Using Nonlinear Methods To Search For Risk Premia in Currency Futures

by

David A. Hsieh

Fuqua School of Business

Duke University

Abstract

This paper uses currency futures prices to test the joint null hypotheses of rational expectations and absence of a time varying risk premium in the foreign exchange market. We find no linear predictability in the logarithm of futures price changes, either using its own past or past interest differentials. Also we establish that there is no nonlinear predictability in log price changes, conditioning on its own past, or past interest rate differentials. Thus, if a time varying risk premium exists in currency futures market, it is not related to its own past or past interest rate differentials.

** Please send proofs and reprint orders to: Professor David A. Hsieh, Fuqua School of Business, Duke University, Durham, NC 27706. Phone: (919)-660-7779.

1. Introduction

        This paper uses currency futures prices to investigate whether a time
varying risk premium exists in the foreign exchange market. Futures prices
are favored over forward prices for several reasons. First, futures contracts
are traded on organized futures exchanges, which report transaction prices.
Forward contracts are traded on the interbank market, which disseminates only
bid and ask quotes, but not transacted prices. As asset pricing models are
concerned with transaction prices, futures are more appropriate for empirical
analysis. Second, forward contracts are quoted at fixed contract length. If
data are sampled at a finer interval than the contract length, there are
overlapping forecasting intervals. Futures contracts mature four times each
year, so that the futures prices of a given contract is a sequence of
expectations of a fixed future spot price. Using the law of iterated
expectations, Hodrick and Srivastava (1987) show that the futures price equals
the expected futures price on the next business day plus a 1-day risk premium,
which is the covariance between the futures price on the next business day and
a term involving the product of the intertemporal marginal rate of
substitution and the nominally riskfree rate of return, conditioning on
current information. The presence of this 1-day risk premium can be tested by
seeing if the futures price is an unbiased predictor of the futures price on
the next business day. This method avoids the overlapping forecasting
intervals in forward contracts. Third, the nature of forward transactions is
such that it requires careful matching of the forward prices to the spot
prices at maturity, while nothing of the sort is needed for futures prices.

        There is one important disadvantage in futures prices. If we use daily
data for 1-month forward contracts, we will be measuring a 1-month risk
premium, should it exist. If we use daily data for futures contracts, we will
be measuring a 1-day risk premium. It should be easier to detect a risk
premium in the 1-month forward contract, since the 1-month risk premium is
presumably larger than the 1-day risk premium.

        The notation for the paper is as follows. Let $F_{t,T}$ be the futures price

(US\$ price of a unit of foreign currency) at date t for delivery at date T ($T \geq t$). The 1-day risk premium is defined as:

$$F_{t,T} = E_t[F_{t+1,T}] + P_t, \tag{1}$$

where $E_t[\ ]$ is the mathematical expectation operator, given all available information at time t. Hodrick and Srivastava (1987) show that P_t is the covariance between $F_{t+1,T}$ and a term involving the intertemporal marginal rate of substitution (conditional on information at time t). If $P_t = 0$, then the futures price is an unbiased predictor of the futures price on the next business day:

$$F_{t,T} = E_t[F_{t+1,T}]. \tag{2}$$

To avoid the problem of Siegel's paradox, this is typically written as:

$$f_{t,T} = E_t[f_{t+1,T}], \tag{3}$$

where $f_{t,T}$ is the natural logarithm of $F_{t,T}$. We can rewrite this as:

$$E_t[f_{t+1,T} - f_{t,T}] = 0. \tag{4}$$

If rational expectations hold, then the expectation can be replaced by its realization plus a forecast error:

$$f_{t+1,T} - f_{t,T} = v_{t+1}, \tag{5}$$

with the restriction that $E_t[v_{t+1}] = 0$.

Typically, this conditional expectation restriction is tested by examining the autocorrelation structure of v_{t+1} as in McCurdy and Morgan (1987), or its covariance with variables known at time t as in Hodrick and Srivastava (1987). While these procedures test the linear correlation of v_{t+1} with variables in the information set at time t, there are even stronger nonlinear restrictions which are testable. In fact, $E_t[v_{t+1}] = 0$ means that the conditional mean of v_{t+1} given the information set at time t must be zero. This implies, for example, that v_{t+1} cannot be generated by either of the two following models:

$$v_{t+1} = g(v_t) + \nu_{t+1}, \tag{6}$$

$$v_{t+1} = G(x_t) + \nu_{t+1}, \tag{7}$$

where g() and G() are nonlinear functions, x_t is a vector of variables known

at time t, and ν_{t+1} an independent and identically distributed (IID) disturbance. This paper utilizes some of the recent developments in nonparametric regressions to test whether g()=0 and G()=0.

It is perhaps appropriate to discuss the relative merits of using a parametric versus a nonparametric approach to test whether g()=0 and G()=0. In order to derive a parametric form for g() and G(), one needs to specify a complete model, which typically requires special functional forms and strong distributional assumptions before a tractable model is obtained, as in the case of Domowitz and Hakkio (1985). While a correctly specified parametric model can give more efficient estimates of g() or G() than a nonparametric model, a misspecified parametric model can lead to inconsistent estimates and hence incorrect inference. In this paper, we have opted for the nonparametric approach.

Before proceeding further, we first discuss the results from previous papers on currency futures. Hodrick and Srivastava (1987) and McCurdy and Morgan (1987) used daily currency futures prices to test whether the futures price is an unbiased predictor of the futures price on the next business day, and both papers concluded that this was not so. Hodrick and Srivastava (1987) founded that futures price changes were linearly predictable using past bases (the basis is the difference between the futures and the contemporaneous spot price). McCurdy and Morgan (1987) found that futures price changes were linearly predictable using its own past in the CD and DM. In addition, the residuals in the BP, DM, and JY regressions were autocorrelated, and those in the BP, DM, and SF regressions were correlated with lagged changes in spot exchange rates.

Unfortunately, both papers utilized data prior to February 22, 1985, when the Chicago Mercantile Exchange has imposed a daily price limit on currency futures contracts. This created a classic case of data truncation, which leads to inconsistent parameter estimates. This paper uses a sample which begins on February 22, 1985, when daily price limits on currency futures

-3-

were removed, and ends on March 9, 1990, totaling 1275 observations. This

avoids the data truncation problem. We use only the BP, DM, JY, and SF, since

the CD is not a very actively traded contract during most of this period.

The remainder of the paper is organized as follows. Section 2 deals

with the test for a linear dependence in futures price changes. It finds,

contrary to McCurdy and Morgan (1987), no evidence of autocorrelation. Thus,

if a time varying risk premium exists, it is not autocorrelated. Section 3

uses a nonparametric technique to test for the presence of additively

nonlinear dependence in futures price changes. Again, it finds no evidence of

additive nonlinearity. It means that, if a time varying risk premium exists,

it is not an additively nonlinear process. Section 4 tests for both linear

and nonlinear correlation of futures price changes with interest rate

differentials. It finds, contrary to Hodrick and Srivastava (1987), that

futures price changes are not linearly correlated with interest rate

differentials. In addition, there is no evidence of nonlinear correlation

with interest rate differentials. This means that, if a time varying risk

premium exists, it is not correlated (linearly or nonlinearly) with interest

rate differentials.

2. Testing for An Autocorrelated Time Varying Risk Premium

We begin with the summary statistics of the log differences of daily

futures settlement prices in Table 1.[1] There is evidence of leptokurtosis,

but no statistically significant autocorrelations. Table 2 examines the

autocorrelation coefficients of the absolute values of price changes, and

confirms that there is evidence of nonlinear dependence, which is consistent

with conditional heteroskedasticity and other types of nonlinearity. This is

similar to the behavior of daily spot exchange rates.

We formally test for the autocorrelation of the forecast error, under

the joint hypotheses of rational expectations and absence of a time varying

risk premium, by estimating the linear regression:

-4-

$$(f_{t+1,T}-f_{t,T}) = \alpha + \beta \ (f_{t,T}-f_{t-1,T}) + w_t. \tag{8}$$

Table 3 reports regressions of log price changes for all four currencies using the first lag. We also tried using additional lags, up to the first five lags. None provided any evidence of serial correlation of log price changes, so these results are not reported in Table 3. There is no indication that log price changes are serially dependent.

These findings contradict those in McCurdy and Morgan (1987), who found statistically significant βs in the CD and DM, autocorrelation in the residuals in the BP, DM, and JY, and correlation of residuals with lagged changes in spot exchange rates in the BP, DM, and SF. There are several possible explanations for the divergent results. In the first place, daily price limits for currency futures contracts were in effect during the sample period in McCurdy and Morgan (1987). Whenever the daily price limits were effective, the data were truncated, which could lead to inconsistent estimates. McCurdy and Morgan (1987) argued that daily limits did not affect their results, which did not change after they removed the observations for which daily price limit were effective. Removal of the truncated data, however, does not correct the problem. It merely converts truncated data to censored data, which also leads to inconsistent estimates.[2] Our study uses currency futures data after the price limits were removed by the Chicago Mercantile Exchange. This avoids the truncation problem altogether.

A second possible explanation of the divergent results is that the estimation techniques are different. We use ordinary least squares (OLS) with heteroskedasticity-consistent standard errors, as in White (1980), Hansen (1982), and Hsieh (1983). Hodrick and Srivastava (1987) showed that this method is appropriate for futures data asymptotically. McCurdy and Morgan (1987) assumes that w_t in (1) follows Bollerslev's (1987) GARCH process:

$$w_t \sim N(0,h_t)$$

$$h_t = \phi_0 + \phi_1 \ h_{t-1} + \phi_2 \ w_{-1}. \tag{9}$$

If the GARCH model is correctly specified, McCurdy and Morgan (1987) may

have a more efficient estimator than OLS asymptotically (although there is no guarantee that this will occur in any given finite sample). If, however, the GARCH model is misspecified, the estimates of β could be inconsistent. This is likely, since the standardized residuals of their GARCH models are strongly leptokurtic.[3] In addition, the use of residuals rather than true disturbances can bias diagnostic tests towards accepting the estimated model, which could be the reason why McCurdy and Morgan (1987) did not reject the GARCH specification.[4]

To ensure that the difference in estimation strategy is not accountable for the divergent results, we estimate equation (1) using the GARCH(1,1) specification, as in McCurdy and Morgan (1987). The maximum likelihood estimates are obtained using the algorithm in Berndt, Hall, Hall, and Hausman (1974). The results for α and β are in the lower panel of Table 3. Two sets of standard errors are reported. The set in parentheses corresponds to the usual standard errors from the inverse of the Hessian matrix, $A = E[-\partial^2 L/\partial\theta\partial\theta']$, where L is the log likelihood and θ the vector of parameters. The set in square brackets corresponds to the robust covariance estimator, $A^{-1} B A^{-1}$, where $B = E[\partial L/\partial\theta \ \partial L/\partial\theta']$. Based on the theory of quasi- or pseudo-maximum likelihood estimation, the robust covariance estimator, under some circumstances, is correct even if the distribution of the data is misspecified. Since the two sets of standard errors are similar, they lead to the same inference, namely, that there is no serial correlation in log price changes. Thus, we are confident that differences in relative efficiencies between OLS and quasi- maximum likelihood are irrelevant to our inference.

The lack of serial correlation in futures price changes leave us with one conclusion. If a time varying risk premium exists, it is not serially correlated. We now proceed to examine the nonlinear predictability of log price changes.

3. Testing for An Additively Nonlinear Time Varying Risk Premium

It is important to realize that the autocorrelation test is a weak test of the existence of a 1-day risk premium. This can be seen as follows. Consider the log price changes:

$$v_{t+1} = f_{t+1,T} - f_{t,T}$$

$$= (\ E_t[f_{t+1,T}] - f_{t,T}\ ) + (\ f_{t+1,T} - E_t[f_{t+1,T}]\ )$$

$$= -p_t + \varepsilon_{t+1}, \tag{10}$$

where by definition $p_t = f_{t,T} - E_t[f_{t+1,T}]$ is the 1-day risk premium and $\varepsilon_{t+1} = (\ f_{t+1,T} - E_t[f_{t+1,T}]\ )$ is the 1-day forecast error. Under rational expectations, ε_{t+1} is not correlated with p_t and ε_t. Thus,

$$Cov(v_{t+1},\ v_t) = Cov(p_t, p_{t-1}) - Cov(p_t, \varepsilon_t). \tag{11}$$

If p_t is serially uncorrelated, and p_t is uncorrelated with ε_t, then log price changes are serially uncorrelated. Thus, the lack of serial correlation in log price changes does not rule out the existence of a 1-day time varying risk premium.

In this section, we use nonlinear time series methods to detect the presence of a risk premium. The assumption of rational expectations and the absence of a time varying risk premium implies that:

$$E_t[v_{t+1}] = 0, \tag{12}$$

where $v_{t+1} = f_{t+1,T} - f_{t,T}$. This would be violated if v_t follows the following process:

$$v_t = g(v_{t-1}) + \nu_t, \tag{13}$$

where $g()$ is a nonlinear function and ν_t an IID disturbance. Hsieh (1989) called this additive nonlinearity, as ν_t enters additively in the equation. An example of an additively-nonlinear model is the threshold autoregression of Tong and Lim (1980).

It is important to note that rational expectations does not preclude all possible classes of nonlinearity. For example, v_t can be generated by a multiplicatively-nonlinear model:

$$v_t = H(v_{t-1})\ \nu_t, \tag{14}$$

where H() is a nonlinear function and v_t an IID noise independent of v_{t-1}, since this satisfies the requirement that $E_{t-1}[v_t] = 0$. [5]

To test whether the log price changes exhibit additive nonlinearity, we estimate the following model:

$$(f_{t+1,T} - f_{t,T}) = g(f_{t,T} - f_{t-1,T}) + w_{t+1}, \qquad (15)$$

where g() is a smooth nonlinear function and w_{t+1} an IID noise. This can be done by nonparametric regression using kernels, splines, series expansions, or nearest neighbor methods. Diebold and Nason (1990) gave a comparison of these methods, and argued that nearest neighbor method should work well for exchange rates. They found no additive nonlinearity in weekly spot rate changes. Meese and Rose (1990) use a modified version of nearest neighbors called locally weighted regression (LWR) suggested by Cleveland (1979). They also found no additive nonlinearity in a variety of spot exchange rates. We decided to use LWR with the tricubic weighting function suggested by Cleveland and Devlin (1988), because simulations in Hsieh (1991) showed that this procedure can detect all of the well-known additively nonlinear model in the time series literature.

There are two approaches to test the null hypothesis that g()=0 and w_{t+1} is IID. Cleveland and Devlin (1988) developed an in-sample test statistic based on the estimated residuals. Unfortunately, the parameters of its distribution are nearly impossible to compute for more than a few hundred data points. We therefore turn to an out-of-sample forecasting procedure. This is implemented as follows. The first three quarters of the data (951 observations) are used as the initial sample. One-step-ahead forecasts from LWR are generated over the remaining one quarter of the data (225 observations), adding successive data points one at a time. These forecasts are compared to those from a naive martingale model, which uses $f_{t,T}$ to predict $f_{t+1,T}$. If the LWR forecasts are more accurate than the martingale model, then there is evidence against the null hypothesis of g()=0 and w_{t+1} IID.

For each exchange rate, we perform 30 forecasting exercises using LWR

-8-

with three different lag lengths (1 through 3) and ten different window

lengths (f=0.1, 0.2, ..., 0.9, 1.0). Then we compare the ability of these

forecasting models to the naive martingale model using three criteria: root

mean squared forecast error (RMSE), mean absolute error (MAE), and the percent

of correct directional forecasts.

Table 4 reports the RMSE for the LWR forecasts divided by the RMSE of

the martingale model. A number in excess of unity indicates that the

martingale model has a smaller RMSE than the corresponding LWR model. This is

true in almost all cases. In the case of the SF, the LWR model can outperform

the martingale model, but the reduction in RMSE is less than 0.3%.

Table 5 reports the MAE of the LWR forecasts divided by the MAE of the

martingale model. The results are similar to those in Table 4. The

martingale model typically has a lower MAE than the LWR models. The only

exception is the SF, where the LWR can outperform the martingale model, but

the reduction in MAE is less than 0.8%.

Table 6 reports the percentage of correct directional forecasts of the

LWR model. This is done as follows. Each time LWR correctly forecasts the

direction of the next day's futures price change, it receives a score of 1.

Otherwise, it receives a score of 0. The percentage of correct directional

forecasts is the total score divided by the number of forecasts. If this

percentage exceeds 0.5, which is the expected value for a random directional

forecast, then LWR has forecasting ability. This would provide evidence

against the null hypothesis that g()=0. The percentages in Table 6 indicate

that, for the most part, LWR is worse than a random directional forecast. The

large majority of percentages are below 0.5. Even for those above 0.5, none

are statistically greater than 0.5.[6]

The negative results of this large data-mining exercise indicate that

there is no evidence against the null hypothesis that g()=0 and w_{t+1} is IID.[7]

In other words, if a time varying risk premium exists, it is not an additively

nonlinear time series.

4. Testing for Correlation of A Time Varying Risk Premium With Other Variables

So far, we have not found any evidence of the presence of a time varying risk premium, since the forecast errors are not predictable based on their own past, either using linear or nonlinear regressions. Next, we consider whether forecast errors are predictable using other variables in the information set at the time the forecast is made. This comes from the fact that the restriction:

$$E_t[v_{t+1}]=0 \tag{16}$$

implies that the conditional mean of v_{t+1} given the information set at time t is zero. One such variable is the interest rate differential, $(f_{t,T}-s_t)/\tau$, where s_t is the natural logarithm of the spot exchange rate at time t, and τ is the number of calendar days remaining to maturity of the futures contract.

It is straight forward to verify that $(f_{t,T}-s_t)/\tau$ is indeed the interest differential. Covered interest arbitrage guarantees that the forward price is related to the spot price by:

$$G_{t,T} = S_t \exp\{(i_{t,T}-i^*_{,T})\tau\}, \tag{17}$$

where $G_{t,T}$ is the forward price at time t for a contract maturing in T periods, τ the number of calendars days between t and T, $i_{t,T}$ the continuously compounded riskfree US interest rate between time t and t+T, and $i^*_{,T}$ the corresponding foreign interest rate. Hence,

$$(i_{t,T}-i^*_{,T}) = \log\{G_{t,T}/S_t\}/\tau. \tag{18}$$

Since Cornell and Reinganum (1981) showed that $F_{t,T}$ is statistically not different from $G_{t,T}$, we can replace the forward price $G_{t,T}$ with the futures price $F_{t,T}$. Thus

$$(i_{t,T}-i^*_{,T}) = \log\{F_{t,T}/S_t\}/\tau = (f_{t,T}-s_t)/\tau \tag{19}$$

is the interest rate differential.

To calculate $(f_{t,T}-s_t)/\tau$, we obtain spot exchange rates from the Board of Governors of the Federal Reserve System, which collects them at 12 noon Eastern time. To match these spot prices as closely as possible, we use the

futures price immediately prior to noon Eastern time, i.e., 11 a.m. Chicago
time. It turns out that several days in each year, the spot currency markets
are closed while the currency futures market are open, or vice versa. These
dates are dropped from the sample. As a result, there are now 1261
observations for all four currencies.

    To test the joint hypotheses of rational expectations and the absence of
a time varying risk premium, i.e, $E_t[v_{t+1}]=E_t[f_{t+1,T}-f_{t,T}]=0$, we estimate the
linear regression:

$$f_{t+1,T}-f_{t,T} = \alpha + \beta \; [(f12_{t,T}-s_t)/\tau] + u_{t+1}, \tag{20}$$

where $f12_{t,T}$ is the logarithm of the noon (Eastern time) futures price, and
test if $\beta=0$. These single market regressions are reported in Table 7. All β
coefficients are negative, which means that a high US interest rate (relative
to the foreign interest rate) tends to precede a fall in the futures price
(depreciation) of foreign currencies. But none are statistically different
from zero at the 1% (two-tailed) level. While some of the β coefficients are
statistically different from zero using a significance level of 5%, interest
rate differentials have very low explanatory power, since the adjusted R^2s for
the regressions are uniformly low. There is little evidence of the violation
of the joint hypotheses of rational expectations and absence of a time varying
risk premium.

    There is slightly more evidence against the joint hypotheses of rational
expectations and absence of a time varying risk premium if we include interest
differentials in the other three currencies. Table 8 reports the results of
the cross market linear regressions:

$$(f_{+1,T}-f_{t,T}) = \alpha + \Sigma \; \beta_j \; x + w_{t+1}, \tag{21}$$

where $x_t=(1/\tau)[f12_{,T}-s \quad f12_{,T}-s \quad f12_{,T}-s \quad f12_{,T}-s]'$ is the vector of interest
differentials. While none of the interest differentials are statistically
different from zero at the 1% significance level, some are different from zero
at the 5% level. However, the joint test of $\beta=[\beta_j]=0$ is not rejected at the

5% level for all four currencies, and the R^2s of the regressions are very low.

Given the fact that the regressions can only explain a very small fraction of the variance of log price changes, it is not clear whether the statistically significant coefficients are meaningful. Whenever we regress one variable on more and more regressors, chances are high that we will find some regressors to be statistically significant. To deal with this issue, we turn to out-of-sample forecasts of log price changes.

The first method to obtain out-of-sample forecasts is rolling regression. We start the forecast at the beginning of 1988. We perform a 1-day out-of-sample forecast using the regression for the data up until the previous day, and roll the regression forward one observation at a time. The RMSEs relative to those of the naive martingale model are reported in Table 9. The results show that the rolling regression performed uniformly poorer than the naive martingale model. This is true both for the single market and the cross market model.

The second method uses LWR instead of the rolling regression. This allows a nonlinear predictor of the following form:

$$f_{t+1,T} - f_{t,T} = G(x_t) + w_{t+1}, \tag{22}$$

where $G()$ is a smooth function, w_{t+1} is IID noise, and x_t is the own interest rate differential in the single market model and the vector of interest rate differentials in the cross market model. The first three quarters of the data prior is used as the initial sample, and one-step-ahead forecasts using LWR are generated for the remaining one quarter of the data, adding one observation at a time. The tricubic weights and f=0.1,..,1.0 are used. As in the previous section, we compare the out-of-sample forecasting performance of LWR to that of the naive martingale model as a test of the null hypothesis that $G()=0$ and w_{t+1} is IID.

The RMSEs are reported in Tables 9, and the MAEs are reported in Table 10. The naive martingale model achieved a lower RMSE and MAE than any of the LWRs, both in the single market and cross market models. Table 11 reports the percent of correct directional forecasts. For the BP, DM, and JY, there is no

-12-

evidence that LWR can outperform a random directional forecast. For the SF, there is slight evidence that LWR can outperform a random directional forecast, since several of the percentages are greater than 0.5 at the 5% significance level.

On the whole, we conclude that, if a time varying risk premium exists, it is not correlated, linearly or nonlinearly, to interest rate differentials.

Our results are in sharp contrast to those in Hodrick and Srivastava (1987), who found not only statistically significant β coefficients, but also very high R^2s. There are several possible explanations. The first has already been mentioned. Daily price limits were in effect for the futures data used in Hodrick and Srivastava (1987). The truncation problem could lead to inconsistent parameter estimates. This problem does not arise in our sample.

The second explanation is that our regressions are different from those in Hodrick and Srivastava (1987), who estimated the following equation:

$$(F_{t+1,T}-S_t)/S_t \quad = \quad \alpha + \zeta \ (F12_{t,T}-S_t)/S_t + u_{t+1}, \tag{23}$$

and tested if ζ = 1, where can be rewritten as:

$$(F_{t+1,T}-F_{t,T})/S_t \quad = \quad \alpha + \beta \ (F12_{t,T}-S_t)/S_t + u_{t+1}, \tag{24}$$

where $\beta = \zeta-1$. To ensure that this difference is not the cause of the divergent results, we report the results of this regression in the lower panel of Table 7. They are similar to those in the upper panel of the same table, and very different from those in Hodrick and Srivastava (1987). While the BP rejects the null hypotheses, the other three currencies do not. Also, the R^2s of the regressions are very low. The differences in the form of the regression cannot account for the divergent results.

The third explanation is the failure of the asymptotic distribution of the GMM estimator to approximate its finite sample distribution. Hodrick and Srivastava (1987) proved the asymptotic properties of the GMM estimator assuming that the number of contracts go to infinity. Since there were only

21 contracts in their data, the finite sample distribution of the GMM estimator may be very different from its asymptotic distribution. This is unlikely to explain the difference between their results and ours, since we also have 21 contracts in our data, and since our estimator is a special case of the GMM estimator.

A fourth possibility is that Hodrick and Srivastava (1987) had more accurate estimates of β, since they used twice as many observations as we do. On closer examination, however, their larger sample size cannot account for the divergent results, as their results did not change when they split the sample into two halves.

5. Conclusion

In this paper, we use currency futures prices to test the joint null hypotheses of rational expectations and absence of a time varying risk premium. We find no linear predictability in log price changes, either using its own past or past interest differentials. This is consistent with Meese and Rogoff (1983), but not with McCurdy and Morgan (1987) and Hodrick and Srivastava (1987). The most likely explanation is that the latter papers used data which suffered from truncation caused by to the presence of price limits on currency futures contracts, while our data do not suffer from the same problem. In addition, this paper also established that there is no <u>nonlinear</u> predictability in log price changes, conditioning on its own past, or past interest rate differentials. We conclude that, if a time varying risk premium exists in currency futures markets, it is not detectable using the methods employed in this paper, since it is neither linearly nor additively nonlinearly dependent on its own past or past interest rate differentials.

References

Berndt, E.K., B.H. Hall, R.E. Hall, and J.A. Hausman, 1974, Estimation and Inference in Nonlinear Structural Models. Annals of Economic and Social Measurement 4, 653-665.

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

Bilson, J., and D. Hsieh, 1989, The Profitability of Currency Speculation. International Journal of Forecasting 3, 115-130.

Bollerslev, T., 1986, Generalized Autoregressive Conditional Heteroskedasticity. Journal of Econometrics 31, 307-327.

Cleveland, W.W., 1979, Robust Locally Weighted Regression and Smoothing Scatterplots. Journal of the American Statistical Association 74, 829-836.

Cleveland, W.S. and S.J. Devlin, 1988, Locally Weighted Regression: An Approach to Regression Analysis by Local Fitting. Journal of the American Statistical Association 83, 596-610.

Cornell, B., and M. Reinganum, 1981, Forward and Futures Prices: Evidence from the Foreign Exchange Markets. Journal of Finance 36, 1035-1046.

Cumby, R., and M. Obstfeld, 1981, A Note on Exchange-Rate Expectations and Nominal Interest Differentials: A Test of the Fisher Hypothesis. Journal of Finance 36, 697-704.

Diebold, F., 1987, Empirical Modeling of Exchange Rate Dynamics (Springer-Verlag, New York).

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

Domowitz, I., and C. Hakkio, 1985, Conditional Variance and the Risk Premium in the Foreign Exchange Market. Journal of International Economics 19, 47-66.

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

Hansen, L., 1982, Large Sample Properties of Generalized Method of Moment Estimators. Econometrica 50, 1029-1054.

Hansen, L., and R. Hodrick, 1980, Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometrics Analysis. Journal of Political Economy 88, 829-853.

Harvey, C., 1988, Commentary on Tests of Unbiasedness in the Foreign Exchange Futures Markets: The Effects of Price Limits. Review of Futures Markets 7, 167-171.

Hodrick, R., 1987, The Empirical Evidence of the Efficiency of Forward and Futures Foreign Exchange Markets (Harwood Academic Publishers, New York).

Hodrick, R., and S. Srivastava, 1987, Foreign Currency Futures. Foreign Currency Futures 22, 1-24.

Hsieh, D., 1983, A Heteroskedasticity-Consistent Covariance Matrix Estimator For Time Series Regressions. Journal of Econometrics 22, 281-290.

Hsieh, D., 1989, Testing for Nonlinearity in Daily Foreign Exchange Rate

Changes. Journal of Business 62, 339-368.

Hsieh, D., 1991, Chaos and Nonlinear Dynamics: Application to Financial Markets. Journal of Finance, 46, 1839-1877.

Kodres, L., 1988, Tests of Unbiasedness in the Foreign Exchange Futures Markets: The Effects of Price Limits. Review of Futures Markets 7, 139-166.

McCurdy, T.H., and I.G. Morgan, 1987, Tests of the Martingale Hypothesis for Foreign Currency Futures. International Journal of Forecasting 3, 131-148.

Meese, R., and K. Rogoff, 1983, Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample? Journal of International Economics 14, 3-24.

Meese, R., and A. Rose, 1990, Nonlinear, Nonparametric, Nonessential Exchange Rate Estimation. American Economic Review 2, 192-196.

Rogalski, R., and J. Vinso, 1978, Empirical Properties of Foreign Exchange Rates. Journal of International Business Studies 9, 69-79.

Tauchen, G., 1985, Diagnostic Testing and Evaluation of Maximum Likelihood Models. Journal of Econometrics 30, 415-443.

Tong, J., and K. Lim, 1980, Threshold Autoregression, Limit Cycles, and Cyclical Data. Journal of the Royal Statistical Society, Series B 42, 245-292.

Westerfield, J., 1977, An Examination of Foreign Exchange Risk Under Fixed and Floating Rate Regimes. Journal of International Economics 7, 181-200.

White, H., 1980, A Heteroskedasticity-Consistent Covariance Estimator and a

Direct Test for Heteroskedasticity. Econometrica 48, 817-838.

Notes:

The results of this paper are the same if we use the last transacted price rather than the
settlement price.

Kodres (1988) used a limited dependent variable framework to treat the limit moves. However,
Harvey (1988) pointed out that Kodres (1988) failed to take into account the conditional
heteroskedasticity in the data. This can also lead to inconsistent estimates.

McCurdy and Morgan (1987) does not state explicitly whether they assume a normal conditional
distribution for the standardized residuals. This is typically the case in GARCH models.

See Tauchen (1985) for a discussion of this point.

The multiplicative nonlinearity can also lead to autocorrelations in the absolute values of t
data.

It is curious that the JY appears to have "reverse" predictability. For example, LWR using 1
lag and f=0.1 predicts the direction correctly 41.85%, which is more than 2 standard errors a
from 50%. This reserse predictability is more apparent than real, for two reasons. First, t

results for the JY are not independent across lags and across the f. Second, the results tur

out to be period specific. If we had used a different period, the proportion of correct

directional forecasts is much closer to 50%.

In particular, there is no low complexity chaos present in daily currency futures prices. Fo

discussion of this point, see Hsieh (1991).