

NBER WORKING PAPER SERIES

**HETEROGENEOUS EXPECTATIONS
AND TESTS OF EFFICIENCY IN
THE YEN/DOLLAR FORWARD FOREIGN
EXCHANGE RATE MARKET**

**Graham Elliott
Takatoshi Ito**

Working Paper 5376

**NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 1995**

The research project for this paper started when the first author was a graduate student and the second author a visiting professor at Harvard University. The earlier draft was circulated under the title "Heterogeneous Expectations and Tests of Rational Expectations." The current version was presented at the Econometric Society 7th World Congress in August 1995 in Tokyo. We would like to thank David Backus, Ralph Bryant, Menzie Chinn, Jeffrey Frankel, Wen-Ling Lin, and Alan Timmerman for comments. We thank Elena Pesavento for her excellent research assistance. All errors are our own. The views expressed here are those of the authors and do not necessarily reflect those of the institutions the authors have been affiliated with in the past and present. This paper is part of NBER's research program in International Finance and Macroeconomics. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

© 1995 by Graham Elliott and Takatoshi Ito. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

HETEROGENEOUS EXPECTATIONS
AND TESTS OF EFFICIENCY IN
THE YEN/DOLLAR FORWARD FOREIGN
EXCHANGE RATE MARKET

ABSTRACT

This paper examines the efficiency of the forward yen/dollar market using micro survey data. We first argue that the conventional tests of efficiency (unbiasedness) of the forward rate or of the survey forecasts do not correspond directly to the zero-profit condition. Instead, we use the survey data to calculate directly potential profits of individual forecasters based on a natural trading rule. We find that although the survey data are not the best predictor of future spot rate in terms of typical mean square forecast error criteria, the survey data can be used to obtain on average positive profits. However, these profits are small and highly variable. We also examine profits generated by a trading rule using regression forecasts, where forward premium is an explanatory variable. These profits are also small and highly variable.

Graham Elliott
Department of Economics
University of California
9500 Gilman Drive
La Jolla, CA 92093-0508

Takatoshi Ito
Senior Advisor, IS 12-908
Research Department
International Monetary Fund
700 19th Street, NW
Washington, DC 20431
and NBER

1. Introduction

The purpose of this paper is two-fold. First, we will argue that an interpretation common in the literature linking biasedness (unbiasedness) of forward exchange rate with the existence (lack, resp.) of unexploited profit opportunity is quite misleading, especially in understanding the magnitude and riskiness of the revealed profit opportunity. An alternative measure, directly measuring potential profits based on a specific trading rule, is proposed. Second, we examine a survey data of exchange rate forecasts to test the extent to which individual forecasters would have made profits. By assuming a trading rule dependant on a sign relationship between the forward rate and the expected exchange rate, we can calculate profits that can be earned if the individuals bet on the direction of the exchange rate changes based on their subjective information. The existence of such a trading strategy will be both a further challenge to the "efficiency" of the market (nonexistence of excess profits beyond risk premium) and also evidence on seriousness (high quality) of the survey responses.

Many papers have been written on the biasedness test of the forward exchange rate, see Froot and Thaler (1990) for a survey. The conventional test asks whether the forward rate at period t for settlement k periods ahead, $f_{t,k}$, is an unbiased predictor of the ex post spot rate, s_{t+k} . If the hypothesis that $s_{t+k} - f_{t,k} = 0$ on average, is rejected, then the forward rate is said to be biased. Most papers find that the forward rate indeed is biased from the ex post spot rate. This is often taken to be a "puzzle," since a naïve betting on the forward would have produced profits. The biasedness is interpreted as evidence of (i) that the expectation was not rational; (ii) that risk premium was present; (iii) that a peso problem (small sample problem) was present; or (iv) a combination of the above.

First of all, note that the conventional test implicitly specifies a naïve trading rule, in such a way that the investor takes the same forward position all the time (always "buy

forward"). A rejection is neither a necessary nor sufficient condition for the existence of profitability. Suppose that the test failed to reject the unbiasedness hypothesis. It could still be possible that other information available to market participants help them predict better using a more sophisticated trading rule and produce economically significant profits. Suppose that the test did reject the unbiasedness. This still may mean that excess profits generated by this "inefficiency" are small and highly variable and thus do not really constitute economically significant evidence of an inefficient forward market.

One can think of more sophisticated trading rules, that are implied from well-known models in the literature on exchange rate behavior. Risk-neutral investors who believe that the spot exchange rate follows a random walk would take a forward short position whenever the forward rate is at a premium over the spot. We refer to this as the random walk trading strategy. Many researchers have found that the forward premium ($f_{t,k} - s_t$) is correlated with ex post forecast errors. This implies that some trading rule using information forward premium would produce more profits compared with the naïve strategy of always buying the forward contract. However, no one to our best knowledge has constructed this trading rule and calculated resulting profits in order to examine the magnitude of resulting profits.¹

In order to improve further on the power of the efficiency test, we propose to use the expectations survey data employed in Ito (1990, 1994). Using the data set, we improve the conventional test in two aspects. First, the expectations data gives us a better measure of the underlying economic model used by market participants, so better measures of potential profits can be identified. Second, the heterogeneous individual expectations data provides a large number of tests for a given sample period.

It has been widely recognized that agents in the market have diverse opinions--that is, heterogeneous expectations--about, say, the future spot rate. Otherwise, it would be

difficult to explain the fact that large volumes of transactions take place every second. Individual survey data also shows that agents display persistent heterogeneity in their forecasts (Ito 1990). We extend these results, and motivate looking at each firm's individual profit strategy rather than, say, profits based on observing only the forward market or average expectations. Individual strategies are to some extent independent rather than being multiple observations of the same excess profits. This gives us substantially more information to test the standard propositions of a zero profit condition.

One of the common criticisms of the use of the survey expectation data is that respondents to questionnaire may not be serious in answering questions. The calculation of how much profits and losses presented later indicate the seriousness of these expectations data. We find that although the expectations data is outperformed in terms of forecast error variance by the random walk model, most of the forecasts outperform the random walk model in terms of profitability.

The actual survey data employed in this study is the JCIF dataset described in Ito (1990), updated so that the data is fortnightly from May 1985 until the end of May 1995. The data consists of forecasts of the future spot yen/dollar exchange rate for 44 companies at the one, three and sixth month horizons. This data set is unique in that regular membership has been maintained for ten years, with few missing observations, except for two companies that are dropped because of too frequent missing observations. See Ito (1990) for further details of this data.²

Profits calculated from the trading rule implied from the survey data, are compared with profits generated from the trading rule based on the random walk model and that on the model using information on forward premium.³

The empirical findings of the paper can be summarized as follows. Whilst some of the firms in the survey would have made statistically significant profits following a simple

trading rule using their respective forecasts, such profits are small in magnitude and highly variable. Profits from the survey-based trading strategy are in general greater than the naïve strategy (always buy forward), or the random walk. However, profits are comparable to the forward-premium based trading strategy. The evidence does not appear to be greatly inconsistent with efficiency of the forward exchange rate market.

The rest of the paper is organized as follows. Section 2 examines the relationship between profits from some trading rule based on the relative position of the forecasts and the forward rate and the usual tests of unbiasedness. Section 3 examines the expectations data directly, giving the heterogeneity results and forecast performance results. Section 4 motivates and details the construction of profits used in this study, and Section 5 examines these profit estimates. The final section concludes.

2. Predicable profits and the naïve trading rule

One of the most enduring puzzles in the international finance literature is the statistical finding of the biasedness of the forward foreign exchange rate as a predictor of the future spot exchange rate. This finding translates to profitability of a naïve trading rule of always buying a currency forward and closing at the exchange rate at maturity in the spot rate market. This section derives the familiar implications of a properly functioning forward market, relating the two main approaches to deriving empirical tests. We then consider the advantages and drawbacks of these tests.

a) Euler equation approach

Hansen and Hodrick (1983) show that in a Lucas (1982) economy, the Euler equation for the representative risk neutral agent can be written

$$E_t [Q_{t+k} (S_{t+k} - F_{t,k})] = 0 \tag{1}$$

where Q_{t+k} is the marginal rate of substitution of money between periods t and $t+k$, S_t is the spot exchange rate in period t , and $F_{t,k}$ is the forward exchange rate set at time t for delivery in period $t+k$. Backus et al. (1994) derive a similar equation to (1) where Q_{t+k} is a pricing kernel. In both papers, the difference between the future exchange rate and the forward rate is viewed as the profit an agent would earn if they were to buy the foreign currency forward and close the position when delivery is due by selling the foreign currency at the future spot exchange rate.

Hansen and Hodrick (1983) further show, conditioning on the dataset and using assumptions of log normality, that this Euler equation can be used to derive a testable relationship between the spot and forward exchange rates, in particular that

$$E_t [s_{t+k} | Z_t] - f_{t,k} = E_t [s_{t+k}^2 | Z_t] + E_t [s_{t+k} q_{t+k} | Z_t] \quad (2)$$

where the lower case letters refer to natural logarithms of the previously defined upper case variables. This relationship says that expected profits from the naïve trading rule should not exceed on average the conditional variance of the spot exchange rate and the conditional covariance of the future exchange rate and the log of the marginal rate of substitution of money, and further should be uncorrelated with information in Z_t . Given that the conditional volatility of the exchange rate and the conditional covariance term are very small, this is often taken to mean that the left hand side of (2) should not be predictable with time t dated information, which can be interpreted as zero expected profits from the naïve trading rule.

b) No arbitrage approach

An alternate derivation relies on using the uncovered interest parity condition (UIP) along with the covered interest parity condition (CIP) (e.g. see McCallum (1994)). Here, it is assumed that CIP holds (it has been argued elsewhere that it is through this

relationship that market participants price forward contracts in the market), and that the log approximation of UIP holds.⁴ Combining these relationships yields the condition $E_t(s_{t+k} - f_{t,k}) = 0$.

This result is of course similar to that derived from the Euler equation approach. It is again a zero profitability condition, if there were expected profits that could be earned from buying the forward contract and settling at delivery we would expect that participants in the market would enter the market arbitraging the anomaly away.

3. Design of empirical tests for the zero-profit condition

a) Conventional tests

After replacing the expected future spot rate by the ex post future exchange rate (and invoking rational expectations), the orthogonality condition implied by both of the approaches has been tested often in the literature, with empirical results at odds with the theoretical null hypothesis. Typical tests of this relationship include directly examining the unconditional mean of such profits, to examine whether or not excess profits could have been made. This corresponds to testing that the naïve trading rule (always "buy forward") profits are zero.

There are two caveats to this line of reasoning. First, a caveat to the implication that non zero profits here imply a market inefficiency is that the measure $(s_{t+k} - f_{t,k})$ underestimates potential profits that could be earned from the forward market in practice. The first caveat is that such a buy and hold strategy represents only one possible strategy in speculating against the forward market to make profits. This measure is popular due to it being the strategy most easily identified by the econometrician. This suggests that, by examining only this simple rule, that the econometrician potentially underestimates the types of profits that would be available to the more sophisticated investors actually active in the market. In particular, when an agent uses some model to make their own forecast,

then potentially they will make profits in excess of that measured by the econometrician using this naïve rule. An example clarifies this statement. Suppose it so happened that $(s_{t+1} - f_{t,1}) = 1$ in even numbered months and -1 in odd numbered months. Average profits for the year are zero for the naïve strategy. A forecaster with perfect foresight, correctly picking the sign of the forward error, always makes a sure profit of 1 unit per month.

A second caveat, unlike the first caveat, suggest that use of the conventional test of a buy-forward rule (testing whether $\{f_{t,k} - s_{t+k}\}$ was statistically significant from zero) as the profit measure may overestimate such potential profits. If the profit measure turns out to be negative, it is interpreted that the reverse (sell-forward) would have made profits. Thus, the conventional test requires that the investor to know beforehand which currency to keep buying forward (and squaring at the respective maturity). Thus, this amounts to knowing which currency will appreciate over the period. Given that the exchange rate is often modelled by a random walk, correctly guessing the sign ex ante is difficult.

b) Trading rule based on the random walk belief

Responding to the caveats cited above, those economists who are familiar with empirical work in the field may respond by saying that "natural" trading rules, which are more sophisticated than the conventional measure, exist. In the literature, it is established that the random walk model of exchange rates are hard to beat even with sophisticated econometric models. (See Meese and Rogoff (1983) and Diebold and Nason (1990).) The random-walk believer, who places an equal probability of appreciation or depreciation, would suggest a trading strategy that sell the currency when its forward is more appreciated than the spot (forward premium), and vice versa. The direction of bets, which stays the same in the conventional test, would now change whenever the forward premium changes the sign. Profits of a trading rule based on the random-walk belief are positive if the ex post spot rate did not appreciate (depreciate) as much as the forward, when the forward

premium is positive (negative, resp). We will construct the measure of such profits based on the random-walk trading rule. c) The regression based model

Alternatively, the conditional mean can be examined (i.e. the orthogonality hypothesis) by regressing ex post profits that would have been realized if this strategy had been employed on information known to the market at time t , i.e.

$$s_{t+k} - f_{t,k} = \beta_0 + \beta Z_t + u_t \quad (3)$$

with the null hypothesis that $\beta_0 = \beta = 0$, or simply $\beta = 0$ (as β_0 non zero may be capturing the effects of the conditional variances and covariances above, see Hodrick (1987)). See Hansen and Hodrick (1980), Cumby (1986), Bekaert and Hodrick (1992)). Such regressions are interesting as rejections of the null hypothesis would define a trading rule that will be more profitable than the naïve rule. However, as far as we are aware the size of such profits, constructed from feasible rules available to the market, has not been previously examined in the literature.

In particular, the following specification is frequently employed

$$s_{t+k} - s_t = \gamma_0 + \gamma (f_{t,k} - s_t) + u_{2t} \quad (4)$$

where the null hypothesis is $\gamma = 1$. This is a special case of the regression in (3) as when $Z_t = (f_{t,k} - s_t)$, then the regressions give numerically equivalent results, where $\hat{\beta} = \hat{\gamma} - 1$ and $\hat{\beta}_0 = \hat{\gamma}_0$. The later version has been more often tested, see Bilson (1981), Fama (1984), Bekaert and Hodrick (1994) and McCallum (1994)). This shows the direct relationship between unbiasedness tests and the zero profitability condition tests.

The finding that γ is different from one in this regression suggests that forecasts can be generated from the regression which will be useful for making profits speculating

against the forward market. We calculate such profits for comparison with the results presented later.

Two forms are employed, first (denoted SR-reg, for regressions with super rationality) we use the full sample to estimate γ_0 and γ , and second (denoted RL-reg, for rolling regressions) we use data only up until time t to obtain time varying estimates of these parameters estimated using data known at the time the forecast was to be made. The forecasts from SR-reg are included to show how much using the whole sample information helps - this regression clearly cannot be used to provide real time forecasts. The results of RL-reg are however available for the construction of a real time profit rule. These profit rules will also be examined in section 5.

The failure of the unbiasedness hypothesis (or zero profit condition) in their confrontation with the data has generated a large literature seeking an explanation. Early arguments, such as measurement errors, have been effectively ruled out (Bekaert and Hodrick (1994)). Arguments that the biases found are due to a risk premium have been considered only a partial explanation, as theoretical and empirically estimated risk premia appear to be too small to generate biases of a large enough size to account for those found empirically. This is true despite the fact that typically the profits from both the naïve trading rule and those implied by these regressions are found in section 5 to be modest and highly variable. Another plausible partial explanation, introduced by McCallum (1994), is that through CIP and UIP and the fact that interest rates are policy variables, the dynamics of the forward rate are dominated by the effects of monetary policy.

A caveat to this analysis is that although we may statistically reject the null hypothesis that the naïve profits $(s_{t+k} - f_{t,k})$ are orthogonal to information available to market participants, it may be that this information is not economically valuable in devising a profitable trading rule. Whilst rejection of the null hypothesis implies that better profit

rules can be constructed, the profits from such rules may turn out to be small and highly variable. Thus statistical rejection is not enough evidence on its own to gauge the alleged inefficiency of the forward market. (see Breen et al (1989) for a similar discussion of this in the context of stock market predictability). This paper attempts to measure such profits specifying a trading rule.

d) The role of expectations

An alternate strategy to identifying potential profits earned from speculating on the forward market would be to use the market participants forecasts to examine potential profits these participants could have earned, rather than the naïve trading rule profits. Similarly, a model of exchange rate determination would provide forecasts that would play the same role. In the following section, we examine forecasts of the future yen/dollar rate by 42 Japanese companies, and in a later section use these to obtain a potentially more precise measure of excess profits that could have been earned from the forward foreign exchange rate market. As will be seen, these profit calculations address the caveats mentioned above. Firstly, all of the information required to calculate the profit rules comes from information known to the participants. Secondly, as we believe that participants have some model of exchange rates, of which the econometrician observes only the forecast, it is conjectured that calculated profits represent a better approximation of the types of profits market players can earn from speculating against the forward market. We examine these excess profits, attempting to explain them by variables we expect to be correlated with a risk premium (in which cases deviations from zero profits are still consistent with the efficiency hypothesis), or due to information known at the time the bets were being undertaken (which would not be consistent with the efficiency hypothesis).

First, however, we examine the forecast data and compare this to a random walk model of the exchange rate.

3. Expectations data and heterogeneity

This section examines the features of the expectations data and in particular previous results of heterogeneity in forecasts, extending the results of Ito (1990) to a longer data set. The findings of heterogeneity motivate the use of individual data in subsequent analysis, as there is apparently information special to each forecast. We also examine the forecasting performance of the survey respondents and simple exchange rate forecasting models.

a) Heterogeneity

Typically, models of the forward market (e.g. Lucas (1982)) assume the existence of a representative agent with some information set, and imply that these identical market participants have equivalent forecasts of future spot rates. However, it may be that agents deviate systematically from each other in their forecast of the future exchange rate due to different information sets, and different models of the exchange rate. Ito (1990), using an earlier version of the data employed here, found that this was indeed the case, that forecasters of the yen/dollar exchange rate systematically deviated from the cross-sectional average forecast. This was tested by examining the mean of this deviation, and regressing on a constant. Specifically, equation (3) of Ito (1990) was

$$s_{t,k}^i - \bar{s}_{t,k} = g_i + u_{3,t}^i \quad (5)$$

where under the null hypothesis of homogeneous information sets, the mean g_i should be zero. If this mean is significantly different from zero, then on average the forecaster predicted greater or smaller depreciations than the group as a whole, hence a finding of heterogeneity.

The result of Ito (1990) that expectations appear to be heterogeneous, is upheld in the updated dataset employed here. We examine the deviations of forecasts from the average of all firms as given in equation (5), and find that many of these unconditional means are

significantly different from zero. This is true for forecasts at all horizons. This indicates that firms consistently deviated from other firms in their setting of forecasts. The results are shown in Table 1, where the maximum and minimum deviations (in terms of mean deviation) from the null hypothesis and their t statistics are reported as a summary of the results. The distribution of the t statistics testing the hypothesis for all of the companies is reported in Figure 1. The results show that 18 of the 42 firms have forecasts at the one month horizon which are (statistically) significantly different from the sample average. At the three and six month horizons, 20 deviate from the sample average (i.e. are statistically significant at the 95 percent confidence level).⁵ Many of these t statistics are quite large, suggesting that these results are strongly supported by the data.⁶ Also included is the deviation from the forward rate from the average expectation. This mean is statistically significant at the 95 percent level for the 1 month horizon, but not at the three or six month horizon. This may only indicate a risk premium with a non zero mean.

The finding of heterogeneity has implications for the analysis of the forward exchange rate market. Typically, in writing down models of the market, we assume that all agents are identical (representative agents), and hence derive the equilibrium relationships for a representative agent. This would involve a single forecast consistent with the model. This is true of the use of the Lucas (1982) model. From these results, it is suggested that this is not an appropriate assumption for this market. A clear reconciliation would be to allow for information sets for each agent to be different, yielding different private forecasts.

b) Forecast performance

An obvious question to ask when examining forecast data is how good the forecasts are. This subsection analyses this question, with a view to provide evidence that this data

is useful in examining the foreign exchange rate market. One measure of forecasts is to examine the standard error of the ex post expectation error, and relate this to some baseline model.⁷ With the disappointing empirical results of typical theoretical models of the exchange rate, we compare here to the random walk model forecast error (this model is often used as a baseline for forecast performance, e.g. see Diebold and Nason (1990)). The results are presented in Table 2 (and the distribution of the full results in Figure 2), along with the average forecast error over the time period for both the cross section average forecast, the forward rate as a forecast, and the forecasts SR-reg and RL-reg described earlier. Each result is divided by the forecast error of the random walk model (which is equivalent to the conditional variance of the spot rate). A value of 1 indicates equivalent performance to the random walk, lower than one a better performance and higher than one a worse performance than the random walk.

The results show that all of the individual firm forecasts have a greater variance of the forecast error than the random walk model, and also greater variance than the forward market. The forecasts from RL-reg have a lower forecast variance than nearly all of the individual forecasters at the one month horizon, but is outperformed by this criterion at the 3 month horizon by 12 firms and at the 6 month horizon by 32 firms. These results tell us that much of the variation in the forecast errors are due to individual heterogeneity, as this variation is not present in the cross sectional average forecast error variance. Also note that the forecasts generated with RL-reg, the typical model employed to demonstrate forward market unbiasedness, are outperformed by the forward rate as a predictor of the future spot rate at all horizons. The superior performance of SR-reg can be dismissed as it uses information from the full sample to make the forecasts.

Contrary to an apparent impression from the result, it is not straightforward that these results indicate that the survey data is outforecasted by the random walk model, or that the survey data is of poor quality. The loss function is of a very special form here, trading off a larger bias for smaller variance equally. It is not clear that this is the correct loss function to attach to the market participants.

Leitch and Tanner (1991) show an example where there is no clear relationship between profits earned by forecasts (i.e. the economic value of the forecasts) and performance measure such as mean square error (in their case forecasts on Treasury Bills). The authors use these results to show why profit maximizing firms will be willing to pay for such forecasts even though they are outperformed in terms of mean square error criteria by simple ARIMA models. Thus it may well be the case that the forecast data here is outperformed by the random walk model, but they are more valuable in terms of earning profits. This will turn out to be the case.

Let us construct an example to illustrate the point on caveats to the conventional test and the random walk trading rule. Consider the situation which confronted a trader on September 10, 1986. The spot rate was at 155.50 yen/dollar after a long spell of yen appreciation which had started early 1985. The 6 month forward rate was 154.69, suggesting marginal further appreciation of the yen. The naïve strategy, assumed in the conventional test, buys the yen forward (or sells the yen forward, depending on the direction picked by the econometrician), regardless of any other information. The random-walk believer, would have the 6 month forecast equal to the contemporary spot rate, 155.50, would sell the yen forward (buy the dollar forward). The actual (ex post) spot rate six months later turned out to be 153.54, a larger appreciation than indicated by the forward market. Those who had bought the yen forward made a profit and those who had sold the yen forward, including the random-walk believer, lost money. In order to see that

this is not a contrived example, profits from such a strategy will be examined in a later section. The random walk model gives a forecast error of less than 2 yen. One of the respondents predicted that the yen would appreciate to 145 Y/\$, a prediction error of 8.5 yen. Yet the prediction of the random walk model used as a trading strategy would have resulted in a loss in the market, whereas the respondent would have made a profit.

Also reported in Table 2 are the proportion of times the forecast correctly forecasts the direction of the spot market. Most of the numbers are close to 50%, however the forecasters tend to beat the regression based forecasts for this measure. The best predictors of direction at all horizons were the individual forecasters, whilst the predictions generated from RL-reg were all below 50%. Indeed these forecasts outpredicted only 4 of the companies in the 1-month horizon, 5 at 3-month, and 4 at 6-month. Notice that the forward rate was slightly above 50% at the longer horizons. These results further indicated that the forecast data contains relevant information.

Secondly, we can use regressions such as those in Froot and Frankel (1989) to evaluate the forecasts.⁸ These regressions derive from regressions such as in equation (4) above, where the forecasting performance of the expectations are examined. This involves running a regression of the ex post depreciation on the forecast depreciation, to examine whether or not this is an unbiased predictor. i.e the regression is

$$s_{t+k} - s_t = \beta_0 + \beta (s_{t,k}^i - s_t) + u_{t+k} \quad (6)$$

where $s_{t,k}^i$ is the value of the survey forecast for s_{t+k} , taken at time t by firm i .

Such results are contained in Table 3 for the maximum and minimum (of the t statistics) for each of the firms (at the one, three and six month horizons), with results also for the cross sectional average, RL-reg and the forward premium in place of the expectation. The individual results are also graphed in Figure 3. These results are

surprising from the point of view of the null hypothesis, that is $\beta=1$. For the expectation data these estimates are clustered around zero and are significantly different from one. The usual result for the forward premium is also seen from these results, the point estimate is negative rather than one. The results of the regression based forecasts accord with those of the individual forecasts.

Elliott (1994) argues that if the stochastic process followed by the spot exchange rate is close to a random walk (in the sense that there is a near unit root and not much dynamics in the spot exchange rate),⁹ then the small sample distribution of the estimate of $\hat{\beta}$ will be around zero, even when the forecasts are optimal. The reason is that most of the signal is removed when looking at the changes in exchange rates. The intuition can be seen by starting with the assumption that the spot rate follows a random walk, and hence the change in the spot rate is a martingale difference sequence with respect to previously dated information. Thus, the right hand side of (6) is unforecastable, we would not expect $\beta=1$ here. In this case, all of the signal, the expected depreciation, is differenced out of the independent variable, and the independent variable is all random error. Similarly, the dependant variable also has the signal differenced out, so it to is all forecast error and no forecast. In the more realistic case of a near random walk, there will be some signal (expected depreciation) mixed in with the random errors. In this case the signal will also be dominated by random error in small samples, as the random error will dominate the signal in the regression and the estimates of β will also be distributed around zero. Thus these results provide neither evidence for or against the expectation data. It is clear from figure 3, however, that there is a systematic difference between the forecast data and the forward data (in each case the outlier in these figures is the coefficient from the regression of the ex post change in the exchange rate on the forward premium).

4. Trading rules

By observing the forecasts of market participants directly, as we do with the JCIF dataset, we are able to examine the potential profits of agents participating in this market directly. This can be an improvement in our examination on the efficiency of the market over the profits from the naïve trading rule as we would expect that such a naïve strategy would not be in place if participants are actively making forecasts as to the future direction of the exchange rate. In the calculation of potential profits, we assume that each agent bets in every single period that their forecast is correct, and that the bets are of the same finite magnitude each period.¹⁰ Thus, if the agent believes that the forward rate undervalues the future level of the spot exchange rate, that agent will take a long position forward. They will take a short position if they believe that the forward rate overvalues the future exchange rate. If the individual firm forecasts coincide with the forward rate then the individual does not trade in the forward market. Thus, under the assumption that if $s_{t,k}^i = f_{t,k}$ that the agent does not trade, profits for the i th company for the period over t to $t+k$ is given by

$$\Pi_{t,k}^i = I(s_{t,k}^i > f_{t,k}) (S_{t+k} - F_{t,k}) + I(s_{t,k}^i < f_{t,k}) (F_{t,k} - S_{t+k}) \quad (7)$$

where $I(\cdot)$ is an indicator function yielding one if the statement inside the bracket is true and zero otherwise. This can be rewritten as

$$\Pi_{t,k}^i = (2 I(s_{t,k}^i > f_{t,k}) - 1) (S_{t+k} - F_{t,k}) \quad (8)$$

We see that the calculated profits are related to the profits in the conventional trading rule, and will be equal to these only if the agent always believes that the forward rate undervalues the future spot exchange rate. The survey-expectation profits calculated in equation (8) equal profits under the random walk trading strategy only when the individual

forecasts are always the current spot exchange rate; and the survey-expectation profits become the regression-based (RL-reg) profits only when the individual forecasts are identical with the forecasts produced by the rolling regression.

It could be argued that the rule of the same sized bet in each period is unrealistic. One way to interpret this profitability is to consider it as profitability at the margin, i.e. we assume that firms hold positions in the two currencies, and that the profit calculations indicate the extra profits these firms could attain at the margin by increasing their exposure in the currency they believe is undervalued by the market through the forward exchange rate. An informal argument can be as follows: the forward market correctly prices the future spot if participants are already at their most profitable position (given their utility functions tradeoff between return and risk), implying that the profit rules we construct should not allow these firms to increase profits by increasing exposure in one currency vis a vis the other.

The heterogeneity of expectations allows us to examine not just a single strategy but instead the strategy of each of these firms, allowing us to gather much more evidence as to the possibility that such profits can be made in the forward market. To embed this into an Euler equation as has been previously examined, we have

$$E_t [Q_{t+k} \Pi_{t,k}^i] = 0 \quad (9)$$

where the i subscripts denote the results that profit rates will differ due to the heterogeneous expectations. We have such an Euler equation for each of the profit strategies, i.e. one for each firm. Substituting for profits using equation (8), and noting that the arguments of the indicator function are all known variables to the agent at time t and so can be conditioned out, then we can write this as

$$(2 I(s_{t,k}^i > f_{t,k}) - 1) E_t^i [Q_{t+k} (S_{t+k} - F_{t,k})] = 0 \quad (10)$$

which is equivalent to equation (1) above except for the multiplicative term at the front and the conditioning on the dataset that pertains to the company (i.e. includes their forecast but not others). Employing the log normality and risk aversion assumptions of Hodrick and Hansen (1983), we can derive the result that

$$\pi_{t,k}^f = -(2I(s_{t,k}^f > f_{t,k}) - 1) \left(\frac{1}{2} E_t^f [s_{t+k}^2] + E_t^f [s_{t+k} \alpha_{t+k}] \right) \quad (11)$$

where we use lower case π as s_{t+k} and $f_{t,k}$ replace the levels terms in the profit calculation in (8). Note that dividing both sides of the equation by the first term yield the result in (2) above. The results show that profits from this strategy should be no more than the right hand side of the relationship in (11) above. This can be tested, in the same way that the conventional trading rule has been previously tested.

We motivate examining individual profits for each firm, rather than a sample average (or using average expectations), as the heterogeneity results suggest that these indeed represent different strategies for earning profits from the forward foreign exchange rate market. Thus we are not considering these forecasts as estimates of a single rationally expected future exchange rate, but rather such estimates based on different information. As each profit strategy is different each involves different risks, so the risk premium implied by each strategy is different. Only for the naïve trading rule strategy does the risk premia relate to the usual definition of a risk premium in the forward market literature. As the individual profit strategies are correlated, so should their individual risk premia.

The different profit rules that will be examined throughout the remainder of the analysis are summarized. With profits given by

$$\pi_{t,k}^e = (2(I(s_{t,k}^e > f_{t,k}) - 1) (s_{t+k} - f_{t,k})) \quad (12)$$

These are:

a) The individual company profit rules where the expectations are from the JCIF data (the cross sectional average of these expectations is also included as AVE).

b) Profits from using forecasts from the unbiasedness regression, SR-reg, where the full sample estimates of γ_0 and γ from equation (4) are used to compute expectations of the exchange rate, $s_{t,k}^e$ which are used to calculate profits according to equation (12).

c) Profits using forecasts from the unbiasedness equation (4) where γ_0 and γ are estimated only with data available up to time t , RL-reg.

d) The naïve trading rule, NT, where the indicator function always is equal to one so $\pi_{t,k} = (s_{t+k} - f_{t,k})$.

e) Profits where $s_{t,k}^e = s_t$ in equation (12), which corresponds to a random walk model or alternatively static expectations, RW.

5. Examining profits

First, we examine the unconditional mean of profits for each of the firms. A summary of these results are given in the Table 4, along with t statistics, the distribution of the t statistics testing for zero profits are given in Figure 4. The profits of each firm are examined individually and the maximum and minimum are reported. The results at the one month horizon have 10 of the 42 firms losing money by marginally increasing their exposure in the direction suggested by their forecasts, whilst the remainder of the firms would have increased profits through such a strategy. In only six cases is the profit statistically significant from zero. At the three month horizon, five firms make statistically significant profits, whilst two have statistically significant losses. At the six month horizon, two of the firms make statistically significant profits and five statistically significant losses.

Often Sharpe ratios are examined in this context to rank risk adjusted mean returns (e.g. Backus et al (1993)). If a robust estimator of the standard deviation of expected returns is employed, then these ratios are simply the t statistic reported above for testing mean returns to zero (as there is no risk free rate applicable here) divided by the square root of the sample size (so here the Sharpe ratio for the profits generated by RL-reg is approximately 0.24, comparable with the results in Backus et al (1993)). Thus the rankings in terms of Sharpe ratios are identical to those based on the t statistics discussed above.

Examining profits from the naïve trading rule (NT), we see that at the one month and six month horizon such profits are significantly different from zero, although at the three month horizon they are not (the fact that they are negative only means that one should buy and hold the local currency (here yen) rather than the foreign currency). In the case of the regression generated forecasts, all regressions result in statistically significant profits at both horizons except for SR-reg at the three month horizon. In the case of regression 2 (which only uses ex ante information), this profit is better than that generated by the best performer of the company forecasters at all horizons. The random walk profits are negative, very close to zero in magnitude, and insignificantly different from zero.

We see that although the random walk model outperformed all other forecasts at in terms of squared forecast error, the random walk model outperforms only 2 of the 42 firms in terms of profitability at the one month horizon, and 3 out of 42 firms at both the three and six month horizons. This is a clear indication that we should not conclude that the forecast data from the survey respondents or the regressions are irrelevant based on their squared forecast error results.

The existence of modest positive profits is not direct evidence against efficient pricing of the forward contract, as such profits may well be simply a compensation for

risk.¹¹ As noted earlier, each trading rule has its own inherent risks and so the risk premium that would apply depends on the rule. However, given that the profit rules are correlated, so will their risk premia. We investigate whether or not these profits are correlated with variables likely to measure risk or alternatively are correlated with common information in all of the forecasters information sets.

It has been typical of this literature to define differences in profits from zero as a risk premium (see Cumby (1986)), and thus use realized profits from a strategy to determine what variables determine the risk premium. Using the survey data, we are able to construct variables which are outside the information set of the market, and hence may enter significantly without the implication that the market inefficiently employs available information. One oft used measure of the risk premium has been the difference between the forward rate and the average expected rate (Frankel and Froot (1987)). Whilst this ignores price risk, and is motivated somewhat by our inability to observe the risk premium, we can at least examine whether or not this variable is correlated with the excess profits derived as above.¹² The regression to be run here is

$$\pi_{t,k}^i = a_0 + a (f_{t,k} - \bar{s}_{t,k}) + u_{5t} \quad (13)$$

If the risk premium dominates the variation in the difference between the forward rate and the average of the expected future spot rates, then we would expect $a=1$ in the regression. The arguments of Bryant (1994) suggest that this variable measures risk with error, a case where the coefficient estimates will be biased towards zero. Whilst this is arguably the correct risk premium when π is measured using the naive trading rule, it is likely to be correlated with the risk premia of the other trading strategies as well.

An alternate specification would be to use the cross section variability in forecasts as a proxy to measure risk. The idea here is that in times of high uncertainty, the risk of

speculating will be considered higher. This measure has both this intuitive appeal and also the appeal that, as this data is not available to the investor as they cannot observe other agents forecasts, then a statistically significant coefficient cannot be attributed to market failure, and so can be attributed to risk. We would expect this variable to be positively correlated with profits if we believe that these extra profits are due to higher risk.

Results from these specifications are presented in Tables 5 and 6 respectively. Figure 5 gives a histogram of the individual firm estimates of α from equation (13). In the case of estimating equation (13), in all but three cases at the one month horizon the coefficients are positive, and lie between zero and one. In most cases, this is statistically significantly different from one. In the 3 month case, most are between zero and one. In the six month horizon regressions, the estimates of the slope coefficient are distributed around one and the t statistics indicate that nearly all point estimates are insignificantly different from one. The R^2 for the regressions are small, generally lower than 10%. The results suggest that such a risk premium may exist, as the estimates are of the correct sign and magnitude, but does not account for much of the variation in profits (which it need not, so long as there is no remaining explainable variation).

*

Similar results are obtained using profits generated using SR-reg and RL-reg forecasts. At the one month horizon, there is virtually no correlation. As the horizon lengthens the coefficient increases towards one. There is a strong negative relation between the naïve trading rule profits and the risk premium.

In the case of the second specification, the individual profits are regressed on the cross section volatility of forecast error, we would expect positive coefficients. This is because a higher cross sectional variance means greater uncertainty, so the risk premium required would be higher if the risk premia captures exchange risk. The results show that

the opposite is apparent, 28 of the 43 firms show a negative relationship between the variance and return. However, only six of these are significantly different from zero in a two tailed 95% test, so the results show that cross sectional variance of forecasts is not a predictor of profits, and hence do not capture the variance in a risk premium. At the three and six month horizon, the point estimates don't appear to be systematic, although more are significant.

An alternative to the possibility that the predictability in the profits data is due to risk is that it is due to informational inefficiencies in the market. This is typically tested by regressing the profits on information in the agents information set at time t . Rejections of the null hypothesis of no predictability indicate better profit rules that could have been formed, although it may be simply that the explanatory variable is correlated with the risk premium. The explanatory variables that have been popular are the (lagged) forward premium, and lagged changes in the spot exchange rate. The results of the regressions are reported in Tables 7 and 8 respectively. In general, for a number of companies these explanators are statistically significant, although for the most part these variables are not statistically significant predictors of profits.¹³ In all cases the R^2 's for these regressions are extremely small, and these variables are not explaining all of the predictability in the profits. Such results obtain using profits from RL-reg. The usual negative coefficients are found when the forward premium is used to predict the naïve trading rule profits.

6. Conclusion

This paper examined the micro survey data on the yen/dollar exchange rate forecasts and calculated profits based on a possible trading rule. The use of the data on survey expectations allows us to identify profits that could be made from this market which could

have been realized ex post. We found that although survey forecasts are worse than random walk predictions in terms of mean square forecast errors (deviations from the ex post spot rate), survey forecasts would have generated mostly positive profits from a trading rule based on the relative position of the subjective forecast to the forward rate. Many of the firms could have earned excess profits. The profits that could have been earned are highly variable and thus there is significant risk in using these strategies. This means that survey forecasts did better than the random walk belief in predicting a correct direction of the future spot exchange rate in relation to the forward rate. The survey-forecasts-based profits were comparable to profits which would have been generated from a trading rule based on forecasts from a regression of spot rate changes with a forward premium as an explanatory variable, and the latter result is typically held to imply thorough rejection of forward market efficiency. The regression-based profits are also highly variable.

The forecast-based profits are correlated with the conventional measure of the risk premium (the difference between the forward and expectation), however this risk premium explains very little of the variance in the profits. It was further found that cross section variance in forecasts, a proxy for uncertainty, was not correlated with the forecast-based profits in a systematic way.

These results back up the belief that although apparent excess profits can be earned from the forward foreign exchange market, their large variation puts in doubt whether these were unexploited profit opportunities, after adjusting for risk. Thus, the economic significance of such profits are in question, as has been similarly argued in the stock market literature in Breen et al. (1989).

Skeptics still may point out the possibility that the forecasters simply gave random forecasts around the current spot rate or forward rate, yielding the results. Whilst this

would result in calculated profits insignificantly different from zero, as found here, it may not explain the other results such as heterogeneity and the correlation of the profits with the conventionally-defined risk premium. The large difference between the profitability behavior of the random walk model forecasts and the respondent forecasts suggest that more is going on in the models of the respondents than static expectations with stochastic reporting.

Bibliography

- Backus, D.K., A.W. Gregory and C.I. Telmer (1993) "Accounting for Forward Rates in Markets for Foreign Currency", *Journal of Finance*, 48, pp. 1887-1908.
- Backus, D.K., S. Foresi and C.I. Telmer (1994), "The Forward Premium Anomaly: Three Examples in search of a Solution", manuscript, NYU. June.
- Bekaert, G. and R.J. Hodrick (1992) "Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets", *Journal of Finance*, 47, pp. 467-509.
- Bekaert, G. and R.J.Hodrick (1993), "On Biases in the Measurement of Foreign Exchange Risk Premiums", *Journal of International Money and Finance*, 12, pp. 115-138.
- Bilson, J.F.O (1981), "The Speculative Efficiency Hypothesis", *Journal of Business*, 54, pp. 435-52.
- Breen, W. L.R. Glosten and R. Jagannathan (1989), "Economic Significance of Predictable Variations in Stock Index Returns", *Journal of Finance*, 44, pp. 1177-1189.
- Bryant, R.C. (1994), "The 'Exchange Risk Premium', Uncovered Interest Parity, and the Treatment of Exchange Rates in Multicountry Models", manuscript, Brookings Institution.
- Cumby, R.E. (1988) "Is it Risk? Explaining Deviations From Uncovered Interest Parity", *Journal of Monetary Economics*, 22, pp. 279-99.
- Dickey, D.A. and W.A.Fuller (1979), "Distribution of Estimators for Autoregressive Time Series with a Unit Root", *Journal of the American Statistical Association*, 74, pp. 427-431.
- Diebold, F.X and J.A. Nason (1990), "Nonparametric Exchange Rate Prediction", *Journal of International Economics*, 28, pp. 315-332.
- Elliott, G (1994), "On Tests of Unbiasedness in the Forward Foreign Exchange Rate Market", manuscript, UCSD.
- Elliott, G., T.J.Rothenberg, and J.H.Stock (1992), "Efficient Tests for a Unit Root", NBER Technical Working Paper #130.
- Fama, E.F. (1984) "Forward and Spot Exchange Rates", *Journal of Monetary Economics*, 14, pp. 319-38.
- Frankel, J.A. and K.A. Froot (1987) "Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations", *American Economic Review*, 77, pp. 133-53.
- Froot, K.A., and J. A., Frankel (1989), "Forward Discount Bias: Is it an Exchange Risk Premium?" *Quarterly Journal of Economics*, 104, pp. 140-61.
- Froot, K.A. and R.H. Thaler (1990, "Anomalies: Foreign Exchange", *Journal of Economic Perspectives*, 4, pp. 179-92.

Hakkio, C.S. (1981), "Expectations and the Forward Foreign Exchange Rate", *International Economic Review*, v22, pp. 663-78.

Hansen, L.P. and R.J.Hodrick (1980) "Forward Exchange Rates as Optimal Predictors of Future Spot rates: An Econometric Analysis", *Journal of Political Economy*, 88, pp. 829-53.

Hansen, L.P. and R.J.Hodrick (1983) "Risk Averse Speculation in the Forward Foreign Exchange Market: An Econometric Analysis of Linear Models", in Exchange Rates and International Macroeconomics, J.A. Frenkel ed. University of Chicago Press.

Hodrick, R.J. (1987) The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets, Harwood Academic Publishers: Switzerland.

Hodrick, R.J. and S. Srivistava (1986) "The Covariation of Risk Premiums and Expected Futures Spot Rates", *Journal of International Money and Finance*, 3, pp. 5-30.

Ito, T (1992) "Foreign Exchange Rate Expectations: Micro Survey Data", *American Economic Review*, 80, pp. 434-49.

Ito, T. (1994), "Short-Run and Long-Run Expectations of the Yen/Dollar Exchange Rate", *Japanese Journal of International Economics*, 8, pp. 119-43.

Leitch, G. and J.E. Tanner (1991), "Economic Forecast Evaluation: Profits versus the Conventional Error Measures", *American Economic Review*, 81, pp. 580-590.

Lucas, R.E. (1982), "Interest Rates and Currency Prices in a Two-Country World", *Journal of Monetary Economics*, 10, pp. 335-60.

McCallum, B. (1994) "A Reconsideration of the Uncovered Interest Parity Relationship", *Journal of Monetary Economics*, v33, pp. 105-32.

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

Stock, J.H. (1991), "Confidence Intervals for the Largest Autoregressive root in U.S. Macroeconomic Time Series", *Journal of Monetary Economics*, 28, pp. 435-459.

White, H. (1980) "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity", *Econometrica*, 48, pp. 817-38.

Notes:

1. Backus et al (1993) do so but use the full sample estimates of the model parameters, which are not available to market participants in real time.
2. Ito is extremely grateful to the Japan Center for International Finance (JCIF), Tokyo, for providing its proprietary data set for this research. The data remains the property of the JCIF.
3. Spot and forward rates are the closing rates in the Tokyo market on the survey date. The six-month forward rates in 1993 and later were not available in Tokyo, so that they were substituted by the New York quotes.
4. The log approximations leave only small approximation errors, so long as the forward and spot rate are not too far apart.
5. All tests in this paper are at the 95 percent confidence level.
6. Each of the regressions is estimated individually by OLS with a White (1980) correction for potential serial correlation and weak heteroskedasticity.
7. This is the square root of the mean square error, i.e. $[\text{T}^{-1}\Sigma(s_{t,t+k}^i - s_{t+k})^2]^{1/2}$.
8. Usually the regression is transformed by subtracting $(s_{t,k}^i - s_t)$ from each side.
9. Here, the largest root in the spot exchange rate is estimated at 0.99, whilst the median unbiased estimate of this root (Stock (1991)) is 0.98. There is a small amount of dynamics in the change in the spot rate (the BIC selector chooses a lag length of one in the augmented Dickey Fuller regression). Both the Dickey Fuller test and the DF-GLS test of Elliott et al (1992) fail to reject the presence of a unit root.
10. There needs to be some limit placed on the bets, as this is an infinitely leveraged market as there is no money down up front. We assume that each company can speculate on the forward market up to a pre-specified credit limit. Alternatively, we can consider such profits as marginal profits earned by a small change in the existing (unknown) exposures of the companies portfolios.
11. It is difficult to say whether or not these are economically significant profits. See Breen et al. (1989), which evaluates the economic significance of profits in the stock market.
12. See Bryant (1994) for an argument of why this does not well represent the risk premium.

13. For the change in the spot as a predictor, we obtain statistically significant coefficients for 3,6 and 4 of the companies at the 1,3 and 6 month horizons respectively. When the forward premium is used as a predictor, the number of firms with statistically significant profits are 1,0 and 1 respectively.

Table 1 : Heterogeneity						
Company	1 Month		3 Month		6 Month	
	mean	t	mean	t	mean	t
maximum	0.0105	3.60	0.0254	5.95	0.0280	6.34
minimum	-0.0099	-4.16	-0.0212	-5.39	-0.0337	-4.84
forward	0.0058	4.80	0.0029	1.05	-0.0085	-1.52

Notes: Means are computed from OLS regressions on a constant over the full data period (see equation 5). The t statistics reported are testing that the mean is zero and employ a White (1980) correction for the variance covariance matrix of the estimate where the spectral density of the variance-covariance matrix of the residuals at frequency zero is calculated using an autoregressive method where the lag length is selected by a Bayesian information criterion (so that standard errors are robust to serial correlation, in part due to the overlapping data). The maximum and minimum deviation (in terms of means) are reported, along with the deviation of the forward rate from the average expected rate.

Table 2 : Prediction Accuracy						
Forecast	1 Month		3 Month		6 Month	
	F'cst error	Direc.	F'cst error	Direc	F'cst error	Direc.
Survey						
maximum	1.2804	.56	1.3050	.59	1.3443	.59
minimum	1.0204	.45	1.0198	.46	1.0011	.47
average	1.0301	.50	1.0112	.51	1.0307	.54
SR-reg	0.9988	.46	0.9962	.48	0.9860	.52
RL-reg	1.0407	.48	1.1037	.48	1.1839	.49
forward	1.0041	.48	1.0097	.51	1.0219	.52

Notes: Numbers reported (F'cst error) are the standard error of the deviation of the forecast from the ex post exchange rate (in logs) divided by the standard deviation of the log change in the exchange rate (the denominator is the forecast error of the random walk model). The maximum and minimum over the 42 companies are reported. The same calculations are made for the other methods of deriving forecasts (see the end of section 4 for a description of these forecasts). The numbers reported in columns labelled Direc. are the percentage of times over the sample where the forecast correctly predicted the direction of the subsequent spot rate movements.

Forecast	1 Month		3 Month		6 Month	
	β_0	β	β_0	β	β_0	β
Survey maximum	-0.01 (-2.78)	0.29 (-2.92)	-0.02 (-2.41)	0.30 (-2.10)	-0.05 (-2.00)	0.38 (-1.70)
minimum	-0.01 (-3.39)	-0.10 (-14.05)	-0.03 (-2.40)	-0.04 (-10.39)	-0.05 (-2.57)	-0.21 (-9.54)
average	-0.01 (-3.16)	0.11 (-3.74)	-0.02 (-2.43)	0.31 (-1.13)	-0.05 (-2.50)	0.08 (-1.38)
RL-reg	-0.00 (-1.14)	0.25 (-4.62)	-0.01 (-0.64)	0.19 (-2.79)	-0.01 (-0.29)	0.24 (-2.65)
Naive	-0.01 (-3.52)	-0.90 (-1.33)	-0.03 (-2.57)	-1.14 (-0.83)	-0.06 (-2.72)	-1.65 (-0.91)

Notes: Point estimates are OLS estimates of the regression in equation (6) in the text (equation (4) for the last row). The t statistics are reported in brackets, for the mean the null hypothesis is zero and for the intercept the null is one. In each case the robust standard error is calculated as per Table 1 notes. The maximum and minimum estimates (based on t statistics) are reported.

Forecast	1 Month		3 Month		6 Month	
	mean	t(H0=0)	mean	t(H0=0)	mean	t(H0=0)
Survey						
maximum	0.007	2.03	0.021	2.50	0.021	2.74
minimum	-0.008	-3.08	-0.028	-4.32	-0.037	-3.59
average	0.006	1.92	0.007	0.81	-0.008	-0.67
SR-reg	0.008	2.81	0.023	1.45	0.052	3.18
RL-reg	0.010	3.62	0.026	2.83	0.045	2.31
Naive	-0.008	-2.81	-0.023	-1.45	-0.039	-2.05
RW	-0.003	-0.90	-0.009	-0.81	-0.021	-1.07

Notes: Point estimates are usual estimates of the mean. In each case the robust standard errors are calculated according to the notes in Table 1. The maximum and minimum are based on t statistics).

Table 5 : Profits on 'Risk Premium'									
f' cast	1 Month			3 Month			6 Month		
	β_0	β	R ²	β_0	β	R ²	β_0	β	R ²
Survey max	-0.01 (-1.55)	0.72 (-1.00)	0.04	-0.00 (-0.00)	1.23 (0.36)	0.06	0.01 (0.86)	1.55 (1.11)	0.11
min	-0.01 (-2.51)	-0.15 (-5.26)	0.00	-0.00 (-0.71)	-0.54 (-3.12)	0.01	-0.03 (-2.64)	-0.05 (-2.71)	0.00
ave	0.00 (1.91)	0.25 (-3.18)	0.00	-0.00 (-0.22)	1.41 (0.67)	0.08	0.00 (0.11)	1.39 (0.71)	0.09
SR-reg	-0.01 (-2.54)	-1.90 (-3.57)	0.00	0.02 (1.86)	0.47 (-0.96)	0.01	0.06 (3.11)	0.46 (-1.19)	0.01
RL-reg	0.01 (3.47)	0.12 (-3.75)	0.00	0.02 (2.47)	0.79 (-0.44)	0.03	0.05 (2.35)	0.61 (-0.99)	0.03
Naive	-0.01 (-2.54)	-0.16 (-4.97)	0.00	-0.03 (-2.57)	-2.14 (-0.83)	0.03	-0.06 (-2.72)	-2.65 (-0.91)	0.07
RW	0.00 (0.29)	-0.69 (-5.65)	0.03	-0.01 (-0.53)	-0.74 (-3.43)	0.02	-0.02 (-0.91)	-0.18 (-1.37)	0.00

Notes: Point estimates are OLS regression coefficients from a regression of profits on a constant and the difference between the forward rate and average expected rate. The t statistics test a null of zero on the constant (β_0) and one on the slope (β) and are corrected for serial correlation as in Table 1. The maximum and minimum are based on point estimates of β .

Table 6 : Profits on cross sectional variance of expectations									
f' cast	1 Month			3 Month			6 Month		
	β_0	β	R ²	β_0	β	R ²	β_0	β	R ²
Survey									
max	-0.02 (-4.15)	27.10 (3.54)	0.03	-0.07 (-4.96)	37.33 (3.76)	0.07	-0.06 (-2.91)	21.46 (2.63)	0.04
min	0.01 (1.76)	-19.40 (-1.78)	0.02	0.05 (2.65)	-25.27 (-2.33)	0.03	0.04 (2.61)	-20.41 (-3.74)	0.03
ave	0.01 (1.90)	-10.84 (-1.35)	0.01	0.01 (0.41)	-0.52 (-0.13)	0.00	-0.03 (-0.87)	9.15 (0.81)	0.01
SR-reg	0.02 (3.08)	-19.07 (-2.34)	0.02	0.06 (2.50)	-33.82 (-2.91)	0.06	0.12 (5.87)	-35.37 (-4.04)	0.13
RL-reg	0.02 (3.58)	-20.30 (-2.48)	0.02	0.06 (4.18)	-32.99 (-3.51)	0.06	0.10 (4.52)	-30.99 (-3.54)	0.13
Naive	-0.02 (-3.08)	19.07 (2.34)	0.02	-0.06 (-2.50)	33.82 (2.74)	0.06	-0.08 (-2.27)	24.50 (1.94)	0.05
RW	-0.00 (-0.47)	-0.93 (-0.12)	0.00	-0.01 (-0.55)	2.78 (0.21)	0.00	-0.02 (-0.57)	-0.48 (-0.04)	0.00

Notes: Reported are OLS coefficients from a regression of profits on a constant and the cross sectional variance of expectations. The t statistics test the constant β_0 and also slope β to zero. Robust standard errors as in earlier tables are employed.

Table 7: Profits on Forward Premium									
F'cast	1 Month			3 Month			6 Month		
	β_0	β	R ²	β_0	β	R ²	β_0	β	R ²
Survey max	0.00 (1.66)	2.87 (2.57)	0.02	0.01 (1.25)	2.69 (1.19)	0.04	0.03 (3.02)	2.70 (2.00)	0.06
min	-0.00 (-0.45)	-2.90 (-1.76)	0.02	-0.01 (-0.78)	-1.79 (-1.33)	0.02	-0.00 (-0.27)	-2.19 (-1.63)	0.04
ave	0.01 (1.35)	-0.89 (-0.58)	0.00	0.01 (1.47)	1.53 (0.82)	0.01	-0.01 (-0.78)	-0.43 (-0.33)	0.00
RL-reg	0.01 (3.38)	0.04 (0.03)	0.00	0.02 (1.88)	-0.98 (-0.64)	0.01	0.04 (1.84)	-0.74 (-0.48)	0.01
Naive	-0.01 (-3.52)	-1.90 (-1.33)	0.01	-0.03 (-2.57)	-2.14 (-0.83)	0.03	-0.06 (-2.72)	-2.65 (-0.91)	0.07
RW	-0.00 (-0.26)	1.65 (0.99)	0.00	-0.00 (-0.47)	1.22 (0.33)	0.01	-0.02 (-0.94)	-0.07 (-0.02)	0.00

Notes : Reported are OLS coefficients from a regression of profits on a constant and the forward premium. The t statistics test the constant β_0 to zero and the slope β to zero. Robust standard errors as in earlier tables are employed.

f' cast	1 Month			3 Month			6 Month		
	β_0	β	R^2	β_0	β	R^2	β_0	β	R^2
Survey max	0.00 (1.27)	0.07 (0.95)	0.00	-0.00 (-0.13)	0.29 (3.29)	0.08	0.01 (0.49)	0.24 (1.99)	0.06
min	0.00 (1.30)	-0.19 (-3.09)	0.03	-0.01 (-1.12)	-0.25 (-2.79)	0.05	-0.01 (-1.07)	-0.22 (-2.86)	0.05
ave	0.01 (2.13)	-0.04 (-0.71)	0.00	-0.00 (-0.09)	-0.18 (-1.52)	0.03	-0.00 (-0.16)	0.02 (0.27)	0.00
SR-reg	0.01 (2.67)	-0.02 (-0.31)	0.00	0.02 (1.63)	-0.08 (-0.56)	0.01	0.05 (2.17)	-0.09 (0.48)	0.01
RL-reg	0.01 (3.44)	-0.03 (-0.52)	0.00	0.03 (3.22)	0.07 (0.93)	0.00	0.05 (2.15)	0.28 (1.27)	0.09
Naive	-0.01 (-2.67)	0.02 (0.31)	0.00	-0.02 (-1.63)	0.08 (0.56)	0.01	-0.03 (-1.19)	-0.06 (-0.32)	0.00
RW	-0.00 (-0.82)	0.03 (0.50)	0.00	-0.00 (-0.25)	0.10 (0.72)	0.01	0.01 (0.25)	0.45 (4.74)	0.19

Notes: As per Table 7 with the lagged change in the spot rate as the regressor (two week change, lagged so that the variable is in the information set of the forecaster).

Figure 1

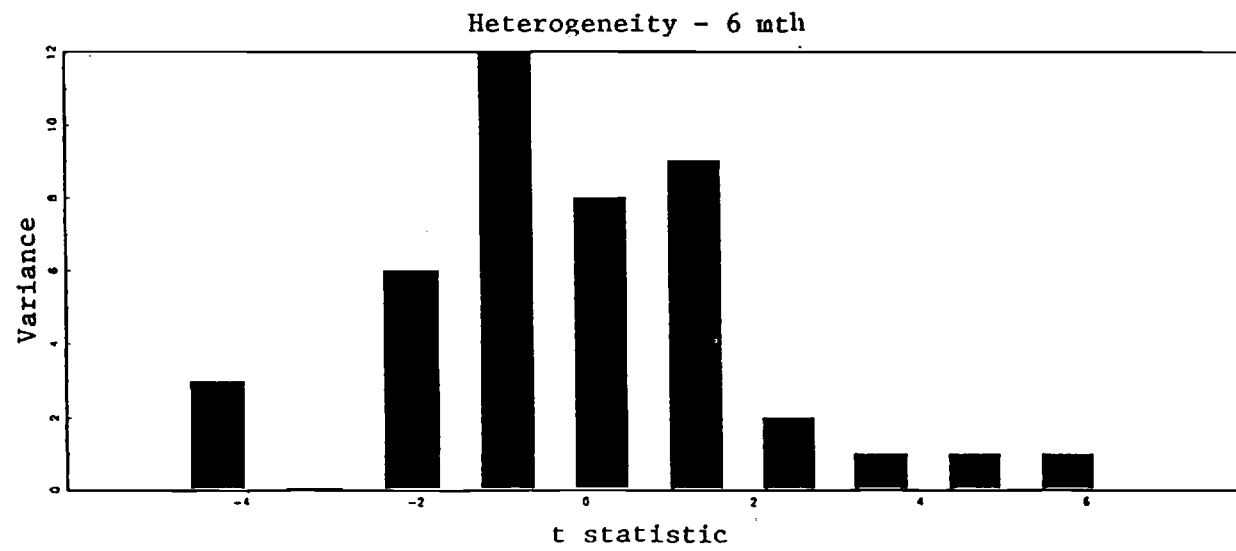
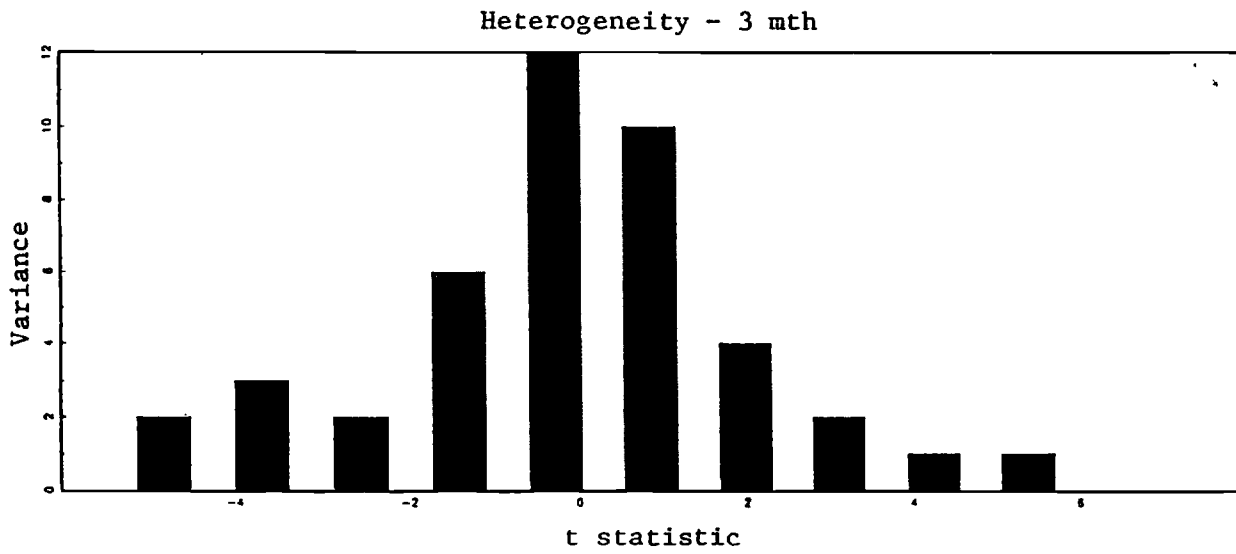
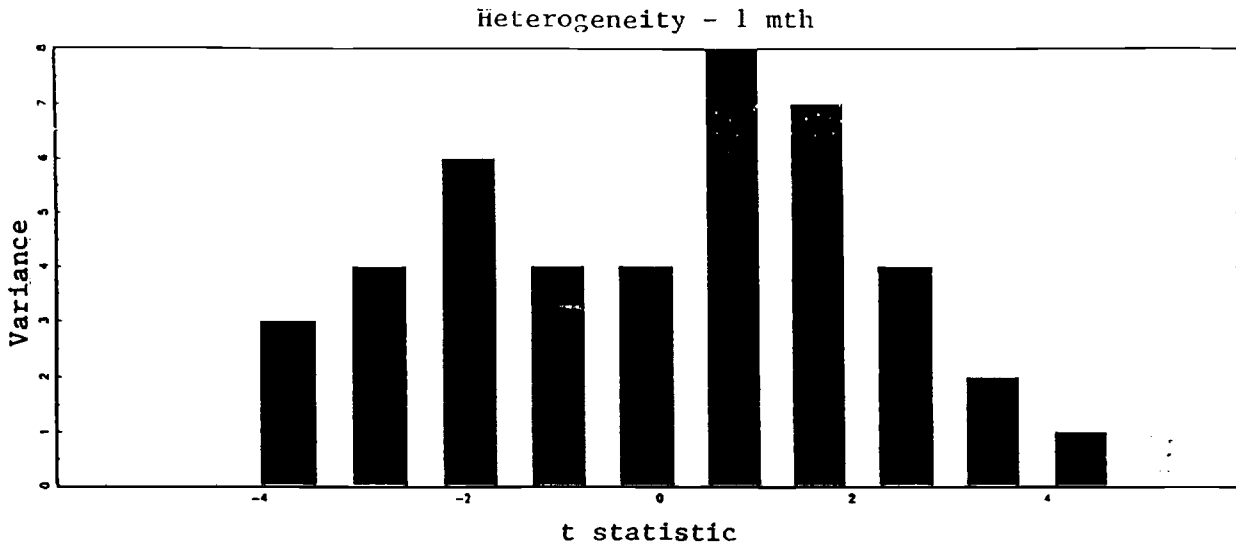




Figure 2

Forecast Variance - 1 mth



Forecast Variance - 3 mth



Forecast Variance - 6 mth

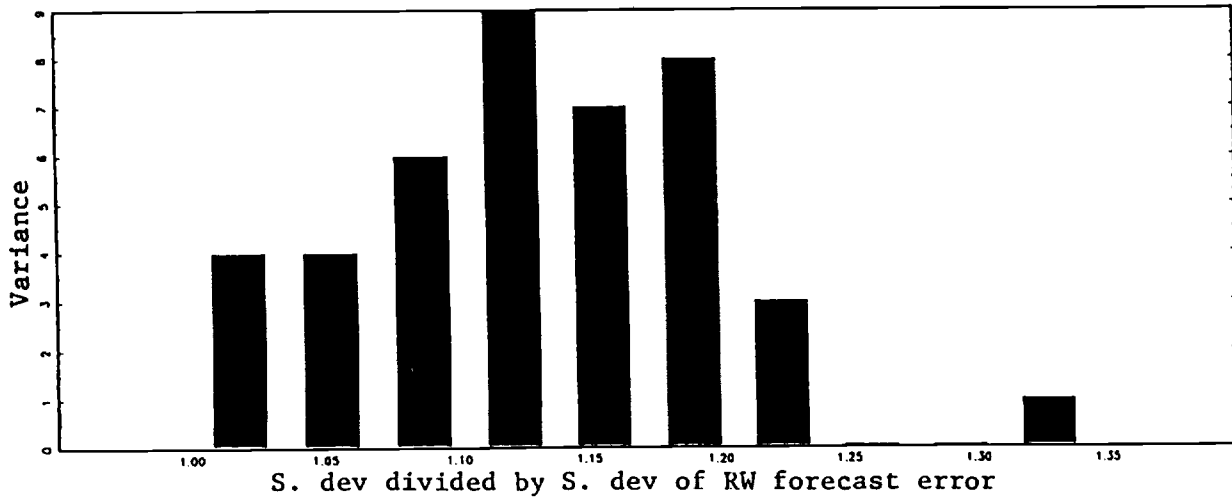




Figure 3

Forecast data to Explain ex post Depreciation - 1 mth



Forecast data to Explain ex post Depreciation - 3 mth



Forecast data to Explain ex post Depreciation - 6 mth



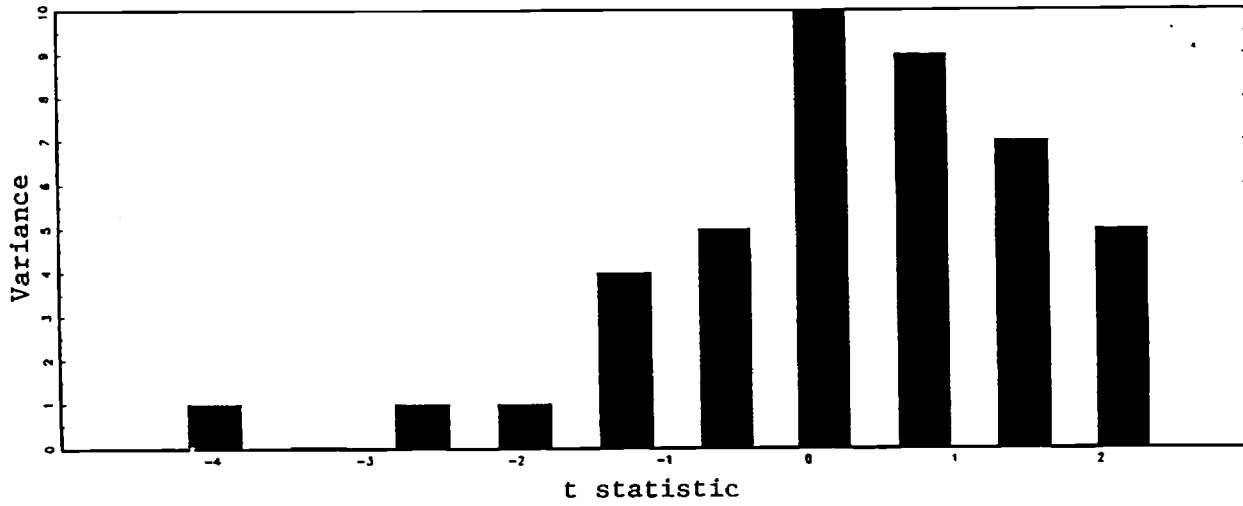


Figure 4

Average Profits - 1 mth



Average Profits - 3 mth



Average Profits - 6 mth






Figure 5

Profits on 'risk prem' - 1 mth



Profits on 'risk prem' - 3 mth



Profits on 'risk prem' - 6 mth

