

Have Exchange Rates Become More Closely Tied? Evidence from a New Multivariate GARCH Model

by Franc Klaassen *

Department of Economics
University of Amsterdam

July 18, 2000

Abstract

We analyze the time-dependence of exchange rate correlations using a new multivariate GARCH model. This model consists of two parts. First, we transform the exchange rate changes into their principal components and specify univariate GARCH models for all components. Second, we use the inverse of the principal components construction to transform the conditional component moments back into those of the exchange rate changes themselves. The model is easy to estimate, as it requires only univariate GARCH estimations. Nevertheless, we find it outperforms the popular constant conditional correlations and factor GARCH models. We show that the major U.S. dollar exchange rates have become more loosely instead of closely tied since the eighties.

Key words: correlations, multivariate models, GARCH, factor models, exchange rates.

JEL classification: C32, C52, F31.

*Department of Economics, University of Amsterdam, Roetersstraat 11, 1018 WB Amsterdam, the Netherlands; tel: +31-20-5254191; fax: +31-20-5254254; E-mail: Klaassen@fee.uva.nl.

I thank Frank de Jong, Harry Huizinga, Theo Nijman, Bertrand Melenberg, Siem Jan Koopman and Bas Donkers for their helpful comments.

An earlier version of this paper has appeared as CentER for Economic Research Discussion Paper No. 9910.

1 Introduction

Correlations are a key determinant of many financial decisions. For instance, investors in stocks need correlation assessments to reduce the riskiness of their portfolios, and correlations between exchange rates are important for internationally trading corporations and banks, as they have to hedge open foreign exchange positions. Several papers examine the correlations between stock returns, for instance, Bertero and Mayer (1990), Koch and Koch (1991), King, Sentana and Wadhvani (1994), Longin and Solnik (1995) and Darbar and Deb (1997). Surprisingly few papers, however, focus on exchange rate correlations. One notable example is Bollerslev (1990), who studies correlations between several European Monetary System (EMS) - U.S. dollar exchange rates. Therefore, in this paper we also focus on exchange rate correlations.

Unlike Bollerslev (1990), however, we do not restrict the correlations to be constant. The reason is that exchange rate correlations are likely time-varying according to economic intuition. For example, suppose the U.K. joins the Exchange Rate Mechanism of the EMS. Then the correlation between the pound-dollar and, say, the mark-dollar exchange rates will rise. Secondly, a change in U.S. monetary policy such as the 1979 Volcker experiment also raises that correlation, since both the pound and the mark will change in the same way against the dollar. Therefore, we allow exchange rate correlations to vary over time. We find that the correlations between eight main U.S. dollar exchange rates have decreased since the eighties, so that exchange rates have become more loosely instead of closely tied.

When modeling high-frequency exchange rates, one has to take account of the well-known conditional heteroskedasticity in such data. The literature suggests various models for that, such as GARCH (see Bollerslev, Chou and Kroner (1992) for an overview), stochastic volatility (see Ghysels, Harvey and Renault (1996)), regime-switching GARCH (see Gray (1996) and Klaassen (1998)) and fractionally integrated GARCH (see Baillie, Bollerslev and Mikkelsen (1996)). For simplicity, we take GARCH, although the approach we will develop also works for any other volatility model.

Since we want to analyze correlations, a univariate GARCH model is insufficient, and a multivariate version is called for. In this paper, we introduce a new multivariate GARCH model that is more suitable for a detailed correlation analysis than existing multivariate GARCH models, as we will explain below. The basic idea of our model stems from the fact that it is the correlations between exchange rates that make multivariate GARCH modeling more difficult than univariate GARCH. Therefore, in the first step of our approach, we remove all unconditional correlations by taking principal components of the exchange rate changes. The conditional mean and variance of each

principal component are specified by a univariate GARCH model. In the second step, the inverse of the principal components construction is used to transform the conditional moments of the principal components into the conditional mean and variance of the exchange rate changes themselves. Since this step requires no further estimation, our indirect approach makes multivariate GARCH estimation as easy as several univariate GARCH estimations.

The remaining part of this introduction presents a brief overview of the literature on multivariate GARCH and explains the contribution of this paper in more detail.

In the GARCH literature so far, extending univariate to multivariate GARCH has been a main endeavour. The reason is that one has to model not only conditional variances, but also all conditional covariances. This can easily lead to an enormous number of parameters. Hence, multivariate GARCH modeling amounts to finding a parsimonious specification of the conditional covariance matrix that does not imply an unacceptable loss of generality.

In this respect, the diagonal model of Bollerslev, Engle and Wooldridge (1988) and the BEKK model of Engle and Kroner (1995) are useful for low-variate systems. However, estimation becomes difficult for higher-variate systems. For instance, in our eight-variate empirical application, one would have to estimate more than a hundred parameters. From a computational point of view, our model is more convenient, as it requires only univariate GARCH estimations.

Another computationally attractive model is the popular Bollerslev (1990) constant conditional correlations model. For our study, however, the model is not suitable, as we want to focus on the dynamics in exchange rate correlations. As indicated above, economic intuition shows that such dynamics are very likely present. This is clearly supported by our data. In this sense, our model is preferable, as it can explain time variation in correlations, leading to a better fit.

A fourth class of existing multivariate GARCH models is factor GARCH; see Diebold and Nerlove (1989), Engle, Ng and Rothschild (1990), Ng, Engle and Rothschild (1992), King, Sentana and Wadhvani (1994) and Fiorentini, Sentana and Shephard (1998). Two reasons behind the success of factor GARCH are that such models are computationally tractable and that, in contrast to the Bollerslev (1990) model, they can capture some time-variation in the conditional correlations. However, the fit of conditional variances and correlations is not as good as that of the model we propose. The explanation will become clear in the next paragraph.

Although our model has practical advantages over existing models, it also has a sound theoretical basis. This stems from the fact that the model is a factor GARCH

model, although not in its traditional form as used above. Usual factor GARCH models are based on the theory that only a few unobserved variables, the factors, govern all exchange rates. Our model, on the other hand, uses as many factors as exchange rates. We demonstrate that this is the reason for the empirical outperformance of our model with respect to usual factor GARCH.

Taking the maximum number of factors in factor GARCH has two advantages. First, this choice turns out to be significantly optimal. Hence, our model solves a major problem of factor GARCH, namely the choice of the number of factors.

The second advantage concerns estimation. To estimate usual factor GARCH models, one commonly takes a two-step estimation method to avoid a complicated simultaneous procedure (see Engle et al. (1990), among others). Correction of the second-step standard errors for first-step estimation inaccuracy, however, is difficult and thus often ignored, leading to biased inference. Since there is no estimation in the second step of our method, we do not have this potentially serious problem.

Our model yields the following conclusions regarding the development of exchange rate correlations over time. First, we find that correlations between the main U.S. dollar exchange rates were decreasing in the years after the first oil shock, were increasing at the end of the seventies, and that they were highest in the eighties.

Second, concerning the central question of the paper, we show that exchange rates have become more loosely instead of closely tied since the eighties. This is caused by the 1992 collapse of the exchange rate mechanism of the EMS, which made several European exchange rates less correlated. Moreover, the EMS - yen correlations have become lower because of the coexistence of more stable EMS - U.S. dollar rates and a long swing in the yen - dollar rate in the nineties.

The plan for the rest of the paper is as follows. In the next section, we introduce our multivariate GARCH model. Section 3 explains why that model is a special factor GARCH model with the maximum number of factors. In section 4 we present the empirical results and analyze the time-variation in the correlations. Section 5 concludes.

2 A New Multivariate GARCH Model

In the first subsection, we develop the new multivariate GARCH model to be used for our study on exchange rate correlations. In subsection 2.2 we explain how to estimate the model. In the final subsection, we examine the implications of our model for the conditional correlations.

2.1 The Model

The basic idea behind the model is based on the fact that it is the correlations between exchange rates that make multivariate GARCH models more complex than univariate ones. Therefore, we first remove the (unconditional) correlations by transforming the exchange rate changes into their principal components. We bring in the GARCH effects through these components instead of directly through the exchange rate changes themselves. In the second step, we then transform the principal component moments into the moments of the exchange rate changes, which we are interested in.

To describe the model, we need the following notation. Let S_t denote the vector of logarithms of I nominal spot exchange rates at time t , where each exchange rate is defined as the domestic currency price of one unit of foreign currency. We concentrate on the I -vector s_t consisting of the (percentage) exchange rate changes $s_{it} = 100(S_{it} - S_{it-1})$. Thus, s_{it} is the depreciation of the domestic currency with respect to currency i . All exchange rate changes up to and including time $t-1$ form the information set I_{t-1} . Finally, we assume that s_t is conditionally normally distributed. Therefore, we only concentrate on its conditional mean and variance.

In the first part of our model, we concentrate on the I -vector of principal components defined by

$$f_t = W' s_t, \quad (1)$$

where the weighting matrix W is the unique (apart from column exchanges) orthogonal $I \times I$ eigenvector matrix of the unconditional variance $V\{s_t\}$. This transforms the correlated exchange rate changes into their (unconditionally) uncorrelated principal components.

To specify $E_{t-1}\{f_t\}$ and $V_{t-1}\{f_t\}$, the mean and variance of f_t conditional on the information set I_{t-1} , we use a standard, univariate AR(1)-GARCH(1,1) model for each principal component separately. We complete the matrix $V_{t-1}\{f_t\}$ by assuming that the off-diagonal elements are zero; this assumption is quite common in the literature (see Engle et al. (1990) and Ng et al. (1992), among others). In summary, we specify the conditional moments of f_t by

$$\begin{aligned} E_{t-1}\{f_{kt}\} &= \mu_k + \theta_k(f_{kt-1} - \mu_k) \\ V_{t-1}\{f_{kt}\} &= \omega_k + \alpha_k(f_{kt-1} - E_{t-2}\{f_{kt-1}\})^2 + \beta_k V_{t-2}\{f_{kt-1}\} \\ Cov_{t-1}\{f_{kt}, f_{lt}\} &= 0, \end{aligned} \quad (2)$$

for principal components $k, l = 1, \dots, I$, $k \neq l$. This completes the first part of the model.

In the second part of the model, we have to transform the conditional moments of the principal components into the ones for the exchange rate changes themselves, as it is the exchange rates that we are mainly interested in. This transformation is straightforward, as (1) and the orthogonality of the weighting matrix W imply that

$$\begin{aligned} E_{t-1}\{s_t\} &= WE_{t-1}\{f_t\} \\ V_{t-1}\{s_t\} &= WV_{t-1}\{f_t\}W'. \end{aligned} \tag{3}$$

This completes the second and final part of our multivariate GARCH model. Hence, the complete multivariate GARCH model is given by (1), (2) and (3).

2.2 Estimation

In this subsection we describe how to estimate our model. The first part of the model, represented by (1) and (2), can be estimated by principal components analysis on the sample covariance matrix of s_t , followed by maximum likelihood estimation of the normal univariate GARCH models for each sample principal component separately. Remarkably, this is all that is needed to estimate the model; the second part of the model, the inverse transformation (3), requires no further estimation, as the weighting matrix W has already been estimated in the first step. Hence, estimation of the full multivariate GARCH system is essentially as simple as several univariate GARCH estimations. This makes our model attractive from a practical point of view, as several other multivariate GARCH models, such as the diagonal and BEKK models mentioned in the introduction, are more difficult to estimate.

2.3 Implications for the Conditional Correlations

The focus of the paper is the development of exchange rate correlations over time. In the introduction we have argued that our model improves over the Bollerslev (1990) constant conditional correlations model in this respect, because our model allows for time-variation in the conditional correlations. However, our model also imposes some structure on the correlations. In this subsection, we examine whether this structure is reasonable.

In our model, the time-variation in the conditional correlations is completely driven by the time-variation in the conditional variances of the principal components. This follows directly from the conditional variance formula in (3) and the diagonality of $V_{t-1}\{f_t\}$.

To see whether such a structure is reasonable, consider the following stylized example. Suppose we have $I=2$ U.S. dollar exchange rate changes, namely the U.K. pound

(s_{1t}) and the German Mark (s_{2t}). Assume that both have unit unconditional variance. This implies that the principal components are

$$\begin{aligned} f_{1t} &= \sqrt{1/2} \cdot s_{1t} + \sqrt{1/2} \cdot s_{2t} \\ f_{2t} &= \sqrt{1/2} \cdot s_{1t} - \sqrt{1/2} \cdot s_{2t} = \sqrt{2} \cdot s_{1t} - f_{1t}, \end{aligned} \quad (4)$$

where the joint component f_{1t} represents the EMS - U.S. dollar exchange rate and the difference component f_{2t} represents the deviation of the U.K. pound from the EMS. Using the variance formula in (3), straightforward calculations show that the conditional correlation between the U.K. pound and the German mark exchange rate changes equals

$$\rho_{t-1}\{s_{1t}, s_{2t}\} = \frac{\frac{1}{2}V_{t-1}\{f_{1t}\} - \frac{1}{2}V_{t-1}\{f_{2t}\}}{\frac{1}{2}V_{t-1}\{f_{1t}\} + \frac{1}{2}V_{t-1}\{f_{2t}\}}. \quad (5)$$

To analyze whether this specification is reasonable, we analyze the effects of two different policy changes. First, suppose the U.K. joins the Exchange Rate Mechanism (ERM) of the EMS. Then the U.K.-EMS component f_{2t} becomes more stable, so that $V_{t-1}\{f_{2t}\}$ falls and the correlation $\rho_{t-1}\{s_{1t}, s_{2t}\}$ rises, as expected.

The second policy change we consider is a change in U.S. monetary policy, which increases the conditional variance of the U.S. dollar versus both EMS currencies. According to the model, the increase in $V_{t-1}\{f_{1t}\}$ raises the intra-EMS correlation $\rho_{t-1}\{s_{1t}, s_{2t}\}$. This is realistic, as both the pound and the mark change in the same way against the dollar after the policy shift.

Although we admit that the previous example is simple, it does show that the restrictions our model imposes on the conditional correlations are quite reasonable. In this sense, our model is preferable over the popular Bollerslev (1990) model. After all, that model restricts the conditional correlations to stay constant, even after important policy changes such as the ones discussed above.

3 Relation with Factor GARCH

In the previous section, we have seen that our model has some advantages over three existing multivariate GARCH models, namely the diagonal, the BEKK and the constant conditional correlations model. In this section we relate the model to the fourth class of existing models, namely factor GARCH. It turns out that our model is a factor GARCH model, albeit not in its traditional form. The usual factor GARCH model assumes that there are only a few factors that govern all exchange rates. In contrast, our model takes as many factors as there are exchange rates. This claim is proved in subsection 3.1.

Although our model uses many more factors than usual factor GARCH, this does not necessarily mean that our model is substantially better. Maybe the inclusion of extra factors does not lead to a much better fit and only complicates the model. In subsection 3.2 we argue that this is not the case.

3.1 A Special Factor GARCH Model

In this subsection we demonstrate that our model of subsection 2.1 is a factor GARCH model with as many factors as exchange rates. We only address the main points of this derivation; the complete derivation is in the appendix.

The central idea of a K -factor GARCH model is that there are K underlying variables, the factors, that govern all I exchange rate changes. More formally, the exchange rate innovation $\varepsilon_t = s_t - E_{t-1}\{s_t\}$ has a systematic and an unsystematic part, where the systematic part is a linear combination of K unobserved factors φ_{kt} :

$$\varepsilon_t = \Lambda\varphi_t + v_t, \quad (6)$$

where $\varphi_t = (\varphi_{1t}, \dots, \varphi_{Kt})'$ is the K -vector of common factors with a time-varying conditional covariance matrix, Λ is the $I \times K$ full-column-rank matrix of factor loadings, and v_t denotes the vector of unsystematic, exchange rate specific changes with a covariance matrix that is constant over time.

There are two problems with a direct implementation of the factor idea. The first problem is that the systematic and unsystematic innovations, φ_t and v_t , are not observed separately, so that Λ is, in general, not directly estimable. As shown by Engel et al. (1990), this problem can be solved by substituting the vector of unobserved factors φ_t by an expression based on an observed K -vector that is closely related to the factors, in the sense that the conditional variance of the k -th component of this factor representing vector is perfectly correlated with that of the k -th factor φ_{kt} (see the end of footnote 10 for a formalization of this). Similar to the existing literature (see Ng et al. (1992), among others), we take K principal components of s_t to form this factor representing vector, and we assume that they are conditionally uncorrelated and that each of them follows a normal AR(1)-GARCH(1,1) model. Hence, the factor representing vector is a K -dimensional subvector of f_t , the vector of all I principal components defined by (1) and modeled by (2). For simplicity of notation, let us denote this subvector of f_t also by f_t , and let W also denote the $I \times K$ full-column-rank matrix of component weights that defines the subvector by $f_t = W's_t$.

The second problem with a direct implementation of the factor idea is caused by a rotational indeterminacy in the factors φ_t in (6); this makes the matrix of factor

loadings Λ unidentified. The appendix shows that this problem is also present after the move to the factor representing vector. To solve the problem, we normalize Λ by $W'\Lambda = I_K$, where I_K is the $K \times K$ identity matrix. This normalization will appear to be crucial for proving the claim that our model is a factor GARCH model with $K = I$ factors.

Having solved both problems, we can derive the commonly-used K -factor GARCH formulas for the two conditional moments of interest:

$$\begin{aligned} E_{t-1}\{s_t\} &= \gamma + \Lambda E_{t-1}\{f_t\} \\ V_{t-1}\{s_t\} &= \Omega + \Lambda V_{t-1}\{f_t\} \Lambda', \end{aligned} \tag{7}$$

where γ and Ω are time-constant parts in the mean and variance, respectively. These moment specifications hold for all $K \in \{1, \dots, I\}$. Note, however, that for $K = I$ the constants γ and Ω are zero. After all, in that case $\Lambda f_t = \Lambda W' s_t = s_t$, because the normalization $W'\Lambda = I_K$ then implies that $\Lambda = (W')^{-1}$.

Although some similarities with our model of section 2 have already become clear, it may not yet be clear that our model exactly equals the I -factor GARCH model. The final link is provided by our factor GARCH normalization $W'\Lambda = I_K$ and the orthogonality of W . They imply that $\Lambda = (W')^{-1} = W$. Hence, relation (7), where γ and Ω are zero due to $K = I$, is the same as the second part of our model given by (3). Because the models for the I principal components are also the same, our model is indeed a special factor GARCH model in which the number of factors equals the number of exchange rates.

3.2 Advantages over the Usual Factor GARCH Model

From the previous subsection we know that our model uses many more factors than usual factor GARCH. In this subsection, we demonstrate that including these extra factors is useful by showing that our model overcomes two important problems with the empirical implementation of usual factor GARCH models. These problems are the choice of the number of factors and the difficult correction of standard errors in the two-step method that is commonly used to estimate factor GARCH.

The first problem is the choice of the number of factors K , or, equivalently, the number of principal components. This problem originates from a trade-off between generality and simplicity. On the one hand, increasing K leads to a more general model, but, on the other hand, it makes the model more complicated.

To alleviate this problem, there are several ad hoc criteria for selecting K (see Bartholomew (1987)). The most popular one is the Kaiser-Guttman rule, which states

that one should select only those principal components that have a larger variance than the average variance of the exchange rate changes. As all other rules, this one yields very few components. For instance, in our eight-variate empirical application, it would select only one.

To investigate whether the neglect of components is serious, we estimate the factor GARCH model for all possible K , using the exchange rate data that we will describe in subsection 4.1. The results, which are described in detail in subsection 4.5, show that using less than I components is strongly rejected. Some components turn out to be essential for a good description of the conditional variances, while other principal components, which do not improve the variance fit much, turn out to be important for the correlation fit. The usual factor GARCH model neglects many of these important components. This demonstrates the dangers involved in the popular rules for choosing the number of factors. According to our results, the correct rule is to use as many factors as possible. Since our model does exactly that and is, nevertheless, easy to estimate, it solves the problem of choosing K in usual factor GARCH.

The second problem with the empirical implementation of the usual factor GARCH model is the difficult correction of standard errors in the two-step method that is commonly used for estimation. To clarify this, we first describe this two-step method. The first step is similar to the first step of our method as described in subsection 2.2. The only difference is the number of univariate GARCH models for the principal components that one has to estimate: I in our model and K for a K -factor GARCH model, as follows from the previous subsection.

The second step in the estimation of usual factor GARCH, however, is essentially different. After substitution of the first step estimates for $E_{t-1}\{f_t\}$ and $V_{t-1}\{f_t\}$ in (7), the usual factor GARCH model requires estimation of the parameters γ , Ω and Λ to obtain estimates for the moments of interest, $E_{t-1}\{s_t\}$ and $V_{t-1}\{s_t\}$.¹ Because one uses only estimates instead of the true values of $E_{t-1}\{f_t\}$ and $V_{t-1}\{f_t\}$, the second step standard errors have to be corrected for the first step estimation inaccuracy. This

¹Most researchers use univariate techniques for this second estimation step. That is, for each exchange rate i , they use maximum likelihood based on conditional normality of s_t with mean and variance implied by the corresponding elements of (7). As Ng et al. (1992) admit, such univariate estimation sacrifices efficiency. The reason for not doing multivariate maximum likelihood is that this would lead to too many parameters to be estimated at once. After all, γ , Ω and Λ have $I + I^2 + I \cdot K$ unknown elements. This may indeed be too much, if one does not take account of all restrictions that the factor GARCH model puts on γ , Ω and Λ . These restrictions are our normalization $W'\Lambda = I_K$, which also implies $W'\gamma = 0$, and the definition of Ω (see below (14), with the additional assumption of a diagonal $V\{v_t\}$ that we use in the empirical section). They greatly reduce the number of free parameters. For instance, for $K = 7$ and $I = 8$, they lead to 16 free parameters instead of 128! Therefore, multivariate estimation is not that difficult, and we prefer it over the univariate techniques used in the literature so far.

is complicated, as Lin (1992) shows. Therefore, many authors do not correct them and use the potentially seriously biased second step standard errors for inference. In this respect, our model is preferable, because the second step is a linear transformation without any estimation (see (3)). Hence, our model involves neither difficult standard error correction, nor the use of biased standard errors.

In summary, the fact that our model employs many more factors than usual factor GARCH is very useful. First, by using the optimal number of factors, the model yields a better fit. Second, estimation is easier than for any other factor GARCH model, as our approach does not require a second estimation step.

4 Empirical Results

In this section we use our multivariate GARCH model to analyze the development of exchange rate correlations over time. First, we describe the data and motivate the choice for our model empirically. Then we estimate the model. In subsection 4.3 we address the central question of the paper, namely whether exchange rates have become more closely tied. Then we check whether the model captures the main characteristics of the data and in subsection 4.5 we compare the fit of our model with some benchmark models, namely the Bollerslev (1990) model with constant conditional correlations and factor GARCH models for all possible numbers of factors.

4.1 Data

We use U.S. dollar exchange rates of $I = 8$ currencies, namely, the Belgian franc, Canadian dollar, French franc, German mark, Italian lira, Japanese yen, Dutch guilder and the British pound. These include all major exchange rates. Moreover, some of them are highly correlated (the EMS rates), while others are much less correlated; this variety allows us to get a fairly complete picture of the behavior of the conditional correlations. We have 1,216 weekly observations for the weekly changes s_t from April 1974 to July 1997. All rates have been obtained from Datastream.

In table 1 we report some descriptive statistics; the notes below the table contain the definitions. The substantial cross-currency correlations in the first panel motivate the use of a multivariate model instead of univariate ones.

In the second panel of table 1, we test for autocorrelation in the exchange rate changes. We find significantly positive first-order autocorrelation in the core EMS exchange rate changes (we always use a significance level of 5%).² For this reason, we

²Our evidence of first-order autocorrelation is in contrast with conclusions of many earlier studies.

have allowed for a first-order autoregressive term in (2), the model for the principal components. Estimates for higher-order autocorrelations are not reported separately, but are combined in Box-Pierce type statistics \tilde{Q}_{10} ; they indicate that higher-order autoregressive terms are unnecessary.

The third panel of table 1 deals with the dynamics of the second moments. The first two rows contain measures for the time-variation in the squared exchange rate changes. Both measures point at conditional heteroskedasticity. The next two rows of panel three contain similar autocorrelation measures, but now regarding the cross products instead of the squares. Since there are seven cross products for each exchange rate series, we have taken the average to save space. The results show clear evidence of time-variation in the conditional covariances. Hence, the data motivate the use of a multivariate GARCH model.

A popular multivariate GARCH model is the Bollerslev (1990) model, which assumes that all conditional correlations for the exchange rate changes are constant over time. In the last row of table 1, we test this restriction as follows. First, we estimate a univariate GARCH model for each series of exchange rate changes and construct conditional correlation estimates by taking the product of the normalized residuals. Then we regress the estimated conditional correlations for time t on a the vector $(1, t, t^2)'$ and test whether the two slope parameters are zero (see the notes below table 1 for further details). The results show that there is clear time-dependence in the conditional correlations. This is not surprising. First, economic intuition tells us that correlations may well be time-varying (see the policy examples in subsection 2.3). Second, Bollerslev (1990) already shows that conditional correlations differ between the pre-EMS and the EMS period. In addition, Andersen, Bollerslev, Diebold and Labys (1999) also find strong empirical evidence of time-varying conditional correlations in exchange rates. This motivates why we use our model instead of the Bollerslev (1990) model, since our model can capture time-variation in the conditional correlations.

4.2 Estimation Results

In this subsection we estimate our multivariate GARCH model. As the second part of this model involves no estimation (see subsection 2.2), we only concentrate on the first part, that is, the principal components construction and the univariate GARCH

For instance, West and Cho (1995) conclude from heteroskedasticity corrected Ljung-Box statistics of orders 10, 50 and 90 that five major U.S. dollar exchange rate changes are serially uncorrelated, with the possible exception of the yen. Indeed, if we had only used the aggregate Box-Pierce type measure Q_{10} in table 1, we would have concluded the same, thereby overlooking the significant first-order autocorrelation in all core EMS exchange rate changes. Hence, our additional check for only first-order autocorrelation is useful.

estimations for each component.

To construct the principal components vector f_t , we define the weighting matrix W in (1) by the matrix of eigenvectors of the sample covariance matrix of s_t . The upper panel of table 2 presents the columns of W , which are the weighting vectors for the principal components. Each of the eight components has a name that indicates the dominating currencies in it. Hence, the components are called EMS, Jap, U.K.-EMS, Ita-EMS, Can, Fra-EMS, Bel-(Ger+Neth) and Neth-Ger. These components have been ordered according to their “explained variance”, that is, their sample variance divided by the sum of the sample variances of the individual exchange rate changes (the “total variance”). The explained variance is commonly used as a measure of importance of the principal components. It shows that the component dominated by the European currencies, the EMS component, is the most important one, explaining 77 percent of the total variance.

The remaining part in the estimation of the model concerns the estimation of the univariate GARCH models in (2) for each principal component. The results, as reported in table 3, are standard. Most importantly, they strongly reflect the presence of conditional heteroskedasticity. According to our model, this is the source of time-variation in the conditional variances as well as correlations of the individual exchange rate changes (see subsection 2.3).

4.3 Have Exchange Rates Become More Closely Tied?

Having estimated our multivariate GARCH model, we can now analyze how exchange rate correlations have evolved over the post-Bretton-Woods period. This is to answer the central question of the paper, namely whether exchange rates have become more closely tied. Note that the conclusions will be in terms of nominal exchange rates. However, they are likely to hold for real exchange rates as well, because prices are fixed in the short run.

In figure 1 we plot the estimated correlations between several dollar exchange rates.³ For the sake of exposition, we have smoothed the actual estimates by an equally weighted moving average using the estimates in the year before and the year after the week under consideration. Despite this smoothing, we still see that the correlations are not constant over time.

From figure 1 one may distinguish three remarkable periods for exchange rate cor-

³The estimates are based on the estimation results for the principal components in subsection 4.2, and the second relation in (3), which specifies the conditional variance of the exchange rate changes as a function of the conditional principal component variances.

relations, roughly spanning the seventies, eighties and the nineties. The seventies are characterized by a decrease in correlation followed by an increase. The decrease may well be caused by the rather autonomous monetary and fiscal responses of governments to the 1974-1975 period of stagflation (see Krugman and Obstfeld (1991)). This policy imbalance, however, caused a steep depreciation of the U.S. dollar, so that Germany and Japan intervened heavily in the foreign exchange market in 1977-1978. Together with the inception of the EMS in 1979, this marks a period of greater coordination, causing the correlations to rise.

The eighties characterize a period of high correlations. This is confirmed by Bollerslev (1990), who finds that correlations between the European currencies were higher during the EMS period than before. In addition, we find that also intercontinental correlations were high. This is mainly caused by the huge swing in the dollar in the eighties. First, the dollar strongly appreciated partly due to the Volcker monetary contraction starting in 1979. In the second half of the eighties, coordinated actions such as the 1985 Plaza agreement brought the dollar down again. Moreover, the 1987 Louvre target zones may also explain the high correlations in the eighties.

The third remarkable period in figure 1 concerns the decrease in the correlations in the nineties. Hence, the main exchange rates have become more loosely instead of closely tied. At first sight, this may seem surprising, as it is often believed that the greater integration of financial markets has increased financial correlations. However, more integration also means that capital can move more freely, which can destabilize exchange rates. This happened in 1992 when the EMS collapsed, leading to a drop in several intra-EMS correlations, as shown by the middle graph of the figure. Furthermore, although European and American markets have become more integrated, Japan is still a relatively independent market. This may be the reason behind the fact that swings in the EMS - U.S. dollar rates have become shorter in the nineties, while the yen - dollar swings are still relatively long (see Klaassen (1999) for empirical evidence). These differences between the European currencies and the yen have also decreased the correlations in the nineties.

With the advent of European monetary unification (EMU), it is likely that the correlations between the participating European currencies will increase again at the end of the nineties. The upward tendency in the Germany - Italy conditional correlations after 1996 may be an indication of this. It will be interesting to analyze whether EMU also affects the correlations between the world's main currencies, namely, the U.S. dollar, yen and euro.

4.4 Diagnostics

The correlation analysis in the previous subsection was based on the multivariate GARCH model of subsection 2.1. In the remaining part of this section, we check empirically whether that model is appropriate for such an analysis. In the current subsection we examine whether it captures the features of the data described in subsection 4.1. In subsection 4.5 we compare the performance of our model with that of the Bollerslev (1990) and factor GARCH models.

To check the specification of our model, we analyze the normalized residuals. They are defined by $\hat{\eta}_t = \hat{V}_{t-1}\{\varepsilon_t\}^{-1/2} \cdot \hat{\varepsilon}_t$, where $\hat{V}_{t-1}\{\varepsilon_t\}^{-1/2}$ is the inverse of the lower triangular Cholesky decomposition of $\hat{V}_{t-1}\{\varepsilon_t\}$ and ε_t is the exchange rate innovation $s_t - E_{t-1}\{s_t\}$. Table 4 presents several test results for them. The i -th column in this table concerns the i -th element of $\hat{\eta}_t$. Unfortunately, we cannot attribute this element to one country, because $\hat{\eta}_{it}$ is a linear combination of the country specific residuals $\hat{\varepsilon}_{1t}, \dots, \hat{\varepsilon}_{it}$. We conclude from the first-order autocorrelations and the Box-Pierce statistics Q_{10} that there is no evidence of remaining autocorrelation in the normalized residuals.

Secondly, the measures for remaining autocorrelation in the squared changes and the cross products show no reason to extend the variance specification of the model.

The final test in table 4 also concerns the variance specification, as it checks whether the normalized residuals are conditionally uncorrelated. This is done by regressing the cross products of the normalized residuals at time t on the vector $(1, t, t^2)'$ and testing whether all three regression coefficients are zero. The difference with the test for short-run autocorrelation in the cross products, as discussed in the previous paragraph, is that the current test has more power against long-run autocorrelation. Moreover, it also tests whether the cross products have mean zero. The results in table 4 again show no serious evidence of misspecification.

It is interesting to observe that the test for zero conditional correlations of the normalized residuals is similar to the test for constant conditional correlations of the exchange rate changes in subsection 4.1. The latter test was clearly rejected, but the test on the residuals of our model is not. Apparently, our model is able to describe the time-varying pattern in the conditional exchange rate correlations quite well. This is the main reason why we prefer our model over the Bollerslev (1990) constant conditional correlations model, as our study is focused on the time-variation in exchange rate correlations. In this respect, our model is also preferable over the usual factor GARCH model, which would be 1-factor GARCH for our data, as argued below. Although that model captures some time-variation in conditional correlations, it does not explain it

completely, as six out of eight zero-conditional-correlation statistics are significant.⁴

4.5 Goodness of Fit

In the introduction we have claimed that our multivariate GARCH model provides a good fit for the conditional exchange rate variances and correlations, at least compared to the Bollerslev (1990) and the usual factor GARCH models with much less than $I=8$ factors. In this subsection we provide evidence for that. We also examine the reasons behind the outperformance by analyzing the variance and correlation fits separately.

To measure the goodness of fit of the models, we use the multivariate normal log-likelihood with conditional mean and variance as estimated by the different models. The “total fit” column of table 5 contains the results. It shows that the log-likelihood of our model, -8,817, is better than the one of the constant conditional correlations model of Bollerslev (1990), which is -10,624.

To compare our model with the usual factor GARCH model, we first have to choose the usual number of factors, or principal components, K . The commonly used Kaiser-Guttman rule states that one should select only principal components that have a larger variance than the average variance of the exchange rate changes (see Bartholomew (1987)). For our data, this rule leads to $K = 1$, as only the variance of the EMS component (11.60, see table 2) exceeds the average variance of 1.89. This is in line with the choice of Diebold and Nerlove (1989), who use about the same exchange rates.

Table 5 demonstrates that our model is preferable over the 1-factor GARCH model, which has a log-likelihood of -10,315, as the likelihood ratio is 2,996.⁵ Hence, we conclude that our model indeed provides a better fit than the popular constant conditional correlations and 1-factor GARCH models. Note that our model also significantly outperforms the other factor GARCH models, as the likelihood ratios in table 5 show.

In the remaining part of this subsection, we investigate the reasons for this outperformance. We first analyze the variance fit and then the correlation fit.

To measure the variance fit, we remove the correlation effects from the log-likelihood by substituting the off-diagonal elements in the estimated conditional variance matrices by zero. The “variance fit” column of table 5 gives these zero-correlation log-likelihoods. Our model (-15,935) somewhat underperforms the constant conditional correlations model (-15,864). This is not surprising. The variance fit of the latter model is entirely

⁴The zero-conditional correlation tests for 1-factor GARCH are 21.34 [with p-value 0.00], 10.36 [0.02], 10.77 [0.01], 19.85 [0.00], 26.26 [0.00], 5.83 [0.12], 4.84 [0.18] and 21.52 [0.00]; see the note below table 4 for the definition of these test.

⁵The K -factor GARCH model is nested in our model, as it follows after restricting the last $K-I$ columns of the matrix of factor loadings in (6), Λ , to zero.

based on univariate GARCH estimations for each exchange rate change and the univariate estimations only have to fit the conditional variance process, while our model is mainly designed to give a good description of the correlation process.

Table 5 also shows that our model outperforms the usual 1-factor GARCH model in terms of variance fit. The reason is that the first principal component is only a single combination of exchange rate changes, and one cannot expect that this would lead to good variance estimates for all exchange rate changes individually.⁶ A good variance fit requires at least five principal components, as table 5 shows. The relevance of the fifth component, the one dominated by Canada, is shown by figure 2. For $K=4$, the conditional variance estimates for the Canadian dollar are almost constant, while only inclusion of the Can component leads to a time-variation pattern that one also finds for univariate AR(1)-GARCH(1,1) on the Canadian dollar.

The correlation fit is the second reason for the relatively good fit of our model. It is measured by the difference between the full and the zero-correlation log-likelihoods, and it is reported in the “correlation fit” column of table 5. It is clear that our model outperforms the constant conditional correlations model. This again supports the conclusion that the assumption of constant conditional correlations is too restrictive for our data.

Table 5 also demonstrates that our model provides a better correlation fit than the 1-factor GARCH model. Moreover, it also outperforms the factor GARCH model with five factors, the number of factors that is at least needed for an acceptable variance fit. Although the final three components do not improve the variance fit, they do yield a better correlation fit. In fact, adding the last component increases the log-likelihood by 420, which is highly significant.⁷ This can be attributed to a better fit of the time-variation in the conditional correlation between the Netherlands and Germany, as figure 3 demonstrates. Only the inclusion of the last component allows the factor

⁶It is interesting to observe that the lack of variance fit of the 1-factor GARCH model is hidden by the full log-likelihood, that is, the quality measure including the conditional correlations, which we have used at the beginning of this subsection. Recall that the full log-likelihood is -10,315, which is much greater than the sum of the log-likelihoods obtained from eight independent univariate AR(1)-GARCH(1,1) models for the exchange rate changes, which is -15,864. Hence, one is tempted to conclude that the 1-factor GARCH model is to be preferred; this is also what Diebold and Nerlove (1989) claim. However, the huge increase in the log-likelihood is entirely due to a better fit of the conditional correlations, and the log-likelihood is very sensitive to that (see also footnote 7). Hence, the log-likelihood of the factor model including the correlations can be a misleading indicator for the quality of the variance fit.

⁷This huge significance (likelihood ratio is 840) is due to the great sensitivity of the log-likelihood to the correlation fit. This is also the reason why $K=1$ at first sight seems to be much better than eight univariate AR(1)-GARCH(1,1) estimations on the individual exchange rates, as shown in footnote 6. It also explains why Bollerslev (1990) gets a highly significant likelihood ratio test of almost 2,000 when testing for zero correlations in his multivariate GARCH model.

GARCH model to capture that since the mid eighties the monetary policy of the Dutch central bank is mainly attributed to keeping the guilder-mark rate stable, so that both currencies move more closely together than before.

In summary, the conclusion from this subsection is that our model results in a better fit than two popular multivariate GARCH models, namely the Bollerslev (1990) model and the usual factor GARCH model. This holds especially for the correlations, which we are particularly interested in.

5 Conclusion

In this paper we analyze exchange rate correlations over time. For that, we introduce a new multivariate GARCH model. It describes the exchange rate changes indirectly through their principal components and assumes that the conditional variances of the components govern the conditional exchange rate correlations. We show that this is quite realistic, both from an economic and empirical point of view. Moreover, the indirect approach implies that the model is very easy to estimate, as it only requires several univariate GARCH estimations to estimate the full multivariate model.

The empirical results show that the model provides a better fit than existing models. First, it outperforms the popular constant conditional correlations model of Bollerslev (1990) with respect to the correlation fit. This is not surprising, as the data show clear evidence of time-variation in the conditional correlations and only our model can capture that. Second, our model provides a better variance and correlation fit than usual factor GARCH models. This is explained by the fact that our model can be viewed as a factor GARCH model with the maximum number of factors and that the factors neglected in usual factor GARCH contain important information for exchange rate variances and correlations.

Given the outperformance qua fit, we use our model to analyze the correlations between eight U.S. dollar exchange rates over the post-Bretton-Woods period. We find that these correlations were highest in the eighties and then decreased in the nineties. Hence, exchange rates have become more loosely instead of closely tied. This originates from the EMS crash in 1992, making several European exchange rates less correlated. Moreover, the EMS - yen correlations have decreased because of the combination of more stable EMS - U.S. dollar rates and a long swing in the yen - dollar rate.

So far, we have concentrated on GARCH in a multivariate setting. However, it is important to realize that our indirect approach via the principal components is not restricted to GARCH. In fact, any univariate model for the principal components can be used to derive a practical multivariate model. This offers a wide range of applications of

our approach. For instance, when analyzing stock or bond return correlations, one can take account of asymmetric volatilities, GARCH-in-mean effects and other deviations from standard GARCH (see Bollerslev et al. (1992)). Furthermore, our approach can form the basis for multivariate extensions of other volatility models, such as stochastic volatility, regime-switching GARCH and fractionally integrated GARCH. This is left for future research.

Appendix: Our Model is a Special Factor GARCH Model

In this appendix we demonstrate that our model of subsection 2.1 is a factor GARCH model with as many factors, K , as exchange rates, I . For that, we first define what we actually mean by the K -factor GARCH model. As in the main text, we concentrate on the conditional mean and variance of exchange rate changes. The final factor GARCH specification of these moments is derived in two stages.

To obtain the first factor GARCH formulation, we split the vector of exchange rate changes s_t into

$$s_t = \mu_t + \varepsilon_t, \quad (8)$$

where $\mu_t = E_{t-1}\{s_t\}$ and ε_t is the innovation. The central idea of the factor model is that ε_t has a systematic and an unsystematic part, where the systematic part is a linear combination of K unobserved factors φ_{kt} :

$$\varepsilon_t = \Lambda\varphi_t + v_t, \quad (9)$$

where $\varphi_t = (\varphi_{1t}, \dots, \varphi_{Kt})'$ is the K -vector of common factors, Λ is the $I \times K$ full-column-rank matrix of factor loadings, and v_t denotes the unsystematic, exchange rate specific change. We assume that $E_{t-1}\{\varphi_t\} = 0$ and $E_{t-1}\{v_t\} = 0$ to ensure $E_{t-1}\{\varepsilon_t\} = 0$. Moreover, let $V_{t-1}\{\varphi_t\}$ denote the time-varying conditional variance of φ_t .⁸ Let $V_{t-1}\{v_t\}$ be the variance of v_t , which we assume constant over time ($V_{t-1}\{v_t\} = V\{v_t\}$), as in Engel et al. (1990). Finally, we have $Cov_{t-1}\{\varphi_t, v_t\} = 0$.

The main effect of the factor model is that it puts structure onto the innovation ε_t . However, as in Engle et al. (1990), the factor idea can also be used to specify the expected exchange rate changes μ_t . This makes μ_t the sum of a systematic part, which is attributed to the factors, and an unsystematic part. More formally,

$$\mu_t = \Lambda\mu_t^\varphi + \mu_t^v, \quad (10)$$

where the systematic part is a linear combination of a K -vector of common sources of expected depreciation, μ_t^φ , and the unsystematic part is an I -vector of exchange rate specific expected depreciations, which we assume constant over time ($\mu_t^v = \mu^v$).

Specifications (8), (9) and (10) lead to the first formulation of the factor GARCH model in terms of the moments of interest:

$$\begin{aligned} E_{t-1}\{s_t\} &= \Lambda\mu_t^\varphi + \mu^v \\ V_{t-1}\{s_t\} &= \Lambda V_{t-1}\{\varphi_t\} \Lambda' + V\{v_t\}. \end{aligned} \quad (11)$$

⁸Note that we do not impose diagonality of $V_{t-1}\{\varphi_t\}$. Diagonality has been commonly used in the literature to help identify Λ . Later on, we will introduce another, very convenient way to identify Λ .

This holds for all $K \in \{1, \dots, I\}$. Note that for $K = I$, the case we are particularly interested in, the parameters μ^v and $V\{v_t\}$ are zero, because in that case $\varepsilon_t(\mu_t)$ is one-to-one related to $\varphi_t(\mu_t^\varphi)$.

The factor model in its current format cannot be estimated because of two problems. The first one is that the systematic and unsystematic innovations, φ_t and v_t , are not observed separately, so that the parameters are, in general, not directly estimable. The second problem is caused by a rotational indeterminacy in the definition of the factors, which makes Λ unidentified. We now solve both problems in turn, so as to derive the second factor GARCH moment specification.

As shown by Engle et al. (1990), the first problem can be solved by substituting the unobserved factors φ_t by an expression based on an observed K -vector that is closely related (but not equal) to the factors in a sense that is explained at the end of footnote 10. Similar to many other papers (for instance, see Ng et al. (1992)), we take K principal components of s_t to form this factor representing vector, and we assume that they are conditionally uncorrelated and that each of them follows an AR(1)-GARCH(1,1) model. Hence, the factor representing vector is a K -dimensional subvector of f_t , the vector of all I principal components described by (1) and (2). For simplicity of notation, let us denote this subvector of f_t also by f_t , and let W also denote the $I \times K$ full-column-rank matrix of component weights that defines the subvector by $f_t = W' s_t$. Using (11), this implies that

$$\begin{aligned} E_{t-1}\{f_t\} &= W' \Lambda \mu_t^\varphi + W' \mu^v \\ V_{t-1}\{f_t\} &= W' \Lambda V_{t-1}\{\varphi_t\} \Lambda' W + W' V\{v_t\} W. \end{aligned} \quad (12)$$

Since $W' \Lambda$ is invertible, we can solve μ_t^φ and $V_{t-1}\{\varphi_t\}$ from these equations and substitute the results in (11). This gives

$$\begin{aligned} E_{t-1}\{s_t\} &= \Lambda(W' \Lambda)^{-1} E_{t-1}\{f_t\} - \Lambda(W' \Lambda)^{-1} W' \mu^v + \mu^v \\ V_{t-1}\{s_t\} &= \Lambda(W' \Lambda)^{-1} V_{t-1}\{f_t\} (\Lambda' W)^{-1} \Lambda' - \Lambda(W' \Lambda)^{-1} W' V\{v_t\} W (\Lambda' W)^{-1} \Lambda' + V\{v_t\}. \end{aligned} \quad (13)$$

The main difference with (11) is that (13) contains only parameters related to the unsystematic innovation v_t , not related to the factors φ_t , as the observable f_t has taken the role of φ_t . Therefore, using the principal components has solved the first problem.

The second problem with (11) is caused by a rotational indeterminacy in the unobserved factors, so that Λ is not identified. That is, if a certain combination of Λ , μ_t^φ and φ_t gives the true conditional moments of s_t , then, for any invertible $K \times K$ -matrix Q , the oblique rotations ΛQ , $Q^{-1} \mu_t^\varphi$ and the oblique factors $Q^{-1} \varphi_t$ yield the same conditional moments. Formula (13) shows this problem again. Since Λ only occurs in the

combination $\Lambda(W'\Lambda)^{-1}$, it is only identified if we can derive its $I \cdot K$ unknown elements from a particular value of $\Lambda(W'\Lambda)^{-1}$, say A . However, this is impossible, since there are only $I \cdot K - K^2$ independent equations in $\Lambda(W'\Lambda)^{-1} = A$.⁹ Therefore, we need K^2 normalizing restrictions on Λ . Considering (13), it is very convenient to use $W'\Lambda = I_K$, where I_K is the $K \times K$ identity matrix.¹⁰ We will see below that this normalization is crucial for proving that our model is an I -factor GARCH model.

Having solved both problems, we can present the second and final factor GARCH formulation, which is commonly used in the literature:

$$\begin{aligned} E_{t-1}\{s_t\} &= \gamma + \Lambda E_{t-1}\{f_t\} \\ V_{t-1}\{s_t\} &= \Omega + \Lambda V_{t-1}\{f_t\}\Lambda', \end{aligned} \quad (14)$$

where $\gamma = (I_I - \Lambda W')\mu^v$ and $\Omega = V\{v_t\} - \Lambda W'V\{v_t\}W\Lambda'$ are the time-constant parts in the mean and variance, respectively. Note that these parts are zero in case of $K = I$, because then μ^v and $V\{v_t\}$ are zero.

Although some similarities with our model of section 2 have already become clear, it may not yet be clear that our model exactly equals the factor GARCH model for $K = I$. The final link is provided by our factor GARCH normalization $W'\Lambda = I_K$. In case of $K = I$, this normalization and the orthogonality of W imply that $\Lambda = (W')^{-1} = W$. Therefore, relation (14), where γ and Ω are zero because of $K = I$, is the same as the second part of our model, that is, formula (3). Because the model for the I principal components is also the same, our model is indeed a special factor GARCH model in which the number of factors equals the number of exchange rates.

⁹The system $\Lambda(W'\Lambda)^{-1} = A$ is equivalent to $(I_I - AW')\Lambda = 0$, where I_I is the identity matrix of dimension I . To compute the rank of $I_I - AW'$, we first note that AW' is idempotent, since $W'A = I_K$. Hence, the rank of $I_I - AW'$ is $r(I_I - AW') = I - r(AW')$. Moreover, $r(AW') = K$, since both A and W' have rank K . Therefore, the rank of $I_I - AW'$ is $I - K$, so that the system $(I_I - AW')\Lambda = 0$ contains exactly $(I - K) \cdot K$ independent equations.

¹⁰This normalization has three interesting characteristics. First, it directly reduces the number of free parameters, which makes estimation simpler. For instance, for $K = 7$ and $I = 8$, it implies that only seven factor loadings have to be estimated instead of 56.

The second characteristic of our normalization is that it is necessary and sufficient. This is in contrast with the sufficient identifying restrictions employed by Sentana (1992) and King, Sentana and Wadhvani (1994), who impose $V\{\varphi_t\} = I_K$ for the conditional variance of the factors to identify Λ (except for column sign).

Finally, our normalization explains in what sense the principal components are ‘‘closely related’’ to the factors. Using $W'\Lambda = I_K$ in the conditional variance of f_t , which is $V_{t-1}\{f_t\} = W'\Lambda V_{t-1}\{\varphi_t\}\Lambda'W + W'V\{v_t\}W$, shows that the conditional variance of each component f_{kt} is perfectly correlated with that of the k -th factor φ_{kt} . This is why the f_{kt} are called ‘‘factor representing portfolios’’ in Engle et al. (1990).

References

- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (1999), "The Distribution of Exchange Rate Volatility," NBER Working Paper No. 6961.
- Baillie, R.T., T. Bollerslev and H.O. Mikkelsen (1996), "Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 74, 3-30.
- Bartholomew, D.J. (1987), *Latent Variable Models and Factor Analysis*, London: Charles Griffin.
- Bertero, E. and C. Mayer (1990), "Structure and Performance: Global Interdependence of Stock Markets around the Crash of October 1987," *European Economic Review*, 34, 1155-1180.
- Bollerslev, T. (1990), "Modelling the Coherence in Short-Run Nominal Exchange Rates: A Multivariate GARCH Approach," *Review of Economics and Statistics*, 72, 498-505.
- Bollerslev, T., R.Y. Chou and K.F. Kroner (1992), "ARCH Modeling in Finance," *Journal of Econometrics*, 52, 5-59.
- Bollerslev, T., R.F. Engle and J.M. Wooldridge (1988), "A Capital Asset Pricing Model with Time-Varying Covariances," *Journal of Political Economy*, 96, 116-131.
- Darbar, S.M. and P. Deb (1997), "Co-Movements in International Equity Markets," *Journal of Financial Research*, 20, 305-322.
- Diebold, F.X. and M. Nerlove (1989), "The Dynamics of Exchange Rates Volatility: A Multivariate Latent Factor ARCH Model," *Journal of Applied Econometrics*, 4, 1-21.
- Engle, R.G. and K.F. Kroner (1995), "Multivariate Simultaneous Generalized ARCH," *Econometric Theory*, 11, 122-150.
- Engle, R.F., V.K. Ng and M. Rothschild (1990), "Asset Pricing with a Factor-ARCH Covariance Structure: Empirical Estimates for Treasury Bills," *Journal of Econometrics*, 45, 213-237.
- Fiorentini, G., E. Sentana and N. Shephard (1998), "Exact Likelihood-Based Estimation of Conditionally Heteroskedastic Factor Models," Paper presented at Econometric Society European Meeting 98, Berlin.
- Ghysels, E., A. Harvey and E. Renault (1996), "Stochastic Volatility," in: Maddala, G.S. and C.R. Rao (eds.), *Handbook of Statistics*, 14, Amsterdam: North Holland, 119-191.
- Gray, S.F. (1996), "Modeling the Conditional Distribution of Interest Rates as a Regime-Switching Process," *Journal of Financial Economics*, 42, 27-62.

- King, M., E. Sentana and S. Wadhvani (1994), "Volatility and Links between National Stock Markets," *Econometrica*, 62, 901-933.
- Klaassen, F.J.G.M. (1998), "Improving GARCH Volatility Forecasts," CentER for Economic Research, Tilburg University, Discussion Paper No. 9852.
- Klaassen, F.J.G.M. (1999), "Purchasing Power Parity: Evidence from a New Test," CentER for Economic Research, Tilburg University, Discussion Paper No. 9909.
- Koch, P.D. and T.W. Koch (1991), "Evolution in Dynamic Linkages across National Stock Indexes," *Journal of International Money and Finance*, 10, 231-251.
- Krugman, P.R. and M. Obstfeld (1991), *International Economics: Theory and Policy*, New York: Harper Collins.
- Lin, W.-L. (1992), "Alternative Estimators for Factor GARCH Models - A Monte Carlo Comparison," *Journal of Applied Econometrics*, 7, 259-279.
- Longin, F. and B. Solnik (1995), "Is the Correlation in International Equity Returns Constant: 1960-1990?" *Journal of International Money and Finance*, 14, 3-26.
- Ng, V., R.F. Engle and M. Rothschild (1992), "A Multi-Dynamic-Factor Model for Stock Returns," *Journal of Econometrics*, 52, 245-266.
- Pagan, A.R. and G.W. Schwert (1990), "Alternative Models for Conditional Stock Volatility," *Journal of Econometrics*, 45, 267-290.
- Sentana, E. (1992), "Identification of Multivariate Conditionally Heteroskedastic Factor Models," London School of Economics, Discussion Paper No. 139.
- West, K.D. and D. Cho (1995), "The Predictive Ability of Several Models of Exchange Rate Volatility," *Journal of Econometrics*, 69, 367-391.
- White, H. (1980), "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity," *Econometrica*, 48, 817-838.

Table 1: Moments of exchange rate changes and autocorrelation tests

	Bel	Can	Fra	Ger	Ita	Jap	Neth	U.K.
Mean	0.00	-0.03	-0.02	0.03	-0.08	0.07	0.02	-0.03
Variance	2.11	0.38	2.06	2.14	2.06	2.11	2.09	2.13
Skewness	-0.26	-0.57	-0.25	-0.14	-0.59	0.53	-0.14	-0.40
Excess kurtosis	1.92	6.78	2.34	1.70	6.38	2.01	1.90	3.00
Cross-currency corr. $\bar{\rho}$	0.71	0.13	0.71	0.72	0.64	0.47	0.73	0.59
Autocorr. s_{it} : ρ_1	0.07* (0.03)	0.01 (0.04)	0.06* (0.03)	0.07* (0.03)	0.02 (0.04)	0.05 (0.04)	0.07* (0.03)	0.04 (0.04)
Autocorr. s_{it} : \tilde{Q}_{10}	14.82 [0.14]	7.35 [0.69]	14.63 [0.15]	14.07 [0.17]	11.78 [0.30]	22.57* [0.01]	13.31 [0.21]	6.05 [0.81]
Autocorr. s_{it}^2 : ρ_1^s	0.09* (0.03)	0.15* (0.03)	0.05 (0.03)	0.04 (0.03)	0.10* (0.03)	0.20* (0.03)	0.07* (0.03)	0.21* (0.03)
Autocorr. s_{it}^2 : Q_{10}^s	48.66* [0.00]	36.53* [0.00]	52.49* [0.00]	57.60* [0.00]	134.20* [0.00]	92.03* [0.00]	56.24* [0.00]	151.82* [0.00]
Autocorr. $s_{it} \cdot s_{jt}$: $\bar{\rho}_1^c$	0.07* (0.03)	0.07* (0.03)	0.07* (0.03)	0.05 (0.03)	0.08* (0.03)	0.05 (0.03)	0.06* (0.03)	0.06* (0.03)
Autocorr. $s_{it} \cdot s_{jt}$: \bar{Q}_{10}^c	57.01* [0.00]	18.46* [0.05]	61.79* [0.00]	57.76* [0.00]	61.95* [0.00]	34.64* [0.00]	55.81* [0.00]	82.21* [0.00]
Constancy of conditional corr.	19.76* [0.00]	21.70* [0.00]	22.34* [0.00]	20.30* [0.00]	21.38* [0.00]	17.94* [0.00]	20.38* [0.00]	22.61* [0.00]

Standard errors in parentheses and p-values in square brackets; * is significant at 5% level.

The correlation measure $\bar{\rho}$ is the average of the sample correlation coefficients of the series under consideration with all seven other series.

The first-order autocorrelation, ρ_1 , is estimated as the slope coefficient in a regression of the change in exchange rate i , s_{it} , on the first lagged change, s_{it-1} , and a constant. The standard errors are based on White's (1980) heteroskedasticity-consistent asymptotic covariance matrix.

\tilde{Q}_{10} denotes a modified Box-Pierce type statistic that combines the first ten autocorrelations. Following Pagan and Schwert (1990), it is defined as the sum of the first ten squared normalized autocorrelation estimates, where the normalizing factors are the heteroskedasticity-consistent standard errors of the autocorrelation estimates. \tilde{Q}_{10} is asymptotically χ_{10}^2 distributed.

The first-order autocorrelation in the squared changes, ρ_1^s , and the Box-Pierce type statistic for the squared changes, Q_{10}^s , are similarly defined as ρ_1 and \tilde{Q}_{10} , respectively, although without the heteroskedasticity correction.

The seven first-order autocorrelations of the cross products $s_{it} \cdot s_{jt}$ ($j \neq i$) are averaged to save space; this average is denoted by $\bar{\rho}_1^c$. The number in parentheses is also the average standard error. Similarly, \bar{Q}_{10}^c denotes the mean of the seven Box-Pierce type statistics of the cross products; its p-value is based on a χ_{10}^2 distribution.

We test for constancy of the conditional correlations $\rho_{t-1}\{s_{it}, s_{jt}\}$ by testing the constancy of $Cov_{t-1}\{\varepsilon_{it}, \varepsilon_{jt}\} / (V_{t-1}\{\varepsilon_{it}\}V_{t-1}\{\varepsilon_{jt}\})^{1/2}$, where ε_{it} is the innovation in a univariate normal-AR(1)-GARCH(1,1) model for s_{it} . The test amounts to regressing the estimated correlation, $\hat{\varepsilon}_{it}\hat{\varepsilon}_{jt} / (\hat{V}_{t-1}\{\varepsilon_{it}\}\hat{V}_{t-1}\{\varepsilon_{jt}\})^{1/2}$, on a constant, t and t^2 , and then computing the Wald statistic for no effect of t and t^2 . For space considerations, we only report the average over the seven possible Wald statistics for each i . The critical values are based on a χ_2^2 distribution.

Table 2: Principal component weights

	EMS	Jap	U.K. EMS	Ita EMS	Can	Fra EMS	Bel G+N	Neth Ger
Belgium	0.41	-0.09	-0.23	-0.22	0.03	-0.40	0.75	0.03
Canada	0.03	-0.05	0.12	0.10	0.99	-0.02	-0.01	-0.01
France	0.40	-0.07	-0.15	-0.05	0.03	0.88	0.19	-0.03
Germany	0.42	-0.06	-0.23	-0.23	0.02	-0.18	-0.44	-0.70
Italy	0.36	-0.19	-0.02	0.89	-0.11	-0.13	-0.02	-0.01
Japan	0.28	0.94	0.14	0.09	0.01	-0.03	0.02	0.01
Netherlands	0.41	-0.07	-0.22	-0.20	0.03	-0.12	-0.46	0.71
U.K.	0.34	-0.22	0.89	-0.19	-0.11	-0.03	-0.00	-0.00
Variance	11.60	1.31	0.85	0.58	0.36	0.22	0.13	0.04
Expl. variance	76.87	8.70	5.65	3.83	2.41	1.43	0.83	0.27

Each column contains the weights of the individual exchange rates changes in the sample principal components. The eight weighting vectors, named according to the dominating currencies, form the weighting matrix W in (1). Hence, W is the matrix of eigenvectors of the sample covariance matrix of the exchange rate changes (normalized at length one, so that the “weights” do not sum to one).

“Variance” denotes the sample variance of a principal component, which is equal to the corresponding eigenvalue.

“Expl. variance” denotes the percentage of the total variance explained by a principal component, that is, the sample variance of the component divided by the sum of the sample variances of the individual exchange rate changes (called the “total variance”).

Table 3: Estimation results for the principal components

		EMS	Jap	U.K. EMS	Ita EMS	Can	Fra EMS	Bel G+N	Neth Ger
Mean	μ	-0.01 (0.08)	0.10* (0.03)	-0.03 (0.03)	-0.04* (0.02)	-0.01 (0.02)	-0.02* (0.01)	-0.00 (0.00)	-0.00 (0.00)
Autocorr.	θ	0.06 (0.03)	0.07* (0.03)	0.09* (0.03)	-0.00 (0.04)	0.02 (0.03)	-0.15* (0.04)	-0.30* (0.03)	-0.25* (0.03)
Cond. var. intercept	ω	0.15 (0.08)	0.08* (0.03)	0.16* (0.05)	0.09* (0.01)	0.04* (0.01)	0.00* (0.00)	0.00* (0.00)	0.00 (0.00)
ARCH	α	0.16* (0.03)	0.13* (0.03)	0.10* (0.03)	0.38* (0.06)	0.16* (0.03)	0.33* (0.05)	0.26* (0.03)	0.02* (0.00)
GARCH	β	0.85* (0.03)	0.82* (0.05)	0.72* (0.08)	0.51* (0.05)	0.74* (0.05)	0.80* (0.02)	0.81* (0.01)	0.98* (0.00)
Log-likelihood		-3131	-1836	-1597	-1128	-1070	-596	-112	652

Standard errors in parentheses; * is significant at 5% level.

The estimated model is (2) without the conditional covariance equation. To start-up the conditional variance, we use a separate parameter, which is not reported. Standard errors are not corrected for the fact that we use only an estimate of the weighting matrix W , because our focus is on the conditional moments of the exchange rate changes, not the intermediate GARCH estimation results for the principal components.

Table 4: Diagnostic statistics for normalized residuals

	$i=1$	$i=2$	$i=3$	$i=4$	$i=5$	$i=6$	$i=7$	$i=8$
Autocorr. $\hat{\eta}_{it}$: ρ_1	0.06* (0.03)	0.02 (0.03)	-0.04 (0.03)	0.05 (0.03)	0.00 (0.03)	-0.00 (0.03)	-0.02 (0.03)	-0.04 (0.03)
Autocorr. $\hat{\eta}_{it}$: Q_{10}	19.83 [0.03]	9.16 [0.52]	10.69 [0.38]	5.42 [0.86]	8.29 [0.60]	18.19 [0.05]	13.33 [0.21]	17.24 [0.07]
Autocorr. $\hat{\eta}_{it}^2$: ρ_1^s	0.01 (0.03)	0.02 (0.03)	-0.00 (0.03)	0.03 (0.03)	0.01 (0.03)	0.06* (0.03)	0.14* (0.03)	0.03 (0.03)
Autocorr. $\hat{\eta}_{it}^2$: Q_{10}^s	11.65 [0.31]	4.05 [0.95]	1.37 [0.99]	2.88 [0.98]	6.52 [0.77]	13.04 [0.22]	26.81 [0.00]	3.71 [0.96]
Autocorr. $\hat{\eta}_{it} \cdot \hat{\eta}_{jt}$: $\overline{\rho_1^c}$	0.03 (0.03)	0.03 (0.03)	0.02 (0.03)	0.02 (0.03)	0.02 (0.03)	0.01 (0.03)	0.03 (0.03)	-0.01 (0.03)
Autocorr. $\hat{\eta}_{it} \cdot \hat{\eta}_{jt}$: $\overline{Q_{10}^c}$	17.00 [0.07]	10.67 [0.38]	6.09 [0.81]	7.77 [0.65]	9.84 [0.45]	11.06 [0.35]	7.41 [0.69]	14.91 [0.14]
Zero conditional correlation	10.77* [0.01]	2.76 [0.43]	3.07 [0.38]	4.26 [0.23]	7.05 [0.07]	3.64 [0.30]	5.65 [0.13]	4.94 [0.18]

Standard errors in parentheses and p-values in square brackets; * is significant at 5% level.

The vector of normalized residuals is $\hat{\eta}_t = \hat{V}_{t-1}\{\varepsilon_t\}^{-1/2} \cdot \hat{\varepsilon}_t$, where $\hat{V}_{t-1}\{\varepsilon_t\}^{-1/2}$ is the inverse of the lower triangular Cholesky decomposition of $\hat{V}_{t-1}\{\varepsilon_t\}$ and ε_t is the exchange rate innovation $s_t - E_{t-1}\{s_t\}$. Hence, $\hat{\eta}_{it}$ is a linear combination of $\hat{\varepsilon}_{1t}, \dots, \hat{\varepsilon}_{it}$, so that, in contrast to $\hat{\varepsilon}_{it}$, $\hat{\eta}_{it}$ does not directly correspond to one country.

All autocorrelation statistics have been defined below table 1, although the standard error of ρ_1 and the value of Q_{10} are no longer corrected for heteroskedasticity.

The test for zero conditional correlation between $\hat{\eta}_{it}$ and the other seven $\hat{\eta}_{jt}$ is similar to the constant correlation test of table 1. However, now we also test for a zero intercept in the regressions involved. Hence, the critical value is based on a χ_3^2 instead of χ_2^2 distribution.

Table 5: Quality of various multivariate GARCH models

Model	TOTAL FIT			VARIANCE FIT		CORREL. FIT	
	log-lik.	change	LR	log-lik.	change	log-lik.	change
Univar. GARCH	-15,864	0	-	-15,864	0	0	0
Const. cond. corr.	-10,624	5420	10,840*	-15,864	0	5240	5240
1-factor GARCH	-10,315	309	-	-16,078	-214	5763	523
2-factor GARCH	-10,248	67	134*	-16,037	41	5789	67
3-factor GARCH	-10,145	103	206*	-16,013	24	5867	103
4-factor GARCH	-9,836	309	618*	-15,958	55	6122	309
5-factor GARCH	-9,796	40	80*	-15,916	42	6120	40
6-factor GARCH	-9,579	217	434*	-15,937	-19	6358	217
7-factor GARCH	-9,237	342	684*	-15,936	1	6699	342
Our model	-8,817	420	840*	-15,935	1	7118	420

A * denotes significance at the 5% level.

The quality measure we use is the log-likelihood based on a normally distributed vector of exchange rate changes with conditional mean and variance as estimated by the different models. In the “total fit” column, the full estimated conditional variance matrix is used to compute this log-likelihood. For the “variance fit” column, the conditional correlations have been substituted by zero. The “correlation fit” column is the difference between the “total fit” and “variance fit” columns.


The “total fit” column also contains the likelihood ratio (LR) for the model against the previous one, if the model includes the previous one as a special case.

“Univar. GARCH” is the model that imposes diagonality of the conditional variance matrix, so that the moments can be estimated by eight univariate GARCH procedures.


“Const. cond. corr.” denotes the Bollerslev (1990) model with constant conditional correlations. It is estimated in two steps. First, we estimate eight univariate GARCH models, and then we derive the conditional correlation estimates.

For the K-factor GARCH models, the conditional mean and variance follow from the multivariate second estimation step (see section 3.2 and footnote 1). For parsimony, we assume that the covariance matrix of the exchange rate specific changes v_t in (6), $V\{v_t\}$, is diagonal, as in Diebold and Nerlove (1989).

Canada versus Germany and Japan



Germany versus Italy, the Netherlands and the U.K.



Japan versus Canada and Germany




Figure 1: Smoothed estimated conditional correlations between dollar exchange rates

Conditional variance of Canada




Figure 2: Effect of Can principal component on the estimated conditional variance of Canada

Conditional correlation between the Netherlands and Germany




Figure 3: Effect of Neth-Ger principal component on the estimated conditional correlation between the Netherlands and Germany