G98/8

Testing the predictive power of
New Zealand bank bill futures rates

Leo Krippner[1]

June 1998

JEL E43, E44, C2

Abstract

The hypothesis that New Zealand 90-day bank bill futures rates are an unbiased predictor of 90-day bank bill rates is tested by applying the single-equation method of Stock and Watson (1993) to quarterly data from 1989 to 1997. The results do not reject the unbiasedness hypothesis for the one and two-quarter-ahead horizons tested. However, the estimated residuals are found to contain significant serial correlation in both cases, which suggests that some degree of information inefficiency might exist. The relative forecasting performance of futures rates against the random-walk is also investigated for weekly horizons up to 26 weeks. The results indicate that futures rates outperform the random-walk over all horizons tested, with the improvement statistically significant for all horizons greater than 1 week.

---

1   Introduction

Bank bill futures contracts provide market participants with the potential for large profits or losses, depending on movements in the underlying bank bill rate relative to the contract rate. Hence, a natural null hypothesis is that futures rates for traded contracts will be an unbiased and information efficient predictor of the bank bill rate on the settlement day of the futures contract. This implicitly assumes a competitive market with rational economic agents, and that any biases associated with transaction costs, liquidity etc. are relatively minor.

Information efficiency, in the context above, essentially means that the market takes all historical information into account when forming predictions. Previous prediction errors are part of the set of historical information, so the presence of any systematic pattern in historical prediction errors might indicate potential information inefficiency.

This paper tests the unbiasedness hypothesis for New Zealand bank bill futures data by applying the method outlined by Stock and Watson (1993) to quarterly data from 1989 to 1997. The one and two-quarters prediction horizons are tested, and the serial correlation of residuals is tested for as one indicator of any potential information inefficiency.

Examples of recent, related work in this area include Mishkin (1988), Campbell and Shiller (1991), Cutherbertson (1996) and Engsted (1996), who test the power of spreads in the term structure for predicting interest rate movements in the United States, United Kingdom, and Denmark, respectively. The results generally conclude that the slope of the term structure is useful for predicting movements in short-term interest rates, but is not always the optimal predictor.

Regardless of whether futures rates are an optimal predictor, they may still contain useful predictive power relative to other methods of forecasting. For example, Deaves (1996) finds that forward rates implied from the term structure generally provide superior forecasts of interest rates compared to professional forecasters, time-series methods, other predictors based on the term structure, and the naive random-walk model.

Hence, this paper also quantifies the relative forecasting performance of future rates against the random-walk for all horizons from 1 to 26 weeks, using the method outlined in Ashley, Granger, and Schmalensee (1980).

2 The general model

A two-variable cointegrated spot-forward system may be represented by the two equations in System (1):

$$R_t = \alpha + \beta F_{t-1} + u_t \tag{1a}$$
$$F_t = F_{t-1} + v_t \tag{1b}$$

where: a and b are the model parameters;
$\quad$ R_t is the spot interest rate;
$\quad$ F_t is the forward interest rate;
$\quad$ u_t and v_t are the model disturbances.

Stock and Watson (1993) provide a method of estimating the parameters in System (1) from a single-equation. Specifically, as outlined in Stock and Watson (1993), and Hamilton (1994), Equation (2) is estimated using feasible generalised least squares (FGLS), with the appropriate order of p to ensure that e_t is uncorrelated with v_t at all leads and lags. This provides estimates that are asymptotically equivalent to those from full information maximum likelihood estimation, and the FGLS estimates have the usual asymptotic related-normal distributions (using the FGLS standard errors[2]). This is contrary to estimates of Equation (2) without the ΔF_{t-1} terms in the summation, which may have non-standard asymptotic distributions, or estimates of Equation (2) using OLS, which may not be of minimum variance.

$$R_t = \alpha + \beta F_{t-1} + \sum_{i=-p}^{+p} \gamma_i \Delta F_{t-1-i} + \varepsilon_t \tag{2}$$

where: $\Delta F_{t-1} = v_t$;
$\quad$ γ_i are the coefficients on the leads and lags of ΔF_{t-1};
$\quad$ e_t is the model disturbance, which is uncorrelated with v_t by construction.

Testing the unbiasedness hypothesis for System (1) requires jointly testing that $\alpha=0$ & $\beta=1$. Testing for potential information inefficiency requires testing the residuals e_t for any significant serial correlation.

---

[2] Hamilton (1994) contains a useful summary of the theory behind FGLS, and the application to models with autoregressive serial correlation in the residuals (which is the Cochrane-Orcutt technique, as used in this paper). The FGLS standard errors for the estimates are the OLS standard errors adjusted using the feasible covariance matrix.

Note that the literature often stylises the risk-neutral information efficient markets hypothesis as:

$$F_{t-1} = E(R_t | \Psi_{t-1}) \tag{3}$$

where: Ψ_{t-1} is the information set available at time t-1, which may include previous prediction errors.

In this paper, the risk-neutral information efficient markets hypothesis has effectively been split into the unbiasedness hypothesis, and the information efficiency hypothesis. Both are interesting questions in their own right: respectively, is there any persistent average difference between the futures rate and the realised spot rate, and does the market make any systematic errors that indicate that the market is not always using all the information available to it?

3. Background and description of the data

Regarding terminology, a 90-day bank bill futures contract is effectively an agreement between two parties to exchange a 90-day bank bill at an agreed rate (the futures rate), on a pre-determined, standardised day in the future (the settlement day or close-out day). In practice, two-way contracts are formed indirectly through the central exchange, and cash is exchanged to settle any net contractual obligation (open position) on the close-out day. The cash amount is the product of the open position, and the difference between the price of a 90-day bank bill calculated using the futures rate and the price of a 90-day bank bill calculated using the average market 90-day bank bill rate on the close-out day (the settlement rate or close-out rate).

New Zealand bank bill futures contracts have been offered for the months of March, June, September, and December since March 1989 (and each month before then, back to January 1987). The close-out day has always been the second Wednesday of the month, and the close-out rate has always been the 90-day bank bill mid-rate recorded in the daily interbank rate-set (BKBM rates) on the close-out day (except for some initial monthly contracts, which used the commercial bill rate).

Daily traded volumes in futures contracts typically ranged from hundreds of contracts per day in early data, to thousands in later data (each contract has a nominal face value of $500,000). Trading in futures contracts with more than 100 trading days to close-out due was relatively thin, so analysis for forecasting horizons greater than 26 weeks is limited by lack of data.

The futures data used in this analysis are all recorded daily closing futures rates associated with the quarterly 90-day bank bill futures contracts settling March 1989 to September 1997; 36 contracts in total. The data are supplied by the New Zealand Futures and Options Exchange, the Sydney Futures and Options Exchange, and the Reserve Bank. The 90-day rate data (daily BKBM rates), used to determine the close-

out rate and to test the random-walk, are supplied by the Reserve Bank. All rates are converted to continuously compounding interest rates (CIRs), applying a basis of 90 days.

A trading day index spanning the entire period is used to organise all of the data, contract by contract. For any given day, futures rates are only recorded for the corresponding contract if trading in that contract actually occurred. Hence, for a given number of days to close-out, some observations are missing. In such cases, the closest observation associated with the same contract is substituted.

Quarterly and two-quarterly data is then extracted to test the unbiased and information efficiency hypothesis for these two horizons. A chart of the actual data is shown below. To test the relative performance of the futures rates predictions to the random-walk, observations are grouped according to days to close-out (DTC), and the relevant data is extracted for the horizon to be tested. Appendix 1 contains a sample of information used for the 70 days to close-out analysis. For each group of given DTC there are 36 observations, since each contract contributes one observation.

Figure 1:   Evolution of 90-day Interest Rates, and Bank Bill Futures Rates
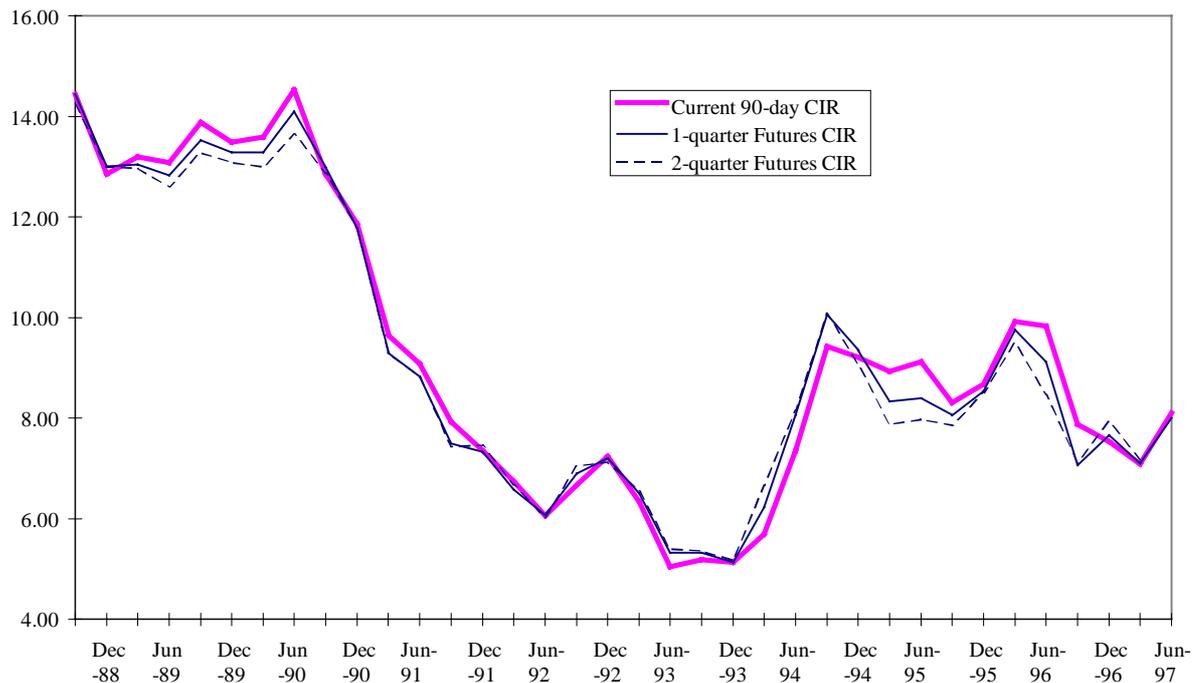
4   Model estimation and results

The valid estimation of Equation (2) requires that close-out rates be cointegrated with lagged futures rates, which is accepted by the data. Specifically, the Dickey-Fuller and augmented Dickey-Fuller results in Table 1 indicate that close-out rates and both quarter and two-quarter ahead futures rates are I(1), and the difference between close-out rates and both sets of lagged futures rates I(0). Only t-ratio tests for OLS regressions including a constant and no trend are reported, but none of the other usual tests gave conflicting results. The number of lags used was chosen by beginning with three lags of the data, and then testing the joint significance of progressively omitting the longer lags, as discussed in Hamilton (1994). Only the test for cointegration between close-out rates and the two-quarter ahead futures rates required augmentation with a single lag.

Table 1 - Unit root and cointegration tests

Unit root test on:

Close-out rate levels	-1.43
Close-out rate first differences	-4.20***
3-month futures rate levels	-1.90
3-month futures rate first differences	-4.38 ***
6-month futures rate levels	-1.46
6-month futures rate first differences	-4.70 ***

Cointegration tests between close-out rate and:

3-month futures rate	-4.88 ***
6-month futures rate	-2.75 *
6-month futures rate (using 1 lag)	-3.70 ***

Note: 5 percent critical values from Hamilton (1994) are -3.00 for unit root tests, and -3.39 for cointegration tests. Ten percent (*), 5 percent (**), and 1 percent (***) levels of significance are indicated.

To find the appropriate order of p in the one-quarter-ahead model, the cross-correlations between the residuals of the OLS levels regression and leads and lags of ΔF_{t-1} are estimated. The results are contained in Table 2, and show strong correlation in the first lead of the lagged forward rate (that is, the contemporaneous term). This is intuitive, since news often causes changes in the general level of the yield curve. Hence, changes in the futures rate and errors in the previous period's prediction of the spot rate are likely to have the same signs, and similar magnitudes.

The correlation results suggest that $p=1$ is appropriate in Equation (2), which reasonably assumes that cross-correlations associated with higher order leads and lags

are zero. Specifically, Equation (2) is estimated by FGLS using the Cochrane-Orcutt technique, allowing for first-order autoregressive correlation. The latter was suggested by analysing the residuals of the initial estimation of Equation (2) by OLS.

Table 2 - Estimation of one-quarter ahead model

Cross-correlation tests

$R_t = \alpha + \beta F_{t-1} + z_t$

$\Delta F_{t-1} = F_{t-1} - F_{t-2}$

Lead / lags	Correlation with z_t
ΔF_{t-3}	-0.11
ΔF_{t-2}	0.06
ΔF_{t-1}	0.19
ΔF_t	0.89 ***
ΔF_{t+1}	0.12

Note: $\Delta F_t = [\Delta F_{t-1}]_{+1}$, and 5 percent critical value is 0.44.

Model estimation

$$R_t = \alpha + \beta F_{t-1} + \gamma_1 \Delta F_t + \gamma_2 \Delta F_{t-1} + \gamma_3 \Delta F_{t-2} + e_t \; ; \; e_t = \phi e_{t-1} + u_t \qquad (3)$$

Coefficient	Estimate	Test statistic
α	-0.38	-1.25
β	1.05	32.88 ***
γ_1	0.88	17.32 ***
γ_2	-0.08	-1.54
γ_3	-0.02	-0.44
β-1	0.05	1.56

Note: 5 percent critical value t(28) = 2.05

Hypotheis Tests	Coefficient	Test Statistic	5 Percent Critical Value
α=0 and β=1	n/a	1.59	F(2,28) = 2.34
ϕ	0.47	(on e_t) 7.13 ***	$\chi^2(1) = 3.84$

Tests on FGLS residuals (u_t)	Test Statistic	5 Percent Critical Value
Breuch-Godfrey test for AR(1)	0.30	$\chi^2(1) = 3.84$
ARCH(1) test for heteroskedasticity	0.01	$\chi^2(1) = 3.84$
Goldfeld-Quant test for heteroskedasticity	1.62	F(16,17) = 2.29
Jarque-Bera test for normality of residuals	0.46	$\chi^2(2) = 5.99$

Similarly, to find the order of *p* in the two-quarter-ahead model, we again check for significant cross-correlations. The results are contained in Table 3, which shows strong correlation in the two contemporaneous terms. In this case, news in both periods will effect similar changes in the forward rate and errors in previous period's predictions.

The cross-correlation results suggest that we use *p*=2 in Equation (1). Specifically, Equation (3) is estimated by FGLS using the Cochrane-Orcutt technique, allowing for first-order autoregressive correlation. The latter was suggested by analysing the residuals of the initial estimation of Equation (3) by OLS.

Table 3 - Estimation of two-quarter ahead model

$R_t = \alpha + \beta F_{t-2} + z_t$

$\Delta F_{t-2} = F_{t-1} - F_{t-2}$

Lead / lags	Correlation with z_t
ΔF_{t-4}	-0.03
ΔF_{t-3}	-0.09
ΔF_{t-2}	0.05
ΔF_{t-1}	0.22 *
ΔF_t	0.72 ***
ΔF_{t+1}	0.60 ***
ΔF_{t+2}	0.03

Note: $\Delta F_t = [\Delta F_{t-1}]_{+1}$, and 5 percent critical value is 0.45.

Model estimation

$R_t = \alpha + \beta F_{t-1} + \gamma_1 \Delta F_{t+1} + \gamma_2 \Delta F_t + \gamma_3 \Delta F_{t-1} + \gamma_4 \Delta F_{t-2} + \gamma_5 \Delta F_{t-3} + e_t \; ; \; e_t = \phi e_{t-1} + u_t$

(4)

Coefficient	Estimate	Test statistic
α	-0.61	-1.13
β	1.09	19.04 ***
γ_1	0.84	10.46 ***
γ_2	1.04	11.96 ***
γ_3	-0.03	-0.39
γ_4	0.08	1.02
γ_5	0.11	1.38
β-1	0.09	1.52

Note: 5 percent Critical $t(26) = 2.06$

Hypotheis Tests	Coefficient	Test Statistic	5 Percent Critical Value

$\alpha=0$ & $\beta=1$	n/a	1.85	F(2,28) = 2.34
ϕ	0.51	(on e_t) 8.13 ***	$\chi^2(1) = 3.84$
Tests on FGLS residuals (u_t)		Test Statistic	5 Percent Critical Value
Breuch-Godfrey test for AR(1)		0.34	$\chi^2(1) = 3.84$
ARCH(1) test for heteroskedasticity		0.10	$\chi^2(1) = 3.84$
Goldfeld-Quant test for heteroskedasticity		1.46	F(15,16) = 2.35
Jarque-Bera test for normality of residuals		2.58	$\chi^2(2) = 5.99$

From the above results, we cannot reject the hypothesis that futures rates are an unbiased predictor of 90-day interest rates for the one and two-quarter-ahead horizons. However, the presence of significant first order serial correlation suggests that the predictions on both horizons may be information inefficient. This cautious proposition is deliberately left open, for reasons explained further in Section 7.

Note also that the diagnostic tests on the FGLS residuals contain no obvious evidence of mis-specification in the model. In particular, the Goldfeld-Quant tests suggests that there is no obviously atypical part of the sample, in terms of underlying heteroskedasticity. Larger variance might have been expected in the first half of the period due to one-off events associated with the disinflation process at that time. The sample is too small to derive meaningful estimates of parameters from sub-samples of the data for comparison.

5 Relative forecasting performance and results

The ratio of RMSE errors is used to compare the relative predictive power of futures rates against the random-walk. The method of Ashley, Granger, and Schmalensee (1980), which jointly tests the relative performance of the absolute mean forecasting error and the variance of forecasting errors, is also used to ascertain the statistical significance of the relative RMSEs. Details of how to conduct the test is contained in Ashley et al. (1980) and Deaves (1996), but essentially futures rate predictions are superior if the absolute mean error and/or the variance of realised forecast errors is significantly less than the random-walk, without either the mean error or variance being significantly worse.

Table 4 contains the results used to compare the relative forecasting performance of futures rates and the random-walk for each forecasting horizon, that is the number of days to close-out (DTC) for the futures contract. Also contained in Table 4 are comparisons of variance in forecast errors over the first half of the sample, compared to the variance in forecast errors over the second half of the sample, as explained below.

Table 4 - Relative forecasting performance

Days to Close-out (DTC)	RMSE Futures (percentage points)	RMSE Random-walk (percentage points)	Futures / RW Ratio	AGS Significance (1)	Futures Period Variance Ratio (2)
5	0.18	0.19	0.93		2.14 *
10	0.23	0.27	0.85	***	1.48
15	0.34	0.38	0.90	***	1.27
20	0.38	0.45	0.84	***	1.77
25	0.44	0.48	0.92	*	1.07
30	0.46	0.48	0.96	**	0.91
35	0.52	0.55	0.94	**	0.98
40	0.63	0.67	0.94	***	0.98
45	0.74	0.76	0.97	**	0.93
50	0.71	0.77	0.92	**	1.08
55	0.75	0.79	0.95	*	1.26
60	0.81	0.87	0.93	*	1.12
65	0.88	0.98	0.91	**	0.86
70	0.95	1.02	0.93	***	1.13
75	0.94	1.07	0.88	***	0.91
80	0.92	1.07	0.86	***	0.81
85	0.99	1.11	0.89	***	0.95
90	0.97	1.09	0.89	***	0.89
95	0.98	1.07	0.91	**	0.95
100	0.98	1.09	0.90	***	0.98
105	1.03	1.13	0.91	***	0.97
110	1.03	1.13	0.91	***	1.08
115	1.03	1.14	0.90	***	0.99
120	0.99	1.09	0.91	***	1.05
125	1.05	1.19	0.88	***	1.14
130	0.95	1.08	0.88	**	1.06

Note (1): * = 10 percent significance AGS
** = 5 percent significance AGS
*** = 1 percent significance AGS

Note (2): See discussion below. 5 percent critical value is 2.22.

The ratio of RMSEs indicate that the futures forecast outperform the random-walk on all horizons, and the AGS results indicate that the out-performance is statistically significant for all horizons greater than 1 week. This is broadly consistent with the results by Deaves (1996), who reports that implied forward rates outperform the random-walk over all forecasting horizons greater than 1 quarter, using quarterly data.

It is also interesting to assess the absolute forecasting performance of futures rates as a function of DTC. According to the efficient markets hypothesis (see, for example Brigham and Capenski (1994)) forecast errors should arise as the result of new information being received by the market. Assuming that new information arrives at a constant rate, the variance of forecast errors would be expected to increase with the forecast horizon. Moving-average autocorrelation between forecast errors would also

be expected to occur when the forecasting horizon exceeds the frequency of observations, since the new information will affect the realised forecast errors in a similar manner on the overlapping days. Hence, the realised forecast errors have been adjusted for expected first-order MA autocorrelation for 65 to 130 DTC. This should give a superior measure of the innovation variance of the underlying random-walk of the spot/futures system.

Expected first-order autocorrelation when the days to close-out is greater than 65 is given in Equation (4), which yields the MA(1) coefficient in Equation (5), which is then used to adjust the realised RMSE as in Equation (6). The latter two equations are derived from formulas contained in Hamilton (1994).

$$\rho_1 = \frac{DTC - 65}{DTC} \tag{5}$$

$$\theta = \frac{1 - \sqrt{1 - 4\rho_1^2}}{2\rho_1} \tag{6}$$

$$\sigma^2 = \frac{RMSE(e_t)}{1 + \theta^2} \tag{7}$$

where: r_1 is the expected first-order autocorrelation;
DTC is the number of days to close-out;
q is the first-order moving-average coefficient;
σ^2 is the underlying innovation variance;
e_t are the forecast errors.

Figure 2:   Mean squared error of futures rate and random-walk predictions

Days to close-out (DTC)
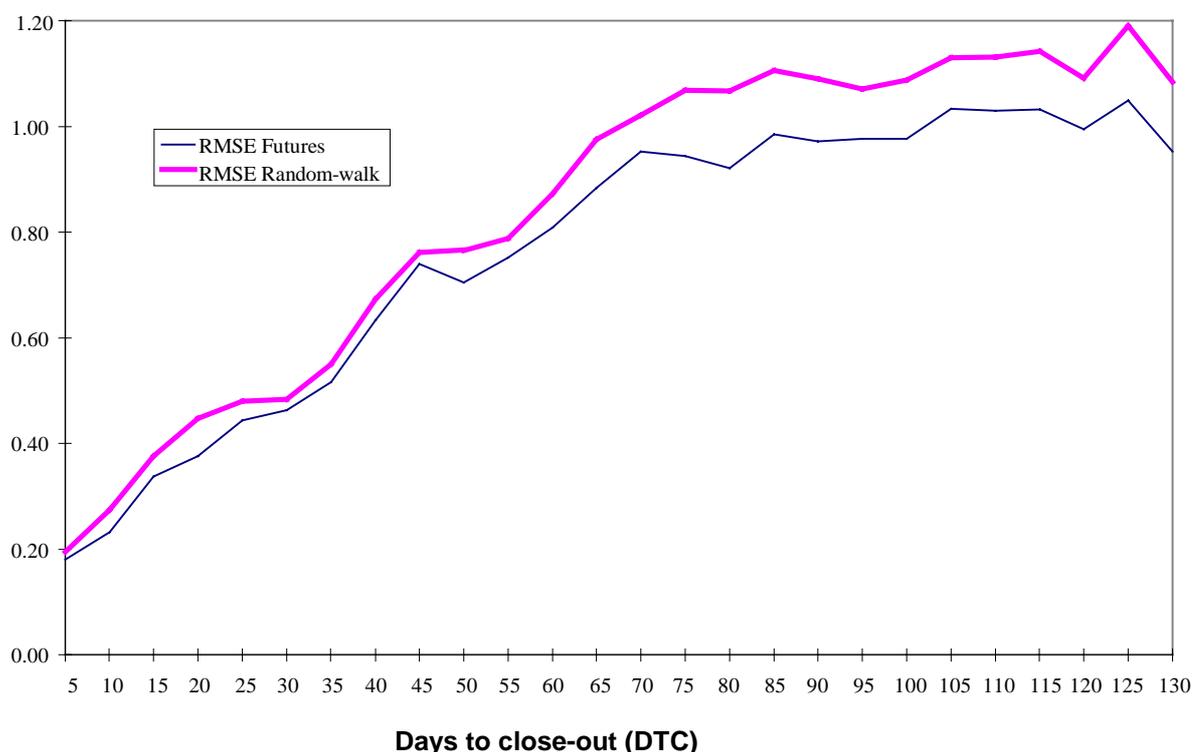
The first point of note is the that both spot futures rates and spot 90-day rates provide predictions with wide confidence intervals, as indicated by the RMSE of the residuals. Even for horizons of only 30 DTC (6 weeks), the 95 percent confidence band is approximately +/- 1 percentage point, rising to approximately +/- 2 percentage points for forecasting horizons of 100 DTC (20 weeks). The difference in the absolute RMSEs is also relatively small: approximately 0.02 percentage points for 30 DTC, and 0.11 percentage points for 100 DTC. The difference in the confidence regions are then only in the order of magnitude of a typical market spread; hence the statistically superior forecasting performance of futures rates would not necessarily translate to something useful or exploitable in practice (since trading would involve "crossing the spread", which would introduce a transactions cost).

The second point of note is that the confidence intervals initially widen consistently with DTC (a chart of variance appears linear up to approximately 80 DTC). This is qualitatively consistent with the efficient markets hypothesis. It is also interesting to note that the confidence intervals tend to remain relatively constant beyond 80 DTC (16 weeks), which is qualitatively consistent with interest rates being a weakly mean-reverting process. This is not necessarily inconsistent with the Dickey-Fuller tests in Table 1; the hypothesis of a random-walk on a quarterly basis is not rejected by the data, but this does not necessarily mean that the process is truly a random-walk. In practice, the level of 90-day interest rates usually stays within a 5 to 9 percent range, so there should be a practical limit to how far any prediction can go wrong.

Finally, the forecast errors are tested to see if there is any difference in the variance of forecast errors over the first half of the sample, compared to the second half of the sample. The results contained in Table 4 indicate that there is no significant difference. Jarque-Bera tests for the normality of the forecast errors (not shown) also showed no evidence of non-normality. As discussed in Section 4, this suggests that there is no obviously atypical sub-samples within the data.

6   Conclusions

The results from this paper show that futures rates provide unbiased predictions of 90-day rates for the one and two-quarter-ahead horizons. However, the futures predictions may be information inefficient, as suggested by significant first-order autoregressive correlation of the estimated residuals.

Regardless of this potential information inefficiency, futures rates do provide superior forecasts of 90-day rates compared to the random-walk for all horizons tested, up to 26 weeks. The superior performance is statistically significant for all horizons greater than 1 week.

The results are best interpreted with reference back to the spot/forward data generating process described in System (1). In that model, changes in the spot rate are dominated by the underlying system innovations (v_t) that move the spot and forward simultaneously by equal amounts. The common effect of those underlying innovations are effectively removed by the construction of the Stock and Watson (1993) method, and that estimation then shows that lagged forward rates do contain some marginal information about changes in the spot rates.

Hence, futures rates would be expected to provide superior predictions of spot rates compared to the random-walk, and this is confirmed when the relative predictive performance is tested. However, because both predictions will still be dominated by the underlying system variance, the relative performance of the futures rates compared to the random-walk would never be expected to be vastly superior, and this is also consistent with the results from the relative forecasting performance. Hence, the statistically superior forecasting performance of futures rates would not necessarily translate to something useful or exploitable in practice.

7   Other comments and possible further work

The evidence for significant first-order autoregressive correlation in the estimated residuals of Equations (2) and (3) suggests that a model that accounts for such correlation might provide superior predictions, compared to simply using futures rates. This proposition would best be tested by comparing out-of-sample predictive performance, which is not possible in this paper due to the relatively small amount of quarterly data available.

The proposition could possibly be tested using implied forward interest rates on a monthly basis, which would allow for a sufficient amount of data. However, even if superior predictions were realised using mid-rate data, this would not necessarily mean that markets were information inefficient, since typical market spreads should also be allowed for. One way of testing the latter is to actually simulate trading rules that use practical market conventions and only historical information, and test if such rules return statistically significant, and practically significant, profits.

It is also worth noting that any potential improvement in the predictive power of forward rates by accounting for the serial correlation of residuals will still only be at the margin. The relative prediction performance results above suggest that the majority of prediction errors arise from the innovations underlying the spot/futures system (v_t in System (1)). Hence, it might be more productive to try and explain more of the latter variance by including other relevant economic variables.

For example, Rae (1997) suggests that prediction errors in the spot/forward exchange rates may arise due to forecasting errors in inflation. If this is the case, and there is any consistent systematic errors in inflation forecasting, including inflation as a variable might help in producing superior predictions of the spot/forward rate. This might provide stronger evidence of any potential market information inefficiencies.

Appendix 1:   Sample information for forecasting performance

Date	Close-out Date	Futures CIR	Close-out CIR	Futures Error	Trading Days to Close-out
7/12/88	Wed-15-Mar-89	14.10	12.85	1.25	70
8/03/89	Wed-14-Jun-89	13.04	13.20	-0.16	70
7/06/89	Wed-13-Sep-89	13.13	13.09	0.05	70
6/09/89	Wed-13-Dec-89	12.57	13.88	-1.31	70
6/12/89	Wed-14-Mar-90	13.46	13.49	-0.03	70
7/03/90	Wed-13-Jun-90	13.23	13.60	-0.37	70

References

Ashley, R, C Granger and R Schmalensee (1980), "Advertising and Aggregate Consumption: An Analysis of Causality," Econometrica 48, 1149-67.

Brigham, E and L Capenski (1994), Financial Management, Theory and Practice, Harcourt Brace College.

Campbell J and R Shiller (1991), "Yield Spreads and Interest Rate Movements: A Bird's Eye View," Review of Economic Studies, 58, 495-514.

Cutherbertson, K (1996), "The Expectations Hypothesis of the Term Structure: The UK Interbank Market," The Economic Journal, 106 (May), 578-592.

Deaves, R (1996), "Forecasting Canadian Short-term Interest Rates," Canadian Journal of Economics, xxxix, number 3 (August), 615-634.

Engsted, T (1996), "The Predictive Power of the Money Market Term Structure," International Journal of Forecasting 12, 289-295.

Hamilton, J (1994), Time Series Analysis, Princeton University Press, 571-629.

Mishkin, F (1988), "The Information in the Term Structure: Some Further Results," Journal of Applied Economics, Vol. 3, 307-314.

Rae, D (1997), "New Zealand's Forward Exchange Market," National Bank of New Zealand Limited Financial Research Paper No. 8 (August 1997).

Stock, J and M Watson (1993), "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems," Econometrica, 61, 783-820.