

On biases in the measurement of foreign exchange risk premiums

GEERT BEKAERT*

Graduate School of Business, Stanford University, Stanford, CA, USA

AND

ROBERT J. HODRICK

*Department of Finance, Kellogg Graduate School of Management,
Northwestern University, Evanston, IL,
and the National Bureau of Economic Research, USA*

The hypothesis that the forward rate is an unbiased predictor of the future spot rate has been consistently rejected in recent empirical studies. This paper examines several sources of measurement error and misspecification that might induce biases in such studies. Although previous inferences are shown to be robust to a failure to construct true returns and to omitted variable bias arising from conditional heteroskedasticity in spot rates, we show that the parameters were not stable over the 1975–89 sample period. Estimation that allows for endogenous regime shifts in the parameters demonstrates that deviations from unbiasedness were more severe in the 1980s. (JEL F31).

This paper reexamines the relation of the forward premium in the foreign exchange market to the expected rate of currency depreciation over the life of the forward contract. For at least ten years, empirical studies of this relation have regressed *ex post* rates of depreciation on a constant and the forward premium. Their null hypothesis is that the slope coefficient is one. Researchers have consistently found point estimates of the slope coefficient that are negative and that are often more than two standard errors from zero. Predicted currency depreciation is therefore *very* different from the forward premium whereas the unbiasedness hypothesis implies that they are equal. An important consequence of this finding is that

* Geert Bekaert gratefully acknowledges partial support for this research from the NBER's Committee on Dissertation Support Awards and from a Sloan Dissertation Fellowship.

expected rate-of-return differentials between foreign investments that are covered to eliminate foreign exchange risk and uncovered investments are large and variable.

One interpretation of these empirical results relies on Fama's (1984) decomposition of the forward premium into the expected rate of depreciation and a risk premium. Finding a negative slope coefficient in an unbiasedness test then can be demonstrated to imply highly variable risk premiums. Another interpretation of the results is summarized in Froot and Thaler (1990), who argue that systematic forecast errors are needed to explain the results. A third position is that of Cornell (1989) who argues that measurement errors in the analysis may be so bad as to render interpretation of the empirical work inappropriate.

This debate supplies the motivation for the paper. We retain the assumption of rational expectations, but we reexamine the unbiasedness hypothesis and address several sources of measurement error and misspecification that might bias the coefficient estimates and thus alleviate the burden on a time varying risk premium as the explanation of the previous empirical results.

The first source of potential bias is measurement error. It is often difficult for researchers to obtain high quality data, and many studies have been rather cavalier in their construction of returns. Cornell (1989) criticizes studies in this area for two reasons. First, many studies fail to use data sampling procedures that observe the market rules governing delivery on foreign exchange contracts. Second, such studies often fail to incorporate transactions costs in terms of bid-ask spreads.¹ Cornell (1989, p. 155) concludes, 'Until the impact of both measurement error and specification error has been more accurately assessed, it is premature to conclude that forward rates are biased predictors of future spot rates.'

The first section of the paper reviews the theory that forms the foundation of the econometric tests. In the second section of the paper, we carefully construct an actual foreign exchange return series that incorporates both the transactions costs inherent in bid-ask spreads and the delivery structure of the market. To preview the empirical results, we find essentially no difference in inference across specifications that are correctly constructed and ones that are incorrectly specified.

The second source of potential bias is conditional heteroskedasticity. The theoretical derivation of an unbiasedness test demonstrates that the conditional variance of the rate of depreciation may enter the equation even if agents are risk neutral. Section III of the paper uses Monte Carlo analysis to determine how this might bias the slope coefficient away from one. We find only a slight bias in the slope coefficient from this analysis.

The last source of bias we investigate is poor small sample properties of the regression equation. Since there have been major regime shifts in monetary and fiscal policies during the sample, the results might not be stable over time.² Thus, running a regression across various regimes could result in a bad estimate of the unconditional covariance between the forward premium and realized currency depreciation since OLS ignores information about the switches in regimes. The fourth section of the paper investigates regime shifts by estimating a bivariate Markov switching model of the rate of depreciation and the forward premium. This builds on the work of Engel and Hamilton (1990).

The final section of the paper provides some concluding remarks and notes the nature of the challenge that the existing empirical results imply for theory.

I. A review of theory

In models of rational maximizing behavior an intertemporal Euler equation dictates investment decisions. The loss in marginal utility from sacrificing a dollar at time t to invest in an asset is equated to the expected gain in marginal utility from holding the asset and selling it at time $t + 1$. Define Q_{t+1} to be the intertemporal marginal rate of substitution of a dollar between period t and $t + 1$, and let R_{t+1} be the dollar return at $t + 1$ on a dollar invested at t . Let $E_t(\cdot)$ denote the conditional expectation. Then, the Euler equation is

$$\langle 1 \rangle \quad E_t(Q_{t+1}R_{t+1}) = 1.$$

Equation $\langle 1 \rangle$ is the foundation of many theoretical and empirical investigations of asset pricing. In the most basic representative agent models, *e.g.*, Lucas (1982), the form of the intertemporal marginal rate of substitution is straightforwardly derived to be

$$\langle 2 \rangle \quad Q_{t+1} = \frac{\rho U'(C_{t+1})\pi_{t+1}}{U'(C_t)\pi_t},$$

which is the agent's discount factor times the ratio of the marginal utility of consumption at time $t + 1$ multiplied by the purchasing power of a dollar at time $t + 1$ to the product of these variables at time t . In general, Q_{t+1} is the ratio of the discounted value of an asset market Lagrange multiplier valued at time $t + 1$ to the value of the multiplier at time t .

Now consider the implications of equation $\langle 1 \rangle$ for the determination of expected returns. In most countries there is an asset that has a certain nominal return. It is common financial terminology to refer to such an asset as the risk-free asset even though the real return on the asset is uncertain. Let the continuously compounded dollar interest rate be i_t . Then, $R_{f,t+1} = \exp(i_t)$ is the nominal risk-free dollar return. Since the dollar denominated risk-free return must satisfy equation $\langle 1 \rangle$, $R_{f,t+1} = [E_t(Q_{t+1})]^{-1}$.

Investing dollars internationally requires conversion into foreign currencies. Let S_t be the dollar price of a unit of foreign currency, in which case $s_{t+1} - s_t = \ln(S_{t+1}/S_t)$ is the continuously compounded rate of depreciation of the dollar relative to the foreign currency. Let the continuously compounded foreign currency interest rate be i_t^* . Then, the dollar return from investing in the foreign currency money market and bearing the foreign exchange risk is $\exp(i_t^*)(S_{t+1}/S_t) = \exp(i_t^* + s_{t+1} - s_t)$. This dollar return must also satisfy equation $\langle 1 \rangle$.

Following Hansen and Hodrick (1983), assume that s_{t+1} and $q_{t+1} = \ln(Q_{t+1})$ are jointly conditionally normally distributed. Then, the interest rate and the expected excess rates of return satisfy the following:

$$\langle 3 \rangle \quad i_t = -E_t(q_{t+1}) - 0.5V_t(q_{t+1}),$$

$$\langle 4 \rangle \quad i_t^* + E_t(s_{t+1} - s_t) - i_t = -0.5V_t(s_{t+1}) - C_t(s_{t+1}, q_{t+1}),$$

where $V_t(\cdot)$ denotes the conditional variance and $C_t(\cdot)$ denotes the conditional covariance.

Let F_t be the forward exchange rate at time t for delivery at time $t + 1$. Then, $f_t - s_t = \ln(F_t/S_t)$ is the continuously compounded forward premium on the foreign currency. To prevent covered interest arbitrage, the return from investing

a dollar in the foreign currency money market and selling this foreign currency return in the forward market must equal the risk-free return on the dollar. Equality of these two returns provides an expression for interest rate parity:

$$\langle 5 \rangle \quad i_t = i_t^* + f_t - s_t.$$

Substituting from equation $\langle 5 \rangle$ into equation $\langle 4 \rangle$ and rearranging, we find

$$\langle 6 \rangle \quad E_t(s_{t+1} - s_t) = (f_t - s_t) - 0.5V_t(s_{t+1}) - C_t(s_{t+1}, q_{t+1}).$$

I.A. Linking the theory to the econometric models

One of the most fundamental issues in international finance is the nature of foreign exchange risk. Empirically, the absence of foreign exchange risk is often equated with the proposition that the forward premium is an unbiased predictor of the rate of depreciation over the life of the forward contract, $f_t - s_t = E_t(s_{t+1} - s_t)$.

A typical regression test of the unbiasedness hypothesis is specified as

$$\langle 7 \rangle \quad s_{t+1} - s_t = \alpha + \beta(f_t - s_t) + \varepsilon_{t+1},$$

and the null hypothesis is $\alpha = 0$ and $\beta = 1$. Under the null hypothesis, $\varepsilon_{t+1} = (s_{t+1} - s_t) - E_t(s_{t+1} - s_t)$ is a rational expectations error term that is orthogonal to time t information.

The theoretical analysis above, which resulted in the derivation of equation $\langle 6 \rangle$, indicates several reasons why the unbiasedness hypothesis cannot literally be derived from an economic model in which agents maximize expected utility. Notice that if the conditional variance and conditional covariance are constant, α would not necessarily be zero, but $\beta = 1$ would be true regardless of the nature of risk aversion.³ When constant conditional variances and covariances are not imposed, equation $\langle 6 \rangle$ implies a β in equation $\langle 7 \rangle$ equal to $1 - \beta_v - \beta_c$, where β_v (β_c) denotes the covariance of $V_t(s_{t+1})$ ($C_t(s_{t+1}, q_{t+1})$) with the forward premium divided by the variance of the forward premium. This decomposition is slightly different from the Fama (1984) decomposition of the continuously compounded forward premium into the expected rate of depreciation plus a risk premium since risk aversion only enters through the covariance term in equation $\langle 6 \rangle$. Before examining the possible effects of movements in the conditional variance on the estimated value of β , we address simple measurement error as a source of potential bias.

II. Measurement errors as a source of bias

Consider first the problem of determining the correct day in the future that the one-month forward rate is predicting. To find the delivery date on a forward contract made today, one first finds today's spot value date, which is two business days in the future for trades between US dollars and European currencies or the Japanese yen. Delivery on a 30-day forward contract occurs on the calendar day in the next month that corresponds to the calendar day of the current month on which spot value is realized if this day in the next month is a legitimate business day. If it is a weekend or a holiday, one takes the next available business day without going out of the month. In this latter contingency, one takes the first previous business day. This rule is followed except when the spot value day is

the last business day of the current month in which case the forward value day is the last business day of the next month (the end-end rule). Unless one matches the forward rate with the appropriate spot rate, a true return on the forward contract is not being calculated and measurement error is introduced into the analysis.

The problem of transactions costs induced by bid-ask spreads is also a potential source of bias in the statistical analysis. When one buys a foreign currency with dollars, one pays the bank's asking price of dollars per foreign currency, and when one sells the foreign currency to the bank for dollars, one receives the bank's bid price. Hence, the dollar return on a forward contract to buy a unit of the foreign currency is the bid price in the future spot market minus the ask price in the current forward market.

It would not be at all surprising if measurement error biased the estimate of β in equation <7> toward zero, but it seems unlikely *a priori* that measurement error could explain the findings of previous research. Nevertheless, to address the concerns raised by Cornell (1989) and to determine how much difference correct sampling procedures make, we obtained daily bid and ask exchange rate data from Citicorp Database Services for the period 1975 to 1989. The data are captured from a Reuter's screen and represent quoted market prices at which someone could have conducted a transaction. We ran several filter tests on the data to check for errors, and, unfortunately, we found a few. These were corrected with observations from the *International Monetary Market Yearbook* or from the *Wall St. Journal*.

Table 1 presents results of estimating equation <7> for correctly sampled data in Panel A and for data that are incorrectly sampled in Panel B. The correctly sampled data follow the market delivery conventions discussed above, and use ask prices for time t and bid prices for future spot rates. The data are overlapping weekly observations. We selected Fridays as the day of the week for the forward buy transaction and matched the future spot rate by determining the correct spot transactions date in the next month that produces spot value on the forward value date. The data in Panel B are sampled incorrectly on every Friday, as in the Harris Bank data employed by Fama (1984) and others, and are all ask rates.

Notice that the point estimates for the correctly sampled data are actually slightly more negative than those from the incorrectly sampled data and their standard errors are approximately 10 per cent higher. But, there are no differences in inference across the two sets of estimates. In all cases the slope coefficients are more than two standard errors below zero. The chi-square statistic that tests the joint significance of the deviation of the three slope coefficients from one is very large in both cases.

Cornell (1989) also argues that errors in the timing of the market prices at time t could bias the estimation of β toward negative numbers. While this is true, his derivation of the bias is incorrect. He specifies the unbiasedness hypothesis as

$$\langle 8 \rangle \quad f_t - s_{t+1} = a + b(f_t - s_t) + \varepsilon_{t+1},$$

and the null hypothesis is $a = b = 0$. Notice that since the β in equation <7> is equal to $1 - b$ in equation <8>, positive bias in b would tend to bias β below one. Cornell illustrates the potential for bias by postulating that the forward rate is measured with error. Thus, the measured forward rate, f_t , deviates from the true forward rate, f_t^* , by a random error, $f_t = f_t^* + e_t$. Since there is no reason

TABLE 1. Tests of unbiasedness: weekly data 1975 to 1989.

Coefficients on regressors					
Currency	Const. (s.e.) Conf. ($\alpha = 0$)	$f_t - s_t$ (s.e.) Conf. ($\beta_1 = 1$)	$f_{t-1} - s_{t-1}$ (s.e.) Conf. ($\beta_2 = 1$)	$\chi^3(3)$ Conf.	R^2
Panel A: Correctly sampled					
Deutsche mark	13.578 (5.076) 0.993	-3.015 (1.243) 0.999			0.026
British pound	-7.956 (2.932) 0.993	-2.021 (0.703) 0.999		31.586 0.999	0.033
Japanese yen	12.821 (3.309) 0.999	-2.098 (0.631) 0.999			0.034
Panel B: Incorrectly sampled on Friday					
Deutsche mark	13.198 (4.591) 0.996	-2.894 (1.142) 0.999			0.028
British pound	-6.484 (2.619) 0.987	-1.878 (0.632) 0.999		33.890 0.999	0.033
Japanese yen	11.567 (2.990) 0.999	-1.884 (0.573) 0.999			0.033
Panel C: Lagged as recommended by Cornell (1989)					
Deutsche mark	11.516 (6.052) 0.999		-2.486 (1.449) 0.984		0.018
British pound	-7.818 (3.560) 0.972		-1.951 (0.828) 0.999	23.633 0.999	0.030
Japanese yen	12.819 (4.000) 0.999		-2.099 (0.718) 0.999		0.034

Notes: Estimation is by Hansen's (1982) GMM with regressors as instruments (a just-identified system). The parameter estimates are consequently OLS for each equation, but the covariance matrix is heteroskedasticity-consistent and allows for the serial correlation of the error terms induced by the overlap in the weekly observations. Exchange rates are dollars per foreign currency and all observations are annualized. The $\chi^3(3)$ statistic is the test of the joint hypothesis that $\beta_1 = 1$ for each equation in Panel A and B or each β_2 in Panel C.

to suppose otherwise, he imposes $\text{cov}(e_t, f_t^* - s_t) = 0$ and $\text{cov}(e_t, f_t^* - s_{t+1}) = 0$. Given these assumptions it is straightforward to show that if the unbiasedness hypothesis holds for the true prices, the estimate of b is given by

$$\langle 9 \rangle \quad b = \frac{\text{var}(e_t)}{\text{var}(e_t) + \text{var}(f_t^* - s_t)}.$$

In comparison, Cornell's equation $\langle 5 \rangle$, which is his analogous derivation of the biased coefficient is

$$\langle 10 \rangle \quad b = \frac{\text{var}(e_t)}{\text{var}(e_t) + \text{cov}(f_t^* - s_{t+1}, f_t^* - s_t)}.$$

Cornell notes that the variance of the measurement error, $\text{var}(e_t)$, will typically be small, but he argues that the covariance in the denominator of equation $\langle 10 \rangle$ will also be small, which makes the potential bias large. But, the correct comparison is in equation $\langle 9 \rangle$. Because the variance of the true forward premium is likely to be much larger than the variance of the measurement error, this source of bias does not seem to be quantitatively important.

Another source of bias in equation $\langle 7 \rangle$ could be from the time t spot exchange rate, which is on both sides of the equation, if, for example, the spot rate at time t is actually measured after the forward rate is set. To determine whether this is a problem, Panel C reports the results of lagging the forward premium by one week, again suggested by Cornell (1989). Inference is again unchanged. The point estimates are similar, and the standard errors are approximately 15 per cent higher. The chi-square statistic that tests the joint significance of the deviation of the three slope coefficients from one is smaller, but it is still well above the critical value of the 0.001 marginal level of significance.

By employing data on the bid-ask spread, we can investigate whether the differences in the point estimates between Panels A and B can be primarily ascribed to misalignment of the data rather than to transactions costs. Rather than having an incorrectly measured forward rate, as above, the source of measurement errors in most studies would seem to arise from use of the wrong future spot rate. The value can be wrong because it is taken on the wrong day, and it can be wrong because it is the ask price instead of the bid price or it is an average of the bid and the ask. Let $s_{t+1}^* - s_t$ be the 'true' regressand, and let $s_{t+1} - s_t$ be the incorrectly sampled regressand. The bias in the 'true' regression coefficient in equation $\langle 7 \rangle$, β^* , that is induced by these measurement errors, is given by

$$\langle 11 \rangle \quad \beta - \beta^* = \frac{\text{cov}(s_{t+1} - s_t^*, f_t - s_t)}{\text{var}(f_t - s_t)}.$$

Without misalignment of the data, $s_{t+1} - s_t^*$ represents the (logarithmic) bid-ask spread in the foreign currency spot market. Since the rates of depreciation and the forward premiums are annualized, we multiply the logarithmic bid-ask spread by 1200 to annualize the percentage spread. The resulting means and standard deviations for the three currencies are 0.567 and 0.348 for the deutsche mark, 0.682 and 0.376 for the pound, and 0.795 and 0.449 for the yen. These transactions costs consequently represent between 50 and 80 basis points which

are non-trivial amounts. Direct calculation of the bias in equation <11> is positive. For the deutsche mark, it is 0.022; for the pound, it is 0.011; and for the yen, it is 0.040. These values are 18 per cent, 8 per cent, and 18.5 per cent of the respective biases of the estimated β s reported in Panels A and B of Table 1. Consequently, we conclude that the less negative estimates for β and the slight increase in statistical significance of the joint test in Panel B are due primarily to sampling the data incorrectly.

III. Omitted variable bias

We now address the importance of conditional heteroskedasticity in currency depreciation in the determination of the value of β in estimates of equation <7>. We conduct Monte Carlo experiments to determine how much the absence of the conditional variance, which is present in equation <6> but not in equation <7>, biases the estimate of β . To the extent that the forward premium is positively correlated with the conditional variance of the rate of depreciation in equation <6>, the true risk premium, which is related only to the covariance term, is relieved of the 'burden' of accounting for negative coefficients since the β in equation <7> would be biased downward even in the absence of risk aversion.

This investigation seems promising for two reasons. First, there is considerable evidence of conditional heteroskedasticity in foreign exchange markets (see, for example, Baillie and Bollerslev, 1989 and 1990). Second, there is some evidence that the squared forward premium is positively correlated with the conditional variance of the rate of depreciation (see Hodrick, 1989).

It is well known that evidence for conditional heteroscedasticity in the foreign exchange market is much stronger with daily or weekly data. Hence, we sample rates of depreciation weekly, and we assume a flat term structure of forward premiums so that one-fourth times the one-month forward premium corresponds to the one-week forward premium. This assumption is motivated by the complexity of the estimation. Table 2 reports our investigation of the dollar–yen rate, and Table 3 contains the investigation of the dollar–DM rate.

Each Panel A of Tables 2 and 3 reports the estimation of a bivariate GARCH-in-mean model for the weekly data. The conditional means and conditional variances of the rate of depreciation and the forward premium are modeled jointly and are allowed to vary over time. The model imposes a constant correlation structure on the innovations in the processes as in Bollerslev (1990). As in equation <6>, we allow the forward premium and the conditional variance of the rate of depreciation to enter the expected rate of depreciation, and we constrain the coefficients to be 1 and -0.5 .⁴ We enter the absolute value of the forward premium into the conditional variance of the rate of depreciation to induce correlation between the two series.

The constrained models serve as the data generating processes for the Monte Carlo experiments. The unconstrained models are presented in each Panel B of Tables 2 and 3. Residual diagnostics for the two models are reported in Panels C. The autocorrelations of the squared residuals divided by their respective conditional variances should be zero. Similarly, the autocorrelations of the product of the residuals divided by the product of the conditional standard deviations should be zero. The Ljung–Box (1978) tests of these restrictions generally do not reject the null hypotheses for the conditional variance models

TABLE 2. Monte Carlo experiments for omitted variable bias: the dollar–yen.

Panel A: Constrained estimates (the data generating processes)						
$s_{t+1} - s_t = 0.064 + 1.000(f_t - s_t) - 0.500h_{1t+1} + e_{1t+1}$ <p style="text-align: center;">(0.049)</p>						
$f_{t+1} - s_{t+1} = 0.002 - 0.0001(s_t - s_{t-1}) + 0.959(f_t - s_t) + e_{2t+1}$ <p style="text-align: center;">(0.001) (0.0006) (0.017)</p>						
$h_{1t+1} = 0.007 + 0.925h_{1t} + 0.064e_{1t}^2 + 0.139 f_t - s_t $ <p style="text-align: center;">(0.008) (0.100) (0.023) (0.569)</p>						
$h_{2t+1} = 0.64^* + 0.916h_{2t} + 0.083e_{2t}^2$ <p style="text-align: center;">(1.6^*) (0.119) (0.074)</p>						
$h_{12t+1} = -0.041(h_{1t+1}h_{2t+1})^{0.5}$ <p style="text-align: center;">(0.031)</p>						
Log-likelihood function = 4208.5						
Panel B: Unconstrained estimates						
$s_{t+1} - s_t = 0.059 - 1.173(f_t - s_t) + 6.544h_{1t+1} + e_{1t+1}$ <p style="text-align: center;">(0.070) (0.618) (5.204)</p>						
$f_{t+1} - s_{t+1} = 0.002 - 0.0001(s_t - s_{t-1}) + 0.960(f_t - s_t) + e_{2t+1}$ <p style="text-align: center;">(0.001) (0.0006) (0.017)</p>						
$h_{1t+1} = 0.009 + 0.924h_{1t} + 0.066e_{1t}^2 + 0.114 f_t - s_t $ <p style="text-align: center;">(0.011) (0.079) (0.024) (0.401)</p>						
$h_{2t+1} = 0.65^* + 0.916h_{2t} + 0.083e_{2t}^2$ <p style="text-align: center;">(1.60^*) (0.119) (0.073)</p>						
$h_{12t+1} = -0.043(h_{1t+1}h_{2t+1})^{0.5}$ <p style="text-align: center;">(0.030)</p>						
Log-likelihood function = 4222.1						
Panel C: Residual diagnostics						
	Unconstrained estimation			Constrained estimation		
	e_{1t}^2	e_{2t}^2	$e_{1t}e_{2t}$	e_{1t}^2	e_{2t}^2	$e_{1t}e_{2t}$
Q3	0.347	7.220	3.919	0.405	7.192	4.304
p-value	0.951	0.065	0.270	0.939	0.066	0.230
Q6	0.912	7.519	28.454	0.890	7.491	29.022
p-value	0.989	0.275	0.0001	0.989	0.278	0.0001
Q12	2.209	8.095	38.998	2.052	8.067	42.502
p-value	0.999	0.778	0.0001	0.999	0.780	0.00001
Pagan–Sabau	16.615	1.802	0.039	14.205	1.799	0.127
	0.0001	0.180	0.844	0.0002	0.180	0.722

Likelihood ratio test: $\chi^2(2) = 27.2$, p-value = <0.001

TABLE 2—Continued.

Panel D: The empirical distribution of β									
Mean: 0.925									
Standard deviation: 0.365									
Quantiles:	1%	5%	10%	25%	50%	75%	90%	95%	99%
	0.015	0.317	0.472	0.696	0.928	1.152	1.374	1.493	1.799

Notes: Observations are weekly for January 1975 to December 1989. The model is the constant conditional correlation model of Bollerslev (1990). The conditional variance of e_{1t} (e_{2t}) is h_{1t} (h_{2t}), and $h_{1,2t}$ is the conditional covariance between e_{1t} and e_{2t} . The parameters in the conditional variance equations were constrained to be positive and the GARCH parameters to be in (0,1). A * indicates 10^{-5} and ** indicates 10^{-10} . Estimation was by maximum likelihood assuming a normal distribution for $e_t = [e_{1t}, e_{2t}]'$. Under very weak conditions, including misspecification of the distribution function (see White, 1982), the vector of parameters Θ is asymptotically normally distributed with covariance matrix $A^{-1}BA^{-1}$, where A is the Hessian form and B the outer product form of the information matrix. These robust standard errors were generally larger than ones based on the usual estimates of the inverse of the information matrix. The Q_j statistics are constructed as in Ljung and Box (1978) and are distributed $\chi^2(j)$, but they are applied to the squared residuals. The Pagan–Sabau (1987) test examines the restriction that the slope coefficient should be one in an OLS regression of a squared residual (or product of the residuals) on a constant and the respective conditional variance (covariance). For the Monte Carlo experiment, we drew standard normal innovations with the RNDNS command in Gauss and constructed e_t using the Cholesky decomposition of H_t , the conditional covariance matrix for e_t . The evolution of H_t and the observations on rates of depreciation and the forward premium were generated recursively from the model.

although the test statistics for the autocorrelations of the cross-residuals are large and have relatively low marginal levels of significance for the yen. We also report the Pagan–Sabau (1987) test. This test examines the restriction that the slope coefficient should be one in an OLS regression of a squared residual (or product of the residuals) on a constant and the respective conditional variance (covariance). The test statistic is constructed in the usual way as the squared ratio of the coefficient estimate minus one relative to the heteroskedasticity-consistent standard error of the estimated parameter. The resulting statistic has a chi-square distribution with one degree of freedom. Generally, the models are reasonably successful in capturing the conditional heteroskedasticity present in the two series although there is some evidence against the models. Likelihood ratio tests of the constrained models versus the unconstrained models, which are chi-square statistics with two degrees of freedom, have a value of 27.2 for the yen and 29.5 for the DM, which coincide with a marginal level of significance of less than 0.001.

In each case we used the constrained GARCH-in-mean model in a Monte Carlo experiment to generate weekly observations on monthly rates of depreciation and monthly forward premiums. We generated 2000 sets of 720 overlapping observations. The simulated data are then used in a regression like equation <7>. The point of the Monte Carlo experiments is to see how much the correlation between the forward premium and the conditional variance biases the regression coefficient in equation <7> when the conditional variance is not included in the regression. The empirical distributions of β arising from these artificial data are given in Panels D of Tables 2 and 3. The means of these distributions are 0.925

TABLE 3. Monte Carlo experiments for omitted variable bias: the dollar-DM.

Panel A: Constrained estimates (the data generating processes)						
$s_{t+1} - s_t = 0.002 + 1.000(f_t - s_t) - 0.500h_{1t+1} + e_{1t+1}$ <p style="text-align: center;">(0.046)</p>						
$f_{t+1} - s_{t+1} = 0.001 - 0.0001(s_t - s_{t-1}) + 0.980(f_t - s_t) + e_{2t+1}$ <p style="text-align: center;">(0.001) (0.0002) (0.009)</p>						
$h_{1t+1} = 0.04^{**} + 0.882h_{1t} + 0.084e_{1t}^2 + 0.948 f_t - s_t $ <p style="text-align: center;">(8.0^{**}) (0.069) (0.046) (0.659)</p>						
$h_{2t+1} = 0.12^* + 0.849h_{2t} + 0.148e_{2t}^2$ <p style="text-align: center;">(0.06^*) (0.067) (0.031)</p>						
$h_{12t+1} = -0.022(h_{1t+1}h_{2t+1})^{0.5}$ <p style="text-align: center;">(0.042)</p>						
Log-likelihood function = 5350.5						
Panel B: Unconstrained estimates						
$s_{t+1} - s_t = 0.193 - 3.435(f_t - s_t) + 5.668h_{1t+1} + e_{1t+1}$ <p style="text-align: center;">(0.081) (1.432) (5.236)</p>						
$f_{t+1} - s_{t+1} = 0.001 - 0.0001(s_t - s_{t-1}) + 0.980(f_t - s_t) + e_{2t+1}$ <p style="text-align: center;">(0.001) (0.0002) (0.009)</p>						
$h_{1t+1} = 3.2^{**} + 0.882h_{1t} + 0.081e_{1t}^2 + 0.981 f_t - s_t $ <p style="text-align: center;">(0.9^{**}) (0.065) (0.043) (0.627)</p>						
$h_{2t+1} = 0.12^* + 0.848h_{2t} + 0.148e_{2t}^2$ <p style="text-align: center;">(0.06^*) (0.067) (0.031)</p>						
$h_{12t+1} = -0.022(h_{1t+1}h_{2t+1})^{0.5}$ <p style="text-align: center;">(0.042)</p>						
Log-likelihood function = 5365.1						
Panel C: Residual diagnostics						
	Unconstrained estimation			Constrained estimation		
	e_{1t}^2	e_{2t}^2	$e_{1t}e_{2t}$	e_{1t}^2	e_{2t}^2	$e_{1t}e_{2t}$
Q3	3.003	4.077	3.481	3.406	4.061	3.328
p-value	0.391	0.253	0.323	0.333	0.255	0.344
Q6	6.148	8.622	11.525	7.656	8.620	12.580
p-value	0.407	0.196	0.073	0.264	0.196	0.050
Q12	11.027	15.448	16.297	13.543	15.420	17.739
p-value	0.527	0.218	0.178	0.331	0.219	0.124
Pagan-Sabau	2.567	1.025	2.861	3.193	1.023	1.994
	0.109	0.311	0.091	0.074	0.312	0.158

Likelihood ratio test: $\chi^2(2) = 29.5$, p-value = <0.001

TABLE 3—Continued.

Panel D: The empirical distribution of β									
Mean: 0.910									
Standard deviation: 1.016									
Quantiles:	1%	5%	10%	25%	50%	75%	90%	95%	99%
	-1.677	-0.773	-0.351	0.259	0.936	1.589	2.160	2.537	3.338

Notes: See the note for Table 2.

for the yen and 0.910 for the DM. These are not very far below 1, and only 0.9 per cent of the observations are below zero for the yen. For the DM, 17.15 per cent of the observations are below zero, and 3.45 per cent are below -1 . Hence, the sample estimates for β of -2.098 for the dollar–yen exchange rate and of -3.015 for the dollar–DM in Table 1 seem unlikely to be drawn from these distributions.

IV. Parameter instability and regime shifts

Given the evolution of international financial markets during our sample and the discussions in the literature of frequent ‘regime shifts’ in monetary and fiscal policy across the countries of our sample, the stability of the parameters of equation <7> is questionable. This section examines the stability of the coefficients.

First, we examine stability tests for a single predetermined break. We choose January 1980 as our break point because this date coincides with the end of Bilson’s (1981) sample, and his paper is often credited with the first published estimates of a statistically significant, negative β in equation <7>. This point is also close to the October 1979 change in monetary operating procedures of the Federal Reserve System.

The first stability test we examine is the Predictive Test which is described in Ghysels and Hall (1990). To develop the test, let the true parameter vector be θ , and consider the three equations to be a system. Hence, θ contains the α and the β for each currency. Let $\theta(T_i)$ denote the estimator of θ for a sample of size T_i . The test statistic is derived from the asymptotic distribution of the vector of orthogonality conditions for the second sample when they are evaluated at the parameter estimates from the first sample. Let $g(T_2, \theta(T_1))$ represent this vector of orthogonality conditions. Hansen (1982) demonstrates that $\sqrt{T_i}g(T_i, \theta) \sim N(0, S_i)$, where S_i is the spectral density of the orthogonality conditions evaluated at frequency zero. Then, the predictive test statistic is $T_2g(T_2, \theta(T_1))'V^{-1}g(T_2, \theta(T_1))$ where V is a consistent estimator of $S_2 + cD_2(D_1'S_1^{-1}D_1)^{-1}D_2'$, D_i' is the gradient of the orthogonality conditions with respect to the parameter vector, and c is T_2/T_1 . Both S_2 and D_2 are evaluated at $\theta(T_1)$. The test statistic has a chi-square distribution with degrees of freedom equal to the number of orthogonality conditions, which is six in this case. The value of the test statistic of 14.543 is larger than the 0.024 critical value. This indicates that the orthogonality conditions for the second sample do not have zero means when evaluated at the parameter estimates from the first sample.

Table 4 also reports parameter estimates for the three currencies. The results in Panel A are for the beginning of 1975 to the first week of 1980 and the ones in Panel B are from the sixth week of 1980 to the end of 1989. In Panel A the β coefficient for the deutsche mark is surprisingly 1.040, and the β for the pound is 1.623.⁵ Given the respective standard errors of 1.313 and 1.162, the estimates are insignificantly different from one. The slope coefficient for the yen is -1.044 with a standard error of 0.907. Hence, the null hypothesis would be rejected at marginal significance levels greater than 0.024. The joint hypothesis that all three coefficients equal one would be rejected at the 0.05 level. In Panel B, the β for the yen is -3.007 , the β for the pound is -4.113 , and the β for the deutsche mark is -0.941 . The evidence against the joint hypothesis that all three coefficients equal one is very strong since the $\chi^2(3)$ has a value of 31.672, which is well beyond the 0.001 critical value of the distribution. The α coefficients for the deutsche mark and the yen have also become more positive while the α coefficient for the pound changes sign from positive to negative.

Table 4 also reports a GMM analogue of seemingly unrelated estimation in which the slope coefficients are constrained to be the same across currencies since Bilson (1981) also constrained his coefficients. The 12 orthogonality conditions are constructed by making each of the three error terms orthogonal to a constant and the three forward premiums, which are the variables on the right-hand sides of the three equations. As with the unconstrained estimation, the standard errors are robust to conditional heteroskedasticity and allow for the overlap in the data. For the first sub-sample, the estimate of the constrained slope coefficient is 0.896 with a standard error of 0.551. For the second sub-sample, the estimate of the constrained slope coefficient is -4.601 with a standard error of 0.594. Clearly, the evidence against the null hypothesis is quite weak in the first sub-sample and extraordinarily strong in the second sub-sample.

Since the system is overidentified, the test of the overidentifying restrictions is a chi-square statistic with eight degrees of freedom. The system would be just-identified if each forward premium entered each equation. Hence, the overidentification test examines the zero restrictions on the coefficients of the forward premiums which are excluded from each equation. The value of the test statistic for the first sub-sample is 14.693 (a confidence level of 0.935), and for the second sub-sample, the value of the test statistic is 7.587 (a confidence level of 0.525). Hence, there is more evidence against the zero restrictions in the first sub-sample than in the second.

Panel C of Table 4 reports a direct test for coefficient stability which is analogous to a Chow test and employs the asymptotic distributions of the coefficients as in Hodrick and Srivastava (1984). Hansen (1982) demonstrates that $\sqrt{T_i}[\theta(T_i) - \theta] \sim N(0, \Omega_i)$. Therefore, under the null hypothesis of no change in parameter values, the difference between the parameter estimates from two non-overlapping samples, $\theta(T_1) - \theta(T_2)$, is also normally distributed with mean zero and variance $\Omega = \Omega_1/T_1 + \Omega_2/T_2$. Consequently, $(\theta(T_1) - \theta(T_2))'\Omega^{-1}(\theta(T_1) - \theta(T_2))$ has a chi-square distribution with degrees of freedom equal to the dimension of θ . The value of the test statistic for the six coefficients is 20.733. This indicates a 0.002 chance that the true coefficients are constant and that sampling error accounts for the differences in the measured values across the two samples.

This evidence on parameter instability suggests that our estimate of the unconditional covariance between the forward premium and rate of depreciation

TABLE 4. Stability tests of unbiasedness: weekly data 1975:2 to 1980:1 and 1980:6 to 1989:52.

Currency	Coefficients on regressors					
	Unconstrained			Constrained		
	Const. (s.e.) Conf. ($\alpha = 0$)	$f_t - s_t$ (s.e.) Conf. ($\beta_1 = 1$)	$\chi^2(3)$ Conf.	R ²	Const. (s.e.) Conf. ($\alpha = 0$)	$f_t - s_t$ (s.e.) Conf. ($\beta_1 = 1$)
	Panel A: First sub-sample					
Mark	3.070 (5.367)	1.040 (1.313)		0.006	3.167 (3.108)	0.896 (0.551)
Pound	0.433 4.791 (5.832)	0.024 1.623 (1.162)	7.845 0.951	0.027	0.692 -0.526 (4.176)	0.150 0.896 (0.551)
Yen	0.589 8.148 (3.160) 0.990	0.409 -1.044 (0.907) 0.976		0.015	0.100 8.685 (2.725) 0.999	0.150 0.896 (0.551) 0.150
	Panel B: Second sub-sample					
Mark	19.338 (7.633)	-4.941 (1.827)		0.054	17.992 (4.433)	-4.601 (0.594)
Pound	0.989 -10.529 (3.222)	0.999 -4.113 (0.921)	31.672 0.999	0.104	0.999 -12.491 (2.754)	0.999 -4.601 (0.594)
Yen	0.999 17.770 (5.723) 0.998	0.999 -3.007 (1.044) 0.999		0.049	0.999 23.329 (4.430) 0.999	0.999 -4.601 (0.594) 0.999
	Panel C: Chi-square tests for stability					
Chow-type test	20.733 0.998				Sample 1 14.693	Sample 2 7.587
Predictive test	14.543 0.976				0.935	0.525

Notes: See also Table 1. The Chow-type test is described in Hodrick and Srivastava (1984) and the Predictive Test is described in Ghysels and Hall (1990). The unconstrained estimation is heteroskedasticity-consistent OLS and the constrained estimation is heteroskedasticity-consistent seemingly unrelated regression with correction for the overlap in the data. The over-identification tests are distributed as $\chi^2(8)$.

might not be a very good one. Because the parameters have apparently changed over time, estimation that assumes constant parameters is inappropriate, and one should allow for endogeneous changes in the parameters during the estimation.

One way that this can be done is to build on the approach of Engel and Hamilton (1990) who develop a Markov switching model for the rate of depreciation of the dollar relative to the deutsche mark, the British pound, and the French franc. Engel and Hamilton (1990) use end of quarter spot exchange rates. They postulate that the rate of depreciation is characterized by two regimes with different means and variances and with constant probabilities of transition between the regimes. They use maximum likelihood estimation and find significant differences in the means and variances of the rates of depreciation for the two regimes.⁶

Engel and Hamilton (1990) also examine the relation of interest rate differentials, which are equivalent to forward premiums because of covered interest arbitrage, to their measures of expected depreciation. In doing so, they encounter an awkwardness in the specification of the unbiasedness hypothesis. Since their model has no autoregressive dynamics other than the Markov process, there are only two expected rates of depreciation. For example, the conditional rate of depreciation when the economy is in state one is the probability of remaining in state one times the state one mean rate of depreciation plus the probability of a transition to state two times the mean rate of depreciation for state two. Since the interest differential is a continuous variable and is highly autocorrelated, it obviously does not fit this two-state characterization. Engel and Hamilton introduce measurement error in the observation of the interest differential to solve this problem.

Our version of the Markov state model overcomes this difficulty by incorporating explicit autoregressive dynamics. We use correctly sampled monthly data as above and simultaneously model the rate of depreciation and the forward premium. The specification of the model retains the two-state Markov process, with transition probability p_{11} (p_{22}) of remaining in state one (state two) given that the economy is in state one (state two), but the conditional means in each state are allowed to depend autoregressively on lagged values of the rate of depreciation and the forward premium as in the following equations for $i = 1, 2$:

$$\langle 12 \rangle \quad \Delta s_{t+1} = a_{11,i} + b_{11,i} \Delta s_t + b_{12,i} f p_t + \varepsilon 1_{t+1,i},$$

$$\langle 13 \rangle \quad f p_{t+1} = a_{21,i} + b_{21,i} \Delta s_t + b_{22,i} f p_t + \varepsilon 2_{t+1,i}.$$

The innovations for regime i are $\varepsilon'_{t+1,i} = (\varepsilon 1_{t+1,i}, \varepsilon 2_{t+1,i})$, and they are assumed to be $N(0, \Sigma_i)$. Hence, there are 20 parameters in the model: the two transition probabilities, the four a_i 's and eight b_i 's, and the six distinct elements of the Σ_i 's.

The parameter estimates are presented in Table 5. Several features of the model are noteworthy. First, notice that the expected rate of depreciation given that the economy is in state 1 is

$$\langle 14 \rangle \quad p_{11}(a_{11,1} + b_{11,1} \Delta s_t + b_{12,1} f p_t) + (1 - p_{11})(a_{11,2} + b_{11,2} \Delta s_t + b_{12,2} f p_t).$$

Hence, the unbiasedness hypothesis requires that $a_{11,i} = b_{11,i} = 0$ and $b_{12,i} = 1$, for $i = 1, 2$. This is clearly rejected in the parameter estimates. The values of $b_{12,1}$ with their standard errors in parentheses are -4.113 (2.036) for the deutsche

TABLE 5. Markov regime switching models.

Coefficient	DM		Pound		Yen	
	Regime 1	Regime 2	Regime 1	Regime 2	Regime 1	Regime 2
p_{ii}	0.984 (0.021)	0.941 (0.061)	0.961 (0.049)	0.961 (0.043)	0.994 (0.021)	0.994 (0.013)
$a_{11,i}$	17.752 (8.618)	3.953 (18.295)	-15.263 (9.099)	-9.166 (3.999)	43.387 (11.875)	5.708 (8.504)
$b_{11,i}$	-0.100 (0.074)	0.023 (0.206)	-0.202 (0.113)	0.271 (0.096)	0.010 (0.104)	0.056 (0.108)
$b_{12,i}$	-4.113 (2.306)	-2.237 (3.138)	-7.906 (2.809)	-0.911 (0.988)	-11.386 (3.766)	-1.373 (1.143)
$a_{21,i}$	-0.192 (0.124)	-2.424 (1.067)	0.096 (0.166)	0.327 (0.294)	-0.531 (0.234)	-0.497 (0.318)
$b_{21,i}$	0.002 (0.001)	-0.019 (0.006)	0.002 (0.003)	0.003 (0.006)	-0.002 (0.002)	-0.0002 (0.006)
$b_{22,i}$	0.942 (0.030)	0.453 (0.163)	0.957 (0.050)	0.846 (0.057)	0.831 (0.063)	0.885 (0.052)
$\sigma_{11,i}^2$	1623.286 (178.709)	1542.575 (411.683)	1749.450 (321.107)	1159.466 (184.847)	1346.185 (250.708)	1554.962 (219.163)
$\sigma_{12,i}$	-1.693 (2.383)	-25.339 (22.296)	8.011 (4.081)	-5.883 (8.678)	-0.940 (2.775)	-3.448 (8.610)
$\sigma_{22,i}^2$	0.385 (0.055)	4.353 (1.068)	0.621 (0.108)	5.221 (0.619)	0.404 (0.068)	4.684 (0.770)

Notes: The model is described in equations <12> to <13> of the text. The parameters are estimated using maximum likelihood with the EM-algorithm. A technical appendix supplies details of this estimation method. Standard errors are in parentheses.

mark, -7.906 (2.809) for the pound, and -11.386 (3.766) for the yen. These are quite far from one. The values of $b_{12,2}$ are also negative, but they are approximately two standard errors from one.

One major difference in the two regimes is the difference in the variances of the forward premiums. The variances of the forward premiums are nine to ten times larger in regime two than in regime one. Figure 1 presents the smoothed probability estimates that the Markov regime is in state one, which is the low variance state. For all three currencies, the probability goes essentially to zero during the 1979-82 period of monetary targeting and increased interest rate variability. After 1982, the probability that the regime is state one is mostly very close to one, and this is the regime with large negative coefficients on the forward premiums in the equations for the means of currency depreciations. These results contrast with those of Engel and Hamilton (1990) who find only one change in regime in 1977 for the pound and only two transitions for the DM, one in 1977 and one at the end of 1984.

As we noted above, when there is more than one regime, OLS may give a bad estimate of the unconditional moments of the true process. We therefore calculate the slope coefficient implied by the Markov switching model as the unconditional covariance of the forward premium and the future rate of depreciation divided by the unconditional variance of the forward premium. The derivation of the

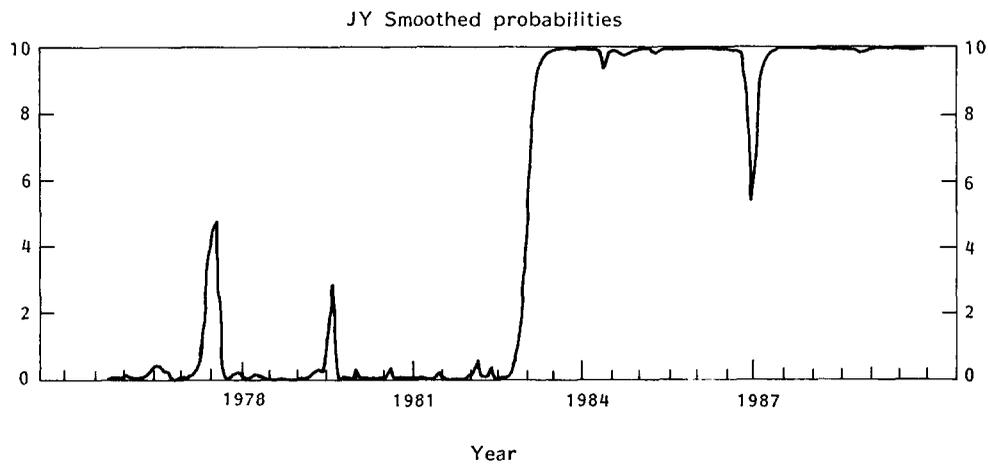
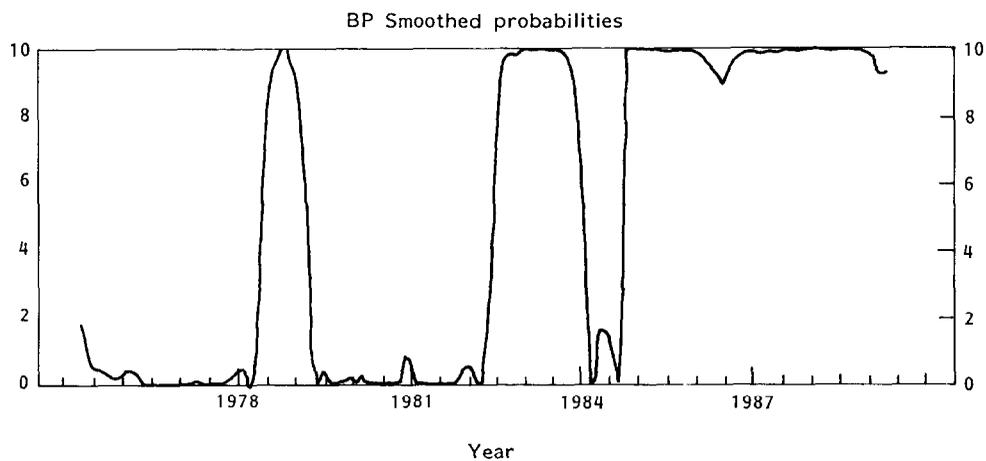
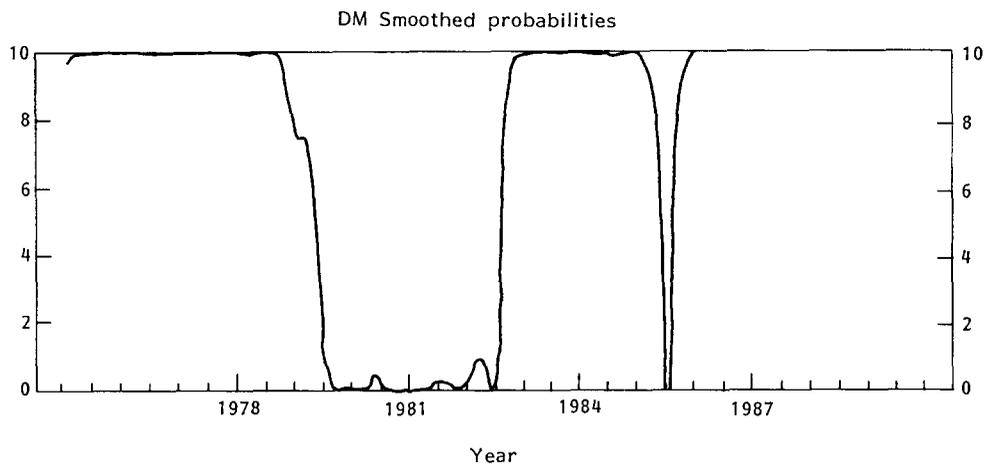


FIGURE I.

unconditional moments is provided in the technical appendix. The point estimates with standard errors in parentheses are -3.389 (1.856) for the deutsche mark, -4.557 (2.314) for the pound, and -6.479 (6.553) for the yen. These implied slope coefficients are more negative but less precisely estimated than the OLS estimates.

V. Conclusions

The paper investigates several sources of bias that could mitigate the burden on a time varying risk premium as the explanation for the consistent rejection of the unbiasedness hypothesis in the relation of the forward premium to the expected rate of currency depreciation. The first source of potential bias is measurement error coming either from incorrect sampling of the data or from failure to account for bid-ask spreads. Both sources were shown to be relatively unimportant.

The next source of bias is an omitted variable problem. There is conditional heteroskedasticity in rates of currency depreciation, and the conditional variance enters the mean rate of depreciation. If the forward premium is correlated with the conditional variance, there is an omitted variable bias in regressions of the rate of depreciation on the forward premium. Monte Carlo experiments demonstrated that this source of bias is also relatively unimportant in explaining the empirical results in the literature.

The last part of the paper investigates instability in the typical regression specification of an unbiasedness test. Formal tests of the stability of the coefficients indicate that the parameters have changed over time. This motivates an investigation of an endogenous regime shifting model. In some respects the results of this model are completely intuitive to someone familiar with the stylized facts of the 1980s. One regime is characterized by a variance for the forward premium that is ten times the variance of the forward premium in the other regime. This regime corresponds to the period of interest rate volatility during the early 1980s. The second regime is characterized by a large negative coefficient on the forward premium in the conditional mean of the rate of depreciation. This regime corresponds to the period in the 1980s in which foreign currencies were at a forward premium relative to the dollar, yet the dollar appreciated throughout the period.

After considering alternative sources of bias, our conclusion is that the evidence against the unbiasedness hypothesis using rational expectations econometrics is very strong. If econometric misspecifications cannot explain away deviations from unbiasedness, the interesting question becomes whether rational expectations economic theory can account for the extremely variable forward market risk premiums implied by the negative slope coefficients. While there are many ways to generate risk premiums, the results of studies by Bekaert (1992) and Canova and Marrinan (1990), who simulate and estimate simple, two-country, general equilibrium models of the cash-in-advance variety, are not very encouraging. The models fail to generate significant variability in risk premiums. Of course, better models of risk can be developed. For instance, both papers assume state and time separable intertemporal preferences whereas nonseparabilities are known to have dramatic effects on equilibrium asset returns.

The particular form of parameter instability documented in this paper suggests a useful alternative direction for future theoretical research. As we showed, the

slope coefficients became more negative in the 1980s, a decade of major changes in both financial markets and fiscal and monetary policies. Rational agents, faced with such a turbulent economic environment, might need time to recognize or understand changes in policy regimes. Such rational ‘learning’ can lead to systematic forecast errors (see, for example, Kaminsky and Peruga, 1988; and Lewis, 1989). If learning is the source of the negative slope coefficients, we might also expect to find negative slope coefficients at the start of the floating exchange rate period in 1973–76. For example, after the breakdown of the Bretton Woods system in February 1973, there was considerable uncertainty about how the international monetary system would evolve. Agents might have been expecting a return to a fixed rate regime and might not have fully understood the monetary policies of the central banks in the beginning of this new era.

To verify this conjecture, we carried out the regression test on weekly data from June 1973 to January 1976.⁷ To determine the comparability of these data with the data used in this paper, we first checked whether this data set mimics the results of the 1975–80 period reported in Table 4. The slope coefficients with standard errors in parentheses are 0.898 (1.252) for the deutsche mark, 1.267 (0.983) for the pound, and -1.147 (0.863) for the yen. These estimates are qualitatively similar to the ones obtained in Table 4. In particular, they imply that for the 1975–80 period, there is no evidence against the unbiasedness hypothesis for the deutsche mark and the pound.

If our conjecture on learning is right, we expect to find negative slope coefficients for the 1973–76 period. The slope coefficients with their standard errors in parentheses are: -1.824 (2.857) for the deutsche mark, -2.730 (0.684) for the pound, and 0.565 (0.301) for the yen. One can interpret the results, at least for the deutsche mark and the pound, as evidence in favor of general equilibrium models that incorporate some form of rational learning about policy regimes. This is a challenging area for future research. Without additional analysis, though, the results could just as easily be interpreted as evidence of market inefficiency as in Froot and Thaler (1990).

Technical appendix 1: the EM-algorithm

In Section IV the following first-order Markov model is estimated:

$$\langle A1 \rangle \quad y_t = A_i x_{t-1} + \varepsilon_{t,i},$$

where $y_t = [\Delta s_t, fp_t]'$, $x_t = [1, y_{t-1}]'$, $\varepsilon_{t,i} \sim N(0, \Sigma_i)$, and $i = 1, 2$. The six coefficients in A_i and the three distinct parameters in Σ_i are drawn from two regimes governed by an unobserved state variable, z_t , which takes on only two values, 1 and 2. The Markov transition probability matrix is therefore fully characterized by p_{11} , the probability of staying in state 1 given state 1, and p_{22} , the probability of staying in state 2 given state 2.

The complete 20-element parameter vector is therefore $\Theta = [\text{vec}(A_1)', \text{vec}(A_2)', \text{vech}(\Sigma_1)', \text{vech}(\Sigma_2)', p_{11}, p_{22}]'$. Let \tilde{y}_T denote a sample of observed data, $\{y_0, \dots, y_T\}$, and let \tilde{z}_T denote the sample of unobserved states, $\{z_0, \dots, z_T\}$. The joint likelihood function of \tilde{y}_T and \tilde{z}_T is $L(\tilde{y}_T, \tilde{z}_T; \Theta)$. Maximum likelihood estimation requires specification of the log-likelihood function of the observed data, $L(\tilde{y}_T; \Theta)$. A computationally convenient estimation method is the EM-algorithm. The method is equivalent to iterating on the normal equations (the first-order conditions for the maximization of the likelihood function). In describing the algorithm, we adopt the notation and approach of Ruud (1991).

The EM-algorithm consists of two steps. In the E-step, one forms the expectation of the log-likelihood function of the observed and unobserved data, $L(\tilde{y}_T, \tilde{z}_T; \Theta)$, conditional on \tilde{y}_T and an initial parameter vector Θ_0 . We denote this as

$$\langle A2 \rangle \quad Q(\Theta, \Theta_0, \tilde{y}_T) = E[L(\tilde{y}_T, \tilde{z}_T; \Theta) | \tilde{y}_T, \Theta_0],$$

where the expectation is taken with respect to the density of the unobserved state variables conditional on the whole sample of observed data and an initial parameter vector Θ_0 . In the M-step, the function $Q(\Theta, \Theta_0, \tilde{y}_T)$ is maximized by choice of Θ , and the argmax yields Θ_1 , which replaces Θ_0 for the next iteration. This recursive procedure converges to the MLE of $L(\tilde{y}_T; \Theta)$, which follows from the results in equations $\langle A3 \rangle$ and $\langle A4 \rangle$:

$$\langle A3 \rangle \quad Q(\Theta_1, \Theta_0, \tilde{y}_T) \geq Q(\Theta_0, \Theta_0, \tilde{y}_T) \Rightarrow L(\tilde{y}_T; \Theta_1) \geq L(\tilde{y}_T; \Theta_0),$$

$$\langle A4 \rangle \quad Q_1(\hat{\Theta}, \hat{\Theta}, \tilde{y}_T) = L_1(\tilde{y}_T; \hat{\Theta}),$$

where $\hat{\Theta}$ denotes the MLE and the subscripts on Q and L denote partial derivatives. Proofs of these results can be found in Ruud (1991) and Hamilton (1990). They guarantee that each step that increases Q also increases L , and that maximization of Q is equivalent to the maximization of L .

The EM-algorithm is a particularly convenient maximization method for this application because of the following two assumptions of our model:

- (a) y_t depends only on z_t and y_{t-1} .
- (b) z_t depends only on z_{t-1} and is independent of the history of y_t .

Let $h_t = (z_t, y'_{t-1})$. Then, the conditional distribution of y_t given h_t is:

$$\langle A5 \rangle \quad f(y_t | h_t) = N(A_{z_t} x_{t-1}, \Sigma_{z_t}).$$

Conditioning on an initial value y_0 , the use of assumptions (a) and (b) together with recursive conditioning leads to:

$$\langle A6 \rangle \quad L(\tilde{y}_T, \tilde{z}_T; \Theta) = \sum_{t=2}^T [\log(f(y_t | h_t, \Theta)) + \log(p(z_t | z_{t-1}))] \\ + \log(f(y_1 | z_1, y_0, \Theta)) + \log(p(z_1))$$

where $p(z_t | z_{t-1})$ can take on four different values: $p_{11}, (1 - p_{11}) = p_{12}, p_{22}, (1 - p_{22}) = p_{21}$. Computing $Q(\Theta, \Theta_0, \tilde{y}_T)$ from $\langle A6 \rangle$ is straightforward. Averaging occurs with respect to the probability of the state variables given the whole sample \tilde{y}_T and some initial Θ_0 . These 'smoothed state probabilities' are easily computed using the filter described in Hamilton (1990).

We obtain:

$$\langle A7 \rangle \quad Q(\Theta, \Theta_0, \tilde{y}_T) = \sum_{z_t=1}^2 \sum_{t=2}^T \log(f(y_t | h_t, \Theta)) p(z_t | \tilde{y}_T, \Theta_0) \\ + \sum_{z_1=1}^2 \log(f(y_1 | z_1, y_0)) p(z_1 | \tilde{y}_T, \Theta_0) \\ \times \sum_{z_t=1}^2 \sum_{z_{t-1}=1}^2 \sum_{t=2}^T \log(p(z_t | z_{t-1})) p(z_t, z_{t-1} | \tilde{y}_T, \Theta_0) \\ + \sum_{z_1=1}^2 \log(p(z_1)) p(z_1 | \tilde{y}_T, \Theta_0),$$

where we have used the unconditional probability of z_t at $t = 1$ as a starting value. Another approach would be to estimate the start-up values, as Hamilton (1990) does.

Maximization of $Q(\Theta, \Theta_0, \tilde{y}_T)$ with respect to Θ is now straightforward, and the first-order conditions give rise to the following estimates for $i = 1, 2$:

$$\langle A8 \rangle \quad A_i = \sum_{t=1}^T y_t x'_{t-1} p(z_t = i | \tilde{y}_T, \Theta_0) \left(\sum_{t=1}^T x_{t-1} x'_{t-1} p(z_t = i | \tilde{y}_T, \Theta_0) \right)^{-1},$$

$$\langle A9 \rangle \quad \Sigma_i = \frac{\sum_{t=1}^T (y_t - A_i x_{t-1})(y_t - A_i x_{t-1})' p(z_t = i | \tilde{y}_T, \Theta_0)}{\sum_{t=1}^T p(z_t = i | \tilde{y}_T, \Theta_0)},$$

$$\langle A10 \rangle \quad p_{11} = \frac{\sum_{t=2}^T p(z_t = 1, z_{t-1} = 1 | \tilde{y}_T, \Theta_0)}{\sum_{t=2}^T (p(z_{t-1} = 1 | \tilde{y}_T, \Theta_0) + \rho - p(z_t = 1 | \tilde{y}_T, \Theta_0))},$$

$$\langle A11 \rangle \quad p_{22} = \frac{\sum_{t=2}^T p(z_t = 2, z_{t-1} = 2 | \tilde{y}_T, \Theta_0)}{\sum_{t=2}^T (p(z_{t-1} = 2 | \tilde{y}_T, \Theta_0) - \rho + p(z_t = 2 | \tilde{y}_T, \Theta_0))},$$

where $\rho = p_{21}/(p_{12} + p_{21})$ arises from the start-up conditions.

The derivatives of Q with respect to A and Σ are found using the matrix-derivative results in Amemiya (1985, pp. 461–462). The estimates obtained in equations $\langle A8 \rangle$ to $\langle A11 \rangle$ constitute the new Θ_0 which is then used to compute smoothed probabilities as input for the next iteration. The iterations are stopped as soon as the maximal element of $|\Theta - \Theta_0|$ is smaller than 10^{-10} .

As Ruud (1991) emphasizes, the score of the likelihood function of the data is readily available in the EM-algorithm, so that an estimator of the information matrix $\hat{\mathcal{F}}$ is easily computed. Standard errors for the parameter estimates are then found by taking the square root of the diagonal elements of $-\hat{\mathcal{F}}^{-1}/T$. The information matrix is estimated by

$$\langle A12 \rangle \quad \hat{\mathcal{F}} = -T^{-1} \sum_{t=1}^T Q_t(\hat{\Theta}, \hat{\Theta}, y_t) Q_t(\hat{\Theta}, \hat{\Theta}, y_t)'$$

This requires the construction of the normal equations evaluated at the optimum.

Technical appendix 2: unconditional moments in the Markov regime switching model

To derive unconditional moments in the Markov regime switching model, it is useful to partition the VAR parameter matrix A_i (equation $\langle A1 \rangle$ in technical appendix 1) as $A_i = [a_i, b_i]$, with a_i representing the constants and b_i the autoregressive parameters. The model can be rewritten as:

$$\langle A13 \rangle \quad y_t = a_i + b_i y_{t-1} + \varepsilon_{t,i}.$$

Using property (b) in the technical appendix and the covariance-stationarity of the y_t process, the unconditional mean of y_t is given by:

$$\langle A14 \rangle \quad E[y_t] = (I - p_1 b_1 - p_2 b_2)^{-1} (p_1 a_1 + p_2 a_2),$$

where $p_1(p_2)$ is the unconditional probability of the first (second) state.

To derive the unconditional variance, we first compute the uncentered second moment. Taking the unconditional expectation of $y_t y_t'$ yields:

$$\langle A15 \rangle \quad E[y_t y_t'] = \mu + p_1 b_1 E[y_{t-1} y_{t-1}'] b_1' + p_2 b_2 E[y_{t-1} y_{t-1}'] b_2' + p_1 \Sigma_1 + p_2 \Sigma_2$$

where μ is a constant given by:

$$\langle A16 \rangle \quad \begin{aligned} \mu = & p_1 (a_1 a_1' + a_1 E[y_{t-1}]' b_1' + b_1 E[y_{t-1}] a_1') \\ & + p_2 (a_2 a_2' + a_2 E[y_{t-1}]' b_2' + b_2 E[y_{t-1}] a_2'). \end{aligned}$$

Denote the vec-operator by $\text{vec}(\cdot)$. Then $\text{vec}(E[y_t y_t'])$ follows from covariance-stationarity and the fact that, if P, Q, R are conformable matrices, $\text{vec}(PQR)$ equals $(R' \otimes P) \text{vec}(Q)$:

$$\langle A17 \rangle \quad \begin{aligned} \text{vec}(E[y_t y_t']) = & (\mathbf{I} - p_1 (b_1 \otimes b_1) - p_2 (b_2 \otimes b_2))^{-1} \\ & \times (\text{vec}(\mu) + p_1 \text{vec}(\Sigma_1) + p_2 \text{vec}(\Sigma_2)). \end{aligned}$$

The unconditional covariance matrix, $C(0)$, is then simply $E[y_t y_t'] - E[y_t] E[y_t]'$.

To derive the covariance between the forward premium and future currency depreciation, we also need to derive:

$$\langle A18 \rangle \quad E[y_t y_{t+1}'] = p_1 (E[y_t] a_1' + E[y_t y_t'] b_1') + p_2 (E[y_t] a_2' + E[y_t y_t'] b_2'),$$

which uses the law of iterated expectations. The first-order covariance matrix is

$$\langle A19 \rangle \quad C(1) = E[y_t y_{t+1}'] - E[y_t] E[y_{t+1}]'$$

Define the index vectors $e1 = [1, 0]'$ and $e2 = [0, 1]'$. The implied slope coefficient β is the unconditional covariance between the forward premium and the future currency depreciation, $e2' C(1) e1$, divided by the unconditional variance of the forward premium, $e2' C(0) e2$.

Note that β is a non-linear function, $f(\Theta)$, of Θ , the vector of 20 parameters estimated with the EM-algorithm. Hence, standard errors for β can be derived from the standard Mean Value Theorem, as:

$$\langle A20 \rangle \quad \sqrt{T}(f(\Theta) - f(\Theta_0)) \sim N(0, \nabla f \Omega \nabla f')$$

where Θ_0 is the true parameter vector and Ω is the variance-covariance matrix of Θ , which is computed with the method described in technical appendix 1. Numerical gradients are used to calculate the gradient of $f(\Theta)$ evaluated at Θ .

Notes

1. Bossaerts and Hillion (1991) argue that use of average bid-ask rates leads to inconsistent parameter estimates in investigations of forward rates and future spot rates. Korajczyk and Viallet (1992) use monthly data and find little difference in inference in their analysis using correctly sampled data versus averaged data. Bossaerts and Hillion (1991) use daily data in which case the problem is more severe.
2. Gregory and McCurdy (1984) were the first to question the stability of coefficients within the context of these studies.
3. See Sibert (1989) and Engel (1992) for additional discussion of the nature of risk premiums in the forward foreign exchange market.
4. Domowitz and Hakkio (1985) estimate a univariate version of the model in Table 3. We abstracted from the conditional covariance term in equation <6>. Under risk neutrality, the conditional covariance in equation <6> is between the rate of depreciation and the inverse of domestic inflation (see equation <2>). Frenkel and Razin (1980) and Kaminsky and Peruga (1990) find this unconditional covariance to be small, but Kaminsky and Peruga stress that the conditional covariance has explanatory power for the rate of change of the exchange rate.

5. Froot and Thaler (1990, p. 182) note that in the large literature testing the unbiasedness hypothesis, most estimates of β are negative. They state, 'A few are positive, but not one is equal to or greater than the null hypothesis of $\beta = 1$.' Clearly, this is not true in the early part of our sample. One reason our results differ from the literature may be that our sample begins in 1975 and many early studies such as Bilson (1981) used data beginning in 1974 or earlier.
6. Engel (1991) extends the two-state Engel–Hamilton model to 18 exchange rates and examines monthly as well as quarterly data.
7. These data are from Data Resources, Inc. and were used by Hansen and Hodrick (1983), who noted that January 1976 corresponds to the date of ratification by the Interim Committee of the IMF of the Rambouillet agreement that formally implemented a system of flexible exchange rates.

References

- AMEMIYA, TAKESHI, *Advanced Econometrics*, Cambridge, MA: Harvard University Press, 1985.
- BAILLIE, RICHARD T., AND TIM BOLLERSLEV, 'The Message in Daily Exchange Rates: A Conditional Variance Tale,' *Journal of Business and Economic Statistics*, July 1989, **7**: 297–305.
- BAILLIE, RICHARD T., AND TIM BOLLERSLEV, 'A Multivariate Generalized ARCH Approach to Modeling Risk Premia in Forward Foreign Exchange Markets,' *Journal of International Money and Finance*, September 1990, **9**: 309–324.
- BEKAERT, GEERT, 'Exchange-Rate Volatility and Deviations from Unbiasedness in a Cash-in-Advance Model,' manuscript, Northwestern University, 1992.
- BILSON, JOHN F.O., 'The Speculative Efficiency Hypothesis,' *Journal of Business*, 1981, **54**: 433–451.
- BOLLERSLEV, TIM, 'Modeling the Coherence in Short Run Nominal Exchange Rates: A Multivariate Generalized ARCH Model,' *Review of Economics and Statistics*, 1990, **72**, 498–505.
- BOSSAERTS, PETER, AND PIERRE HILLION, 'Market Microstructure Effects of Government Intervention in the Foreign Exchange Market,' *Review of Financial Studies*, 1991, **4**: 513–541.
- CANOVA, FABIO, AND JANE MARRINAN, 'Nominal Profits, Risk and Uncertainty in Foreign Exchange Markets,' manuscript, Brown University, 1990.
- CORNELL, BRADFORD, 'The Impact of Data Errors on Measurement of the Foreign Exchange Risk Premium,' *Journal of International Money and Finance*, March 1989, **8**: 147–157.
- DOMOWITZ, IAN, AND CRAIG S. HAKKIO, 'Conditional Variance and the Risk Premium in the Foreign Exchange Market,' *Journal of International Economics*, August 1985, **19**: 47–66.
- ENGEL, CHARLES M., 'Can the Markov Switching Model Forecast Exchange Rates?,' unpublished manuscript, University of Washington, July 1991.
- ENGEL, CHARLES M., 'On the Foreign Exchange Risk Premium in a General Equilibrium Model,' *Journal of International Economics*, May 1992, **32**: 305–319.
- ENGEL, CHARLES M., AND JAMES D. HAMILTON, 'Long Swings in the Dollar: Are They in the Data and Do Markets Know It?,' *American Economic Review*, September 1990, **80**: 689–713.
- ENGLE, ROBERT F., 'Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation,' *Econometrica*, July 1982, **50**: 987–1007.
- FAMA, EUGENE F., 'Forward and Spot Exchange Rates,' *Journal of Monetary Economics*, November 1984, **14**: 319–338.
- FRENKEL, JACOB A., AND ASSAF RAZIN, 'Stochastic Prices and Tests of Efficiency of Foreign Exchange Markets,' *Economic Letters*, 1980, **6**: 165–170.
- FROOT, KENNETH A., AND RICHARD H. THALER, 'Anomalies: Foreign Exchange,' *The Journal of Economic Perspectives*, Summer 1990, **4**: 179–192.
- GHYSELS, ERIC, AND ALASTAIR HALL, 'Are Consumption-Based Intertemporal Capital Asset Pricing Models Structural?,' *Journal of Econometrics*, July/August 1990, **45**: 121–139.
- GREGORY, ALLAN W., AND THOMAS H. MCCURDY, 'Testing the Unbiasedness Hypothesis in

- the Forward Foreign Exchange Market: A Specification Analysis,' *Journal of International Money and Finance*, December 1984, **3**: 357–368.
- HAMILTON, JAMES D., 'Analysis of Time Series Subject to Changes in Regime,' *Journal of Econometrics*, July/August 1990, **45**: 39–70.
- HANSEN, LARS PETER, 'Large Sample Properties of Generalized Method of Moments Estimators,' *Econometrica*, July 1982, **50**: 1029–1054.
- HANSEN, LARS PETER, AND ROBERT J. HODRICK, 'Risk Averse Speculation in the Forward Foreign Exchange Market: An Econometric Analysis of Linear Models,' in Jacob A. Frenkel, ed., *Exchange Rates and International Macroeconomics*, Chicago, IL: University of Chicago Press, 1983.
- HODRICK, ROBERT J., 'Risk, Uncertainty, and Exchange Rates,' *Journal of Monetary Economics*, May 1989, **23**: 433–459.
- HODRICK, ROBERT J., AND SANJAY SRIVASTAVA, 'An Investigation of Risk and Return in Forward Foreign Exchange,' *Journal of International Money and Finance*, April 1984, **3**: 5–29.
- KAMINSKY, GRACIELA, AND RODRIGO PERUGA, 'Credibility Crises: The Dollar in the Early Eighties,' Discussion Paper 88-15, University of California, San Diego, 1988.
- KAMINSKY, GRACIELA, AND RODRIGO PERUGA, 'Can a Time-Varying Risk Premium Explain Excess Returns in the Forward Market for Foreign Exchange?,' *Journal of International Economics*, February 1990, **28**: 47–70.
- KORAJCZYK, ROBERT A., AND CLAUDE VIALLET, 'Equity Risk Premia and the Pricing of Foreign Exchange Risk,' *Journal of International Economics*, November 1992, **33**: 199–219.
- LEWIS, KAREN K., 'Can Learning Affect Exchange-Rate Behavior,' *Journal of Monetary Economics*, January 1989, **23**: 79–100.
- LJUNG, G. M., AND GEORGE E. P. BOX, 'On a Measure of Lack of Fit in Time Series Models,' *Biometrika*, March 1978, **65**: 297–303.
- LUCAS, ROBERT E. JR., 'Interest Rates and Currency Prices in a Two-Country World,' *Journal of Monetary Economics*, November 1982, **10**: 335–360.
- PAGAN, ADRIAN R., AND H. C. L. SABAU, 'Consistency Tests for Heteroskedastic and Risk Models,' Australian National University, working paper, 1987.
- RUUD, PAUL A., 'Extensions of Estimation Methods using the EM Algorithm,' *Journal of Econometrics*, September 1991, **49**: 305–342.
- SIBERT, ANNE, 'The Risk Premium in the Foreign Exchange Market,' *Journal of Money, Credit and Banking*, February 1989, **21**: 49–65.
- WHITE, HALBERT, 'Maximum Likelihood Estimation of Misspecified Models,' *Econometrica*, January 1982, **50**: 1–26.