficient models

The three regression methods discussed so far assume that the coefficients are fixed over the estimation period. This may be especially misleading for the constant term because we saw earlier that volatilities appear to have fallen over time – especially exchange rate volatilities. Hence, we run the same regression but allow the coefficients to vary through time. This is done using a Kalman filter, in which the coefficients follow random walks.

4. TESTING METHOD

Our testing approach is to compare the forecasting ability of each model as if they were being used in real time by traders. The database starts in January 1986 while the testing itself starts 150 weeks later in November 1988. In the first step, each forecasting model is estimated on data up to November 1988. Hence we calculate the historical mean up to that point, do the OLS and LAD regressions using data from Jan 1986 to Nov 1988, etc. Then a set of 4, 12, 26, and 52 week-ahead forecasts are calculated using these models. Forecast errors are then calculated and stored. Next, we add one more week to the sample, re-estimate all the models, calculate a new set of forecasts, and keep repeating the process until all the data up to April 1997 has been used. Hence, we calculate out-of-sample forecasts for the 433 weeks from Nov 1988 to April 1997. This procedure is carried out for each of the ten variables being tested.

The forecasting models are compared using two measures of forecast accuracy: the mean squared error (MSE) and mean absolute error (MAE). They are defined as

$$MSE = \frac{1}{433} \sum (\sigma_t - \sigma_t^F)^2$$

$$MAE = \frac{1}{433} \sum |\sigma_t - \sigma_t^F|$$

where σ_t^F is the forecast volatility at time t (this forecast would have been prepared some weeks earlier). Both measures led to a similar ranking of forecasting models and so only the MSE results are reported here.

5. RESULTS

We found that that the results for the bills (30-day and 90-day rates) were very similar to each other. Similarly, forecasts of bond volatility were similar to each other, as were the exchange rate volatility forecasts. Hence we only report the results for the 90-day and 10-year rates, the USD and the TWI.

Summary statistics are shown in Tables 1(a) – 1(d). Each table shows the average (MSE) forecast error for each method *relative to* the no change (random walk) model. For reference, we also show the MSE of the no change model, measured in percentage terms.

The most obvious conclusion to draw from the tables is that the **no change model** is universally bad across all variables and over all forecasting horizons. The average forecast error can be halved by choosing one of the alternatives. This confirms what was seen earlier – that there is a pattern in the volatilities and thus the assumption of random volatility leads to poor predictions. A second conclusion is that the **simple historical average** is a poor predictor and it is worse for the exchange rates than for the interest rates. In sum, the two simplest methods – putting zero weight on past history (the no change model) and putting *equal* weight on each week in history (the historical average) – are poor forecasting methods.

The two moving average models use some historical information but discard it after a certain point. For instance, the **3-month moving average** places equal weight on each week for the past three months and forgets everything that happened before then. On balance, the moving average models are reasonable predictors of volatility but their performance relative to the other methods gets worse at longer forecasting horizons. The **one-year moving average** typically gives better forecasts than the three-month average and this probably reflects the mean-reversion that we saw earlier.

When forecasting interest rates, the **regression methods** provide slightly better forecasts than the simple methods, especially at longer forecast horizons. But this conclusion does not hold for the exchange rates. The best predictor of exchange rate volatility is the one-year moving average. The four regression methods give similar results to each other although the changing-coefficients model performs erratically: sometimes it forecasts very well, other times it does poorly. It is not a robust forecasting method.

To check whether these conclusions are sensitive to the estimation period, the first few years of the sample were dropped and the analysis repeated. This time the database starts in June 1989 and the forecasting was done from January 1990 to April 1997. This period was chosen because it covers the period over which the Reserve Bank placed more emphasis on the exchange rate when setting monetary policy. It therefore covers the period over which volatilities were stable and removes the structural break from the early part of the sample. The results are in Tables 2(a) – 2(d).

Most of the conclusions hold true over the shorter sample except that the **regression methods** suddenly come into their own. The regression methods become the clear winners for both interest and exchange rates and at all forecasting horizons. Their forecast errors are typically half the errors of the worst performing method. Again, there is no clear difference between the regression methods themselves. The changing coefficients model remains touchy and, on average, there does not seem much advantage in using VAR or LAD models compared with the simpler – and more robust – OLS approach.

The conclusion that regression methods work best over the second part of the sample is revealing. Three of the four regression methods assume a stable process driving volatility and this is clearly not always true. Changes in the behaviour of financial markets and changes in the way policy is conducted can lead to structural shifts in relative volatilities. Regression methods are poor at handling these shifts while simple moving averages do better because they discard the now-irrelevant historical information. However, between the breaks when the policy *framework* is stable the regression methods are very good at picking the patterns in the data and use these patterns well when forecasting. The trick is knowing when a fundamental structural shift has occurred and progressively putting more emphasis on moving average methods.

To help put the results in perspective, Figure 3 graphs 90-day rate volatility along with one-month ahead forecasts produced by the OLS and one-year moving average methods. It is clear that OLS produces significantly better forecasts in this case, and that there is enough variation in the forecasts to make them non-trivial. Given that these are out-of-sample forecasts, OLS does surprisingly well.

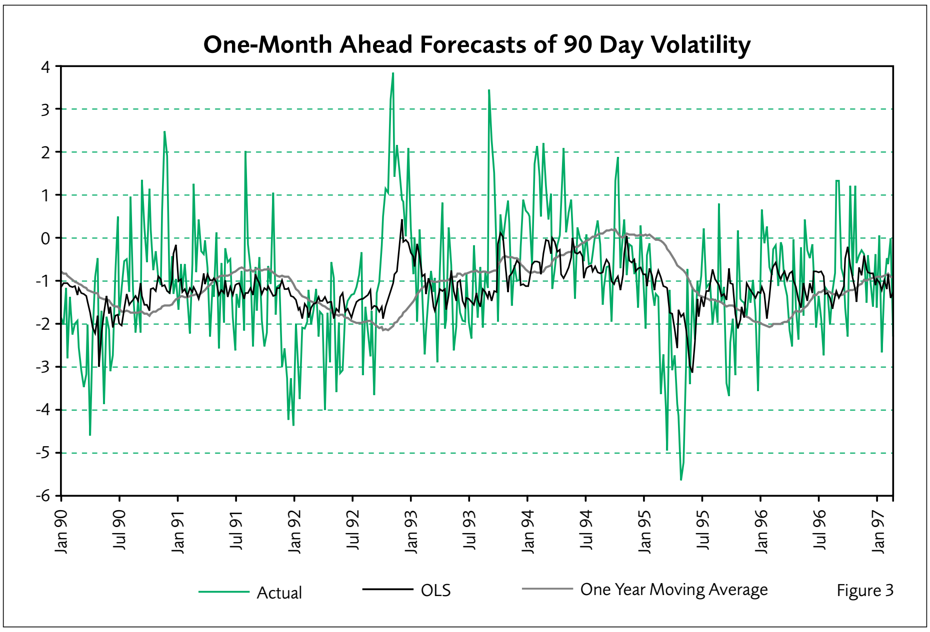

Figure 3

6. SUMMARY

Option pricing depends crucially on good forecasts of volatility. There is no shortage of alternative models that have been suggested in the academic literature or used in the markets. The most flexible models will always fit the data better within sample, but it is not clear which work best when forecasting out of sample. This paper compared eight forecasting methods over a variety of forecasting horizons, up to one year out of sample.

In short, the preferred method for forecasting *volatility* is simple OLS. This implies that a GARCH model (or one of its cousins) would be the best method for forecasting *returns.* But this conclusion comes with a proviso. If there is a clear structural break in the policy framework, such as giving up on monetary targeting or moving to exchange rate targeting, then more emphasis should be placed on moving average models because they place less weight on historical information.

REFERENCES

Bachelier, Louis (1990), 'Thèorie de la spéculation', *Ann. Sci. École Norm. Sup. 17,* 21-86.

Black, F. and M. Scholes (1973), 'The Pricing of Options and Corporate Liabilities', *Journal of Political Economy 81,* 637-654.

Brailsford, T. J. and R. W. Faff (1996), 'An Evaluation of Volatility Forecasting Techniques', *Journal of Banking and Finance,* 20, 419-438.

Hull, J. and A. White (1987), 'The Pricing of Options on Assets with Stochastic Volatilities', *Journal of Finance,* 42, 281-300.

Merton, R. C. (1976). 'Option Pricing when Underlying Stock Returns are Discontinuous', *Journal of Financial Economics,* 3, 125-144.

Table 1a
Forecast Comparison: 90 Day Rate

	1 month ahead	3 months ahead	6 months ahead	12 months ahead
Simple Methods				
No change	1.00	1.00	1.00	1.00
Historical Average	1.24	0.98	1.18	0.76
3-month average	0.75	0.71	1.02	0.71
1-year average	0.86	0.73	0.92	0.58
Regression Methods				
OLS	0.71	0.73	1.00	0.66
LAD	0.69	0.69	0.93	0.62
VAR	0.76	0.76	1.02	0.67
Changing coefficients	0.77	0.97	..	..
MSE of 'No Change' model	**2.53**	**3.33**	**2.90**	**4.80**

Table 1b
Forecast Comparison: 10 Year Rate

	1 month ahead	3 months ahead	6 months ahead	12 months ahead
Simple Methods				
No change	1.00	1.00	1.00	1.00
Historical Average	0.68	0.60	0.52	0.51
3-month average	0.45	0.67	0.63	0.69
1-year average	0.47	0.61	0.58	0.57
Regression Methods				
OLS	0.46	0.58	0.52	0.52
LAD	0.47	0.61	0.54	0.53
VAR	0.47	0.58	0.52	0.52
Changing coefficients	0.47	0.64	0.61	..
MSE of 'No Change' model	**1.86**	**2.08**	**2.37**	**2.29**

Table 1c
Forecast Comparison: USD

	1 month ahead	3 months ahead	6 months ahead	12 months ahead
Simple Methods				
No change	1.00	1.00	1.00	1.00
Historical Average	0.99	0.86	0.95	0.98
3-month average	0.67	0.65	1.11	0.60
1-year average	0.64	0.56	0.74	0.61
Regression Methods				
OLS	0.69	0.73	0.88	0.87
LAD	0.71	0.71	0.84	0.80
VAR	0.69	0.69	0.86	0.86
Changing coefficients	0.71	0.64	0.79	0.61
MSE of 'No Change' model	**1.63**	**1.94**	**1.09**	**1.89**

Table 1d
Forecast Comparison: TWI

	1 month ahead	3 months ahead	6 months ahead	12 months ahead
Simple Methods				
No change	1.00	1.00	1.00	1.00
Historical Average	0.89	0.97	0.93	0.99
3-month average	0.57	0.62	0.56	0.57
1-year average	0.53	0.57	0.53	0.53
Regression Methods				
OLS	0.66	0.81	0.80	0.84
LAD	0.72	0.89	0.88	0.92
VAR	0.66	0.73	0.76	0.82
Changing coefficients	0.75	0.82	0.81	0.79
MSE of 'No Change' model	**1.58**	**1.50**	**1.58**	**1.61**

Table 2a
Forecast Comparison: 90 Day Rate

	1 month ahead	3 months ahead	6 months ahead	12 months ahead
Simple Methods				
No change	1.00	1.00	1.00	1.00
Historical Average	0.86	0.64	0.53	0.45
3-month average	0.76	0.72	0.75	0.70
1-year average	0.88	0.72	0.64	0.53
Regression Methods				
OLS	0.69	0.61	0.53	0.44
LAD	0.69	0.62	0.54	0.44
VAR	0.77	0.61	0.52	..
Changing coefficients	0.70	0.68	0.69	0.62
MSE of 'No Change' model	2.69	3.64	4.25	4.97

Table 2b
Forecast Comparison: 10 Year Rate

	1 month ahead	3 months ahead	6 months ahead	12 months ahead
Simple Methods				
No change	1.00	1.00	1.00	1.00
Historical Average	1.04	0.93	0.87	0.76
3-month average	0.64	0.70	0.66	0.66
1-year average	0.70	0.66	0.63	0.52
Regression Methods				
OLS	0.65	0.63	0.54	0.47
LAD	0.71	0.69	0.58	0.50
VAR	0.73	0.63	0.56	0.48
Changing coefficients	0.63	0.63	0.58	0.53
MSE of 'No Change' model	1.64	1.89	2.08	2.36

Table 2c
Forecast Comparison: USD

	1 month ahead	3 months ahead	6 months ahead	12 months ahead
Simple Methods				
No change	1.00	1.00	1.00	1.00
Historical Average	1.50	1.36	1.31	1.56
3-month average	0.62	0.65	0.62	0.59
1-year average	0.61	0.55	0.48	0.54
Regression Methods				
OLS	0.57	0.53	0.47	0.52
LAD	0.61	0.59	0.52	0.62
VAR	0.65	0.57	0.50	0.57
Changing coefficients	0.69	0.74	0.73	0.66
MSE of 'No Change' model	1.53	1.70	1.79	1.58

Table 2d
Forecast Comparison: TWI

	1 month ahead	3 months ahead	6 months ahead	12 months ahead
Simple Methods				
No change	1.00	1.00	1.00	1.00
Historical Average	1.21	1.24	1.18	1.35
3-month average	0.56	0.63	0.58	0.58
1-year average	0.54	0.55	0.51	0.52
Regression Methods				
OLS	0.53	0.55	0.52	0.54
LAD	0.55	0.57	0.53	0.56
VAR	0.59	0.56	0.54	0.57
Changing coefficients	0.63	0.67	0.66	0.61
MSE of 'No Change' model	1.44	1.41	1.48	1.42

The National Bank of New Zealand Limited has provided financial services to New Zealand individuals and businesses since 1873. As a registered trading bank, it has a wide network of branches throughout the country, with corporate and commercial offices in Auckland, Christchurch and Wellington, as well as representatives in Tokyo, Hong Kong and New York.

The National Bank is a wholly owned subsidiary of Lloyds TSB Plc. which has more than 2,700 offices and branches in 40 countries around the world.

In 1987 the Bank merged its operations with those of its subsidiary, South Pacific Merchant Finance Limited, one of New Zealand's leading merchant banks.

The acquisition of the Rural Bank in 1992 established the Bank as the leading financier to the New Zealand rural sector.

As well as the full retail services offered by its branch network, the Bank provides a comprehensive range of New Zealand dollar related treasury products, foreign exchange dealing and corporate funding packages. Through its membership of the Lloyds TSB Group, the Bank is also ideally placed to offer a complete range of trade finance services both domestically and internationally, specialising in the provision of New Zealand domiciled foreign currency accounts.

FRP No. 1 July-94 David Rae
The Exchange Rate and Inflation: Are Trade Weights Appropriate?

FRP No. 2 July-94 Arthur Grimes
The Determinants of New Zealand Bond Yields

FRP No. 3 Feb-95 Capital Markets Grp & Economics Div
Inflation Linked Bonds: A New Zealand Guide

FRP No. 4 Nov-95 Arthur Grimes
Outlook for the NZD/AUD

FRP No. 5 May-96 Brendan O'Donovan & David Rae
New Zealand's House Prices

FRP No. 6 June-96 Adrian Orr & David Rae
Exchange Rate Behaviour Under Inflation Targets

FRP No. 7 Nov-96 Brendan O'Donovan, Adrian Orr & David Rae
A World Interest Rate Model (Wirm)

FRP No. 8 Aug-97 David Rae
New Zealand's Forward Exchange Market

FRP No. 9 Aug-97 David Rae
Forecasting Volatility