CHAPTER 16

Forecasting realignments: The case of the French franc in the exchange-rate mechanism

Bruce Mizrach

I utilize a new model of Markov switching with probit transitions to characterize exchange-rate mechanism. For the French franc, I show that the probit model accurately anticipates realignments of the central parity. Used as a risk measure, the probit model also is a good proxy for volatility. When risk is high, we make less accurate forecasts. Adopting a rule of thumb, not to predict when the probit risk estimate is more than two standard deviations above the mean, we can dramatically improve forecast accuracy.

Our knowledge concerning exchange rates does not seem proportional to the effort. Despite being one of the most widely studied financial variables, very little is known about the behavior of the spot exchange rate. Foreign exchange rates seem to defy simple equilibrium relationships such as interest parity.[1] Atheoretical models have fared no better. Once estimated, the data quickly expose models to be nothing more than curve-fitting exercises.[2] Nonlinear modeling has emerged as a possible solution to this conundrum.

One branch of this new literature has adopted empirical models for the higher moments of exchange rates. Exchange-rate returns are generally nonnormal (in the Gaussian sense).[3] Leptokurtic (fat-tailed) distributions, with clustering of the large errors, has motivated the application of the generalized autoregressive conditional heteroscedasticity (GARCH) model of Engle (1982) and Bollerslev (1986) to exchange rates. The GARCH models have faltered largely on two grounds. A compelling theoretical explanation for the volatility clustering is

An earlier version of this paper was presented at the American Statistical Association Meetings in San Francisco, August 8–12, 1993. I would like to thank Doug Patterson for helpful comments.

[1] See, e.g., Froot and Thaler (1990) and the resolution proposed in Mizrach (1995c).

[2] An exhaustive survey of models from the 1970s can be found in Meese and Rogoff (1983).

[3] See, e.g., Hsieh (1988).

needed.[4] The GARCH effects seem also not to be present in the mean, and therefore are of little use for forecasting.

Nonlinear modeling also has been motivated by institutional changes. Since March 1979, nearly all major European currencies have been traded within target zones. Krugman (1991) has shown that target zones will introduce nonlinearities into exchange rates even if the bands never change.

Several researchers have tried to make sense of the nonlinearities using time-series techniques. Engel and Hamilton (1990) used a Markov-switching model. Meese and Rose (1990, 1991), Diebold and Nason (1990), and Mizrach (1992) used nonparametric techniques. None of these papers has produced robust out-of-sample forecast improvements. The random walk is often a better predictor than sophisticated nonlinear technology.

This paper tries to get a better understanding of the failures of the nonlinear models. I show that the spot exchange rate is indeed predictable most of the time. At a few crucial junctures, coinciding with exchange-rate realignments, all models, linear and nonlinear, forecast poorly. I show that a new class of model introduced previously (Mizrach 1995a) can help us understand when forecasting is likely to be difficult. The Mizrach model produces a daily time series of the probability of realignment. Adopting a rule of thumb, not to forecast when the risk is more than two standard deviations above average, we can dramatically improve our overall accuracy.

The paper is organized as follows. I begin with the naive benchmark for the exchange rate, the unit root. I then turn to linear and nonlinear time-series models, including Mizrach's probit-Markov model. Using the probit's risk estimates, I repeat the modeling exercise using the rule of thumb. A summary and conclusion follow.

1 The random walk

This section is devoted to the benchmark model for all exchange-rate forecasting exercises – the random walk. I begin by describing the data used in the analysis of the unit root.

1.1 *Data*

For purposes of empirical illustration, I look at the French franc, German deutsche mark (Fr/DM) exchange rate over the period March, 13, 1979, to September, 11, 1992. This spans the creation of the exchange-rate mechanism (ERM) in Europe to the suspension of the Italian lira and the British pound from the ERM. I have 3,417 daily observations on European Currency Unit exchange rates from the 14:30 fix in Basle, which I convert into a cross rate.

[4] See Mizrach (1996) for some work along these lines.

Table 16.1. *Unit-root tests*

	Mar. 13, 1979 to Sept. 11, 1992	No. of observations
Coefficient	−5.002E-04	3,416
t-Statistic	(2.06)	

Table 16.2. *Descriptive statistics*

Δs_t	Full sample	Omit regime changes
Mean	11.4573-05	5.935E-05
SD	1.997E-03	1.253E-03
Kurtosis	324.694	26.916

ERM exchange rates float in 2.25 percent bands around a central parity. The central parity is occasionally realigned. With the franc, the parity has been rest on seven occasions.[5]

1.2 *Unit-root tests*

Let s_t denote the log of the spot exchange rate, and let Δs_t be the log difference. I regress the difference on the lag of the level. A coefficient significantly less than zero would reject the null hypothesis of a unit root.

The relevant critical values are the Dickey-Fuller statistics, not the usual t-statistics.[6] In a large sample, a 5 percent one-sided critical value is 2.86, and a 1 percent critical value is 3.43. Despite a t-ratio of 2.06, the coefficient is not significantly different from zero. This evidence supports the conventional wisdom that it is hard to reject the random walk as a statistical description for the spot exchange rate.

1.3 *Analysis of unconditional moments across regimes*

I next looked at the unconditional moments of the first differences of the Fr/DM exchange rate. The daily movements are quite small, but rates can still be quite volatile. The kurtosis for the sample as a whole is over 300.

I wanted to examine whether the regime changes were influencing these results. I omitted 66 influential data points, five days around either side of a

[5] On August, 4, 1993, the franc passed below its old ERM floor of 3.4305 Fr/DM. Bands of 15 percent have been introduced around the old parity, so technically, the franc has not devalued.

[6] For more discussion on this issue, see Mizrach (1993) in which the critical values is bootstrapped because of the nonnormality and serial correlation in the data. In an experimental design a bit different from this one, the appropriate critical values are found to be 3.43 and 5.06.

devaluation. These results appear in the second column of Table 16.2. The daily changes are much smaller because I have removed the large devaluations. Even more notable is that the kurtosis falls by a factor of more than 10, to 26.916.

2 Alternative models for the exchange rate

I turn next to several different approaches to model the exchange rate. I then analyze the residuals to see which, if any, of the models does the best job of depicting the important devaluation episodes.

2.1 *Linear time series*

The most straightforward time-series approach is a Box-Jenkins model. Using the Akaike information criterion, I fit an AR(3) model to the first differences.

2.2 *Near-neighbor models*

Nonparametric approaches are appealing because they can provide meaningful statistical inference when very little is known about the series' fundamentals or distribution. Recent efforts at nonparametric modeling include Meese and Rose (1990, 1991), Diebold and Nason (1990), and Mizrach (1992).

A technique that is amenable to our application is nearest-neighbor methods. The idea is to find neighbors near the current realization of the independent variables. With locally weighted regression, one then fits a regression surface to the neighboring dependent variables. I selected a model with five neighbors and used least-squares weights to estimate the exchange-rate changes. I denote this model as 5-NN.

3 Analysis of the residuals

In this section, I analyze the residuals of the two models to evaluate their performance in explaining the critical devaluation episodes.

The first column of Table 16.3 lists the dates of the six devaluation episodes. The next column is the percentage of the variance in the raw data due to the 66 realignment observations. Under the columns for the two models, AR(3) and 5-NN, are the percentages of the sum of squared residuals for the same dates. Even though they comprise only 66 observations out of 3,416, or less than 2 percent of the total, the six devaluation episodes explain 61 percent of the variance.

The AR(3) model does not explain these sudden devaluations. In the residuals of the AR(3) model, 58 percent of the sum of squared residuals is due to the six episodes.

Table 16.3. *Analysis of variance*

		Percentage of sum of squared residuals	
Date	%Variance	AR(3)	5-NN
9/17-10/1/179	0.37	0.39	0.20
9/24-10/8/81	13.81	14.35	13.93
6/7-6/21/82	25.24	25.16	24.79
3/15-3/30/83	12.54	8.73	12.59
3/24-4/10/86	8.63	8.81	8.49
1/5-1/19/87	0.83	0.83	0.73

The NN-model leaves just as much information behind in the residuals as do the linear AR models. Sixty-one percent of the sum of squared residuals is in these devaluation episodes.

It seems that if we are to make much progress, we need to uncover something that helps us predict realignments.

4 A probit-Markov model of devaluation risk

In Mizrach (1995a), I introduce a new type of Markov-switching model. Unlike conventional switching models, the probability of a change in regime varies smoothly throughout the sample.

The model has two parts. In the first part, as with conventional switching models, one specifies models for the conditional mean in both regimes. In our exchange-rate context, they are the within band and devaluation regimes. Within the band, the exchange rate is mean reverting. Outside of the band, I find that devaluations are proportional to the cumulative departure from purchasing-power parity.

I describe the probit part of the model in greater detail because that is what I will use in this section. I link the devaluation risk to a constant term, z_{1t}, and two state variables. The first is the position of the exchange rate within the band. Define

$$z_{2t} = (s_t - \underline{s})/(\bar{s} - \underline{s}), \tag{16.1}$$

where $[\underline{s}, \bar{s}]$ are the lower and upper bounds of the target zone.

The second variable is based on the yield curve. During several devaluation crises, the term structure has become steeply negatively sloped. For example, on March, 15, 1983, five days prior to a realignment of the franc, the French 3-month $i_i^{3/12}$, 1-month $i_i^{1/12}$ spread, denoted here as

$$z_{3t} = \log(1 + i_t^{3/12}) - \log(1 + i_t^{1/12}) \tag{16.2}$$

was -46.00. I also add a constant term, defining $z_t = (1, z_{2t}, z_{3t})$.

Table 16.4. *Risk estimates prior to realignment*

Date	Average risk	Peak risk
Sept. 24, 1979	8.56	9.11
Oct. 5, 1981	32.11	37.80
June 14, 1982	35.31	43.70
March 21, 1983	99.99	99.99
April 7, 1986	19.67	25.91
Jan. 12, 1987	14.72	19.75

To ensure that the risk remains on [0, 1], I make a probit transformation

$$p_t = \int_{-\infty}^{\gamma z_t} \left(\sqrt{2\pi}\right)^{-1} \exp(-t^2/2)dt \equiv \Phi(\gamma z_t). \tag{16.3}$$

In a fully specified model for the French-German interest differential, I obtain implicit market estimates of the potential devaluation risk. I find $\hat{\gamma} = (-1.563, 0.367, -17.177)$. I then compute a risk measure series, $\hat{p}_t$.

In Table 16.4, I look at the risk just prior to realignment. The first column contains the average risk in the five days prior to realignment. In the second column, I have the peak risk, which is almost always the day before the devaluation.

These risks should be compared relative to a mean risk of devaluation of 8.3 percent with a standard deviation of 6.2 percent. In five of the seven realignments, a risk two standard deviations above the mean (20.7 percent) was observed prior to a devaluation.[7]

Now we'll see whether this risk model can be useful in fitting the exchange-rate data.

5 Model evaluation

In regression exercises, I discovered that the risk model was not very precise in detecting the exact day of realignment. If you predicted a large change in the spot rate every day in which the risk was significantly above its mean, you would forecast very poorly. I chose instead to look at a rule of thumb where the risk measure provided information on when *not* to forecast.

In Table 16.5, I look at the sample mean squared errors (MSE) for the linear and nonparametric models. The random walk (a no-change forecast) is included as a benchmark. Note that the MSE for the nonlinear 5-NN model is less than 0.2 percent better than the random walk.

[7] On July, 30, 1993, prior to the widening of the bands, the franc's risk estimate was at 24.5%.

Table 16.5. *Model evaluation*

Model	MSE	M-stat
Random walk	4.01E-06	
5-NN	3.97E-06	0.03
AR(3)	3.95E-06	1.16
Risk AR(3)	1.37E-06	2.14

To make a formal comparison, I use the robust forecast comparison introduced previously (Mizrach 1995b), which I designate the M-stat in the table. This statistic has very weak population assumptions that can readily handle the kurtosis I found in Section 1. It has an asymptotic normal distribution, which is a good approximation in a sample of this size.

The last line of the table is a forecast rule of making no prediction when risk is more than two standard deviations above its mean (20.7 percent in our sample). Using this rule of thumb, 56 observations are eliminated, but the MSE improves almost threefold. The M-stat shows that the improvement is statistically significant.

6 Conclusion

The idea that exchange rates are unpredictable needs to be qualified. The fixed exchange rates of the ERM are difficult to predict only at times of realignment. These regime changes, which contribute to the characteristic GARCH effects, are also predictable. We were able to improve our forecast accuracy almost 300 percent by limiting our predictions to those days in which the risk of realignment was not significantly higher than average.

REFERENCES

Bollerslev, T. (1986). Generalized autoregressive conditional heteroscedasticity, *J. Econ.* **31**, 307–27.

Diebold, F. X., and Nason, J. (1990). Nonparametric exchange rate prediction? *J. Int. Econ.* **28**, 315–32.

Engle, R. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of U.K. inflation, *Econometrica* **50**, 987–1008.

Engel, C., and Hamilton, J. D. (1990). Long swings in the dollar: are they in the data and do Markets know it? *Am. Econ. Rev.* **80**, 689–713.

Froot, K., and Thaler, R. (1990). Anomalies: Foreign Exchange, *J. Econ. Perspect.* **3**, 179–92.

Hsieh, D. (1988). The statistical properties of daily foreign exchange rates: 1974–83, *J. Int. Econ.* **24**, 129–45.

Krugman, P. (1991). Target zones and exchange rate dynamics, *Q. J. Econ.* **56**, 669–82.

Meese, R., and Rogoff, K. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *J. Int. Econ.* **14**, 3–24.
Meese, R., and Rose, A. (1990). Nonlinear, nonparametric, nonessential exchange rate estimation, *Am. Econ. Rev. Pap. Proc.* **80**, 192–96.
(1991). An empirical assessment of nonlinearities in models of exchange rate determination. *Rev. Econ. Stud.* **80**, 608–19.
Mizrach, B. (1996), Mimeo, Rutgers University. Learning and conditional heteroscedasticity in asset returns, Working Paper, The Wharton School,
(1992). Multivariate nearest-neighbour forecasts of EMS exchange rates, *J. Appl. Econ.* **7** (Supp.), S151–63.
(1993). Mean reversion in EMS exchange rates, Working Paper No. 93-01, Federal Reserve Bank of New York.
(1995a). *J. of International Money and Finance* **14**, 641–57. Target zones models with stochastic realignments: An econometric evaluation, Working Paper No. 93-02. Federal Reserve Bank of New York.
(1995b). Rutgers U. Working Paper 95-24. Forecast comparison in L_2, Working Paper, The Wharton School.
(1995c). Mimeo, Rutgers University. Uncovering interest parity in the ERM, Working Paper, Federal Reserve Bank of New York.

CHAPTER 17

Daily returns in international stock markets: Predictability, nonlinearity, and transaction costs

Steve Satchell and Allan Timmermann

Evidence of nonlinear components in daily returns is compared across 12 national stock-market indexes and a world index over the period 1980–92. A nonparametric nearest-neighbor algorithm is applied to identify nonlinear patterns in the daily returns and update embedding vectors one step ahead to generate recursive forecasts. The economic value of these nonparametric recursive forecasts is measured by using the forecasts in a simple trading strategy that explicitly accounts for transaction costs. The transaction costs are important factors in determining the stochastic properties of the daily returns series in the markets under consideration and probably are sufficiently large to rule out economic profits from trading on the basis of nonparametric forecasts.

1 Introduction

Nonlinear analysis is now securely placed as an important technique in applied economics. A possible procedure, and the one taken in this paper, is to test if a time series is independent and identically distributed, and, if this state is rejected, to investigate whether the series is chaotic.[1] The latter analysis can be based on an examination of the correlation integral as a function of the embedding dimension. In some cases, indications of chaos have been found although the statistical methodology is weakened by small sample problems (Ramsey, Sayers, and Rothman 1990).

What has not been done in most of these analyses is to link the identification of nonlinear patterns with the short-term predictability of the series and assess

Comments from two anonymous referees and from participants at the conference on nonlinear dynamics and economics at the European University Institute in Florence are gratefully acknowledged.

[1] See Brock et al. (1991), Frank, Gencay, and Stengos (1988), Frank and Stengos (1988), Hsieh (1989, 1991), and Sheinkman and LeBaron (1989). Alternatively, one could test for linearity and, if rejected, proceed to model the detected nonlinearity as a nonlinear stochastic process.