Title
Heterogeneous Expectations and Tests of Efficiency in the Yen/Dollar Forward Exchange Rate Market

Authors
Elliott, Graham
ITO, TAKATOSHI

Publication Date
1998-02-01

98-06

UNIVERSITY OF CALIFORNIA, SAN DIEGO

DEPARTMENT OF ECONOMICS

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

BY

GRAHAM ELLIOTT

AND

TAKATOSHI ITO

DISCUSSION PAPER 98-06
FEBRUARY 1998

Heterogeneous expectations and tests of efficiency in the yen/dollar forward exchange rate market

Graham Elliott,[a] Takatoshi Ito[b*]

[a]Department of Economics, University of California, San Diego, La Jolla, CA 92093-0508
[b]Institute of Economic Research, Hitosubashi University, Kunitachi Tokyo 186, Japan

January 1995, This revision February 1998

Abstract

This paper examines the efficiency of the forward yen/dollar market using micro survey data. Conventional tests of unbiasedness do not correspond directly to the zero-profit condition. Instead, we use the survey data to calculate 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 rates 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. Similar results are found when we examine profits generated by a trading rule using regression forecasts. The profits are found to be correlated with risk type variables but not other available information.

Key Words: Foreign exchange rate; Expectations; Forward rate; and Efficient markets.

JEL classification: F31, G14, G15

Acknowledgment: We thank J. Frankel, A. Timmermann, A. Melino, an anonymous referee and seminar participants at the Federal Reserve Board of Governors for their useful comments. Elena Pesavento provided excellent research assistance. We are grateful to the Japan Center for International Finance for their kindly providing us an access to their survey data. Views expressed in the paper are ours.

1. Introduction

The unbiasedness hypothesis for the forward foreign exchange rate states that the forward rate for settlement k periods ahead ($F_{t,k}$) is unbiased predictor of the k-period ahead spot rate (S_{t+k}). The conventional test asks whether the forward premium (in logs), ($f_{t,k}$-s_t), is an unbiased predictor of ex post depreciation (s_{t+k}-s_t). The frequent rejection of this test is commonly referred to as a "puzzle", as it suggests the possibility of the existence of unexploited profit opportunities. (Froot and Thaler (1990), Lewis (1995) give surveys).

The risk premium is typically considered the most likely source of the puzzle, but attempts at resolution of the regression results through this avenue require risk premia that are far larger than most researchers would accept or find (see Bekaert (1995)). In this paper we skip over the typical regression tests and focus directly on the 'excess profits' that can be earned[1]. We calculate such potential profits using the regressions and using measures of potential profits based on expectations of market participants, finding that the excess profits have a large variance and are (statistically) small. The excess profits are found to be correlated with usual proxies for risk but not with other variables. We also find that although the random walk model outperforms the expectations data in terms of mean square forecast error, excess profits that would be derived using the expectations data typically exceeds that generated using the random walk forecasts.

The data set in this study, used in Ito (1990) and updated, contains individual forecasts of 42 companies on future spot yen/dollar rate at the one, three and six month horizons, fortnightly from May 1985 to May 1996. This data set is unique in that regular membership has been maintained for more than ten years, with few missing observations[2]. A strength of this paper is that we are able to bypass the problem that the econometrician does not observe the forecasters' model in examining efficiency. As the data contain heterogeneous individual expectations, this particular data set provides a large number of tests for a single sample period. A potential weakness is that respondents may not give serious answers, but profit calculations presented later indicate the seriousness of the data in our sample.

Section 2 examines the relationship between profits from trading rules and regression tests of unbiasedness. Section 3 examines the expectations data, giving results on heterogeneity and forecast performance. 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. Predictable profits and the naive trading rule

We review the familiar implications of an efficiently functioning forward market, relating this to the derivation of empirical tests.

A. Zero profit condition

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

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

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

---

[1] Others have examined profits from filter rules, e.g. Dooley and Shafer (1984), Sweeney (1986), Levich and Thomas (1993) and from regressions Bilson (1981), Hodrick and Srivastava (1984).

[2] The original data has 44 companies (see Ito (1990)), 2 are dropped due to frequent missing observations.

t+k. Backus et al. (1996) derive a similar equation to (1) where Q_{t+k} is a (yen) 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 US dollar forward and close the position when delivery is due by selling the US dollar at the future spot exchange rate.

Hansen and Hodrick (1983) further show, conditioning on the data Z_t 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

$$E[s_{t+k}|Z_t] - f_{t,k} = -\tfrac{1}{2}Var[(s_{t+k} - s_t)|Z_t] - Cov[q_{t+k}(s_{t+k} - s_t)|Z_t] \qquad (2)$$

where s_t=ln(S_t), $f_{t,k}$=ln($F_{t,k}$) and q_{t+k}=ln(Q_{t+k}). In the absence of a risk premium, expected profits from the naive trading rule, i.e. when one currency is always bought forward, should not exceed on average the conditional variance of the spot exchange rate plus 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 except through the correlation of Z_t with these variables. Given that the conditional volatility of the exchange rate and the conditional covariance term are small relative to the variation (s_{t+k} - $f_{t,k}$), this is often taken to mean that the left hand side of (2) should not be predictable at time t. This can be interpreted as zero expected profits from the naive trading rule.

B. Unconditional tests

After replacing the expected future spot rate by the ex post future exchange rate (and invoking rational expectations), the orthogonality condition implied above has been tested often in the literature, with empirical results at odds with the theoretical null hypothesis. One test of this relationship is to examine directly the unconditional mean of such profits, to examine whether or not excess profits are available, i.e. that naive trading rule[3] profits are zero.

There are two caveats to such tests. First, profits from the naive trading rule underestimate potential profits that could be earned from the forward market in practice, as this represents only one possible strategy in speculating against the forward market. This measure is popular due to it being the strategy most easily identified by the econometrician, but potentially underestimates profits available to the more sophisticated investors[4]. A second caveat suggests that use of the unconditional test of a buy-forward rule as the profit measure may overestimate such potential profits. If average profits turn out to be negative, it is interpreted that the reverse (sell-forward) would have made profits. But this requires that the investor know beforehand which currency will on average appreciate over the period. Given that the exchange rate is near to a random walk, correctly guessing the sign ex ante is difficult.

---

[3] Throughout, exchange rates are yen per dollar. Buying forward refers to buying dollars forward, which yields a profit if $s_{t+k}>f_{t,k}$.

[4] 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 naive strategy. A forecaster with perfect foresight, correctly picking the sign of the forward error, always makes a sure profit of 1 unit per month.

C. Trading rule based on the random walk belief

More sophisticated trading rules can be considered. A common baseline model is the random walk model of exchange (See Meese and Rogoff (1983), Diebold and Nason (1990)). This model places an equal probability on appreciation or depreciation, suggesting a trading strategy that sells dollars forward when the forward yen rate is at a premium over the spot rate and vice versa. The direction of bets, which stays the same in the naive strategy, would now change whenever the forward premium changes sign.

D. The regression based model

Alternatively, the conditional mean of $(s_{t+k}-f_{t,k})$ can be examined 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.

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

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 (1988), and Bekaert and Hodrick (1992). Such regressions are interesting as rejections of the null hypothesis would define a trading rule that could be more profitable than the naive rule. In particular, the following specification is frequently employed

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

where the null hypothesis is $\gamma=1$. This is a special case of (3) with $Z_t=(f_{t,k}-s_t)$ where $\beta = \gamma-1$ and $\beta_0 = \gamma_0$. This version has been more often tested, see Bilson (1981), Fama (1984), Bekaert and Hodrick (1993) and McCallum (1994). This shows the relationship between unbiasedness tests and the zero profitability condition tests. The finding that γ is different from one in this regression[5] suggests that forecasts can be generated from the regression which will be useful for making profits speculating against the forward market.

Two forms of this regression 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 although 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.

A caveat to this analysis is that although we may statistically reject the null hypothesis that the naive 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 in that profits from such rules may turn out to be small and highly variable. Statistical rejection is not enough evidence on its own to gauge inefficiency of the forward market. (see Breen et al (1989) in the context of stock market predictability).

E. <u>The role of expectations</u>

An alternate strategy to identifying potential profits earned from speculating on the forward market would be to use the forecasts of market participants directly. We consider these forecasts as the outcome of an unobserved model of exchange rate determination. We examine forecasts of the future yen/dollar rate by 42 Japanese companies, and use these to obtain a potentially more precise measure of excess profits that could have been earned.

---

[5] See Lewis (1995) for a recent survey and discussion of potential explanations.

These profit calculations address the caveats mentioned above. All of the information required to calculate the profit rules comes from information known to the participants. Also, as we believe that participants have some model of exchange rates, of which the econometrician observes only the forecast, the calculated profits represent a better approximation of potential profits market players can earn from speculating against the forward market.

3. **Expectations data and heterogeneity**

This section examines previous results of systematic heterogeneity in forecasts and also the performance of the expectations data. Findings of heterogeneity motivate the use of individual data in subsequent analysis; apparently there is information special to each forecast.

A. Heterogeneity

The survey data show large differences in opinion amongst market participants, especially during the 1985-87 period. Differences of 15 to a maximum of 30 yen per dollar are recorded at the one month horizon, and 25 to a maximum of 70 yen (most are 25-40 yen) at the six month horizon. Agents also deviate systematically from each other in their forecast of the future exchange rate. Ito (1990), using a subsample of the data employed here, found 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

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

where $s_{t,k}^{i}$ are the individual (log) forecasts of s_{t+k}, $\bar{s}_{t,k} = \frac{1}{42}\sum_{i=1}^{42} s_{t,k}^{i}$ and under the null hypothesis of nonsystematic deviations in expectations, the mean g_i should be zero.

The result of Ito (1990) that expectations appear to be systematically heterogeneous is upheld in the updated data set employed here. This is true for forecasts at all horizons. 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. A histogram of the t statistics testing the hypothesis for all of the companies is reported in Figure 1. Seventeen of the forty-two firms have forecasts at the one month horizon which are (statistically) significantly different from the sample average. At the three- (and six-month horizons), 13 (and 12 respectively) deviate from the sample average (i.e. are statistically significant at the 95 percent confidence level).

B. Forecast performance

An obvious question to ask when examining forecast data is how ‘good’ the forecasts are. One measure of the ‘goodness’ of forecasts is to examine the standard error of the ex post expectation error. We normalize these statistics by the random walk forecast error. The results are presented in Table 2 (and the distribution of the full results in Figure 2), along with the other forecasts described in Section 2. 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.

Nearly all of the individual firm forecasts have a larger variance of the forecast error than the random walk model (2 exceptions), and also greater variance than the forward market (none has a lower forecast variance at the 1-month horizon, 3 firms are better at the 3-month and 5 are better at the 6-month horizon). The forecasts from RL-reg have a lower forecast variance than nearly all of the individual forecasters at the one month horizon, but these forecasts are outperformed on this criterion at the 3-month horizon by 8 firms and at the 6-month horizon by

37 firms. Also note that the forecasts generated with RL-reg are outperformed by the forward rate at all horizons.

It is not straightforward that these results indicate that the survey data are outforecasted by the random walk model, or that the survey data are 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. Thus it may well be the case that the forecast data here are outperformed by the random walk model, but they are more valuable in terms of earning profits (see Leitch and Tanner (1991) in the context of Treasury Bill forecasts). This will turn out to be the case.

Also reported in Table 2 are the proportion of times the forecast correctly picks the direction of the spot market. Most of the numbers are close to 50%, with regression based forecasts superior to the individual forecasts. Forecasters and regression based forecasts for this measure are better for longer maturities.

4. Trading rules

By observing the forecasts of market participants directly, we are able to examine the potential profits of agents participating in this market directly. In calculating 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[6]. Thus if the agent believes that the forward rate undervalues the dollar, that agent will take a long position forward. They will take a short

---

[6] 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.

position if they believe that the forward rate overvalues the future value of the dollar. Profits for the ith company for the period over t to t+k are given by

$$\Pi^i_{t,k} = I(s^i_{t,k} > f_{t,k})(S_{t+k} - F_{t,k}) + I(s^i_{t,k} < f_{t,k})(F_{t,k} - S_{t+k}) = (2I(s^i_{t,k} > f_{t,k}) - 1)(S_{t+k} - F_{t,k}) \quad (6)$$

where I(.) is an indicator function yielding one if the statement inside the bracket is true and zero otherwise. We motivate examining individual profits for each firm, rather than a sample average (using average expectations), as the heterogeneity results suggest that these indeed represent different strategies for earning profits from the forward foreign exchange rate market.

The different profit rules that will be examined throughout the remainder of the analysis are summarized. With log profits given by $\pi^e_{t,k} = (2(I(s^e_{t,k} > f_{t,k}) - 1)(s_{t+k} - f_{t,k})$ 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 are used to compute $s^e_{t,k}$.

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 naive trading rule, NT, where $\pi^e_{t,k} = (s_{t+k} - f_{t,k})$.

e) Profits where $s^e_{t,k} = s_t$, which corresponds to a random walk model, RW.

5. Examining profits

Unconditional means of calculated profits are given in Table 3, along with t statistics and a histogram containing the t statistics testing for zero profits are given in Figure 3. The profits of each firm are examined individually and the maximum and minimum are reported. Only 8 of the

42 firms lose money by investing in the direction suggested by their one month ahead forecasts (5 are statistically significant from zero at the 95% level). At the 3-month horizon, 5 firms make statistically significant profits (2 have statistically significant losses). At the 6-month horizon, no firms make statistically significant profits (2 make 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 naive trading rule (NT), we see that at the one month and six month horizon such profits are insignificantly different from zero (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 all horizons except for SR-reg at the three month horizon. In the case of RL-reg (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 in terms of average squared forecast error, the random walk model outperforms only 2 of the 42 firms in terms of profitability at the one- and three-month horizons, and none at the six month horizon.

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. We investigate whether or not these profits are correlated with variables likely to measure risk (following Cumby (1988)) or alternatively are correlated with common information in all of the forecasters information sets through regressing profits on each of the variables separately.

The survey data enables construction of variables outside the information set of the market, and hence may enter significantly without the implication that the market inefficiently employs available information. The most commonly used measure of the risk premium has been the difference between the forward rate and the average expected rate (Frankel and Froot (1987)). If the "risk premium" was measured without error by $(f_{t,k} - \bar{s}_{t,k})$, then we would expect a coefficient of 1 in the regression of profits on a constant and this risk proxy. Bryant (1994) suggests that as $(f_{t,k} - \bar{s}_{t,k})$ contains a large high frequency component to its variation it measures risk with error, a case where the coefficient estimates will be biased towards zero. Results from this regression are presented in Table 4. Figure 4 gives a histogram of the slope coefficients. At the 1-month horizon, 34 of the coefficients are positive, and lie between zero and one, and in most cases are statistically significantly different from one. In the 3- and 6-month cases, most are between zero and one. In the six month horizon regressions, 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, 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.

An alternate specification would be to use the cross section variability in forecasts as a proxy to measure risk. We assume that in times of high uncertainty, the risk of speculating will be considered higher. 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 this regression are presented in Table 5. The results show no clear relationship, with point estimates distributed around zero and mostly statistically insignificant.

It is possible that the predictability in the profits data not due to risk is due to informational inefficiencies in the market. This is 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. Popular explanatory variables are the (lagged) forward premium, and lagged changes in the spot exchange rate. Regression results are reported in Tables 6 and 7 respectively. For a number of companies these explanators are statistically significant, but in most cases these variables are not statistically significant predictors of profits[7]. In all cases the R^2's for these regressions are extremely small. Similar results obtain using profits from RL-reg.

---

[7] 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.

6. Conclusion.

This paper examined 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 small positive profits from a trading rule based on the relative position of the subjective forecast to the forward rate. The profits that could have been earned are highly variable and thus there is significant risk in using these strategies. 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 often held to imply rejection of forward market efficiency. The regression-based profits are also highly variable. The high variability of profits questions the economic significance of the profits.

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

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, it does 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 random noise.

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, 1887-1908.

Backus, D.K., S. Foresi and C.I. Telmer, 1996, Affine Models of Currency Pricing, NBER WP #5623, June.

Bekaert, G.,1995, The Time Variation of Expected Returns and Volatility in Foreign-Exchange Markets, Journal of Business and Economic Statistics, 13, 397-408.

Bekaert, G. and R.J. Hodrick,1992, Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets, Journal of Finance, 47, 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, 115-138.

Bilson, J.F.O.,1981, The Speculative Efficiency Hypothesis, Journal of Business, 54, 435-452.

Breen, W. L.R. Glosten and R. Jagannathan, 1989, Economic Significance of Predictable Variations in Stock Index Returns, Journal of Finance, 44, 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, 279-299.

Diebold, F.X. and J.A. Nason,1990, Nonparametric Exchange Rate Prediction, Journal of International Economics, 28, 315-332.

Dooley, M. and J. Shafer, 1984, Analysis of Short-run Exchange Rate Behavior: March 1973 to November 1981 in: D. Bigman and T.Taya, eds., Exchange Rate and Trade Instability: Causes, Consequences and Remedies (Ballinger: Cambridge, MA) 43-69.

Fama, E.F.,1984, Forward and Spot Exchange Rates, Journal of Monetary Economics, 14, 319-338.

Frankel, J.A. and K.A. Froot,1987, Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations, American Economic Review, 77, 133-153.

Froot, K.A. and R.H. Thaler,1990, Anomalies: Foreign Exchange, Journal of Economic Perspectives, 4, 179-192.

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, 88,829-853.

Hansen, L.P. and R.J.Hodrick, 1983, Risk Averse Speculation in the Forward Foreign Exchange Market: An Econometric Analysis of Linear Models in: J.A. Frenkel ed., Exchange Rates and International Macroeconomics (University of Chicago Press, Chicago) 113-152.

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. Srivastava,1984, The Covariation of Risk Premiums and Expected Future Spot Rates, Journal of International Money and Finance, 3, 5-30.

Ito, T.,1990, Foreign Exchange Rate Expectations: Micro Survey Data, American Economic Review, 80, 434-449.

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

Levich, R. and L.R. Thomas,1993, The Significance of Technical Trading-Rule Profits in the Foreign Exchange Market: A Bootstrap Approach, Journal of International Money and Finance, 12, 451-474.

Lewis, K.K.,1995, Puzzles in International Finance Markets in: G. Grossman and K. Rogoff eds., Handbook of International Economics (North Holland, Amsterdam),1913-1971.

Lucas, R.E., 1982, Interest Rates and Currency Prices in a Two-Country World, Journal of Monetary Economics, 10, 335-360.

McCallum, B.,1994, A Reconsideration of the Uncovered Interest Parity Relationship, Journal of Monetary Economics, 33,105-132.

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.

Sweeney, J., 1986, Beating the Foreign Exchange Market, Journal of Finance, 41, 163-182.

White, H.,1980, A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity, Econometrica, 48, 817-838.

Table 1. Heterogeneity						
Company	1 Month		3 Month		6 Month	
	Mean	t	Mean	t	Mean	t
Maximum	0.0047	3.04	0.0210	4.31	0.0282	6.50
Minimum	-0.0082	-4.24	-0.0209	-5.76	-0.0330	-4.67
Forward	0.0052	4.76	0.0010	0.36	-0.0088	-1.70

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. The variance of the estimate is $(B_T^{-1} J_T B_T^{-1})$ where $B_T = \frac{1}{T}\Sigma x_t^2$ and here $x_t = 1$ so $B_T = 1$, J_T is the spectral density of $\hat{u}_t^*$ (residual from (5)) at frequency zero (divided by 2π) estimated by running the autoregression $a(L)\hat{u}_t^* = \varepsilon_t$, where the lag length is selected by a Bayesian information criterion and setting $J_T = \hat{a}(1)^{-1}\hat{\sigma}_\varepsilon^2 \hat{a}(1)^{-1}$ where $\hat{a}(1)$ = sum of coefficients in $\hat{a}(L)$. (This is done 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	
	Forecast Error	Direction	Forecast Error	Direction	Forecast Error	Direction
Survey						
Maximum	1.3801	0.5709*	1.2621	0.6250**	1.2912	0.6345***
Minimum	1.0110	0.3831	0.9987	0.3359	0.9882	0.3529
Average	1.0252	0.5479	0.9911	0.5469	1.0026	0.4580
SR-reg	0.9894	0.6054**	0.9719	0.7305***	0.9915	0.6723***
RL-reg	1.0224	0.5642	1.0596	0.6466**	1.1593	0.7156***
Forward	1.0087	0.5019	1.0206	0.5234	1.0192	0.5966***

Notes: Numbers reported in columns labeled Forecast 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 labeled Direction are the percentage of times over the sample where the forecast correctly predicted the direction of the subsequent spot rate movements, and the statistical significance of the text of the null hypothesis that each forecast has a 50% chance of being correct at the 10%, 5%, and 1% levels are denoted by 1, 2, and 3 asterisks respectively (using non-overlapping data).

Table 3. Properties of Profits - Unconditional Means						
Forecast	1 Month		3 Month		6 Month	
	Mean	$t(H_0=0)$	Mean	$t(H_0=0)$	Mean	$t(H_0=0)$
Survey						
Maximum	0.007	2.50	0.017	2.45	0.021	1.56
Minimum	-0.009	-3.58	-0.025	-3.56	-0.032	-2.37
Average	0.006	2.00	0.012	1.40	-0.005	-0.31
SR-reg	0.010	3.61	0.026	2.15	0.044	2.50
RL-reg	0.010	3.81	0.022	1.86	0.043	3.03
Naive	-0.005	-1.64	-0.014	-1.07	-0.035	-1.39
RW	-0.001	-0.25	-0.004	-0.32	-0.018	-0.77

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 4. Profits on 'Risk Premium'									
Forecast	1 Month			3 Month			6 Month		
	β_0	β	R^2	β_0	β	R^2	β_0	β	R^2
Survey									
Maximum	-0.01 (-1.80)	0.67 (-1.20)	0.03	0.00 (0.60)	0.92 (-0.23)	0.04	0.02 (0.84)	1.48 (0.83)	0.11
Minimum	0.00 (1.20)	-0.21 (-4.55)	0.00	0.00 (0.15)	-1.01 (-4.05)	0.05	-0.03 (-2.60)	-0.51 (-1.57)	0.01
Average	0.01 (2.03)	0.09 (-3.30)	0.00	0.01 (1.35)	0.43 (-0.91)	0.01	0.00 (0.10)	0.77 (-0.27)	0.03
SR-reg	0.01 (3.36)	0.15 (-3.98)	0.00	0.03 (2.03)	0.26 (-1.33)	0.00	0.05 (2.82)	0.57 (-1.11)	0.02
RL-reg	0.01 (3.21)	0.24 (-3.19)	0.00	0.02 (3.34)	0.96 (-0.11)	0.06	0.05 (3.86)	0.57 (-1.21)	0.03
Naive	-0.00 (-1.14)	-0.28 (-5.43)	0.01	-0.03 (-2.33)	-4.09 (-1.88)	0.09	-0.05 (-2.31)	-2.22 (-0.94)	0.05
RW	0.00 (0.81)	-0.70 (-6.39)	0.04	-0.00 (-0.23)	-1.22 (-4.93)	0.07	-0.02 (-0.82)	-0.18 (-1.61)	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 where $x_t = [1,(f_{t,k} - \bar{s}_{t,k})]'$ and $\hat{u}_t^* = \hat{u}_{5t}\, x_t$. The maximum and minimum are based on point estimates of β.

Table 5. Profits on Cross Sectional Variance of Expectations									
Forecast	1 Month			3 Month			6 Month		
	β_0	β	R^2	β_0	β	R^2	β_0	β	R^2
Survey									
Maximum	-0.01 (-1.92)	19.69 (1.94)	0.01	-0.07 (-4.38)	41.57 (3.23)	0.08	-0.07 (-3.53)	28.99 (2.96)	0.06
Minimum	0.02 (3.06)	-26.33 (-2.88)	0.03	0.05 (3.01)	-30.68 (-2.54)	0.04	0.05 (1.99)	-22.00 (-2.03)	0.04
Average	0.01 (1.43)	-7.74 (-0.65)	0.00	0.00 (0.00)	10.26 (0.64)	0.00	-0.04 (-1.45)	17.18 (1.42)	0.02
SR-reg	0.01 (2.14)	-5.38 (-0.54)	0.00	0.04 (1.75)	-10.33 (-0.82)	0.01	0.10 (5.10)	-29.26 (-3.09)	0.08
RL-reg	0.02 (4.24)	-31.48 (-3.33)	0.04	0.07 (4.33)	-43.23 (-3.58)	0.09	0.10 (4.75)	-28.31 (-3.08)	0.09
Naive	-0.02 (-3.41)	33.11 (3.53)	0.04	-0.06 (-2.94)	42.99 (3.24)	0.08	-0.09 (-3.10)	32.08 (2.33)	0.08
RW	-0.01 (-1.07)	12.76 (1.25)	0.01	-0.02 (-1.03)	15.34 (0.91)	0.01	-0.03 (-1.08)	8.57 (0.50)	0.01

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 6. Profits on Forward Premium									
Forecast	1 Month			3 Month			6 Month		
	β_0	β	R^2	β_0	β	R^2	β_0	β	R^2
Survey									
Maximum	0.00 (1.15)	1.93 (1.15)	0.01	0.02 (2.69)	3.15 (2.64)	0.05	0.03 (2.27)	2.58 (2.45)	0.06
Minimum	-0.00 (-1.24)	-3.84 (-2.76)	0.04	-0.01 (-1.37)	-2.39 (-1.37)	0.03	0.01 (0.39)	-1.79 (-1.35)	0.03
Average	0.01 (1.29)	-0.23 (-0.17)	0.00	0.01 (1.16)	-0.49 (-0.25)	0.00	-0.01 (-0.50)	-0.39 (-0.27)	0.00
SR-reg	0.01 (2.55)	-1.27 (-0.94)	0.00	0.02 (2.12)	-0.54 (-0.33)	0.00	0.05 (2.94)	0.34 (0.29)	0.00
RL-reg	0.01 (3.03)	-0.43 (-0.32)	0.00	0.02 (3.23)	0.25 (0.19)	0.00	0.03 (3.06)	-1.09 (-1.19)	0.02
Naive	-0.01 (-3.28)	-3.84 (-2.49)	0.04	-0.03 (-2.33)	-4.09 (-1.88)	0.09	-0.05 (-2.31)	-2.22 (-0.94)	0.05
RW	-0.00 (-0.84)	-1.49 (-0.90)	0.01	-0.01 (-0.81)	-1.82 (-0.74)	0.02	-0.03 (-0.96)	-0.86 (-0.31)	0.01

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.

Table 7. Profits on Lagged Change in Spot									
Forecast	1 Month			3 Month			6 Month		
	β_0	β	R^2	β_0	β	R^2	β_0	β	R^2
Survey									
Maximum	0.00 (0.35)	0.10 (1.59)	0.01	0.00 (0.49)	0.23 (2.25)	0.06	0.02 (1.06)	0.18 (0.89)	0.03
Minimum	0.00 (1.08)	-0.19 (-3.14)	0.03	0.01 (1.14)	-0.19 (-2.78)	0.04	-0.01 (-1.17)	-0.29 (-2.64)	0.08
Average	0.01 (1.96)	-0.06 (-1.08)	0.00	0.01 (0.73)	-0.12 (-1.11)	0.01	-0.00 (-0.19)	-0.06 (-0.57)	0.00
SR-reg	0.01 (3.50)	0.05 (0.69)	0.00	0.02 (2.24)	-0.07 (-0.65)	0.01	0.04 (2.04)	0.09 (0.50)	0.01
RL-reg	0.01 (3.63)	-0.04 (-0.76)	0.00	0.02 (3.44)	0.13 (2.00)	0.02	0.05 (2.60)	0.27 (1.11)	0.07
Naive	-0.00 (-1.40)	0.08 (1.01)	0.01	-0.01 (-0.85)	0.10 (0.72)	0.01	-0.03 (-1.08)	-0.14 (-0.47)	0.02
RW	0.00 (0.01)	0.09 (1.17)	0.01	0.00 (0.15)	0.12 (1.01)	0.01	-0.00 (-0.18)	0.07 (0.40)	0.00

Notes: As per Table 6 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).

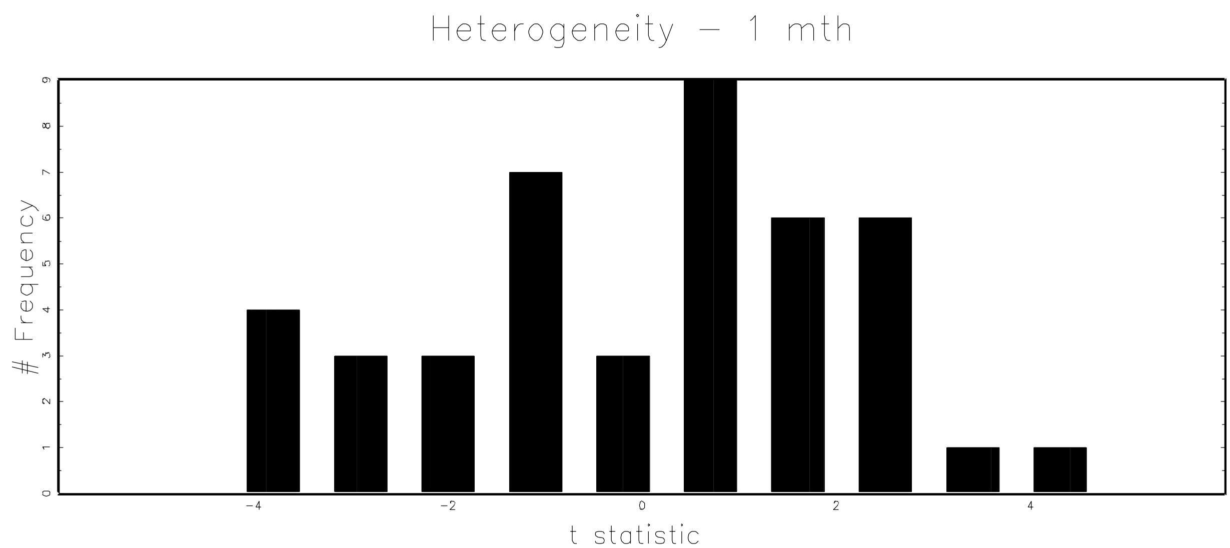
Heterogeneity – 1 mth
Frequency
t statistic

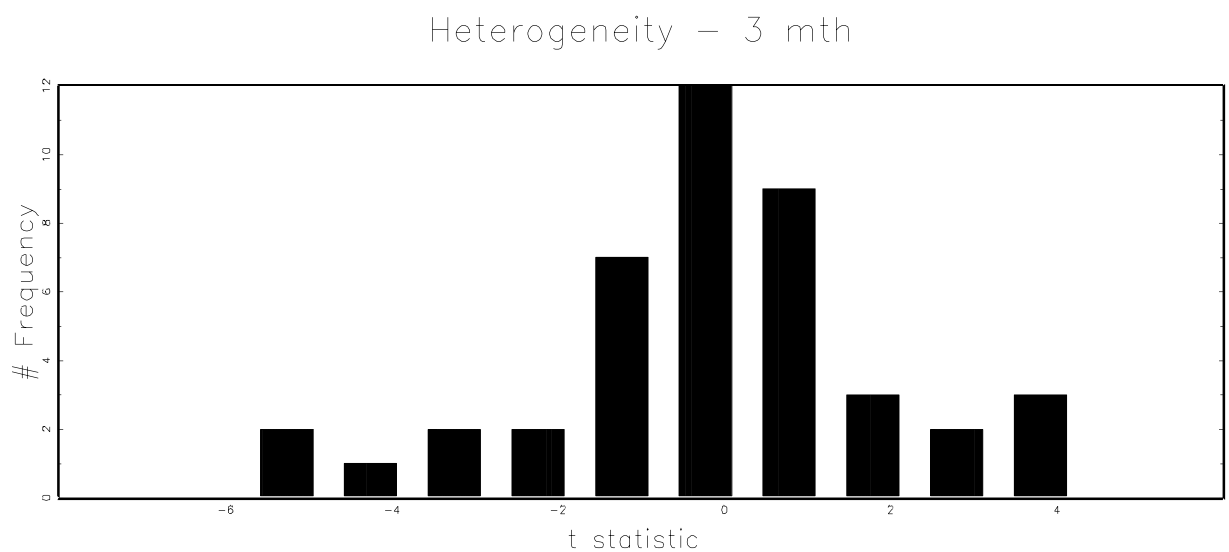
Heterogeneity – 3 mth
Frequency
t statistic

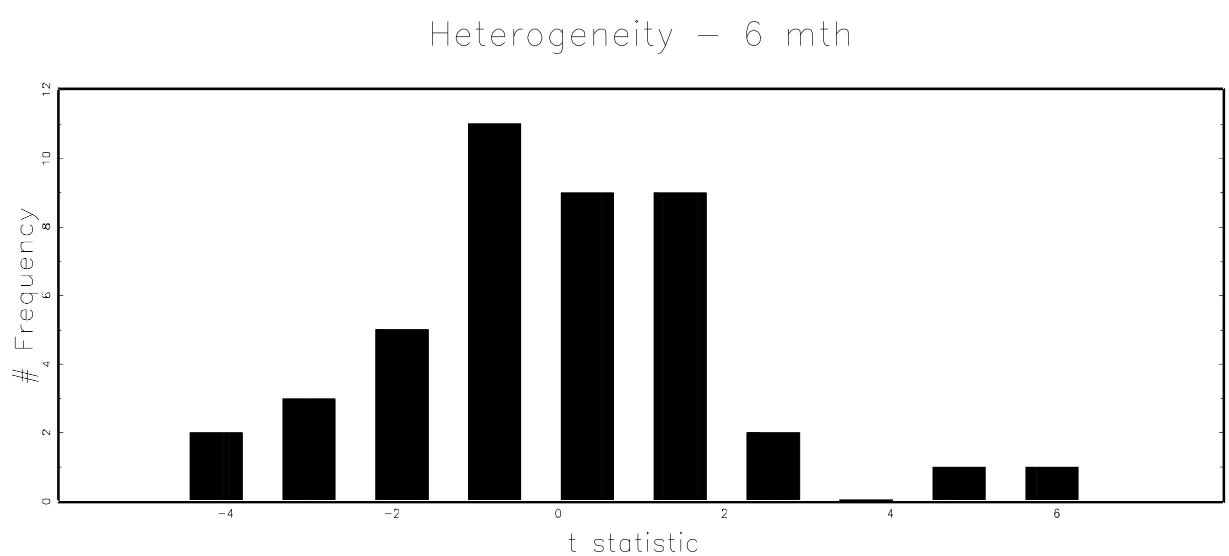
Heterogeneity – 6 mth
Frequency
t statistic

Forecast Variance – 1 mth

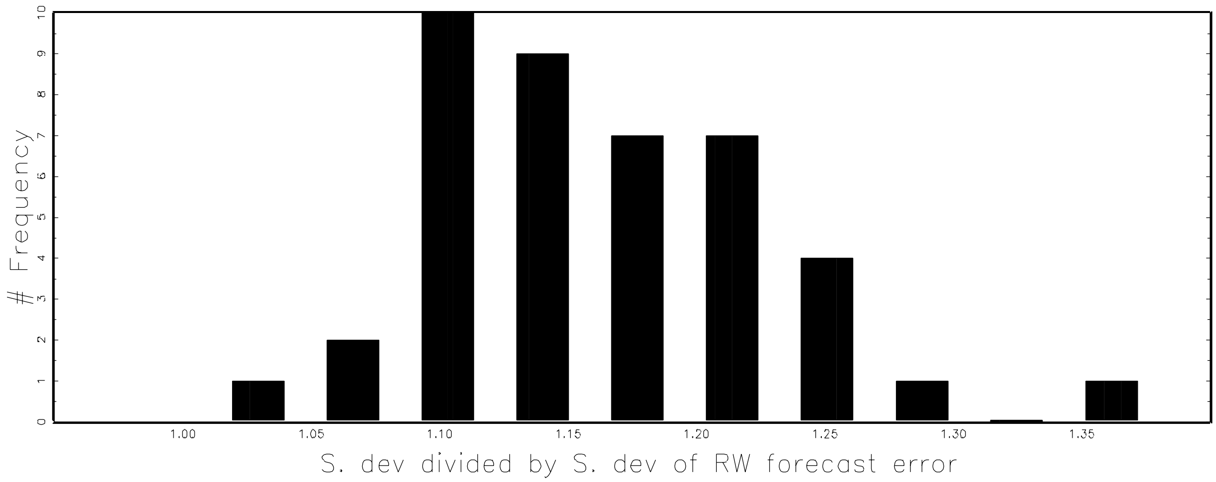

Forecast Variance – 3 mth

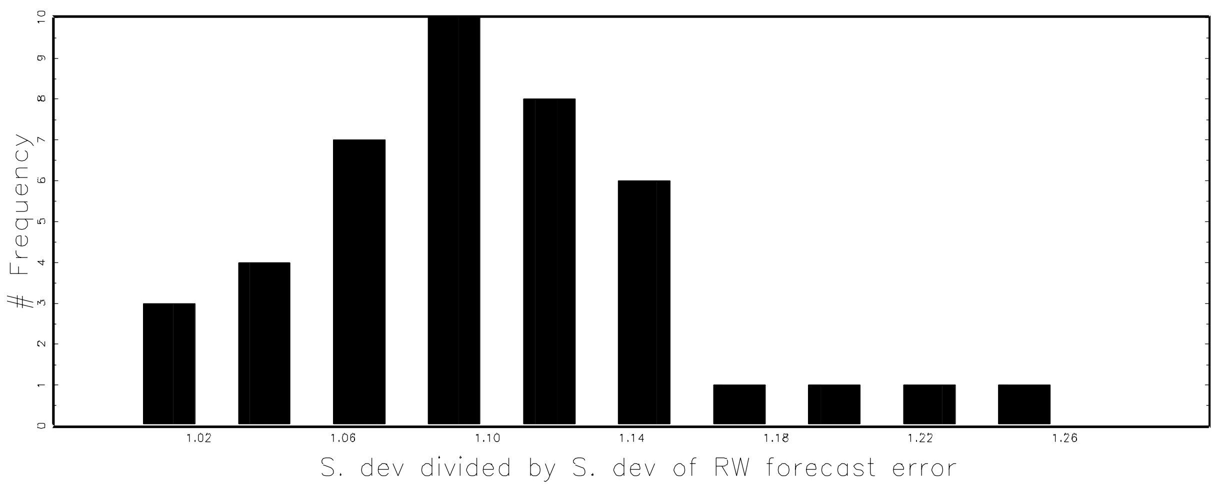

Forecast Variance – 6 mth

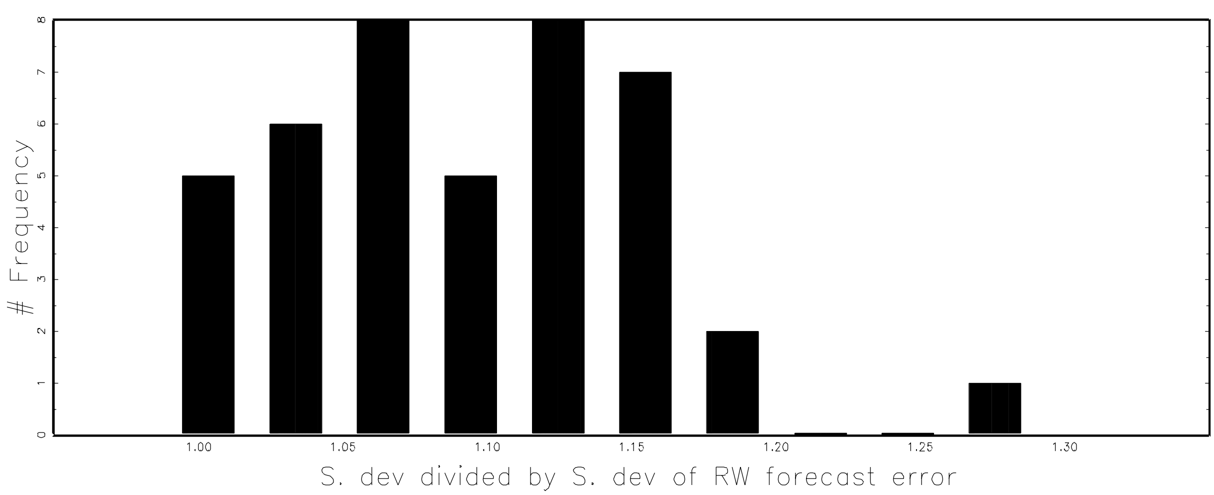

Forecast data to Explain ex post Depreciation – 1mth

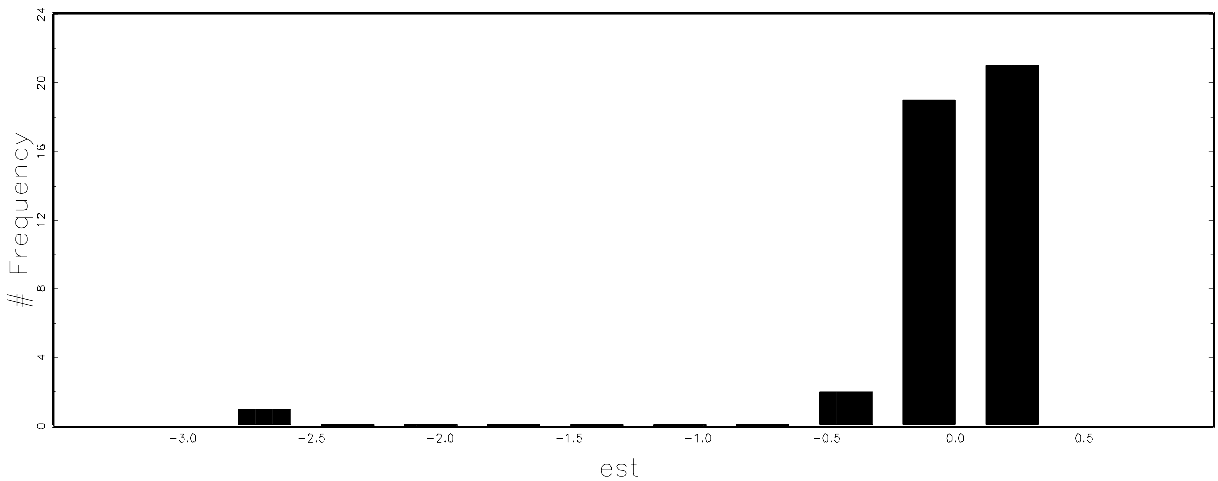

Forecast data to Explain ex post Depreciation – 3mth

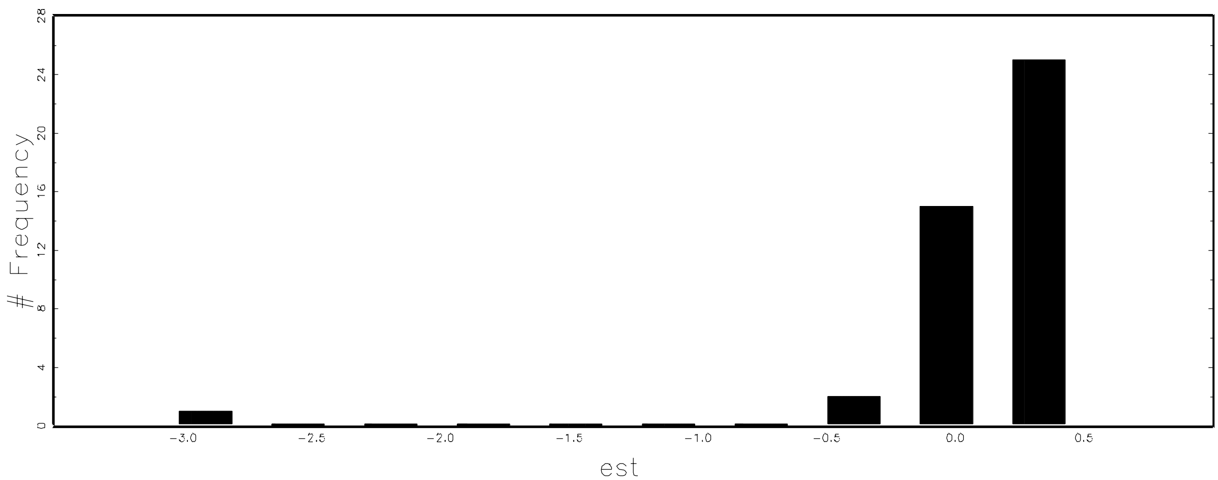

Forecast data to Explain ex post Depreciation – 6mth

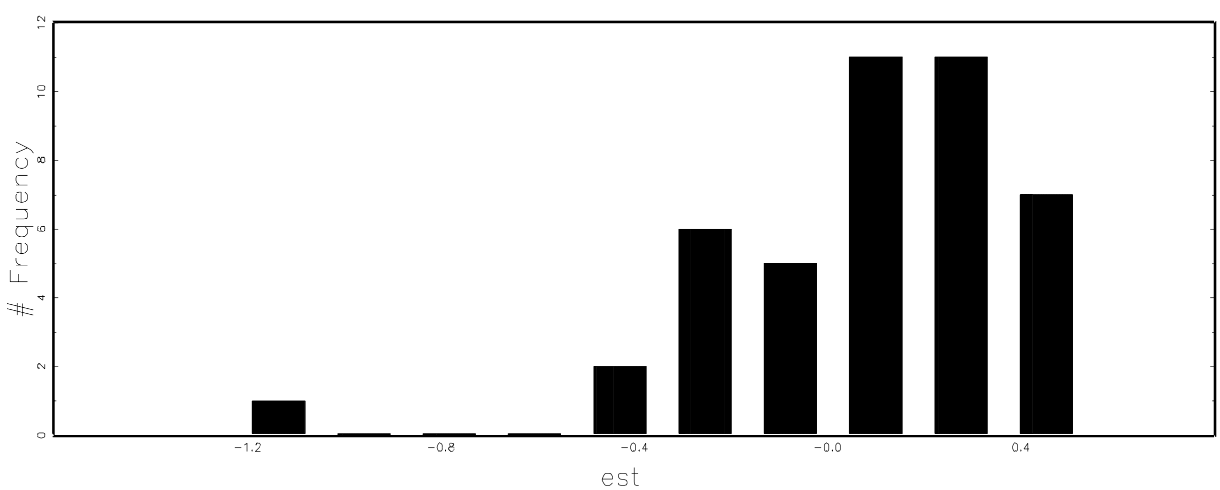

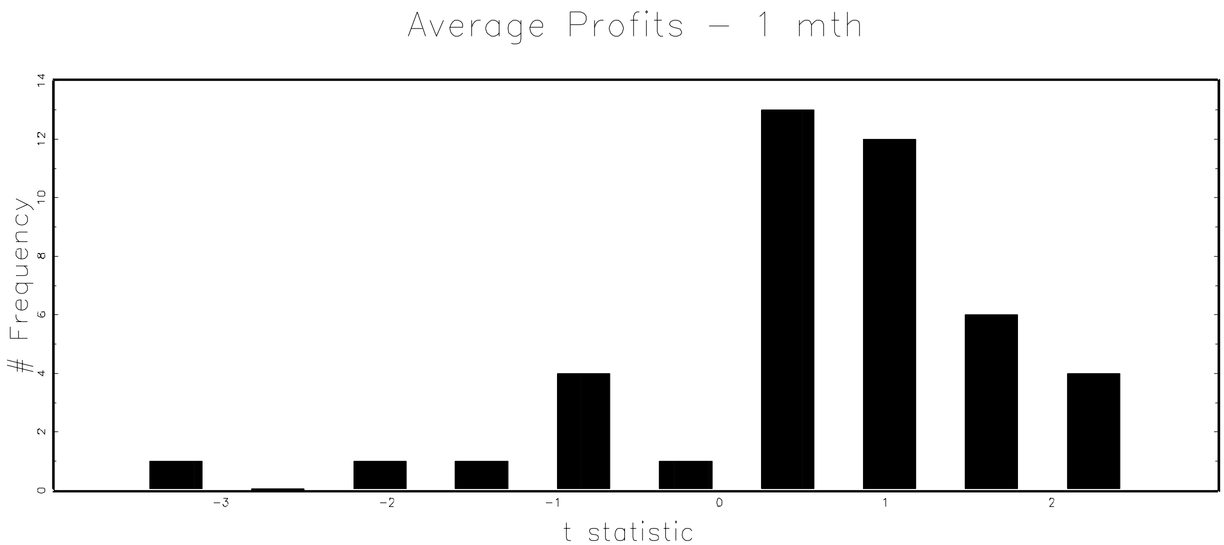
Average Profits – 1 mth
Frequency
t statistic
0
2
4
6
8
10
12
14
−3
−2
−1
0
1
2

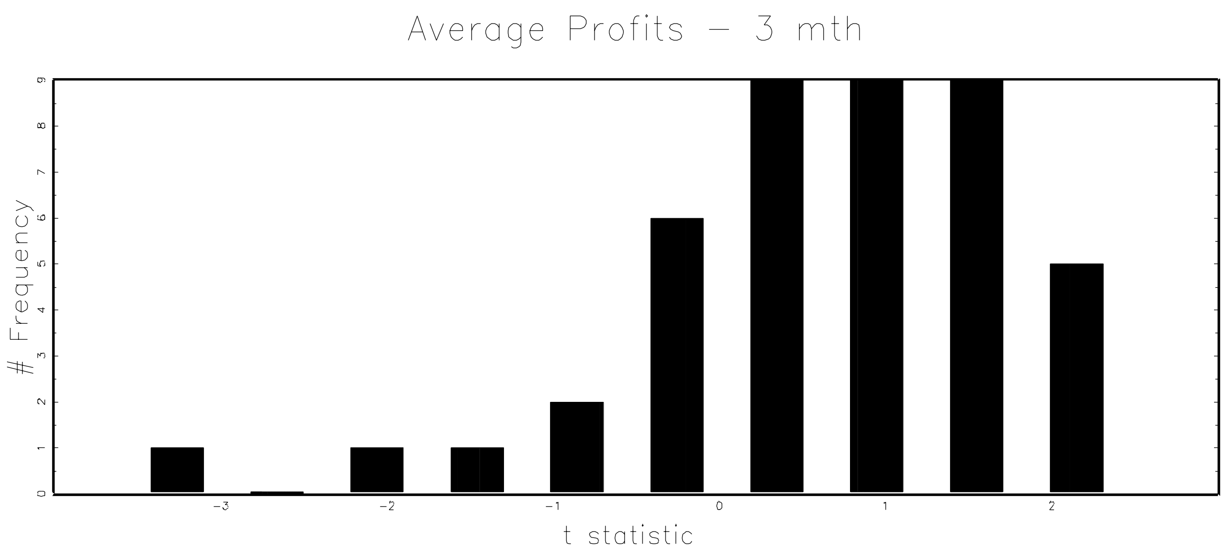
Average Profits – 3 mth
Frequency
t statistic
0
1
2
3
4
5
6
7
8
9
−3
−2
−1
0
1
2

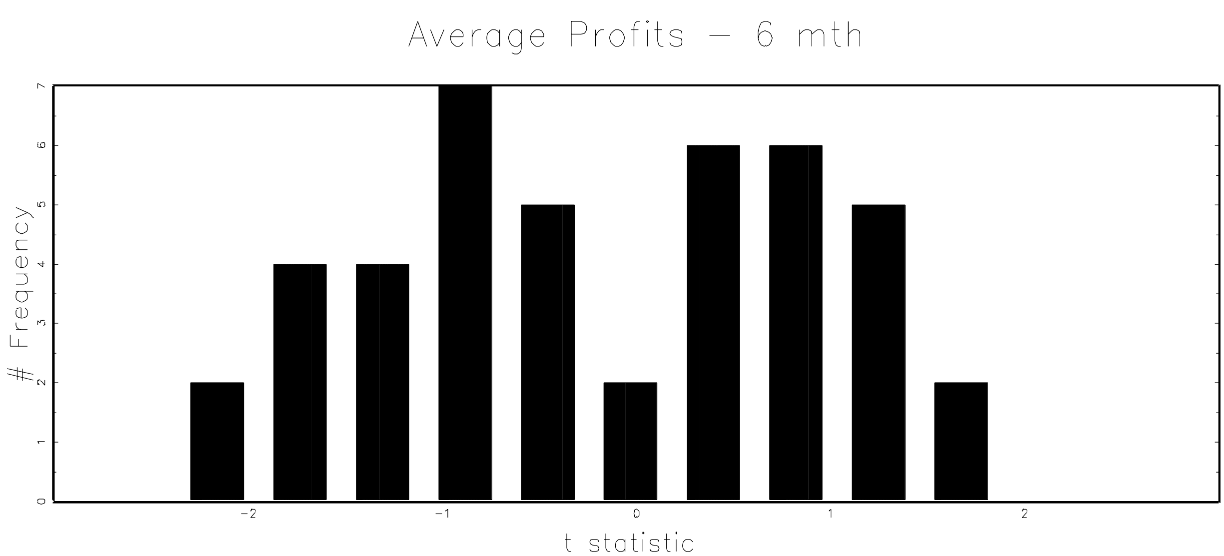
Average Profits – 6 mth
Frequency
t statistic
0
1
2
3
4
5
6
7
−2
−1
0
1
2

Profits on 'risk prem' – 1 mth

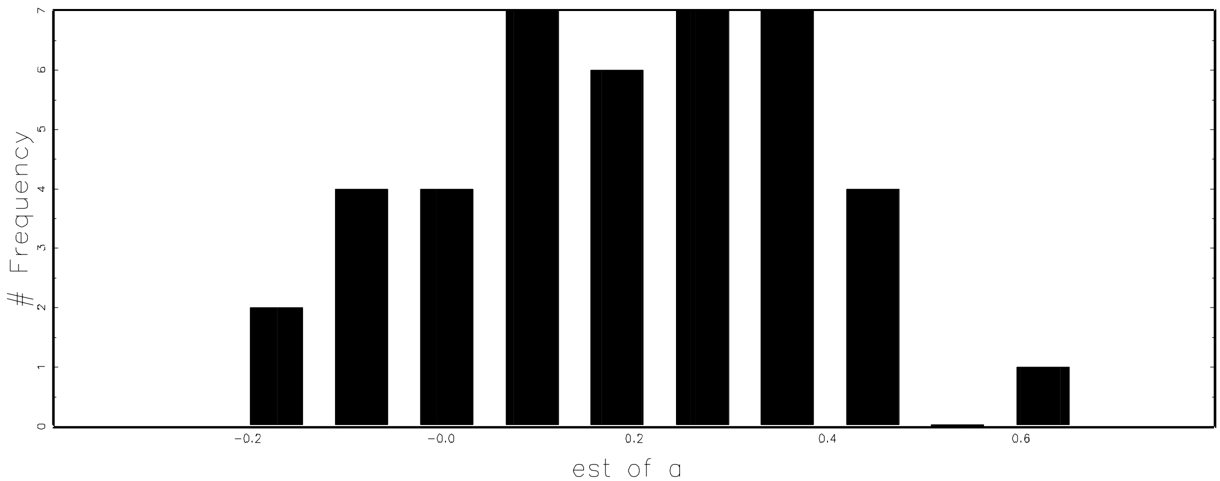

Profits on 'risk prem' – 3 mth

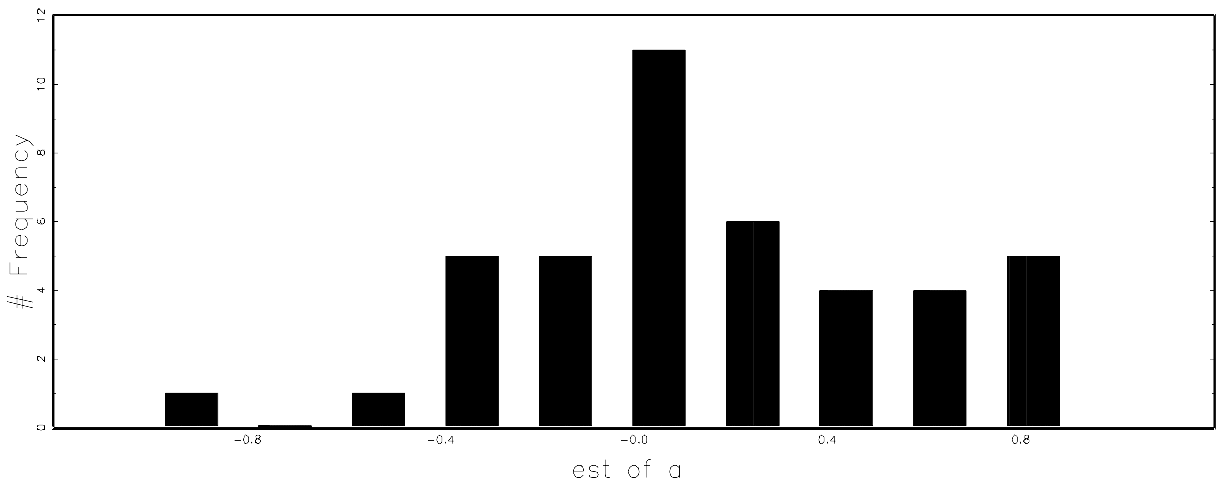

Profits on 'risk prem' – 6 mth