The Distribution of Exchange Rates in the EMS

Charles Engel and Craig S. Hakkio
Research Division, Federal Reserve Bank of Kansas City, 925 Grand Boulevard, Kansas City, MO 64198-0001, U.S.A.

Exchange rates of currencies in the Exchange Rate Mechanism (ERM) of the European Monetary System (EMS) are characterized by long periods of stability interrupted by periods of extreme volatility. The periods of volatility appear at times of realignments of the central parities and at times when the exchange rate is within the ERM bands. We begin by considering a procedure for finding outliers based on measuring distance as a quadratic form. The evidence suggests that the exchange rates of the EMS can be described by a mixture of two distributions. We therefore model the exchange rate as switching between two distributions—one that holds in stable times and the other that holds in volatile times. In particular, we use Hamilton's Markov-switching model. In addition, we extend Hamilton's model by allowing the probability of switching from one state to another to depend on the position of the exchange rate within its EMS band. This model has the interesting implication that near the edge of the band, large movements—either realignments or large jumps to the centre of the band—are more likely if the move to the edge of the band has been precipitous.

KEYWORDS: EMS; ERM: exchange rates; Markov-switching model

SUMMARY

EMS exchange rates are characterized by long periods of stability interrupted by periods of extreme volatility. Moreover, the periods of volatility do not occur at random points in time. Rather, they persist for a period of time. This result is based on two very different econometric approaches—one non-parametric and the other parametric. Because there is persistence in the volatile regime, we employ a Hamilton Markov-switching model rather than an independent switching model or a diffusion-jump model.

Secondly, we find that the probability of remaining in the volatile or stable regimes depends on how the exchange rate approaches the edge of the exchange rate band. If the exchange rate jumps to the edge—in a volatile move—the probability of remaining in the volatile regime is greater. However, if the exchange rate moves to the edge of the band in a smooth way, or if the exchange rate simply bumps along the edge of the band, then the probability is greater that the exchange rate will remain in the stable regime. Intuitively, if the exchange rate approaches the outer band in a volatile way, exchange rates are likely to remain volatile. This reflects the fact that exchange rates are often realigned in these cases. However, when central banks fight off realignments, exchange rates often remain at the edge of the band for a period of time—that is, they remain in the stable regime.

INTRODUCTION

Exchange rates of currencies in the Exchange Rate Mechanism (ERM) of the European Monetary System (EMS) are characterized by long periods of stability interrupted by periods of extreme volatility. The periods of volatility appear at times of realignment of the central parities in the ERM, and at times when the exchange rates are within the ERM bands. In this paper, we argue that EMS exchange rates, like many exchange rates, are not normally distributed. In addition, the distribution of EMS exchange rates is quite different from that of more freely floating rates. We model the exchange rate as switching between two distributions—one that holds in stable times and the other that holds in volatile times.

We begin with a preliminary examination of the data, comparing the behaviour of weekly changes in the French franc/DM and Italian lira/DM exchange rates to the behaviour of the dollar/DM and the Japanese yen/DM exchange rates. We show that the exchange rates under the EMS and floating systems behave considerably differently. In addition, we show that a single normal distribution does not adequately describe EMS exchange rates.

More specifically, the first section considers a procedure for determining outliers. 'Outliers' is the term used by statisticians to describe observations that do not conform to the pattern suggested by the majority of observations. In this paper, we will say that outliers come from the volatile period. The procedure is based on measuring distance as a quadratic form, as in the recent work of Hadi (1992, 1993). EMS exchange rates have more outliers than floating exchange rates, but EMS outliers are less volatile than floating rate outliers. This evidence suggests that the exchange rates of the EMS can be described by a mixture of two distributions. Furthermore, EMS outliers tend to cluster together.

The outliers are not merely realignment dates. Sometimes there are large week-to-week changes of exchange rates within the ERM bands. This is not surprising for the lira, which had very wide bands during much of the time it was participating in the ERM, but it is also true, to some extent, for the franc. Conversely, sometimes when a realignment of the central parity occurred, the new bands overlapped with the old bands, so that a jump in the exchange rate was not necessary.

We proceed from these observations to estimate a stochastic process for EMS exchange rate changes assuming that exchange rates are distributed as a mixture of two normal distributions. Furthermore, the preliminary examination of the data showed that there is some persistence in the 'state' of the exchange rate. While realizations from the volatile distribution are infrequent, they are more likely to occur if the exchange rate was in the volatile state in the previous period than if it was not. Consequently, we let the probability that the exchange rate is drawn from the 'volatile' or 'stable' distribution in any period depend on which distribution the exchange rate was drawn from in the previous period. So, Hamilton's (1989) Markov-switching model is applied to the EMS exchange rates. (Subsequent to our first two revisions of this paper, we received Mundaca (1994), which fits a version of the Hamilton model estimated for several Scandinavian currencies that forces regime breaks at times of realignments. It assumes constant switching probabilities between states.)

We then extend the Hamilton model by allowing the probability of switching from one state to another to depend on the position of the exchange rate within its EMS band (as opposed to Hamilton's assumption of constant switching probabilities). We find that the probability of staying in the volatile regime is higher when the exchange rate is near the top of its band. This reflects the fact that, around times of realignments (which have always been upward realignments for these exchange rates), the exchange rates are volatile. But, interestingly, we also find that the probability of staying in the stable regime is also higher when the exchange rate is near the top of its band—assuming the exchange rate was already in the quiet regime.

A common interpretation of the data is that the EMS exchange rates are more volatile near the edge of the bands (see, for example, Bertola and Baballero, 1992, p. 527). However, our finding shows that it is not accurate to say that the exchange rate is more likely to have a large jump—either because of realignment or a large intervention that pushes the exchange rate toward the centre of the band—when the exchange rate approaches the edge of the band. It matters how the exchange rate approached the boundaries of the target zone. If it moved there gradually, then large changes in the exchange rate are unlikely. If it moved there precipitously, then further volatility is likely. The fact that the probability of a realignment depends on the regime rather than depending on predictive variables in a linear way could explain the general lack of success of previous model of realignments.

(See Svensson, 1992 and Rose and Svensson, 1993. Note, however, that Chen and Giovannini, 1993 find some limited success explaining realignments using the distance of the exchange rate from the centre of the band and the length of time since the previous realignment.)

There is a connection between our approach and the literature that investigates the frequent occurrence of outliers in exchange rate changes in floating exchange rate regimes. See, for example, Baillie and Bollerslev (1989), Boothe and Glassman (1987), Diebold (1988), Koedijk *et al.* (1990) and Koedijk *et al.* (1992). As in that literature, we emphasize the non-normality of the exchange rate. However, we concentrate on the peculiar features of the exchange rates in the ERM, as opposed to the aforementioned studies, which focus on floating exchange rates.

It might seem natural to model the exchange rate in the EMS as a mixed diffusion-jump process. Akgiray and Booth (1988) and Tucker and Bond (1988) model floating exchange rates in this manner, and Ball and Roma (1993) apply the diffusion-jump model to EMS exchange rates. However, an aspect of the exchange rates in the EMS, alluded to above, is that volatility tends to cluster. That is, it is not true that EMS exchange rates can be described as having periods of stability broken by jumps caused by realignments. Instead, there seem to be periods of volatility around the time of realignments, and occasionally at other times. Furthermore, these periods of volatility often last for more than one week at a time. Thus, it seems natural to describe the exchange rate as being drawn from two regimes—a stable regime and a volatile regime—with some persistence in the regimes. So, the likelihood of a volatile period is increased if the current state is volatile. Indeed, effectively we reject a discrete time version of the jump-diffusion model because we find statistically significant evidence of persistence in the states.

One possible way to model clustering of volatility is with GARCH, as in Baillie and Bollerslev (1989). In the EMS, however, the clusters of volatility appear abruptly, as if there were a sudden switch in regimes. So, it is more natural to segment the time series into periods of small and large changes in exchange rates. It is possible that ARCH could provide additional explanatory power within the context of a regime-switching model (that is, regime-switching and ARCH are not mutually exclusive). Hamilton and Susmel (1992) developed such a model, which they labelled SWARCH, but we do not pursue that here. Nieuwland *et al.* (1994) combined a GARCH model of variances with the jump-diffusion model to capture this clustering of volatility that we observe.

We investigate the data using quadratic distance measures, and introduce the Hamilton–Markov switching model and provide estimates of the stochastic process for EMS exchange rates. The Hamilton model is extended by allowing time-varying transition probabilities. The early methods provide algorithms for choosing periods of highly volatile exchange rates. We compare the estimated periods of volatility from each of the three later methods. We then focus on the time-series properties of the exchange rates in the EMS with particular emphasis on modelling the outliers and the clustering of volatility. We step aside from the target-zone literature, as has much of the literature cited above. Finally, we discuss, from the perspective of portfolio selection, why the distribution of exchange rates is important.

A DISTANCE APPROACH FOR IDENTIFYING OUTLIERS

In this section we use a statistical procedure to determine which observations are outliers. We also compare the behaviour of EMS and non-EMS exchange rates. Let e denote the percentage change in the exchange rate, measured as $100[ln(E_t) - ln(E_{t-1})]$. In this paper, we focus on two EMS exchange rates—the franc/DM and the lira/DM—and two floating exchange rates—the dollar/DM and the yen/DM. We use weekly data: noon buying rates, from New York City, on Tuesday (Wednesday if Tuesday is unavailable). The sample period is 20 March 1979 to 8 September 1992 (for the lira) and to 23 October 1993 (for the franc, dollar and yen). The first observation was chosen since the EMS was formed on 13 March 1979. 8 September 1992 was the last observation for the lira, since Italy withdrew from the ERM on 17 September 1992.

The approach for identifying outliers was developed by Hadi (1992, 1993). Those observations, which come from the tails of the empirical distribution of e, are said to be outliers, while those observations that come from the centre are said to come from the stable period. We will say that observations from the stable period come from $F(e)$. The basic idea of Hadi is simple: all observations that are 'close' together belong to $F(e)$, and all other observations are outliers. 'Close' is defined by measuring distance as a quadratic form, which should be distributed χ^2.

The Rule for Classifying Observations

How would someone find an outlier if they knew there was only one outlier? First, for all observations form two groups: one group is the single observation and the other group is all other observations. Then, calculate the distance from the single observation to the centre of all other observations. The centre is calculated as the mean (or median); distance is measured relative to the dispersion of the data (the variance). The single outlier is then that observation for which the distance is maximized.

A similar approach could be used to find multiple outliers. Suppose we knew there were exactly two outliers. Consider all pairs of observations, and let the remaining T-2 observations come from $F(e)$. Then, calculate the distance between each pair of observations and the centre of $F(e)$. The pair of outliers is then that pair for which the distance is maximized. As long as the number of outliers is known in advance, the procedure for finding outliers is straightforward, but time-consuming.

However, if the number of outliers is unknown, the problem is more difficult. One must determine how many outliers exist and which observations are outliers. The discussion of the solution presented here is intuitive; see Hadi (1992, 1993) for a more detailed discussion.

The procedure to find outliers is iterative.

Step 1. Begin with a two element set, denoted $\mathcal{F}_2(e)$, that is most likely *not* to contain any outliers. ('$\mathcal{F}$' is used as a reminder of 'F,' and 2 as a reminder that it is a two element set.)

Step 2. Increase the size of $\mathcal{F}_2(e)$ to about half the observations. In this step, let r denote the number of observations in the set. Rank the observations by the distance between the observation and the centre of $\mathcal{F}_r(e)$, relative to a measure of dispersion. Continue until $\mathcal{F}_r(e)$ has half the observations.

Step 3. For this step, let r again denote the number of observations in $\mathcal{F}_r(e)$. Rank the observations by the distance between the observation and the centre of $\mathcal{F}_r(e)$, and test whether observation $r+1$ is an outlier. Observation $r+1$ is then compared with a $\chi^2_{\alpha/n}$ statistic, because distance is measured as a quadratic form of normals.
(a) If the distance is greater than $\chi^2_{\alpha/n}$, set $F(e) = \mathcal{F}_r(e)$ and the remaining observations are outliers. Stop.
(b) If the distance is less than $\chi^2_{\alpha/n}$, increase the size of $\mathcal{F}$ to $r+1$, denoted $\mathcal{F}_{r+1}(e)$. Repeat Step 3.
(c) If $\mathcal{F}_r(e)$ contains all the observations, there are no outliers. Set $F(e)$ = all observations and there are no outliers. Stop.

The Distribution of EMS Exchange Rates

Table 1 summarizes the exchange rate data according to the distance rule. The table shows the number of observations and the standard deviation for observations from the stable period and for the outliers. Two facts are clear. First, EMS exchange

Table 1. Summary statistics based on the distance rule.

	French franc	Italian lira	US dollar	Japanese yen
$F(e)$: the stable period				
Number of observations	723	681	761	759
Standard deviation	0.20	0.29	1.54	1.32
Outliers				
Number of observations	40	23	2	4
Standard deviation	1.85	1.36	9.13	0.58

rates have many more outliers than floating exchange rates. Secondly, the standard deviation of EMS exchange rates is much less than the standard deviation of floating exchange rates, for both the stable period and the volatile period (outliers).

Figures 1 and 2 show when the outliers occur. The vertical lines represent realignments in the French franc/DM exchange rate and the Italian lira/DM exchange rate. There were only two dollar/DM outliers (24 September 1985 and 15 September 1985) and four yen/DM outliers (11 December 1979, 9 December 1980, 11 August 1981, and 15 September 1992). Two conclusions follow from the figures. First, not all outliers occur at the time of a realignment. Secondly, some realignments do not correspond to an outlier. That is, outliers are neither necessary nor sufficient for a realignment.

We now turn to testing whether outliers and observations belong to the $F(e)$ cluster. We use a non-parametric runs test. Define an indicator variable that equals 0 if $e \in F(e)$ and 1 if e is an outlier. The null hypothesis is that the 0s and 1s occur randomly over time. That is, outliers do *not* cluster. The results strongly suggest that the observations for the franc (p value is 0) do cluster, while the lira, dollar, and yen do not cluster (p value of 0.76, 0.93 and 0.73). So, floating rates do not exhibit clustering of outliers, while the results are mixed for EMS exchange rates.

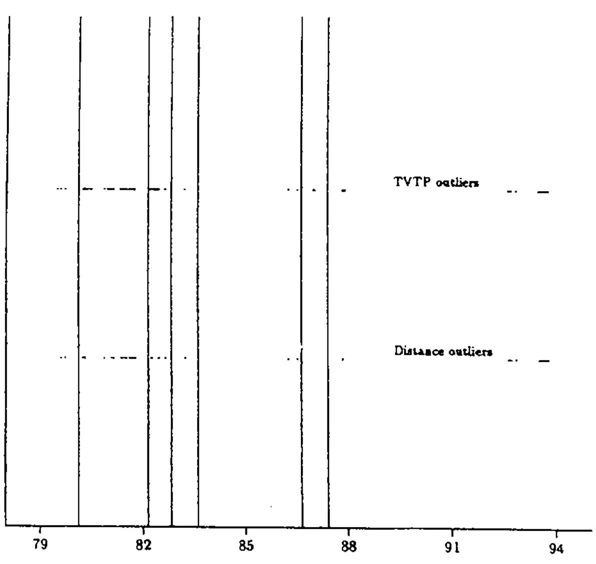

Figure 1. Franc/DM exchange rates from the volatile period. *Note:* The vertical lines represent realignments.

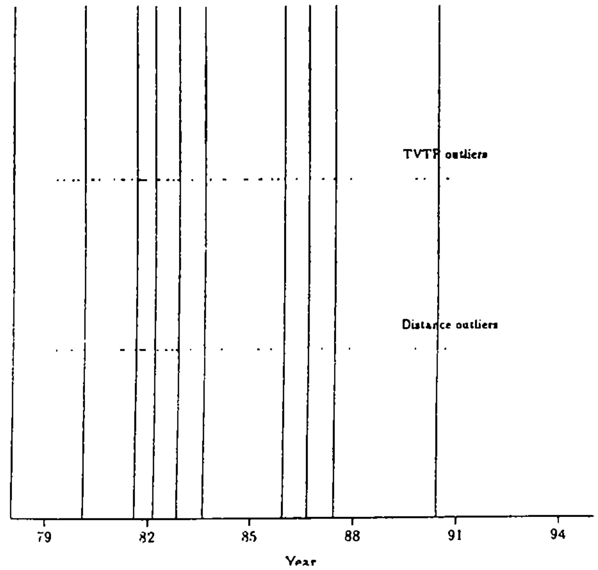

Figure 2. Lira/DM exchange rates from the volatile period. *Note:* The vertical lines represent realignments.

To summarize, floating exchange rates, like the dollar/DM and yen/DM, behave differently than EMS exchange rates. Outliers are more frequent in the EMS, but they are also less volatile. EMS exchange rate changes seem to fall into two distinct categories: very large changes and very small changes. There are few realizations that fall in the middle ground. This contrasts with floating rates, which exhibit a wide range of outcomes. We conclude that the EMS exchange rates are drawn from a mixture of two distributions—one which has a high variance and one with a low variance. Since there is some evidence that volatile periods tend to cluster together, we are led to model the EMS exchange rates with Hamilton's (1989) Markov switching model.

THE SWITCHING MODEL

The previous section implies that the stochastic process for exchange rates in the EMS can be based on a mixture of probability distributions. In addition, since outliers tend to cluster, we estimate Markov-switching models for EMS exchange rates, as opposed to diffusion-jump processes.

The Rule for Classifying Observations

We assume there are two possible states of the world: the 'stable' state (s) and the 'volatile' state (v).

The stable state occurs most of the time, while the volatile state occurs less frequently. We do not exogenously choose which periods are stable and volatile. Instead, we assume that the exchange rate is described by a mixture of normal distributions. The parameters of the distribution are the mean and variance in the stable state, μ_s and σ_s^2, and the mean and variance in the volatile state, μ_v and σ_v^2. In addition, there are parameters that determine the probability of the stable and volatile state occurring. These parameters are all estimated to maximize the likelihood function.

We could model the probability of each state occurring at any date as being independent of the state that the exchange rate was in during the previous period. In this case, there is a fixed probability π of state s occurring, and a probability $1 - \pi$ of state v occurring. This model is essentially the discrete time analogue of the mixed process estimated by Akgiray and Booth (1988) and Tucker and Pond (1988). It implies that the probability that the exchange rate is in the volatile state in period t is independent of whether the exchange rate was in the stable state or the volatile state in period t-1. If this were accurate, then the exchange rate changes would not exhibit volatility clustering.

However, as discussed earlier, it appears that a volatile week is more likely to occur if the previous week was also volatile and a stable week is more likely if the previous week was stable. If this is true, Hamilton's (1989) Markov-switching model would be appropriate. In this case, we let π_s be the probability that a stable week is followed by a stable week, and π_v be the probability that a volatile week is followed by a volatile week. We do not impose the restriction that $\pi_s = 1 - \pi_v$, as would the independent switching model. We will present tests of this restriction.

In the next section we estimate a time-varying transition probability model. To nest the model of this section in the model of the next section, it is convenient to parameterize the model slightly differently than does Hamilton. We actually estimate the parameters θ_s and θ_v, defined by

$$\pi_s = \frac{e^{\theta_s}}{1 + e^{\theta_s}}$$

$$\pi_v = \frac{e^{\theta_v}}{1 + e^{\theta_v}}$$

The parameters of the model—$\mu_s, \mu_v, \sigma_s^2, \sigma_v^2, \theta_s, \theta_v$—are estimated to maximize the likelihood using the procedure described in Hamilton (1989) and Engel and Hamilton (1990).

The Distribution of EMS Exchange Rates

Table 2 presents the parameter estimates. The first thing to note is that for all currencies the variance of the exchange rate in the volatile state, σ_v^2, is much larger than the variance in the stable state, σ_s^2. So, it is natural to think of the two states as being high variance and low variance states.

The mean of the exchange rate in the volatile state (high-variance) is positive and larger than the mean in the stable state (low variance). This is because the mark tends to be revalued relative to the other currencies during realignments.

The mean change in the nominal exchange rate is small but positive in the stable state. This is because there is a slight tendency for currencies to depreciate against the mark when the exchange rate remains within the bands, but the average change is very small.

The probability of staying in the stable state in period $t+1$ if the exchange rate is in the stable state in period t, given by π_s, is very high for all currencies. This is because the periods of stability are relatively long-lasting. In fact, the expected length of stay in state s is given by $1/(1 - \pi_s)$. For the franc/DM rate, $\pi_s = 0.953$, so the stable state lasts 21 months.

We cannot say with certainty whether the exchange rate change on any date was drawn from the volatile distribution or the stable distribution. We can follow the practice, however, of classifying an observation as being drawn from the volatile

Table 2. Markov switching model.

	French franc		Italian lira	
	Stable	Volatile	Stable	Volatile
μ	0.009	0.376	0.001	0.561
	(0.007)	(0.133)	(0.010)	(0.120)
σ^2	0.030	1.610	0.062	0.851
	(0.001)	(0.257)	(0.002)	(0.145)
θ	3.010	0.709	2.400	−0.236
	(0.170)	(0.262)	(0.150)	(0.257)
π	0.953	0.670	0.917	0.441

distribution if the Bayesian probability of it being from that distribution is greater than 0.50. Given the parameters of the distribution (we treat the parameter estimates as being the true parameters in these calculations), the 'filter probability' that a given observation comes from state s is the probability that the exchange rate at time t is from state s, conditional on all observations in the sample up until time t. Using this classification, all French franc/DM realignments occurred during the volatile period; two Italian lira/DM realignments occurred during the stable period and seven realignments occurred during the volatile period.

As in previous sections, we test whether observations that come from the volatile distribution (specifically, observations for which the filter probability of coming from the volatile state is greater than 50%) tend to cluster. Using the runs test, we conclude that observations from the volatile period do tend to cluster.

For this model, we can test the hypothesis more specifically. The Markov switching model is distinguished from the independent switching model because it does not impose the condition that $\pi_v + \pi_s = 1$. Using the equation for π_v and π_s, this condition is equivalent to $\theta_s + \theta_v = 0$. Using a likelihood ratio test, we can reject this hypothesis with a p-value of 0.0 for the franc and a p value of 0.028 for the lira. So, there is strong evidence against the simple independent switching model.

In this section we have assumed that the probability of moving from one state to another is constant. That does not seem like an adequate description of the situation in the EMS, where often speculative attacks make it seem much more likely than usual that an observation will be from the volatile distribution this week if it is near the outer band. In the next section, we move beyond the Hamilton model to allow the transition probabilities to vary over time in a way that takes into account the likelihood of a realignment.

A TIME-VARYING SWITCHING MODEL

The previous section assumed that the probability of switching from the stable state to the volatile state (and the volatile state to the stable state) is constant. In this section, we generalize this assumption: we let the probability of switching states depend on how far the exchange rate is from the upper band. For example, suppose the exchange rate is in the volatile state and very close to the upper band. The time-varying transition probability model allows the probability of staying in the volatile state to be greater in this case (than if the exchange rate were at the central parity). In other words, if the exchange rate is likely to be more volatile when it is near the upper band, then we want to allow the probability of being in the volatile regime to be greater.

The Rule for Classifying Observations

Estimating a time-varying transition probability model is very similar to estimating a constant transition probability model. In the previous section, π_s was constant; in this section, we let $\pi_s = \pi_s(z_{t-1})$ and $\pi_v = \pi_v(z_{t-1})$. More specifically, let

$$\pi_s(z_{t-1}) = \text{Prob}[\text{state}_t = \text{stable} \mid \text{state}_{t-1} = \text{stable},\ z_{t-1}] = \frac{e^{\theta_s^0 + \theta_s^1 z_{t-1}}}{1 + e^{\theta_s^0 + \theta_s^1 z_{t-1}}}$$

$$\pi_v(z_{t-1}) = \text{Prob}[\text{state}_t = \text{volatile} \mid \text{state}_{t-1} = \text{volatile},\ z_{t-1}] = \frac{e^{\theta_v^0 + \theta_v^1 z_{t-1}}}{1 + e^{\theta_v^0} + \theta_v^1 z_{t-1}}$$

Eight parameters are estimated: $\mu_s, \mu_v, \sigma_s^2, \sigma_v^2, \theta_s^0, \theta_s^1, \theta_v^0, \theta_v^1$.

Our estimation procedure extends the work of Hamilton (1989) to allow for these time-varying transition probabilities. Other recent studies have examined time-varying transition probabilities in a Markov-switching framework. Diebold *et al.* (1992), Filardo (1994), and Zhu (1993) discussed some of the econometric issues in estimating these models.

The Distribution of EMS Exchange Rates

The results are presented in Table 3. z_t is the percentage distance from the upper band. More specifically, $z_t = 100[ln(upper_t) - ln(E_t)]$, where *upper* equals the upper band and E equals the actual exchange rate. As it turns out, the estimates of $\mu_s, \mu_v, \sigma_s^2, \sigma_v^2$ are similar to those reported in Table 2 for the constant probability model.

Table 3. Time varying transition probability model.

	French franc		Italian lira	
	Normal	Volatile	Normal	Volatile
μ	0.009	0.353	0.001	0.563
	(0.007)	(0.122)	(0.010)	(0.117)
σ^2	0.029	1.508	0.063	0.840
	(0.001)	(0.222)	(0.002)	(0.142)
θ^0	3.193	0.588	2.592	0.050
	(0.254)	(0.314)	(0.192)	(0.283)
θ^1	−0.489	−1.107	−0.193	−0.143
	(0.218)	(0.270)	(0.097)	(0.196)
π	0.961	0.643	0.930	0.512

We find θ_v^1 and θ_s^1 are both negative. It is easy to show that $sign(\partial\pi/\partial z_{t-1}) = sign(\theta^1)$. This means that the probability of staying in state *v*, if we are already in it, is greater the closer the exchange rate is to the top of the band (z_{t-1} smaller). In other words, it is more likely that a volatile realization will be followed by a volatile realization when the exchange rate is far from the central parity.

This result is somewhat similar to the result of Ball and Roma (1993). in their jump-diffusion process, they find that the probability of a jump is greater the farther the exchange rate is from the central parity. There are two reasons why our result is different—one minor and one significant. The minor reason is that we model the probability as depending on the distance from the top of the band. We prefer that formulation because, for the lira, the width of the band changed during our sample. Hence, as the band narrows, being close to the top of the band corresponds to a smaller distance from the centre.

The more significant difference is that we are not modelling the unconditional probability of the volatile state as a function of the position in the exchange rate band. Rather, in addition to depending on this location of the exchange rate in the band, we model the probability of the volatile state as being conditional on the current state—that is, our model is of the *Markov-switching*, rather than the independent switching, variety. We capture the idea of volatility clustering with the Markov-switching model, because being in the volatile state makes it more likely to remain in the volatile state. This property is not in the independent switching model, or the jump-diffusion model. So, our finding that θ_v^1 is negative does not mean simply that the volatile state is more likely when the exchange rate is at the top of the band. It means that it is more likely *conditional* on the exchange rate being in the volatile state.

This distinction is important because we also find that θ_s^1 is negative. That means that, conditional on being in the stable state, remaining in the stable state is more likely when the exchange rate is near the top of the band. EMS exchange rates have experienced long spells near the top of the band, as central banks fight off realignment. These spells are often broken eventually by realignments, but sometimes the exchange rate drops back toward the central parity. This result is not inconsistent with θ_v^1 being negative. We can interpret the two findings simply as saying that a transition from the stable state to the volatile state is more likely to occur when the exchange rate is near the centre of the band, but a volatile state is more likely to follow a volatile state near the top of the band.

We later discuss the restrictions placed on the model if the switching probabilities are independent of the state but dependent on the position within the band. One of the restrictions is that $\theta_v^1 = \theta_s^1$. This condition relates to the discussion above. The independent switching model, and the jump-diffusion model, insist that either the volatile or the stable state is more likely irrespective of the state's last period. So, the finding in that model that the volatile state is more likely near the top of the band would hold even if the state last period were stable. The fact that we find that both θ_v^1 and θ_s^1 are negative is inconsistent with that conclusion. We later reject that conclusion—that is, we reject the discrete time version of the Ball and Roma model.

As in previous sections, a date is classified as being in the volatile state if the filter probability of the volatile state is greater than 0.50. Figures 1 and 2 show which observations come from the volatile distribution. All French franc/DM realignments occurred during the volatile period; two Italian lira/DM realignments occurred during the stable period and seven realignments occurred during the volatile period.

As in previous sections, we test whether observations that come from the volatile distribution tend to cluster together. Using the runs test, we conclude that observations from the volatile period do tend to cluster.

Our specific test for volatility clustering is testing the Markov-switching model against the model in which the probability of the volatile (or stable) state can depend on z_{t-1}, but does not depend on the state in the previous period. This independent switching model imposes $\pi_s(z_{t-1}) = 1 - \pi_v(z_{t-1})$ for all realizations of z_{t-1}. This condition is equivalent to $\theta_v^0 = -\theta_s^0$ and $\theta_v^1 = -\theta_s^1$. We reject that restriction with a p-value of 0 for the franc. This is not surprising since θ_v^0 and θ_s^0 are estimated to be of the same sign as θ_v^1 and θ_s^1. However, we cannot reject the hypothesis for the lira, the p-value is 0.13. The reason is that θ_v^0 and θ_v^1 are insignificant, so the test cannot reject that $\theta_v^0 = -\theta_s^0$ and $\theta_v^1 = -\theta_s^1$.

HOW DO THE THREE RULES COMPARE?

We have presented three different rules for choosing which observations come from the stable period and which observations come from the volatile period. To simplify notation, we define two distributions: observations from the stable period come from $F(e)$, and outliers and observations from the volatile period come from $G(e)$. This section compares the results from the three rules.

Figures 1 and 2 show the observations assigned to $G(e)$ for the distance rule and the time-varying parameter model. While the rules do not always pick the same observations to put into $G(e)$, they do seem to pick out many of the same observations.

Table 4 can be used to compare how the three rules assign franc/DM observations to $F(e)$ and $G(e)$. The table is actually three two-way tables. In the table, FTP denotes the 'fixed transition probability' switching model and TVTP denotes the 'time varying transition probability' switching model. For example, the FTP/distance subtable compares the FTP rule with the distance rule. It shows that the two rules assign observations to $F(e)$ and $G(e)$ in a similar way. Of the 723 observations that belong to $F(e)$ based on the distance rule, 694 were put in the $F(e)$ distribution and only 29 were put into the $G(e)$ distribution when using the FTP rule. Similarly, of the 40 observations that belong to $G(e)$ based on the distance rule, all 40 were put in the $G(e)$ distributions when using the FTP rule. The other subtables are read in a similar way.

Table 4. Comparison of rules for classifying franc/DM outliers.

		Distance		FTP	
		$F(e)$	$G(e)$	$F(e)$	$G(e)$
FTP	$F(e)$	694	0		
	$G(e)$	29	40		
TVTP	$F(e)$	683	0	681	2
	$G(e)$	40	40	13	67

Since the table is several two-way tables, we can test whether the allocations of observations by the three rules are related. We can calculate Fisher's exact test. The probability is 0 for each pair. Therefore, even though the rules are different, they are basically picking the same outliers.

It would be possible to construct similar tables for the other exchange rates. However, they look very similar to Table 4. In addition, Fisher's exact test for association between the rules yields a probability of 0. As with the franc, the rules are basically picking the same outliers.

IMPLICATIONS OF TWO DISTRIBUTIONS FOR EMS EXCHANGE RATES

As Boothe and Glassman (1987) pointed out, there are at least two reasons to investigate the nature of the stochastic process underlying exchange rate movements. First, the distribution of the exchange rate matters for statistical tests of open economy models. Secondly, theoretical models in international finance often rely upon specific distributional assumptions. We take up the second argument in this section. We contrast the optimal portfolio choice of an investor when the exchange rate is normally distributed and when it is drawn from a mixture of normal distributions. This section is not a full-fledged study of intertemporal asset demand, as in the study of Park *et al.* (1993) for the jump-diffusion process. It is set in a two-period framework, and is intended to illustrate the importance of the distributional assumption for asset choice.

We look at the asset choice of an individual with a one-period horizon. We assume that the investor has an exponential utility function. We choose this function because it is well known that, with this function, and with returns distributed normally, the investor who maximizes expected utility will choose a mean-variance efficient portfolio. We will contrast the optimal portfolio choice for this investor when returns are normally distributed with the optimal portfolio choice when returns are drawn from a mixture of normal distributions. In this case, the optimal choice of the expected utility maximizer cannot be characterized simply as the choice of a mean-variance efficient portfolio.

The individual's utility function exhibits constant absolute risk aversion

$$U(C) = -e^{-aC}.$$

Suppose at time 0 the agent allocated wealth to maximize expected utility at time 1. Because the agent consumes all wealth at time 1, expected utility of consumption at time 1 simply equals expected utility of wealth. Let the asset choice be between a foreign and domestic security: $W_0 = b + Sb^*$.

The assets could be denominated in nominal terms—so utility depends on nominal wealth, and S is the nominal exchange rate. In that case, b is denominated in the home currency and b^* in the foreign currency. Alternatively, investors might only consume the home good. Then, wealth is denominated in terms of that good. b^* is in terms of the foreign good and S is the price of the foreign good in terms of the domestic good. Under this interpretation, S is the real exchange rate.

Assume that neither security pays interest and that their prices are fixed at one in terms of their own currency (or good). Then wealth in period 1 is given by

$$W_1 = b + S_1 b^* = W_0 + \delta\hat{b},$$
$$\delta = \frac{S_1 - S_0}{S_0},$$
$$\hat{b} = S_0 b^*.$$

The asset choice problem becomes one of choosing $\hat{b}$ to maximize

$$E_0(-e^{-a(W_0+\delta\hat{b})}).$$

Since e^{-aW_0} is constant, the problem reduces to maximizing

$$-E_0 e^{-a\delta\hat{b}}.$$

If δ is distributed normally with mean μ and variance σ^2, then using the moment generating function for a normal distribution, we can write

$$E_0 e^{-a\delta\hat{b}} = e^{-a\hat{b}\mu + (a\hat{b})^2\sigma^2/2}.$$

The first-order condition is $\mu - \hat{b}a\sigma^2 = 0$. If we define $R = aW_0 =$ the degree of relative risk aversion, and $\lambda^N = \hat{b}/W_0 =$ the share of the portfolio held in the risky foreign asset when returns are normally distributed, then

$$\lambda^N = \frac{\mu}{R\sigma^2}. \tag{1}$$

Now, consider the asset choice when δ is distributed according to the switching model. We assume that there are two possible states of the world. Changes in the exchange rate are distributed normally conditional on the state of the world. The stable state occurs most of the time. The exchange rate has low variance in the stable state. The volatile state occurs less frequently and is characterized by higher variance.

The parameters of the distribution are the mean and variance in the stable state, μ_s and σ_s^2, and the mean and variance in the volatile state, μ_v and σ_v^2. In addition, there are parameters that determine the probability of the stable and volatile state occurring. Let π be the fixed probability of state s occurring and $1 - \pi$ be the probability of state v occurring.

With these definitions, we now have

$$E_0 e^{-a\hat{b}\delta} = \pi e^{-a\hat{b}\mu_s + (a\hat{b})^2\sigma_s^2/2} + (1-\pi)e^{-a\hat{b}\mu_v + (a\hat{b})^2\sigma_v^2/2}.$$

The first-order condition can be written

$$\begin{aligned} 0 = {} & \pi_n(\mu_s - R\lambda\sigma_s^2)e^{-R\lambda\mu_s + (R\lambda)^2\sigma_s^2/2} \\ & + (1-\pi_n)(\mu_v - R\lambda\sigma_v^2)e^{-R\lambda\mu_v + (R\lambda)^2\sigma_v^2/2}. \end{aligned} \tag{2}$$

This first-order condition cannot be solved analytically for λ except in some special cases. Notice that the choice of λ depends on all five of the parameters of the distribution: $\mu_s, \mu_v, \sigma_s^2, \sigma_v^2, \pi$. In particular, λ cannot be expressed in terms of simply the mean and variance of δ. The mean of δ is given by

$$\mu = E_0\delta = \pi\mu_s + (1-\pi)\mu_v \tag{3}$$

and the variance of δ is given by

$$\sigma^2 = var_0(\delta) = \pi\sigma_s^2 + (1-\pi)\sigma_v^2 + (1-\pi)\pi(\mu_s - \mu_v)^2. \tag{4}$$

We will consider a series of examples that illustrates how the portfolio choice differs from the one chosen under the normal distribution, which depends only on the mean and variance.

Example 1

Suppose that the variance of the exchange rate in each state was zero, so that $\sigma_s^2 = 0$ and $\sigma_v^2 = 0$. Note that this does not mean that the exchange rate is not risky. The exchange rate still has a positive variance because it can jump between its value in the stable state, μ_s, and its value in the volatile state, μ_v.

If $\mu_s > 0$ and $\mu_v > 0$, then $\lambda = \infty$. Recall that the safe domestic asset pays off zero with certainty. The foreign asset, in domestic currency terms, would always have a higher pay-off under this scenario, whether the state was normal or abnormal. Investors would always be better off with the foreign asset. The return on the foreign asset stochastically dominates the return on the domestic asset. So, investors would want to take as large a long position in the foreign asset as possible, and short the domestic asset.

However, if investors had simply applied mean-variance analysis to this problem, following the portfolio choice dictated by Equation (1), they would have ended up choosing a vastly different portfolio. The optimal share of the foreign asset in this case would be given by

$$\lambda = \frac{\pi\mu_s + (1-\pi)\mu_v}{R[(1-\pi)\pi(\mu_s - \mu_v)^2]}.$$

This value of λ is finite. Clearly the mean-variance solution is far from the optimal choice.

Example 2

Now consider the special case in which $\mu_s = 0$ and $\sigma_s^2 = 0$. This would correspond to the case in which the exchange rate was absolutely fixed in the stable state. From the first-order condition, Equation (2), we can derive in this case that

$$\lambda = \frac{\mu_v}{R\sigma_v^2}.$$

Given that $\mu_v = \frac{1}{1-\pi}\mu$ and $\sigma_v^2 = \frac{1}{1-\pi}\sigma^2 - \frac{\pi}{(1-\pi)^2}\mu^2$, we can write

$$\lambda = \frac{\mu}{R(\sigma^2 - \frac{\pi}{1-\pi}\mu^2)}.$$

Therefore, $\lambda > \lambda^N = \mu/R\sigma^2$.

In this case, the investor's optimal choice of λ is always greater than λ^N. So there is a sense in which the variance overstates the volatility of the exchange rate. The investor is always willing to hold more of the risky asset than he or she would in a mean-variance efficient portfolio.

Example 3

Table 5 shows the optimal choice of λ for variety of parameters. In constructing the table, the parameters were chosen so that the unconditional mean $\mu = 0.10$ and the unconditional variance $\sigma^2 = 0.04$. The risk aversion parameter, R, was chosen to be 5. With these values, according to Equation (1), the mean variance optimal portfolio is always 0.50.

The parameters were chosen to approximate what we might see in the EMS. The probability of the volatile state is low compared with the probability of the stable state. But the mean change in the exchange rate is much larger in the volatile state, and the variance is chosen to be ten or twenty times larger in the volatile state than in the stable state.

Examination of Table 5 shows that there is a wide range of values for the optimal λ Often the optimal portfolio is very different from the mean-variance optimal one (that is, λ is very different from 0.50.)

We conclude that the variance is not a sufficient statistic to summarize the volatility of the exchange rate. An investor who based his decision on the rule of minimizing the variance for a given mean would end up choosing a very different portfolio from the

Table 5. Portfolio share for risky asset exchange rate follows a switching process $R = 5$, $\mu = 0.10$, $\sigma^2 = 0.04$.

π_n	μ_n	μ_a	σ_n^2	σ_a^2	λ
0.646	−0.010	0.30	0.004	0.04	0.819
0.821	0.057	0.30	0.012	0.12	0.660
0.970	0.073	1.0	0.012	0.12	1.211
0.893	0.077	0.30	0.012	0.24	0.521
0.974	0.076	1.0	0.012	0.24	1.210
0.911	0.080	0.30	0.02	0.20	0.545
0.980	0.082	1.0	0.02	0.20	0.821
0.953	0.090	0.30	0.02	0.40	0.441
0.983	0.085	1.0	0.02	0.40	0.795
0.959	0.092	0.30	0.028	0.28	0.497
0.989	0.089	1.0	0.028	0.28	0.638
0.979	0.096	0.30	0.028	0.56	0.414
0.991	0.092	1.0	0.028	0.56	0.621
0.989	0.098	0.30	0.036	0.36	0.487
0.997	0.097	1.0	0.036	0.36	0.538
0.995	0.099	0.30	0.036	0.72	0.429
0.997	0.098	1.0	0.036	0.72	0.529

individual who maximized expected utility, taking into account the fully parametrized distribution.

CONCLUSIONS

There are six conclusions that can be drawn from this paper.

1. The distribution of exchange rates is different in a system of floating exchange rates and a system of fixed but adjustable rates.
2. EMS exchange rates can be described by a mixture of two distributions: one for the stable period and one for the volatile period.
3. Realignments generally come from the volatile distribution. However, not all volatile observations are realignments.
4. Observations from the volatile distribution cluster together.
5. The standard Hamilton switching model needs six parameters to describe the distribution of EMS exchange rates, while the time varying model needs eight parameters: μ_s, μ_v, σ_s^2, σ_v^2, θ_s^0, θ_s^1, θ_v^0, θ_v^1. But even though it takes more than two parameters to describe the distribution of EMS exchange rates, the added complexity is worth it because a two-parameter distribution is clearly inadequate.
6. There is evidence that the behaviour of the exchange rate near the edge of the EMS band depends on the nature of the behaviour of the exchange rate as it approaches the band. If the exchange rate rises to the top of the band in a gradual, stable manner, large changes in the exchange rate are unlikely. However, if it jumps to the edge of the band rapidly, further volatility is likely either through a realignment or a large move back toward the centre of the band.

ACKNOWLEDGEMENTS

The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Kansas City nor of the Federal Reserve System. We thank Tim Schmidt for excellent research assistance. Bob Chirinko, Andrew Filardo, and seminar participants at the Federal Reserve Bank of Kansas City, the University of Florida, the University of Washington, Clemson University, and the System Committee on International Economic Analysis provided helpful comments.

REFERENCES

Akgiray, V. and Booth, G. G. (1988) Mixed diffusion-jump process modeling of exchange rate movements. *Review of Economics and Statistics*, **70**, 631–637.

Baillie, R. T. and Bollerslev, T. (1989) The message in daily exchange rates: a conditional variance tale. *Journal of Business and Economic Statistics*, **7**, 297-305.

Ball, C. A. and Roma, A. (1993) A jump diffusion model for the European monetary system. *Journal of International Money and Finance*, **12**, 475–492.

Bertola, G. and Caballero, R. J. (1992) Target zones and realignments. *American Economic Review*, **82**, 520–536.

Boothe, P. and Glassman, D. (1987) The statistical distribution of exchange rates: empirical evidence and economic implications. *Journal of International Economics*, **22**, 297–319.

Chen, Z. and Giovannini, A. (1993) The determinants of realignment expectations under the EMS: some empirical regularities. National Bureau of Economic Research, Working Paper no. 4291.

Diebold, F. X. (1988) *Empirical Modeling of Exchange Rates*. Springer-Verlag: New York.

Diebold, F. X., Lee, J.-H. and Weinbach, G. (1992) Regime switching with time-varying transition probabilities. Working paper, Department of Economics, University of Pennsylvania.

Engel, C. M. and Hamilton, J. D. (1990) Long swings in the dollar: are they in the data and do markets know it? *American Economic Review*, **80**, 689–713.

Filardo, A. J. (1994) Business-cycle phases and their transitional dynamics. *Journal of Business and Economic Statistics*, **12**, 299–308.

Hadi, A. S. (1992) Identifying multiple outliers in multivariate data. *Journal of the Royal Statistical Society, Series B*, **54**, 761–771.

Hadi, A. S. (1993) A modification of a method for the detection of outliers in multivariate samples. *Journal of the Royal Statistical Society, Series B*, forthcoming.

Hamilton, J. D. (1989) A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, **57**, 357–384.

Hamilton, J. D. and Susmel, R. (1992) Autoregressive conditional heteroskedasticity and changes in regime. Working paper, Department of Economics, University of California, San Diego.

Koedijk, K. G., Schafgans, M. M. A. and de Vries, C. G. (1990) The tail index of exchange rate returns. *Journal of International Economics*, **29**, 93–108.

Koedijk, K. G., Stork, P. A. and de Vries, C. G. (1992) Differences between foreign exchange regimes: the view from the tails. *Journal of International Money and Finance*, **11**, 462–473.

Mundaca, B. G. (1994) Probabilities of realignments in a target zone: a switching regime approach. Graduate School of Business, Stanford University.

Nieuwland, F. G. M. C., Verschoor, V. F. C. and Wolff, C. C. P. (1994) Stochastic trends and jumps in EMS exchange rates. *Journal of International Money and Finance*, **13**, 699–727.

Park, K., Ahn, C. M. and Fujihara, R. (1993) Optimal hedged portfolios: the case of jump-diffusion risks. *Journal of International Money and Finance*, **12**, 493–510.

Rose, A. and Svensson, L. E. O. (1993) European exchange rate credibility before the fall. National Bureau of Economic Research, Working Paper no. 4495.

Svensson, L. E. O. (1992) An interpretation of recent research on exchange rate target zones. *Journal of Economic Perspectives*, **6(4)**, 119–144.

Tucker, A. L. and Pond, L. (1988) The probability distribution of foreign exchange price changes: tests of candidate processes. *Review of Economics and Statistics*, **70**, 638–647.

Zhu, Z. (1993) Econometric analysis of a dirty float with no arbitrage: the DM-$ exchange rate after the Louvre accord. Manuscript, Department of Economics, University of Michigan.