John F. O. Bilson
Hoover Institution, University of Chicago, and National Bureau of Economic Research

The "Speculative Efficiency" Hypothesis*

This paper proposes and implements a direct test of the hypothesis that profits from foreign exchange speculation are not significantly different from zero. A system of forecasting equations for the forward rate bias is estimated using Zellner's seemingly unrelated regression model. A speculative portfolio model is constructed based upon these estimates, and a postsample simulation of the model is conducted over 1980. On the basis of these results, the hypothesis that the mean of the distribution of speculative profits is zero is rejected.

Introduction

The hypothesis that forward prices are the best unbiased forecast of future spot prices is often presented in the economic and financial analysis of futures markets. In economic analysis, the hypothesis often appears under the guise of rational expectations, while in the finance literature the term "market efficiency" is more generally employed. It is important to note, however, that best unbiased forecasting by the forward price is not a necessary component of either a rational expectations or an efficient markets approach. It is easy, for example, to construct a model in which market expectations are rational, in the sense of being based upon the forecasts implicit in the model itself, but in which forward prices are not equal to expected future spot prices because of transactions costs, information costs, and risk aversion.[1] In the same manner and for the same reasons, it is possible to

*This research is part of the National Bureau of Economic Research Program in International Studies and is financed in part by a grant from the National Science Foundation. The views expressed are solely the responsibility of the author and are not necessarily shared by the sponsoring organizations. The author extends thanks to Greg Smith and Lyle Pai for research assistance and to Jacob Frenkel, Michael Mussa, and Arnold Zellner for helpful comments on an earlier draft.

1. The conditions under which the equality between forward rates and expected future spot rates will be disturbed by a risk premium have been discussed in a number of recent papers. For a useful discussion of the issue, see Fama and Farber (1979) and Frankel (1979).

(*Journal of Business*, 1981, vol. 54, no. 3)

0021-9398/81/5403-0005$01.50

construct a model in which markets are efficient in the sense of removing any opportunities for risk-adjusted excess returns but in which there is a predictable bias in the forward price forecast. For these reasons, this paper considers the hypothesis independently of its implications for rational expectations or market efficiency and in order to stress this fact, the term "speculative efficiency" is used to characterize the state envisaged under the hypothesis. If a market is subject to efficient speculation, the supply of speculative funds is infinitely elastic at the forward price that is equal to the expected future spot price. The expected future spot price is a market price determined as the solution to the underlying rational expectations macroeconomic model. Although the paper is primarily concerned with testing this hypothesis in the foreign exchange market, the methodology introduced in the paper is of general application to all futures markets.

Given that speculative efficiency is not directly related to either market efficiency or rational expectations, it is necessary to begin with some motivation for the investigation of the topic. First, the speculative efficiency hypothesis provides a strong and simple theory of the determination of forward prices and it has, for this reason, been widely employed in theoretical and empirical investigations of the determinants of the value of a currency.[2] Second, the issue is important in the evaluation of the net worth of foreign exchange forecasting services since the market for these services is based upon the assumed absence of an adequate, free, market forecast.[3] Third, the bias in the forward rate forecast is a critical determinant of the cost of covering foreign exchange exposure. Recent studies have suggested that if the forward rate is an unbiased forecast, the correct long-run hedging strategy is to not cover foreign exchange exposure because transactions costs are greater in the forward market than in the spot market.[4] Finally, the case for government intervention in the foreign exchange market is often based upon the judgment that an unregulated foreign exchange market sets forward prices at levels which are unrelated to underlying economic conditions, particularly during periods of disorderly market conditions. Statistical evidence that forward rates are unbiased forecasts of future spot rates would presumably weaken this case for intervention.

Perhaps the strongest evidence on the importance of the hypothesis is the very large and rapidly growing number of papers addressed to the

2. For a survey of models that assume that forward rates equal expected future spot rates, see Bilson (1979).

3. Goodman (1979) and Levich (1979) survey the performance of commercial forecasting services. *Euromoney* magazine also provides a less formal annual evaluation.

4. See Aliber (1978). In addition, the uncertainty of foreign currency denominated cash flows may make a fully hedged position risky.

TABLE 1 **The Standard Deviation of the Rate of Depreciation and the Forward Premium**

Country	$\sigma_{\Delta s}$	σ_x	$\sigma_{\Delta s}/\sigma_x$
Canada	14	2	7.9
U.K.	30	4	8.4
Belgium	30	5	5.9
France	28	4	7.3
Germany	32	2	14.1
Italy	26	9	3.0
Netherlands	29	3	9.1
Switzerland	40	3	11.8
Japan	32	4	7.3

NOTE.—$\sigma_{\Delta s}$ = standard deviation of 1-month rate of depreciation. σ_x = standard deviation of 1-month forward premium. Definitions and sources of exchange rate data are described in the Data Appendix. Units of measurement are annual percentage changes.

topic.[5] However, the large number of papers also demonstrates that the issue is not one that is easy to resolve through an empirical investigation. Different studies tend to reach different conclusions even when investigating the same set of exchange rates over roughly the same sample period. This lack of consensus in the econometric research is reflected in an even wider variance of opinion in less formal evaluations that often form the basis for policy analysis.[6] This is clearly not a desirable state of affairs.

The main reason why a strong empirical test of the speculative efficiency hypothesis is difficult to perform is that the variance of the change in the exchange rate is very large relative to the variance of the forward premium. This fact is illustrated in table 1, which lists the standard deviation of the 1-month rate of depreciation, the 1-month forward premium, and their ratio. These statistics are calculated for all of the major currencies over the period from July 1974 to January 1980.

The large difference between the standard deviation of the actual rate of depreciation and the standard deviation of the forward premium is a natural consequence of the fact that the forward premium is an estimate of the expected rate of depreciation. Any estimate of the mean of a series will always have a smaller standard deviation than that of the raw series. The difficulty is that the large variance in the actual rate hinders tests of the hypothesis that the forward premium is an unbiased

5. A short list of recent contributions includes Cornell (1977), Frenkel (1977), Geweke and Feige (1979), Hansen and Hodrick (1980), Hakkio (1980*b*), Longworth (1979), and Tryon (1979).

6. As an example, consider the following statement by Richard Cooper (1976): ". . . The forward exchange rates do a very poor job of forecasting future spot rates, and they do a worse job under flexible rates than they did under fixed rates. Under flexible rates, they have been strongly biased as well as inaccurate estimators. This means that there have been substantial opportunities for speculative profit" (p. 577).

estimate of the true expected rate because the confidence limits for the true expected rate are quite large.

The conventional test of the speculative efficiency hypothesis involves a regression of the actual rate of depreciation on the forward premium.[7] The regression equation is:

$$\Delta s_t = \beta_0 + \beta_1 x_{t-1} + u_t, \tag{1}$$

where Δs_t is the actual rate of depreciation, x_t is the forward premium, and u_t is an error term that may be serially correlated. Tests of the hypothesis that the forward premium is an unbiased forecast typically involve testing the joint hypothesis $\beta_0 = 0$, $\beta_1 = 1$, and $E(u_t, u_{t-i}) = 0$ for $i \neq 0$. Extensions of the analysis include the addition of other regressors that are available in period $t - 1$ including, for example, lagged rates of depreciation, lagged values of the forward premium, or lagged forecast errors. In these extended models, the speculative efficiency hypothesis includes the requirement that these additional variables are jointly insignificant in the regression equation. For the most part, this paper will be concerned with the simple formulation in equation (1), although one extended model will be considered below.

The ordinary least-squares estimates of equation (1) are presented in table 2. For expositional purposes, the estimating equation is written as

$$\Delta s_t - x_{t-1} = \beta_0 + (\beta_1 - 1) x_{t-1} + u_t. \tag{2}$$

This form has the advantage that the value of both the constant and the slope coefficient are zero under the null hypothesis. In addition, the R^2 statistic yields an estimate of the fraction of the variance in the forecast error that is accounted for by the lagged value of the forward premium.

These results, which are typical of others reported in the literature, yield mixed evidence on the speculative efficiency hypothesis. On the one hand, it is not possible to reject the hypothesis that the constant and slope coefficients are zero in all but two of the cases studied, and this failure to reject is particularly strong in the case of the major currencies—the pound sterling, the deutsche mark and the Swiss franc. On the other hand, significant serial correlation is indicated in three of the nine cases studied, and the hypothesis on the coefficients is rejected in the cases of the Italian lira and the Dutch guilder. Critics of the market efficiency hypothesis would also point to the fact that all of the estimated coefficients bar the pound sterling on the forward premium are closer to -1 than to zero. In other words, the sample evidence is also not inconsistent with the hypothesis that the forward premium is unrelated to the actual rate of depreciation.

All in all, these results demonstrate that the conventional tests of speculative efficiency in the foreign exchange market are not very

7. The methodology behind this test is due to Fama (1975).

TABLE 2 **OLS Estimates of the Equation**
$\Delta s_t - x_{t-1} = \beta_0 + (\beta_1 - 1) x_{t-1} + u_t$

Country	β_0	$\beta_1 - 1$	R^2	SE	D-W	$F(2,71)$
Canada	−4.010	−1.804	.050	14.17	2.06	2.35
	(2.05)	(.92)				
U.K.	1.928	−.372	.002	30.52	1.57†	.20
	(5.21)	(.99)				
Belgium	5.270	−.973	.027	29.88	2.18	1.29
	(3.86)	(.69)				
France	.407	−1.849*	.061	27.72	2.32	2.62
	(4.39)	(.86)				
Germany	6.737	−1.208	.007	31.91	2.33	.37
	(6.29)	(1.65)				
Italy	−7.428	−1.372*	.178	26.44	1.50†	8.04*
	(4.73)	(.35)				
Netherlands	7.285*	−2.741*	.086	29.38	2.32	3.52*
	(3.61)	(1.05)				
Switzerland	11.32	−1.184	.010	40.11	1.99	.48
	(8.52)	(1.38)				
Japan	3.917	−1.665	.049	32.50	1.57†	1.89
	(4.03)	(.86)				

NOTE.—Standard errors are in parentheses beneath the coefficients. An asterisk indicates that the estimated coefficient is significantly different from zero at the 5% significance level. The F-statistic tests the joint hypothesis: $\beta_0 = \beta_1 = 0$. The 5% critical value for the F-distribution with 2 and 71 degrees of freedom is 3.12. An asterisk indicates a calculated F-statistic which exceeds this value. A dagger indicates that the estimated Durbin-Watson statistic (D-W) falls below the lower bound for this statistic, thus indicating the presence of significant first-order autocorrelation in the residuals. The estimation period includes 4-weekly observations from July 1974 to January 1980.

strong. There are a number of approaches that one could take to increase the strength of the tests. First, higher frequency data could be used since spot and forward exchange rates are published on a daily basis. This approach encounters econometric difficulties, however, since overlapping forecasts induce serial correlation in the residuals. Hakkio (1980*a*) and Hansen and Hodrick (1980) have proposed econometric solutions to this problem and have successfully tested the speculative efficiency hypothesis on weekly data. Second, it is not necessary to restrict the tests to one maturity since forward rates are typically available for maturities of 1, 3, 6, and 12 months. Joint tests of the efficiency of the entire term structure of forward rates, as undertaken by Hakkio (1980*b*), also promise some increase in econometric efficiency.

There is a difficulty, however, with both of these approaches. Although they effectively increase the number of observations, the new observations are not greatly different from those included in studies using either monthly data or one maturity. Since the precision of the estimates depends upon the sample variance of the forward premium as well as on the number of observations, the gain in efficiency resulting from an increase in the number of observations is likely to be small.

For this reason, this study adopts an alternative approach in which the observations are aggregated both across countries and across time. Since the mean rate of depreciation has typically shown greater variance across countries than across time, this approach does offer a means of increasing the sample variance of the forward premium. Specifically, the following equation will be estimated:

$$\Delta s_{nt} - x_{nt} = \beta_0 + \beta_1 x_{nt} + u_{nt}, \tag{3}$$

where n is an index of currency (or country) and t is an index of time. The estimation procedure that is appropriate for this pooled cross-section time-series regression clearly depends upon the nature of the error structure.

The first, simplest, and least realistic description of the error structure is the classical assumption that the errors are spatially and temporally uncorrelated. This assumption leads directly to the following ordinary least-squares estimates:

$$\underset{(1.18)}{\Delta s_{nt} - x_{nt} = 3.123} - \underset{(.18)}{0.822}\, x_{nt} + u_{nt}, \tag{4}$$

$$R^2 = .029,\ F(2{,}664) = 8.46,\ \mathrm{SE} = 29.94.$$

The gain in efficiency from the cross-sectional restrictions on the coefficients is very evident in these results. For example, the standard error of the slope coefficient, which averaged .97 in the individual country regressions, is reduced to .18 in the pooled regression. The slope coefficient is now estimated with sufficient precision for the hypothesis β_1 equals unity to be rejected by the sample data with a high degree of confidence. The F-statistic tests the joint hypothesis that both the constant and the slope coefficient are zero. It is again possible to reject this hypothesis since the critical value of the F-distribution with 2 and 600 degrees of freedom is approximately 4.60 at the 1% level of significance. The cross-section evidence consequently offers a firm rejection to the hypothesis that the forward premium is an unbiased forecast of the future spot rate.

An alternative hypothesis is that the forward premium is purely determined by nonspeculative factors and is consequently unrelated to the future rate of depreciation. In this case, the null hypothesis is that the slope coefficient is not significantly different from minus unity. This alternative hypothesis cannot be rejected in this sample. Another approach to the alternative hypothesis is the joint test: $\beta_0 = 0$ and $\beta_1 = 0$. This model may be stated as the hypothesis that the expected rate of depreciation (a) is zero and (b) is unrelated to the forward premium. The calculated F-statistic for this hypothesis is $F(2{,}664) = 3.61$, which is above the 5% critical value but below the 1% for these degrees of freedom. Hence even this extreme alternative is more consistent with the sample data than the speculative efficiency hypothesis.

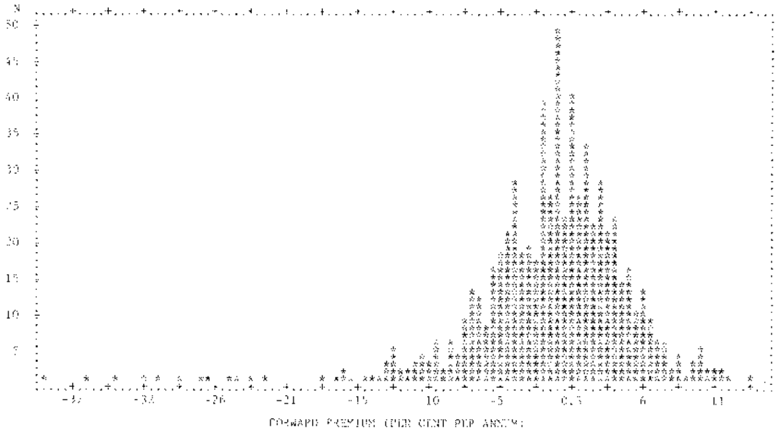

FIG. 1.—Interval histogram for forward premium. Nine countries, July 1974–January 1980.

In any regression analysis, there is always a possibility that the results are dominated by a small number of outlying observations. This is certainly the case in the present instance because the distribution of the forward premium, shown in figure 1, is highly skewed. Before proceeding, it is consequently worthwhile to check if the preceding results can be attributed to these extreme observations. This is done by dividing the observations into two groups—those less than 10% in absolute value and those greater than 10% in absolute value. Entering both variables in the regression resulted in the following estimates.

$$\underset{}{\delta s_{nt} - x_{nt}} = \underset{(2.53)}{2.986} + \underset{(0.53)}{0.091}\, x_{nt}^{\text{SMALL}} - \underset{(0.26)}{1.235}\, x_{nt}^{\text{LARGE}} + u_{nt}, \tag{5}$$

$$R^2 = .041,\ F(2{,}663) = 11.453,\ \text{SE} = 29.89.$$

These results lead to a partial rehabilitation of the speculative efficiency hypothesis since they demonstrate that the negative results in the previous tests were due to the large outlying observations. For values of the forward premium in the range of $\pm 10\%$, the forward premium does appear to be an unbiased forecast of the future actual rate of depreciation. The large value of the standard error of the coefficient on the x_{nt}^{SMALL} variable demonstrates, however, that the predictive power of this variable is quite small.

As far as the large values are concerned, the weight given to the forward premium is not only negative but also greater than unity. This implies that these values, on average, do not even capture the direction of change. More to the point, the statistics suggest that values of the forward premium in excess of 10% should be ignored for forecasting purposes. In terms of the debate over the efficiency of the forward exchange rate, the existence of "destabilizing" speculation, and the presence of "disorderly market conditions," the results suggest that episodes do exist in times of extreme stress in which the forward rate is not an unbiased forecast of the future spot rate. A risk averse speculator—private or public—would appear to be justified in intervening in the market when the forward premium exceeds 10%. For the most part, however, the forward rates are unbiased forecasts.

Generalized Least-Squares (GLS) Estimation

The previous ordinary least-squares (OLS) estimates have the virtue of being unbiased under the reasonable condition that the forward premium is uncorrelated with the future forecast error. These estimates are not very efficient, however, because they do not make use of the strong contemporaneous correlation in the forecast errors. On theoretical grounds, there are two major reasons why one would expect strong contemporaneous correlation in the forecast errors. In the first place, all of the exchange rates are expressed relative to the U.S. dollar, so that shocks originating in the American monetary sector will influence all of the exchange rates. Second, there have been a number of formal and informal arrangements for exchange rate coordination, including the "snake," the "tunnel," and the recently created European Monetary System, that insure a high correlation between the forecast errors for the European currencies. These considerations are incorporated into the estimation through the assumption that $E(u_{nt}, u_{n't'}) = \sigma_{nn'}$ if $t = t'$, zero otherwise.[8] The results from this estimation procedure are given below.

$$\Delta s_{nt} - x_{nt} = \underset{(1.38)}{-2.434} - \underset{(0.16)}{0.844}\, x_{nt}^{\mathrm{SMALL}} - \underset{(0.15)}{1.322}\, x_{nt}^{\mathrm{LARGE}} + u_{nt}. \tag{6}$$

Apart from the significant increase in the precision of the estimates, the most noticeable feature of the results reported in equation (6) is that the weight given to x_{nt}^{SMALL} is now larger and significantly different from zero at the 5% level. These estimates consequently do not support the speculative efficiency hypothesis even for moderate values of the forward premium. It is also noticeable, however, that the large values exhibit an even greater bias and that there is still good reason for distinguishing between small and large values of the premium.

Given that the results from the two estimation procedures are quite different, it is necessary to address the question concerning which procedure is most appropriate. On econometric grounds, the GLS estimates are clearly superior since they make use of information that is not employed by the ordinary least-squares estimation. There is,

8. These assumptions lead to the "seemingly unrelated regression model" introduced by Zellner (1962). Frankel (1980) has also employed this model in testing the hypothesis that forward rates are unbiased forecasts of future spot rates.

however, another reason for preferring the GLS estimates. In assessing the forecasting efficiency issue, one is essentially examining if there is an alternative forecasting rule that is statistically superior to the forward rate. Presumably, the object of the analysis, then, is to determine if there were opportunities for speculative profit during the sample period if this alternative rule were known. For a completely risk neutral speculator, either estimate will suffice since both have the desirable property of being unbiased and risk neutrality implies that the variance of the estimates is not important. On the other hand, a risk neutral speculator will be concerned about the precision of the estimates and she will also be interested in the covariance matrix of the forecast errors. The reason is, of course, that some knowledge of this covariance can be employed to diversify the risk from foreign exchange speculation.

It is at this point that we can introduce an interesting feature of the GLS estimates: while the OLS procedure chooses the forecasting rule that minimizes the total variance of the forecast, the GLS procedure essentially minimizes the nondiversifiable variance of the forecast error.[9] Since the speculator is only interested in the nondiversifiable risk, the GLS procedure is clearly the most appropriate procedure to use. In the following section of the paper, the speculative issue will be examined in greater detail. Before moving to this topic, however, there is one additional restriction that should be imposed upon the previous estimates. The results included in the previous section included a constant term. Since there is little reason to presume that there is a constant bias in the forecast on theoretical grounds and since the estimated constant term is not significantly different from zero, the equation was reestimated with this parameter constrained to be equal to zero. The revised estimates are presented below:

$$\Delta s_{nt} - x_{nt} = \underset{(0.15)}{-0.741}\, x_{nt}^{\text{SMALL}} - \underset{(0.12)}{1.280}\, x_{nt}^{\text{LARGE}} + u_{nt}. \tag{7}$$

As one would expect, these estimates are not greatly different from the preceding ones.

The Speculation Issue

The preceding results have demonstrated that an alternative forecast would have been statistically superior to the forward premium during the sample period. These results do not, of course, indicate the exis-

9. The GLS procedure transforms the forecasting equation so that the forecast error for currency n is uncorrelated with all other contemporaneous forecast errors. It is consequently not possible to reduce the risk of a position in currency n by taking offsetting positions in other currencies. This point may also be interpreted through equation (10) below. This equation may be interpreted as the standard GLS formula for the vector of regression coefficients. In this case, however, the right hand variables are the expected profits and the endogenous variable is the position taken in the currency.

tence or nonexistence of "excess" profits from foreign exchange speculation since, first, the forecasting rule and its performance were not known to market participants during the sample period and second, the risk of forecasting on the basis of the rule has not been measured. These issues will be examined in this section of the paper.

The first task is to restate the previous results from the viewpoint of a speculative strategy. This will allow a direct test of the existence of the ex post speculative profits from a portfolio speculative strategy involving all of the currencies in the sample.

The analysis begins with a description of the speculative strategy. The profit from a portfolio speculative position may be defined as

$$\pi = q' r, \tag{8}$$

where q is a column vector whose elements are the dollar value of the forward purchases of the nth currency in period $t - 1$, and r is the realized dollar profit per period. On the basis of the previous results, r is presumed to be distributed $N(\bar{r}, \Omega)$, where $\bar{r}$ is defined as $-.741\ x_{nt}^{\text{SMALL}} - 1.280\ x_{nt}^{\text{LARGE}}$, and Ω is the covariance matrix of the forecast errors.

Following basic financial theory, the speculator is presumed to select the elements of q in order to minimize the variance of profits for a given expected level of profits. The analysis is couched in terms of profits rather than rates of return because speculative gambles do not require the placement of capital. (Margins may be required, but these can be placed in the form of Treasury bills which earn a market rate of return. Large corporations may deal in the interbank market without margin requirements.) This feature of the market also implies the absence of any summation or liquidity constraints. Finally, since the speculator can trade on either side of the market, there are also no nonnegativity constraints on the elements of the q vector.

With these simplifications, the speculator's problem can be stated as

$$\min_{q} L = q'\Omega q + \lambda(q'\bar{r} - \pi^*). \tag{9}$$

In equation (9), λ is a Lagrange multiplier defined as the marginal increase in variance for a marginal increase in target profit, π^*. The solution to this problem is

$$\hat{q} = \Omega^{-1}\bar{r}[\bar{r}'\Omega^{-1}\bar{r}]^{-1}\,\pi^*. \tag{10}$$

Equation (10) simply states that the optimal positions are proportional to the level of expected profit and that the factor of proportionality is a function of the covariance matrix and the vector of expected returns. If equation (10) is substituted back into the definition of the variance of profits, the "efficient frontier" equation is obtained:

$$\pi^* = k\ \sigma(\pi^*), \tag{11}$$

where $k = (\bar{r}'\Omega^{-1}\bar{r})^{.5}$. In words, the efficient frontier is a ray from the origin. This result follows from the fact that the speculator can avoid both profit and risk by not entering the market and that, as demonstrated in equation (10), the system is homogeneous of degree 1 in its nominal magnitudes.

From the point of view of testing, equation (11) is important because it allows a test of the existence of statistically significant profits independently of the scale of the position. If the model is correct, the random variable $(\pi - \pi^*)/\sigma(\pi^*)$ should be distributed $N(0,1)$. This hypothesis will be tested below.

Before proceeding to this test, it is worthwhile to discuss the operational aspects of the speculator's problem. If the speculator is to implement the strategy described above, the regression system must first be estimated in order to obtain the regression parameters and the covariance matrix of the forecast errors. The fact that widely available information—spot and forward rates—can be used to "beat" the forward rate must be recognized and the parameters of the forecasting model must be known before the strategy can be implemented. Consequently, in-sample tests of speculative efficiency are biased against the hypothesis because they assume that this information—the regression parameters and the covariance matrix—is available to speculators before they have been estimated. It is as if a regression of equity prices on a time trend were used to test for the existence of speculative profits in the stock market. As many followers of this approach have found, there is a world of difference between knowing the trend over the last 6 months and the trend over the next 6 months.

In the estimation of the forecasting model, a strong attempt was made to avoid this type of in-sample bias. The cross-sectional restrictions on the coefficients and the GLS estimation procedure allowed for two parameters to be estimated from a total of 666 independent observations. While this procedure may help to limit in-sample bias, it does not protect the model against parameter drift resulting from the accumulation of knowledge in the market. As the bias in the forward rate becomes known, it is likely that forces will be set into motion which will lead to the elimination, or at least the reduction, in the bias.

For these reasons, the tests of the model presented below will be based purely upon its postsample speculative performance. The regression model was estimated in February 1980 on data running up to January 18, 1980. The regression results were distributed as an NBER Working Paper (no. 474) in April 1980. The postsample results utilize both the regression parameters and the variance-covariance matrix estimated in this working paper. The only new information included in the tests are spot and forward exchange rates, which are widely and freely available to all market participants. Although the model was estimated using one observation per 4-week period, the postsample portfolios were constructed on a weekly basis from February 1, 1980 to January 16, 1981. Since the gamble is realized 4 weeks after it is made, the series of actual profits will follow a third-order moving average process. Although the serial correlation in the series does not bias the mean of the series, it will bias the standard estimates of the variance.[10]

The overlapping forecast problem will be dealt with below. Before taking on this task, however, the major conclusions from the study may be derived from a less formal evaluation of the performance of the model. For this purpose, the most important statistics are presented in table 3. The first series in the table, SRE, is the standardized expected profit, which is defined as the expected profit divided by the expected standard deviation. The average value of this variable is .929. Using the two standard deviations rule, an SRE of unity implies 95% confidence limits ranging from −1 to +3: for every dollar the speculator expects to make, the confidence limits range from a loss of $1.00 to a gain of $3.00. This is an extremely favorable risk/return trade-off for a single gamble, but it becomes even more so as the speculator engages in a number of independent trials. For example, if the speculator divides his initial dollar into two 50-cent gambles, the 95% confidence limits fall to −.41 to 2.82. Under the same conditions, the 95% confidence limits for four drawings from the distribution range from zero to two.

The same point can be made with the CRE series in table 3. This series is simply the cumulative value of the standardized expected profit series. After the first five trials, the CRE has a value of 2.85, which is significantly different from zero. After 50 trials, the CRE has increased to 46.57 standard deviations from zero. This discussion should serve to emphasize the true meaning of risk in foreign exchange speculation. If the speculator knew that the forecasting model was correct, then the risk from speculation is small. The situation is exactly akin to that of a casino operating a roulette wheel. The true source of risk is that the forecasting model may be wrong, because of some in-sample bias, or that it may change in the postsample period as the market evolves. Most of the studies of this problem have not considered this source of risk because they have been based upon in-sample statistics.

From this perspective, then, the important question is the extent to which the model accurately forecasts ex post bias in the forward exchange rate. The short answer to the question is very well. After the year of postsample trials, the cumulative standardized expected profit was 46.5 standard deviations from zero. The cumulative standardized actual profit, CRA, was 42.8 standard deviations from zero. Out of the total of 50 trials, the speculator would have only made 10 losses and

10. Hansen and Hodrick (1980) discuss this problem and propose a solution for the case in which the right hand side variables are not orthogonal to the forecast error at all leads and lags. Hakkio (1980*b*) also investigates the issue.

TABLE 3 Postsample Speculative Performance

Date	SRE	SRA	CRE	CRA
20180	.64	1.38	.00	.00
20880	.67	.91	.64	1.38
21580	.64	1.47	1.32	2.29
22280	.91	.76	1.95	3.77
22980	.55	1.21	2.87	4.53
30780	.56	.35	3.43	5.74
31480	.63	.18	3.99	6.09
32180	.63	.85	4.63	6.28
32880	.67	.89	5.26	7.13
40480	1.07	2.15	5.93	8.03
41180	1.22	2.03	7.01	10.18
41880	1.42	1.85	8.23	12.22
42580	1.15	1.95	9.66	14.07
50280	1.17	1.21	10.81	16.02
50980	1.22	1.42	11.98	17.23
51680	.94	2.09	13.21	18.66
52380	1.12	.65	14.15	20.75
53080	.79	2.40	15.28	21.41
60680	.68	3.29	16.08	23.81
61380	.81	2.77	16.76	27.10
62080	.94	1.43	17.58	29.88
62780	.91	1.06	18.52	31.31
70380	.95	.24	19.94	32.37
71180	.94	−.21	20.39	32.63
71880	.88	.63	21.33	32.41
72580	1.85	1.66	22.22	33.04
80180	1.69	1.14	24.07	34.71
80880	1.12	.46	25.76	35.85
81580	1.03	−.23	26.89	36.31
82880	1.24	−.40	27.92	36.09
82980	1.21	−.14	29.16	35.68
90580	1.21	.44	30.38	35.54
91280	1.65	1.07	31.59	35.98
91980	1.51	1.53	33.25	37.05
92680	1.41	.63	34.76	38.59
100380	.71	.11	36.17	39.22
101080	.69	.15	36.88	39.34
101780	.82	−.05	37.58	39.50
102480	.94	−.78	38.41	39.44
103180	.52	−.30	39.35	38.65
110780	.52	1.86	39.87	38.35
111480	.48	−.12	40.39	40.22
112180	.58	.35	40.88	40.21
112880	.61	.35	41.46	40.56
120580	.63	−2.36	42.07	40.91
121280	.54	1.41	42.71	38.54
121980	.62	1.18	43.26	39.96
122680	.63	.95	43.88	41.14
10281	1.15	1.40	44.52	42.10
10981	.90	−.66	45.67	43.50
11681	.84	.84	46.57	42.84

NOTE.—Date = month/day/year; SRE = standardized expected profit $[\pi^*/\sigma(\pi^*)]$; SRA = standardized actual profit $[\pi/\sigma(\pi^*)]$; CRE = ΣSRE; CRA = ΣSRA.

only one of those was greater than unity in absolute value. More importantly, the cumulative actual profit series never became negative throughout the sample period.

The final position does not, however, give a complete picture of the speculative performance over the year. Two distinct periods can be described. In the first half of the year, actual profits ran well ahead of speculative profits so that, by June 27, cumulative actual profits were 31 standard deviations while cumulative expected profits were only 17 standard deviations from zero. During the subsequent 6 months, the model did not perform as well as expected: by December 12, 1980 CRE has increased to 42.7 while CRA had only increased to 38.5. In the last month of the postsample, the two series appear to be following the same general trend. In purely factual terms, the reason for the excess profits in the early part of the period was a set of extremely profitable positions in the Italian lira. In part, the poor performance also reflects a set of positions in this currency. This experience suggests that the practice of constraining the coefficients to be the same for all currencies may be a costly one and that more work needs to be done on the imposition of cross-currency restrictions.

For a more formal test of the postsample accuracy of the profit forecasts, the following regression was estimated:[11]

$$\text{SRA} - \text{SRE} = \beta_0 + \beta_1 t + \Theta_1 v_{t-1} + \Theta_2 v_{t-2} + \Theta_3 v_{t-3} + v_t. \quad (12)$$

A third-order moving average process was imposed in this regression in order to take account of the overlapping forecast problem. If the model is correct, the bias, $\beta_0 + \beta_1 t$, should not be significantly different from zero for any value of t. Using an iterative estimation procedure, the following estimates of the equation were obtained.

$$\begin{aligned} \text{SRA} - \text{SRE} = {} & \underset{(0.302)}{0.5624} - \underset{(0.009)}{0.2437}\, t + \underset{(0.246)}{0.4786}\, v_{t-1} \\ & + \underset{(0.264)}{0.0442}\, v_{t-2} + \underset{(0.254)}{0.0746}\, v_{t-3} + v_t, \end{aligned}$$

$$\text{SE} = 0.94,\ R^2 \text{ (adjusted)} = .138,\ F(4{,}43) = 2.88,\ \text{D-W} = 1.94.$$

Standard errors are given in parentheses beneath the coefficients. The most important result from this regression is that the time trend bears a

11. Because of the absence of a transfer function computer program at the Hoover Institution, the moving average model was estimated by an iterative regression technique. Initially, the residuals were expressed in purely autoregressive form and the "whitened" residuals were estimated. The lagged values of these whitened residuals were then entered into the regression equation as additional independent variables. The parameter estimates stabilized after four iterations in which new estimates of the residuals were obtained from the regression. In order to implement the procedure, the first four observations were lost.

significantly negative coefficient, thus suggesting that the profitability of the model is declining through time. It is also noticeable, however, that the constant term is positive so that the model was underestimating actual profits in the earlier part of the sample period.

Both of these factors can be combined in the form of a standardized bias measure: $B/\sigma(B)$, where B is defined as $\beta_0 + \beta_1 t$. As expected, this measure is positive in the earlier part of the sample period, but it declines throughout and exceeds two in absolute value in the last 3 months of the sample. At the end of the sample, the estimated bias was $-.78$, with a t-statistic of 2.53. If this bias is compared with the actual standardized expected profit of 0.83, the conclusion reached is that the predictable opportunities for speculative profit have been eliminated.

The final issue to be discussed is the variance of the profit forecast errors. Since both SRA and SRE are measured in standardized units, the variance of profit should be unity. On the basis of the regression estimates of equation (12), the actual variance of the profit forecast errors is estimated to be 1.09. Additional tests revealed no tendency for the variance of profits to increase through time.

Conclusion

In terms of general methodology, the main objective of this paper has been to introduce a direct test of nonzero speculative profit which relies on a combination of Zellner's seemingly unrelated regression model and modern portfolio theory. This testing procedure made use of the distinction between diversifiable and nondiversifiable forecast errors and demonstrated, through the use of this distinction, that predictable profits could be generated by data in which the variance of the total forecast error is large relative to the variation in the forecast. This methodology appears to have wide applicability for many other issues in finance and economics.

As far as the specific problem of speculative efficiency is concerned, the results lead to a firm rejection of the hypothesis that there are no predictable profits to be made from foreign exchange speculation. Through the use of a very simple forecasting rule, in conjunction with freely available information, predictable profits were realized in a simulated foreign exchange portfolio. These profits may be eliminated by transactions costs, or they may provide an adequate compensation for risk, and hence it is not possible to make inferences about the efficiency of the foreign exchange market on the basis of these results. There are, however, a number of points that can be made. In the first place, the profit/risk ratio appears to be too large to be accounted for in terms of risk aversion: there are very few people who would not participate in the game of drawing balls from an urn containing dollar returns with a mean and standard deviation of unity. This high profit/

risk ratio also precludes explanations based upon transactions costs; in any case, the tax advantages from currency speculation may be sufficiently large to offset these costs. In the second place, there are few theories of risk which posit that the risk premium is a function of the level of the forward premium on a currency. Finally, the model suggests that the forward premium is too volatile relative to the conditional forecast of the rate of appreciation of the exchange rate. If anything, risk aversion would tend to make the forward premium less volatile than the conditional forecast.

One explanation that may be consistent with the data is that forward rates are predominantly determined by interest rates which are, in turn, a reflection of the anticipated rate of inflation. For a given growth rate of the money supply, higher inflation rates reduce the quantity of real money balances and create upward pressure on the exchange rate. At the same time, the high interest rates induce a discount on the currency in the forward market. Under these conditions, an appreciation of the exchange rate would be correlated with a discount in the forward market. To the extent that monetary policy accommodates the inflation, both the appreciation of the exchange rate and the discount in the forward market would be mitigated and the results reported in the paper would emerge. At the present time, this interpretation of the results is simply a conjecture.

Data Appendix

The data on spot rates and 1-month forward rates were taken from the Harris Bank Weekly Bulletin, published by Harris Bank and Trust, Chicago. These data are collected every Friday, and every fourth observation was used to represent a 1-month contract. A printout of the actual data is available from the author on request.

References

Aliber, Robert Z. 1978. *Exchange Risk and Corporate International Finance*. London: Macmillan.

Bilson, John F. O. 1979. Recent developments in monetary models of exchange rate determination. *I.M.F. Staff Papers* 26, no. 2 (June): 201–23.

Bilson, John F. O. 1980. The speculative efficiency hypothesis. Working Paper no. 474. National Bureau of Economic Research, April.

Cooper, Richard N. 1976. "Comment." *Brookings Papers on Economic Activity*, no. 3, p. 577.

Cornell, Bradford. 1977. Spot rates, forward rates and exchange market efficiency. *Journal of Financial Economics* 5 (August): 55–65.

Fama, Eugene F. 1975. Short term interest rates as predictors of inflation. *American Economic Review* 65, no. 3 (June): 269–82.

Fama, Eugene F., and Farber, Andre. 1979. Money, bonds and foreign exchange. *American Economic Review* 69, no. 4 (September): 639–49.

Frankel, Jeffrey. 1979. The diversifiability of exchange risk. *Journal of International Economics* 9, no. 3 (August): 379–93.

Frankel, Jeffrey. 1980. Tests of rational expectations in the forward exchange market. *Southern Economic Journal* 46, no. 4 (April): 1083–1101.

Frenkel, Jacob A. 1977. The forward exchange rate, expectations, and the demand for money: the German hyperinflation. *American Economic Review* 67, no. 4 (September): 653–70.

Geweke, John, and Feige, Edgar. 1979. Some joint tests of the efficiency of markets for forward foreign exchange. *Review of Economics and Statistics* 61, no. 3 (August): 334–41.

Goodman, Stephen H. 1979. Foreign exchange rate forecasting techniques: implications for business and policy. *Journal of Finance* 34, no. 2 (May): 415–27.

Hakkio, Craig S. 1980*a*. Expectations and the forward exchange rate. Working Paper no. 439, National Bureau of Economic Research, January.

Hakkio, Craig S. 1980*b*. The term structure of the forward premium. Working Paper no. 426, National Bureau of Economic Research, January.

Hansen, Lars Peter, and Hodrick, Robert J. 1980. Forward exchange rates as optimal predictors of future spot rates: an econometric analysis. *Journal of Political Economy* 88, no. 5 (October 1980): 829–53.

Levich, Richard M. 1979. Analyzing the accuracy of foreign exchange forecasting services: theory and evidence. Working Paper no. 336, National Bureau of Economic Research, April.

Longworth, David. 1979. Tests of the efficiency of the Canadian-U.S. exchange market, 1971–1978. Working Paper, Bank of Canada, May.

Tryon, Ralph. 1979. Testing for rational expectations in foreign exchange markets. International Finance Discussion Papers, no. 139, Federal Reserve Board of Governors, May.

Zellner, Arnold. 1962. An efficient method for estimating seemingly unrelated regressions and tests for aggregation bias. *Journal of the American Statistical Association* 57 (June): 348–68.