THE PERFORMANCE OF EXCHANGE RATE FORECASTS

Philip W. Lowe
Robert G. Trevor*

Reserve Bank of Australia

Research Discussion Paper

8609

July 1986

* The views expressed herein and any remaining errors are our own and should not be attributed to our employer.

ABSTRACT

Since the floating of the Australian dollar the forecasting of exchange rate movements has become more difficult and received much more attention. As a result, some participants in the foreign exchange market have, on a number of occasions, come under criticism for their inability to predict exchange rate movements. This paper seeks to evaluate these criticisms through an examination of exchange rate forecasts made by market participants (as published in the Australian Financial Review from March 1985 to December 1985). The accuracy of the \$A/US\$ forecasts is compared with that of forecasts generated from a number of simple forecasting rules as well as forecasts of the US\$/Yen exchange rate. In general, the simple forecasting rules provide superior forecasts to those provided by the individual market participants. However, under some criteria, the mean of the individual participants' forecasts may be preferred to these simple forecasting rules. Further, the comparison of the US\$/Yen forecasts with the \$A/US\$ forecast shows the former to be generally more accurate.

TABLE OF CONTENTS

THE PERFORMANCE OF EXCHANGE RATE FORECASTS

Philip W. Lowe and Robert G. Trevor

> "The personal traits shared by most spot traders are often quite an explosive mixture ... Their most elusive feature - and the one they are mostly paid for - is their "instinct" to pick the trends or swings in the market ... Hand in hand with what could be called a trader's "second sight", superstitution also appears to have a niche in the forex market ... For example, one spot trader for inspiration peers at his dealing screen through a pair of black lace woman's panties."
>
> Australian Finanacial Review
> 12 December 1985, p.17

1. Introduction

Since the floating of the Australian dollar in December 1983 the domestic foreign exchange market has grown considerably. The last two and a half years have seen the number of licensed dealers increase from around 10 to over 70 and the volume of business transacted increase many fold. Paralleling this growth has been an increase in interest in the behaviour and operation of the domestic foreign exchange market. This growth has seen the market and its reaction to economic and political events become, in many eyes, one of the principal barometers of the performance of the Australian economy.

Despite the importance of the foreign exchange market the ability of its participants to provide reasonable forecasts of future exchange rates has been questioned widely. Headlines such as "Cain Lashes Forex 'Whiz Kids'",[1] "BT Charts by the Light of the Moon"[2] and "Scepticism of Currency Forecasts Grows"[3] and comments such as "predicting the exchange rate level of the $A even one week ahead is rapidly becoming a game of Russian roulette"[4] have appeared frequently in the financial press over the last twelve months.

Notwithstanding these headlines, and the theory of efficient markets which suggests that all available information regarding the future exchange rate is incorporated into the current exchange rate, foreign exchange dealers and corporate advisers have little choice but to continue making and publishing foreign exchange forecasts.

---

1. Australian Financial Review, 4 July 1986, p9.
2. Australian Financial Review, 10 June 1986, pl.
3. The Australian, 28 February 1986, Special Report, p3.
4. Australian Financial Review, 21 November 1985, pl.

While the accuracy of say a one week ahead or one month ahead forecast may not significantly influence the profitability of a trader (or his/her employer) who opens and closes positions over a period of minutes, the accuracy of such forecasts may be important to the corporate customers of the forecaster. The increased volatility of the Australian dollar since the float has made the need for such forecasts, and the consequent decisions regarding the timing of foreign exchange transactions, of increased importance to many corporate treasurers. Therefore, a failure of the forecaster to provide reasonable forecasts should be of some concern to corporate treasurers.

Any analysis of the forecasting performance of various forecasters is made difficult by the lack of a relatively long, publicly available, and consistent series of forecasts. However, during 1985 the Australian Financial Review (AFR) surveyed various foreign exchange dealers about their expectations of exchange rate movements over a one week horizon. This paper uses these forecasts to throw some light on the above criticisms.

Section 2 of the paper discusses the nature of the AFR's survey while Section 3 defines a number of benchmark forecasting models against which the performance of the forecasts published in the Australian Financial Review are compared.[5] The results of this comparison for the A/US exchange rate forecasts are presented in Section 4, while Section 5 compares the performance of these forecasts with that of forecasts of the US$/Yen exchange rate. Section 6 compares the performance of individual forecasters with a "no change" forecast model and with the "group mean" forecasts. The principal conclusions of the paper are drawn together in Section 7.

2. The Forecasts

The AFR commenced its weekly survey of the forecasts of foreign exchange dealers on 13 March 1985. The survey continued until 18 December 1985 when the AFR ceased publishing the forecasts.[6] Over this period of 41 weeks a total of 406 individual forecasts were provided on various exchange rates. Under the terms of the AFR survey a number of foreign exchange dealers were asked each Wednesday, to predict four pieces of currency-related data for the coming week. The survey participants were asked to estimate:

---

5. These forecasts will be referred to as the "AFR survey forecasts".

6. The results of a similar survey are now published on the "Reuters screen", just after page FEWV.

(1) the range of the \$A/US\$ exchange rate over the following week;

(2) the \$A/US\$ hedge settlement rate[7] (HSR) current on the following Wednesday;

(3) the US\$/Yen exchange rate applying at the close of trading in Tokyo on the following Wednesday; and

(4) the \$A/US\$ three-month forward margin applying on the following Wednesday.

The results of the survey were published in the AFR each Thursday. The maximum number of forecasters included in the survey on any one week was twelve. Participation in the survey was on a revolving basis, with the two forecasters providing the week's poorest forecasts of the hedge settlement rate being dropped from the following week's survey and replaced by two forecasters from the AFR's master list of forecasters. This master list consisted of all dealers with a licence to deal in the domestic foreign exchange market. A total of 49 different forecasters provided at least one forecast over the sample period with the maximum number of forecasts provided by any one forecaster being 35. Only two dealers supplied more than 20 forecasts.

The data on the forecasts for the \$A/US\$ hedge settlement rate and the US\$/Yen exchange rate will be used in the current analysis.[8] Figure 1 is a plot of the actual (solid line) and forecasted (broken line) \$A/US\$ exchange rate. It shows that there was a small net depreciation of the Australian currency over the sample period. The graph also shows that there were frequent changes in the direction of movement. As a result, a forecaster who was consistently pessimistic about the value of the \$A need not have supplied forecasts superior to those generated from a simple no change model.

---

7. The hedge settlement rate is fixed each day. It is calculated by taking the average of the midpoints of the Australian dollar/U.S. dollar offer/bid rates of the ten contributing banks on the OZZU Reuters page at 9.45 a.m. The rate is used to close off all hedge contracts maturing on that day. It is quoted as 1\$A=xUS\$. A fall in this rate is thus a depreciation of the Australian dollar against the U.S. dollar.

8. The responses to the other two questions will not be used in this study. It is difficult to evaluate forecasts of trading ranges (e.g., is a range that is too narrow preferable to one that is too wide), and there is little weekly movement in forward margins.

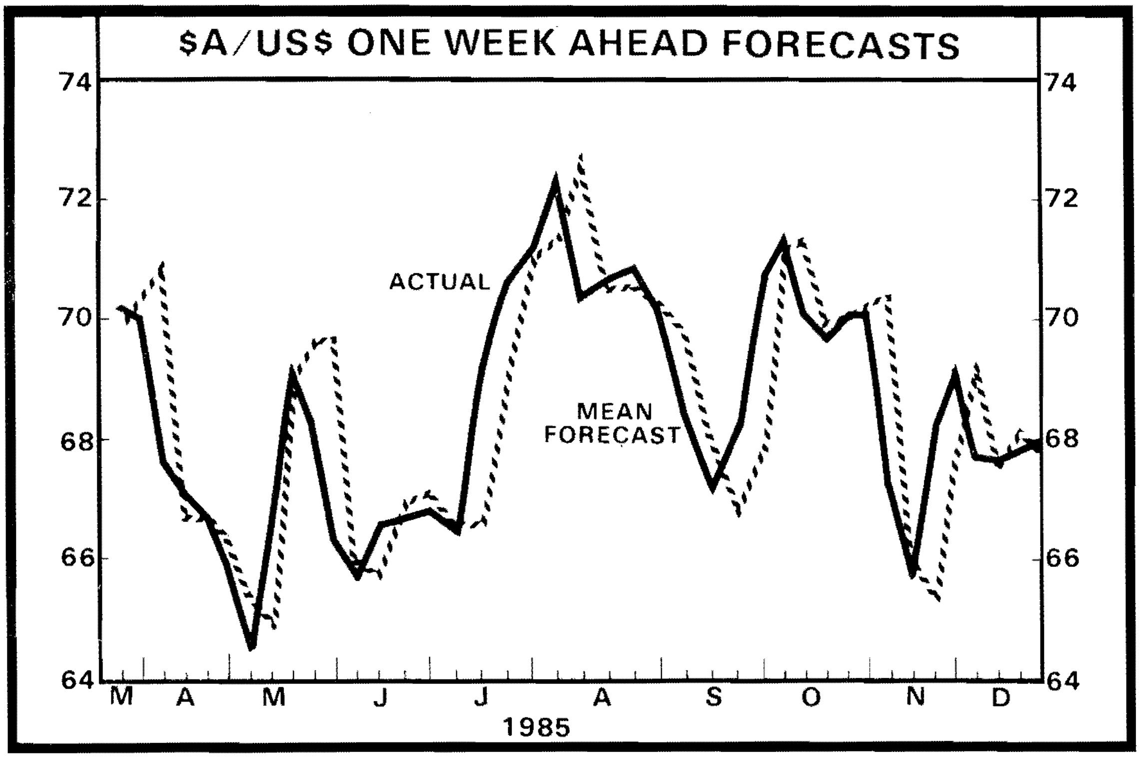

Figure 1

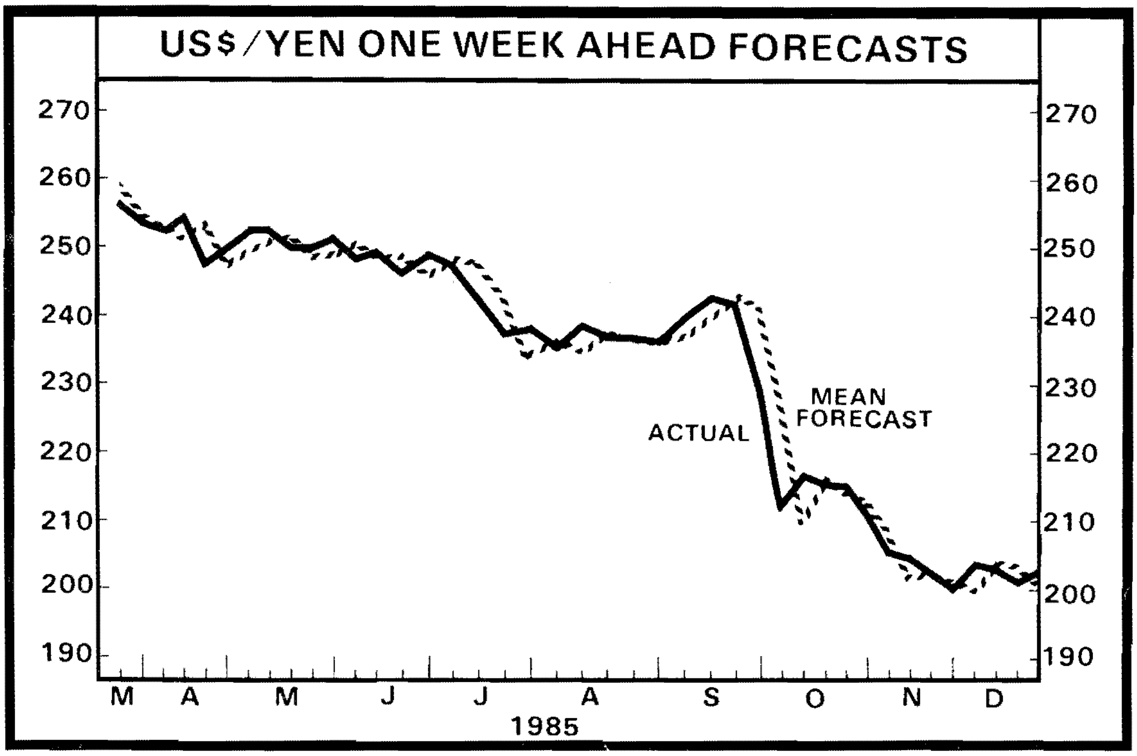

Figure 2

The graph of the US$/Yen exchange rate, presented in Figure 2, shows that the Yen appreciated against the US$ over the sample period. The bulk of the depreciation of the US$, however, occurred in late September 1985. Notwithstanding this overall weakening of the US$ there were many weeks over which the US$ strengthened against the Yen.

3. Benchmark Models and Evaluation Criteria

In order to evaluate the accuracy of economic forecasts it is necessary to define a number of benchmark forecasting models against which the forecasts can be compared. These benchmark models may take many forms including naive forecasting rules, econometric forecasting equations, and time series models. Their common feature, however, is that they should only use information available at the time the forecast was made. This paper compares the AFR survey forecasts to forecasts generated by each of the above three types of benchmark forecasting models. It has, however, not been our intention to develop optimal forecasting equations. Rather, our aim has been to specify a number of relatively simple forecasting models, the forecasts from which we can compare with the forecasts published in the AFR.

(a) Naive Forecasting Rules

Two naive forecasting rules are employed in this paper. The first of these says that the best forecast of the future exchange rate is the current exchange rate. That is

$$E^f_{t+1} = E_t$$

where E^f_{t+1} is the forecast of the exchange rate in a week's time and E_t is the current exchange rate. This forecasting rule, known as the no change or random walk rule, has been found to perform relatively well in providing out of sample exchange rate forecasts.[9]

The second naive benchmark model assumes that the expected change in the exchange rate is equal to the previous change. That is:

$$E^f_{t+1} - E_t = E_t - E_{t-1}$$

9. For example Meese and Rogoff (1983) found that a no change model forecasted a number of major-country exchange rates as well as a number of structural and time series models.

This rule can be re-expressed as,

$$E^f_{t+1} = 2E_t - E_{t-1}$$

Since this forecasting model involves extrapolating the most recent change in the exchange rate into the future, it is known as the <u>extrapolative expectations model</u>.

(b) <u>Econometric Forecasting Equations</u>

Recent years have seen the development of a large number of models attempting to explain the behaviour of bilateral exchange rates. These models characteristically have explanatory variables such as relative money supplies in the two countries, relative interest and inflation rates and relative growth in national incomes. These models are, however, of limited use in explaining exchange rate behaviour over a period as short as a week. This is a result of data problems and the fact that such models have been developed to explain exchange rate movements over relatively long periods of time.

Of the above variables typically included in exchange rate models, the interest rate differential is most likely to influence weekly movements in exchange rates. This influence can arise through two important links between exchange rates and interest rates. Firstly, portfolio balance models of exchange rates suggest that an increase in the interest differential in favour of the foreign currency will, ceteris paribus, strengthen the foreign currency as the demand for foreign assets increases. Alternatively, under interest rate parity theory an increase in the interest differential will indicate an expected depreciation of the foreign currency. Given these links between interest rates and exchange rates and the fact that the interest differential is one of the few variables observed sufficiently frequently to be included in a weekly forecasting model, the following equation was specified:

$$E^f_{t+1} = a + b_1E_t + b_2(r^*_t - r_t)$$

where $(r^*_t - r_t)$ is the current three month interest differential (i.e., the gap between U.S. and Australian interest rates).[10]

10. These interest rates are the U.S. 90-day prime bankers' acceptances rate and the Australian 90-day bank accepted bill rate. These rates are more readily available than yields on instruments of a week's maturity.

As a result of the lack of other appropriate variables, forecasting equations with the current hedge settlement rate and a current forward margin as explanatory variables were also estimated. The forward margin should, in some sense, act as a proxy for the variables that we are unable to observe. Given the unavailability of one week forward margins, equations were estimated using 15, 30 and 90-day margins (F_t^{15}, F_t^{30} and F_t^{90} respectively).

In estimating the forecasting equations, only data available at the time the forecasts were made are used. With the Australian dollar being floated on 12 December 1983 and the AFR survey commenced on 13 March 1985, there are 65 weekly observations with which to estimate the initial forecasting equations (allowing for lags). The choice of benchmark models from these forecasting equations was made on the basis of the sample fit over these first 65 observations. Model selection was not based on the quality of the fit over the entire period from December 1983 to December 1985 (107 observations) as this information was only known ex-post. Hence, a market participant could not have used this information in the preparation of his/her forecasts.

Table 1 reports the two selected benchmark forecasting equations. They were estimated by ordinary least squares over the first 65 observations. These equations were used to provide the first ex-ante one step ahead forecasts of the \$A/US\$ hedge settlement rate (that is the forecast of the hedge settlement rate for 20 March 1985 made on 13 March 1985). Each week the sample period was updated by one observation and the equations were re-estimated using the augmented sample. These new estimates were then used to generate a new ex-ante one step ahead forecast. This updating, re-estimating and forecasting procedure was continued until forecasts had been generated for each of the 41 weeks for which forecasts were published in the AFR.

Table 1
Initial Forecasting Equations

$$E_{t+1} = \underset{(-0.81)}{-2.00} + \underset{(0.03)}{1.02}\, E_t + \underset{(0.08)}{0.15}\, (r_t^* - r_t)$$

$\bar{R}^2 = 0.9548$ Durbins h = -0.017

$$E_{t+1} = \underset{(2.41)}{-2.71} + \underset{(0.03)}{1.03}\, E_t + \underset{(0.01)}{0.02}\, F_t^{30}$$

$\bar{R}^2 = 0.9556$ Durbins h = 0.025

Note: Standard errors appear in brackets

(c) Time Series Models

An alternative to using a simple forecasting rule or an econometric forecasting equation to produce forecasts is to use a model derived solely from the information contained in the history of the exchange rate itself. Two such models, known as time series models, were selected as benchmark forecasting equations for this paper. The first of these is the simplest possible such model. That is:

$$E^f_{t+1} = a + b\ E_t$$

This model is slightly less constrained than the no change model. In the no change model the intercept (a) is set to zero, the coefficient on the current exchange rate (b) is set to one and coefficients on lags of the exchange rate are all set to zero. In this simple time series model the coefficients on lagged exchange rates remain set at zero, however, the intercept and the coefficient on the current exchange rate are estimated rather than imposed. This simple time series model can be further generalised by estimating the coefficients on the lagged values of the exchange rate. Such a model constitutes the second time series benchmark model used in this paper. To remove subjectivity from the lag specification, a computer procedure was used to select the lags to be included in the model.[11] The only decision required by the forecaster concerns the significance levels to be used in selecting the lags.

As with the econometric forecasting models, the first ex-ante one step ahead forecast was generated from an equation estimated over the period December 1983 to March 1985. Following the calculation of this forecast, the updating, re-estimating and forecasting procedures were repeated to generate a complete set of 41 forecasts.

When the first 65 observations were used to estimate the second time series model (i.e., the computer-generated, or "optimised" model)[12] the first and

---

11. Namely, the STEPAR method of SAS's PROC FORECAST. This method fits a time trend model to the series and then fits an autoregressive process to the de-trended series using a backwards-stepping method to select parameters. That is, the least significant parameter is removed from the model, which is then re-estimated. This process continues until only significant autoregressive parameters remain. The chosen level of significance was 5 per cent.

12. This model is "optimised" in the sense that it minimises the sum of squared residuals over the estimation period. This period always includes the first 65 observations which are not part of the forecast sample. Thus, it will not necessarily be the "optimal" forecasting equation over this period.

fourth lags proved to be significant. Subsequent re-estimations saw the second through seventh lags become significant on various occasions. In the following discussion the first of these time series models will be referred to as the "restricted" time series model and the second as the "optimised" time series model.

(d) Measures of Forecast Performance

Three criteria are used to compare the relative forecast performance of the various models. The first of these is the mean absolute error (MAE) of the forecasts, which is defined as

$$MAE = \frac{1}{n} \sum_{i=1}^{n} \left| E^{f}_{t+1,i} - E_{t+1,i} \right|$$

where n is the number of forecasts.

This criterion provides a measure of the average forecast accuracy of the AFR survey and benchmark forecasts. The second standard of comparison is the mean square error (MSE) of the forecasts, where

$$MSE = \frac{1}{n} \sum_{i=1}^{n} (E^{f}_{t+1,i} - E_{t+1,i})^2$$

This measure penalizes large forecast errors to a greater extent that does the MAE criterion.

The third measure used to compare the performance of the various forecasts is the percentage of forecasts which predict the correct directional movement. This criterion is important since the timing of the conversion of domestic currency into foreign currency (and visa versa) often depends critically on the expected direction of movement of the exchange rate.

For the purpose of evaluating the performance of the AFR survey forecasts the mean and the median of each week's forecasts were calculated. These forecasts are referred to as the "group mean" and "group median" forecasts respectively.

4. Forecast Performance of the Group Mean and Median A/US Forecasts

Table 2 reports the mean absolute errors and the mean square errors for the group mean and median forecasts and for the six benchmark forecasting models.

Table 2
Forecast Accuracy

Model	Mean Absolute Error	Mean Square Error
AFR Survey		
- Group Mean	0.9838	1.9906
- Group Median	1.0138	2.0183
Benchmark Models		
- No Change	1.0551	1.7852
- Extrapolative	1.3485	2.8509
- Interest Differentials	1.0874	1.8669
- Forward Margin	1.1320	2.0245
- Restricted Time Series	1.0315	1.7787
- Optimised Time Series	1.3322	2.4217

Comparison of the relative forecasting performance of the group mean and the median forecasts shows the mean forecast to be slightly superior to the latter, having both a smaller MAE and MSE. This superiority of the mean forecast over the median is consistent with recent studies of the forecasting accuracy of US forecasters (see Zarnowitz (1982b) and Hafer (1984)).

This superiority stems from the fact that the mean value gives some weight to all pieces of information, while the median may ignore certain pieces of information. For example, suppose a forecaster in the foreign exchange market gains last minute sole access to a piece of information which suggests that the exchange rate is likely to fall significantly over the next week. As a result he predicts a sharper fall in the exchange rate than he was previously predicting. This lower prediction will reduce the value of the mean forecast, but may well leave the median forecast unchanged. Since the mean makes more extensive use of such information, it is not surprising that the mean forecast series proves to be superior to that of the median.

Given the superior performance of the group mean forecast it, rather than the group median, will be used to compare the performance of the AFR survey forecasts with that of the benchmark forecasts.

A comparison of the MAE figures in Table 2 shows the group mean AFR survey forecast to have been more accurate than any of the six benchmark forecasting models. All the benchmark models produced a mean absolute error of greater than one US cent while the group mean forecast produced a mean absolute error of only 0.9838 US cents. Of the six benchmark models the no change model and the restricted time series model produced the most accurate forecasts.

Given that the optimised time series model is less restrictive than the restricted model, its inferior performance is perhaps a little surprising. The explanation, however, lies in the fact that the models were estimated using "out of sample" data. The inferior performance of the more general model suggests that the structure of the time series of exchange rates changed through time. That is, the relationship between the exchange rate and its lagged values was different over the period December 1983 to March 1985 to that over the period March 1985 to December 1985. Had the forecasting equations been estimated using "in sample" data the general model would have out-performed the one lag model (and the no change model). However, as discussed in Section 3 the use of "in sample" data is not appropriate in estimating benchmark forecasting equations, because these data were not available to market participants at the time they prepared their forecasts.

While the group mean forecast series performed better than all the benchmark forecasting equations using the MAE criterion, it proved to be superior to only three out of the six benchmark forecasts using the MSE criterion. The no change model, the restricted time series model and the interest differentials model all produced forecasts with a lower MSE than the mean AFR survey forecast. The optimised time series model and the extrapolative expectations model again proved to be the worst performers.

The third criterion used to evaluate the performance of the various forecasts is the percentage of forecasts predicting the correct direction of movement. Table 3 reports these percentages for the group mean and median forecasts and for the benchmark forecasting models. The table shows that the group mean forecast produced the highest percentage of forecasts in the correct direction (65.9%). Of the benchmark models the extrapolative model was the only one to produce forecasts in the correct direction on more than 50 per cent of the 41 weeks. The poorest performer on this criterion was the optimised time series model.

The fact that approximately two thirds of the group mean forecasts predicted the correct direction of the hedge settlement rate movement suggests that movements in the rate differ from a pure random walk. The result indicates that the market often had some information regarding the direction of movement of the spot rate. This information, however, seems to consist of little more than the previous movement in the spot rate, as the extrapolative model performs almost as well as the AFR survey forecasts. Such a result would be consistent with a predominant use of charting methods in the formation of expectations in the foreign exchange market.

Table 3
Predicting Direction of Movement

Model	Predictions in wrong direction	Predictions in correct direction		
		Under prediction	Over prediction	Totalϕ
	%	%	%	%
AFR Survey				
- Group Mean	34.1	51.2	14.6	65.9
- Group Median	36.6	39.0	24.4	63.4
Benchmark Models				
- No Change	n.a.	n.a.	n.a.	n.a.
- Extrapolative	39.0	22.0	39.0	61.0
- Interest Differentials	51.2	31.7	17.1	48.8
- Forward Margin	51.2	26.8	22.0	48.8
- Restricted Time Series	51.2	36.6	12.2	48.8
- Optimised Time Series	58.5	29.3	12.2	41.5

n.a. not applicable - the "no change model" predicts an unchanged rate.
ϕ Components may not sum to "Total" due to rounding.

The above comparisons of the performance of the mean AFR survey forecast series with the benchmark forecasts, suggests that on average the mean AFR survey forecast provided a slightly better forecast of the hedge settlement rate on the following Wednesday than did any of the benchmark forecasts. However, the MSE results suggest that on a number of occasions when the exchange rate moved significantly, the mean forecast not only failed to pick the size of the movement but also failed to pick the direction of the movement.

Any conclusion, therefore, regarding the overall performance of the mean of the AFR survey forecasts must rest on an assessment of the costs of making occasional large forecast errors. The occasional large errors are likely to be of major concern if transactions dependent on the forecasts are executed relatively infrequently. If positions are taken or altered continuously on the basis of the forecasts, the effect of the occasional large error may well be offset by the consistently relatively good forecasts. If, however, such transactions are undertaken infrequently, forecasts which exhibit a relatively low mean absolute error but relatively high mean square error may well prove to be inferior to forecasts with a relatively high mean absolute error and low mean square error. This is particularly the case for highly risk averse transactors.

5. A Comparison of the A/US and US$/Yen Forecasts

This section compares the forecasting accuracy of the A/US and US$/Yen forecasts. To facilitate comparison, the forecast errors have been converted to percentage forecast errors.

Table 4 reports the mean absolute percentage errors and the mean square percentage errors of the forecasts, as well as statistics on the directional performance of the two forecasted exchange rates. The table shows that the weekly percentage movements in the A/US were, on average, larger than those in the US$/Yen rate (1.54 per cent compared to 1.22 per cent). In line with these larger movements in the A/US, the mean absolute percentage error was larger for the A/US rate than for the US$/Yen rate. However, as a percentage of the no change forecast error, the mean absolute percentage error for the A/US was marginally less than the mean absolute percentage error of the US$/Yen forecasts. This was the case for both the mean and median forecasts.

A comparison of the directional performance of the forecasts of the two exchange rates shows little difference between the forecasts. For both exchange rates, the group mean and median forecasts predict the exchange rate's moving in the correct direction just less than two thirds of the time.

Table 4
Comparison of the Performance of the A/US and US$/Yen Forecasts

1. Mean Absolute Percentage Error

Forecast	A/US (Hedge Settlement Rate)	US$/Yen (Tokyo Close)
Group Mean	1.44	1.16
Group Median	1.47	1.18
No Change	1.54	1.22

2. Mean Square Percentage Error

Forecast	A/US (Hedge Settlement Rate)	US$/Yen (Tokyo Close)
Group Mean	4.27	3.10
Group Median	4.33	3.10
No Change	3.81	3.26

3. Predicting Direction of Movement

Forecast	A/US (Hedge Settlement Rate) %	US$/Yen (Tokyo Close) %
Group Mean		
Incorrect Direction	34.1	39.0
Correct Direction		
Underprediction	51.2	51.2
Overprediction	14.6	9.8
Total	65.9	61.0
Group Median		
Incorrect Direction	36.6	36.6
Correct Direction		
Underprediction	39.0	43.9
Overprediction	24.4	19.5
Total	63.4	63.4

While the above results from the MAE and direction of movement criteria suggest that there is little to distinguish between the performance of the A/US and US$/Yen forecasts, this conclusion is not supported by a comparison of the mean square errors. Using the MSE criterion, the A/US forecasts are less precise than the US$/Yen forecasts. While the MSE of the US$/Yen forecasts is some 5 per cent lower than that for the no change forecasts, the MSE of the A/US forecasts is 12 per cent higher than that for the comparable no change forecasts. These results indicate that on average the A/US was not significantly more difficult than the US$/Yen rate for the Australian foreign exchange dealers to forecast. However, the standard deviation of the forecast errors was substantially larger for the A/US forecasts than for the US$/Yen forecasts. This suggests that the level of uncertainty may have been greater in the domestic market than it was in the US$/Yen market.

Further insight into the extent of relative uncertainty in the two markets can be gained by calculating the standard deviation of each week's individual forecasts and then averaging these standard deviations over the entire period. If a market is characterised by a high level of uncertainty the forecasts of a future exchange rate by individual forecasters should be relatively dispersed. In contrast, if there is little uncertainty as to the future exchange rate the individual forecasts should be relatively tightly distributed around the mean forecast.

Table 5 reports summary information regarding the standard deviation of the individual forecasts. The table shows that the average standard deviation of the A/US forecasts was some 26 per cent higher than that of the US$/Yen forecasts (0.9264 for the A/US compared to 0.7615 for the US$/Yen rate). In addition, of the 41 weeks for which forecasts were published, the standard deviation of the A/US forecasts was greater than that of the US$/Yen forecasts on 28 occasions. These results again suggest that over the sample period, uncertainty with respect to future exchange rate movements may have been greater in the A/US market than in the US$/Yen market.

Table 5
Standard Deviation of Weekly Forecasts

	A/US (Hedge Settlement Rate)	US$/Yen (Tokyo Close)
Standard Deviation		
Average	0.9264	0.7615
Maximum	2.0530	1.8595
Minimum	0.2282	0.2135

6. Performance of Individual Forecasters

The inclusion of different forecasters in the sample each week makes difficult any accurate comparison of the relative forecasting performance of individual forecasters. As a result, the approach adopted here is not to directly compare individual forecasters with one other, but to compare their performance against that of the group mean forecasts and the no change benchmark model forecasts. Further, given that it is difficult to evaluate the performance of a forecaster when he/she has provided very few predictions, only those forecasters who provided more than five forecasts are included in the following analysis. This reduces the number of forecasters from 49 to 30.

A comparison of the performance of individual forecasters' predictions with that of a no change model is made by calculating two ratios for both the A/US and US$/Yen forecasts. The first of these, A_i/ANC_i, is the ratio of the mean absolute error made by forecaster i to the mean absolute error which would have been made had the forecaster provided a no change forecast. If the ratio is greater than unity then the forecaster provided, on average, less accurate forecasts than those provided by the no change model. If the ratio is less than unity the forecaster's performance was superior to that of the no change model.

The second ratio, S_i/SNC_i, is the ratio of the MSE of forecaster i's forecasts to the MSE which would have been made had no change forecasts been supplied. Its interpretation is similar to that of the A_i/ANC_i ratio.

Similar ratios are calculated to compare the individual forecasters' performance with that of the group mean forecasts. The ratio A_i/AGM_i is the ratio of the MAE made by forecaster i to the MAE of the group mean forecasts calculated using only those weeks for which forecaster i provided a forecast. The second ratio S_i/SGM_i [13] is similar to A_i/AGM_i, however, it uses the MSE criterion. If the ratios for an individual forecaster are greater than unity, then the forecaster provided consistently worse forecasts than those provided by the group mean forecasts. Conversely, if the ratios are less than unity the forecaster provided consistently better forecasts than the group mean forecast.

Figures 3(a) and 3(b) below show the distributions of the A/US forecast A_i/ANC_i and S_i/SNC_i ratios respectively. The most noticeable feature of both histograms is that they are skewed to the right of unity. This indicates that on average most forecasters did worse in predicting the following Wednesday's hedge settlement rate for the A/US than did the no change model. In fact, of the 30 forecasters, 24 would have produced forecasts with a lower mean absolute error had they predicted no change in the hedge

---

13. Zarnowitz (1982a) uses a similar ratio to compare the accuracy of individual forecasters included in a NBER-ASA survey with the group mean forecast.

Figure 3(a): Distribution of the A_i/ANC_i Ratios for the A/US Hedge Settlement Rate

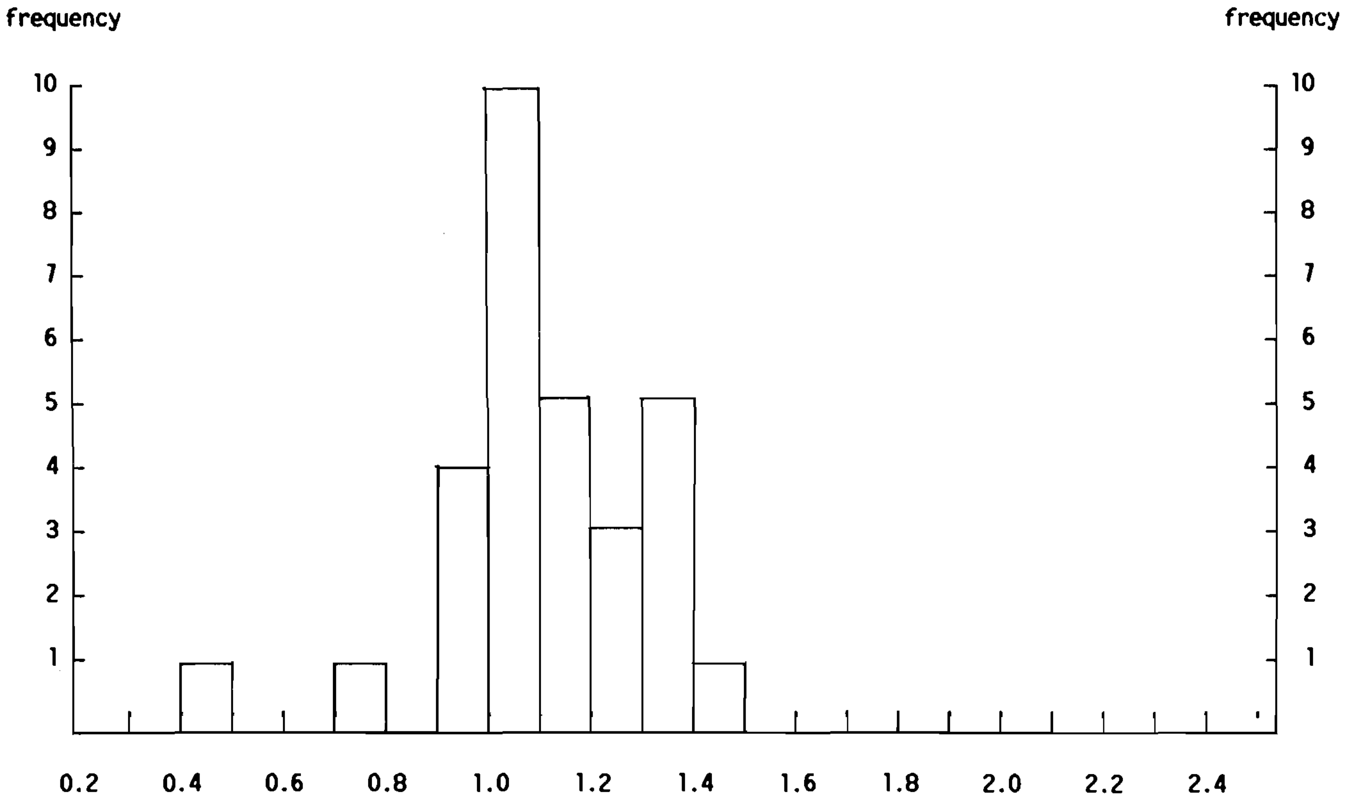

Figure 3(b): Distribution of the S_i/SNC_i Ratios for the A/US Hedge Settlement Rate

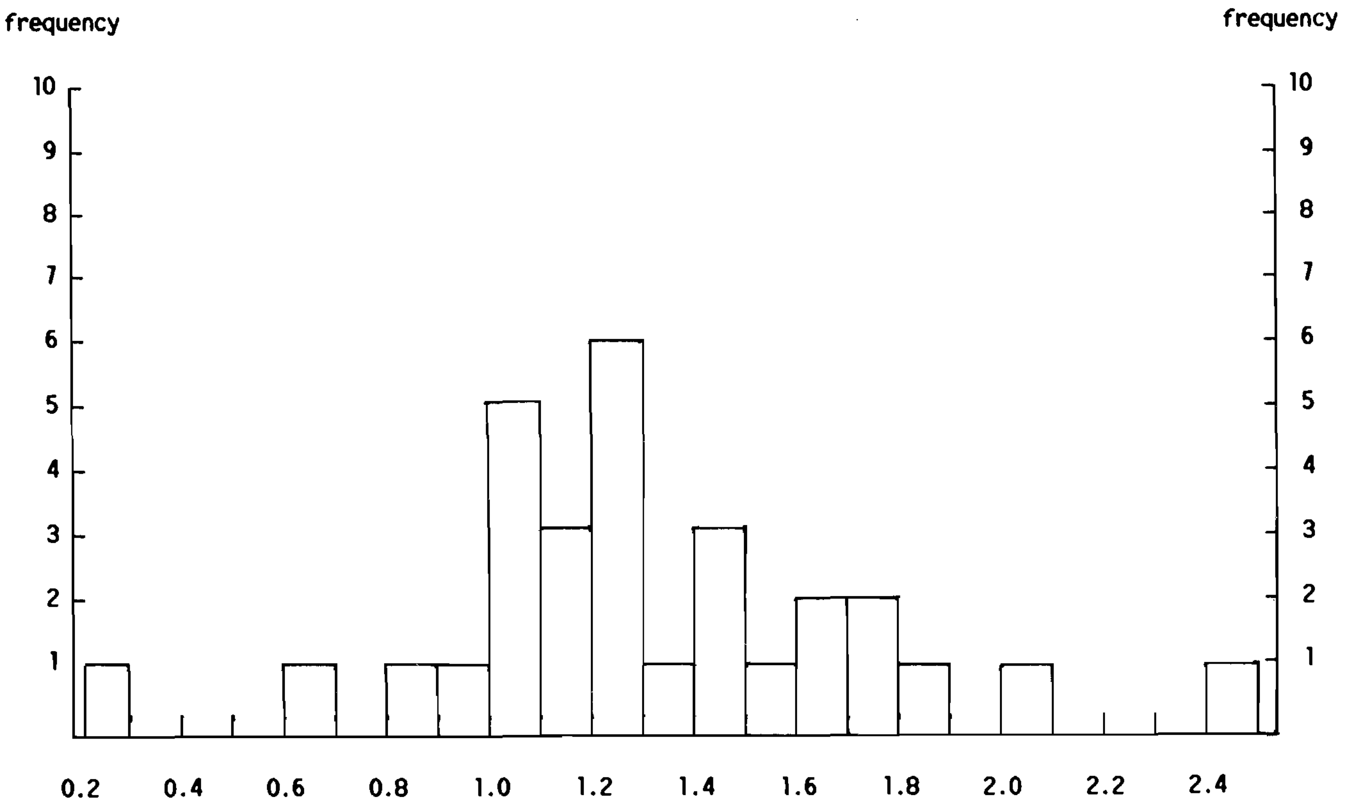

settlement rate. Further, only four out of the 30 forecasters were able to produce forecasts with a lower mean square error than the no change forecasts. Three of these four dealers were also members of the group of six who had an A_i/ANC_i of less than unity. However, all three were in the survey for less than ten weeks.

The distributions of the A_i/ANC_i and S_i/SNC_i ratios for the US$/Yen forecasts are shown in Figures 4(a) and 4(b) respectively. While both distributions are again skewed to the right of unity, the number of forecasters producing forecasts superior to the no change model is higher than was the case for the A/US forecasts. Of the 30 forecasters, 12 were able to produce forecasts with a lower MAE than that produced by the no change model, while 11 were able to produce forecasts with a lower MSE. Only six dealers fell into both categories. Of these six, four were in the survey for more than ten weeks.

While the A/US and US$/Yen group mean forecasts were able to produce a lower MAE than the no change forecasts, this was not the case for the bulk of the individual dealers. Clearly, most forecasters provided less accurate forecasts than the group mean forecasts. That this is the case can be seen from Figures 5(a) and 5(b) which show the distribution of the A/US A_i/AGM_i and S_i/SGM_i ratios respectively and Figures 6(a) and 6(b) which show the distributions of the same ratios for the US$/Yen forecasts.

Figure 5(a) shows that of the 30 forecasters only four were able to provide A/US forecasts with a lower MAE than the group mean forecasts while Figure 5(b) shows that only seven of the 30 forecasters were able to beat the group mean forecasts using the MSE criterion. In this case, the seven include the four with an A_i/AGM_i ratio of less than unity. However, all four were in for less than ten weeks. The picture is much the same for the US$/Yen forecasts with only four of the 30 forecasters beating the group mean forecasts on the MAE criterion, and five beating the group mean forecasts on the MSE criterion. The overlap between these two groups contained three dealers, one of whom was in the sample for more than ten weeks.

The fact that most forecasters find it difficult to perform more accurately than the group mean is hardly surprising given that in a series of group averages errors of opposite sign tend to cancel out.

Figure 4(a): Distribution of the A_i/ANC_i Ratios for the US$/Yen Close Forecasts

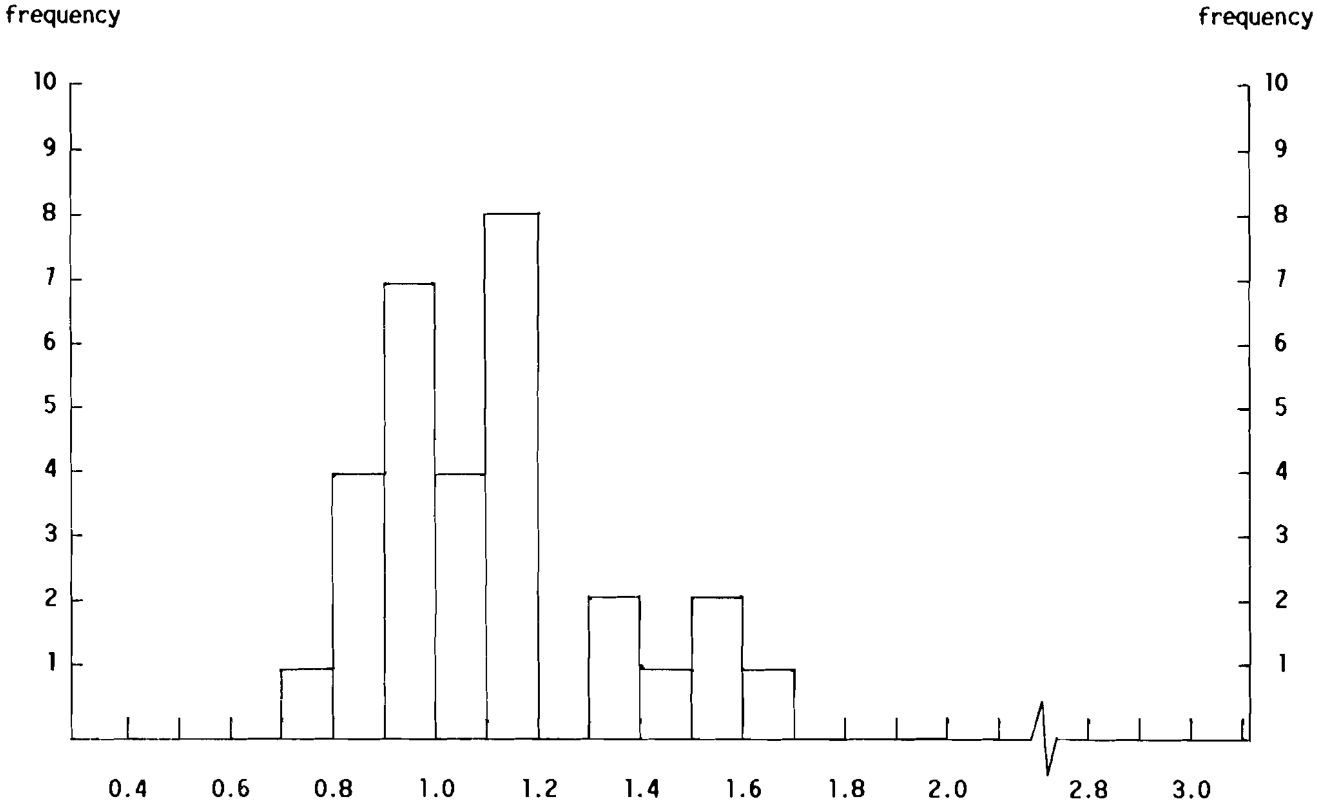

Figure 4(b): Distribution of the S_i/SNC_i Ratios for the US$/Yen Close Forecasts

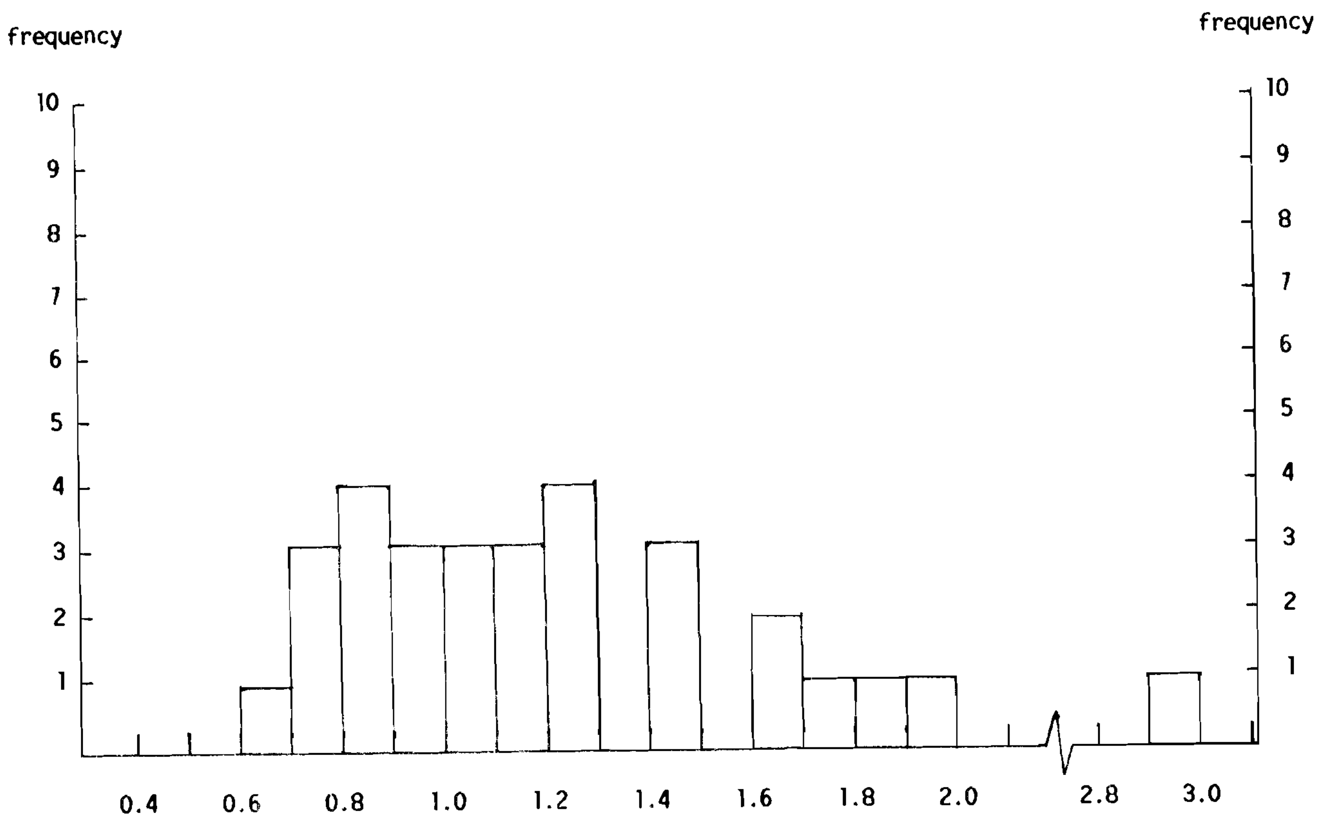

Figure 5(a): Distribution of the A_i/AGM_i Ratios for the A/US Hedge Settlement Rate

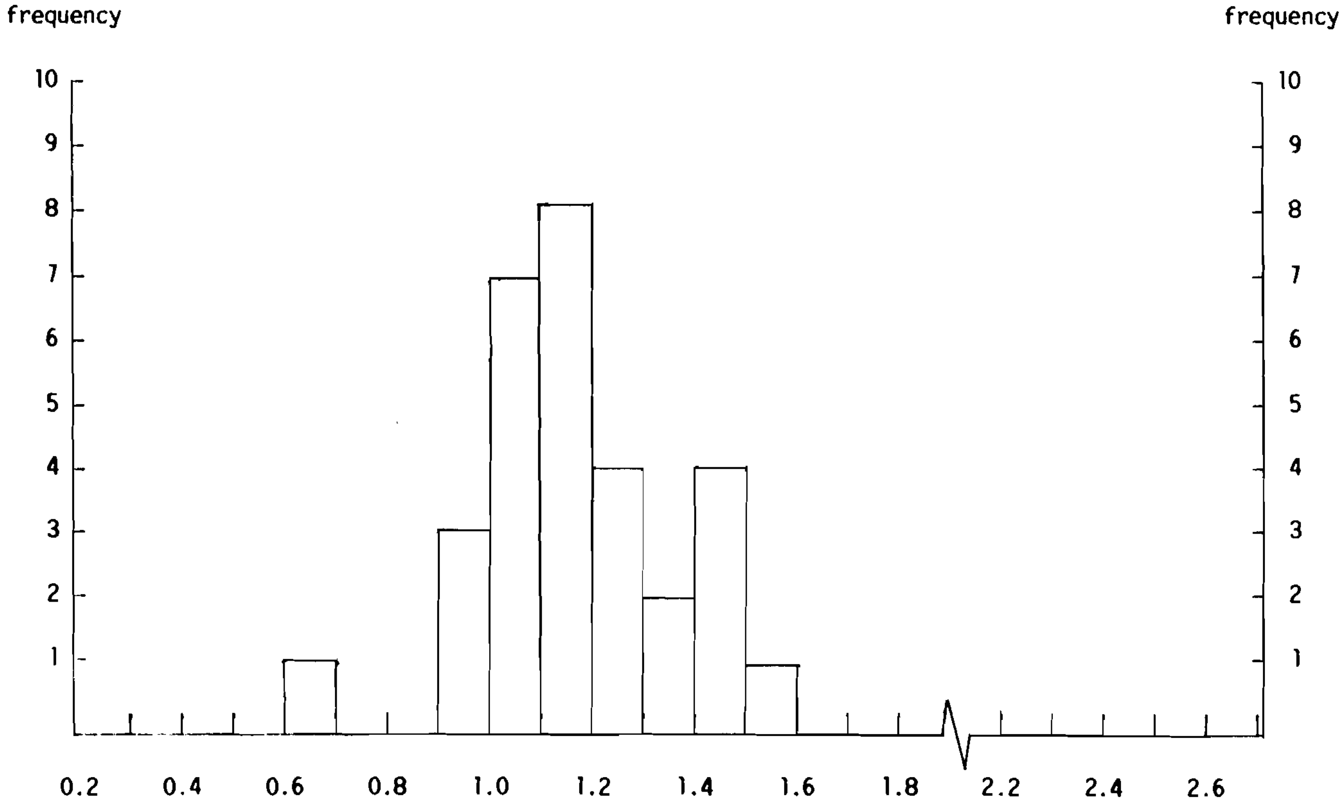

Figure 5(b): Distribution of the S_i/SGM_i Ratios for the A/US Hedge Settlement Rate

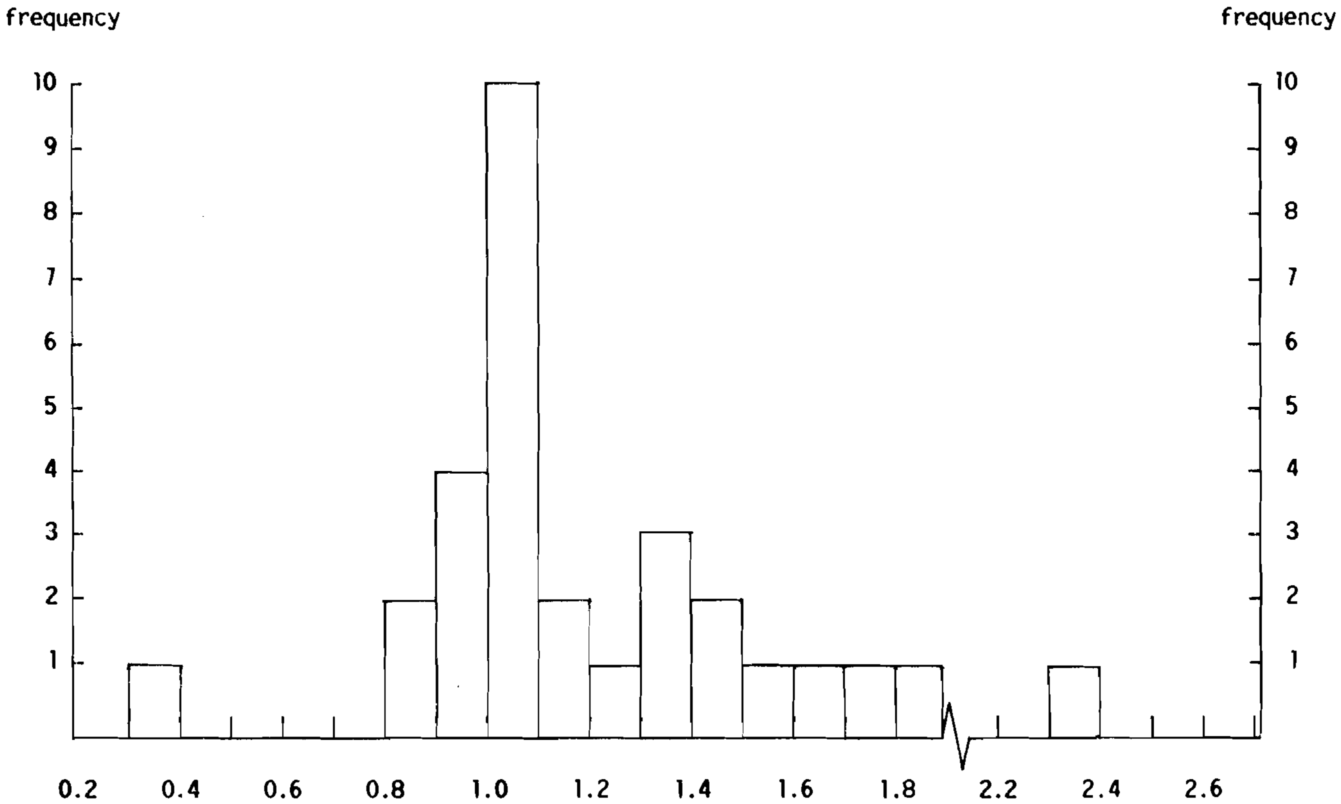

Figure 6(a): Distribution of the A_i/AGM_i Ratios for the US$/Yen Tokyo Close Forecasts

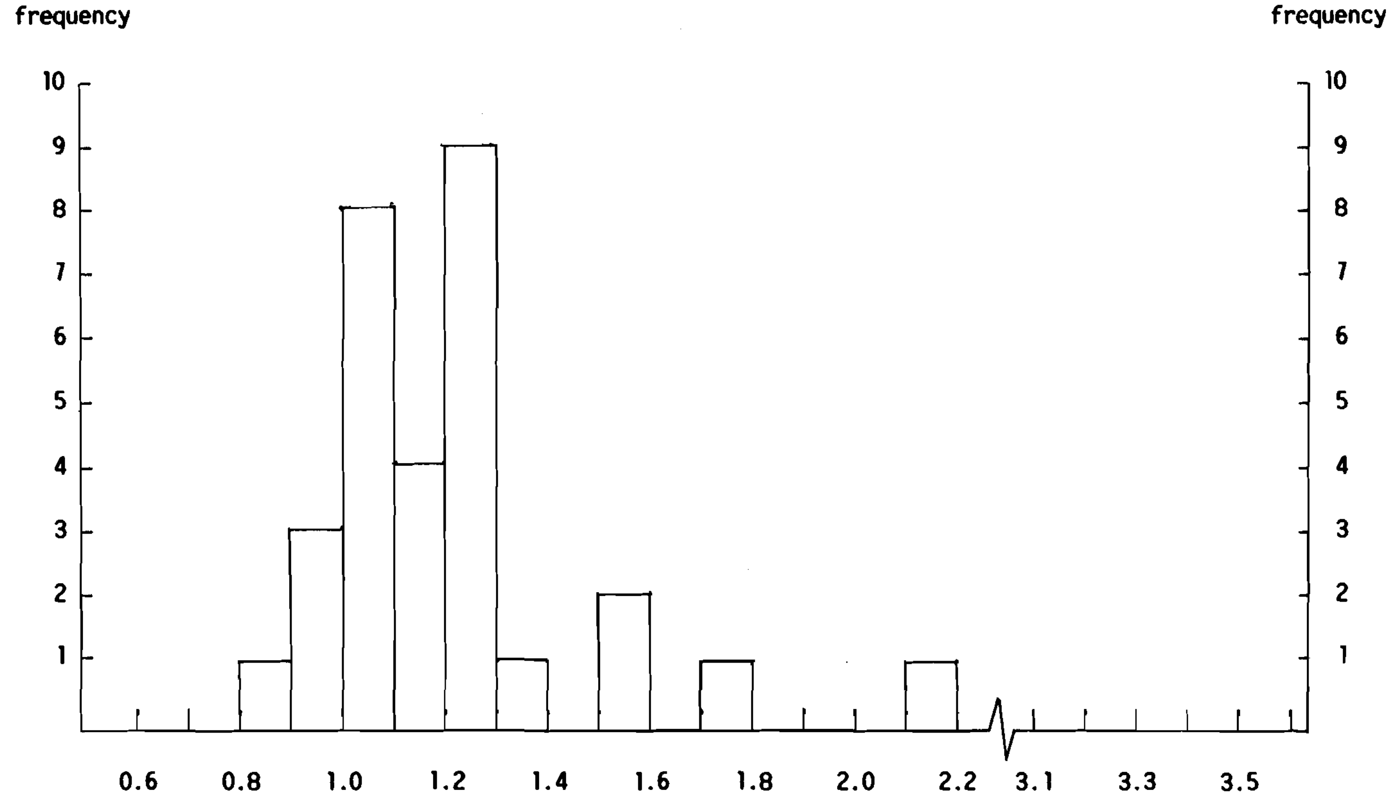

Figure 6(b): Distribution of the S_i/SGM_i Ratios for the US$/Yen Tokyo Close Forecasts

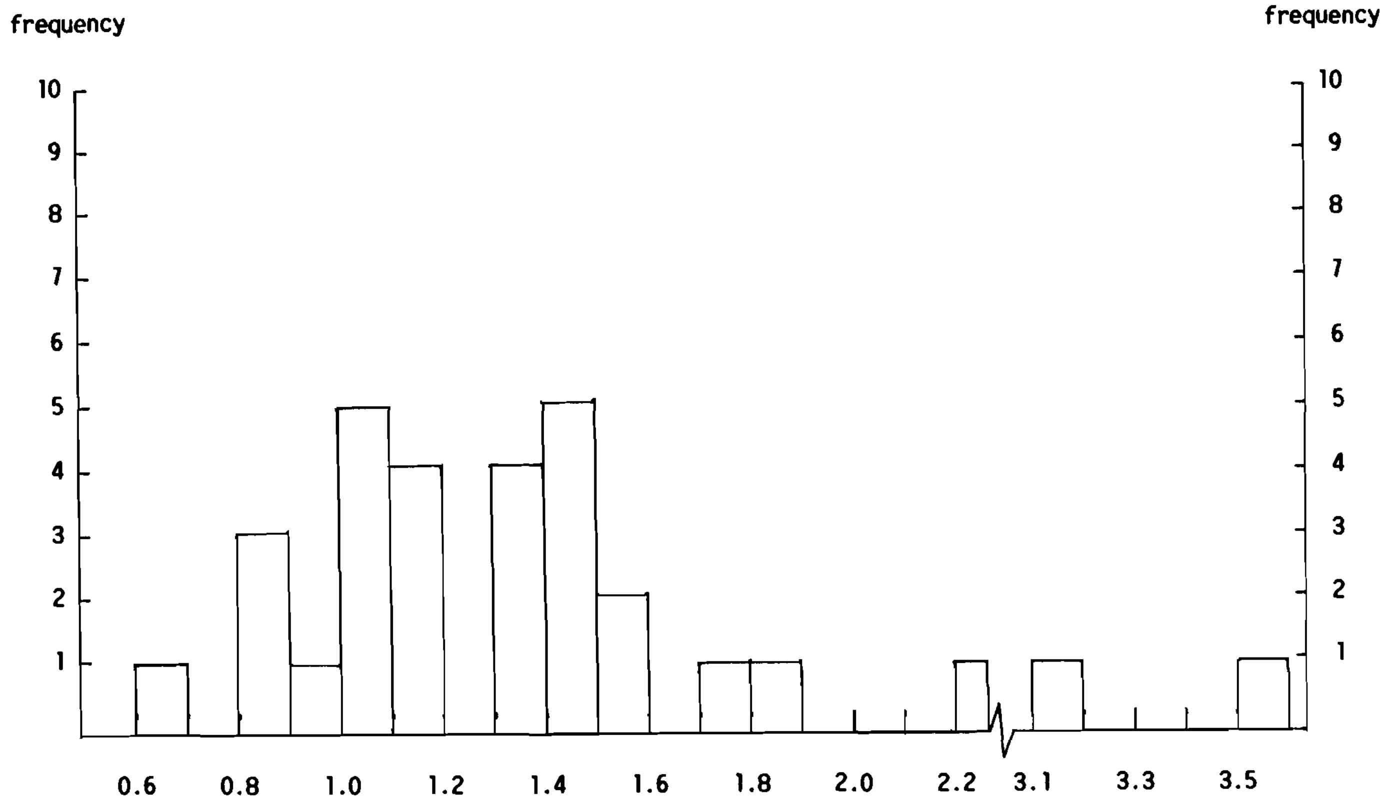

Table 6 reports the Spearman rank correlation coefficients between the ratios for the A/US and US$/Yen forecasts.

Table 6
Spearman Rank Correlations Coefficients

		A/US		US$/Yen	
		A_i/AGM_i	S_i/SGM_i	A_i/AGM_i	S_i/SGM_i
A/US	A_i/AGM_i	1.0000	..	..	..
	S_i/SGM_i	0.8861 (0.0001)	1.0000	..	..
US$/Yen	A_i/AGM_i	-0.1755 (0.3535)	-0.1871 (0.3222)	1.0000	..
	S_i/SGM_i	0.0687 (0.7181)	0.0523 (0.7838)	0.7526 (0.0001)	1.0000

		A/US		US$/Yen	
		A_i/ANC_i	S_i/SNC_i	A_i/ANC_i	S_i/SNC_i
A/US	A_i/ANC_i	1.0000	..	..	..
	S_i/SNC_i	0.8438 (0.0001)	1.0000	..	..
US$/Yen	A_i/ANC_i	-0.2256 (0.2306)	-0.1845 (0.3291)	1.0000	..
	S_i/SNC_i	0.0879 (0.6443)	0.0652 (0.7322)	0.7608 (0.0001)	1.0000

Note: The figures in brackets are the marginal significance levels - i.e., the minimum level of significance required to reject the null hypothesis that the column variable is uncorrelated with the row variable.

As expected, forecasters who do well using the MAE criterion also do well using the MSE criterion. The correlations, however, suggest that forecasters who provide the best A/US forecasts do not provide the best US$/Yen forecasts. In fact, two of the four correlations between the same ratios for the different exchange rates are negative, although insignificant at the 10 per cent level of significance (i.e., the marginal significance levels are greater than .01). No forecaster was able to provide forecasts superior to the group mean forecasts in both markets on both the MAE and MSE criteria. Two forecaster were, however, able to produce forecasts with a lower MSE than the group mean forecasts for both the A/US and the US$/Yen forecasts. However, of these two forecasters only one was able to dominate the no change model (on both the MSE and MAE criteria) in the prevision of A/US forecasts. This forecaster, however, provided less than ten forecasts.

7. Conclusions

The evidence presented in this paper suggests, on the basis of published data, forecasters of the exchange rate have much about which to be modest. The vast majority of individuals providing the one week ahead A/US hedge settlement rate forecasts published in the Australian Financial Review would have produced forecasts with a lower mean absolute error (or mean square error) had they supplied forecasts based on a simple rule such as no change. Notwithstanding this poor performance of the individual forecasts, the group mean forecast series performed relatively well, predicting the correct directional movement approximately two thirds of the time and outperforming all six benchmark models on the mean absolute error criterion. The group mean series, however, performed less well on the mean square error criterion having a larger mean square error than three of the benchmark models. Of these benchmark models, the no change model and the time series model with one lag (the restricted model) produced the lowest errors.

Like the group mean forecasts of the A/US, the group mean forecasts of the US$/Yen rate produced a lower mean absolute error than the no change model. However, unlike the A/US forecasts the group mean US$/Yen forecasts also outperformed the no change model on the mean square error criterion. Further, the average standard deviation of individual weekly forecasts was greater for the A/US forecasts than for the US$/Yen forecasts. These results suggest that over the sample period the market for Australian dollars may have been characterised by greater uncertainty than was the US$/Yen market.

REFERENCES

Hafer, R.W. (1984). Comparing Time Series and Survey Forecasts of Weekly Changes in Money: a Methodological Note. Journal of Finance Vol. 39 pp.1207-1213.

Meese, R.A. and Rogoff, K. (1983). Empirical Exchange Rate Models of the Seventies: Do they Fit out of Sample? Journal of International Economics Vol. 14 p.3-24.

Zarnowitz, V. (1982a). Expectations and Forecasts From Business Outlook Surveys National Bureau of Economic Research Working Paper No. 845.

Zarnowitz, V. (1982b). The Accuracy of Individual and Group Forecasts from Business Outlook Surveys. National Bureau of Economic Research Working Paper No. 1053.