Bank of England

Discussion Papers

Technical Series

No 13

The arch model as applied to the study of international asset market volatility

by

R R Dickens

March 1986

TS 04 611831 INT.BANK 01 COPY

MR CLEWS
HO-4

No 13

The arch model as applied to the study of international asset market volatility

by

R R Dickens

March 1986

The object of this Technical Series of Discussion Papers is to give wider circulation to econometric research work predominantly directed towards revising and updating the various Bank models. Any comments should be sent to the author at the address given below.

The author has been indebted to Professor Charles Goodhart who initiated this research project, has overseen its progress, and provided helpful comments. A discussion with Professor Robert Engle was much appreciated, as were comments received from Professor Peter Robinson. Ongoing discussions with Dr David Barr have been invaluable. The usual disclaimers apply. The views expressed are the author's own and not necessarily those of the Bank of England.

Issued by the Economics Division, Bank of England, London, EC2R 8AH to which requests for individual copies and applications for mailing list facilities should be addressed; envelopes should be marked for the attention of the Bulletin Group.

©Bank of England 1986
ISBN 0 903312 80 8
ISSN 0263-6123

CONTENTS

THE ARCH MODEL AS APPLIED TO THE STUDY OF INTERNATIONAL ASSET MARKET VOLATILITY

1 Introduction

The aim of this paper is to set out how Engle's [1982] autoregressive conditional heteroskedastic (ARCH) model has been adopted for the investigation of UK and international asset market volatility. It pertains to the results reported in Dickens [1985, 1986], which are the first two papers stemming from a continuing project attempting both to describe and to explain yield and price volatility in asset markets in the UK and in other major industrial countries. The first of the two papers concentrates on describing volatility in selected UK asset markets, while the second paper both describes and compares asset market volatility in the UK, US, Germany, Japan, Italy and France.[1]

The rest of the paper is ordered as follows: Section 2 presents the ARCH model; in Section 3 practicalities related to the application of the ARCH model are discussed - they are of both general and data set specific natures; in Section 4 the "outlier" problem is addressed; in Section 5 the ARCH variance estimator is compared with a conventional estimator - the moving variance about moving mean (MVAMM) estimator; Section 6 presents a brief conclusion. A selective review of the literature which derives the theoretically expected behaviour of interest rates and asset

---

1 The data series investigated for each country include both a long-term and a short-term interest rate, and a broad share price index. The US$/£ exchange rate is also included in both papers.

prices based on the market efficiency hypothesis is presented in the Appendix.

2 The ARCH model

The "underlying motivation" for the ARCH model was to relax the empirically questionable assumption of constant variance imposed by conventional econometric estimation techniques on the disturbances of the equations they are employed to estimate. Engle [1982, pp 988] observes that, "The standard approach of heteroskedasticity is to introduce an exogenous variable x_t which predicts the variance ... This standard solution to the problem seems unsatisfactory, as it requires a specification of the causes of the changing variance, rather than recognising that both conditional means and variances may jointly evolve over time. Perhaps because of this difficulty, heteroskedasticity corrections are rarely considered in time series data."

Instead Engle proposes the ARCH regression model:

(1) $y_t | I_{t-1} \sim N(x_t \beta, h_t)$,

(2) $h_t = h(\varepsilon_{t-1}, \varepsilon_{t-2}, \ldots, \varepsilon_{t-p}, \alpha)$,

(3) $\varepsilon_t = y_t - x_t \beta$

where, for β a vector of unknown parameters, $x_t \beta$ is a linear combination of lagged endogenous and exogenous variables included in the information set I_{t-1} and used to describe the mean of the dependant variable. The residual of the conditional mean equation ε_t is a zero mean serially uncorrelated process, which Engle defines as an ARCH process. The variance function h_t could be generalised to include other contents of the information set, although in practice Engle employs the more restricted form:

(4) $h_t = \alpha_0 + \alpha_1 \varepsilon^2_{t-1} + \alpha_2 \varepsilon^2_{t-2} + \ldots + \alpha_p \varepsilon^2_{t-p}$.

Engle [1982] proposes a Lagrange multiplier procedure to test for a pth order ARCH process. In the test the current period's OLS residual ε_t - which measures the dispersion of the sample datum around the conditional mean - is squared and regressed on an intercept and $\varepsilon^2_{t-1}, \varepsilon^2_{t-2}, \ldots, \varepsilon^2_{t-p}$. The sample size times the R^2 of this regression is asymptotically distributed as chi-squared with p degrees of freedom if the null hypothesis ($\alpha_i = 0$, all i) is true.

Engle [1982] also presents an iterative maximum likelihood (ML) estimator to jointly estimate the two equations - the conditional mean and variance specifications - of the ARCH regression model. He shows that if lagged dependent variables are not included in the conditional mean equation and the process under consideration is stationary, then OLS is still the best linear unbiased estimator of the ARCH model equations, although the ML estimator is non-linear and therefore more efficient. However, "if there are lagged dependent variables in x_t, the standard errors as conventionally computed will not be consistent, since the squares of the disturbances will be correlated with squares of the x's". (Engle [1982, pp 994]). An indication of the efficiency gains of the ML estimator is given in Engle [1982, Sections 6 and 7], where it is shown that they can be "very large".[1]

Engle [1982, pp 989-990] lists three reasons for the attractiveness of the ARCH model:

1 Engle, Hendry and Trumble [1985] investigate the small sample properties of ARCH estimators and tests.

(1) "Econometric forecasters have found that their ability to predict the future varies from one period to another. McNees [17, pp 52][1] suggests that, 'the inherent uncertainty or randomness associated with different forecast periods seems to vary widely over time'. He also documents that, 'large and small errors tend to cluster together (in contiguous time periods)'. This analysis immediately suggests the usefulness of the ARCH model where the underlying forecast variance may change over time and is predicted by past forecast errors".

(2) "By the simplest assumptions, portfolios of financial assets are held as functions of the expected means and variances of the rates of return. Any shifts in asset demand must be associated with changes in expected means and variances of the rates of return. If the mean is assumed to follow a standard regression or time-series model, the variance is immediately constrained to be constant over time."

(3) "The ARCH specification might ... be picking up the effect of variables omitted from the estimated model. The existance of an ARCH effect would be interpreted as evidence of misspecification, either by omitted variables or through structural change. If this is the case, ARCH may be a better approximation to reality than making standard assumptions about the disturbances, but trying to find the omitted variable or determine the nature of the structural changes would be even better".

---

1 McNees, S.S, "The Forecasting Record for the 1970s", New England Economic Review, pp 33-53, September/October 1979.

3 The Application of the ARCH Model

3.1 Non-negativity and Stationarity Conditions

To be a sensible specification the coefficients of the ARCH variance function - equation (4) - need to satisfy some non-negativity and stationary conditions. If any of the coefficients are negative, then a single large residual could produce a negative conditional variance estimate; while if their sum exceeds unity, then the process is unstable and will eventually produce infinite variances (ie the unconditional variance is infinite).[1]

It is possible that unrestricted estimation of the variance function will satisfy these conditions. However, regarding the non-negativity condition Engle [1980, pp 8] points out that it seems unlikely that the "numerical optimization" used in the estimation of the unrestricted specification would return only positive coefficients. Engle [1982, pp 1002; 1980, pp 8] imposes a two parameter model to ensure that the non-negativity condition is satisfied:

$$(5)\quad h_t = \alpha_o + \alpha_1 \left[\sum_{i=0}^{p-1} (p - i)\varepsilon^2_{t-i-1} \Big/ \sum_{j=0}^{p-1} (p - j) \right]$$

The summation variable imposes linearly declining coefficients on the variance function. Such a specification is consistent with economic agents progressively discounting past information, although its strictness could impose an unnecessarily high loss of information relative to the unrestricted specification. In a

---

1 See Engle, Lilien and Robins [1984, pp 13].

discussion with Professor Engle he stated that he was now using more than one such summation variable, although each with a different lag length, to allow a more flexible lag structure. When two or more variables with different lag lengths are included, the overall lag structure will still have linear segments, however, it is free to take various forms and can even approximate Almon type characteristics. This method is adopted in Dickens [1986], while, except when used as a supplementary method for determining the order of the ARCH processes, the more restrictive weighting system was used in Dickens [1985].

3.2 Conditional Mean Specifications and Martingale Behaviour

The major advantage of the ARCH approach over conventional variance estimators (ie MVAMM type estimators) is that it measures dispersion around the conditional mean rather than about the sample mean. However, the preference for relatively high frequency weekly data in the author's study of asset market volatility has made the collection of data an expensive process and effectively ruled out, at least in the first stages of the exercise, the possibility of econometric modelling of the conditional means of each series. Instead autoregressive (AR) models are fitted to the series which, as it turns out, is not necessarily a second best approach. The literature which derives the theoretically expected behaviour of interest rates and asset prices based on the market efficiency hypothesis, suggests that approximate martingale behaviour[1] of such series, especially when high frequency data are employed, is reasonable. In this case the majority of the change in the series from this week to the next will be unanticipated on the basis of currently available information. The Appendix contains a partial review of this literature. It is by no means all encompassing, although it is considered sufficient in the present context.

3.3 A Supplementary F Test For Indentifying the Order of the ARCH Model

The test for a pth order ARCH process involves regressing the squared residuals from the conditional mean equation against a constant and the first p lags of themselves. The R^2 from this equation multiplied by the sample size is asymptotically distributed as chi-squared with p degrees of freedom under the null hypothesis of no significant ARCH process. ARCH statistics for orders 1 up to 52 were calculated for the weekly asset market series employed in Dickens [1985, 1986].

A feature of the results which is inherent in the nature of the ARCH test, is that, where several orders were found to produce significant ARCH statistics, the test does not indicate which order is the most "appropriate". Consider the example where, of the orders tested, only the first lag of the squared residuals is significant - significance

---

1 A sequence x_t follows a martingale process if:

$$x_t = x_{t-1} + \varepsilon_t, \quad \text{where } E(\varepsilon_t) = 0,\ E(\varepsilon_t, \varepsilon_s) = 0 \text{ all } t = s$$

In words, it is a process whose increments over fixed intervals are mean zero and serially uncorrelated (ie linearly independent). Random walk processes also satisfy these conditions although, generally, they also require a constant variance, and constancy of all higher movements if the distribution is non-normal. See Chatfield [1975, pp 39-41] for a discussion of random walk processes.

measured by the t values on the lagged squared residual variables in the test equation. The first order ARCH statistic would obviously reject the null hypothesis of no ARCH process. However, because the first lag of the squared residuals is also included in all higher order ARCH tests, significant ARCH statistics are likely to be returned for more than just the first order. In practice cases were frequently found where a handful of significant lower orders of the lagged squared residuals series produced a large enough R^2 that in excess of 52 higher order lags - on extra years worth of weekly data - did not need to add any explanatory power for the R^2's of the higher order equations to still return significant ARCH statistics.

In effect, in attempting to identify the order of an ARCH process, the ARCH statistic provides one with the equivalent information that the autocorrelation plot does when attempting to identify the order of an autoregressive (AR) process in a time series. A first order ARCH process can return spurious ARCH test results that suggest higher order processes exist, in the same way that the autocorrelation plot for a series which follows a first order AR process exhibits significant higher order autocorrelation coefficients which gradually die away as the order increases. The equivalent of the partial autocorrelation plot, which is used to identify the order of an AR process, is therefore required to identify the order of the ARCH process.

The individual coefficients on the lagged squared residuals in the ARCH test equations are very similar to the partial autocorrelation coefficients used in identifying AR processes. This suggests that having identified the orders for which significant ARCH statistics are returned, the actual order chosen is the one for which the highest ordered lagged squared residual series returns a significant coefficient as measured by its t value. A problem with this approach is that over the number of orders tested it is quite likely, even at the 1 per cent level of significance, that spuriously significant t values will arise. For example, each test will still have a 1 per cent probability of rejecting a correct null hypothesis, however, assuming independence of each test, if n tests are carried out then the probability of rejecting at least one correct null hypothesis increases to $1-0.99^n$ (ie the probability of accepting correct null hypothesis for all tests is 0.99^n). As 52 order are tested there is a 40.7 per cent probability of rejecting at least one correct null hypothesis.

Instead the usual F test for the relevance of additional regressors has been applied to the ARCH test results.[1] The F test is applied firstly by taking the lowest order which returns a significant ARCH statistic at the preferred level of significance - usually the 1 per cent level - as the initial base equation, and then testing against this subsequently higher order equations which are both significant in terms of their ARCH test statistics <u>and</u> for which the highest ordered squared residual series returns a significant t value. If a higher order equation returns a significant F statistic (ie the extra lags of the squared residuals in this equation add significant explanatory power over and above that given by the lags in the base equation) it becomes the new base equation. In this way a significant lagged squared residual must be sufficiently so as to outway the non-significance of the early lags, if any, between itself and the

1 See Kmenta [1971, pp 370-371] for a description of the test. David Barr has pointed out that because the test is applied to an equation including lags on the dependent variable as regressors the actual statistic calculated has an unknown distribution. However, in large samples it will approximate the F distribution. The calculated statistic multiplied by the number of constraints to be tested will also be asymptotically distributed as chi-squared. As the sample sizes used in Dickens [1985, 1986] were generally in the region of 900 per series, there does not seem to be a problem with the test used.

previous highest ordered lag which returned a significant t value. This procedure is repeated until the highest relevant order is determined.

4 The "Outlier" Problem

In Dickens [1985] extreme values were identified in the increments of three of the four UK asset market series investigated. These "outliers" were particularly noticeable because for each series they were of the same sign and caused the sample distributions of the differenced series to be significantly skewed. Even without these observations the sample distributions were markedly leptokurtic (ie long tailed and more peaked) relative to the normal distribution.[1] However, Fama [1977, pp 26] presents evidence of leptokurtosis in daily stock market returns and reports the work of Mandelbrot which indicates that it is reasonable to expect series of the sort being investigated to display leptokurtosis.[2] It therefore seems that it is more the skewness of these observations than their extreme distance from the sample mean (ranging from 5 to 8 standard deviations away the respective sample means) that identify them as "outliers".

Extreme values are also identified in the international asset market series used in Dickens [1986]. Again the kurtosis statistics reported for the residuals of the autoregressive conditional mean models fitted to the differenced international

---

1 See Yule and Kendall [1965, pp 164] for the definition of leptokurtosis and a general discussion of kurtosis.

2 Mandelbrot, B, "The Variation of Certain Speculative Prices", Journal of Business, pp 395-419, October 1963. See also Taylor [1985, pp 719] for evidence of leptokurtosis in the distributions of daily returns from assets futures markets.

series all rejected normality, while most sample distributions are significantly skewed.

Potential explanations for the outliers in the UK data are canvassed in Dickens [1985, pp 12-18]. It is concluded there that the most favourable explanation is that the data for each series are the product of more than one parent distribution. In the simple two distribution case, the hypothesis is that there is an "underlying" population which generates most of the observations, and a secondary "outlier" population which generates the outlying observations.[1] The underlying distribution is seen as being symmetric, although probably leptokurtic, and arising from "normal" exogenous shocks and economic behaviour. It is hard to prejudge the characteristics of the population distribution for the hypothetical outlier process because only a handful of observations were identified as outliers in the sampled data, although its elements are seen as the result of large exogenous shocks or the tendency for endogenous variables to behave discontinuously.

Casual empiricism suggests that the outliers observed in the UK data are closely linked with policy interventions, in terms of both policy adjustments and regime changes. Examples of relationships between outliers and policy actions are the floating of sterling in 1972 and the first budget of the Conservative Government in June 1979, which were both associated with over two percentage point jumps in UK three month interbank interest rates,

---

1 The secondary population might be responsible for some non-outlier observations within the body of the sample distribution, similarly the outlying observations may be the joint product of both the underlying and outlier populations.

as were the re-introduction of the minimum lending rate (MLR) in January 1985 and a one and three quarter percentage point increase in the MLR in November 1973.

If the two (or more) process hypothesis is accepted, it does not provide an unambiguous answer for handling the outliers, although it would appear to enable bounds to be put on the variance point forecasts. If the probability of observing a single outlier in any one period is both low and constant,[1] then the variance estimates obtained when the outliers are excluded will closely approximate the underlying variance levels.[2] Alternatively, if the probability of observing an outlier is time dependent - in particular, if they tend to cluster together - then their influence should be included in the variance estimates. The exclusion and inclusion of the outliers would therefore provide the lower and upper bounds, respectively, for the point forecasts of the variance.

The proper inclusion of the outliers under the two distribution hypothesis is a complicated business. It would involve the separate modelling of the variance of the outlier process - assuming the latter was not constant - and the relevant

---

1 This implies that the outliers come from the interaction of two distributions. The first related to the possibility of observing an outlier in any one period, and the second determining the value of the outlier given one will occur. Also, for this example, the variance of the process can be either constant or time dependent, as long as the probability of observing an outlier is sufficiently low its impact will be small.

2 This also assumes that the ARCH variance estimates are intended to represent the markets' variance forecasts, or, at least, the best forecasts the markets could have, and that the market can distinguish outlying observations.

covariances. However, as already mentioned, insufficient observations have been identified as coming from the hypothetical outlier process to infer anything about either the probability of an outlier occurring, or the distribution from which they are drawn. The usual practice is for the outliers to be included with the rest of the data and the ARCH model fitted to the full data set. This assumes that the skewness caused by the outliers will not seriously upset the fitting of the conditional mean models and the ARCH test, both of which are based on normality assumptions.

A slightly more sophisticated method of handling the outliers has been experimented with in Dickens [1986]. It provides a means of allowing the data to reweight the extreme observations.[1] In its simplest form it involves adding one dummy variable to the ARCH test equation for each lagged dependent variable in the equation. The dummy variables would be zero in periods where "normal" observations occurred in the corresponding lagged dependent variables, and contain the actual data observations where it was a specified number of standard deviations from the sample mean (eg 3 sigma). If a dummy returned a coefficient of equal but opposite sign to the coefficient on its respective lagged dependent variable, it would be equivalent to excluding the "outliers"; while a zero coefficient would give them full weight in the variance calculations. It would be hoped that the coefficients would lie somewhere between the extremes.

The obvious generalisation of this would be to include separate sets of dummy variables for small outliers (eg those from 2.5 to 4

1 This was suggested by Professor Engle.

sigma from the sample mean), moderate outliers (eg those in the 4 to 6 sigma interval), and large outliers (eg those outside the 6 sigma interval). Given that the outliers within each sigma interval could be of opposite sign, it would be their absolute values that would be included in the dummy variables.

5 Comparison of ARCH and MVAMM Variance Estimators

The difference between the specifications of the ARCH and MVAMM variance estimators is reduced considerably when they are applied to the first difference of the asset market series. Barr [1984, pp 4] applies a conventional variance estimator to the first differenced series in his investigation of exchange rate variability. The argument being that, on the basis of the "fairly common finding that the exchange rate follows a random walk", all of the change in the series is unanticipated.

The reasons for the similarity are probably clearer if the estimators themselves are considered. The MVAMM estimator based on a moving window of historic data of length k for both the mean and variance is given by (6), where V_t is the one step ahead variance forecast:

$$(6) \quad V_t = \sum_{i=1}^{k} [(DR_{t-i} - (\sum_{j=i}^{k} DR_{t-j}/k))^2/(k-1)]$$

where $DR_t = R_t - R_{t-1}$, for R_t the level of the asset market series.

This k period moving variance estimator is quite similar to the ARCH estimator if DR_t is approximately mean zero for the value of k chosen <u>and</u> if the conditional mean equation fitted to DR_t [1] in the ARCH approach explains only a little of the week-to-week changes in the series (ie R_t approximates a martingale process). If the first of these holds then (6) reduces to (7); while if the second holds, and so the residuals of the conditional mean model for the differenced series will be approximately equal to DR_t, then (8) is an approximate statement of the ARCH variance estimator.

$$(7)\ V_t \cong \sum_{i=1}^{k} mDR_{t-i}^2 \quad \text{, where } m = 1/(k-1)$$

$$(8)\ h_t \cong \alpha_o + \sum_{i=1}^{p} \alpha_i\ DR_{t-i}^2$$

The similarity between the estimators is strong, although three potentially important differences remain: the ARCH variance function contains a constant; the weights in (8) depend on i unlike the constant weight in (7); p is chosen on the basis of the ARCH test, while there is no obvious method for choosing k.

The impact of the first two of these differences between the estimators depends on the estimated values of the constant and the weights in the ARCH specification. All the ARCH variance functions estimated for the international data series had positive constant terms, and the sums of their weights (the α_i in (8))

1 In Dickens [1985, 1986] the mean models were fitted to the first differences in the asset market series because their levels were found to be non-stationary.

were all in the range 0.15 to 0.90 and generally close to the average of 0.52 for all series,[1] and therefore only the implications of this combination of results are discussed here.[2]

The minimum ARCH variance estimate for any period, given the imposition of non-negative coefficients, is the value of the constant in the equation, while the minimum bound of all MVAMM estimates is zero. In relatively tranquil periods where only small increments are observed, such as the 1960s, the ARCH variance estimates will therefore generally exceed the MVAMM estimates by the extent to which the constant exceeds zero. The second difference, which arises because the sum of the weights in the ARCH variance functions are less than unity, means that the ARCH estimates will be less volatile than the MVAMM estimates.

The difference is also particularly noticeable when outliers are encountered. When a large increment enters the moving window of data for both estimators, their variance estimates increase considerable, sometimes more than doubling with one large observation. The initial increases are often of quite similar magnitude as the earlier weights in the estimated ARCH equations

---

1 The numbers quoted apply to the unrestricted ARCH variance functions reported in Dickens [1986]. Summing the weights on the restricted equations is more difficult and so has not been done, although it seems reasonable to expect that they will not differ markedly from the results for the unrestricted equations.

2 Engle, Lilien and Robins [1984, pp 12-14] encounter non-stationary ARCH variance functions (ie the sum of the weights exceed unity) in their investigation of time varying risk premia in the term structure.

were often close in value to the constant weight in the MVAMM estimator.[1] The difference appears once the large increment starts passing through the respective windows of the estimators. The ARCH estimates will decay reasonably rapidly because of the linearly declining lag structure and return smoothly to the earlier level,[2] while the MVAMM estimates will stay at the higher level until the observation drops out of the end of the window, at which point the estimates will step down by a similar magnitude as the initial increase.

This has assumed only one large increment. If a second one appeared while the first was still in the windows, then both estimates would step up by similar amounts. However, not only would the ARCH estimate be starting from a lower base before the second large increment appeared - how much lower depending on how far the first observation had passed through the window - but the different weighting systems would ensure that by the time the first increment reached the end of the window the difference between the two estimators could be very large. Such a case occurred for the UK share price index series investigated in Dickens [1985], where the MVAMM estimates were over twice the

---

1 If both estimators have 13 week lag lengths - which is exactly the case for the MVAMM estimators and approximately the case for the ARCH estimators reported in Dickens [1985] - and given that the sum of weights on the estimated ARCH equations is often in the region of 0.5, then the first ARCH weight is 0.071 while all MVAMM weights are 0.077.

2 Even in the unrestricted ARCH equations the estimates would still generally show a tendency to decay as the large observated passed through the lag structure, although it would not be smooth and they would pop up (down) on occasions when significant positive (negative) coefficients were encountered in the lag structure.

size of the ARCH estimates for the last several weeks of 1975 before the first increment drop out the end of the windows. Only as the number of large increments increases will the impact of the different weighting systems diminish.

Prior to obtaining knowledge of the ARCH test results, the choice of window length for the MVAMM estimator was an arbitrary one, with potential arguments available for both long and short lag lengths. Large differences in order could give the variance series produced by the two estimators quite different profiles, especially in the face of large observations.

6 Conclusion

In conclusion, to the extent that the cost of applying the ARCH model is not prohibitive, the extra rigour it provides for describing the time varying nature of the variance of a series is a most welcome advance in an area which has been noticeable for its lack of technological innovations.

Appendix: Theoretically Expected Behaviour of Interest Rates and Asset Prices - A Selective Review

There is a belief in some quarters that market efficiency should impose martingale behaviour on interest rates and asset prices. Possibly underlying this belief is the Samuelson/Mandelbrot[1] result which proposes that market efficiency implies the following sequence is a martingale:[2]

(2) $({}_{t+j}F_{t}, {}_{t+j}F_{t+1}, \ldots, {}_{t+j}F_{t+j-1}, R_{t+j})$

where R_{t+j} is the one-period spot rate applying in period t+j; ${}_{t+j}F_{t+i}$ is the forward rate on one-period loans that prevails at time t+i for loans to be made at time t+j (i < j).

The Samuelson/Mandelbrot result applies to forward rates. However, intuitive arguments have been presented asserting that market efficiency requires spot rates to follow a random walk. Pesando [1979, pp 45] points out that such reasoning has been advanced by Poole [1976, pp 476] and Phillips and Pippenger [1976, pp 11]. If their arguments were substantive, then the significant serial correlation reported for many of the differenced asset market series employed in Dickens [1985, 1986] would be evidence against the market efficiency hypothesis. Possibly the most poignant comment in this respect was made by

---

1 Samuelson, P A, "Proof that Properly Anticipated Prices Fluctuate Randomly", Industrial Management Review, pp 41-49, Spring 1965.

Mandelbrot, B, "Forecasts of Future Prices, Unbiased Markets and Martingale Models", Journal of Business, 39 (Special Supplement), January 1966.

2 See Sargent [1972, pp 74] and Pesando [1979, pp 458] for statements of the market efficiency hypothesis in this context, while Sargent also discussed the Samuelson/Mandelbrot result.

Taylor [1985, pp 730], who observed that, "Any attempt to disprove the efficient market hypothesis is difficult because academics cannot yet describe the evidence which would suffice to reject it conclusively."

On the basis of the literature reviewed which theoretically derives the autoregressive implications of market efficiency for spot rates, it would seem that empirical evidence for or against spot rates following a martingale process does not shed unambiguous light on the market efficiency hypothesis debate. It would appear that the finding of correlation in the increments of such series could indicate that one or any combination of the following apply, inter alia: (1) risk premia exist and are autocorrelated; (2) equilibrium returns are autocorrelated; (3) the cost of acquiring such information exceeds any benefits its limited explanatory power provides; (4) frictions and non-competitive elements exist in the relevant market. The papers considered are Sargent [1972], Leroy [1982], Pesando [1979], Sims [1984] and Begg [1983]. Their results are certainly not all encompassing, although are considered sufficient for the present purpose.

Sargent [1972, pp 85] concludes for one-period spot rates, even when they are assumed to follow a discrete (covariance) stationary stochastic process with finite variance,[1] "that our two hypotheses [joint hypotheses making up the market efficiency hypothesis[2]], even

---

1 This is a more restictive assumption than required to derive the Samuelson/Mandelbrot result, for which spot rates can come from a distribution with infinite variance.

2 Joint hypotheses are involved in tests of market efficiency because "a particular model of market equilibrium is examined simultaneously with the question of market efficiency". (Pesando [1979, pp 457]).

in the special form of this section, do not imply things that are commonly thought to be their implications.

"Thus, they do not imply that the spot one-period rate, R_t, follows a random walk, which would mean that $R_t - R_{t-1}$ is serially uncorrelated. They do not imply that the j-period spot rate follows a random walk for any finite j. Moreover, the 'fair game' property built into the model clearly does not mean that spot rates cannot be described by a stable stochastic difference equation."

The intuitive argument Sargent gives for his result is that, if market efficiency only requires that forward rates follow a martingale, then, as long as market efficiency as defined ensures no opportunities arise for excessively profitable arbitrage, it seems reasonable that this requirement will be met, but no more, other than by pure coincidence.

Leroy [1982, pp 186] is in agreement with Sargent when he states that, "Martingales play a central role in expectations models of asset prices, ... However, as will be seen below, the variable which follows a martingale is sometimes not the economic variable of primary interest, but rather a related variable. For example, under the expected present-value model stock prices plus dividends follows a martingale, but stock prices itself does not. Also, under the expectations hypothesis of the term structure of interest rates, no interest rate series itself follows a martingale; rather, it is implied future rates that do so."

Subject to a linear approximation Pesando [1979] derived (3) for interest rates on non-coupon bonds and (4) for coupon rates:

$$(3)\quad E(R_{n,t} \mid I_{t-1}) - R_{n,t-1} = \frac{1}{n}\,[E({}_{t+n-1}f_{1,t} \mid I_{t-1}) - R_{1,t-1}]$$

$$(4)\quad E(R_{n,t} \mid I_{t-1}) - R_{n,t-1} = \frac{(1-\lambda)}{(1-\lambda^{n})}\,[(1-\lambda)\,{}_{t}f_{1,t-1} + \lambda(1-\lambda)\,{}_{t+1}f_{1,t-1} +$$

$$\ldots + \lambda^{n-2}(1-\lambda)\,{}_{t+n-2}f_{1,t-1} + \lambda^{n-1}E({}_{t+n-1}f_{1,t} \mid I_{t-1}) - R_{1,t-1}]$$

where $R_{n,t}$ is the spot rate on an n-period non-coupon bond in (3) (coupon bond in (4)) in period t; I_{t-1} is the information set available at t-1 (the current period); ${}_{t+i}f_{1,t}$ is the market's expectation at time t of the one-period spot rate in period t+i, $i \geq 0$; $\lambda = 1/(1+R)$ where R is the "representative" or "normal" one-period rate.

The implications Pesando [1979, pp 460] draws from (3) for n-period non-coupon spot rates are:

(1) The RHS of (3) approaches zero as n gets large, and so yields on long-term bonds will "follow (approximately) a martingale sequence".

(2) This approximate martingale behaviour can be expected if "the bond market is efficient and if term premiums, should they exist, are time-invariant. If time-varying term premiums do exist, then the change in the long-term bond rate can vary

predictably with the change in the corresponding term premium without contradicting the efficient market hypothesis".[1]

(3) Such behaviour for long-rates is then "grounded ultimately in the well-documented role of expectations as a key determinant of the term structure, together with empirical studies [3, 7][2] that suggest that term premiums may well be time invariant".[3]

(4) The martingale approximation for long rates, "is less restrictive than the random walk model, since it does not require that successive changes in long-term interest rates be independently and identically distributed over time".

---

1 Time varying risk premia also upset the Samuelson/Mandelbrot result. If such premia exist but are time invariant the result no longer holds exactly although remains largely intact: instead of following a martingale process forward rates follow a semimartingale process which still ensures uncorrelated increments. See Sargent [1972, pp 78] and Pesando [1979, pp 460] for discussions of this point, and Sims [1984, pp 267] and Doob [1953, pp 294] for a discussion of semimartingale processes.

2 McCulloch, J H "An Estimate of Liquidity Premium", Journal of Political Economy, 83, pp 95-120, February 1975.

Pesando, J E "Alternative Models of the Determination of National Interest Rates", Journal of Money, Credit and Banking, 8, pp 209-18, May 1976.

3 Time invariant term premiums do not upset the result because, while they appear in the expectations equation that underlies the market efficiency hypothesis and which states n-period spot rates as a discounted function of future expected one-period spot rates, they drop out once this equation is differenced in the manipulations to obtain (3) and (4).

(5) "The proposition that short-term rates follow a random walk can be obtained <u>only by direct assumption</u> ... If and only if the equilibrium return on Treasury bills [the example of a non-coupon short rate chosen] follow a random walk will the one-period rate in an efficient market follow a random walk. There is, however, no requirement of efficient market theory that the equilibrium return - hence the 90-day bill rate - behave in such a fashion."[1]

(6) "This result requires only that the short-term rate not be 'too' nonstationary, so that its expected value n periods in the future is not dramatically different from its latest value."

Of the term on the RHS of (4), Pesando [1979, pp 461-463] observes:

(1) "For <u>coupon</u> bonds, whose yields are generally employed in empirical research, the result is less clear ... Although the size of the term is not immediately obvious (more on this below), it does suggest that the <u>anticipated</u> component of the change in the long-term bond rate could be significant".

(2) "Rough calculations with Canadian data indicate that, at least for the period 1961 I to 1976 IV, the proportion of the

---

1 With particular reference to the Phillips and Pippenger [1976] random walk assertion, Pesando [1979, pp 457] warned against such intuitive reasoning, pointing out that earlier studies of stock price behaviour that similarly equated random walk behaviour of stock prices with market efficency, "typically failed to make explicit their assumptions that equilibrium returns are constant over time. Only under this assumption does evidence on the autocorrelation of successive one-period returns bear directly on the question of market efficiency".

change in long-term Canadian yields that has been anticipated is of relatively minor importance ... only 1.75 per cent of the variance of the change in the long-term rate could be assigned to the anticipated component. The vast majority - 98.25 per cent - of the variance of the long-term rate thus represents, under the joint hypothesis being investigated, the receipt of new information.

In a highly technical paper Sims [1984] also states that the intuitive argument for martingale behaviour of asset prices and interest rates is overstated. In this context he observes that, "careful examination of competitive general equilibrium models of behaviour under uncertainty shows, as emphasised by R E Lucas, Jr [1978] and by Stephen F Leroy [1973],[1] among others, that (1) [Sims' (1) is the martingale specification] emerges from such models only under extremely restrictive assumptions". (Sims [1984, pp 1]).

Based on "very general assumptions"[2] he derives that, "durable good prices and interest rates will in a frictionless competitive market, show such approximate martingale behaviour. More

---

1 Lucas, R E, Jr "Asset Prices in an Exchange Economy", Econometrica, 46, pp 1429-1449, 1978.

Leroy, S F, "Risk Aversion and the Martingale Property of Stock Prices", International Economic Review, 14, pp 436-446, 1973.

2 "In effect, we assume that a competitive market equilibrium exists and has 'realistic' charactistics ... The harder problem of deriving existence of equilibrium with realistic price behaviour from assumptions about individual behaviour is sidestepped". (Sims [1984, pp 3]).

precisely, what is shown is that the linear regression of $P_{t+s}-P_t$ on X_t [X_t the information set that becomes available at t, the current time period] and lagged X_t's, predicted by (1) to yield an R^2 of zero, instead has an R^2 converging to zero as s goes to zero. Instead of price changes being unpredictable, price changes over small intervals are very nearly unpredictable.

"Thus if one wishes to interpret an econometric test of (1) as a test of the importance of frictions and noncompetitive elements in a market, one ought to carry out the test with s 'small'. Also, the fact that a price is set in a competitive market with few frictions and seems to fit (1) reasonably well for, say, an s of one week, does not mean that (1) should be expected to work well also for an s of one year". (Sims [1984, pp 2-3]).

Sims terms the approximate martingale behaviour "instantaneous unpredictability". A process is defined as being instantaneously unpredictable, if and only if:

$$(5) \quad E_t[(P_{t+v}- E_t(P_{t+v}))]^2/E_t[(P_{t+v}- P_t)^2]$$
$$\rightarrow 1 \text{ almost surely as } v \rightarrow 0$$

"In words, for an instantaneously unpredictable process prediction error is the dominant component of changes over small intervals. Of course, for a martingale with finite second moments, the ratio in (2) [our (5)] is exactly 1". (Sims [1984, pp 5])

As a slight qualification, Sims does note that the conclusion of instantaneous unpredictablity for increments of high frequency data can be upset if information does not arrive continuously, causing discontinuous jumps, and if "the size or timing of the jumps flows 'non-smoothly' in a certain sense". (Sims [1984, pp 23]). This would seem to have some relevance to the discussion of "outlying" observations contained in Section 4 of this paper.

On a different tack, Begg [1983] argues that the "fundamental theorem" underlying the usual statement of the market efficiency hypothesis, and employed by Mishkin[1] to derive that the n-period yield should follow a random walk, is wrong even under certainty-equivalence. The problem he sees with the fundamental theorem, which "asserts that the redemption yield on an n-period bond is an average of the expected future one-period interest rates over the remaining life of the bond" (Begg [1983, pp 45]), is that it "is based on a quite misleading approximation, namely that the interest rate at which coupon payments are instantaneously reinvested is the same as the average return over n-periods when the bond was initially purchased". (Begg [1983, pp 53]). This assumption would also seem to underlie the work of Sargent [1972], Pesando [1979] and Sims [1984].

Begg suggests that the statement of asset market equilibrium as the equality of single period holding yields overcomes this problem. The implication of this drawn out by Begg [1983, pp 53] for the random walk hypothesis is that, "If it is anticipated that the short rate will change in the future, bond yields and bond

---

1 Mishkin, F S, "Efficient Markets Theory: Implications for Monetary Policy", Brookings Papers on Economic Activity, pp 707-752, 1978.

Mishkin, F.S, "Is the Preferred Habit Model of the Term Structure Consistent with Financial Market Efficiency?", Journal of Political Economy, Vol.88, pp 406-411, 1980.

prices will not follow a random walk. Since short rates are often close to a random walk this assertion is compatible with empirical findings that redemption yields are not far from a random walk but the analysis of this section identifies circumstances in which the random walk model of bond prices should be expected to fail, namely when the market anticipated changes in the short rate".

REFERENCES

Barr, D G, [1984], "Exchange Rate Variability: Evidence for the Period 1973-1982," Bank of England Discussion Paper No 11, November.

Begg, D K H, [1983], "Rational Expectations and Bond pricing: Modelling the Term Structure with and without Certainty Equivalents," The Economic Journal (Conference Papers), pp 45-58, May.

Chatfield, C, [1975], The Analysis of Time Series: Theory and Practice, London, Chapman and Hall.

Dickens, R R, [1985], "Variability in Some Major UK Asset Markets since the mid-1960s: An Application of the ARCH Model", paper presented at the Money Study Group Conference, Brasenose College, Oxford, 11-13 September.

Dickens, R R, [1986], "International Comparison of Asset Market Volatility: A Further Application of the ARCH Model", forthcoming.

Doob, J L, [1953], Stockastic Processes, New York, Wiley.

Engle, R F, [1980], "Estimates of the Variance of US Inflation Based Upon the ARCH Model", University of California, San Diego, Discussion Paper 80-14, July.

Engle, R F, [1982], "Autoregressive Conditional Heteroscedasticity with Estimates of the Variances of United Kingdom Inflation", Econometrica, Vol. 50, No. 4, pp 987-1007, July.

Engle, R F, Lilien, D M, and R P Robins, [1984], "Estimating Time Varying Risk Premia in the Term Structure: The ARCH-M Model", University of California, San Diego, Discussion Paper 85-17, December.

Engle, R F, Hendry, D F, and Trumble, D, [1985], "Small-sample Properties of ARCH Estimators and Tests", Canadian Journal of Economics, Vol 18, No 1, pp 66-93, February.

Fama, E F, [1977], Foundations of Finance, Oxford, Blackwell.

Kmenta, J, [1971], Elements of Econometrics, New York, MacMillan.

Leroy, S F, [1982], "Expectations Models of Asset Prices: A Survey of Theory, "The Journal of Finance, Vol 37, No 1, pp 185-217, March.

Pesando, J E, [1979], "On the Random Walk characteristics of Short- and Long-Term Interest Rates In an Efficient Market", Journal of Money, credit and Banking, Vol 11, No 4, pp 456-466, November.

Phillips, L and J Pippenger, [1976], "Preferred Habital Vs. Efficient Market: A Test of Alternative Hypotheses", Federal Reserve Bank of St Louis, Review, 58, pp 11-19, May.

Poole, W, [1976], "Rational Expectations in the Macro Model", Brookings Papers on Economic Activity, 2, pp 463-505.

Sargent, T J, [1972], "Rational Expectations and the Term Structure of Interest Rates", Journal of Money, Credit and Banking, Vol 4, No 1, pp 74-97, February.

Sims, C A, [1984] "Martingale-Like Behaviour of Prices and Interest Rates", University of Minnesota, Centre for Economic Research Department of Economics, Discussion Paper No 205, October.

Taylor, S J, [1985], "The Behaviour of Futures Prices Over Time," Applied Economics, 17, pp 713-734.

Yule, G U, and M G Kendall, [1965], An Introduction to the Theory of Statistics (fourteenth edition), London, Griffin.

Bank of England Discussion Papers

	Title	Author
1-5, 8, 11-14, 16-17, 20-22	*These papers are now out of print, but photocopies can be obtained from University Microfilms International (see below).*	
6	'Real' national saving and its sectoral composition	C T Taylor A R Threadgold
7	The direction of causality between the exchange rate, prices and money	C A Enoch
9	The sterling/dollar rate in the floating rate period: the role of money, prices and intervention	I D Saville
10	Bank lending and the money supply	B J Moore A R Threadgold
15	Influences on the profitability of twenty-two industrial sectors	N P Williams
18	Two studies of commodity price behaviour: Interrelationships between commodity prices Short-run pricing behaviour in commodity markets	 Mrs J L Hedges C A Enoch
19	Unobserved components, signal extraction and relationships between macroeconomic time series	T C Mills
23	A model of the building society sector	J B Wilcox
24	The importance of interest rates in five macroeconomic models	W W Easton
25	The effects of stamp duty on equity transactions and prices in the UK Stock Exchange	Mrs P D Jackson A T O'Donnell

Papers presented to the Panel of Academic Consultants[(a)]

	Title	Author
8	International monetary arrangements the limits to planning*	P M Oppenheimer
9	Institutions in the financial markets: questions, and some tentative answers*	M V Posner
10	The arguments for and against protectionism*	M Fg Scott The Hon W A H Godley
14	The usefulness of macroeconomic models*	Prof W H Buiter T F Cripps Prof Angus Deaton Prof A P L Minford M V Posner
15	Factors underlying the recent recession*	G D N Worswick Dr A Budd
17	Why do forecasts differ?*	Prof M J Artis
19	Bank lending, monetary control and funding policy*	Prof A D Bain
20	The economics of pension arrangements*	Prof Harold Rose J A Kay
22	Monetary trends in the United Kingdom	Prof A J Brown Prof D F Hendry and N R Ericsson
23	The UK economic recovery in the 1930s	G D N Worswick P N Sedgwick Prof Michael Beenstock Dr Forrest Capie Prof Brian Griffiths
24	Employment, real wages and unemployment in the United Kingdom*	Prof J R Sargent Sir Bryan Hopkin

Technical Series

	Title	Author
1	The consumption function in macroeconomic models: a comparative study*	E P Davis
2	Growth coefficients in error correction and autoregressive distributed lag models	K D Patterson
3	Composite monetary indicators for the United Kingdom; construction and empirical analysis*	T C Mills
4	The impact of exchange rate variability on international trade flows*	G Justice
5	Trade in manufactures*	A C Hotson K L Gardiner
6	A recursive model of personal sector expenditure and accumulation*	E P Davis
7	A dynamic 'translog' model of substitution technologies in UK manufacturing industry	D J Asteraki
8	Financial asset portfolio allocation by industrial and commercial companies*	Mrs P D Jackson
9	Some properties of the Bank model*	G P Dunn N H Jenkinson I M Michael G Midgley
10	Asset demands and interest rate setting equations in imperfect markets*	C J Green
11	Exchange rate variability: evidence for the period 1973 — 1982*	D G Barr
12	The development of expectations generating schemes which are asymptotically rational	K D Patterson
13	The arch model as applied to the study of international asset market volatility	R R Dickens

* These papers are no longer available from the Bank, but photocopies can be obtained from University Microfilms International, at White Swan House, Godstone, Surrey, RH9 8LW.

(a) Other papers in this series were not distributed.