THE TIME-SERIES PROPERTIES OF THE RISK PREMIUM IN THE YEN/DOLLAR EXCHANGE MARKET

FABIO CANOVA

Department of Economics, Box B, Brown University, Providence, RI 02912 and Department of Economics University of Rochester, Rochester, NY 14627, USA

AND

TAKATOSHI ITO

Department of Economics, Hitotsubashi University, Kunitachi-shi, Tokyo, 186 Japan; Department of Economics University of Minnesota, Minneapolis, MN 55455; and National Bureau of Economic Research, Cambridge, MA, 02138

SUMMARY

In this paper a VAR model is employed to construct a measure of the conditional expectations of the future yen/dollar spot rate. This measure allows us to examine the dynamics of an *ex-ante* time-series for the risk premium in the market. The VAR model produces 'better' forecasts than the survey responses for turbulent periods such as 1981–1982 and 1984–1985. The VAR-generated expectations are then used to construct a risk premium time-series. This risk premium series seems to be more reliable than the ones obtained using either survey data on expectations of the future spot exchange rate or the *ex-post* realized spot exchange rate. Tests on the risk premium series suggest that a risk premium was present, but that it was virtually constant throughout the sample. The conditional variance of the risk premium changed over time, but its unconditional distribution seemed stable across subsamples. Despite these features, the volatility of the series was substantial and varied considerably throughout the sample.

1. INTRODUCTION

The size and the variability of the risk premium in foreign exchange markets are of crucial importance in analysing issues concerning the efficiency of these markets. Several authors (for example, Frankel, 1986; Frankel and Meese, 1987) have argued that, theoretically, the risk premium should be small and approximately constant. The experience of the 1980s in several foreign exchange markets indicates that the *ex-post* bias of the forward rate is very large and volatile. Since the *ex-post* forward bias is the sum of a forecast error and of a risk premium, many researchers find it hard to attribute the empirical features of the forward bias to a risk premium, and conclude that there are persistent patterns in the forecast errors made by agents. Some have interpreted this evidence as an indication that a 'peso problem' may have occurred in the 1980s (see e.g. Krugman, 1987). Others have used this result to question the rationality of investors' expectations (see e.g. Frankel and Froot, 1987).

Since the risk premium is unobservable, a crucial step in investigating the efficient market hypothesis is the construction of (a proxy for) the risk premium. The empirical evidence on the properties of risk premium proxies is somewhat contradictory. Hansen and Hodrick (1980, 1983), Hodrick and Srivastava (1984, 1986), Cosset (1984), Domowitz and Hakkio (1985), and Cumby (1988) suggest that the risk premium is important in explaining the *ex-post* bias of the forward rate and the existence of complex pattern of heteroskedasticity and nonlinearities in

the structure of the forecast error. Frankel (1982), Frankel and Engel (1984), and Giovannini and Jorion (1988), find little support for the hypothesis that the risk premium played a key role in explaining the forward bias.

Recently, Frankel and Froot (1987), Froot and Frankel (1989), Dominguez (1986) and Ito (1990) have used survey data on expectations of future spot rates to construct an *ex-ante* measure of the risk premium. All of them find that forecast errors explain most of the size and of the variability of the forward biases, and that agents have 'wishful expectations' or are systematically wrong in their forecasts.

Although obtaining direct estimates for expectations is an advantage of the survey method, a consistent data source for a long enough time span is not easily available.

This paper provides an alternative way of constructing an *ex-ante* time-series for the risk premium. We employ a vector autoregressive (VAR) forecasting model to construct a measure of market expectations for the future spot rate as a *k*-step ahead linear forecast, conditional on the amount of information available at the time of expectation formation. The risk premium is then defined as the difference between the model-generated expected spot and the forward rate.

The VAR-constructed measure of the risk premium is useful in several respects. It allows us to make statements concerning the correlation of the risk premium with variables in the information set, its predictability given the forward premium, and its relationship with the expected change in the spot rate (see Fama, 1984; Hodrick and Srivastava, 1984, 1986, for similar exercises using realized values under the maintained hypothesis of rational expectation). Also, our measure of the risk premium would exactly replicate the one produced using survey data, if agents used projection techniques in computing their expectations of future spot rates. Therefore, comparing the two risk premium series may help to identify certain possible deviations from a VAR-rational behaviour.

To examine the statistical features of the risk premium series in detail, and to test some interesting hypotheses, we employ a time-varying coefficient (TVC) model of the Doan, Litterman and Sims (1984) variety. This statistical model is quite general, it includes as a special case several nonlinear time-series specifications recently employed to model the first two moments of financial data, it has a recursive conditional Gaussian structure and, if used as a data-generating mechanism, it produces non-normalities in the generated time-series (see Canova, 1990, for these properties).

In this paper we concentrate on the analysis of the risk premium in the yen/dollar market. Canova (1991) extends this analysis to several other markets. The yen/dollar market in the 1980s offers an interesting laboratory to study the properties of the risk premium series. From the third week of 1981 to the 44th week of 1982, the yen depreciated from 199 yen per dollar to 276 yen per dollar, a depreciation of 38 per cent. This 2-year spell of sharp yen depreciation took place in the presence of a large yen forward premium. (Note that the forward premium is equal to the interest rate differential between the two countries because of *covered* interest parity. Ito (1986) gives details for the covered interest parity between the yen and the US dollar.) The 3-month interest rate for the dollar-denominated asset was about 10 per cent higher than the one for the yen-denominated asset in 1981, and about 5 per cent higher in 1982. Similarly, from 1986 to 1987 the forward premium on the yen was very close to zero and the yen strongly appreciated from 200 to 130 yen per dollar.

All these observations suggest that *uncovered* interest parity did not hold in the period of sharp yen depreciation (appreciation) and that deviations of the forward rate from the *ex-post* realized spot rate were large, predictable and volatile.

Our analysis shows that the risk premium in the yen/dollar market is significantly different

from zero, has a mean which changes over subperiods and is highly volatile. The series shows nonlinearities, but we find little evidence for time variation and for structural changes in the coefficients of the model. Hence, apart from affecting the volatility of the series, the events of the 1980s had a negligible impact on the time-series properties of the risk premium in this market.

The rest of the paper is organized as follows: the next section presents the data set and the VAR model used to construct the measure of expectations and the risk premium. Section 3 describes the features of our measure of expectations, compares it to the expectations contained in survey data, and outlines some of the empirical features of the risk premium series. Section 4 tests several hypotheses concerning the risk premium series and compares the results with the ones existing in the literature. Concluding remarks are presented in section 5.

2. THE DATA, THE MODEL AND THE FORECASTING PROCEDURE

The VAR model considered in this paper includes five variables: the (log of) spot exchange rate, two 3-month interest rates and two (log of) stock price indices. All data are weekly samplings (at Wednesday) of daily figures from 1979 to 1987. The exchange rates are measured in yen per dollar at the closing level of the New York market. The stock price index in the United States is the Standard & Poor 500 and the Japanese one is Nikkei 225, that is a weighted average of 225 stocks traded in the Tokyo market.[1] For the short-term dollar denominated interest rate, the offshore (Eurodollar) 3-month rate is used, and for yen denominated interest rate the Gensaki rate is used (see Ito, 1986, for reasons for using the Gensaki rate.)

Many structural models of international finance have identified these financial variables as important ingredients in explaining the dynamic features of the spot rate, although researchers do not agree on the causal relationship among variables. Since a VAR model treats all variables as endogenous, it avoids biases due to *ad-hoc* assumptions or restrictive specifications. Also, since the major purpose for using a VAR model is to derive a measure for the expected spot rate, the lack of identifying restrictions does not pose any problem.

Only financial variables which are observed at a weekly (or daily) frequency are included in the system. Variables such as GNP, inflation or trade deficits are not included because they are not observed frequently enough, and financial variables quickly respond to announcements of those variables anyway. Including the logarithm of the spot rate, together with the level of the two interest rates, contains information of the forward rate, since the (log of) forward rate is the (log) spot rate plus interest rate differential.[2] Logarithms of the stock indices are also computed to ensure that all the variables in the system have comparable units.

The forecasting model we use takes the form:

$$X_t = A_t(l) \times X_{t-1} + c + e_t \qquad e_t \sim (0, M) \tag{1}$$

$$b_t = b_{t-1} + u_t \qquad u_t \sim (0, N_t) \tag{2}$$

where $X_t = [X_{1,t}, X_{2,t}, X_{3,t}, X_{4,t}, X_{5,t}]'$ is a 5×1 vector of endogenous variables; $A_t(l)$ is a polynomial matrix in the lag operator of dimension 5×5 for each $l = 1, ..., P$, where P indicates the lag length of the system; c is a 5×1 vector of intercepts; u_t and e_t are innovations in the coefficients and the variables of the model and b_t, the stacked version of the matrix

---

[1] We also looked at the New York Stock Exchange composite (NYSE) index and at Tosho, the Tokyo Stock Exchange composite index as alternatives, but the empirical results were not affected by the choice of variables.

[2] In other words, adding the forward rate to our system would cause a severe multicollinearity.

$A_t(l)$, is of dimension $25P \times 1$, and comes from the AR(1) companion form of (1)

$$Z_t = Z_{t-1}^l b_t + w_t \qquad (3)$$

The system (1)–(2) is in a standard state space form and (2) represents the law of motion of the coefficients.

To make the forecasting model operative we need to specify b_0, N_0 and the link between N_t and N_0. For b_0, we assume a prior mean of one on the first own lagged coefficient in each equation and of zero on all other coefficients. This choice for the prior mean comes from observing that standard unit root tests (Dickey and Fuller, 1981) on the variables of the system do not reject the nonstationary null hypothesis in any case. We specify $N_0 = \mathrm{diag}\,[n_0(i, l, j),$ $i = 1, ..., 5;\ j = 1, ..., 5;\ l = 1, ..., P)]$ where the diagonal elements of the matrix have the form

$$n_0(i, l, j) = h(l) f(i, j) \sigma(i) / \sigma(j) \qquad (4)$$

$h(l)$ regulates the concentration of the prior distribution around the median as the lag length increases, $f(i, j)$ is a matrix of weights describing the importance of variable j in equation i and $\sigma(i)$ and $\sigma(j)$ are the standard errors of the innovations in variables i and j. In Appendix A of Canova and Ito (1987) we showed, using four different criteria, that the best forecasting model for the yen/dollar spot rate which employs these five variables requires lags 1 through 4 plus the eighth. In the framework of this paper we use this information by selecting $L = 8$ and specifying the functional form for $h(l)$ to be $h(l) = 1/l$, $l = 1, ..., 8$. Also, since the variables of the system seem to be explained primarily by their own lags, we choose the matrix $f(i, j)$ to be symmetric and to have the off-diagonal elements equal to $0 \cdot 05$. Here, $\sigma(i)$ and $\sigma(j)$ are estimated from univariate autoregressions of each of the five variables on eight lags. Finally, we assume that $N_t = \gamma N_0$, where γ is a parameter controlling how much of the uncertainty in X_t is due to innovations in the coefficients. After some experimentation, we chose $\gamma = 0 \cdot 1 \times 10^{-6}$.

Given these choices for b_0 and N_0, a recursive least-square ('rolling regression') procedure is used to update the coefficients of the model and to compute forecasts conditional on the amount of information available at that point in time. The procedure is repeated using data from 1979:1 up to 1987:38 and forecasts 13 periods (3 months) ahead are recursively computed for the period 1980:39–1987:51.

Two important facts need to be noted. First, contrary to standard procedures which use rational expectations as a maintained hypothesis, we explicitly compute expectations from the model. Standard procedures would not yield any meaningful time-series for the risk premium, since forecast errors sneak into the risk premium series. Second, our procedure implies that rational agents form their forecasts by taking recursive linear projections on the available information set. In other words, we approximate conditional expectations with time varying linear projections.[3] If agents actually used nonlinear schemes in forecasting, the forecast errors of the VAR are bound to be heteroskedastic and possibly time-varying. As a diagnostic for our approximation procedure we present tests for heteroskedasticity in the forecast error of the VAR (Table I) and plot the point estimates of the risk premium series with a one-standard error band (Figure 2).

Efficiency in the yen/dollar market implies that the risk premium is the difference between the forward rate and the expected future spot rate. Since the VAR yields a k-step ahead

---

[3] There are other approximating procedures available which do not make this approximation error. See, e.g., Diebold and Nason (1990), Gallant, Tauchen, and Hseh (1988).

forecast of the spot rate, the risk premium, measured as deviation from uncovered interest parity, is easily constructed from the model.[4]

Let $S_t = X_{1,t}$ be the log of the yen/dollar spot exchange rate; $E_t S_{t+k}$ be the expected value at t of the log of the spot rate at $t+k$, $RJA_{t,t+k} = X_{2,t}$ and $RUS_{t,t+k} = X_{3,t}$ be the level of the k-period Gensaki and Eurodollar interest rates; and let $F_{t,t+k}$ be the (log of) the forward rate quoted at t for transaction at $t+k$. The model implies that

$$E[S_{t+k} \mid H_x(t)] = \delta E[Z_{t+k} \mid H_z(t)] = \delta[\Pi_{i=1}^{k} Z_t' \hat{B}_{t+i-1}]$$

where $\hat{B}_{t+i-1} \equiv E_t B_{t+i-1}$ and $\delta = [1, 0, 0, ..., 0]$. Therefore, $Y_t = E_t S_{t+k} - F_{t,t+k}$ is the constructed series for the risk premium, where $H_x(t) = H_z(t)$ represents the information set available at t.

Since covered interest parity (CIP) implies that the *forward premium* equals the interest rate differential

$$FP_{t,k} \equiv F_{t,t+k} - S_t = RJA_{t,t+k} - RUS_{t,t+k}. \tag{5}$$

and since the *expected change* in the spot exchange rate is defined as

$$EX_{t,k} \equiv E_t S_{t+k} - S_t \tag{6}$$

while the forecast error is given by

$$FE_{t,k} \equiv S_{t+k} - E_t S_{t+k} \tag{7}$$

the *risk premium* can also be computed as

$$RP_{t,k} \equiv E_t S_{t+k} - F_{t,t+k}$$
$$= EX_{t,k} - FP_{t,k}, \tag{8}$$

where, in (8) the risk premium is decomposed into the sum of the expected changes in the spot rate $EX_{t,k}$ and the forward premium $FP_{t,k}$.

As is well known the (*ex-post*) forward bias, $FB_{t,k}$, i.e. the discrepancy between the forward rate and the *ex-post* spot rate, can be decomposed into a risk premium and a forecast error

$$FB_{t,k} = S_{t,k} - F_{t,t+k} = (S_{t+k} - E_t S_{t+k}) + (E_t S_{t+k} - F_{t,t+k})$$
$$= FE_{t,k} - RP_{t,k}.$$

or

$$RP_{t,k} = FE_{t,k} - FB_{t,k} \tag{9}$$

If uncovered interest parity (UIP) holds, the interest differential equals the expected change in the spot rate: $E_t S_{t+k} - S_t = RJA_{t,t+k} - RUS_{t,t+k}$. Hence, $FP_{t,k} = EX_{t,k}$, and $RP_{t,k} = 0$. If we assume rational expectations, $S_{t+k} = E_t S_{t+k} + e_t$, where e_t is white noise. Hence, according to the conventional wisdom, the risk premium is the *ex-post* forward bias contaminated by a white noise. Conventional analysis, however, fails to give insights on risk premium behaviour, either if rational expectations fail to hold, or if forecast errors have large variances. Indeed, the recent literature has emphasized these situations.

---

[4] In a different approach to a similar issue, Ito (1988) tested uncovered interest parity as a cross-equation restriction constraint on a VAR model.

3. OVERVIEW OF THE RISK PREMIUM TIME-SERIES

Plots of the forward, spot and expected spot rates are presented in Figure 1. We concentrate on the behaviour of these three variables over three periods. In 1981, and again in 1984, the forward rate was persistently below the realized spot rate, while our measure of expected spot rate was in between the two. The first period of divergence, 1981:14–1982:1, agrees with findings of Ito (1984) using monthly data, and of Frankel and Froot (1987) using survey data. A plausible explanation of the divergence is that the lifting of capital controls in Japan, which occurred at the end of 1980, affected the behaviour of Japanese investors, so that the forward rate was a bad predictor for the expected spot rate (see Ito, 1986). The divergence of 1984, though smaller in size, is still relevant. The forward rate was consistently below the realized spot rate by about 10 per cent. The biasedness of the forward rate in this period might be due to a series of adjustments occurred in the European Monetary System in 1984, which caused uncertainty and instabilities in many exchange markets including the yen/dollar market.

Finally, from mid-1985 to mid-1986, the forward rate was constantly above the realized spot rate with our measure of the expected spot rate being almost the same as the forward rate. Therefore, for a year following the Plaza agreement of October 1985, forecast errors, not the risk premium, seems to explain most of the forward bias.

The VAR forecasts for the expected spot rate may not be particularly accurate. In Figure 2 we plot the expected spot rate together with a one standard error band. The band is significantly large and seems to increase as more data points are added to compute forecasts. Three important elements of uncertainty can contribute to the forecast error: a specification error, an error due to time variation in the coefficients, and an error due to the innovations
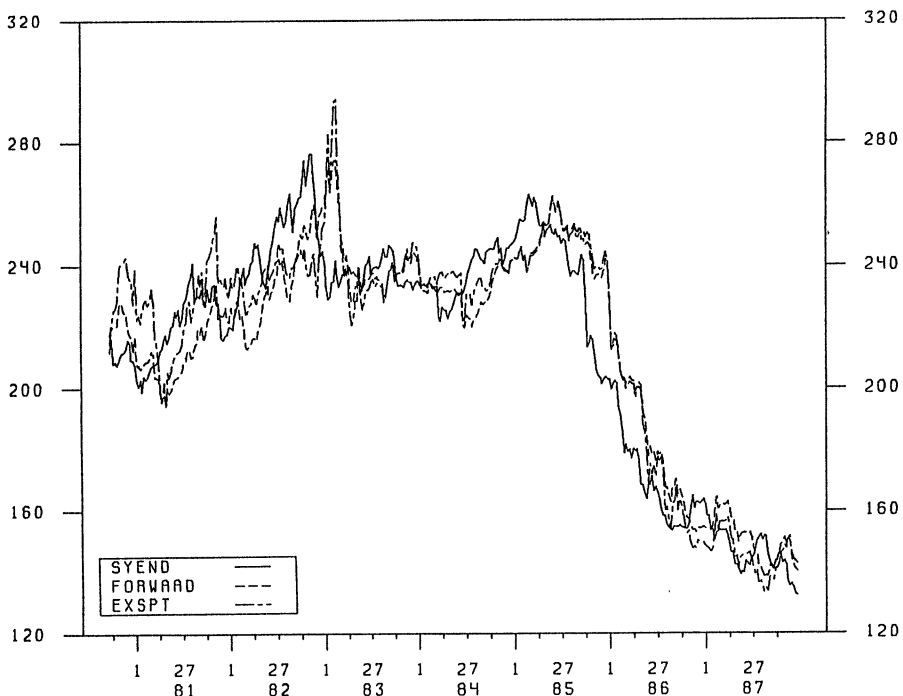

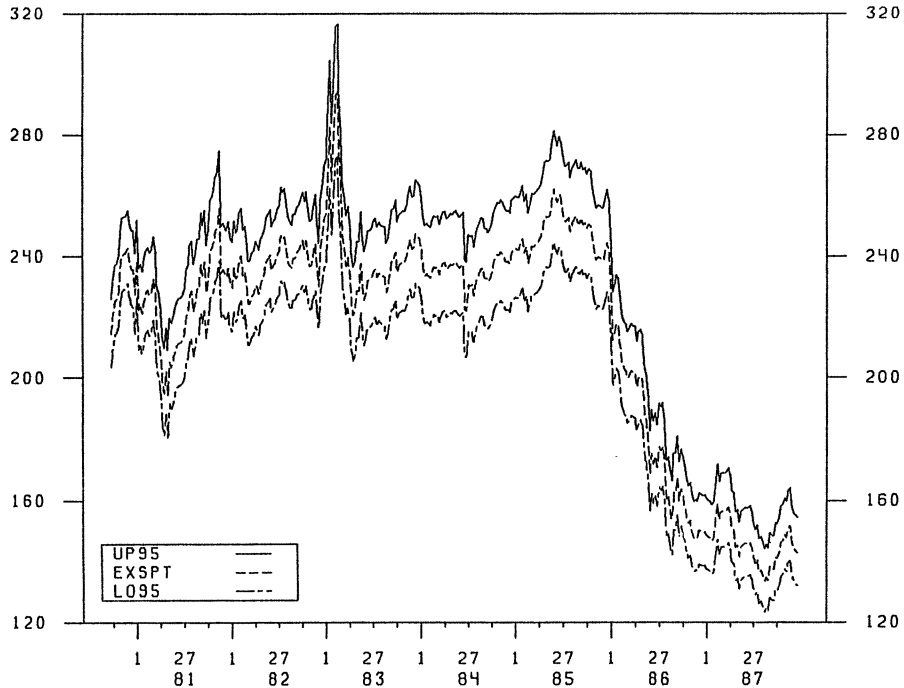Figure 1. Spot, forward, and expected spot

Figure 2. Point forecasts and 95 per cent band

in the variables of the system. Table I indicates that the specification error due to the linear approximation of a possibly nonlinear function is small, and it is not crucial in determining the size of the band around the point forecast of the spot rate. Canova (1991) finds that the error due to time variation in the coefficients of this model is relatively small and can be neglected as a first approximation. Therefore, the large size of the band seems essentially due to innovations in the variables of the system.

Frankel and Froot (1987) found that, for the time span 1981–1985, New York traders consistently expected an appreciation of the yen (from about 15 per cent in 1981 down to 6 per cent in the late 1985). According to their survey data, investors were willing to sacrifice

Table I. Heteroskedasticity tests on forecasts errors, significance levels

ARCH (8)	BP (8)	White (16)
0·987	0·192	0·297

Note: ARCH (8) refers to a test where the prewhitened squared forecasted errors are regressed on eight lags. Prewhitening is done regressing the forecast errors on eight lags. BP (8) is a Breush–Pagan test where the prewhitened squared forecast errors are regressed on four lags of the forecast errors and four lags of the expected spot rate. White (16) is a White test where the regressors are four lagged forecast errors, four lagged squared forecast error and four lagged expected spot rates and four lagged square expected spot rates.

Figure 3. Risk, forward premium, and expected change in spot

higher effective returns on the yen in order to hold dollars. This behaviour generates a risk premium on the dollar both in appreciation and depreciation phases, and implies that most of the movements in the risk premium are induced by movements in the forward premium.

In Figure 3 we plot the forward premium and the point estimates of the risk premium and of the expected change in the spot rate in annualized percentages. Figure 3 indicates that the forward discount was relatively stable, ranging from $11 \cdot 32$ to $-3 \cdot 64$ per cent over the sample, and that movements in the risk premia are closely related to movements in the expected change in the spot rate. The contemporaneous correlation between these variables is above $0 \cdot 9$ in every sample we considered. Also note that, according to our measure of expectation, and contrary to the survey data of Frankel and Froot, the yen was expected to depreciate up to 1982.

It is interesting to compare the recursive VAR forecast errors with the survey forecast errors. Frankel and Froot report that, for a 13-week horizon for the (weekly) sample period from 1981:6 to 1985:12, the survey data collected by the *Economist* indicate an expected average depreciation of the dollar of 12.66 per cent per year. According to the VAR-generated forecasts, the expected depreciation of the dollar for that period was only $2 \cdot 33$ per cent per year on average (with a standard error of $15 \cdot 7$) that is much closer to the actual depreciation of $4 \cdot 37$.[5] Our results therefore indicate that if agents had used mechanical methods to generate

---

[5] Survey data that Frankel and Froot use contain only the median of the respondents, and a measure of dispersion among respondents is not available. This prevents a more extensive comparison between the two procedures. Also, given the way the forecasting model is chosen, our estimates are the best possible under the mean square error criteria.

forecasts of future variables, they could have improved their predictions. In this sense survey data do not seem to produce a reliable risk premium series.

Figure 3 also confirms characteristics found by Fama (1984) and Hodrick and Srivastava (1986). First, the risk premium is negatively correlated with the expected appreciation of the yen (our measure of risk premium is the negative of theirs). Second, the variance of the risk premium series (206·49) is larger than the variance of the expected change in the spot rate (175·91) and the covariance between the risk premium and the expected depreciation of the yen (186·52) is larger than the covariance of the forward premium with the realized change in the spot rate (33·55).

Next, we discuss the statistical features of the risk premium series which are summarized in Tables II and III. Figure 3 indicates that the unconditional distribution of the risk premium

Table II. Statistics for the risk premium

	Sample			
	1981–1987	1981–1982	1982–1985	1985–1987
Mean	2·92	19·18	0·58	−8·74
Standard deviation	14·36	10·87	9·09	8·31
t-stat mean = 0	4·00	18·16	0·79	−11·27
Skewness test	0·29	0·64	0·03	0·84
Kurtosis test	0·00	0·00	0·00	0·00
Kiefer–Salmon test	0·00	0·03	0·00	0·01
Autocorrelation function				
Lag 1	0·96	0·86	0·92	0·91
Lag 4	0·86	0·51	0·61	0·67
Lag 8	0·70	0·16	0·12	0·54
Lag 13	0·58	−0·09	0·00	0·41
Lag 18	0·53	−0·29	0·19	0·18
Lag 26	0·48	−0·00	0·17	−0·00
Total variance	205·95	117·07	82·11	68·50
Cross-correlation risk premium/ expected change in spot rate				
Lead 13	0·48	−0·16	−0·05	0·46
Lead 8	0·61	0·07	0·09	0·58
Lead 4	0·80	0·43	0·58	0·69
Lead 1	0·93	0·80	0·90	0·91
Lag 0	0·97	0·93	0·98	0·99
Lag 1	0·95	0·85	0·91	0·90
Lag 4	0·85	0·59	0·62	0·66
Lag 8	0·69	0·29	0·13	0·52
Lag 13	0·56	−0·00	−0·01	0·37
Cross-correlation risk premium/ forward premium				
Lead 13	0·66	0·22	0·38	−0·53
Lead 8	0·63	0·23	0·22	−0·43
Lead 4	0·55	0·13	0·12	−0·25
Lead 1	0·47	0·00	0·03	−0·10
Lead 0	0·45	−0·04	−0·00	−0·05
Lag 1	0·41	−0·14	−0·04	−0·02
Lag 4	0·33	−0·34	−0·07	0·05
Lag 8	0·29	−0·39	−0·01	0·19
Lag 13	0·32	−0·22	0·10	0·33

Table III. Diagnostic tests for nonlinearities in the risk premium, P-values

	1981:1–1987:51	1981:14–1982:40	1982:40–1985:40	1985:41–1987:51
F test all lags $= 0$	0·0003	0·16	0·00	0·003
F test all coefficients $= 0$	0·00	0·0001	0·00	0·00
χ^2 (13)	0·00	0·52	0·00	0·00
Brock–Deckert	4·35	1·54	2·67	2·39

may have changed over time. The volatility of risk premium shows a declining trend in the first 2 years of sample, presents more persistent cyclical movements until 1985 and is again volatile around a negative mean in 1986 and 1987. The autocovariance function (ACF) for the whole sample decays slowly. For the subsample 1982:41–1985:40 the ACF shows cycles of about 18 weeks. For the sample 1985:41–1987:51 the ACF is still positive after 25 lags, in contrast with 10 lags of the first subsample. The overall sample means of the process is 2·96 which, at an average 216 yen per dollar, corresponds to an average risk premium on the yen of about 0·75 per cent per quarter. For the three subsamples the means are 19·18, 0·58 and −8·74 respectively, with the mean for the 1982–1985 period being insignificantly different from zero at the 5 per cent significance level. Therefore, there was a quarterly risk premium on the yen of about 5 per cent quarter in the first subsample and an average quarterly risk premium on the dollar of about 2 per cent in the last subsample. The standard deviation of the series is 14·36, so that a 95 per cent confidence band around the overall mean for a Gaussian process is (−25·18; 31·10), which is approximately the band of oscillation of the series. Note also that the standard deviation declines in the last two subsamples, indicating a substantial reduction in uncertainty surrounding the risk premium. However, because the standard errors are large, differences in the means across subsamples are statistically irrelevant, even though they appear to be economically significant.

In order to check for the presence of fat tails in the unconditional distribution, we examine the size of the kurtosis of the historical distribution of the risk premium in various samples and we test whether the estimated kurtosis differs from the one of a Gaussian distribution. The test rejects the hypothesis that the distribution is normal in all samples. An alternative test based on Kiefer and Salmon (1983) also rejects the null hypothesis of normality of the empirical distribution of the risk premium.

A similar test for the skewness of the process indicates the existence of a positive and significant skewness only in the 1982–1985 sample. This result could be interpreted as supporting Fama's (1984) idea that the uncertainty regarding the direction of US Government policies in 1982–1983 may have induced a 'peso problem' in the market.

The existence of three distinct episodes, in which tranquillity and turbulence alternate, suggests the presence of conditional heteroskedasticity in the series. To check this, we first compute a diagnostic for ARCH residuals in the series by regressing the squared residuals of a AR(8) regression of the risk premium series on a constant and on 13 lags. An F-test for the null hypothesis of zero lagged coefficients and a LM test for the insignificance of the regression R^2 is strongly rejected for the whole sample and for the last two subsamples (see Table III). Finally, the Brock and Deckert (1988) test for nonlinearities rejects the null hypothesis of white noise in the normalized residuals of this model.

In sum, the conditional distribution of the risk premium series appears to be non-stationary.

The first two moments of the unconditional distribution may also be changing over time. In addition, the tails of the unconditional distribution are fatter than the ones of a Gaussian distribution. Since fat tails and nonstationary behaviour may be connected, the next section considers a model which can generate these features in the estimated time-series.

4. TESTING THE PROPERTIES OF THE RISK PREMIUM SERIES

In this section we test several hypotheses concerning the properties of the risk premium series. Hodrick and Srivastava (1984), Domowitz and Hakkio (1985), Kaminsky and Peruga (1989), Engel and Rodrigues (1989), among others, have tested some of these properties using *ex-post* measures of the risk premium and alternative econometric techniques.

The existence of nonstationarities and fat tails in the generated risk premium series poses some problems in estimation. A common way to proceed is to use a (quasi-)differencing filter to induce stationarity in the data and estimate the constructed series using a version of an ARCH-M model (see Engle, Lilien, and Robbins, 1987; Domowitz and Hakkio, 1985). Although ARCH models have often proved to be useful instruments in estimating time-series with some form of heteroskedasticity, we tackle the estimation problem from a different point of view in this paper.

Our approach is Bayesian in spirit. We specify a time-varying coefficient (TVC) model for the risk premium. The model retains linearity of coefficients and variables in the model structure, but has the potential to generate nonstationarities and heteroskedasticity in the estimated time-series by means of time variation in the coefficients. The approach is flexible and can be adapted to a rich class of situations without requiring data transformations and/or the introduction of complex nonlinearities in the model structure (as would be the case with the ARCH-M model). Canova (1990) shows how generic TVC-AR models with a rich enough prior parameterization are able to induce general patterns of conditional heteroskedasticity and time variation in the unconditional moment structure of the model.

Let Y_t be the risk premium series. The model we propose is the following:

$$Y_t = B_t' W_t + \varepsilon_t \qquad \varepsilon_t \sim (0, \sigma^2) \tag{10}$$

$$B_t - B_0 = G(B_{t-1} - B_0) + v_t \qquad v_t \sim (0, \Omega_t) \tag{11}$$

where W_t is a $P + 1$ vector including P lagged Y_t values and a 1, B_0 is the unconditional mean of B_t, $B_t = [B_t^*, c_t]$, c_t is the intercept of the model, G is a $(P + 1) \times (P + 1)$ symmetric matrix, v_t and ε_t are innovation processes which are uncorrelated at all leads and lags. Equation (10) describes the law of motion of the coefficients of the model.

In a standard Bayesian approach one first provides a prior probability distribution for the parameters and then integrates to find the posterior of the parameters. In our approach the prior has fixed parameters (as in Doan, Litterman, and Sims, 1984) and we choose their values so as to numerically maximize the likelihood function of the data (see Engle and Watson, 1987, for this approach). Under the assumption that the prior of the parameters is flat in a given hypercube, our procedure chooses the vector of parameters which comes closest in producing the mode of the posterior distribution of the parameters. In this context the numerical search for the peak of the likelihood can be interpreted as an approximate numerical integration in the region around the mode of the posterior distribution.

Without restrictions the number of the coefficients in (10) is large. To decrease the dimensionality of the model we link the free coefficients in B_0, G and Ω_t to a set of parameters

which fully describe our prior. We assume:

$$G = \lambda_0 \times I$$

$$B_0 = \lambda_1 \qquad \text{if } l = 1 \tag{12}$$

$$= 0 \text{ otherwise} \tag{13}$$

$$\Sigma_0 = \text{diag}\{\sigma_{011}\} \tag{14}$$

$$\sigma_{011} = (\lambda_2/l)^{\lambda 3} \qquad \text{if } l = 1, 2, \ldots, P$$

$$= \text{varc} \qquad \text{if } l = P + 1 \tag{15}$$

$$\Omega_0 = \Sigma_0 - \Sigma_0 S[\sigma_f^2 I - S\Sigma_0 S']^{-1} S' \Sigma_0 \tag{16}$$

$$S = \lambda_4 \times [1, 1, 1, \ldots 1] \tag{17}$$

$$\Omega_t = \lambda_6 \times \lambda_5 \times \Omega_0 \tag{18}$$

A brief discussion to motivate our choices for the prior is in order. The log spectrum of the risk premium series is very flat except for a mild peak around frequency 0. Therefore, a low-order polynomial may suffice to generate a transfer function with the required properties. An AR(1) model with the first coefficient substantially smaller than 1 could be the most reasonable choice in this case. We embed this knowledge in our prior by choosing the mean of the process to be λ_1 on the first lag, and zero otherwise where λ_1 is an unknown parameter. Also, we let the coefficients of the model decay towards the mean at the rate $(I - G)$. The prior covariance matrix of the coefficients is Σ_0. Since no information about the correlations across coefficients is available, we assume that Σ_0 is diagonal, with λ_2 representing the tightness of the variance on the coefficients and λ_3 is a decay parameter. λ_2 and λ_3 determine the concentration of the prior variance of the coefficients around the prior median as l increases. Also, depending on the choice of λ_3, past information is discounted at faster or slower rate. The variance of the intercept is separately parameterized. In (16) we scale down this prior covariance matrix to take into account the uncertainty regarding the correct prior specification for the mean. To do this we assume that a linear combination of the coefficients is arbitrarily close to 0 with λ_4 controlling the size of the variance (σ_f^2) of the restriction. The parameter λ_5 controls the amount of time variation injected in the conditional variance of the coefficients at each date. For $\lambda_5 = 0$, the coefficients are constant over time. Finally, λ_6 is a dummy parameter controlling for possible structural breaks in the coefficients.

Since the model is in a state space form, estimates of the coefficients of the model can be obtained recursively with the Kalman filter algorithm. However, since the Kalman filter requires the λ values to be known, we run the algorithm over a coarse grid of parameters and compare the fit of the model for each set of parameters using the likelihood statistics. Doan, Litterman, and Sims (1984) showed that the likelihood statistics can be written as

$$L = (T/2) \times \log\left[1/T \times \sum_t (\varepsilon_t)^2 / (v_t/v_t^*)\right] \tag{19}$$

where $v_t = \sigma^2(1 + Y_{t-1}\eta_t Y_{t-1})$; $\eta = G'\Omega_{t-1}G + \Omega_t$; v_t^* is the geometric mean of v_t, and where $\varepsilon_t = Y_t - \hat{B}_t' W_t$. Since the likelihood statistic is a weighted average of one step-ahead recursive forecast errors, it naturally measures the forecasting performance of the model for a given set of parameters. Therefore, a choice of parameters which do not agree with the data will produce large forecast errors and a low value for the likelihood.

This framework of analysis also provides a way of testing several interesting hypotheses. A

null hypothesis of the absence of a risk premium, for example, would amount to the restriction that the prior mean of the coefficients and of the intercept is zero, and their prior variance is highly concentrated around the median. A null hypothesis of the constancy of the risk premium, on the other hand, requires the same type of restrictions on the coefficients of the model only (the variance of the constant is unrestricted).

Two other interesting hypotheses concern the existence of a slow but continuous time variation and/or the existence of structural breaks in the coefficients of the model. These two hypotheses are important in the context of the risk premium literature because they allow us to gauge how the risk premium is linked to variables that change slowly over time (productivity, population, etc.) and how sensitive it is to those macroeconomic events (Fed changing operating procedures, the Group of Five deciding to drive down the value of the dollar, etc.) that produced substantial changes in financial markets.

We chose to evaluate the likelihood function over a 10-value grid for each of the parameters for a total of 8^{10} combinations. The range for λ_0, λ_1, λ_4 is [0,1], for λ_2 is [0·05,2], for λ_3 is [0,2], for λ_5 is [0,0·1 × 10^{-4}] for λ_6 is [0·1,10] and for varc is [0·005,100]. Equidistant points in each range are chosen. The four hypotheses can be cast in terms of grid choices in the following way: the null hypothesis of the absence of a risk premium restricts $\lambda_1 = 0$, $\lambda_2 = 0·05$, varc = 0·005; the null hypothesis of constant risk premium restricts $\lambda_1 = 0$, and $\lambda_2 = 0·05$; the null hypothesis of no time variation in the coefficients restricts $\lambda_5 = 0$ and the null hypothesis of no structural changes restricts $\lambda_6 = 1$. In all these cases the remaining parameters are left unrestricted.

Tests of these null hypotheses can be undertaken by comparing the unrestricted model with each of these restricted specifications using the Schwarz criterion.[6] Since all specifications are nested, a comparison of the likelihoods across specifications determines the contribution of certain parameters to the forecasting performance of the model. Additional insights can be obtained by comparing the estimated moving-average representation for the risk premium generated by each specification. We report the 'best' parameter vector in Table IV while the likelihood values, the value of the Schwarz criterion and the results of the hypotheses testing are presented in Table V.

Table IV. Optimal estimates of the parameters

	Sample			
	1981–1987	1981–1982	1982–1985	1985–1987
λ_0	0·92	0·58	0·69	0·73
λ_1	0·00	0·00	0·00	0·00
λ_2	0·05	0·05	0·1	0·05
λ_3	0·00	1·20	0·50	0·00
λ_4	1·00	1·20	1·00	1·00
λ_5	0·1 × 10^{-6}	0·1 × 10^{-5}	0·1 × 10^{-6}	0·5 × 10^{-7}
Var. const.	20	20	20	20
λ_6	Break (1)	1·0	1·0	
	Break (2)	1·00003	1·0002	

---

[6] The Schwarz criterion selects specification 1 if $\log(L_1 - L_2) \geq 2(p_1 - p_2) \times T$ where p_i is the number of parameters in model i, T is the number of observations and L_i is the likelihood for model i.

Table V. Likelihood values for the various models

Model	Dates	Likelihood value	Schwarz criterion
Unrestricted		399·10	12011·46
No risk premium		140·23	2840·35
Risk premium constant		399·10	12016·46
No time variation		398·72	12026·46
Breaks (1): Jump in Time Variation	1982:38–1982:41 1985:38–1985:41	399·08	12009·86
Breaks (2): segmentation in Time Variation	1982:40–1985:40 1985:41–1987:51	399·17	11866·85
Breaks (3): segmentation in all the parameters	1982:40–1985:40 1985:41–1987:51	387·76	10987·76

Several features deserve comment. First, a prior mean of zero on all the coefficients is the one that best agrees with the data. Second, the 'optimal' value of λ_2 is the lowest of the grid, indicating that the prior variance should be highly concentrated around the median. Third, the optimal value of the variance of the intercept is large. Given these observations, a test of the absence of a risk premium is rejected, but a test of its constancy is not rejected. Fourth, since the estimated value for λ_5 is small, time variation in the coefficients is small and accounts for no more than 0·1 per cent of the standard error of the series. When we test for the contribution of λ_5 to the forecasting ability of the model we find that it is minimal, and that the restricted and unrestricted specifications do not differ significantly.

Given the changes in the unconditional mean and variance of the series over subsamples noted in the previous section, we suspect the existence of two structural breaks in the series: one around 1982:40 (when the Fed abandoned the money supply targeting) and another around 1985:40 (when the Group of Five agreed to pressure down the value of the dollar).

Regime changes can be modelled in the framework of this section in several ways. If we believe that a regime shift temporarily displaces the law of motion for the coefficients one could let λ_6 be different from a unity only for a few periods around 1982:40 and 1985:40. Alternatively, a regime shift can permanently change the law of motion of the variance of the coefficients of the model. In this case, λ_6 should be set to be one for the period of 1980:39–1982:40, while it is freely estimated in each of the other two subsamples. Finally, a regime shift can cause the entire specification of the model change across subsamples. In this case all the parameters of the model should be allowed to be separately estimated in each sample. Table IV reports the optimal estimated values of λ_6 for the first two specifications and the estimated values of $\lambda_0, \lambda_1, \ldots, \lambda_5$ for the three separate subsamples.

The results indicate that in the first two specifications the optimal value for λ_6 is close to, or exactly equal to, one, while the values of the λ vector estimated in the three subsamples are very close to the estimates obtained using the entire sample. As far as forecasting performance is concerned, the gains obtained when the parameters are recomputed after 1982:40 and after 1985:40 are negligible.

Figure 4 plots the estimated MA representation for the risk premium under four different specifications. In the model estimated using the entire sample and $\lambda_6 = 1$, which appears in the upper panel of Figure 4, a unit innovation in the risk premium creates oscillatory response up to 52 weeks, with cycles of about 13 weeks. This recurrent and strong cyclical behaviour indicates the possible presence of unstable elements and some kind of seasonalities in the sample under consideration. The estimated MA representation with no time variation (third
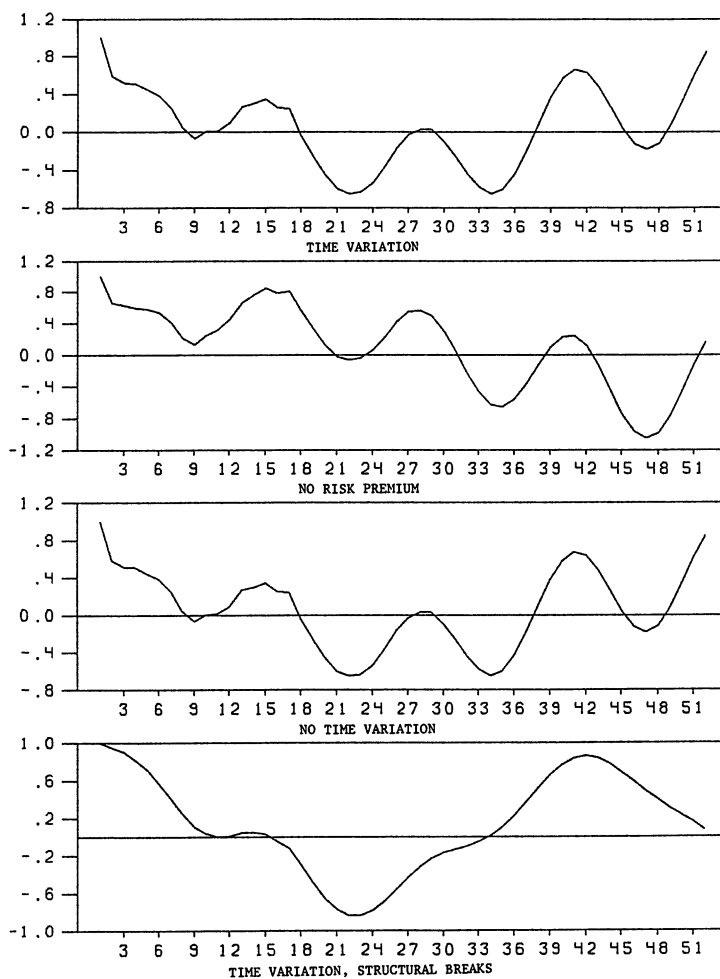
Figure 4. MA representation

panel) is identical in shape and in size to the one in the first panel, while the one corresponding to a model with no risk premium (second panel) has a much stronger 26-week cycle and a declining trend. Finally, by comparing the top panel with the bottom panel, which reports the estimated MA representation for a model where all the parameters are re-estimated after 1982:40 and after 1985:40, one can see that, despite the negligible differences in terms of forecasting performance, the two specifications produce substantial differences in the risk premium dynamics.

In sum, the results show that over the period 1980:39–1987:51, the risk premium series was approximately constant and its unconditional distribution was stable in the sense that the coefficients of the model exhibited neither smooth variations nor structural breaks at certain specific dates. While these results may be due to the large standard error surrounding our measure of risk premium, they provide insight on the properties and on the behaviour of risk premium in the 1980s which constrasts the evidence obtained using survey data.

Finally, we compare our findings with the ones existing in the literature. Our results confirm

those of Hodrick and Srivastava (1984) in detecting the presence of heteroskedasticity, but contrasts with their finding of significant time variation in the risk premium series. Also, the high correlation between the risk premium and the expected spot rate change suggests that a VAR-generated forecast is a more reliable predictor than the forward premium in forecasting the risk premium series. As compared with Fama (1984) and Hodrick and Srivastava (1986) our estimates suggest a much lower β-coefficient for a regression of the realized changes in the spot rate on the forward premium and smaller estimates of the differences between the variances of the forward premium and of the expected changes in the spot rate. Finally, as in Domowitz and Hakkio (1985) the null hypothesis of no risk premium is rejected and, in addition, we demonstrate the presence of substantial volatility in the risk premium series for the 1979–1987 period.

5. CONCLUSION

In this paper a VAR model was employed to construct a measure of the conditional expectations of the future yen/dollar spot rate. This measure allowed us to examine the dynamics of an *ex-ante* time series of risk premium in the market.

The VAR model produced 'better' forecasts than the survey responses for turbulent periods such as 1981–1982 and 1984–1985. The VAR-generated expectation was then used to construct a risk premium time-series. This risk premium series seems to be more reliable than the ones obtained either using survey data or the *ex-post* realized spot exchange rate.

Tests on the risk premium series suggest that a risk premium was present, but that it was virtually constant throughout the sample. The conditional variance of the risk premium changed over time, but its unconditional distribution seems stable across subsamples. Despite these features, the volatility of the series was substantial and varied considerably throughout the sample.

REFERENCES

Brock, W., and W. Deckert (1988), 'A general class of specification tests: the scalar case' (mimeo).
Canova, F. (1990), 'Are exchange rates really random walks?', Brown University, Working Paper 90-12.
Canova, F. (1991), 'An empirical analysis of ex-ante profits from forward speculation in foreign exchange markets', *Review of Economics and Statistics*, (forthcoming).
Canova, F., and T. Ito (1987), 'On the time varying risk premium in the yen/dollar exchange market', University of Minnesota, Working Paper No. 244.
Cosset, J. (1984), 'On the presence of risk premiums in foreign exchange markets', *Journal of International Economics*, **16**, 139–154.
Cumby, R. (1988), 'Is it risk? Explaining deviations from uncovered interest parity', *Journal of Monetary Economics*, **22**, 279–300.
Cumby, R., and M. Obstfeld (1983), 'International interest rate and price level linkages under flexible exchange rate: a review of recent evidence', in J. F. O. Bilson and R. C. Marston (eds), *Exchange Rates: Theory and Practice*, University of Chicago Press, Chicago, pp. 121–151.

Dickey, D., and W. Fuller (1981), 'Likelihood ratio statistics for autoregressive time series with a unit root', *Econometrica*, **49**, 1057–1072.

Diebold, F., and J. Nason (1990), 'Nonparametric exchange rate prediction', *Journal of International Economics*, **28**, 315–322.

Doan, T., R. Litterman, and C. Sims (1984), 'Forecasting and conditional projections using realistic prior distributions', *Econometric Reviews*, **3**, 1–100.

Dominguez, C. (1986), 'Are foreign exchange forecasts rational: new evidence from survey data', *Economic Letters*, **22**, 267–271.

Domowitz, I., and C. Hakkio (1985), 'Conditional variance and the risk premium in the foreign exchange market', *Journal of International Economics*, **19**, 47–66.

Engel, C., and T. Rodrigues (1989), 'Tests of internation CAPM with time varying covariances', *Journal of Applied Econometrics*, **4**, 119–138.

Engle, R., D. Lilien, and R. Robins (1987), 'Estimating time varying risk premium in the term structure: the Arch-M model', *Econometrica*, **55**, 391–407.

Engle, R., and M. Watson (1987), 'The Kalman filter: application to forecasting and rational expectations models', in T. Bewley (ed.), *Advances in Econometrics*, 5th World Congress, Vol. 1, Cambridge University Press, Cambridge, pp. 245–283.

Fama, E. (1984), 'Forward and spot exchange rates', *Journal of Monetary Economics*, **14**, 319–328.

Flood, R. (1988), 'Asset prices and time varying risk', IMF Working Paper No. 8842.

Frankel, J. (1982), 'In search of the exchange risk premium: a six currency test assuming mean–variance optimization', *Journal of International Money and Finance*, **1**, 255–274.

Frankel, J. (1986), 'The implications of mean–variance optimization for four questions in international macroeconomics', *Journal of International Money and Finance*, **5**, S53–75.

Frankel, J., and C. Engel (1984), 'Do asset-demand functions optimize over mean and variance of real returns? A six currency test', *Journal of International Economics*, **17**, 309–323.

Frankel, J., and K. Froot (1987), 'Short-term and long-term expectations of the yen/dollar exchange rate: evidence from survey data', *American Economic Review*, **77**, 133–156.

Frankel, J., and R. Meese (1987), 'Are exchange rates excessively volatile?' *NBER Macroeconomic Annual*, pp. 118–153.

Froot, K., and J. Frankel (1989), 'Forward exchange rate bias: is it an exchange risk premium', *Quarterly Journal of Economics*, **102**, 139–161.

Gallant, R., G. Tauchen, and D. Hsieh (1988), 'On fitting recalcitrant series: the dollar–pound exchange rate', University of Chicago (mimeo).

Geweke, J., and E. Feige (1979), 'Some joint tests of the efficiency of markets for forward foreign exchange', *Review of Economics and Statistics*, **61**, 334–341.

Giovannini, A., and P. Jorion (1987), 'Interest rates and risk premia in the stock market and in the foreign exchange market', *Journal of International Money and Finance*, **6**, 107–124.

Giovannini, A., and P. Jorion (1988), 'Foreign exchange risk premia volatility once again', *Journal of International Money and Finance*, **7**, 11–113.

Hansen, L. P., and R. Hodrick (1980), 'Forward exchange rates as optimal predictors of future spot rates: an econometric analysis', *Journal of Political Economy*, **88**, 829–853.

Hansen, L. P., and R. Hodrick (1983), 'Risk averse speculation in the forward foreign exchange market: an econometric analysis of linear models', in J. Frenkel (ed.), *Exchange Rates and International Macroeconomics*, University of Chicago Press, Chicago, pp. 113–152.

Hodrick, R. (1989), 'Risk, uncertainty and exchange rates', *Journal of Monetary Economics*, **23**, 433–460.

Hodrick, R., and S. Srivastava (1984), 'An investigation of risk and return in forward foreign exchange', *Journal of International Money and Finance*, **3**, 5–29.

Hodrick, R., and S. Srivastava (1986), 'The covariation of risk premium and expected future spot exchange rate', *Journal of International Money and Finance*, **5**, S5–S26.

Ito, T. (1984), 'On the interdependence of exchange rates and interest rates between the U.S. and Japan', University of Minnesota, Working Paper No. 192.

Ito, T. (1986), 'Capital controls and covered interest parity between the yen and the dollar', *Economic Studies Quarterly*, **37**, 223–241.

Ito, T. (1988), 'Use of (time domain) vector autoregression to test uncovered interest parity', *Review of Economics and Statistics*, **70**, 296–305.

Ito, T. (1990), 'Foreign exchange rate expectations: micro survey data', *American Economic Review*, **80**, 434–449.

Kaminsky, G., and R. Peruga (1989), 'Risk premium and the foreign exchange market', *Journal of International Economics*, **28**, 47–70.

Karajezyc, R. (1985), 'The pricing of forward contract for foreign exchange', *Journal of Political Economy*, **93**, 346–368.

Kiefer, N., and M. Salmon (1983), 'Testing normality in econometric models', *Economic Letters*, **11**, 123–127.

Krugman, P. (1987), 'Exchange rates and international adjustments', NYU Center for Japan–US Business and Economic Studies, Working Paper No. 39.

Mark, N. (1985), 'On time varying risk premia in foreign exchange markets: an econometric analysis', *Journal of Monetary Economics*, **16**, 3–18.