Bank of England

Discussion Papers

Technical Series

No 41

News effects in a high frequency model of the sterling-dollar exchange rate

by

C A E Goodhart (a)

S G Hall (b) (c)

S G B Henry (b)

and

Bahram Pesaran (b)

June1991

No 41

News effects in a high frequency model of the sterling-dollar exchange rate

by

C A E Goodhart (a)

S G Hall (b) (c)

S G B Henry (b)

and

Bahram Pesaran (b)

June 1991

(a) London School of Economics (LSE).
(b) Bank of England.
(c) London Business School (LBS).

The object of this Technical Series of Discussion Papers is to give wider circulation to econometric research work in the Bank's Economics Division, and to invite comment upon it; any comments should be sent to the authors at the address given below.

The views expressed in this paper are those of the authors and do not necessarily represent those of the Bank of England, the LSE or the LBS. The authors would like to thank the Bulletin Group for editorial assistance in producing this paper.

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. *(Telephone: 071-601-4030)*

ISBN 1 85730 005 X
ISSN 0263-6123

Summary

This paper uses an extremely high frequency data set on the dollar-sterling exchange rate to investigate the impact of news events on the very short-term movements in exchange rates. The data set is a continuous record of the quoted price for the exchange rate on the Reuters screen. As such it records some 130,000 observations over an eight week period. The paper investigates the time-series properties of the data using orthodox regression models, and then by making allowance for a time varying conditional variance. The conclusions vary significantly in moving to this more sophisticated model. The exercises are repeated now incorporating news announcement effects, letting these affect the level of the exchange rate and then the conditional variance process. Again it is found that the conclusions are radically altered in moving to the increasingly sophisticated model.

1 Introduction

A number of recent papers have explored the behaviour of the foreign exchange markets using extremely high frequency data, Goodhart (1990), Goodhart and Figliuoli (1991) being two examples. This paper extends their work by investigating the response of the foreign exchange market to news events within this very high frequency setting. The data consists of the continuous time series of prices quoted on the Reuters' FXFX page for the sterling-dollar rate over the period Sunday 9 April to Monday 3 July 1989. The Reuters' service provides data on a number of bilateral rates, and the data sets vary with the level of activity for each currency, but in each case there is a very large number of observations (for the sterling-dollar rate there are some 135,978 observations, while the deutschemark-dollar rate has approximately 250,000 observations).

The very size of the data sets have tended to limit the nature of the exercises which have been based on them so far, these being limited to descriptive analysis and OLS regression analysis. This paper seeks to remedy this and it provides an extension to the analysis previously used by including the behaviour of the second moment of the data set, using the GARCH-M framework of Engle, Lilien and Robins (1987). Maximum likelihood estimation of a GARCH-M model is used to assess the presence of ARCH effects in the data, and to test whether the conditional variance has a significant effect on the behaviour of the exchange rate at very high frequencies. A further extension investigates the effects of news events within this GARCH-M framework, where news may potentially affect either the level or the variance of the exchange rate.

The plan of the paper is as follows. Section 2 describes the data set used, and section 3 outlines the GARCH-M model and our estimation strategy. Section 4 then describes the empirical results, and section 5 offers conclusions.

2 The data set

As noted earlier the data is a continuous record of the prices quoted on the REUTERS FXFX screen over the period 9 April to 3 July 1989. These prices are invitations to trade rather than records of actual trades, and no public record is available of actual traded prices. In fact this feature is true of almost all exchange rate data, although in the very high frequency context this does give rise to special problems, most notably input errors,and our treatment of these is described further below. In addition to the direct measures of the exchange rate the Reuters' screen provides a useful record of the timing and availability of news to the markets in its FXNB page which contains a record of news announcements and timings.

Figure 1
An example of an outlier, 17 May 1989 around 01:43

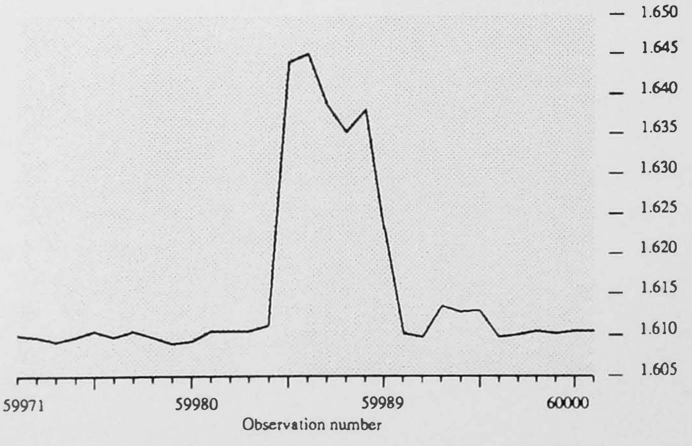

Although the data was prefiltered for errors it was necessary to remove one further class of input errors which had a particularly interesting shape and pattern. Figure 1 shows an example of this type of error, this one occurring on 17 May at around 1:43.

The exchange rate which had been trading at around 1.6090 suddenly jumped to 1.6450 and remained at around that level for around 10 consecutive bids before falling back to its original level. The explanation for this seems to be that the traders are generally concerned only with the last two digits of a bid, often copying the first three from the previous bid. Once an input error occurs in the first few digits this can easily be copied for a period of time until the error is noticed and corrected. Of course the existence of these errors does not mean that trades actually take place at the quoted price, and so the observations do not represent a real change in price but merely an input error. We have therefore filtered the data to remove such anomalies and, perhaps not surprisingly, the overall properties of the data in terms of skewness and Kurtosis, were quite radically altered once this had been done.

In allowing for news effects, the focus in this paper is on two particular news events. These are the announcement on Wednesday 17 May of the US Trade figures, which were considerably better than expected, and the rise of 1% in UK base interest rates on Wednesday 24 May. These two events are graphed in Figures 2 and 3 and, as is clear, both events caused a fairly obvious and long-lasting change in the level of the exchange rate (in terms of the number of observations affected).

Figure 2
The effect of the US trade figures, 17 May 1989 around 13:30

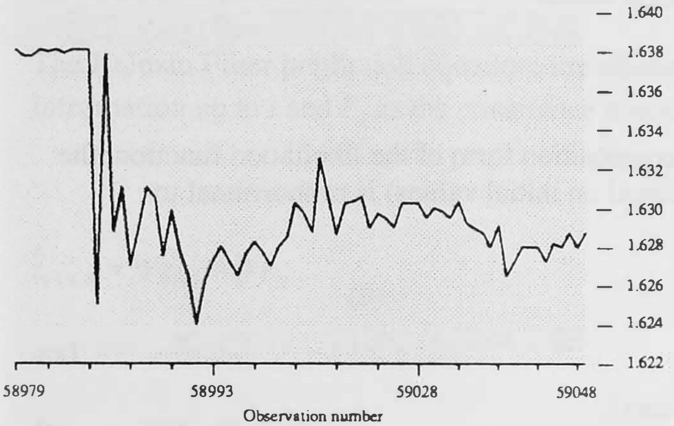

Figure 3
The effect of a rise in UK interest rates, 24 May 1989 around 09:47

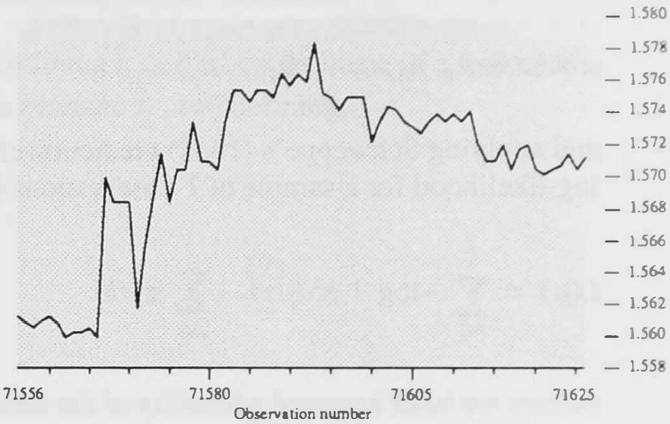

3 Estimation techniques

This section first gives a brief account of the GARCH-M model and then outlines the strategy adopted for efficiently estimating the model using a very large data set.

A The GARCH-M Model

Engle, Lilien and Robins (1987) suggest an extension of Engle's (1982) ARCH model whereby the conditional first moment of a time series itself becomes a function of the conditional second moment, which follows an ARCH process:

$$y_t = \alpha' x_t + \delta h_t^2 + \varepsilon_t \quad \varepsilon_t \mid \Omega_{t-1} \sim N(0, h_t^2) \tag{1}$$

$$h_t^2 = \gamma_0 + \sum_{i=1}^{n} \gamma_i \ \varepsilon_{t-i}^2 + \kappa' z_t \tag{2}$$

where x_t and z_t are vectors of weakly exogenous conditioning variables. Engle, Lilien and Robins (1987) term this kind of model ARCH-in-mean or ARCH-M. Note that the h_t^2 is the conditional variance of ε_t

formed at period t based on the information available up to period t-1 (a σ-field). It may also be noted in passing that, by assumption, (2) is a non-stochastic equation.

A further extension of the ARCH formulation, which imposes smoother behaviour on the conditional second moments, has been suggested by Bollerslev (1986). In Bollerslev's GARCH formulation, the conditional second moments are functions of their own lagged values as well as the squares and cross-products of lagged forecast errors. Bollerslev did not consider the GARCH-M extension although this is fairly obvious and was subsequently used in Bollerslev, Engle and Wooldridge (1988). Thus, for example, the GARCH-M (n, p) formulation of the above model would consist of (1) and

$$h_t^2 = A_0 + \sum_{i=1}^{n} A_i \, \varepsilon_{t-i}^2 + \sum_{i=1}^{p} B_i \, h_{t-i}^2 + \kappa' z_t \qquad (3)$$

where the B_i and A_i are coefficients.

Stacking all of the parameters of the system into a single vector

$$\mu = (\alpha, \delta, A_0, A_1, \ldots, A_n, B_1, \ldots, B_p)$$

and applying Schweppe's (1965) prediction error decomposition form of the likelihood function, the log-likelihood for a sample of T observations (conditional on initial values) is proportional to

$$L(\mu) = \sum_{t=1}^{T} -\log \mid h_t^2(\mu) \mid - \sum_{t=1}^{T} \varepsilon_t^2 / h_t^2 \qquad (4)$$

(where we have assumed normality of the forecast errors).

Although the analytic derivatives of (4) can be computed (see Engle, Lilien and Robins, 1987) variable-metric algorithms which employ numerical derivatives are simpler to use and easily allow changes in specification. Under the usual regularity conditions (Crowder, 1976), maximization of (4) will yield maximum likelihood estimates with the usual properties.

Numerical maximization of such a function using data sets of 100-200 observations is now fairly routine, but raises problems when dealing with the data set as large as that described in section 2. For this data set each function evaluation for the maximization routine involves passing through the complete set of 130,000 observations, and a typical maximization will involve many function evaluations. In such a situation solution times can be reduced enormously if good starting values can be supplied to the solution algorithm. So two strategies were used to achieve this. First an OLS model was estimated and this provided a well fitting set of initial values for the structural parameters. Second a GARCH-M model was estimated using only the first 10,000 observations on the assumption that this model should fit the whole data set tolerably well and so provide good starting values for all the parameters.

B OLS estimation

Even the use of OLS is not straightforward for such a large data set as most computers are unable to hold the complete data set in memory at one time. So a set of recursive updating formulas have to be used which allow the OLS results to be built up one observation at a time. The easiest way to do this is by use of the Kalman Filter equations.

To describe this, consider the familiar state space formulation with the appropriate Kalman Filter equations for the univariate case. Following Harvey (1987), let

$$y_t = \delta' z_t + \alpha' x_t + \varepsilon_t \quad (5)$$

be the measurement equation, where y_t is a measured variable, z_t is the state vector of unobserved variables, δ and α are parameters and $\varepsilon_t \sim$ NID$(0,\Gamma_t)$ and x_t is fixed. The state equation is then given as,

$$z_t = \Psi z_{t-1} + \beta' w_t + \omega_t \quad (6)$$

where Ψ and β are parameters, w_t is fixed and $\omega \sim$ NID$(0,Q_t)$.

The only departure from the standard state space form given in many textbooks is in the introduction of x_t into the measurement equation and w_t into the state equations. This is not an important elaboration as long as both x_t and w_t are known at time t. As far as OLS estimation is concerned x_t and w_t can be set to zero in (5) and (6) respectively.

The Kalman Filter prediction equations are obtained by defining $\hat{z}_t$ as the best estimate of z_t based on information up to t and P_t as the covariance matrix of the estimate $\hat{z}_t$, and defining;

$$\hat{z}_{t\mid t-1} = \Psi \hat{z}_{t-1} + \beta' w_t \quad (7)$$

and

$$P_{t\mid t-1} = \Psi P_{t-1} \Psi' + Q_t \quad (8)$$

Once the current observation on Y_t becomes available these estimates can be updated using the following equations,

$$\hat{z}_t = \hat{z}_{t\mid t-1} + P_{t\mid t-1}\, \delta (y_t - \alpha' x_t - \delta' \hat{z}_{t\mid t-1}) / (\delta' P_{t\mid t-1} \delta + \Gamma_t) \quad (9)$$

and

$$P_t = P_{t\mid t-1} - P_{t\mid t-1}\, \delta \delta' P_{t\mid t-1} / (\delta' P_{t\mid t-1}\, \delta + \Gamma_t) \quad (10)$$

Equations (6)-(10) jointly represent the Kalman Filter equations. These equations may be used to generate standard OLS results recursively by setting x_t and w_t equal to zero, Ψ equal to the identity matrix, Q_t equal to a matrix of zeros and setting Γ equal to unity. In this formulation, the z_t are the OLS parameters and δ are the variables in the model. These equations may then be used for OLS estimation on very large data sets, holding only one observation from the data set in computer memory at a time, and using the recursive formulas to build up the full sample OLS results.

If we define the one step ahead prediction errors as,

$$v_t = y_t - \delta' \hat{z}_{t\mid t-1} - \alpha' x_t \quad (11)$$

then the concentrated log likelihood function is proportional to,

$$\log(L) = \sum \log(f_t) + N \log(\sum v_t^2/Nf_t) \tag{12}$$

where $f_t = \delta' P_{t\,|\,t-1}\,\delta + \Gamma_t$ and $N{=}T{-}k$ where k is the number of periods needed to derive estimates of the state vector.

That is to say the likelihood function may be expressed as a function of the one step ahead prediction errors suitably weighted. In order to obtain standard OLS results we do not need to refer to the likelihood function. In the next section we explain how to write the GARCH-M model in state space form and in that case (12) is the form of the concentrated log-likelihood function which needs to be maximized.

C The GARCH-M Model in State Space Form

A comparison of equations (5) and (1) and (6) and (3) will quickly show that the GARCH-M model is already very close to a state space representation, all we need to do is to reinterpret some of the notation of the state space model. For simplicity we will only consider the GARCH(1,1) model but higher order GARCH models can be treated in an exactly analogous fashion. To begin with, in the state equation (6) we interpret w_t to be a 2x1 vector which is made up of a constant and ε_{t-1}^2, this equation is then identical to the GARCH equation (3) where we identify z_t with h_t^2. The z_t term in the measurement equation is then seen as the variance term which is identical to the variance term in equation (1), we then only need to make explicit allowance for the structure of the variance of the measurement equation and this is done simply by equating Γ_t with z_t. The state space model for a stochastic GARCH-M model may then be written out as,

$$y_t = \delta' z_t + \alpha' x_t + \varepsilon_t \quad \varepsilon_t \sim NID\,(0,\hat{z}_{t\,|\,t-1}) \tag{13}$$

and

$$z_t = \Psi z_{t-1} + A_0 + A_1\,\varepsilon_{t-1}^2 + \omega_t \quad \omega_t \sim NID(0,Q_t) \tag{14}$$

This model would be identical to the equations (1) and (3) when Q_t=0 and P_0=0, that is when the GARCH equation is non stochastic with certain initial conditions. This can be seen by noting that if Q_t=0 and P_0=0, then the updating equations, (9) and (10), are non-operational and $f_t = \Gamma_t = z_t$ and the likelihood function becomes equivalent to (4). For a more general description and formulation of a stochastic GARCH-M model see Hall (1991).

D GARCH-M and NEWS

Within the framework of the basic GARCH-M model outlined in (1) and (3), news effects may have two quite distinct effects; either news may effect the level of the exchange rate directly through (1) as an independent news effect, or it may affect the variance of the exchange rate through the GARCH process (3), and then only affect the level through the effect of the variance on the mean via some notion of a risk premia effect. Both of these possibilities may be investigated within the GARCH-M model, the first by including specific news variables in the main equation (1) and the second by including similar effects in the GARCH equation (3). Conventional likelihood ratio or Wald tests may be constructed to test for the significance of these effects.

A further complication arises over the detailed specification of the news dummy effects which may vary not only with the form of the news effect but also with the specification of the rest of the equation for the exchange rate. Although this paper investigates the properties of univariate models of the exchange rate only (mainly because of the lack of any suitable explanatory variable at such high frequencies), this still leaves open the question of whether the exchange rate should be modelled in first difference terms or in levels. Clearly the crucial factor here is whether the exchange rate at such high frequencies is a stationary variable. We do not impose an *a priori* view on this question but we use the following basic formulation for the model which allows for tests of stationarity at each stage. The exchange rate is assumed to be given by the process:

$$\Delta E_t = \alpha_0 + \alpha_1 E_{t-1} + \sum \beta_i \, \Delta E_{t-i} + \gamma N \qquad (15)$$

where E_t is the log of the exchange rate and N is a news effect. This equation takes the same form as the augmented Dickey Fuller regression. Hence if the exchange rate is a random walk and, therefore, non-stationary α_1 would tend to equal zero and the equation reduces to a first difference specification. If on the other hand the exchange rate is actually a stationary process, then α_1 will be significantly less than zero, and the specification becomes a permissible reparameterisation of a normal 'levels' equation in E. The news effect will operate differently, or at least need to be specified differently, to have the same effect depending on which of the two formulations actually dominates. So for example if the news effect changes the level of the exchange rate permanently, then if α_1 is significantly negative so that the equation is actually a 'levels' one, the news effect will need to be represented as a permanent change in the intercept (ie be a variable of the form . .0 0 0 1 1 1 . . .). But if $\alpha_1 = 0$ so that the equation is in first differences then the news effect should operate as a once for all jump (ie a variable like 0 0 0 1 0 0 0. . .). As a practical procedure we specify (15) with a number of lags in the news effect so that the equation can suggest a sensible parameterisation of this effect. The important point being that the specific form of news dummy used varies between the different examples of the exchange rate equation reported in the next section. These specific differences are noted as they arise.

4 Empirical results

(a) Without News

The first set of results are for models without news. Two versions are reported in Table 1. The first is a conventional OLS estimate of equation (15), but where $\gamma=0$. The second is a GARCH-M model so extends (15) to

$$\Delta E_t = \alpha_0 + \alpha_1 E_{t-1} + \sum \beta_i \, \Delta E_{t-i} + \delta h_t^2 \qquad (15')$$

where h^2 is the conditional variance as dictated by a GARCH model (ie equation (3) with $\kappa = 0$), and where the GARCH is assumed to be (1,1).

This set of results is a preliminary to assess the overall adequacy of the model, and to get insight into whether the high frequency data is stationary or non stationary. Table 1 gives the two sets of results for our model with four lagged difference terms. (The choice of four lags is arbitrary, though this seemed to be fairly adequate given the time-series properties of the data. This assumption is maintained in later examples.)

Table 1

The model without news effects

Parameter	OLS MODEL		GARCH MODEL	
α_0	0.00009	(1.3)	0.000056	(1.8)
α_1	-0.000057	(1.4)	-0.000036	(1.9)
β_1	-0.54	(202.7)	-0.569	(252.9)
β_2	-0.29	(97.34)	-0.29	(116.1)
β_3	-0.17	(56.25)	-0.14	(56.82)
β_4	-0.08	(29.36)	-0.058	(25.9)
GARCH parameters				
A_0	—		0.000000027	(1.4)
A_1	—		0.89	(1768.3)
B_1	—		0.208	(1.41)
δ	—		0.00044	(0.037)
SEE	0.0009875		0.7	
SKEW	0.0801		0.64	
KURT	61.09		58.52	
CORRELOGRAM				
1	-0.005		0.02	
2	-0.122		0.011	
3	-0.0219		0.0003	
4	-0.041		-0.01	
8	-0.017		-0.012	
12	0.0043		-0.005	
16	-0.0006		-0.0045	

SEE standard error
Skew Coefficient of skewness (centered on zero, calculated on scaled residuals)
KURT Coefficient of Kurtosis (centered on three, calculated on scaled residuals)
t statistics in parenthesis

Table 2

The preferred model with news effects in the level

Parameter	OLS MODEL		GARCH MODEL	
α_0	0.00009	(1.3)	-0.000092	(3.0)
α_1	-0.000056	(1.4)	-0.000059	(3.2)
β_1	-0.55	(202.7)	-0.564	(251.2)
β_2	-0.30	(97.42)	-0.30	(117.6)
β_3	-0.17	(56.55)	-0.16	(62.60)
β_4	-0.08	(29.23)	-0.059	(26.4)
γ_1	0.008	(14.3)	0.016	(89.2)
γ_2	0.005	(9.51)	0.013	(37.7)
GARCH parameters				
A_0	—		0.000000017	(3.4)
A_1	—		0.89	(1921.8)
B_1	—		0.149	(3.3)
δ	—		0.00153	(0.038)
SEE	0.0009864		0.7	
SKEW	0.0906		0.79	
KURT	61.03		52.91	
CORRELOGRAM				
1	-0.006		0.017	
2	-0.012		0.016	
3	-0.0216		0.0099	
4	-0.0416		-0.016	
8	-0.0180		-0.012	
12	0.0043		-0.005	
16	-0.0005		-0.0045	

SEE standard error
Skew Coefficient of skewness (centered on zero, calculated on scaled residuals)
KURT Coefficient of Kurtosis (centered on three, calculated on scaled residuals)
t statistics in parenthesis

The OLS results suggest that, even at this high frequency, the exchange rate is a non-stationary process (the t statistic on α_1 is well below the Dickey-Fuller critical value). The coefficients on the lagged difference terms are however highly significant and almost exactly sum to unity. There is some sign of non-normality in the errors, as the coefficient of Kurtosis is fairly high suggesting that there is occasionally a large outlier in the data. In turn the residuals seem to be reasonably free from serial correlation, although Box-Pierce tests revealed significant effects. (This is a by-product of the very large number of observations being used, as any degree of serial correlation with a sufficiently large data set will tend to be significant.) Thus, apart from the problem given by large outliers, the OLS assumptions about the error process are not seriously violated. The GARCH estimation results are very similar for the exchange rate equation itself. There are also very well determined auto regressive effects in the conditional variance, as indicated by the B_1 parameter which is highly significant. There is no evidence that the conditional variance affects the level of the exchange rate, however, since the parameter δ is close to zero and insignificant. [The properties of the unscaled error process are similar to the OLS residual, we note that the changing variance may well be the cause of the high Kurtosis.]

(b) With News

Next the effects of the news variables are incorporated, assuming first that the effect of news is on the exchange rate itself. In examining the effects of news in the two periods discussed above, a number of forms for the news dummy were tried and in each case the form of effect chosen had the same general form. This was a sequence ... 0, 0, 0, 1, -1, 1, 0,0, . . ., which allows for a considerable rebound after the initial news shock, but this rebound is quickly reversed. Note that this form of dummy has a permanent effect on the exchange rate if the level term (α_1) in the exchange rate equation is zero, (so that the equation is in effect a difference equation) but that it will have only a temporary effect if the equation is a levels equation ($\alpha_1 \neq 0$). Table 2 gives the results for both the OLS and GARCH-M estimation procedures, and the two sets of results present an interesting contrast.

In both sets of results the news effects are significant, but the rest of the model changes in an important way. The OLS results are similar to those in Table 1, in that α_1 is not significant so that the model suggests that the news has a permanent effect on the exchange rate. The GARCH-M results however suggest that the exchange rate tends to return to some stable level and that the jump which occurs at the time of the news event is subsequently eroded. The initial jump is also seen as being much larger in the GARCH-M results (almost twice the size) and much more significant. The mean lag of the GARCH-M model suggests that half the news effect will be eroded over the subsequent 16,600 offers to trade which appear on the screens. This would generally be a period of around four to five days, so this model suggests that the effects of these news events tended to disappear over the following week. It is interesting that only the GARCH-M model has this implication. In turn these results suggest that the OLS assumption of a constant variance causes big errors to occur around the times of these news events, which tends to give distorted results precisely because the error variance is not allowed to change.

Table 3

The preferred model with news effects in the variance term

Parameter	OLS MODEL	GARCH MODEL	
α_0		0.0020	(52.2)
α_1		-0.00111	(47.6)
β_1		-0.484	(175.1)
β_2		-0.25	(150.1)
β_3		-0.13	(56.52)
β_4		-0.046	(15.7)
γ_1		0.013	(19.7)
γ_2		0.0056	(2.0)
γ_3		0.0088	(27.7)
GARCH parameters			
A_0		0.000000051	(47.7)
A_1		0.25	(136.3)
B_1		0.090	(61.1)
δ		0.00102	(0.010)
C_1		0.86(E-6)	(13.1)
C_2		0.82(E-8)	(10.3)
SEE		2.34	
SKEW		0.202	
KURT		68.96	
CORRELOGRAM			
1		-0.038	
2		0.004	
3		0.002	
4		-0.16	
8		-0.0055	
12		0.0087	
16		0.0048	

SEE standard error
Skew Coefficient of skewness (centered on zero, calculated on scaled residuals)
KURT Coefficient of Kurtosis (centered on three, calculated on scaled residuals)
t statistics in parenthesis

It is entirely plausible that news affects not only the level of the exchange rate but also its variance. This would of course be partly captured by the standard GARCH formulation in that once a large error had occurred it would tend to persist through the lagged dependent variable. This GARCH model will however always miss the initial announcements period, and so it is possible that the omission of direct news effects from the variance process would bias the estimates of the GARCH parameter upwards. The reason is that, if the true explanation for the increase in the variance is missing, this will tend to be explained by finding serial correlation. A better way of accounting for this feature however is to add the news dummies into the GARCH equation so that the initial increase in variance is captured correctly. It should be remembered that the market knows in advance the date when news will be announced and so it is plausible that an anticipated increase in uncertainty could occur in advance of the news announcement itself. To include this in the model the GARCH equation (3) is augmented with news dummies, with parameters C_1 and C_2. The model may then be re-estimated with these extra effects in the GARCH process. The results of this final exercise are presented in Table 3.

We have respecified the dynamic structure of the news dummies in the exchange rate equation in the light of the new dynamics of the model. The first news event is now simply represented by a single dummy γ_1 which is equal to minus one as the news is announced and zero everywhere else. The second event is represented by two dummies, γ_2 —which equals one as the news is announced—and γ_3 which equals 1 followed by -1 in the subsequent periods to allow for a dynamic response to the announcement.

The results reported in Table 3 show a number of important differences from those given in Tables 1 and 2. The most notable are that the size of the parameters in the GARCH equation have fallen substantially, as both A_1 and B_1 are much smaller, which suggests that the degree of serial correlation is much reduced (although it is still significant). The size of the lagged level affect (α_1) has also increased dramatically which suggests that the exchange rate will return to a given level after a shock much more quickly. Both C_1 and C_2 are positive which suggests that both sets of news effects act to increase the variance, as we would expect. Although α_1 is significantly different from zero it is still a very small coefficient which suggests that the dynamic process is still very close to a random walk. Nonetheless its significance suggests that, at least over a reasonably short period of about 8 weeks, the exchange rate does tend to return to an equilibrium level. Thus, the effects of the two news announcements used in these regressions were reversed over the following few days.

Conclusion

We have investigated the effects of introducing news effects into an extremely high frequency model of the sterling-dollar exchange rate. Initial estimates without news effects find a near integrated conditional variance process with a unit root exchange rate equation. The introduction of news effects into the model changes this result dramatically. In the case when news is allowed to affect the level of the exchange rate the results suggest that the level of the exchange rate is not a random walk but that it is a stable process. Allowing the news effects to enter the conditional variance process of the exchange rate strengthens this result dramatically and also changes the parameterisation of the GARCH process from being one which is very close to an integrated one to one which is clearly stable.

Our conclusion then is that inference about the time series properties of very high frequency data can be seriously distorted by a small number of major shocks. When these shocks are included in the analysis the point estimates of parameters can change significantly as can the inferences which can be drawn about model properties. In the case presented here the evidence seems fairly strong that the news effects we analyse influence both the level and the uncertainty of the exchange rate but, for this sample period, these effects are not permanent. In this case both the amount of volatility and the level of the exchange rate tend to return to their pre-announcement values.

References

Baba Y, Engle R F, Kraft D F and Kroner K F, (1987), 'Multivariate Simultaneous Generalized ARCH' Mimeo University of California San Diego.

Bollerslev T, (1986), 'Generalized Autoregressive Conditional Heteroskedasticity', Journal of Econometrics, vol 31, no 3, pages 307-28.

Bollerslev T, Engle R F and Wooldridge J M, (1988), 'A Capital Asset Pricing Model with Time-varying Covariances', Journal of Political Economy, vol 96, no 1, pages 116-31.

Crowder M J, (1976), 'Maximum Likelihood Estimation for Dependent Observations', Journal of the Royal Statistical Society, Series B, vol 38, no 1, pages 45-53.

Engle R F, (1982), 'Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of UK Inflation' Econometrica, vol 55, no 4, pages 987-1,008.

Engle R F, (1987), 'Multivariate GARCH with Factor Structures—Cointegration in Variance', paper presented at a conference on time-varying variances in money and finance, University of California at San Diego, April.

Engle R F, Lilien D and Robins R, (1987), 'Estimating Time Varying Risk Premia in the Term Structure: The Arch-M Model', Econometrica, vol 55, no 2, pages 391-407.

Griliches Z,(1983), 'Introduction: Data and Econometrics—The Uneasy Alliance' Chapter 25, 'Handbook of Econometrics' vol 33 (ed) Griliches Z and Intriligator M D North Holland.

Goodhart C, (1990), "'News' and the Foreign Exchange Market", forthcoming in the Proceedings of the Manchester Statistical Society.

Goodhart C and Figiuoli L, (1991), 'Every Minute counts in Financial Markets', forthcoming in the Journal of International Money and Finance.

Hall S G, (1991), 'A Note on the Estimation of GARCH-M Models Using the Kalman Filter', Bank of England working paper.

Harvey A C, (1987), 'Applications of the Kalman Filter in Econometrics' in Bewley T F.(ed) 'Advances in Econometrics. Fifth World Congress' vol 1, Econometric Society Monograph No.13 Cambridge University Press.

Schweppe F, (1965), 'Evaluation of Likelihood Functions for Gaussian Signals', IEEE Transactions on Information Theory, vol 11, no 1, pages 61-70.

Van Praag B, (1983), 'The Population-Sample Decomposition in Minimum Distance Estimation' Unpublished, Presented at the Harvard-MIT Econometric Seminar.

White H, (1982), 'Maximum Likelihood Estimation of Misspecified Models', Econometrica Vol 50, no 1, pages 1-25.

Bank of England Discussion Papers

	Title	Author
1-5, 8, 11-14, 16-17, 19-22, 31	*These papers are now out of print, but photocopies can be obtained from University Microfilms International* [(a)]	
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
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
26	An empirical model of company short-term financial decisions: evidence from company accounts data	Mrs G Chowdhury C J Green D K Miles
27	Employment creation in the US and UK: an econometric comparison	I M Michael R A Urwin
28	An empirical model of companies' debt and dividend decisions: evidence from company accounts data	Ms G Chowdhury D K Miles
29	Expectations, risk and uncertainty in the foreign exchange market: some results based on survey data	M P Taylor
30	A model of UK non-oil ICCS' direct investment	E J Pentecost
32	The demographics of housing demand; household formations and the growth of owner-occupation	M J Dicks
33	Measuring the risk of financial institutions' portfolios: some suggestions for alternative techniques using stock prices	S G F Hall D K Miles
34	An error correction model of US consumption expenditure	I R Harnett
35	Industrial structure and dynamics of financial markets; the primary eurobond market	E P Davis
36	Recent developments in the pattern of UK interest rates	D K Miles
37	Structural changes in world capital markets and eurocommercial paper	J G S Jeanneau
38	Stockbuilding and liquidity: some empirical evidence for the manufacturing sector	T S Callen S G B Henry
39	The relationship between employment and unemployment	M J Dicks N Hatch
40	Charts and fundamentals in the foreign exchange market	Mrs H L Allen M P Taylor
41	The long-run determination of the UK monetary aggregates	S G Hall S G B Henry J B Wilcox
42	Manufacturing stocks; expectations, risk and co-integration	T S Callen S G Hall S G B Henry
43	Instability in the euromarkets and the economic theory of financial crises	E P Davis
44	Corporate governance and the market for companies: aspects of the shareholders' role	J Charkham
45	Stock-flow consistent income for industrial and commercial companies: the UK experience	K D Patterson
46	The money transmission mechanism	D K Miles J B Wilcox
47	Monetary aggregates in a changing environment: a statistical discussion paper	R D Clews Ms J E C Healey Glenn Hoggarth C R Mann
48	A model of manufacturing sector investment and employment decisions	J W Lomax
49	A simple model of the housing market	M J Dicks
50	An industrial approach to financial instability	E P Davis
51	International financial centres—an industrial analysis	E P Davis
52	A model of ICCs' dividend payments	J W Lomax
53	The determination of average earnings in Great Britain	M A S Joyce
54	Cross-border savings flows and capital mobility in the G7 economies	Shelley Cooper

Technical Series

	Title	Author
1-11, 14, 20, 23	*These papers are now out of print, but photocopies can be obtained from University Microfilms International* [(a)]	
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
15	International comparison of asset market volatility: a further application of the ARCH model	R R Dickens
16	A three sector model of earnings behaviour	D J Mackie
17	Integrated balance sheet and flow accounts for insurance companies and pension funds	Raymond Crossley
18	Optimal control of stochastic non-linear models	S G Hall I R Harnett M J Stephenson
19	A multivariate GARCH in mean estimation of the capital asset pricing model	S G Hall D K Miles M P Taylor
21	Modelling of the flow of funds	D G Barr K Cuthbertson
22	Econometric modelling of the financial decisions of the UK personal sector preliminary results	D G Barr K Cuthbertson
24	Modelling money market interest rates	J S Flemming D G Barr
25	An independant error feedback model of UK company sector asset demands	D G Barr K Cuthbertson
26	A disequilibrium model of building society mortgage lending	S G Hall R A Urwin
27	Balancing the national accounts: an asymptotically maximum likelihood approach using trends	G P Dunn D M Egginton
28	Testing a discrete switching disequilibrium model of the UK labour market	S G Hall S G B Henry M Pemberton
29	The Bank of England Model 1989: recent developments and simulation properties	F J Breedon A J Murfin S H Wright
30	A data-based simulation model of the financial asset decisions of UK, 'other' financial intermediaries	D G Barr K Cuthbertson
31	The demand for financial assets held in the UK by the overseas sector: an application of two-staged budgeting	D G Barr K Cuthbertson
32	A note on the estimation of GARCH-M models using the Kalman Filter	S G Hall
33	Modelling the sterling effective exchange rate using expectations and learning	S G Hall
34	Modelling short-term asset holdings of UK banks	D G Barr K Cuthbertson
35	A Monte Carlo study of alternative approaches to balancing the national accounts	D M Egginton
36	Sterling's relationship with the dollar and the deutschemark: 1976–89	A G Haldane S G Hall
37	Using and assessing CBI data at the Bank of England	Bahram Pesaran C B Wright
38	A system approach to consumption and wealth	S G Hall K D Patterson
39	Exchange rate equations	Helen Allen Brian Henry Bahram Pesaran
40	Inventories, technological change, financial effects and risks	Keith Cuthbertson David Gasparro
41	News effects in a high frequency model of the sterling-dollar exchange rate	C A E Goodhart S G Hall S G B Henry Bahram Pesaran

(a) 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