NBER WORKING PAPER SERIES

CAN THE MARKOV SWITCHING MODEL FORECAST EXCHANGE RATES?

Charles Engel

Working Paper No. 4210

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 1992

I would like to thank Frank Diebold and Jim Hamilton for useful suggestions. Much of the work for this paper was completed while I was a visiting scholar at the Federal Reserve Bank of Kansas City. The views expressed in this paper are not necessarily those of the Federal Reserve Bank of Kansas or the Federal Reserve System. This paper is part of NBER's research program in International Finance and Macroeconomics. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

NBER Working Paper #4210
November 1992

CAN THE MARKOV SWITCHING MODEL FORECAST EXCHANGE RATES?

ABSTRACT

A Markov-switching model is fit for eighteen exchange rates at quarterly and monthly frequencies. This model fits well in-sample at the quarterly frequency for many exchange rates. By the mean-squared-error or mean-absolute-error criterion, the Markov model does not generate superior forecasts at a random walk or at the forward rate. There appears to be some evidence that the forecast of the Markov model are superior at predicting the direction of change of the exchange rate.

Charles Engel
Department of Economics
University of Washington
Seattle, WA 98195
and NBER

The behavior of the dollar against certain currencies appears to be forecastable. In particular, the exchange rate appears to follow long swings -- it drifts upward for a considerable period of time, and then switches to a long period with downward drift. Engel and Hamilton (1990) found that the dollar/mark, dollar/pound and dollar/French franc exchange rates can be described well by Hamilton's (1988, 1989) Markov switching model.

This paper investigates whether the Markov switching model is a useful tool for describing the behavior of floating exchange rates more generally. The time series properties of eighteen exchange rates in the post-Bretton Woods period, including eleven non-U.S. dollar exchange rates, are investigated. In general, the model of Engel and Hamilton (EH) does not clearly outperform the random walk model or the forward exchange rate in out-of-sample forecasts. The mean-squared-errors and mean-absolute-errors of the forecasts of the segmented trends model tend not to be significantly lower than those of a zero drift random walk, a random walk with drift or the forward rate.[1] There is some evidence, however, that the segmented trends model is superior to its competitors in forecasting the direction of change of exchange rates.

The paper focuses on exchange rates that are measured at the end of each quarter. The estimation period is 1973-1986, and the post-sample forecast period is 1986-1991. The model was also estimated for the period 1973-1988, and forecasts were constructed for the post-sample period.

---

[1] The terms segmented trends model and Markov switching model are used interchangeably in this paper.

There is no natural frequency to measure changes in exchange rates. If the exchange rate process follows a two-state Markov switching process on monthly data, for example, it will not generally then also follow a two-state Markov switching process on quarterly data, and vice-versa. So, the models were also estimated on monthly data. The general conclusions about the forecasting ability of the segmented trends model was not affected by the choice of estimation period, but the model estimated on quarterly data does seem to perform better in sample than the monthly model.

Section 1 reviews briefly the Markov switching model, and techniques for estimation. Section 2 presents estimates of the model on quarterly changes in logs of exchange rates. Section 3 discusses criteria for evaluating the forecasting performance of the model. Section 4 presents the results of the forecasting contest between the Markov switching model, the random walk with and without drift and the forward exchange rate.

Section 5 discusses estimates of a Markov switching model on monthly data, and compares this model to the one estimated from quarterly data. Section 6 concludes the paper by speculating on the meaning of the less than spectacular forecasting performance of the Markov-switching model.

1. The Model and Its Estimation

This section outlines the Markov switching model and discusses its estimation. An excellent survey is available in Hamilton (1991b).

Changes in the log of the exchange rate are distributed normally with mean μ_i and variance σ_i^2 in each of two possible states of the world ($i = 1,2$). The state at time t is determined randomly, and depends only on the state at were in state i at time t-1 is p_i.

This is a simple version of the more general Markov switching model described by Hamilton (1989, 1991b). The exchange rate could be allowed to follow more general stochastic processes in each state. There can be more than two states. The p_i could vary over time. Moreover, the exchange rate need not be modeled as a univariate process. It could be one element of a multivariate Markov switching process.

The simple univariate model described above, in which the log of the exchange rate follows a random walk with drift in each of two states (the segmented trends model), was found by EH to provide a good description of exchange rate behavior. It outperformed a random walk both in and out of sample. The model still allows a variety of behavior in the time series of exchange rates. In particular, large values of the p_i generate the "long swings" in exchange rates which EH argue are characteristic of dollar exchange rates. Furthermore, the parsimony of the parameterization (only six parameters need to be estimated for each exchange rate: μ_i, σ_i and p_i, i = 1,2) promises good forecasting properties for the model.[2]

The parameter vector $\Theta = (\mu_1,\mu_2,\sigma_1,\sigma_2,p_1,p_2)$ can be estimated by maximum likelihood methods. The sample likelihood is a function of the observed values of the changes in the logs of exchange rates $(y_1,y_2,\ldots,y_T)$. The states $(s_1,s_2,\ldots s_T)$ are unobserved, and the econometrician must draw inferences about the probability of s_t=1 or 2 based on the observed data.

[2] The unconditional distribution of changes in the log of the exchange rate in this model is a mixture of normal distributions. Baillie and Bollerslev (1989) find evidence that daily changes in the log of the exchange rate are not normally distributed, exhibiting skewness and leptokurtosis. Both of these characteristics could arise from a distribution that was a mixture of normals.

The maximum likelihood estimation in this paper is performed using the EM algorithm described by Hamilton (1990). Hamilton (1989, 1991b) shows how "smoothed" inferences can be drawn about the probability that $s_t = 1$ or 2, given θ and $\underline{y}_T$ (where $\underline{y}_T = (y_1, y_2, \ldots y_T)$). An initial guess at θ allows one to calculate these probabilities. New guesses at the parameters can be calculated from the formulas

$$\hat{\mu}_j = \frac{\sum_{t=1}^{T} y_t \cdot p(s_t=j|\underline{y}_T;\hat{\theta})}{\nu + \sum_{t=1}^{T} p(s_t=j|\underline{y}_T;\hat{\theta})}$$

$$\hat{\sigma}_j^2 = \frac{\beta + .5\cdot\nu\cdot\hat{\mu}_j + \sum_{t=1} (y_t - \hat{\mu}_j)^2 \cdot p(s_t=j|\underline{y}_T;\hat{\theta})}{\alpha + \sum_{t=1}^{T} p(s_t=j|\underline{y}_T;\hat{\theta})}$$

$$\hat{p}_j = \frac{\sum_{t=2}^{T} p(s_t=j, s_{t-1}=j|\underline{y}_T;\hat{\theta})}{\hat{\rho}_j - p(s_1=j|\underline{y}_T;\hat{\theta}) + \sum_{t=2}^{T} p(s_{t-1}=j|\underline{y}_T;\hat{\theta})}$$

where

$\hat{\rho}_1 = (1-\hat{p}_2)/[(1-\hat{p}_1)+(1-\hat{p}_2)]$ and $\hat{\rho}_2 = 1-\hat{\rho}_1$. These new values of $\hat{\theta}$ can be used to calculate new smoothed probabilities. Final estimates of θ are produced when this iterative procedure satisfies some convergence criteria.

Ignore for the moment α, β and ν, and consider the meaning of these formulas for parameter estimation. The numerator of the estimator of μ_j weights each observation by the probability that that observation is drawn from state j. The denominator divides by the number of observations expected to have come from state j. Similar interpretations can be given to the other two formulas. In the denominator of the equation for $\hat{p}_j$ a correction is made for the difference between the unconditional probability that the state at time 1 is j ($\hat{\rho}_j$) and the smoothed probability that it is j.

In estimating any mixture of normal distributions, maximum likelihood estimation is not consistent. An infinite value for the likelihood can be achieved by having the mean of one state equal the value of any of the observations, with the variance of that state equal to zero. Hamilton (1991a) proposes a pseudo-Bayesian solution to this problem that allows us to impose our prior belief that neither of the states is likely to have zero variance. These estimators effectively impose priors that we have ν additional observations from each regime that take on the value zero, and 2α observations that σ_j^2 equals β/α, where ν, β and α are parameters chosen by the econometrician to reflect his priors.

In practice, a given data set may by pure chance have a few observations at disparate dates that take on very similar values. With a limited number of observations, if β/α is set to a low value, the likelihood will be maximized by taking those few data points to represent observations from state 2, with p_2 equal to zero. Such an occurrence can be avoided by choosing β and α so that β/α and α are large. There is the usual tension, however, between imposing enough prior information to get sensible estimates and imposing too much so that the data are not allowed to speak for themselves.

2. Estimates on Quarterly Data

This section investigates the behavior of the U.S. dollar exchange rate relative to the Japanese yen, the U.K. pound, the Canadian dollar, the Italian lira, the French franc, the Swiss franc and the German mark; and the behavior of the yen and pound against the remaining currencies. The exchange rate changes are measured on a quarterly basis in this section. (Monthly changes are taken up in section 5).

The data were compiled by Data Resource Inc., and measured as the average of the bid and ask rates on the last day of the period. The log changes are multiplied by 100 to express things in percentage terms.

The first observation is for the second quarter of 1973, so the first quarterly change is observed at the end of the third quarter of that year. The model is estimated through the first quarter of 1986. (Later in this section we will consider the same model, estimated through the first quarter of 1988.)

Estimates for the first sample period are reported in Table 1. These estimates were undertaken with the priors $\alpha = .1$, $\beta = .5$, and $\nu = .05$ imposed. The maximum likelihood estimates show that the two states generally differ not only in mean, but in variance as well. The standard errors of the coefficient estimates are reported in parentheses.

For many of the currencies, there seems to be evidence of long swings in the exchange rate. That is, the probability of staying in a state once you are in it (p_i) is large. This seems to be true for all U.S. dollar exchange rates except for the U.S. dollar/Canadian dollar rate, for all yen rates, and for the pound/Canadian dollar rate. On the other hand, the exchange rates for those countries that are in close geographical proximity -- the U.S. dollar/Canadian dollar and the pound against European currencies -- show much lower p_i.

The estimation procedure constrains p_i to lie between zero and one. For the pound/French franc and pound/Swiss franc rates, the likelihood is maximized when one of these probabilities equals zero. Hence, one of the states is very short-lived. Standard errors are constructed under the assumption that p_2 is constrained to equal zero.

EH note that while a simple random walk is nested in the segmented trends model $((\mu_1 = \mu_2)$ and $(\sigma_1 = \sigma_2))$, under the null hypothesis two problems arise which invalidate the usual approaches for establishing asymptotically consistent tests of that null. The first is that under the null, p_1 and p_2 are not identified. The second is that the derivative of the likelihood with respect to μ_i and σ_i is identically zero under the null hypothesis.

EH suggest two alternative null hypotheses. Each represents a stochastic process which in turn nests the random walk. The first is simply $\mu_1 = \mu_2$. Under this null hypothesis, the two states are different only by their variances, so the exchange rate essentially follows a random walk with heteroskedastic errors.

Tests of this hypothesis are reported in Table 2. The null can be rejected at the 5% level for all currencies of countries not in close geographic proximity. It cannot be rejected for any others except the pound/Swiss franc rate. (The 5% critical value is 3.84).

The second null proposed by EH is $p_1 + p_2 = 1$, so that the distribution of s_t is independent of s_{t-1}. EH note that large values of p_1 and p_2 result in exchange rate movements that are characterized by long swings. They strongly reject $p_1 + p_2 = 1$ in favor of the segmented trends model for the U.S. dollar against the French franc, English pound and German mark.

Table 2 shows that the null hypothesis of $p_1 + p_2 = 1$ is indeed rejected for those three currencies. In fact, it is rejected at the 5% level for the exchange rates of all countries that are not in close proximity, with the exception of the dollar/Swiss franc rate. It is not rejected for any of those countries which are neighbors.

The conclusion that can be reached here is that the segmented trends model fits well in-sample for the thirteen exchange rates of non-neighboring countries. In particular, it must outperform a random walk because it outperforms generalizations of the random walk. In section 4, we will consider the forecasting ability of the segmented trends model for these thirteen exchange rates.

Local Maxima

Estimation of the Markov switching model often encounters multiple local maxima. Occasionally the situation might arise in which the estimated parameters of the global maximum are much less in accord with one's priors than the estimates from one of the other local maxima. For example, in one particular sample, by chance it might happen that the change in the log of the exchange rate is very nearly the same on several different dates. This might lead to global maximum likelihood estimates which assign all of those observations to one state that has a very low variance. However, it seems unlikely that the MLE has captured the true data generating process in this case, since one might believe it is not probable that that particular realization would recur frequently in a large sample.

One way of dealing with this problem is to impose stronger priors. In this case, that means larger values for α and β. That is the approach taken when the model was reestimated over the 1973-1988 period. Another approach is to examine the properties of the parameter estimates from some local maxima whose that generate a value for the likelihood near that of the global maximum.

It was noted above that for the countries in close proximity that the global maximum estimates showed little evidence of long swings. We cannot reject $p_1 + p_2 = 1$ for any of these countries, and for most of them one of the states was very short-lived. Is this just an artifact of our sample? Were there any local maxima that found greater evidence of long swings? In fact, for these exchange rates, the answer is no. Of the local, non-global maxima, the sum of p_1 and p_2 was less than 1.15 in all cases except for the pound/French franc. Even in that case, the standard errors were high enough that the null hypothesis of $p_1 + p_2 = 1$ could not be rejected.

Estimation over 1973-1988

The model was reestimated using two additional years of data, so that the final in-sample date is the first quarter of 1988. This corresponds to the estimation period in Engel and Hamilton. In this case, α was increased to 1, and β to 5, while ν was still set equal to .05.

The reason for reestimating over this period is that by choosing the same estimation period as EH it is clear that the sample period is not chosen in any way that would bias the performance of the out-of-sample forecasts. The drawback is that we are left with only three years (12 quarters) of data for assessing the out-of-sample forecasting ability of the model.

With only one exception, the parameter estimates over the longer period match those in the shorter period quite closely. The exception is the U.S. dollar/Canadian dollar rate. For this exchange rate, we find much stronger evidence of long swings: $p_1 = .908$ and $p_2 = .978$. We can strongly reject $p_1 + p_2 = 1$, and $\mu_1 = \mu_2$.

3. Standards for Measuring Forecastability

We would like to know whether the Markov switching model forecasts well out of sample. To a large extent, the question is a relative one -- does the model do well compared to some alternative? A minimal standard would be that the model forecast better than some naive alternative, such as a random walk.

This is precisely the standard used by Meese and Rogoff (1983) in their famous paper which found that virtually no economic model of the exchange rate could outforecast a zero-drift random walk. Since that paper appeared, many authors have used the zero-drift random walk specification as the standard for comparison in measuring forecasting performance (see, for example, Diebold and Nason (1990)).

Engel and Hamilton (1990) argue that the random walk with drift is a more reasonable standard of comparison when the drift term is estimated to be significantly different from zero. Indeed, suppose the drift term is significant within sample, but by some measure the driftless random walk outperforms the random walk with drift in the out of sample period. That suggests that the data generating process in the post-sample period is no longer a random walk with the same drift term as in-sample. In fact, if the data generating process has switched from a random walk with drift to a driftless random walk -- and if the underlying state (drift or no drift) switches according to a Markov process -- then the exchange rate can be described exactly by the segmented trends model.

So, first, if the drift is significant in sample, then the appropriate out of sample standard is the random walk with drift. While for a particular sample the zero-drift random walk may perform better out of sample, one must avoid the fallacy of choosing the out-of-sample standard based on its out-of-sample performance. Second, if the zero-drift random walk does significantly outperform the random walk with drift out of sample (but the drift term was significant in sample), there is prima facie evidence that the exchange rate can be described by a model similar to the segmented trends model.

One of the purposes of this section of the paper is to determine whether the random walk with drift or the driftless random walk is the appropriate standard. The first column of Table 3 presents t-statistics for the test of whether the mean change in the log of the exchange rate is zero for quarterly data from 1973:3 to 1986:1. The mean is significant at the 5 per cent level for only four of the thirteen exchange rates, and at the 10 per cent level for only one additional rate. The second and third columns of Table 3 show LM tests for the hypothesis that the mean of the change in the log of the exchange rate changed from the 1973:3-1986:1 to the 1986:2-1991:1 period, under the normality assumption. The first column performs the test assuming equal variances in the two periods, and the second assumes different variances. In both sets of tests, the null of the same mean drift is rejected only for the dollar/lira rate at the 5 per cent level (although the yen/lira statistic is very nearly large enough to reject at this level.) The evidence from Table 3 is that it probably does not make much difference whether the random walk with drift or the driftless random walk is used as the standard for out-of-sample forecasting. This conclusion will be further supported shortly by evidence on the out-of-sample forecasting properties of these two alternatives.

We will also compare the forecasting ability of the Markov model to that of the forward exchange rate. Under some assumptions -- for example, if there are risk-neutral speculators and markets are efficient -- the forward rate should be the best possible predictor of the future exchange rate. So, it is a natural standard of comparison.

Once we decide on the alternative with which to compare the Markov switching model, we must decide what measure of forecasting ability is best. The most frequently used measure is the mean-squared-error of the forecasts. Using this measure implies that the forecaster has a quadratic loss function defined over the forecasts. It is, however, difficult to justify such a loss function on the basis of a particular economic decision facing some individual. From an economist's standpoint, a more reasonable approach would be to postulate some objective for an economic agent, and then measure the value of different forecasts in terms of achieving the maximum of that objective. For example, West, Edison and Cho (1992) assume that agents wish to choose consumption levels and asset shares at each point in time to maximize a function of the mean and variance of next period's wealth. The best forecast on that criterion is not necessarily one that minimizes the mean square error of forecasts.

Rather than undertake the difficult exercise of West, Edison and Cho, here we will look at a few different measures of forecasting ability. The forecasts are judged on the basis of their ability to minimize the mean squared forecast error and the mean absolute forecast error, and by their ability to predict the direction of change of the exchange rate.

One argument for looking at the direction of change of the exchange rate is that it may actually not be a bad proxy for a utility-based measure of forecasting performance. Leitch and Tanner (1991) find that the direction of change criteria is the best proxy among several (including mean squared error and mean absolute error) for choosing forecasts of interest rates on their ability to maximize expected trading profits.

It is also possible to think of important circumstances under which the direction of change criteria is exactly the right one for maximizing welfare of the forecaster. One example is that central banks under pegged exchange rate systems often are interested only in the direction of change in the exchange rate. They might need to intervene to support the currency if it is expected to depreciate, regardless of the size of the expected depreciation.

So, we will compare the mean squared error, the mean absolute error and the direction of change of the forecasts of the Markov model to that of the random walk (with and without drift) and the forward rate.

It is interesting at this point to compare the forecasts of the random walk with drift to the random walk with zero drift. For each of the thirteen exchange rates, Table 4 reports the mean squared error of forecasts for one-, two- and four-quarter changes in the exchange rate over the 1986:2-1991:1 period. Table 5 reports the mean absolute error for the same exchange rates.

The random walk with drift is the better forecaster in terms of mean squared error at all horizons for the dollar/yen, dollar/Swiss franc, dollar/mark, and yen/Swiss franc exchange rates. The zero drift random walk is better for all horizons for the other nine exchange rates. Using the mean absolute error criterion, the random walk with drift is better at all horizons only for the dollar/Swiss franc. However, there are five exchange rates for which one model is better at the one-quarter horizon and the other is better at the four-quarter horizon.

Simply comparing the values of the mean-squared or mean-absolute forecast errors does not give us any idea of the significance of the difference.

Diebold and Mariano (1991) propose a statistic for comparing forecasts that is asymptotically N(0,1). It can be used in our context to test whether the mean squared error of one forecast is better than another, or whether one forecast has a significantly lower mean absolute error.[3]

Table 4a reports the Diebold-Mariano statistics for the mean-squared error of the forecast of the random walk with drift compared to the driftless random walk, while Table 5a reports the analogous statistic in the case of mean-absolute errors. At the one-month horizon, the difference in the mean-squared error is not significant at the 10% level for any exchange rate. The same is true at the 5 per cent level for mean-absolute errors, and only in the case of the dollar/pound is the zero-drift random walk significantly better at the 10 per cent level. At the two-quarter horizon, the zero-drift random walk has a significantly lower mean-squared error only for the dollar/lira, and a significantly lower mean-absolute error for the dollar/French franc and dollar/lira. At the four-quarter horizon, we can reject the null of no difference in mean squared error at the 5 per cent level only for the dollar/lira level, and at the 10 per cent level, the dollar/French franc, dollar/pound and yen/pound. In terms of the mean absolute error, the zero drift random walk does significantly better than the random walk with drift only for the dollar/lira. In general, then, there is not much of a significant difference between the forecasting ability of the random walk with drift and the random walk without drift models.

Table 6 allows us to check whether the random walk with drift model gets the direction of change correct more than half the time. Since the zero-drift

[3] Following Diebold and Mariano, the Newey-West weighting scheme is used in constructing the measure of the variance of the difference in forecast accuracy.

random walk predicts no change, this is, in a sense, another check of the relative forecasting ability of the two specifications. The random walk with drift correctly forecasts the direction of change more than half the time for the dollar/yen, dollar/Swiss franc, dollar/mark, yen/French franc, yen/lira, yen/pound and yen/mark. At the one quarter horizon there are 260 forecasts across the 13 exchange rates. The random walk with drift forecasts the direction correctly about half the time -- 132. If we can consider each of those forecasts to be independent of the other, then clearly the random walk with drift is not significantly better than a coin toss at forecasting the exchange rate. The same is true at the other horizons -- it forecasts correctly 126 out of 247 two-quarter-ahead forecasts and 110 out of 221 four-quarter-ahead forecasts.

So, we conclude that we will compare the forecasting ability of the Markov switching model against three alternatives: the random walk with drift, the zero-drift random walk and the forward rate; using three measures of forecasting accuracy: significant differences in mean squared error, mean absolute error and accuracy of direction of change.

4. Forecasts from Quarterly Model

Table 4 examines the mean squared error for out-of-sample forecasts. The model is estimated through the first quarter of 1986, and the forecast errors are calculated for the five-year period from the second quarter of 1986 to the first quarter of 1991.

First, consider the comparison of the forecasts of the segmented trend to the random walk with drift. The segmented trend model improves on the random walk in only about half of the cases. For the one-quarter ahead forecasts, the segmented trends model does better for seven of the thirteen exchange rates. It outforecasts the random walk with drift for five exchange rates at the two-quarter horizon, and six at the four-quarter horizon. Inspection of Table 4 reveals that there are sometimes large differences in the mean squared forecast errors for the two models, but not consistently favoring one of the models. Table 4a examines the significance of the differences using the Diebold-Mariano statistics. Most of the differences in mean-squared error are not significant. There is only one currency for which the Markov model does significantly better (the dollar/Swiss franc) and only a few for which the random walk does significantly better at any horizon.

The performance of the segmented trends model against the forward rate is quite similar -- each wins the forecasting contest about half the time. The Markov model does significantly better for the dollar/Swiss franc, and the forward rate has a significantly lower mean squared error for the dollar/lira and yen/lira.

Generally, the random walk with zero drift outperforms all other models in terms of mean squared forecast errors. For example, in the four-quarter ahead forecasting contest, the zero-drift random walk wins nine of the thirteen contests (the segmented trend wins three, and the random walk with trend wins one). At a one-quarter horizon, the zero-drift random walk outperforms the segmented trends model for seven of the thirteen exchange rates, and at the other horizons for ten of the exchange rates.

However, the variance of the forecast errors is quite large. Hence, by the Diebold-Mariano statistics, the zero-drift random walk significantly outperforms the segmented trends model only for the dollar/lira rate.

Not surprisingly, for the five exchange rates for neighboring countries (for which the Markov model was not much different than a random walk in-sample), the forecasts of the segmented trends model and the random walk with drift are very similar.

In almost all cases, the global maximum likelihood estimates of the segmented trends model do no worse than the local maxima in mean-squared forecast errors.

Table 5 reports the mean absolute errors of the forecasts, and Table 5a reports the Diebold-Mariano statistics. Qualitatively there is essentially no difference in the conclusions reached using the M.A.E. criterion as compared to M.S.E.

Table 6 reports the count of how many times each model forecasted in the correct direction. Clearly, the zero-drift random walk does not belong in this comparison. An analogy to the no-change model in this case would be the coin-toss model.

In the five-year post-sample period, there are twenty quarters to predict. For all but two of the currencies, the segmented trend gets the direction correct more than half of the time. (It misses for the dollar/lira and dollar/yen.) Out of the 260 one-quarter ahead forecasts, the Markov model forecasts the direction correctly 147 times. Treating these as independent forecasts, the null hypothesis that the model can predict direction no better than a coin toss can be rejected with a p-value of 1.7%.

The random walk with drift and the forward rate do not do nearly so well at the one-quarter ahead horizon. The random walk model gets direction correct more than half of the time for only six currencies, while the forward rate achieves this level for only five currencies. The random walk gets the direction correct for 132 out of 260 forecasts, and the forward rate for 121 out of 260 forecasts.

The results are not as impressive at the two and four-quarter horizons, although they still favor the Markov model. At both horizons, the Markov model gets the direction of change right more than half the time: 128 out of 247 two-quarter ahead forecasts, and 114 out of 221 four-quarter ahead forecasts. The p-values for these outcomes, assuming independence of all of the forecasts, are .283 and .319, respectively. It also outperforms the random walk (126 correct at the 2-quarter horizon, and 110 at the 4-quarter horizon) and the forward rate (114 and 98).

So, when the size of the forecast error is de-emphasized, the segmented trends model seems more attractive.

<u>Estimates for 1973-1988</u>

The results of the forecasting contest for this period are much the same as for the longer forecasting period. There is virtually a dead heat between the segmented trends model and the random walk with drift in terms of mean-squared error. The segmented trends model does slightly worse against the zero-drift random walk in this period. However, given the shortness of the post-sample period the difference should be attributed to the particular set of realizations over this period.

At the one-quarter horizon, the segmented trend model again does better than the coin-toss model in predicting the direction of change. The p-value for the null hypothesis of the coin-toss model is .082. The random walk and the forward rate predict the direction of change correctly less than half the time at all horizons. At the two- and four-quarter horizons, the Markov model also gets the direction wrong more than half the time over this period.

although it does better than either the random walk or the forward rate.

4. Monthly Model

Currencies are traded virtually continuously around the globe. A discrete-time stochastic process is only an approximation to the underlying continuous-time process. A problem similar to the aliasing problem could arise for the Markov model.

Suppose y_t follows a two-state segmented trend model when measured over some time interval. Then suppose changes in y_t are measured only every k intervals. For example, if y_t follows a segmented trend on monthly data, and is measured quarterly, k = 3. On the longer intervals, y_t also follows a segmented trend, but with many more states. The number of states depends on how many sub-periods y_t was in state 1 (0, 1, ...,k), and the state in the first sub-period and last sub-period of the aggregated time interval. If y_t follows a two-state segmented trend model on monthly data, it will follow an eight-state segmented trend model on quarterly data.

We might nonetheless expect that in the case where the p_j are large that the choice of time interval will not make much difference in terms of identifying the current state, and in terms of out-of-sample forecasting. In this section, the same exchange rates over the same time periods examined in the previous sections are considered with monthly data.

In one way, the estimates on the monthly data do not seem too incompatible with the estimates on the quarterly data. We can reject $p_1 + p_2 = 1$ for many of the same currencies over the 1973-1986 and 1973-1988 estimation periods. This null is rejected now more generally for the pound against the European currencies, particularly in the longer sample period, where it is rejected for all of them.

The null hypothesis of $\mu_1 = \mu_2$ is rejected much less frequently on monthly data. Over both forecast intervals, this null is rejected for fewer than half the currencies. This suggests that the Markov model on monthly data may primarily be a model of heteroskedasticity.

The monthly model can be compared to the quarterly model by asking how similar their predictions are for the probability of being in state 1. Plots of the smoothed probabilities of being in state 1 (where the label for the state is chosen arbitrarily) provide some insight into the nature of the estimates of the two models.

Figures 1a, 2a and 3a plot the smoothed probabilities of $s_t = 1$ estimated on quarterly data through March 1988 for the dollar/mark, dollar/pound and dollar/French franc exchange rates. These estimates correspond to the ones reported in EH. The vertical lines in the charts represent quarters in which the probability switches from greater than .5 to less than .5, or vice-versa. Hence, the periods between vertical lines can be thought of as the estimated time periods in which the exchange rates are in state 1 or state 2.

Figures 1b, 2b and 3b plot the probabilities of $s_t = 1$ estimated on the monthly data over the same time period. (Note that in the quarterly data state 1 was normalized to be the one with the positive mean, while in the monthly data it was normalized to be the one with the negative mean.) The points of switching are remarkably similar for the monthly and quarterly data for the U.K. pound. However, there is much less correspondence for the mark and the franc. The pictures for the mark and franc are more typical than the one for the pound in representing the other fifteen exchange rates. In these charts, the mark and the franc switch states less frequently when estimated on monthly data than on quarterly data. There is, however, no such general tendency for the other exchange rates, over either period of estimation.

Another point of comparison is the out-of-sample forecasts. In particular, the monthly model was used to construct forecasts for the same time interval as those from the quarterly model. The forecasts from the monthly model were regressed on those from the quarterly model. We would expect that if the forecasts were close that the intercept in the regression would be close to zero, the slope coefficient would be close to 1, and the R^2 would be high. With only a few exceptions, the intercept term is found to be significantly different from zero, the slope is significantly different from one (and not significantly different from zero), and the R^2 is low. The forecasts are not very similar.

The models can be compared by the mean-squared-errors of out of sample forecasts over the same time interval. However, over the first estimation period, there were only five currencies for which the null hypotheses of $p_1 + p_2 = 1$ and $\mu_1 + \mu_2 = 1$ were rejected for both frequencies. Thus, there is a very limited set of comparisons we can make for cases in which the Markov model performs well in-sample. The quarterly model out-forecasts the monthly model for three of these five currencies. Over the second estimation period, there are only four currencies for which the segmented trend model performs well in-sample at both frequencies. The quarterly model out-forecasts the monthly model in two of these cases.

There seems to be little to gain from going to the monthly data. At the monthly frequencies, the estimated Markov model is not significantly different from a model of heteroskedasticity (that is, a model whose states are distinguished only by the variance of exchange rate changes). There is no tendency for forecasts of the monthly model to be superior to those of the quarterly model.

5. Conclusions

The results of this paper seem to indicate that the Markov switching model does not provide particularly good forecasts of exchange rates. The forecasts generated from this model do not generally have lower mean-squared error or mean-absolute-error than a random walk (with or without drift) or the forward exchange rate. There is weak evidence that the segmented trends model outperforms its rivals in predicting the direction of change of the exchange rate. All of this evidence comes in spite of the fact that the Markov model significantly outperforms generalized versions of the random walk in-sample.

The zero-drift random walk has done quite well in out-of-sample forecasting contests. It was the winner in Meese and Rogoff's (1983) contest between exchange rate models, and in Diebold and Nason's (1990) test of the forecasting ability of non-linear models. However, neither of those papers used the Diebold-Mariano statistic to test the significance of the difference in forecasting performance. The more recent study of Chinn and Meese (1992) finds that models that include an error-correction term significantly outperform the no-change hypothesis at longer horizons.[4]

If there are long swings in the exchange rate, then even if the parameter estimates of the Markov model are not very precise, it might be expected to perform well in terms of getting the direction of change correct. That is

[4] Mark (1992) finds similar results at long horizons, but does not include a measure of the significance of out-of-sample performance results.

because there will tend to be runs in one direction and then the other in changes in the exchange rate. Figures 1a, 2a, and 3a show that the runs do not have to persist for a long time before the Markov model concludes that the state has shifted. Hence, it will miss the direction of change for a short period of time around the date at which the regimes shift, but will tend to get the direction of change correct during long periods of time in which the exchange rate drifts in one direction.

Even though this paper examines many currencies, there is a sense in which the out-of-sample forecast performance of the models may be closely related. The Louvre accord of March 1987 seems to have stabilized all exchange rates. As it happens, this accord comes just one year after the end of our first estimation period (and one year before the end of our second estimation period.) If there was a regime switch at this point to a new state that is characterized by low variances and not much drift in exchange rates, the Markov model estimated during our sample would not perform well out of sample. While a change in regime is bad news for most models, it is exactly in the spirit of the Markov switching model. Perhaps the Markov model will perform better in the future, allowing for a third state.

References

Baillie, Richard T. and Tim Bollerslev, "The Message in Daily Exchange Rates: A Conditional-Variance Tale," Journal of Business and Economic Statistics.

Chinn, Menzie and Richard Meese, (1992). "Banking on Currency Forecasts: How Predictable is Change in Money?" working paper, University of California, Santa Cruz.

Diebold, Francis X. and Roberto S. Mariano, (1991). "Comparing Predictive Accuracy I: An Asymptotic Test," Institute for Empirical Macroeconomics, Federal Reserve Bank of Minneapolis, Discussion Paper no. 52.

Diebold, Francis X. and James A. Nason, (1990). "Nonparametric Exchange Rate Prediction," Journal of International Economics 28, 315-332.

Engel, Charles M. and James D. Hamilton, (1990). "Long Swings in the Dollar: Are They in the Data and Do Markets Know It?" American Economic Review 80, 689-713.

Hamilton, James D., (1988). "Rational Expectations Econometric Analysis of Changes in Regime: An Investigation of the Term Structure of Interest Rates," Journal of Economic Dynamics and Control 12, 385-423.

Hamilton, James D., (1989). "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," Econometrica 57, 357-384.

Hamilton, James D., (1990). "Analysis of Time Series Subject to Changes in Regime," Journal of Econometrics 45, 39-70.

Hamilton, James D., (1991a). "A Quasi-Bayesian Approach to Estimating Parameters for Mixtures of Normal Distributions," Journal of Business and Economic Statistics 9, 27-39.

Hamilton, James D., (1991b). "Estimation, Inference, and Forecasting of Time Series Subject to Changes in Regime," manuscript, University of Virginia.

Leitch, Gordon and J. Ernest Tanner, (1991). "Economic Forecast Evaluation: Profits versus the Conventional Error Measures," American Economic Review 81, 580-590.

Mark, Nelson C., (1992). "Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability and Overshooting," working paper, Ohio State University.

Meese, Richard and Kenneth Rogoff, (1983). "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?" Journal of International Economics 14, 3-24.

West, Kenneth D., Hali J. Edison and Dongchul Cho, (1992). "A Utility Based Comparison of Some Models of Exchange Rate Volatility," working paper, University of Wisconsin.

Table 1

Parameter Estimates for Segmented Trends Model

Quarterly, 1973:3-1986:1

(α=.1, β=.5, ν=.05)

Currency	μ_1	μ_2	p_1	p_2	σ_1^2	σ_2^2
USCA	-0.81297	-0.21505	0.710930	0.163457	5.33011	0.68901
	(0.40800)	(0.39907)	(0.200119)	(0.361071)	(1.50279)	(0.51157)
USFR	3.20287	-3.69747	0.821723	0.885774	10.9337	28.6117
	(1.01668)	(1.23946)	(0.120082)	(0.079598)	(4.07552)	(8.40175)
USIT	0.42973	-2.96907	0.866254	0.954205	3.12269	37.0035
	(0.54709)	(1.06553)	(0.108056)	(0.052007)	(1.26906)	(8.96167)
USJA	-1.46271	9.19055	0.893270	0.613107	14.3615	11.4970
	(0.90233)	(3.04496)	(0.064839)	(0.195890)	(3.97635)	(13.2517)
USSW	2.20500	-6.09130	0.894640	0.491165	45.5362	4.46869
	(1.25306)	(1.20831)	(0.075991)	(0.219126)	(10.2276)	(3.29118)
USUK	2.64354	-3.81335	0.892561	0.919611	10.6791	19.6765
	(0.82392)	(1.07884)	(0.085266)	(0.071844)	(3.54918)	(5.63793)
USWG	3.25496	-1.52475	0.859764	0.928004	15.8625	38.9234
	(1.23438)	(1.60705)	(0.136305)	(0.090800)	(7.39103)	(10.5356)
JACA	8.77097	-1.57612	0.704687	0.875286	13.5429	14.0444
	(1.60512)	(0.78973)	(0.150420)	(0.646017)	(7.47560)	(3.94569)
JAFR	3.58878	-5.51198	0.940906	0.731959	13.5845	4.52678
	(0.61267)	(0.73186)	(0.042641)	(0.151683)	(3.26396)	(2.12485)
JAIT	-5.30669	3.56015	0.733310	0.972229	2.79205	23.9144
	(0.93512)	(0.74361)	(0.214448)	(0.028113)	(2.00113)	(5.05191)
JASW	1.99968	-0.85551	0.960745	0.975514	4.80924	34.2684
	(0.59900)	(0.98371)	(0.052939)	(0.027275)	(1.99166)	(8.07182)
JAUK	-4.33880	4.03509	0.659601	0.886454	8.53950	19.1604
	(1.51176)	(1.37454)	(0.150418)	(0.095993)	(3.93052)	(6.39424)
JAWG	1.51739	-6.90914	0.974017	0.736699	21.4477	4.51320
	(0.70844)	(1.01764)	(0.027176)	(0.216006)	(4.68867)	(3.16459)
UKCA	2.52869	-6.85585	0.853455	0.693462	12.9848	6.49192
	(0.66632)	(0.70796)	(0.062965)	(0.125515)	(3.48845)	(2.59656)

Table 1
(continued)

Currency	μ_1	μ_2	p_1	p_2	σ_1^2	σ_2^2
UKFR	0.98397	-5.59244	0.819681	0	15.8641	26.5625
	(0.88491)	(8.02732)	(0.322459)	(---------)	(4.47681)	(35.4442)
UKIT	0.57283	1.91421	0.784616	0.276690	31.3830	2.61992
	(0.93841)	(0.95555)	(0.171082)	(0.322431)	(8.57979)	(2.37041)
UKSW	-2.63951	3.03195	0.846922	0	35.0052	1.98700
	(1.01336)	(1.28425)	(0.114527)	(---------)	(7.77181)	(2.85848)
UKWG	-1.44420	0.29711	0.864952	0.298741	28.9078	0.57115
	(0.83510)	(0.38260)	(0.081470)	(0.243325)	(6.54449)	(0.51855)

Table 2

Wald Tests of Markov Switching Model

(Chi-Square(1) Statistics)

Quarterly, 1973:3-1986:1

Currency / Null:	$p_1 + p_2 = 1$	$\mu_1 = \mu_2$
USCA	0.104	0.910
USFR	17.735	21.180
USIT	34.145	7.777
USJA	6.675	18.023
USSW	2.584	26.893
USUK	36.161	29.288
USWG	15.518	4.858
JACA	11.200	51.792
JAFR	16.135	100.724
JAIT	10.179	61.768
JASW	189.890	6.038
JAUK	8.993	49.467
JAWG	9.982	45.255
UKCA	13.910	106.691
UKFR	------	0.705
UKIT	0.026	0.867
UKSW	------	11.190
UKWG	0.378	3.461

Table 3

Tests of Significance of Drift Terms
Quarterly

73:3-86:1 vs 86:2-91:1

Currency	null: $\mu_1=0$	$\mu_1=\mu_2$, $\sigma_1^2 = \sigma_2^2$	null: $\mu_1=\mu_2$
USFR	-1.32	1.80	1.69
USIT	-2.59	4.18	3.98
USJA	0.91	0.09	0.07
USSW	0.79	0.15	0.13
USUK	-1.53	1.73	1.48
USWG	0.07	0.86	0.76
JACA	1.65	0.46	0.41
JAFR	2.56	1.52	1.49
JAIT	3.59	3.47	3.77
JASW	-0.06	0.03	0.03
JAUK	2.38	1.10	1.30
JAWG	0.94	0.65	0.66
UKCA	-0.58	0.05	0.05

(Statistics in the first column have t(50 d.f.) distributions, and those in the second two columns have χ^2(1 d.f.) distributions.)

Table 4

Post-Sample Mean Squared Forecast Errors

(Estimated quarterly, 1973:3-1986:1; Forecast 1986:2-1991:1)

Currency	Model	Horizon: 1	2	4
USFR	Random Walk	43.83	97.69	202.32
	RW - 0 Mean	40.39	81.08	138.01
	Forward Rate	41.45	86.14	154.01
	Markov Model	40.12	84.09	158.28
USIT	Random Walk	46.74	114.09	250.90
	RW - 0 Mean	38.14	76.78	112.76
	Forward Rate	42.16	92.27	170.44
	Markov Model	50.41	129.97	313.30
USJA	Random Walk	48.80	95.99	138.19
	RW - 0 Mean	50.04	100.47	149.64
	Forward Rate	49.81	101.06	159.64
	Markov Model	53.58	117.01	135.04
USSW	Random Walk	57.09	110.63	190.63
	RW - 0 Mean	58.83	118.76	217.92
	Forward Rate	58.81	120.24	222.17
	Markov Model	55.34	103.16	174.74
USUK	Random Walk	44.54	95.86	190.46
	RW - 0 Mean	41.46	81.70	128.23
	Forward Rate	45.24	97.45	189.13
	Markov Model	46.11	103.91	182.48
USWG	Random Walk	50.09	100.13	168.53
	RW - 0 Mean	50.27	100.94	171.42
	Forward Rate	48.99	98.13	162.07
	Markov Model	46.56	87.26	137.70
JACA	Random Walk	47.36	97.54	160.30
	RW - 0 Mean	46.17	91.70	120.27
	Forward Rate	47.43	98.58	167.62
	Markov Model	52.39	120.04	175.09
JAFR	Random Walk	28.68	82.03	248.20
	RW - 0 Mean	25.94	65.08	168.10
	Forward Rate	26.63	70.82	195.54
	Markov Model	24.96	77.78	275.82

Table 4
(continued)

Currency	Model : Horizon:	1	2	4
JAIT	Random Walk	32.49	96.88	305.37
	RW - 0 Mean	25.30	61.76	158.85
	Forward Rate	28.10	77.27	221.88
	Markov Model	34.77	102.29	323.13
JASW	Random Walk	23.17	57.91	162.68
	RW - 0 Mean	23.19	58.11	163.59
	Forward Rate	25.00	65.77	184.60
	Markov Model	22.86	57.97	190.80
JAUK	Random Walk	23.09	55.11	145.52
	RW - 0 Mean	21.06	42.79	70.72
	Forward Rate	23.74	56.86	139.54
	Markov Model	22.10	53.02	163.54
JAWG	Random Walk	24.85	65.03	180.53
	RW - 0 Mean	23.84	59.23	155.95
	Forward Rate	25.07	65.17	175.16
	Markov Model	26.79	73.99	213.97
UKCA	Random Walk	38.23	73.72	114.44
	RW - 0 Mean	38.12	72.60	109.12
	Forward Rate	38.59	75.24	128.54
	Markov Model	37.08	77.71	113.00

Table 4a

Diebold-Mariano Statistics for Comparison of Mean Squared Forecast Errors

(Estimated quarterly, 1973:3-1986:1; Forecast 1986:2-1991:1)

Currency	Model : Horizon:	1	2	4
USFR	RW vs. RW-0 mean	1.11	1.70	1.77
	Markov vs. RW	-0.68	-0.92	-0.92
	Markov vs. RW-0 mean	-0.06	0.23	0.43
	Markov vs. Forward Rate	-0.15	-0.15	0.08
USIT	RW vs. RW-0 mean	1.55	2.14	2.30
	Markov vs. RW	2.27	2.51	3.58
	Markov vs. RW-0 mean	1.79	2.28	2.59
	Markov vs. Forward Rate	1.88	2.65	3.16
USJA	RW vs. RW-0 mean	-0.53	-0.60	-0.40
	Markov vs. RW	0.86	1.41	-0.18
	Markov vs. RW-0 mean	0.61	0.94	-0.35
	Markov vs. Forward Rate	0.66	1.04	-1.18
USSW	RW vs. RW-0 mean	-0.65	-0.92	-0.79
	Markov vs. RW	-0.74	-2.47	-3.35
	Markov vs. RW-0 mean	-0.83	-1.46	-1.23
	Markov vs. Forward Rate	-0.91	-2.68	-4.47
USUK	RW vs. RW-0 mean	0.95	1.39	1.72
	Markov vs. RW	0.22	0.41	-0.20
	Markov vs. RW-0 mean	0.86	1.74	1.62
	Markov vs. Forward Rate	0.62	0.37	-0.22
USWG	RW vs. RW-0 mean	-1.00	-1.47	-1.43
	Markov vs. RW	-0.86	-1.26	-0.88
	Markov vs. RW-0 mean	-0.88	-1.31	-0.94
	Markov vs. Forward Rate	-0.76	-1.11	-0.76
JACA	RW vs. RW-0 mean	0.27	0.41	0.84
	Markov vs. RW	0.68	1.13	0.54
	Markov vs. RW-0 mean	0.78	1.16	1.52
	Markov vs. Forward Rate	0.68	0.98	0.22
JAFR	RW vs. RW-0 mean	0.65	0.97	1.04
	Markov vs. RW	-0.65	-0.31	1.06
	Markov vs. RW-0 mean	-0.18	0.93	1.34
	Markov vs. Forward Rate	-0.35	0.63	1.94

Table 4a
(continued)

Currency	Model : Horizon:	1	2	4
JAIT	RW vs. RW-0 mean	1.14	1.36	1.27
	Markov vs. RW	2.01	1.20	1.44
	Markov vs. RW-0 mean	1.33	1.57	1.42
	Markov vs. Forward Rate	1.55	2.24	2.24
JASW	RW vs. RW-0 mean	-0.25	-0.46	-0.46
	Markov vs. RW	-0.13	0.01	0.82
	Markov vs. RW-0 mean	-0.14	-0.02	0.82
	Markov vs. Forward Rate	-0.79	-0.68	0.18
JAUK	RW vs. RW-0 mean	0.52	0.90	1.66
	Markov vs. RW	-0.37	-0.37	1.67
	Markov vs. RW-0 mean	0.23	0.71	1.83
	Markov vs. Forward Rate	0.09	-0.53	2.20
JAWG	RW vs. RW-0 mean	0.67	0.93	0.92
	Markov vs. RW	1.16	1.72	1.75
	Markov vs. RW-0 mean	1.08	1.50	1.39
	Markov vs. Forward Rate	0.70	1.14	1.21
UKCA	RW vs. RW-0 mean	0.09	0.27	0.36
	Markov vs. RW	-0.24	0.42	-0.10
	Markov vs. RW-0 mean	-0.22	0.75	0.21
	Markov vs. Forward Rate	-0.48	0.29	-1.02

Table 5

Post-Sample Mean Absolute Forecast Errors

(Estimated quarterly, 1973:3-1986:1; Forecast 1986:2-1991:1)

Currency	Model : Horizon:	1	2	4
USFR	Random Walk	5.54	8.92	12.05
	RW - 0 Mean	5.26	7.87	10.05
	Forward Rate	5.43	8.16	10.34
	Markov Model	4.77	8.19	10.26
USIT	Random Walk	5.89	9.59	13.43
	RW - 0 Mean	5.29	7.64	8.92
	Forward Rate	5.66	8.58	10.64
	Markov Model	6.07	10.25	15.06
USJA	Random Walk	5.65	8.24	9.93
	RW - 0 Mean	5.63	8.32	10.20
	Forward Rate	5.74	8.45	10.74
	Markov Model	6.04	9.15	9.44
USSW	Random Walk	6.09	9.22	11.72
	RW - 0 Mean	6.25	9.63	12.38
	Forward Rate	6.35	9.71	12.96
	Markov Model	5.90	8.70	11.16
USUK	Random Walk	5.57	8.29	11.43
	RW - 0 Mean	5.19	7.64	9.76
	Forward Rate	5.64	8.39	11.48
	Markov Model	5.56	9.30	10.77
USWG	Random Walk	5.88	8.75	10.63
	RW - 0 Mean	5.88	8.79	10.70
	Forward Rate	5.94	8.65	10.74
	Markov Model	5.58	8.28	9.88
JACA	Random Walk	5.52	8.19	10.56
	RW - 0 Mean	5.50	7.77	8.99
	Forward Rate	5.54	8.20	10.85
	Markov Model	5.69	9.04	10.78
JAFR	Random Walk	3.61	6.22	11.45
	RW - 0 Mean	3.75	6.45	10.52
	Forward Rate	3.50	5.96	10.48
	Markov Model	3.76	6.34	12.44

Table 5
(continued)

Currency	Model : Horizon:	1	2	4
JAIT	Random Walk	4.01	6.81	13.22
	RW - 0 Mean	3.91	6.48	10.28
	Forward Rate	3.85	6.36	11.53
	Markov Model	4.20	7.32	13.99
JASW	Random Walk	3.47	5.99	10.54
	RW - 0 Mean	3.46	5.99	10.55
	Forward Rate	3.65	6.34	11.01
	Markov Model	3.38	6.10	11.76
JAUK	Random Walk	3.62	5.42	9.31
	RW - 0 Mean	3.15	4.39	6.51
	Forward Rate	3.66	5.41	8.36
	Markov Model	3.68	5.35	9.94
JAWG	Random Walk	3.43	5.72	10.78
	RW - 0 Mean	3.62	5.85	10.31
	Forward Rate	3.68	6.07	10.79
	Markov Model	3.58	6.09	11.71
UKCA	Random Walk	5.12	7.14	9.16
	RW - 0 Mean	5.08	7.00	9.06
	Forward Rate	5.20	7.28	9.85
	Markov Model	5.17	7.29	8.79

Table 5a

Diebold-Mariano Statistics for Comparison of Mean Absolute Forecast Errors

(Estimated quarterly, 1973:3-1986:1; Forecast 1986:2-1991:1)

Currency	Model : Horizon:	1	2	4
USFR	RW vs. RW-0 mean	1.18	2.42	1.43
	Markov vs. RW	-1.74	-0.83	-1.10
	Markov vs. RW-0 mean	-1.23	0.41	0.12
	Markov vs. Forward Rate	-1.37	0.03	-0.05
USIT	RW vs. RW-0 mean	1.38	2.44	2.06
	Markov vs. RW	1.28	2.34	3.20
	Markov vs. RW-0 mean	1.43	2.51	2.28
	Markov vs. Forward Rate	1.35	2.52	3.06
USJA	RW vs. RW-0 mean	0.12	-0.20	-0.22
	Markov vs. RW	0.78	1.34	-0.54
	Markov vs. RW-0 mean	0.73	1.05	-0.42
	Markov vs. Forward Rate	0.72	0.92	-1.21
USSW	RW vs. RW-0 mean	-0.89	-0.99	-0.52
	Markov vs. RW	-1.15	-2.81	-2.30
	Markov vs. RW-0 mean	-1.29	-2.09	-0.98
	Markov vs. Forward Rate	-1.57	-2.87	-4.51
USUK	RW vs. RW-0 mean	1.67	1.23	1.05
	Markov vs. RW	-0.22	1.05	-0.34
	Markov vs. RW-0 mean	0.99	2.28	0.77
	Markov vs. Forward Rate	0.60	0.99	-0.48
USWG	RW vs. RW-0 mean	-0.44	-1.62	-0.96
	Markov vs. RW	-1.02	-0.88	-0.59
	Markov vs. RW-0 mean	-1.02	-0.94	-0.63
	Markov vs. Forward Rate	-1.09	-0.73	-0.63
JACA	RW vs. RW-0 mean	0.05	0.65	0.89
	Markov vs. RW	0.32	0.95	0.21
	Markov vs. RW-0 mean	0.31	1.37	1.49
	Markov vs. Forward Rate	0.21	0.87	-0.08
JAFR	RW vs. RW-0 mean	-0.40	-0.24	0.36
	Markov vs. RW	0.34	0.16	0.89
	Markov vs. RW-0 mean	0.01	-0.09	0.65
	Markov vs. Forward Rate	0.38	0.42	1.10

Table 5a
(continued)

Currency	Model : Horizon:	1	2	4
JAIT	RW vs. RW-0 mean	0.18	0.24	0.79
	Markov vs. RW	1.19	1.73	1.84
	Markov vs. RW-0 mean	0.48	0.61	0.99
	Markov vs. Forward Rate	0.78	1.53	1.61
JASW	RW vs. RW-0 mean	0.89	0.17	-0.16
	Markov vs. RW	-0.37	0.17	0.91
	Markov vs. RW-0 mean	-0.33	0.18	0.93
	Markov vs. Forward Rate	-0.51	-0.32	0.51
JAUK	RW vs. RW-0 mean	1.27	1.40	1.44
	Markov vs. RW	0.20	-0.18	1.45
	Markov vs. RW-0 mean	1.20	1.08	2.01
	Markov vs. Forward Rate	0.88	-0.11	2.48
JAWG	RW vs. RW-0 mean	-1.26	-0.34	0.48
	Markov vs. RW	0.64	0.85	1.23
	Markov vs. RW-0 mean	-0.14	0.34	0.87
	Markov vs. Forward Rate	0.18	0.03	0.61
UKCA	RW vs. RW-0 mean	0.44	0.60	0.14
	Markov vs. RW	0.10	0.24	-0.53
	Markov vs. RW-0 mean	0.22	0.56	-0.30
	Markov vs. Forward Rate	-0.32	0.01	-1.58

Table 6
Was Forecast Change the Correct Direction?

(Estimated quarterly, 1973:3-1986:1; Forecast 1986:2-1991:1)

Count of correct forecast changes

Currency	Model : Horizon (# of forecast periods)	1 (20)	2 (19)	4 (17)
USFR	Random Walk	8	5	6
	Forward Rate	6	6	8
	Markov Model	13	10	10
USIT	Random Walk	7	5	6
	Forward Rate	7	5	6
	Markov Model	7	5	6
USJA	Random Walk	10	10	10
	Forward Rate	10	11	10
	Markov Model	9	9	10
USSW	Random Walk	12	13	11
	Forward Rate	10	11	9
	Markov Model	13	13	11
USUK	Random Walk	7	7	6
	Forward Rate	7	7	6
	Markov Model	11	7	9
USWG	Random Walk	11	13	12
	Forward Rate	12	12	11
	Markov Model	12	12	9
JACA	Random Walk	10	9	7
	Forward Rate	10	9	7
	Markov Model	11	9	7
JAFR	Random Walk	13	12	9
	Forward Rate	13	12	8
	Markov Model	13	11	8

Table 6
(continued)

Currency	Model : Horizon (# of forecast periods)	1 (20)	2 (19)	4 (17)
JAIT	Random Walk	13	12	8
	Forward Rate	13	12	8
	Markov Model	13	12	8
JASW	Random Walk	8	9	9
	Forward Rate	5	3	4
	Markov Model	12	9	8
JAUK	Random Walk	11	11	11
	Forward Rate	11	11	11
	Markov Model	11	9	11
JAWG	Random Walk	13	12	7
	Forward Rate	9	7	5
	Markov Model	11	11	6
UKCA	Random Walk	9	8	8
	Forward Rate	8	8	5
	Markov Model	11	11	11

FIGURE 1A, DOLLAR/FRANC, QUARTERLY

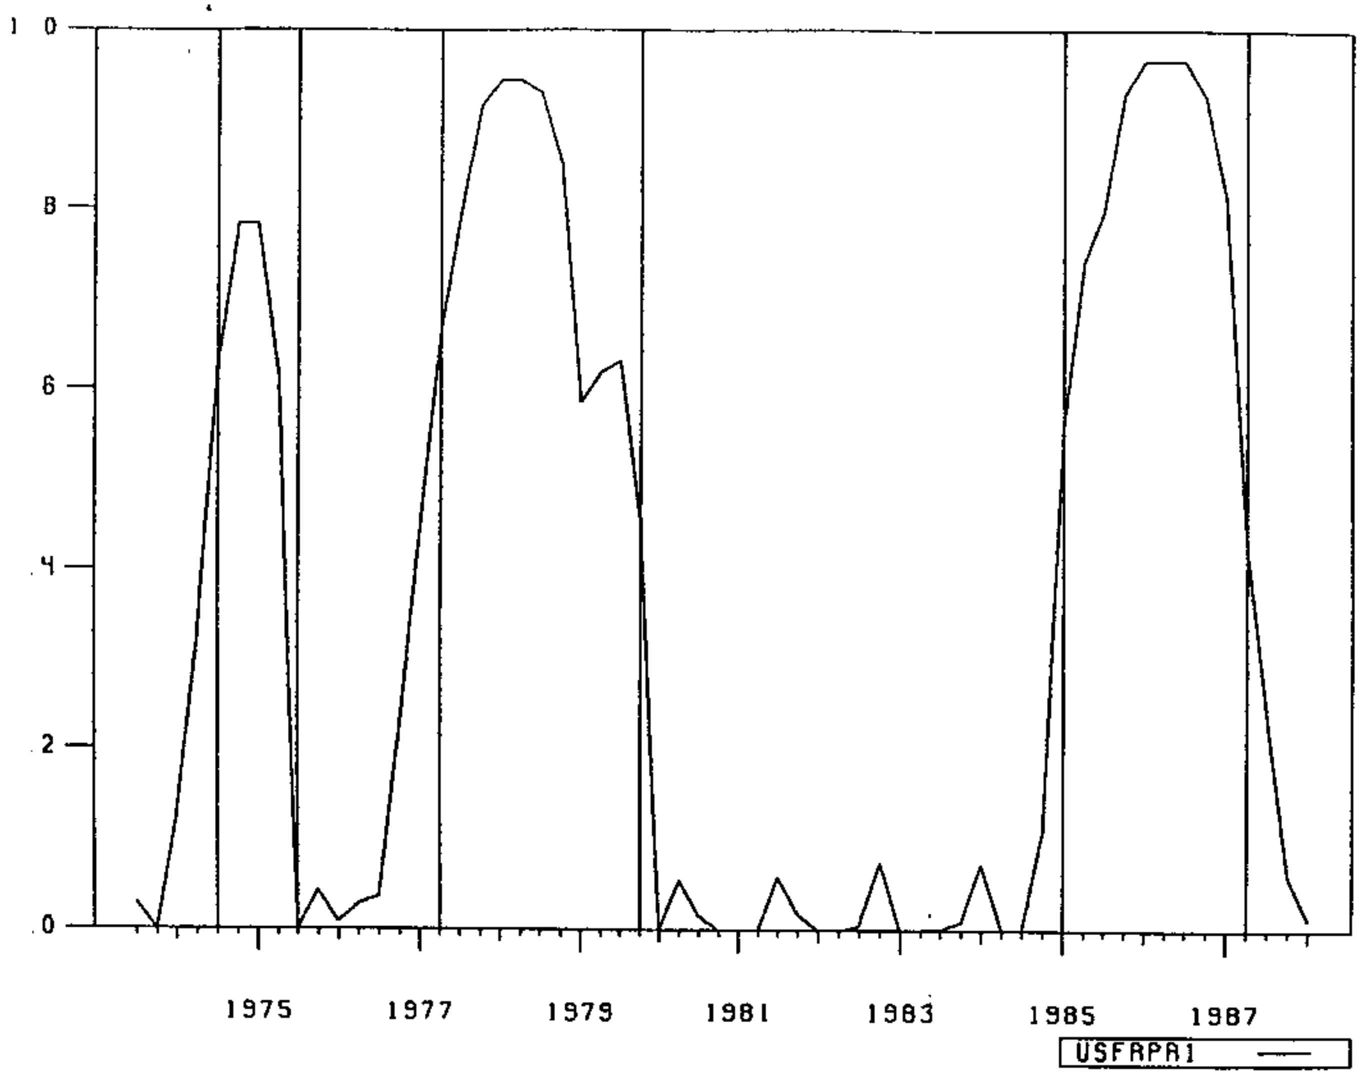

FIGURE 1B, DOLLAR/FRANC, MONTHLY

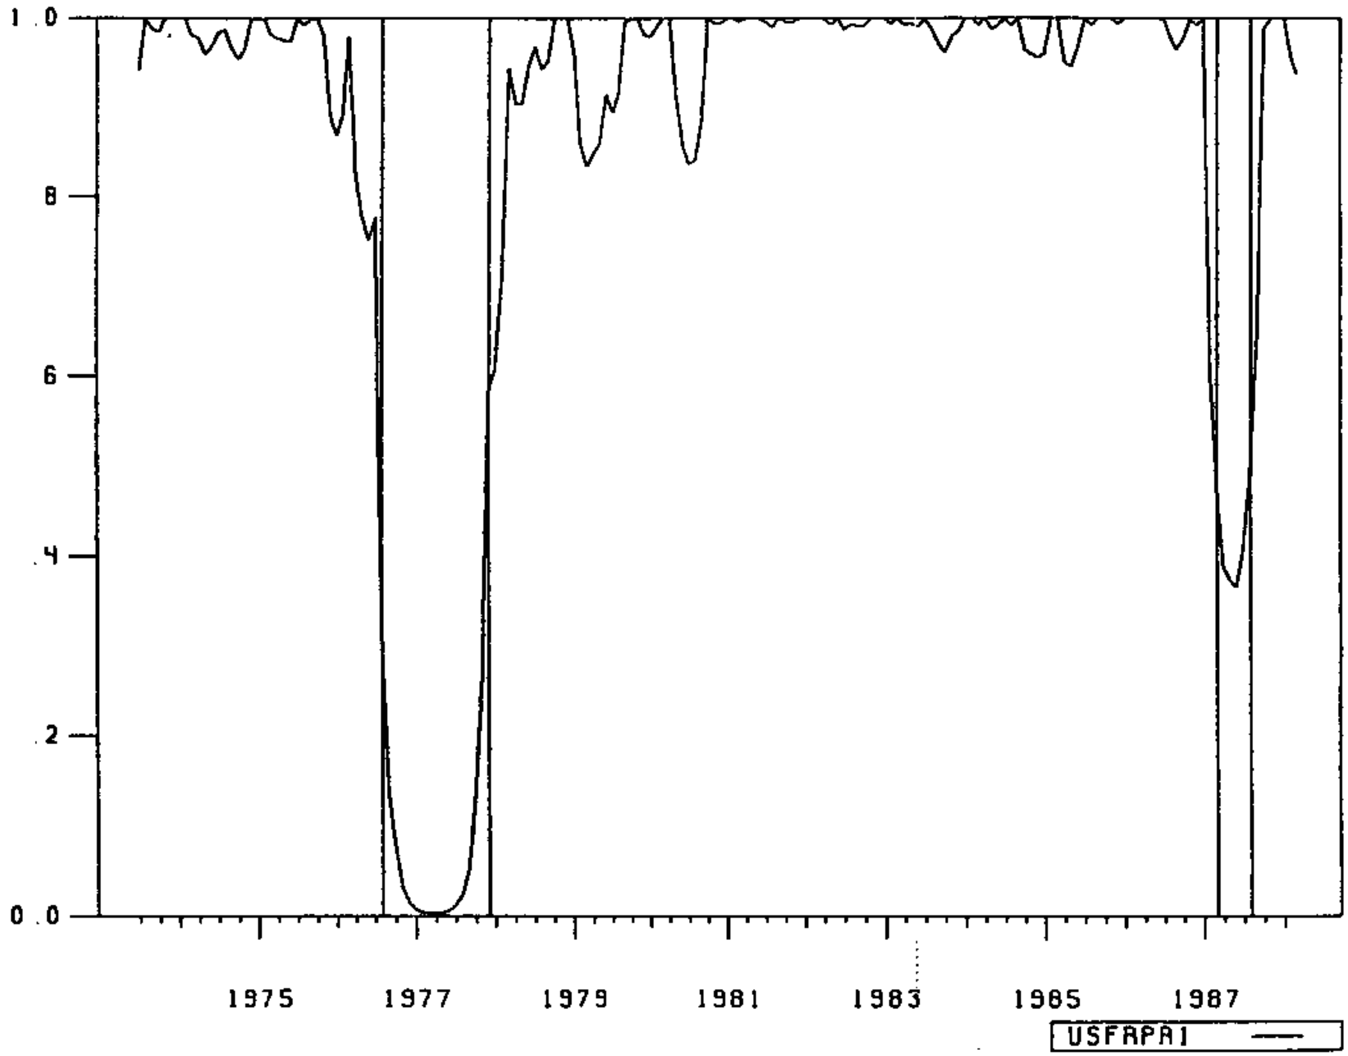

FIGURE 2A, DOLLAR/POUND, QUARTERLY

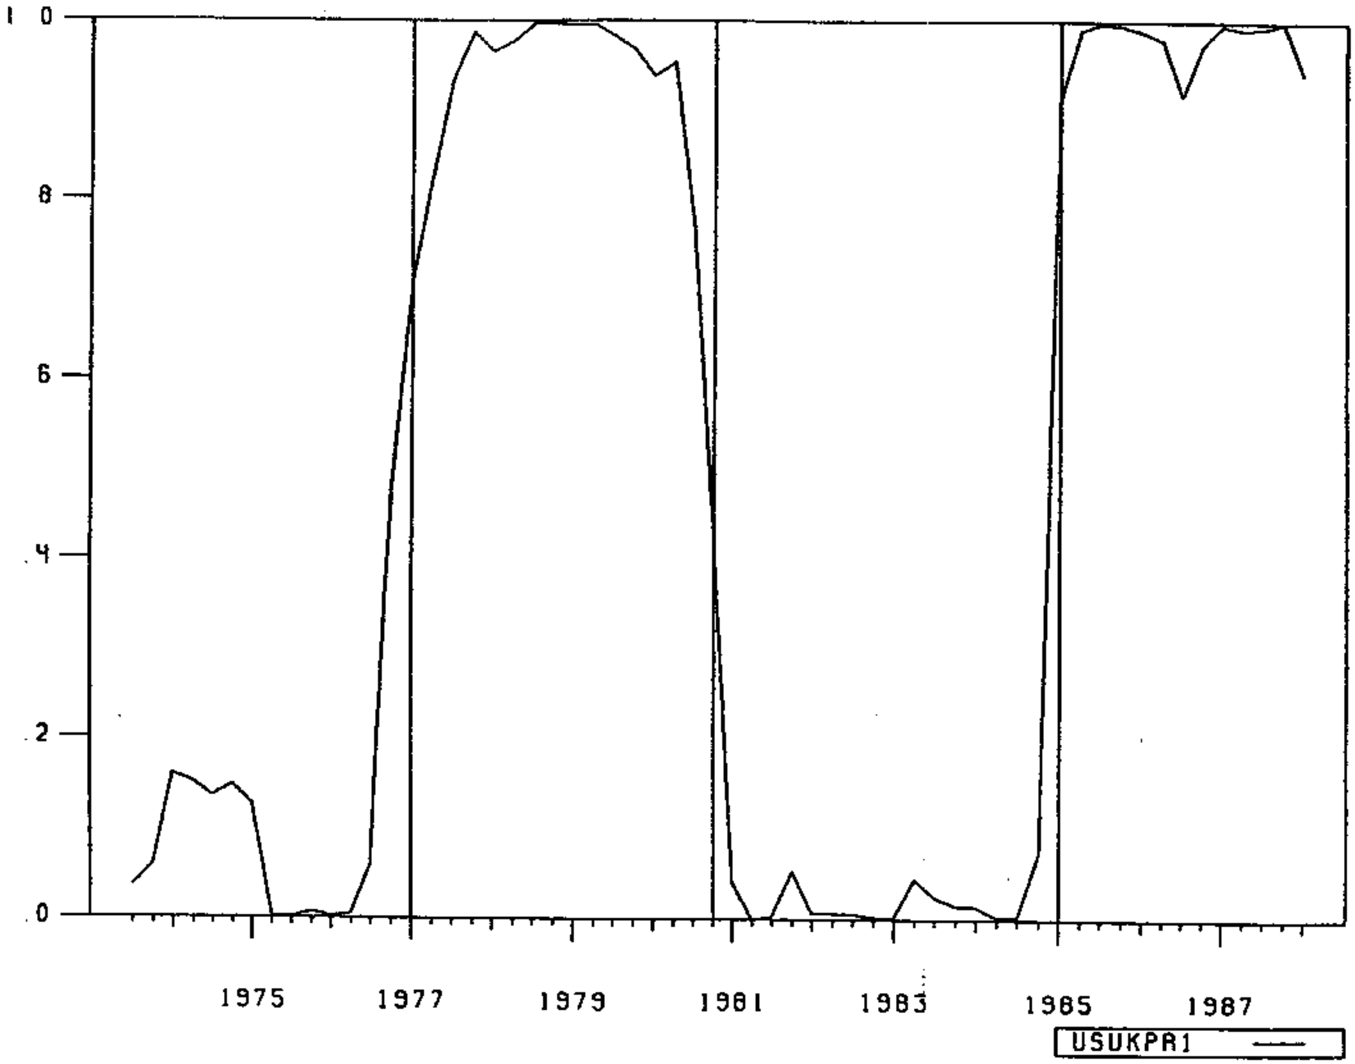

FIGURE 2B, DOLLAR/POUND, MONTHLY

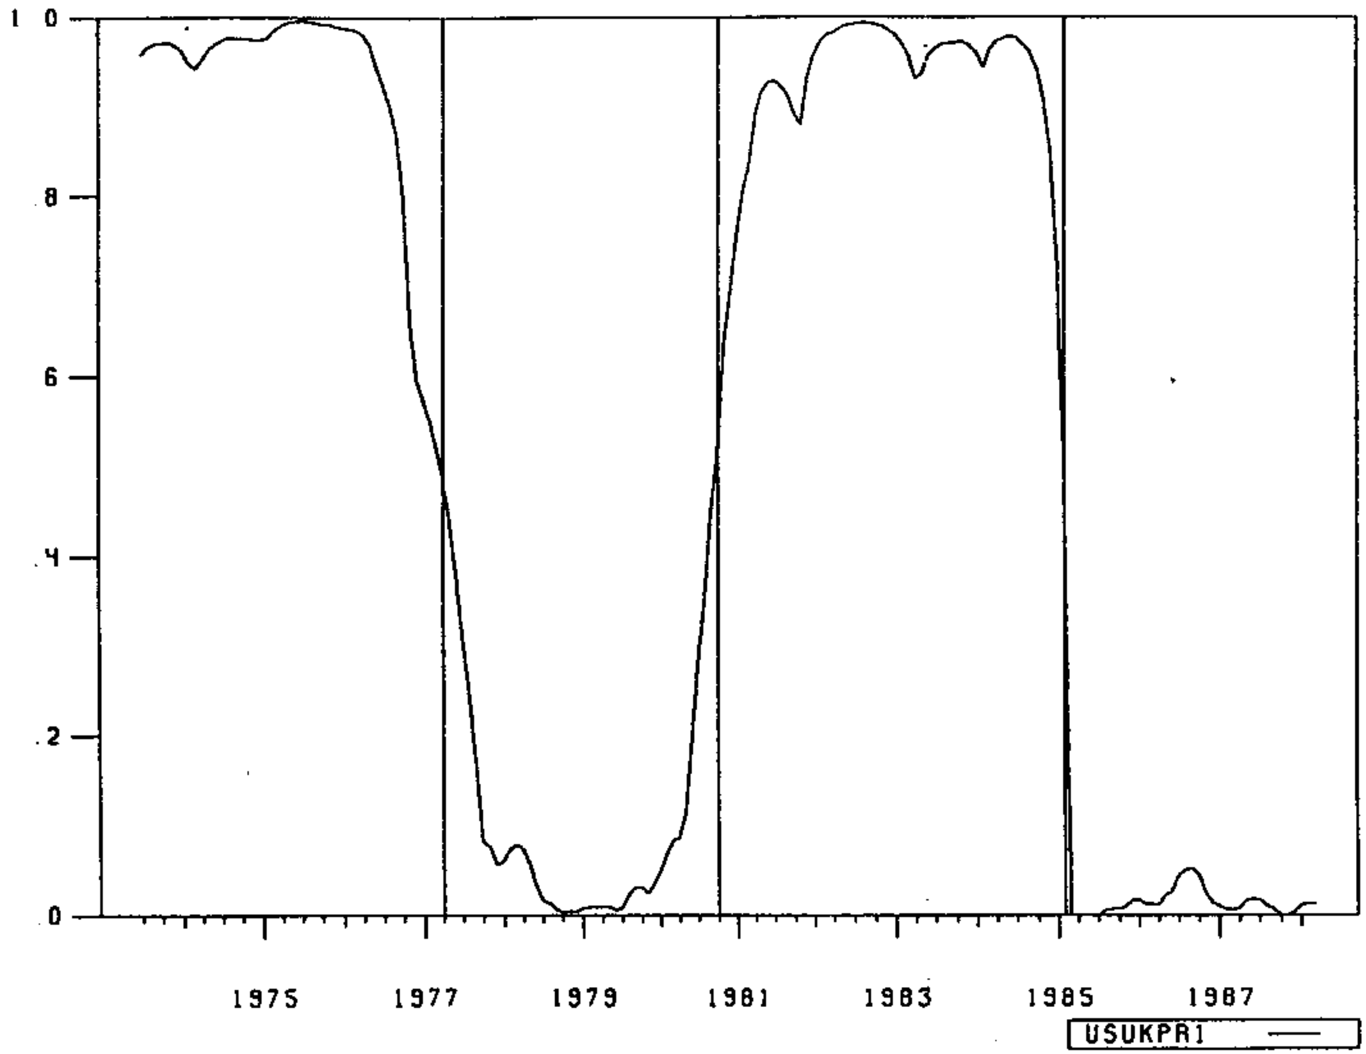

FIGURE 3A, DOLLAR/MARK, QUARTERLY

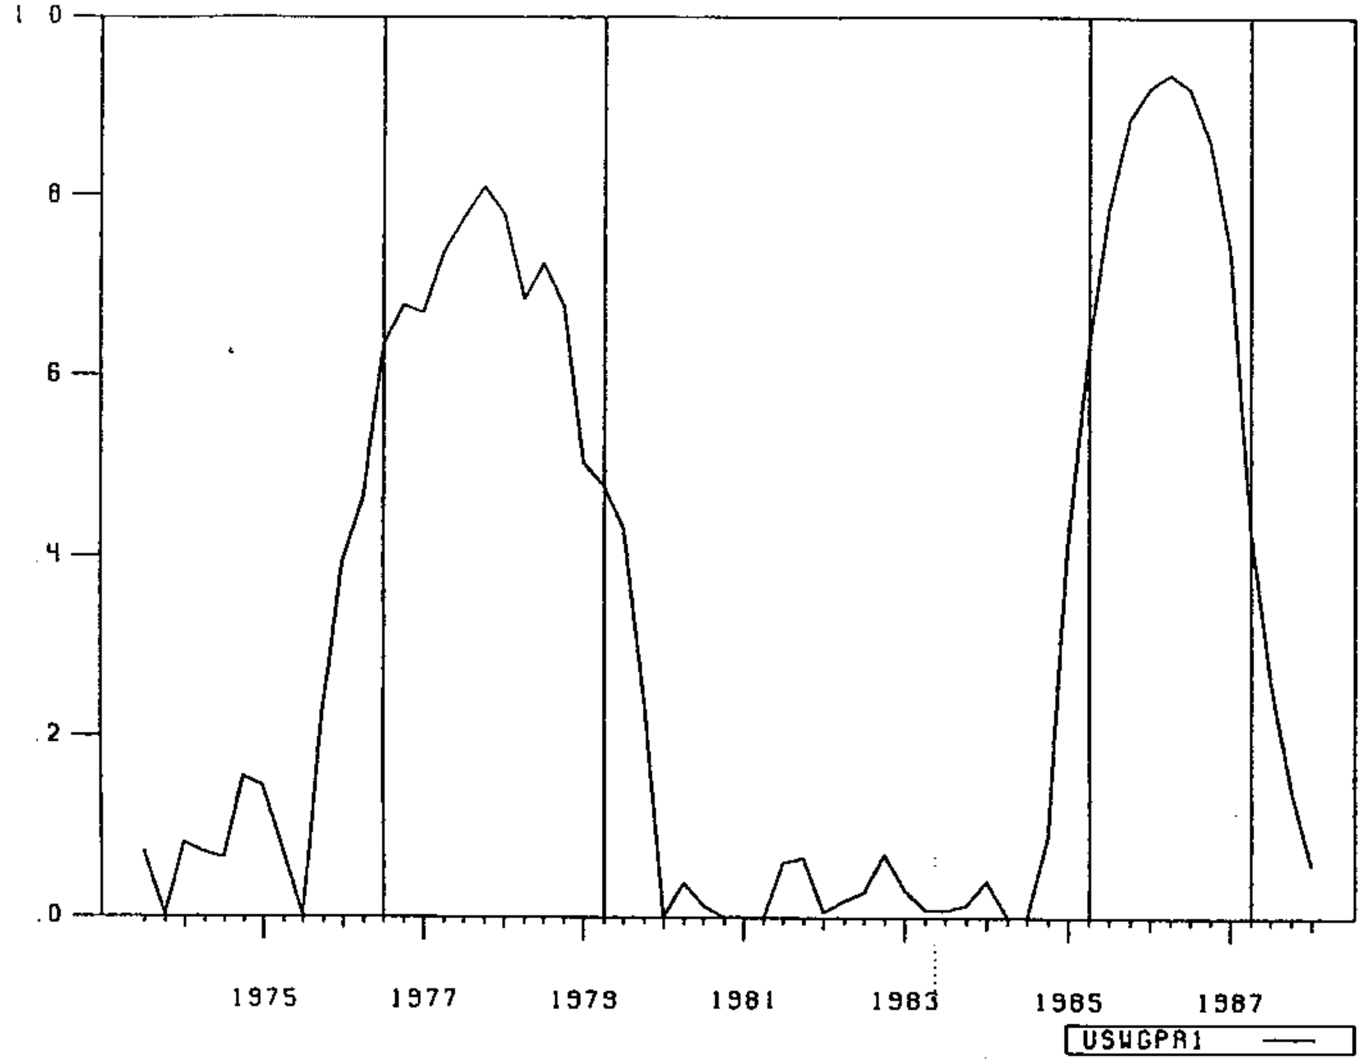

FIGURE 3B, DOLLAR/MARK, MONTHLY