ELSEVIER

Journal of Empirical Finance 8 (2001) 695–704

Journal of EMPIRICAL FINANCE

www.elsevier.com/locate/econbase

The bias of tests for a risk premium in forward exchange rates

George Tauchen *

Department of Economics, Duke University, Social Science Building, Durham, NC 27708, USA

Abstract

The pure expectations theory of unbiased forward exchange rates predicts that the slope coefficient in a regression of the change in the spot rate on the difference between the current forward and spot rates should equal unity. In the recent empirical work by Fama, the estimates of this coefficient turn out to be negative in all regressions for nine major industrialized nations. This paper demonstrates that under the expectations theory, the sampling distribution of the regression estimator of this coefficient is upward-biased relative to unity and strongly skewed to the right. The likelihood of negative values is essentially zero. Thus, the estimator is biased in a direction opposite to what is observed. Since the observed estimates lie far out in the thin left-hand tail of the estimator's sampling distribution, the evidence against the hypothesis of unbiased forward rates is much stronger than previously believed.

JEL classification: 431
Keywords: Expectations theory; Risk premium; Forward exchange rates

1. Introduction

In a recent paper, Fama (1984) reports the evidence of the relationship between spot and forward exchange rates that is generally discrediting towards the expectations theory of unbiased forward rates. In particular, Fama finds that the difference between the current forward and spot rates is an upward-biased estimate of the

* Tel.: +1-919-660-1812; fax: +1-919-684-8974.
E-mail address: get@econ.duke.edu (G. Tauchen).

PII: S0927-5398(01)00042-1

subsequent change in the spot rate. Fama's evidence suggests that the bias in the forward rate is such that when the forward rate exceeds the spot rate, then the spot rate tends to decline on the average, whereas the expectations theory predicts that the spot rate will increase on the average in this situation. The empirical finding holds true for the exchange rates (relative to the US dollar) of a large group of major industrialized nations. Fama interprets the evidence as being generally favorable towards the presence of a time-varying risk premium in forward rates.

The purpose of this paper is to examine the small sample properties of the estimator used by Fama and other workers in related papers (Bilson, 1981; Frankel, 1982; Levich, 1982). In Section 3 below, it is shown that the sampling distribution of the key regression coefficient used in these studies is closely related to the sampling distribution of the estimate of the autoregressive coefficient for an AR(1) model with the root near unity. As is well known in the time series literature, the latter sampling distribution displays several unusual features that are not accounted for in the conventional asymptotic theory and have important implications for statistical inference in finite samples. The sampling distribution of the regression coefficient also can be expected to display these unusual features, and this is verified in the Monte Carlo sampling experiments. Interestingly though, the small sample bias of the regression coefficient is in a direction opposite to what is actually observed. That is, the results below suggest that if the expectations hypothesis was true, then one would expect to find the reverse of what Fama found, i.e., one would expect to find that the forward–spot differential appears to underestimate, not overestimate, the subsequent change in the spot rate. The evidence presented here on the small sample properties of the estimator thus suggests that the case against the expectations hypothesis is much stronger than Fama's results alone indicate.

The fact that autoregressive roots at or near unity may be important for empirical modeling exchange rates has also been noted by Meese and Singleton (1982). These authors find that for three exchange rates (Switzerland, Canada and Germany), there is not enough evidence to reject the null hypothesis of one unit root in AR(2) models for the logs of the spot and forward rates. This paper extends and elaborates upon their work by presenting explicit evidence on the small sample properties of an important test for no-risk premium when the autoregression for the spot rate has a root near unity.

2. Evidence on the expectations theory

Let s_t and f_t denote the logarithms of the spot and one-period forward exchange rates, respectively. Fama points out that the forward rate f_t can be thought of as comprised of two terms, one of which represents the forecast of the spot rate embodied in the forward rate and the other represents a possibly time-varying risk premium. This decomposition of the forward rate can be written as

$$f_t = E_t[s_{t+1}] + p_t, \tag{2.1}$$

where $E_t[\ ]$ represents the rational expectations operator conditional on information available at time t and p_t is the risk premium.

Under one version of the expectations theory of unbiased forward prices, the premium p_t equals a constant, $p_t = \bar{p}$, independent of t. That is, the forward rate equals the expected spot rate plus a time-invariant risk premium. Under an even stronger version of the expectations theory, the premium is zero for all t and the forward rate equals the expected spot rate.

To test for the presence of a time-varying risk premium, Fama estimates the following commonly used regression relating the change in the spot rate to the forward–spot differential

$$s_{t+1} - s_t = \alpha + \beta(f_t - s_t) + u_{t+1}. \tag{2.2}$$

Here, α and β are the parameters to be estimated and u_{t+1} is the error term. Either form of the expectation hypothesis implies $\beta = 1$. Thus, estimates of β that are far from unity in either direction can be interpreted as evidence against the expectations theory and in favor of the time-varying risk premium. Fama also reports estimates of the equation

$$f_t - s_{t+1} = \alpha' + \beta'(f_t - s_t) + u'_{t+1}. \tag{2.3}$$

The logic behind estimating both equations is that Eq. (2.2) tells us about the ability of the forward–spot differential $f_t - s_t$ to predict the direction of change in the spot rate, while Eq. (2.3) tells us about the systematic time variation in the risk premium. However, as Fama notes, the parameter estimates for Eqs. (2.2) and (2.3) are directly linked by the identities $\hat{\alpha}' = -\hat{\alpha}$ and $\hat{\beta}' = 1 - \hat{\beta}$, as can be seen by adding the equations. There is therefore no more information in Eq. (2.3) than there is in Eq. (2.2). In what follows, we shall only concentrate on Eq. (2.2), remembering that all the conclusions apply with equal force to Eq. (2.3).

Using the data on nine exchange rates sampled at four-week intervals for the period of August 1973–December 1982, Fama finds that the estimates of β in Eq. (2.2) are negative in all cases (Table 2, p. 326). The most negative estimate of β is -1.58 (Belgium) while the least negative is -0.29 (Japan). Generally speaking, the t-statistics for the test of $\beta = 1$ turn out to be around 2 or 3 in magnitude. This evidence tends to discredit the expectations hypothesis, as one would expect the point estimates to cluster around unity and the t-statistics to be smaller in magnitude. The negative point estimates imply that when the forward rate is above the spot rate, then the spot run tends to decline, whereas the expectations hypothesis predicts that the spot rate will increase by an amount equal to the forward–spot differential.

Shown below are the regression results for the US dollar/UK pound exchange rate which confirm Fama's findings for a subset of his sample period:

$$s_{t+1} - s_t = -\underset{(.0044)}{.0017} - \underset{(1.22)}{3.69}\,(f_t - s_t), \tag{2.4}$$

$R^2 = 0.17$, S.E.E. $= 0.029$
$N = 45$ observations
Period of fit: July 1978–December 1981
(Friday closings sampled at four-week intervals)

The exchange rates are in unit pounds per dollar and f_t is the one-month forward rate. The data source is the annual issues of the IMM Yearbook and a few issues of the *Wall Street Journal*. As in Fama's results, the coefficient of the forward–spot differential is negative and several estimated standard deviations are below the maintained value of unity under the expectations hypothesis.

3. Small sample properties of $\hat{\beta}$

The usual statistical justification for inferences based on the regression in Eq. (2.2) is the conventional asymptotic theory for the time series. There are, however, reasons to suspect the accuracy of the asymptotic approximations in this case and to believe that there are important small sample biases inherent in the regression. The grounds for these suspicions concern the special characteristics of the error term and the statistical properties of the explanatory variable.

It is frequently pointed out that the expectations theory ensures that the error term in Eq. (2.2) must be well behaved in the sense that u_{t+1} is serially uncorrelated and uncorrelated with the explanatory variable. This conclusion follows easily from $E_t[u_{t+1}] = 0$ under the expectations theory and the fact that s_{t-j} and f_{t-j} for $j \geq 0$ are in the information set for forming $E_t[\;]$. In general, however, the expectations theory does not imply that u_{t+1} is uncorrelated with future values of the spot or forward rates. Indeed, one can expect in general that u_{t+1} and $(f_{t+1} - s_{t+1})$ are correlated since in period $t+1$, traders will use the information in u_{t+1} when it becomes available to update the forward–spot differential $f_{t+1} - s_{t+1}$. Thus, there is feedback from the error term into future values of the explanatory variable, and so the explanatory variable is only predetermined and not strictly exogenous. Furthermore, the explanatory variable is strongly serially correlated. Fama reports that the first order correlation coefficient for the $f_t - s_t$ series range between 0.65 and 0.87 for the nine countries in his data set. For the US/UK data described above, the coefficient is 0.77. The combined effects of lack of strict exogeneity and autocorrelation in the right-hand variable cause the estimates of β to have finite sample properties different from what the asymptotic theory suggests.

To begin the investigation of the small sample properties of $\hat{\beta}$, suppose the true process generating the spot exchange rate s_t is

$$s_{t+1} = \lambda s_t + \varepsilon_{t+1}, \tag{3.1}$$

where $E_t[\varepsilon_{t+1}] = 0$ and $|\lambda| < 1$. Typically in empirical work, estimates of λ in Eq. (3.1) turn out to be close to unity and sometimes statistically indiscernible from unity. In what follows, we shall think of λ as being near but not equal to unity, i.e., $\lambda = 1 - \eta$ where η is a small positive number, and thus s_t is stationary in the levels. With λ near unity, then s_t, though stationary in levels, is "close" to a random walk, and this is what leads to the conclusion that $\hat{\beta}$ has nonstandard sampling distribution under the expectations hypothesis. The same conclusion will hold when s_t is nonstationary in levels but stationary in differences although the modeling is more elaborate in this case without adding additional insight.

Consider the least squares estimator of β in model Eq. (2.2):

$$\hat{\beta} = \frac{\sum_{t=1}^{T} (f_t - s_t)(s_{t+1} - s_t)}{\sum_{t=1}^{T} (f_t - s_t)^2}, \tag{3.2}$$

where T denotes sample size. For notational simplicity, we only consider the origin-forced estimator of the slope coefficient. The conclusions are similar for the case in which an intercept is included in the regression. (Moreover, the estimated intercept is usually very small and insignificant anyway, so there is very little difference between the origin-forced and the other estimator.) The expectations hypothesis implies that the forward rate is $f_t = \lambda s_t$. Substituting this and Eq. (3.1) into the formula for $\hat{\beta}$ yields,

$$\hat{\beta} = \frac{(\lambda - 1)\sum_{t=1}^{T} s_t[(\lambda - 1)s_t + \varepsilon_{t+1}]}{(1-\lambda)^2 \sum_{t=1}^{T} s_t^2},$$

which can be written as

$$\hat{\beta} = 1 - \frac{1}{(1-\lambda)} \left(\frac{\sum_{t=1}^{T} s_t \varepsilon_{t+1}}{\sum_{t=1}^{T} s_t^2} \right).$$

The coefficient of $-1/(1-\lambda)$ in this expression is exactly equal to the sampling error $\hat{\lambda} - \lambda$ for an OLS estimation of λ in autoregression Eq. (3.1). That is, $\hat{\lambda}$

$= \left(\sum_{t=1}^{T} s_t s_{t+1}\right) / \left(\sum_{t=1}^{T} s_t^2\right)$ and $\hat{\lambda} - \lambda = \left(\sum_{t=1}^{T} s_t \varepsilon_{t+1}\right) / \left(\sum_{t=1}^{T} s_t^2\right)$. Thus, the last expression for $\hat{\beta}$ simplifies to

$$\hat{\beta} = 1 - (1-\lambda)^{-1}(\hat{\lambda} - \lambda).$$

Under the expectations hypothesis, $\beta = 1$, and so

$$\hat{\beta} - \beta = -(1-\lambda)^{-1}(\hat{\lambda} - \lambda). \tag{3.3}$$

According to Eq. (3.3), the sampling error in $\hat{\beta}$ about its maintained value of unity is equal to $-(1-\lambda)^{-1}$ times the sampling error of $\hat{\lambda}$ about the true λ. Therefore, since λ is assumed to be positive, the sampling distribution of $\hat{\beta}$ will be the same as that of λ, except for a sign reversal and a change in scale.

The distribution of $\hat{\lambda} - \lambda$ for an autoregression has been studied extensively in the time series literature. For the extreme case $\lambda = 1$, the conventional asymptotic distribution theory is invalid and the asymptotic distribution of $\hat{\lambda} - \lambda$ is not normal but rather is highly skewed to the left (Fuller, 1976, Chaps. 8 and 9). Sampling experiments show that in finite samples, $\hat{\lambda}$ underestimates the true value of unity and that large negative values of $\hat{\lambda} - 1$ are much more likely than are positive values. When $|\lambda| < 1$, on the other hand, the conventional asymptotic theory applies and in large samples, $\hat{\lambda}$ will be approximately normally distributed. However, if λ is near unity, then the conventional asymptotic theory, though valid, provides a good approximation to the actual distribution of $\hat{\lambda} - \lambda$ only for very large sample sizes well in excess of those available in ordinary practice. That is, if λ is near unity and the sample size is moderate, then the sampling distribution of $\hat{\lambda} - \lambda$ can be expected to be more like that predicted by the case $\lambda = 1$ than like that predicted by the conventional asymptotic theory. In particular, $\hat{\lambda}$ will be downward-biased somewhat and its distribution will be skewed to the left (Phillips, 1977).

By Eq. (3.3), we can expect the sampling distribution of $\hat{\beta}$ to also have these unusual characteristics when λ is near unity, except with a sign reversal. Specifically, we can expect $\hat{\beta}$ to overestimate somewhat the true value $\beta = 1$ under the expectations hypothesis. Furthermore, the sampling distribution of $\hat{\beta}$ will be skewed to the right relative to unity, with large values above unity being more likely than small positive or negative values.

To verify these predictions about the distribution of $\hat{\beta}$, Monte Carlo sampling experiments were run in which the spot exchange rate was assumed to be generated by the following estimate of Eq. (3.1) (with an intercept):

$$s_{t+1} = \underset{(.041)}{.049} + \underset{(.054)}{.935}\, s_t + e_{t+1}, \tag{3.4}$$

$$N_{\mathrm{obs}} = 45, \quad R^2 = 0.88$$

Period = July 1978–December 1981, at four-week intervals

The regression was run using the US/UK exchange rate data described earlier. This estimated equation satisfied the usual diagnostic tests, and in particular, additional lags of s_t beyond lag 1 contributed essentially no additional explanatory power. It is worth noting that by using Table 8.5.2 of Fuller (1976, p. 373), the null hypothesis $\lambda = 1$ cannot be rejected, though given the imprecision of the estimate, there is clearly a range of λs that would also be consistent with these empirical results. Generally speaking, the results of the sampling experiments were insensitive to the value of λ as long as it was near unity.

To generate pseudo-data on the exchange rate s_t for the sampling experiments, a method is needed for generating innovations e_t in autoregression Eq. (3.4). One method for generating the e_t's would be to assume normality and draw e_t randomly from $N(0,\hat{\sigma}_e^2)$, where $\hat{\sigma}_e^2$ is the estimated residual variance. However, because the normality assumption is questionable for financial market data (Clark, 1973; Tauchen and Pitts, 1983), an alternative approach based on the bootstrap (Freedman, 1981) was employed. Specifically, the innovations e_t were drawn randomly from the sample distribution of the residuals $\hat{e}_t$ in the estimated autoregression Eq. (3.4). This was done by generating for each t a random digit $\tau(t)$ between 1 and 45 (the sample size in Eq. (3.4)) and then setting $e_t = \hat{e}_{\tau(t)}$. Note that any number of e_t's (not just 45) can be generated in this manner. The spot rates s_t were generated recursively through the autoregression Eq. (3.4) and the forward rates were set to the rational forecasts, $f_t = 0.049 + 0.935 s_t$, of s_{t+1}. The autoregression was initialized by drawing the first value from the implied stationary distribution of s_t.

The method just described was used to generate $T+1$ pseudo-observations on s_t and f_t and the coefficient estimates $\hat{\alpha}$ and $\hat{\beta}$ were obtained for the regression

$$s_{t+1} - s_t = \alpha + \beta(f_t - s_t) + u_{t+1} \qquad t = 1,2,\dots,T \tag{3.5}$$

using the pseudo-data. Ten thousand replications of the estimation were calculated for each of two values of the econometric sample size T in Eq. (3.5). One of the values, $T = 45$, is the sample size for the regression Eq. (2.4) reported above while the other, $T = 122$, is the sample size in Fama's empirical work.

Columns (1) and (2) of Table 1 contain summary statistics on the sampling distribution of $\hat{\beta}$ for each of these two sample sizes. Both of the sampling distributions have the general characteristics predicted by the analysis above. As expected, the estimate $\hat{\beta}$ is upward-biased somewhat relative to unity, median = 1.11 for $T = 45$. The bias is not much smaller in column (2), median = 1.08 for $T = 122$, despite of more than doubling the sampling size. As also expected, the distributions are skewed to the right. In column (1), for example, the 1st and 99th percentiles were not at all symmetrically placed about unity, but rather the first percentile (0.39) is much closer to unity than is the 99th (2.15). In column (2), the distribution of $\hat{\beta}$ is more concentrated about unity which is to be expected given the larger sample size, but it is still clearly right-skewed and the percentiles are not symmetrically placed about unity.

It is interesting to contrast the empirical estimates of β to what is suggested from the sampling experiments. In Eq. (2.4), the coefficient estimate of β is

Table 1
Sampling distributions of $\hat{\beta}$

Percentile (%)	(1) $N_{obs} = 45$	(2) $N_{obs} = 122$
100 (maximum)	3.18	3.01
99	2.15	1.81
95	1.79	1.54
90	1.63	1.43
75	1.37	1.25
50 (median)	1.11	1.08
25	0.89	0.92
10	0.70	0.80
5	0.59	0.73
1	0.39	0.61
0 (minimum)	0.03	0.42

10,000 replications were used for each distribution.

-3.69. However, in column (1) of Table 1, we see that if the expectations hypothesis was true and if Eq. (3.1) generate the data, then 75% of the time the observed $\hat{\beta}$ would be above 0.89, and with 99% of the time it would be above 0.39. The smallest $\hat{\beta}$ obtained in the 10,000 replications is 0.03. The observed value of -3.69 is so far out in the tail that the prob–val for the expectations hypothesis $\beta = 1$ is essentially zero. Similarly, in Fama's results, the estimates of β range between -1.59 and about -0.29, whereas column (2) of Table 1 indicates that negative values for $\hat{\beta}$ are exceedingly unlikely under the expectations hypothesis with the 122 observations that Fama had. What is particularly discrediting to the expectations hypothesis is that the observed $\hat{\beta}$ lies way out in the "thin" left-hand tail of the sampling distributions.

These calculations, though informative, are still an approximation to the sampling distribution of the estimator used in the actual practice. The reason is that the univariate autoregression Eq. (3.1) is only an approximation to the actual process generating the spot rate. Even though an AR(1) model like Eq. (3.1) can provide a good fit to the univariate distribution of the spot rate, one can still expect in general that other variables will help predict the spot rate. Put another way, the error ε_{t+1} in Eq. (3.1) satisfies $E[\varepsilon_{t+1} \mid s_{t-j},\ j \geqslant 0] = 0$, while in general, $E[\varepsilon_{t+1} \mid s_{t-j},\ z_{t-j},\ j \geqslant 0] \neq 0$ where, z_t represents all other variables available to traders. The expectations theory predicts that the forward rate will incorporate these other variables, $f_t = \mathrm{E}[s_{t+1} \mid s_{t-j},\ z_{t-j},\ j \geqslant 0]$. Thus, in general, one cannot expect $f_t = \lambda s_t$ and $f_t - s_t = (\lambda - 1)s_t$ to hold exactly, as in the case in Eqs. (3.2) and (3.3) where $E_t[\varepsilon_{t+1}] = 0$ was assumed. There is no way to take this into account unless one has a complete model of the spot rate that is known to be true, and this seems unattainable presently.

Nevertheless, the calculations above are, arguably, a reasonably close approximation to the actual sampling distribution of $\hat{\beta}$. What is important about the

AR(1) model in Eq. (3.1) for s_t with λ near unity is that the explanatory variable $f_t - s_t = (\lambda - 1)s_t$ in Eq. (2.2) is highly serially correlated, and although cov$(f_t - s_t, u_{t+1}) = 0$, there is feedback from the disturbance to the future of the explanatory variable such that cov$(f_{t+1} - s_{t+1}, u_{t+1}) < 0$. The combined effects of these two factors—strong autocorrelation in $f_t - s_t$ and the negative correlation between u_{t+1} and $f_{t+1} - s_{t+1}$—are what produce the upward bias in $\hat{\beta}$ and the skewness to the right in the sampling distribution. (In general, the skewness is in a direction opposite the sign of the correlation.) As noted earlier in actual data, $f_t - s_t$ is strongly autocorrelated and so the first factor is present. Secondly, under the strong expectations hypothesis $s_{t+1} - f_t = u_{t+1}$, and a regression of $f_{t+1} - s_{t+1}$ on this variable gives,

$$f_{t+1} - s_{t+1} = \underset{(.00046)}{-.00041} - \underset{(.014)}{.059}\,(s_{t+1} - f_t),$$

$$R^2 = 0.29, \quad \text{S.E.E.} = 0.0031$$
$$N_{\text{obs}} = 45$$

which was estimated using the UK/US data described above. The negative slope coefficient suggests that cov$(f_{t+1} - s_{t+1}, u_{t+1}) < 0$, and so the second factor is very likely present also. Thus, the true sampling distribution of $\hat{\beta}$ is probably very similar to that shown in Table 1.

4. Conclusion

Because of the autocorrelation in the explanatory variable and a lack of strict exogeneity, there are important small sample biases in the regression

$$s_{t+1} - s_t = \alpha + \beta(f_t - s_t) + u_{t+1}.$$

This paper has shown how to use resampling methods to quantify the biases and take them into account in making inferences about expectational theories of forward rates. It turns out that the biases are such that regression estimates $\hat{\beta}$ should tend to overestimate the maintained value $\beta = 1$ if the expectations hypothesis was true. Large deviations in $\hat{\beta}$ above unity should not be surprising, while moderate deviations below unity are unlikely and negative values for $\hat{\beta}$ exceedingly unlikely. Since in actual empirical work, estimates of $\hat{\beta}$ turn out to be negative, the conclusion should be that the evidence is very damaging towards the theory of unbiased forward exchange rates.

Acknowledgements

I am grateful to Ron Gallant, Kent Kimbrough and Dudley Wallace for helpful comments. First draft, July 1985, revised draft, September 1987.

References

Bilson, J.F.O., 1981. The speculative efficiency hypothesis. Journal of Business 54, 435–451, July.

Clark, P.K., 1973. A subordinate stochastic process model with finite varianice for speculative prices. Econometrica 41, 135–155, January.

Fama, E.F., 1984. Forward and spot exchange rates. Journal of Monetary Economics 14, 319–338, November.

Frankel, J.A., 1982. In search of a risk premium: a six-currency test assuming mean variance optimization. Journal of International Money and Finance 1, 255–274, December.

Freedman, D., 1981. Bootstrapping regression models. Annals of Statistics 9, 1218–1228, December.

Fuller, W., 1976. The Statistical Analysis of Time Series. Wiley, New York.

International Money Market (IMM), Yearbooks: 1978–1982. Chicago Mercantile Exchange, Chicago.

Levich, R.M., 1982. On the efficiency of markets for foreign exchange. In: Dornbusch, R., Frenkel, J.A. (Eds.), International Economic Policy: Theory and Evidence. Johns Hopkins Univ. Press, Baltimore, MD, pp. 246–266.

Meese, R.A., Singleton, K.J., 1982. On unit roots and the empirical modeling of exchange rates. Journal of Finance, 1029–1035, September.

Phillips, P.C.B., 1977. Approximations to some finite sample distributions associated with a first-order stochastic difference equation. Econometrica 45, 463–486, March.

Tauchen, G.E., Pitts, M., 1983. The price variability–volume relationship on speculative markets. Econometrica 51, 485–505, March.