IMF *Staff Papers*
Vol. 40, No. 3 (September 1993)

Portfolio Performance of the SDR and Reserve Currencies

Tests Using the ARCH Methodology

MICHAEL G. PAPAIOANNOU and TUGRUL TEMEL*

In evaluating their foreign exchange exposure, international investors often compare actual portfolios with those calculated under the assumption that the variability of returns on various currency assets is time invariant. This paper uses autoregressive conditional heteroskedastic (ARCH) models to test that assumption. For major reserve currencies, including the SDR, we find evidence that the variances of returns do vary over time and that ARCH models that specify changing variances are superior to models that assume constant variance. By incorrectly assuming a constant variability of returns, the error introduced is smaller with the SDR than with any other national currency. [JEL G11, C22]

IN CALCULATING optimal fixed-income portfolios, the standard assumption is that historical volatilities of returns from multicurrency investments in government securities remain constant over time. It is an assumption that is often used both by investors to determine the "best" allocation of their wealth in terms of fixed-income investment instruments and by some central banks to manage their external reserve assets. The same assumption is also made in the capital asset pricing model and

* Michael G. Papaioannou is an Economist in the Treasurer's Department and holds a Ph.D. from the University of Pennsylvania. Tugrul Temel, a doctoral candidate in applied economics at the University of Minnesota, was a summer intern in 1991 in the Treasurer's Department.
The authors would like to thank Lawrence K. Duke, Szeina Lurie, Inci Otker, Anthony Richards, Orlando Roncesvalles, George Tavlas, George Tsetsekos, and Michael Wattleworth for helpful comments and suggestions.

modern portfolio theory when calculating the diversification benefits from investing in multiple assets; it also enters into the Black-Scholes model when calculating option pricing formulas. The implications of incorrectly assuming a constant variability of returns are a loss of efficiency in terms of portfolio allocations and inappropriate decisions in terms of hedging the associated risk. Recent empirical work has shown that large changes in equity returns and exchange rates, using high-frequency data, tend to be followed by large changes of sign; that is, the variance of returns and exchange rates is not constant, and the traditional assumption of constancy in many standard models of portfolio allocation does not hold. Very few studies, however, have addressed the issue of returns from multicurrency investments in fixed-income assets.

A suitable model for dealing with nonconstant volatilities of returns is the ARCH model.[1] This paper attempts to model total fixed-income investment returns for a U.S. dollar–based investor as a univariate autoregressive conditional heteroskedastic process, as an alternative to the more common portfolio model that assumes the variance of returns from multicurrency-denominated assets to be stationary or time invariant. The ARCH model expresses the current conditional variance of a time series of total returns as a linear function of its past residuals squared. ARCH models can be used to predict changes in the conditional variance of asset returns on different currencies but not the changes in the unconditional mean value of asset returns. Knowledge of the statistical properties that generate the observed rates of return on various investments—in particular how the process depends on developments in the financial markets—is important for characterizing asset return volatility and for modeling optimal portfolios of multicurrency assets.

Traditional models assume a constant one-period forecast variance and covariance matrix for the returns on different investments. This assumption can be generalized to introduce a new class of stochastic processes called ARCH processes. These processes imply serially uncorrelated errors with zero means and nonconstant variances conditional on past returns. To test whether the error terms follow an ARCH process, the autocorrelation of the squared ordinary-least-squares (OLS) residuals has been calculated. After the detection of ARCH effects based on the autocorrelation statistics, a factor ARCH model can be used to reduce the number of parameters to be estimated.

Our empirical results show that the variability of returns from investing in various currencies and currency composites is affected by lagged residuals, although the number of lags differs among currencies. The

[1] See Engle and Rothschild (1992).

estimated constant terms in the ARCH specification, α_0, are significant and nonnegative for all currencies, indicating a positive autonomous conditional volatility for the returns on these assets. Further, the estimated parameters of the lagged squared residuals in the ARCH process, α_i, exhibit strong statistical significance for returns on all currency assets, indicating that the stochastic processes generating these returns have nonstationary variances. Investment in SDRs, however, provides greater insulation to investors than investing in any other currency with one lag of ARCH effects, as SDR returns demonstrate the lowest transmission of unexpected weekly shocks into volatility.[2] In these circumstances and as already noted, an ARCH specification is a more efficient approach for modeling these asset returns than is the OLS approach. Finally, the Lagrange multiplier statistic for first-order ARCH specification is significant for all currencies, which also indicates that the conditional variances are not constant through time.

In addition, the results of several diagnostic checks on the distributional properties of total returns indicate that they deviate from the normal distribution. In particular, the kurtosis statistics on the probability distribution of returns on all currency investments are substantially larger than those from a standard normal distribution, indicating that these returns are leptokurtic (fat tailed). This result is further supported by Bera-Jarque tests, which indicate not only a highly significant nonnormality of the probability function for returns on these currencies but also a significant likelihood that the variance of these returns is not constant over time. The BJ test for normality[3] is highly significant for the deutsche mark, the pound sterling, the French franc, the Japanese yen, the European currency unit (ECU), the SDR, and gold with ARCH(1, 1) as well as for the U.S. dollar with ARCH(6, 2). Only the BJ test for the average composite index with ARCH(1, 1) is not statistically significant.

I. Brief Literature Review

The uncertainty of speculative prices has been observed to change through time (Mandelbrot (1963) and Fama (1965)). The tendency of large (small) price changes in high-frequency financial data to be followed by other large (small) price changes is often called "volatility clustering."

[2] Although the estimated coefficient α_1 for the average composite index (ACI) is lower than that for the SDR, the sum of the estimated coefficients α_i ($i = 1, \ldots, 5$) for the ACI process is greater than the estimated coefficient α_1 for the SDR.

[3] For the use of this test, see Table 2.

One specification that has emerged for characterizing such changing variances is the ARCH model (Engle (1982)) and various extensions of it. In his seminal paper, Engle suggests that one possible parametrization for variances is to express them as a linear function of squared past values of the errors of the model. With financial data, the ARCH model captures the tendency for volatility clustering, and numerous empirical applications of the ARCH methodology to asset return variances and covariances have already appeared in the literature. ARCH effects have generally been found to be highly significant in equity markets—for example, highly significant test statistics for ARCH models have been reported for various individual stock returns (Engle and Mustafa (1989)).

In a series of papers, the ARCH model has been analyzed, generalized, extended to the multivariate context, and used to test for time-varying risk premia in financial markets. These papers include Engle (1983) and Engle and Kraft (1983). The ARCH model was extended to the multivariate framework in Kraft and Engle (1983). Diebold and Nerlove (1989) use a multivariate approach exploiting factor structure, a method that facilitates tractable estimation via a substantial reduction in the number of parameters to be estimated. The factor structure captures commonality in the volatility of movements in financial data over time. Engle, Ng, and Rothschild (1990) suggest using the factor ARCH model to describe the conditional covariance matrix of excess asset returns. One- and two-factor ARCH models have been applied to the pricing of treasury bills, with the results showing reasonable stability over time.

The importance of ARCH models in finance comes partly from the direct association of variance and risk and from the fundamental trade-off between risk and return. The ARCH-M (in mean) model developed by Engle, Lilien, and Robins (1987) expresses the conditional mean as a function of the conditional variance of the ARCH process. It thus provides a tool for estimating the (possibly) linear relationship between the returns and variances of a given investment portfolio.

Another interesting observation in financial data is that the unconditional price and returns distributions tend to have fatter tails than the normal distribution (they are leptokurtic) (Mandelbrot (1963) and Fama (1965)). Ignoring the fat tails and the time-varying variances could lead to erroneous detection of abnormal returns (De Jong, Kemna, and Kloeck (1990)). Stock returns tend to exhibit nonnormal unconditional sampling distributions, both in the form of skewness and excess kurtosis (Fama (1965)). For many financial time series, the leptokurtosis cannot be fully explained. The conditional normality assumption in ARCH generates some degree of unconditional excess kurtosis, but it is typically less than adequate to account fully for the fat-tailed properties of the data. One solution to the kurtosis problem is the adoption of conditional distributions with tails fatter than the normal distribution. Skewness and kurtosis are shown to be important in characterizing the conditional density function of returns on stocks (Engle and Gonzales-Rivera (1989)).

In addition to the leptokurtic distribution of stock return data, Black (1976) noted a negative correlation between current returns and future volatility. The linear generalized ARCH(p, q) model, where the variance depends only on the magnitude and not on the sign of the estimated residuals, cannot capture this negative relationship since the conditional variance is linked only to past conditional variances and squared innovations; hence, the sign on returns does not affect the volatilities. This limitation of the standard ARCH formulation is one of the primary motivations for the exponential GARCH(p, q) model by Nelson (1991). In this type of generalized ARCH model, the volatility depends not only on the magnitude of the past surprises but also on their corresponding signs.

Bollerslev, Engle, and Wooldridge (1988) first estimated a multivariate GARCH(p, q) process for returns on bills, bonds, and stocks where the expected return was assumed to be proportional to the conditional covariance of each return from a fully diversified or market portfolio. They found that conditional covariances are quite variable over time and are a significant determinant of the time-varying risk premia. Bollerslev (1990) also applied a multivariate GARCH model to short-run nominal exchange rates. Hall, Miles, and Taylor (1988) similarly undertook a multivariate GARCH-M estimation of the capital asset pricing model. Pagan and Schwert (1990) compared several statistical models for monthly stock return volatility and showed the importance of nonlinearities not captured by conventional ARCH or GARCH models. They also showed the nonstationarity of the volatility of stock returns. Vries (1991) compared stable and GARCH methods of modeling financial data and showed that the unconditional distribution of variables generated by a GARCH-like process, which models the clustering of volatility and exhibits the fat-tail property, can be stable. Further, Bollerslev (1987) extended the ARCH and GARCH models, called an integrated GARCH model, to explain the persistence of variance over time.

II. Estimation and Testing for ARCH Effects

Autoregressive conditional heteroskedasticity models account for the empirical observation that large changes in asset returns tend to be followed by further large changes in these returns, although the sign is unpredictable; also, small changes in asset returns tend to be followed by small changes, thus leading to periods of persistently high or low volatility. We utilize the ARCH model by assuming that the conditional mean, $E(y_t|y_{t-1})$, of each asset return, y_t, is linearly unpredictable, while its conditional variance, h_t, is predictable by an ARCH model. It is well known that the unconditional distribution of ARCH processes have thicker tails, even though their conditional distributions are normal.[4]

The first-order autoregressive model, AR(1), for y_t, combined with ARCH(1, 1) errors, can be written as[5]

$$y_t = \phi_0 + \phi_1 y_{t-1} + \epsilon_t, \tag{1}$$

$$\text{where } E(\epsilon_t|y_{t-1}) = 0,$$

$$\text{var}(\epsilon_t|y_{t-1}) = E(\epsilon_t^2|y_{t-1}) = h_t, \text{ and}$$

$$h_t = \alpha_0 + \alpha_1 \epsilon_{t-1}^2. \tag{2}$$

For the regression to be stable, $|\phi_1|$ is assumed to be less than one. To ensure that h_t is positive and that the process is stationary, we must have $\alpha_0 \geq 0$ and $1 \geq \alpha_1 \geq 0$. Intuitively, α_0 can be interpreted as a "normal" stationary element in the series on asset returns and α_1 as the heteroskedastic coefficient, which shows the correlation between changes in returns and their corresponding volatilities in the previous period. As an illustration, for a U.S. dollar–based investor, y_t would denote asset returns in U.S. dollars obtained by investing in various currency assets; it would be defined as the product of the three-month treasury bill rate for each asset multiplied by the percentage change of the corresponding U.S. dollar exchange rate, all expressed on a weekly basis.[6]

The generalization of the ARCH(p, q) model for asset return y_t is as follows:

$$y_t = \phi_0 + \phi_1 y_{t-1} + \cdots + \phi_p y_{t-p} + \epsilon_t, \tag{3}$$

where $\phi_p(L) = \sum_{i=1}^{p} \phi_i L^i$ is the polynomial autoregressive lag operator. Note that

$$y_t \sim N[\phi_p(L)y_t, h_t] \tag{4}$$

$$\text{with } h_t = \alpha_0 + \sum_{i=1}^{q} \alpha_i \epsilon_{t-i}^2, \tag{5}$$

[4]See Engle (1982).

[5]That is, an AR(1) model means that y_t depends on y_{t-1} but not on earlier $(y_{t-2}, y_{t-3}, \ldots)$ observations. An ARCH(1, 1) model means that y_t depends on y_{t-1} but not on earlier $(y_{t-2}, y_{t-3}, \ldots)$ observations and that h_t depends on ϵ_{t-1}^2 but not on earlier squared errors $(\epsilon_{t-2}^2, \epsilon_{t-3}^2, \ldots)$.

[6]In calculating returns, we have omitted capital gains or losses resulting from changes in the prices of bills, on the assumption that such short-term securities generally have relatively small price variations over the course of a week.

where ϵ_t follows a white noise process defined by $E(\epsilon_t) = 0$ for all t, $E(\epsilon_t \epsilon_s) = 0$ for $t \neq s$, and $E(\epsilon_t^2) = \sigma^2$ for all t. The conditional variance of ϵ_t, h_t, is allowed to vary over time and is a linear function of past residuals squared.

Testing for ARCH effects requires a test of the null hypothesis that the conditional variance h_t is constant over time:

$$H_0\colon \alpha_1 = \alpha_2 = \cdots = \alpha_q = 0 \text{ (no ARCH effects);}$$

$$H_1\colon \alpha_1 \neq \alpha_2 \neq \cdots \neq \alpha_q \neq 0 \text{ (ARCH effects exist).}$$

To test the null hypothesis H_0, we have utilized the following procedure: first, the residuals $\hat{\epsilon}_t$ are estimated from equation (3) using ordinary least squares; second, the number of lags of $\hat{\epsilon}_t$ is determined, using all available information, through maximization of its log-likelihood function; and third, estimates of $\alpha_0, \ldots, \alpha_q$ are obtained by applying OLS to

$$h_t = \alpha_0 + \sum_{i=1}^{q} \alpha_i \hat{\epsilon}_{t-i}^2.$$

In order to test for an ARCH process of order q, we form the regression

$$E(\hat{\epsilon}_t^2) = h_t = \alpha_0 + \sum_{i=1}^{q} \alpha_i \hat{\epsilon}_{t-i}^2, \tag{6}$$

where $\hat{\epsilon}_t^2$ is the square of the OLS residuals obtained from equation (3).[7] To test for higher-order ARCH specifications and the joint significance of the coefficients attached to equation (6), the Lagrange multiplier (LM) statistic is calculated as $T \cdot R^2$ (where R^2 is the coefficient of determination and T is the sample size), which has an asymptotic $\chi^2(q)$ distribution. Significantly large values of the statistic will lead to the rejection of the null hypothesis H_0 of homoskedasticity in favor of an ARCH process of order q.

Finally, if ARCH effects are detected, maximum likelihood estimation rather than OLS should be applied,[8] and the additional assumption of (conditional) normality of returns should be employed in the third step of the above procedure.

III. Empirical Results

In this section, the univariate ARCH model using weekly data is estimated for a U.S. dollar–based investor. The data set contains returns on short-term securities denominated in the U.S. dollar, the deutsche

[7]See Diebold (1988, p. 19); Greene (1990, p. 417).
[8]See Engle (1982, pp. 996–99).

Table 1. *ARCH(p,q) Estimation Results for Total Returns, Weekly Data, 1982:5–1991:50*

$(y_t = \phi_0 + \phi_1 y_{t-1} + \cdots + \phi_p y_{t-p} + \epsilon_t \text{ and } h_t = \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \cdots + \alpha_q \epsilon_{t-q}^2)$

	U.S. dollar	Deutsche mark	Japanese yen	Pound sterling	French franc	SDR	ECU	Average composite index	Gold
$\hat{\phi}_0$	0.001	0.126	0.135	0.118	0.136	0.082	0.095	0.016	−0.059
	(2.889)*	(2.262)*	(2.406)*	(2.038)*	(2.458)*	(2.883)*	(1.714)*	(1.939)*	(−0.062)
$\hat{\phi}_1$	1.191	0.331	0.347	0.311	0.332	0.330	0.302	0.362	0.248
	(27.026)*	(8.087)*	(7.394)*	(7.679)*	(8.119)*	(7.500)*	(7.134)*	(7.829)*	(5.454)*
$\hat{\phi}_2$	−0.291	...	...	...	...	...	...	−0.105	...
	(−4.597)*							(−2.094)*	
$\hat{\phi}_3$	0.212	...	...	...	...	...	...	0.102	...
	(4.184)*							(1.999)*	
$\hat{\phi}_4$	0.009	...	...	...	...	...	...	−0.039	...
	(0.211)							(−0.816)	
$\hat{\phi}_5$	0.055	...	...	...	...	...	...	...	...
	(1.565)								
$\hat{\phi}_6$	−0.187	...	...	...	...	...	...	...	...
	(−8.734)*								
$\hat{\alpha}_0$	0.000	1.355	1.25	1.374	1.270	0.369	1.163	0.018	3.329
	(11.199)*	(12.767)*	(15.332)*	(14.310)*	(14.431)*	(13.072)*	(10.963)*	(5.368)*	(22.932)*
$\hat{\alpha}_1$	0.538	0.144	0.17	0.181	0.197	0.114	0.203	0.106	0.194
	(9.665)*	(2.473)*	(4.742)*	(3.870)*	(3.107)*	(2.002)*	(3.065)*	(1.726)*	(4.369)*
$\hat{\alpha}_2$	0.162	...	...	...	...	...	...	0.074	...
	(2.944)*							(1.467)**	
$\hat{\alpha}_3$	...	...	...	...	...	...	...	0.108	...
								(1.452)**	
$\hat{\alpha}_4$	...	...	...	...	...	...	...	0.130	...
								(2.179)*	
$\hat{\alpha}_5$	...	...	...	...	...	...	...	0.037	...
								(0.708)	
Log likelihood	2,234.24	−845.77	−829.02	−857.23	−841.03	−503.10	−755.69	158.17	−1,081.92
Lagrange multiplier	112.72*	10.99*	6.09*	8.18*	19.15*	4.45*	10.21*	16.69*	0.60

Note: * and ** denote statistically significant *t*-statistics at the 0.05 and 0.10 levels, respectively.

mark, the Japanese yen, the pound sterling, the French franc, the SDR, the ECU, an average composite index (ACI), and gold. The ACI represents the average share of currencies in official holdings of foreign exchange reserves of IMF member central banks for the 1981–87 period. Our sample spans the period February 1982 to December 1991. The data set is described in the Appendix.

For given values of past-realized investment returns, α_0 and α_i ($i = 1, \ldots, 5$) were estimated for the total return on each of the nine types of currency assets (see Table 1). All estimated coefficients, $\hat{\alpha}_0$ and $\hat{\alpha}_i$, are nonnegative, thus ensuring that a positive variance is obtained. Furthermore, the summation of α_i's ($i \geq 1$) is always less than one, indicating that the estimated variance is finite. The optimal lag length, q, is determined as the value that maximizes the log-likelihood function in a grid search over lags 1 to 6. To maximize the likelihood function, an iterative procedure based on the method of BHHH[9] was used. As can be seen from Table 1, the optimal lag length is *one* for the deutsche mark, the pound sterling, the French franc, the Japanese yen, ECU, SDR, and gold; *two* for the U.S. dollar; and *five* for the average composite index. The difference in the optimal lag length between the U.S. dollar and the rest of the employed currencies, except the average composite index, may be attributed to the fact that U.S. dollar returns comprise only interest earnings on the respective assets, without taking into account any counteracting movements in the exchange rate as is the case with the other currencies. The ARCH coefficients for almost all currency returns are significantly greater than zero, as indicated by the asymptotic t-statistics given in parentheses.

The specification test for qth-order ARCH disturbances is based on the qth-order autocorrelation of the squared residuals. The presence of first-order ARCH effects is accepted at the 5 percent level for the deutsche mark, the Japanese yen, the pound sterling, the French franc, SDR, and ECU returns. The null hypothesis of no second-order ARCH effects is rejected for returns on the U.S. dollar; that of no fifth-order ARCH effects is rejected for average composite investment returns. These results appear to provide statistical evidence that the variances of the respective currency returns are not independent over time, in that they are related through the square of past residuals. With regard to gold returns, the evidence is inconclusive, since the t-statistic of the corresponding ARCH coefficient is statistically significant but the LM statistic suggests the absence of ARCH effects.

[9]The Berndt and others (1974) iterative algorithm.

Table 2. *Preliminary Data Analysis on Total Returns*[a]

	Skewness	Kurtosis[b]	Bera-Jarque	Augmented Dickey-Fuller
U.S. dollar	0.998	10.506	4817.583*	−20.929*
Deutsche mark	0.281	3.133	9.197*	−17.363*
Japanese yen	0.899	3.190	220.964*	−22.918*
Pound sterling	0.443	3.218	65.608*	−12.918*
French franc	0.074	3.264	14.496*	−17.542*
SDR	−0.226	3.081	4.519*	−12.293*
ECU	−0.209	3.282	10.524*	−16.671*
Average composite index	0.070	3.245	2.689	−12.381*
Gold	0.651	3.254	1158.280*	−22.367*

Note: * indicates that statistic is significant at the 5 percent level.
[a] The number of observation is 496.
[b] For the calculation of the kurtosis statistic, see Diebold (1988, pp. 8–11).

The significance tests for the α_i parameters and the LM statistics are corroborated by several additional diagnostic tests on the distributional properties of total returns (see Table 2). These statistics indicate that most currency returns, except for the ECU and the SDR, display a negative skewness. This characterizes a distribution of returns in which negative changes are less than positive ones. A small negative skewness indicates that large negative changes are outnumbered by positive ones of smaller magnitude. Moreover, all of the currency returns considered here exhibit high leptokurtosis, which commonly appears in high-frequency exchange rate data.[10] As mentioned above, when the conditional distribution is normal with ARCH disturbances, its unconditional distribution is known to be leptokurtic,[11] implying that an ARCH model could partly eliminate the associated leptokurtosis. The kurtosis statistics for all currency returns are significantly larger than three, which is inconsistent with a standard normal distribution of total returns on these assets.

In addition, the Bera-Jarque (BJ) test statistic[12] for the deutsche mark, the pound sterling, the French franc, the Japanese yen, ECU, SDR, and gold with ARCH(1, 1) is highly significant. Furthermore, the BJ tests for the U.S. dollar with ARCH(6, 2) and for the average composite index with ARCH(4, 5) are also significant. The resulting high negative values

[10] See Diebold (1988, p. 5).
[11] See Engle (1982); Diebold (1988, pp. 8–11); Kroner and Claessens (1991, p. 136).
[12] The Bera-Jarque test statistic is approximately distributed as a central $\chi^2(2)$ under the null hypothesis of normality in the underlying distribution of returns (see Hendry (1989, pp. 32–33)).

for the augmented Dickey-Fuller tests imply that the characteristic roots of the Taylor expansion of the unconditional variance of ϵ_t^2 lie outside the unit circle for all currency returns.[13]

The finding that the LM statistic for the testing of ARCH effects is significant for all currency returns (except gold) suggests that investment in these currencies will result in increased estimated variances—and therefore in increased risk during periods of large unexpected shocks to returns and in diminished risk during periods of relative stability. This result indicates that it might be desirable to split the sample into periods of high and low variability of returns and examine whether ARCH effects persist. The values for the parameter α_0 range from zero for U.S. dollar returns to 3.329 for gold, and all are statistically significant. The parameter α_i ($i = 1, \ldots, 5$) represents the relative contribution of the previous period's squared error of returns to their conditional variance in the current period. The parameter values of α_1 range from 0.106 for the average composite index to 0.538 for U.S. dollar returns, and all are statistically significant at the 5 percent level. In the case of the U.S. dollar, this means that about 54 percent of a week's disturbances is carried over into the following week. Among currency returns displaying one lag of ARCH effects, returns on the SDR demonstrate the lowest transmission of weekly deviations from the mean into volatility of SDR returns, while ECU returns had the highest. This result suggests that shorter-term investment periods are likely to prove more "successful" for the SDR than for the ECU and the rest of the currencies, in terms of achieving a predicted risk-return trade-off. The result for the ECU may be attributed to the fact that the currencies in its basket have higher volatilities than those in the SDR basket.

The constant term, ϕ_0, takes into account a possible nonstationary element (drift), like a time trend, in total returns. Under the null hypothesis, currency returns would allow an autoregressive structure with constant-variance residuals. Under the alternative, an ARCH structure on the residuals is imposed. During the sample period, the constants on all currency returns, except that of gold, were positive, and only those on gold returns were not significantly different from zero. These terms may be interpreted as mean weekly returns. Thus, the intercept (con-

[13]The augmented Dickey-Fuller test is a stationarity test, in which the coefficient of the lagged dependent variable is tested to see whether it is equal to or greater than unity (unit-root test). A coefficient that is less than one indicates a stationary time series. If the augmented Dickey-Fuller statistic is significant, the existence of a characteristic root outside the unit circle cannot be rejected, and therefore the respective autoregressive series is considered nonstationary. The statistic is robust to ARCH specification.

stant) for deutsche mark returns indicates that the annualized average weekly returns on deutsche mark investments were 10.28 percent.[14] Further, since for all currency returns ϕ_1 is less than one, it would seem that all returns are stable. Note that U.S. dollar returns exhibit a different pattern of generating process from other currency returns owing to the fact that returns on U.S. assets consist only of the respective interest earnings. Such dollar returns at time t are found to be affected by a six-lag structure, although the coefficients on the fourth and fifth lags are not statistically significant. However, the coefficient on the sixth lag of the dollar returns process is highly significant, indicating an apparent cyclicality to U.S. interest rates of a similar periodicity (a six-week cycle). Since the sum of the six coefficients in the lag structure is 0.989, U.S. dollar returns also appear to be stable.

For most currency returns, the lagged squared errors contribute sizably to the conditional variance, as indicated by the significant α_i coefficients, which imply that currency returns are generated by processes with nonconstant unconditional variance. Because the principal assumption of constant variances of the error terms in an OLS estimation is violated, the latter methodology should not be used as an estimation procedure for currency returns. Furthermore, the coefficient structure in the ARCH model of U.S. dollar returns is monotonically decreasing, indicating that a squared error from the previous week has a greater effect on current conditional variance than a squared error from two weeks ago.

An important observation on the statistical properties of the stochastic process followed by SDR returns is that the conditional variance of the SDR returns displays an optimal lag length of one, as only the coefficient on the first-period (squared) residual is statistically significant. Thus, the data indicate an ARCH process of order one. The estimated SDR returns have finite variance since the coefficient on the lagged residuals is less than one. The resulting ARCH effects for the SDR returns are further supported by their distributional properties of high leptokurtosis and the highly significant Bera-Jarque tests indicating a nonnormal distribution

[14]The implied annualized average weekly returns are calculated as

$$\left[\left(1 + \frac{\bar{y}}{100}\right)^{52} - 1\right]100, \qquad \text{where } \bar{y} = \phi_0 \Big/ \sum_{i=1}^{6} \phi_i.$$

Accordingly, the implied annualized average weekly returns for U.S. dollar investments are 4.84 percent, Japanese yen investments 11.34 percent, pound sterling investments 9.30 percent, French franc investments 11.15 percent, SDR investments 6.56 percent, ECU investments 7.33 percent, ACI investments 1.23 percent, and gold 0.96 percent, during the period February 1982 to December 1991.

for these returns. These empirical findings indicate that the conditional variance of the SDR returns is not constant over time and that optimization of a portfolio that includes the SDR requires the use of techniques such as the ARCH methodology to measure risk or variance accurately.

IV. Conclusions

This paper provides empirical evidence that would permit investors, including central banks, to make more accurate calculations of optimal strategies in the management of their reserves. The gains in accuracy depend directly on how the variance, or risk, measures of the returns on various currency investments change over time. To test whether the variability of returns is constant, we employ the ARCH model, which was further used to estimate the variability of the returns on various currency investments.

On the basis of weekly data for returns on nine types of currency and gold investments for the period February 1982 to December 1991, we show that the variances of total returns on all currencies and gold change significantly over time. Evidence of positive and highly significant autocorrelation of the weekly returns on those assets and of their predictable level of variability is not consistent with the assumption of constant (time-invariant) risk that is often used in asset allocation models. The empirical results presented in this paper generally suggest that for a U.S. dollar–based investor an ARCH model that takes into account the nonstationarity of return variability provides more accurate estimates of the variability of total returns on multicurrency investments than ordinary least squares models.

In particular, the sign of the intercept in the ARCH specifications is nonnegative for all currencies, indicating that the variability of asset returns would increase in the absence of any influence from the previous period's returns. Another interesting empirical result is that the volatility of returns from investing in the deutsche mark, the pound sterling, the French franc, the Japanese yen, the ECU, the SDR, and gold is affected by the square of the residuals lagged one period, while the variability of returns on U.S. dollar assets is affected by the square of the residuals lagged two periods and that of average composite index by the square of the residuals lagged five periods. This result means that for individual currencies the future effect of changes in variability lasts only a short period and diminishes over time. The variability of returns on U.S. investments may be attributed to the fact that they comprise only interest earnings on the respective assets. Interest rates, as opposed to exchange rates, tend to exhibit longer persistence: they are influenced by longer lags. For a currency-composite investment, however, the effect lasts for a significantly longer period.

The returns on investments in the SDR and ECU follow an ARCH process of order one, and their probability distributions appear to be highly leptokurtic. These results are broadly similar to those for individual currencies. In particular, the effect of the past-period error on the variance of SDR returns is smaller than that found for any of the other currency returns, largely owing to the built-in diversification of currencies in the SDR basket. In contrast, ECU returns exhibit the highest transmission of unexpected weekly shocks into returns volatility during the sample period. The difference between the statistical results for the SDR and those for the ECU could be attributed to the fact that the ECU contains more volatile currencies than those in the SDR. Further study of the SDR and ECU returns might shed some light on more fundamental differences between these two composite reserve assets.

APPENDIX

Calculation of Asset Returns

This Appendix describes the various calculations used to develop our returns data.

Asset Returns in Different Currencies

Since exchange rates are expressed bilaterally against the U.S. dollar, our analysis is dollar based. The same procedure could be repeated for the mark-, yen-, or franc-based investors. We define dollar-based returns as follows:

$$y_{t,k} = [(1 + i_{t-1,k})(1 + e_{t,k}) - 1]100,$$

where $e_{t,k} = (E_{t,k} - E_{t-1,k})/E_{t-1,k}$; and

$y_{t,k}$ = U.S. dollar–based returns on investments in currency k at time t;

$i_{t-1,k}$ = three-month treasury bill rate, on a weekly basis, for the country corresponding to currency k at time $t-1$;

$e_{t,k}$ = rate of appreciation or depreciation of currency k against the U.S. dollar at time t;

$E_{t,k}$ = spot rate, defined as U.S. dollar per unit of local currency k at time t.

The source for exchange rate and interest rate data for the various countries is the IMF Treasurer's Department data base.

Gold Returns in Different Currencies

The formula used in the calculation of returns on gold is

$$g_t = [(G_t - G_{t-1})/G_{t-1}]100,$$

where g_t and G_t denote, respectively, the weekly percentage change in the U.S. dollar price of gold and the level of the U.S. dollar price of gold at time t. Further, we define

$$GR_{t,k} = g_t + e_{t,k}\,100,$$

where $GR_{t,k}$ denotes the weekly gold returns in U.S. dollars by investors of currency k at time t. The source for gold prices is the IMF Treasurer's Department data base.

Average Composite Exchange and Interest Rates

The hypothetical currency basket consists of seven currencies: the U.S. dollar (US), the pound sterling (PS), the deutsche mark (DM), the French franc (FF), the Swiss franc (SF), the Netherlands guilder (NG), and the Japanese yen (JY). A prespecified amount[15] of these currencies and an average of their exchange rates (the noon quotations from London in terms of U.S. dollar) for the 1981–87 period are used to calculate the weight of each currency in the average composite exchange and interest rate.

Calculation of Weights for the Composite Exchange and Interest Rates

	Amount of currency in the basket	Exchange rate (noon quotation from London)		Weights = amount/rate	
U.S. dollar	0.7217	$(US/US =)$	1.000	$(w1 =)$	0.721654
Pound sterling	0.0118	$(PS/US =)$	0.420	$(w2 =)$	0.028120
Deutsche mark	0.2775	$(DM/US =)$	1.974	$(w3 =)$	0.140601
French franc	0.0576	$(FF/US =)$	4.560	$(w4 =)$	0.012631
Swiss franc	0.0423	$(SF/US =)$	1.781	$(w5 =)$	0.023759
Netherlands guilder	0.0225	$(NG/US =)$	2.142	$(w6 =)$	0.010526
Japanese yen	12.7213	$(JY/US =)$	202.870	$(w7 =)$	0.062706

The formulas by which the average composite exchange rates (ACE) and the average composite interest rates (ACI) are computed follow:

$$\text{ACE}_t = [w1(US/US)_t + \cdots + w7(JY/US)_t]$$

$$\text{ACI}_t = [w1(USI)_t + w2(PSI)_t + \cdots + w7(JYI)_t]$$

where $(USI), (PSI), \ldots, (JYI)$ denote three-month treasury bill interest rates in the United States, United Kingdom, . . . , Japan, respectively.

[15] It represents the average share of currencies in total official holdings of foreign exchange reserves for the 1981–87 period.

REFERENCES

Berndt, E.K., and others, "Estimation Inference in Nonlinear Structural Models," *Annals of Economic and Social Measurement*, Vol. 3 (1974), pp. 653–65.

Black, Fisher, "Studies in Stock Price Volatility Changes," *Proceedings of the 1976 Business Meeting of the Business and Economic Statistics Section, American Statistical Association* (1976), pp. 177–81.

Bollerslev, Tim, "A Conditional Heteroskedastic Time Series Model for Speculative Prices and Rates of Return," *Review of Economics and Statistics*, Vol. 69 (1987), pp. 542–47.

———, "Modelling the Coherence in Short-Run Nominal Exchange Rates: A Multivariate Generalized ARCH Model," *Review of Economics and Statistics*, Vol. 72 (1990), pp. 498–505.

Bollerslev, Tim, Robert F. Engle, and Jeffrey M. Wooldridge, "A Capital Asset Pricing Model with Time-Varying Covariances," *Journal of Political Economy*, Vol. 96 (1988), pp. 116–31.

De Jong, Frank, Angelien Kemna, and Teun Kloeck, "The Impact of Option Expirations on the Dutch Stock Market" (unpublished, Erasmus University, 1990).

Diebold, Francis X., *Empirical Modeling of Exchange Rate Dynamics* (Berlin: Springer-Verlag, 1988).

Diebold, Francis X., and Marc Nerlove, "The Dynamics of Exchange Rate Volatility: A Multivariate Latent Factor ARCH Model," *Journal of Applied Econometrics*, Vol. 4 (1989), pp. 1–21.

Engle, Robert F., "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica*, Vol. 50 (1982), pp. 987–1007.

———, "Estimates of the Variance of U.S. Inflation Based on the ARCH Model," *Journal of Money, Credit, and Banking*, Vol. 15 (1983), pp. 286–301.

Engle, Robert F., Victor Ng, and Michael Rothschild, "Asset Pricing with a Factor-ARCH Covariance Structure," *Journal of Econometrics*, Vol. 45 (1990), pp. 213–37.

Engle, Robert F., and Dennis F. Kraft, "Multiperiod Forecast Error Variances of Inflation Estimated from ARCH Models," in *Applied Time Series Analysis of Economic Data*, ed. by A. Zellner (Washington: Bureau of the Census, 1983), pp. 293–302.

Engle, Robert F., David M. Lilien, and Russell P. Robins, "Estimating Time-Varying Risk Premia in the Term Structure: The ARCH-M Model," *Econometrica*, Vol. 55 (1987), pp. 391–407.

Engle, Robert F., and Chowdhury Mustafa, "Implied ARCH Models from Options Prices" (unpublished; University of California, San Diego, 1989).

Engle, Robert F., and Gloria Gonzales-Rivera, "Semiparametric ARCH Models" (unpublished; University of California, San Diego, 1989).

Engle, Robert F., and Michael Rothschild, eds., "ARCH Models in Finance," *Journal of Econometrics*, Vol. 52 (April/May 1992).

Fama, Eugene F., "The Behavior of Stock Market Prices," *Journal of Business*, Vol. 38 (1965), pp. 34–105.

Greene, William H., *Econometric Analysis* (New York: Macmillan, 1990).

Hall, Steven G., David K. Miles, and Mark Taylor, "A Multivariate GARCH in Mean Estimation of the Capital Asset Pricing Model," Discussion Paper Technical Series No. 19 (London: Bank of England, 1988).

Hendry, David F., *PC Give: An Interactive Econometric Modelling System, Version 6.0/6.01* (Great Britain: Bidolles Ltd., Guildford and King's Lynn, January 1989).

Kraft, Dennis F., and Robert F. Engle, "Autoregressive Conditional Heteroskedasticity in Multiple Time Series" (unpublished; University of California, San Diego, 1983).

Kroner, F. Kenneth, and Stijn Claessens, "Optimal Dynamic Hedging Portfolios and the Currency Composition of External Debt," *Journal of International Money and Finance*, Vol. 10 (1991), pp. 131–48.

Mandelbrot, Benoit B., "The Variation of Certain Speculative Prices," *Journal of Business*, Vol. 36 (1963), pp. 394–419.

Nelson, Daniel B., "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica*, Vol. 59 (1991), pp. 347–70.

Pagan, Adrian R., and G. William Schwert, "Alternative Models for Conditional Stock Volatility," *Journal of Econometrics*, Vol. 45 (1990), pp. 267–90.

Vries, Casper G., "On the Relation Between GARCH and Stable Processes," *Journal of Econometrics*, Vol. 48 (1991), pp. 313–24.