

Can implied forward mortgage rates predict future mortgage rates? – recent New Zealand experience

David Tripe* Bingru Xia[†] Leigh Roberts[‡] [§]

November 14, 2011

Abstract

Retail mortgage rate data for the last 13 years in New Zealand indicates that implied forward mortgage rates have only limited power to predict later spot mortgage rates. The low correlation of the forward rates and the future spot rates may in part arise from thin futures and forward markets in interest rates in New Zealand for anything longer than short term contracts.

While the pattern of mortgage yield curves has varied substantially over those 13 years, the accumulated or future value of a putative deposit of one dollar with a bank offering the same term rates as the mortgage rates shows relatively little variation over this period. In the wake of the uncertainties following the global financial crisis, the

*Massey University, Palmerston North, New Zealand

[†]Massey University, Palmerston North, New Zealand

[‡]School of Economics and Finance, Victoria University, Wellington

[§]Thanks to Dawn Lorimer and Peter Thomson for helpful discussions. Responsibility for the contents of this paper naturally remains with the authors.

relatively stable pattern of these accumulated values probably provides the best means of prediction of New Zealand mortgage yield curves, at least in the short term.

The framework used for dealing with data in this paper could be applied to yield curves based on further families of interest rates; to exchange rates; to analyses of run-off data, as in cohort and longevity analysis; and for claims payments run-off in insurance, as well as in many other contexts.

1 Introduction

For many households, their home mortgage is one of the key ways in which they interact with the financial system. New Zealand households are no exception in this regard, with home mortgage loans comprising more than 90% of household debt. They are also very important to New Zealand banks, with more than 50% of credit extended being accounted for by home lending (and more than 60% of credit being to households).¹

The products offered by banks to satisfy this demand for home mortgage lending are thus of particular importance. When the New Zealand banking system was deregulated in the mid-1980s and many of the restrictions previously applied to bank loan portfolios were removed, banks took advantage of the opportunity to expand their mortgage lending business. As the original (1988) Basel capital accord came into effect during the early 1990s, banks also noted the benefits of lower capital requirements applying to such lending, which provided them with further encouragement to concentrate on this area.

This process was however not without its challenges. During the early years of the deregulated environment, the New Zealand economy was characterised by relatively high inflation, which was in turn reflected in relatively high home mortgage interest rates. This put stress on borrowers, with floating mortgage interest rates peaking at 20.5% in June 1987, and rates not falling below 10% until May 1992 (and even then rates have risen above 10% from time to time during subsequent years). At this stage the banks were effectively limited to offering a single mortgage product, a floating rate loan which could generally have its rate changed at one months notice, and which was typically priced relative to short-term interest rates (the 90-day bank bill rate in particular). Borrowers had little choice but to accept the rates that the banks imposed on them. This stress on borrowers was compounded by the variability in interest rates, which mean that a loan that was affordable at the time it was taken out might be much less affordable at interest rates that increased significantly.

¹Statistical data have been obtained from the tables on the Reserve Bank of New Zealands web-site www.rbnz.govt.nz.

While our data prior to 1998 is incomplete, some idea of the volatility of interest rates in the 1990s may be gleaned from the early behaviour of the 1 and 3 year fixed rates shown in Figure 1 on p. 9.

By the early 1990s the banks had come to appreciate that the 90-day bank bill rate was not the pricing basis for all of their funding. In particular, they had significant liabilities that were at zero interest rates, viz. capital and some transactions account balances; some at low fixed rates, passbook savings accounts for example; and yet more funding at relatively low but not especially market sensitive rates, such as other types of savings accounts. The returns on this pool of fixed rate funding were thus particularly sensitive to changes in the general level of interest rates, and banks began to explore the opportunities to make fixed rate loans (and to invest in other fixed rate assets) to reduce the volatility in their net interest incomes.

Offering fixed rate loans thus eased problems for both the banks and their borrowers, and during the 1990s the proportions of loans at fixed rates steadily increased, while the range of fixed rate products (and the range of maturities for which loans could be fixed) steadily increased. By June 1998 (which was the date at which the Reserve Bank of New Zealand started reporting and publishing relevant statistics), 47.3% of loans by number, and 61.6% of loans by value were at fixed rates. Fixed rate loan exposure to this extent was well beyond what would be covered by hedges from the natural structure of their liabilities, so banks were by this time making significant use of swaps to hedge their fixed rate loan portfolios. This meant, in turn, that they generally sought to impose a charge on borrowers wanting to repay their fixed rate loans in advance of their maturity, to cover the costs of unwinding the swaps.

These factors all combined to contribute to a complex decision for borrowers: should they choose a fixed rate or should they float; and if they chose a fixed rate, what was the best term for them to choose? A common approach to answering this question has been simply to choose whatever appeared cheapest – in other words, to go for the lowest headline rate. Because of bank pricing choices, floating rates were often higher than short-term fixed rates, with some banks foregoing some margin for the opportunity to retain

borrowers' business for the duration of the fixed rate period. There has also been a preponderance of negatively-sloping yield curves (off the back of high short-term rates), which has made longer term rates cheaper.

None of this was helped by limitations in the financial education of both borrowers and some of the bank staff who might be called on to advise them. There was a general lack of understanding as to why interest rates at different maturities should differ, and what should be interpreted from interest rate differences.

Our goal ultimately is to explore the relationships between the interest rates at different fixed maturities, to establish whether there might be any advantage to borrowers from preferring specific maturities. Within the limited scope of this paper, the immediate aim is however more modest. Rather than attempting to compare multitudinous possible strategies of revoking one mortgage contract to commence another, including the comparison of fixed and floating rates on the one hand and fixed rates of varying length on the other, we restrict ourselves to considering the extent to which implied forward mortgage rates can predict future spot mortgage rates.

For anyone with the relevant education in Economics or Finance, the expected outcome should be straightforward: any particular fixed rate should be able to be decomposed into a spot fixed rate for a shorter maturity and one or more forward rates to cover the balance of the fixed rate periods.

Given sufficiently liquid and deep interest rate markets, one expects that the forward mortgage rate at one time would largely determine the spot rate later on; or more precisely, one expects the forward rate to conform to the market's expectations as to the future spot rate. Indeed the futures price of a good is the expected future spot price under the risk neutral measure (see *i.a.* Hull (2000, p. 511)); and a popular approach for fixed income theory in finance, which subsumes many of the common short interest rate models, is the Heath Jarrow Morton model, in which the forward interest rates adapt over time to give future spot interest rates, again in a risk neutral framework (e.g., Baxter Rennie (1996)).

A risk neutral probability measure is intended to price securities in a highly liquid market in which there is no arbitrage possible. The latter condition is clearly not valid for mortgages in New Zealand: in particular there are no conventional forward contracts available for mortgages in New Zealand, and certainly no futures exchange for them. This point might assume less importance if the banks in New Zealand were using transparent and deep swap markets for sourcing their funds, but global interest rate markets have been unsettled since the global financial crisis (GFC); and nor are they likely to settle down quickly now given the recent ructions in the Eurozone. A final caution could be that there is not necessarily a close connection between the risk neutral measure Q and the real probability measure P , whatever the latter means in the wake of the GFC. The standard means of changing between the P and Q measures is to use the basic Merton model for default rates (McNeil, Frey Embrechts 2005, p. 402), which model assumes relative stability in the market place.

While there is certainly literature on the extent to which forward rates are able to forecast future spot rates, not much has been written on the analysis of mortgage rates in New Zealand, presumably because only relatively recently has the data been gathered in an accessible form. One month ahead forward rates for 3 month interest rates and exchange rates for New Zealand are investigated in Ha Reddell (1998), but the interest rates at which they look are those available to overseas investors, whence the connection with exchange rates. While concluding that markets have forward rates that are expected values of spot rates, they discuss reasons for which those expected values and future spot rates could be consistently out of line. They find little or no predictive power of forward interest or exchange rates for future spot rates, although they give the caveat that they are looking at the short term, and such unbiasedness may become more apparent over the longer term. Also given in Ha Reddell (1998) are references to further papers exploring the unbiasedness of forward rates for future spot rates, some in the New Zealand context; but the results are mixed.

Liu (2001) applies a non-linear error correction mechanism to model mortgage rates' deviations from long run equilibria, but concludes that this offers no forecasting improvement over simple regressions using government bond rates as regressors. In any case, recent financial conditions, especially the

GFC, have hardly been conducive to successful use of models assuming mean reversion to equilibrium.

The seminal framework for testing for the unbiasedness of exchange rate forecasts in Meese Rogoff (1983) is updated by Evans Lyons (2005), who claim to have found evidence that forward exchange rates do better than a random walk in forecasting spot rates in the short term. Their methodology relies on microeconomic variables, and the proportion of variance explained by the forward rates does not seem high (the R^2 in the regressions); and they maintain that the drivers of the exchange rates are macroeconomic in nature, not all of it in the public domain. Although written in an exchange rate setting, their paper is a clear indication of the difficulties inherent in modelling yield curves in general, let alone considering the unbiasedness of forward rate forecasts for future spot rates.

Poole (2005) discusses the behaviour of the yield curve towards the end of the period in which the Federal Reserve lowered its target rate by 25 basis points every 3 months over some 3 years (from 2002-2005 roughly). He emphasises the role often played by publicly announced policy decisions, and the Fed's decisions on the target rate in particular, in affecting market expectations, thereby altering the shape of the yield curve overnight. This ties in with the finding in this paper that the mortgage rates seem influenced mainly by contemporaneous events in the financial and economic environment, rather than from forward rates a year or two earlier.

While the official cash rate (OCR) and the term interest rates offered by New Zealand banks, *inter alia*, are expected to be relevant to the determination of mortgage rates, we restrict our analysis to the mortgage rates themselves, together with quantities defined directly from the (mortgage) yield curve, viz. forward interest rates and accumulated or future values. We stand thereby in some danger of missing the boat in a statistical or perhaps economic sense: it may be that causality runs from the OCR to the floating rate, and we should be modelling the relationship between floating and fixed rates, as well as the political and economic impacts on the OCR. But one has to start somewhere; and looking at graphs of the data and running elementary statistical models in a preliminary attempt to understand what is going on is a reasonable thing

to do. Indeed, to use covariates and complicated models may have its place; but our idea is instead to look closely at the data itself to see how far we can get. Even from the cursory literature review above, it is clear that the choice of covariates and modelling techniques is vast, and subject to much data-mining.

The next section of the paper outlines the data that we use for our study, and the methods employed to calculate the forward rates and future or accumulated values. Taking our cue from Figure 5 on p. 13, the next section investigates the data from three points of view: the first is the comparison of the forward rates and the future spot rates, from which we infer that there is little predictive power of the former for the latter; the second is fitting splines to look at the behaviour of the yield curves over the three principal subperiods; and the third is a look at the data through the lens of wavelets fit to the data. Section 4 discusses the practical implications of our results, identifies some further research issues, and concludes.

2 Data

2.1 The overall picture

The mortgage rate data was downloaded from the Reserve Bank of New Zealand web site, and comprised the floating mortgage rate and a series of quoted fixed mortgage rates for new loans, for each month from June 1998 to May 2011 inclusive.

The particular loan durations available to us were six months; 1 year; 2 years, etc., up to 5 years. The data is graphed in Figure 1. As indicated there, two of the time series (for 1 year and 3 year mortgage rates) are in fact available prior to 1998; but our analysis focuses on the period from 1998 to the present.

The data is most conveniently presented as in Table 1, in which the spot rates




Figure 1: Extended time series of mortgage rates

of the mortgage yield curve are the mortgage rates, labelled as $mr(1/2)$, $mr(1)$, \dots , $mr(5)$.

2.2 Description of the data

Figure 2 on p. 11 gives a less cluttered view of the mortgage rate movements over time, being restricted to the floating rate, together with the one and three year spot rates.

While the pattern of the data in Figure 2 is hardly transparent, one can tentatively identify some features of the data. Prior to mid 2004, the 3 year rates generally exceeded the 1 year rates. From approximately mid 2004 to mid 2008, the 1 year rate is generally the higher, and the floating rate is well above any of the spot mortgage rates, to judge from Figure 1. Both graphs indicate a steeply rising yield curve in the aftermath of the GFC.

	Floating	6 months	1 year	2 years	3 years	4 years	5 years
	mr(ft)	mr(1/2)	mr(1)	mr(2)	mr(3)	mr(4)	mr(5)
Jun 1998	11.17	10.00	9.86	9.82	9.58	9.43	9.22
Jul 1998	10.01	8.85	8.91	9.01	8.96	8.97	8.80
Aug 1998	8.93	7.77	7.91	8.28	8.45	8.56	8.53
Sep 1998	8.09	6.82	7.22	7.85	8.08	8.17	8.18
...							
Feb 2011	6.50	6.30	6.47	6.69	7.18	7.53	7.78
Mar 2011	6.14	6.08	6.08	6.54	7.05	7.46	7.73
Apr 2011	6.10	5.93	6.04	6.54	7.05	7.46	7.73
May 2011	6.07	5.94	6.02	6.52	7.05	7.46	7.73

Table 1: Floating and fixed mortgage rates (% p.a.)

These impressions as to the shape of the yield curve are borne out by splines fitted to the yield curves. Figure 15 on p. 29 in the Appendix presents the spline fitted by ordinary least squares to the yield curve over the entire 13 years; while the following Figures 16, 17 and 18 indicate the average yield curve over the subperiods 1998 to mid 2004; mid 2004 to mid 2008; and mid 2008 until 2011, the post GFC period, respectively. The middle period indeed exhibited a slightly falling yield curve; and the final period a very strong upwards trend. This is discussed further in §3.2.

Figure 3 indicates that the floating mortgage rate is strongly influenced by the OCR over the period in question. The focus of this paper however lies with the fixed rates.

2.3 Alternative tabular and graphical views of the data

While clearly less convenient typographically than Table 1, in many ways it is preferable to set out the data as in Table 2, in which the horizontal time axis is consistent from one row to the next; otherwise stated, the columns in Table 2 correspond to the same calendar month of maturity, and the rows

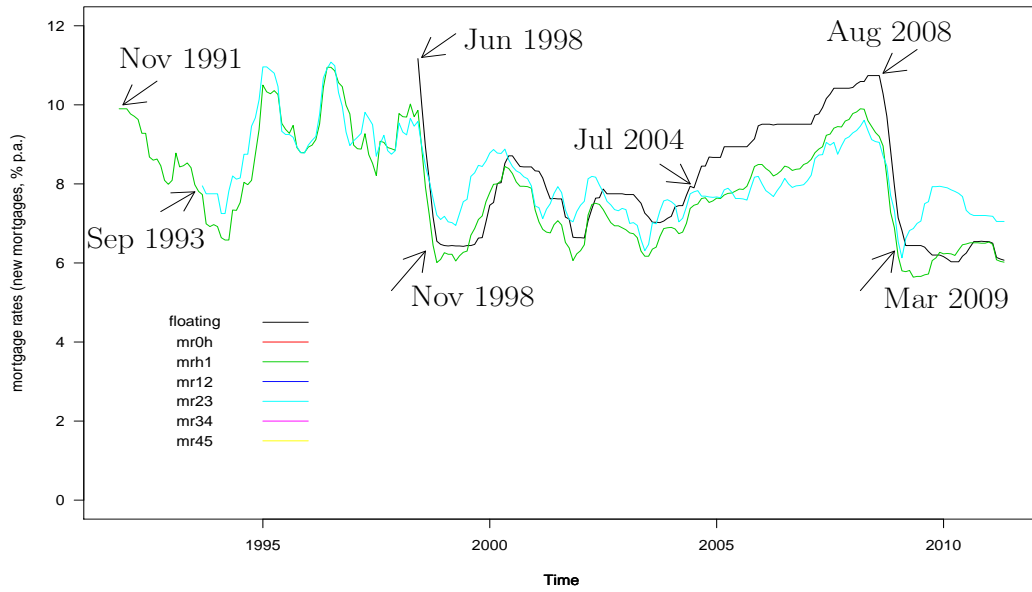


Figure 2: Extended time series of 1 and 3 year mortgage rates




Figure 3: floating rate for new loans compared with the OCR




Figure 4: Fixed rates data availability

and columns are measured in consistent time units. Zeroes in the body of the table indicate empty data cells.

The layout and generation of the data in this fashion can be represented graphically in Figure 4, in which the various colours along diagonals correspond to the mortgage rates of different durations.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	...	23	24	25	26
Jun 1998	0	0	0	0	0	10.00	0	0	0	0	0	9.86	0	0	0	0	0	9.82	0	0	0
Jul 1998		0	0	0	0	0	8.85	0	0	0	0	0	8.91	0	0	0	0	0	9.01	0	0
Aug 1998			0	0	0	0	0	7.77	0	0	0	0	0	7.91	0	0	0	0	0	8.28	0

Table 2: Fixed mortgage rates (% p.a.)

The advantage of the layout in Table 2 is that it facilitates identification and modelling of the data generating process. Possible factors impacting on a cell in Table 2 are the calendar month of maturity of mortgage (from the cell vertically above); the values of the previous cohort of mortgages at similar duration (from the top left cell): and a factor arising from belonging

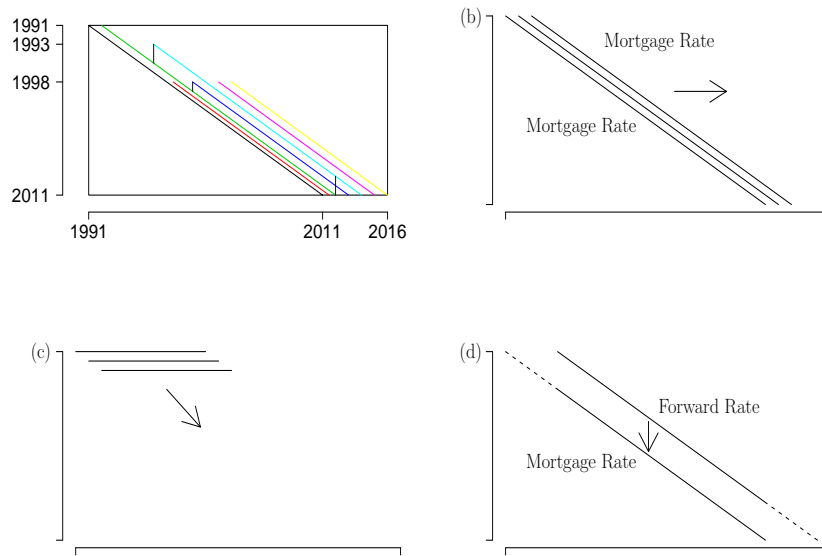


Figure 5: Data availability and methods of analysis

to a particular cohort (from the cell to the immediate left), possibly from a contemporaneous event such as 9/11 or a sudden fall in the stock market (acting horizontally).

The transparency of setting out the data as in Table 2 is enhanced when the data is concentrated for small values of ‘time’ or ‘duration’ along the horizontal axis. This will often be the case. Forward prices in many markets, for example, are typically available for maturities separated by one month for the first 6 months, and less frequently thereafter. Given consistent units of time on the horizontal and vertical axes, the data on the right of the table would be relatively sparse. A similar point can be made for yield curves, since the estimated yields are often highly uncertain as one moves to the longer durations.

For our data, an influence acting vertically in Table 2 would perhaps be less likely: cells containing data are well separated vertically. This point assumes some importance, for it is the vertical influence that we are trying to pin down in this paper.

The availability of fixed mortgage rates data is summarised in Figure 4 on p. 12. The 1 year rates commence in late 1991, and the 3 year rates in 1993; the other rates, as listed in Table 1 on p. 10, commence in 1998. The vertical lines indicate that when considering causality, the time series involved need to be truncated. In order to test for whether the 3 year rate causes the 1 year rate, for instance, those parts of the vectors between the vertical lines are the vectors of maximum length that one could use. The other example shown in this Figure is the truncation of the 1 year and 2 year rates for testing whether the 2 year rates cause the 1 year rates a year later. The truncation of vectors to this end is discussed in greater detail below.

The impact of factors on the data from ‘different directions’ in Table 2 can be summarised in the Figure 5 on p. 13, the first subgraph of which is a reproduction of Figure 4.

Figure 5(b) indicates schematically a time series analysis from 1998 to 2011, with 6 vectors to be analysed, corresponding to $mr(1/2), mr(1), \dots, mr(5)$. These vectors are expected to be more closely correlated the stronger the contemporaneous influence on a yield curve. Unusual, or unusually strong, influences on the yield curve could for instance arise from the GFC; or a particular event such as 9/11; or a sudden drop in stock values, and such events impact horizontally: they should be manifest at least partly through enhanced correlations between these 6 vectors.

Figure 5(c), on the other hand, looks at each yield curve separately, attempting to pick up the extent to which a particular mortgage rate impacts on that same mortgage rate for the next month. Given a certain stickiness in mortgage rates over time, arguably present in say Figure 8 on p. 18, the correlations between the yield curve vectors could be strong.

The data from the perspective of Figure 5(c) is indicated in the 3D graph in Figure 6, from which the steeply rising yield curve since the GFC is apparent. While the figure is hardly transparent, indicated clearly is the high variability of yield curve behaviour over the 13 year data period. The analogous graph for the one year forward rates is shown in Figure 7 on p. 17. There is some indication that the forward rates are more volatile than the spot rates, as

expected; but the graphs are too crowded to be much more informative.

In fact we shall analyse these yield curves using accumulated future values (see §2.4 on p. 16), separating the 13 years of data into parts determined by the arrows in Figure 1 on p. 9. If one were wanting to predict future mortgage rates on the basis of past data, the best approach is probably to use the splines fit to the future values in order to obtain a window of future values for the predicted yield curve.

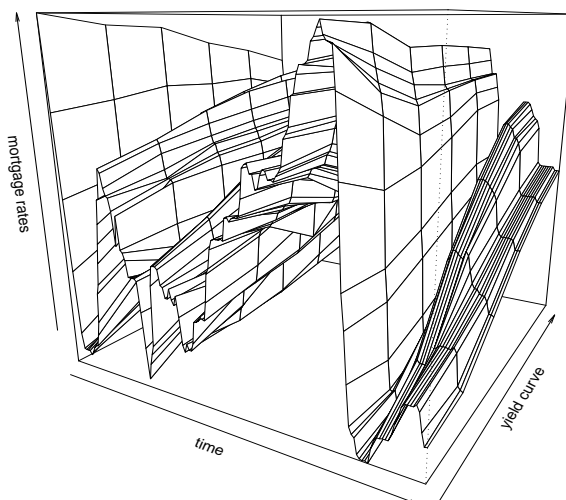


Figure 6: Progression of yield curve over time, based on Fig 5(c)

Finally, in Figure 5(d), we examine the possibility of forward interest rates determining the future spot rate. Correlations here will be between the vectors akin to those in Figure 5(b), but omitting the earliest values of the spot rate vector and the later values of the forward rate vector: when using the forward rate vector $fr(1, 2)$ to predict the one year mortgage rate $mr(1)$, for example, one omits the first 12 values of $mr(1)$, and the final 12 values of $fr(1, 2)$.

2.4 Forward mortgage rates and future values

From the spot rates for a given yield curve are calculated forward rates on the one hand; and future values or accumulated values on the other. The forward rate $fr(2, 3)$, for instance, is calculated as

$$e^{3 \times mr(3)} = e^{2 \times mr(2)} \times e^{fr(2,3)}$$

or, for the particular month September 1998 say,

$$e^{3 \times .0808} = e^{2 \times .0785} \times e^{fr(2,3)}$$

from Table 2. The mortgage rates in Table 1 are in fact compounded semi-annually; but given the uncertainties in the overall situation to be modelled, there seems little harm in making the simplifying assumption that they are continuously compounded.

Similarly, the forward rate $fr(2, 4)$ is calculated as

$$e^{4 \times mr(4)} = e^{2 \times mr(2)} \times e^{2 \times fr(2,4)}$$

so that all rates, whether spot rates $mr(n)$ or forward rates $fr(n_1, n_2)$, are rates p.a. One expects the forward rates to be more volatile than the spot rates (e.g., Hull (2000, p. 94)), and to some extent this seems to be borne out here. A comparison of Figure 8 on p. 18 with Figure 9 on p. 19 indicates a greater volatility of the one year ahead forward rate from 2004 and 2008, although the evidence at other times is less clear. We shall in the main concentrate on one year ahead forward rates, and sometimes abbreviate $fr(n, n + 1)$ to $fr(n)$.

For the future or accumulated values, corresponding to the accumulated value of a dollar deposited in a bank paying term interest rates equal to the mortgage rates, one defines

$$fv(t) = e^{t \times mr(t)}$$

In addition to Table 1, we have then analogous tables of data containing forward interest rates $fr(n_1, n_2)$ and future values $fv(n_1)$, all for each month over the same 13 year period.

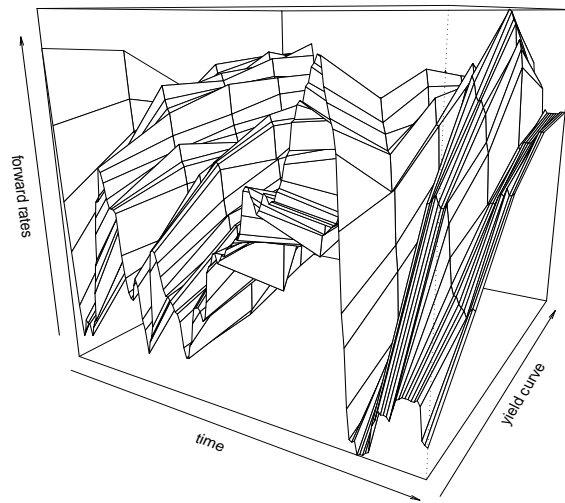


Figure 7: Progression of forward rates over time, based on Fig 5(c)

3 Analysis

Proceeding in the inverse order to that given in Fig 5, we first examine the extent to which forward rates forecast spot rates, based on Fig 5(d); we graph the yield curves for the principal subdivisions of the data in the following section, basing our analysis on Fig 5(c); and finally we look at the data from a time series perspective, with Fig 5(b) in mind. We do not fit conventional time series models in the latter section, because of the irregularity and obvious non-stationarity of the data; but we do investigate the spectra of the various components of the time series, showing the importance of the low frequencies of between 1 and 2 year periods over all subperiods, which is not evident from a casual glance at the data.

3.1 Do forward rates predict spot rates?

Basing our investigation in this section on Fig 5(d), we illustrate by means of an example. We are for instance investigating the extent to which the




Figure 8: Mortgage rates

forward rate $fr(2,3)$ for September 1998 impacts on the 1 year spot rate $mr(1)$ for September 2000; and the way in which the forward rate $fr(2,4)$ for September 1998 impacts on the 2 year spot rate $mr(2)$ for September 2000; etc. Restricting ourselves to the one year ahead forward rates, we consider for instance how the $fr(1,2)$, $fr(2,3)$, $fr(3,4)$ and $fr(4,5)$ rates in September 1998 impact on the one year spot rates $mr(1)$ in the months September 1999, 2000, 2001 and 2002 respectively.

Figure 8 on p. 18 plots the spot rates $mr(\cdot)$, reproduced from Figure 1 on p. 9; while the 1-year forward rates $fr(n, n+1)$ for $n = 1, 2, 3$ and 4, along with $mr(1)$, are plotted in Figure 9. The forward rates are more volatile than the spot rates, as expected.

Of more use than Figure 9 is a graph translating the forward rates so that they are superimposed on the mortgage rates that they purport to be forecasting, which leads to Figure 10 on p. 19. This graph is in its turn sufficiently cluttered that we turn to Figure 11, which highlights the graph of $fr(1,2)$ translated forward by one year; and that of $fr(3,4)$ translated forward by




Figure 9: Implied (one year) forward rates




Figure 10: Implied (one year) forward rates: translated




Figure 11: Implied (one year) forward rates: translated

three years.

Even a casual glance at Figure 11 suffices to show that the predictive power of the 1 year forward rates $fr(1, 2)$ is not particularly good.

The main features of the $mr(1)$ graph are picked up by the contemporaneous forward rate, not the translated forward rate. The situation is much the same for $fr(3, 4)$ and the other 1 year forward rates: the translated forward rates have little correlation with the spot rates, whereas the contemporaneous forward rates are highly correlated with the spot rates.

To revert to our example, $fr(1, 2)$, $fr(2, 3)$, $fr(3, 4)$ and $fr(4, 5)$ rates in September 1998 are highly correlated with $mr(1)$ in September 1998, and seem to have little predictive power for $mr(1)$ in the months September 1999, 2000, 2001 and 2002 respectively.

Our tentative conclusion about the forward rates failing to predict the spot

rates is borne out by looking at the correlations. The six mortgage rate time series are highly correlated, as seen in Table 3, with the correlations falling away roughly linearly as we move across the yield curve.

The correlation between the spot rates $mr(1)$ and the translated forward rates $fr(1, 2)$, as illustrated in Figures 5(d) and 10, is -0.07 . In like vein, the correlation between the translated $fr(2, 3)$ rates and the $mr(1)$ rates two years ahead is -0.42 . This value is easily significantly different from zero using conventional tests and significance levels, but it has the wrong sign. In any case, the conventional tests are based on underlying normality and stationarity, so that one should not necessarily read too much into such ‘significance’.

	mr(ft)	mr(1/2)	mr(1)	mr(2)	mr(3)	mr(4)	mr(5)
mr(ft)	1.00	0.98	0.95	0.83	0.70	0.53	0.37
mr(1/2)	0.98	1.00	0.99	0.91	0.80	0.64	0.49
mr(1)	0.95	0.99	1.00	0.95	0.85	0.71	0.56
mr(2)	0.83	0.91	0.95	1.00	0.96	0.87	0.76
mr(3)	0.70	0.80	0.85	0.96	1.00	0.97	0.90
mr(4)	0.53	0.64	0.71	0.87	0.97	1.00	0.98
mr(5)	0.37	0.49	0.56	0.76	0.90	0.98	1.00

Table 3: Correlations between the spot rates

3.2 Yield curve modelling by splines

Ideally one would consider the ‘shape’ of the yield curve by looking at the level, slope and curvature. Because of the paucity of data along each yield curve, however, we have simply fitted quadratic splines separately to the spot rates, the accumulated or future values, and to the (one year) forward rates; and we do this for the whole 13 years of data, as well as for each of the 3 subperiods identified in Figure 1 on p. 9, and as described in §2.2. The single knot is taken at the 2 year mark; and the fit is by OLS for the spot and forward rates. For fitting splines to the future values, weighted least squares are used, with weights inversely proportional to distance from the origin:

thus $fv(1/2)$ was given initial weight 2, $fv(1)$ was given initial weight 1, etc; following which the weights were reset to sum to unity. All splines fitted have 99% confidence intervals surrounding the individual points; regarding the regression residuals at those points as independent (and assuming that the fitted values are accurate, since the uncertainty in parameter estimates is not reflected in the bounds shown) means that the confidence intervals shown correspond roughly to the 95% confidence intervals for the whole curve shown.

The spline fitted to the spot rates over the whole period is shown in Figure 15 on p 29, while those fitted over the subperiods are shown in Figures 16, 17 and 18 (graphs of spline fits are relegated to the Appendix). The impressions as the shape of the yield curve over the three subperiods inferred in §2.2 are borne out: the middle period, from mid 2004 to end 2008, exhibits a slightly falling yield curve, albeit with wide confidence intervals; both the initial period, from mid 1998 to mid 2004, and the period since the GFC, exhibit strongly rising yield curves, in accord with the graphs in Figures 1 and 2 on pp. 9 and 11; and as far as one can tell at least, with Figure 6 on p. 15.

Although not reproduced here, one obtains similar pictures from the splines fitted to the forward rates, with the forward rates exhibiting similar slope to the spot rates, albeit with somewhat exacerbated trends and wider confidence intervals, as expected since forward rates should be more volatile than spot rates.

The spline fits to the future values, on the other hand, are models of propriety compared with those fitted to the spot and forward rates. The confidence bounds are narrow; and, despite the variability of the individual future values $fv(t)$, the splines over the whole period and the subperiods are close to each other. While one may well be less than enthusiastic about forecasting yield curves in general, and especially from rates that look anything like those in Figure 6 on p. 15, the forecasting of future values looks far more promising than forecasting spot or forward rates per se.

3.3 Wavelet analysis

As for the time series analysis, based on Figure 5(b), we fit the LA(8) maximal overlap wavelets to the vector of $mr(1)$ values, using the wavelets package in R (Aldrich 2008). The depth of fit was 4 levels. The data was extended by circularity for the fit at the edges, which looks undesirable given the wide disparity of beginning and end values of the data, as seen at the bottom of Figure 14 on p. 26. But using the reflected data (reflected at the latest time point, not the earliest) led to similar graphs of details and smooths, and similar numbers in Table 4.

As is typical of undecimated wavelet analysis, the wavelet coefficients can be seen roughly to follow the details graph at the same level (Nason 2008): compare Figures 13 and 14. The smooths graph is also given for completeness, in Figure 12. The highest frequency (level 1) in theory picks up frequencies with periods between 2 months (for the the Nyquist frequency) and 4 months, although the leakage in the finite width wavelet filters means that this fails to be exactly so in practice. In like vein, the period corresponding to level 4, viz. the lowest frequency, lies between 16 and 32 months: 2 years say. This is borne out by the top graph in Figure 13, in which the rough cycles have a period of about two years. This is some indication that fitting wavelets to a greater depth would not produce much more information; and in any case there is not so much data, and what there is clearly has three distinct subperiods of different characterists, so that one would hesitate to fit to a greater depth.

The most striking aspect of Figure 13 is the strength of the low frequency signals. Note before proceeding that the first subperiod (mid 1998 to mid 2004) roughly corresponds to the time indices 1-75 in Figures 13 and 14; the second subperiod roughly to 75-130; and the final subperiod, viz. post GFC, to 130-156.

Apart from the beginning and end of the data, the high frequency detail is non-zero only when the $mr(1)$ vector changes value very sharply, but the impact is slight. There is a hint of this at approximately time indices 30 and

40, and again at the sharp drop at about time 130. But these sharp changes impact more strongly at level 2, roughly in the 4 to 8 month period range. Given the circularity assumption, the sharp change at the beginning and end of the data is manifest in large non-zero values of the high frequency detail, and the slowly varying behaviour of the data just after inception and just before the end ensures that the impact of the end points stays strong at all frequencies. But the overall impact on a casual observer is that despite the jagged appearance of the original data, the low frequencies dominate, and are centred on periods of about 1 year and 2 years.

These suppositions are borne out by breaking down the spectrum into the 4 levels. Table 4 shows the composition of the spectrum, as measured by the power or sums of squares of the wavelet coefficients at the different levels. The figures shown are proportions of the total power attributable to that level, taken from a single wavelet fit to the whole period. The impact of the high frequency is small, even in the final subperiod when the large negative value at the end of the data contributes substantially to the proportion of the power borne by the highest frequency detail.

proportional spectrum (%)			frequency			
period		time index	high	...	low	
Jun 1998	May 2011	1-156	7	13	22	57
Jun 1998	Jul 2004	1-73	9	17	25	48
Aug 2004	Mar 2009	74-129	1	3	18	77
Apr 2009	May 2011	130-156	11	16	19	54

Table 4: Importance of high and low frequencies over the subperiods

4 Conclusion

Given sufficiently deep futures and forward markets in interest rates, a marked divergence of a spot rate from an anticipatory forward rate should send a clear signal to market participants; and in the absence of such perfect markets, then such behaviour could still be observed if banks used deep overseas




Figure 12: Wavelet smooths




Figure 13: Wavelet details




Figure 14: Wavelet coefficients

swap markets in sourcing their funds. In fact, however, at least for the last 13 years in New Zealand, forward mortgage rates have had little to say about future spot mortgage rates.

The dominant impact on the yield curve seems to be from simultaneous influences, with forward rates strongly correlated with current spot rates. In calmer times, it might be that market expectations, as expressed in forward rates, would be correct, at least on average. But times have hardly been normal, with the GFC being the obvious example of an event with enormous repercussions during the investigation period; and 9/11 being another. But one could argue that it is unlikely that times will ever be so settled that market expectations will drift towards equilibrium values.

A key factor in spot mortgage rates is the retail interest rates payable by the banks, in turn determined presumably by the wholesale rates that the banks face. Banks would prefer to use options to hedge against interest rate risk;

but in the absence of a liquid options market, they generally prefer to swap liabilities and assets into floating rates, and the mortgage rate yield curves can be expected to be influenced by the international swap markets.

Further research will centre on clarifying the interrelationship between mortgage rates, swap rates and banks' base rates; and on the changes wrought by the GFC on the way in which banks source their funds. The situation regarding mortgages in Australia would be of interest, especially if it were possible to get data over a longer period than in New Zealand.

This paper has been an attempt to model mortgage rates in a relatively pure fashion, in that there has been a deliberate attempt to eschew the use of covariates and statistical constructs such as mean reversion to equilibria. As noted in the introduction, this approach could be self defeating, since term rates offered by banks to depositors; the OCR; and swap rates, *inter alia*, are all expected to influence the mortgage rate yield curve. On the other hand, to broaden the analysis to incorporate covariates would not necessarily be an easy task, because it is by no means clear which covariates would be optimal, and how causality would run. In any case, covariates need to be forecast for making predictions. In our relatively rudimentary attempt to model mortgage rates we have established the 'shape' and characteristics of the data; we have established that the theoretically expected influence of forward mortgage rates on future spot rates is absent; we have concluded that the principal influence on the yield curve over those 13 years in New Zealand seems to be from contemporaneous events; and we have shown that the strongest cyclical presence in the data seems to be from cycles of period approximately 1 and 2 years, despite the jagged nature of the data hiding such features from a casual observer. To hope for more in the aftermath of the GFC could be considered optimistic.

References

Aldrich, E. (2008). The wavelets package. In R.

- Baxter, M. Rennie, A. (1996). *Financial Calculus: An Introduction to Derivative Pricing*, CUP.
- Evans, M. D. D. Lyons, R. K. (2005). Meese-Rogoff redux: micro based exchange rate forecasting, *American Economic Review* **95**(2): 405–414.
- Ha, Y. Reddell, M. (1998). What do forward interest and exchange rates tell us?, *Reserve Bank of New Zealand Bulletin* **61**(2): 129–147.
- Hull, J. (2000). *Options, futures and other derivatives*, 4th edn, Prentice Hall.
- Liu, Y. (2001). Modelling mortgage rate changes with a smooth transition error-correction model, *Working Paper 23*, Bank of Canada.
- McNeil, A. J., Frey, R. Embrechts, P. (2005). *Quantitative risk management*, Princeton University Press.
- Meese, R. Rogoff, K. (1983). Empirical exchange rate models of the seventies, *Journal of International Economics* **14**: 3–24.
- Nason, G. P. (2008). *Wavelet methods in statistics with R*, Use R!, Springer.
- Poole, W. (2005). Understanding the term structure of interest rates, *Federal Reserve Bank of St Louis Review* **87**(5): 589–595. Speech to the Money Marketeers, NY, June 14 2005.

Appendix A: spline fits to the spot rates




Figure 15: Spline fit to mortgage rates, 1-156 = Jun 1998 - May 2011




Figure 16: Spline fit to mortgage rates, 1-73 = Jun 1998 - Jun 2004




Figure 17: Spline fit to mortgage rates, 74-129 = Jul 2004 - Feb 2009




Figure 18: Spline fit to mortgage rates, 130-156 = Mar 2009 - May 2011

Appendix B: spline fits to the future values




Figure 19: Spline fit to accumulated values, 1-156 = Jun 1998 - May 2011




Figure 20: Spline fit to accumulated values, 1-73 = Jun 1998 - Jun 2004




Figure 21: Spline fit to accumulated values, 74-129 = Jul 2004 - Feb 2009




Figure 22: Spline fit to accumulated values, 130-156 = Mar 2009 - May 2011