

A utility-based comparison of some models of exchange rate volatility

Kenneth D. West

Department of Economics, University of Wisconsin, Madison, WI 53706, USA

Hali J. Edison

Board of Governors of the Federal Reserve System, Washington, DC 20551, USA

Dongchul Cho*

Department of Economics, Texas A&M University, College Station, TX 77843, USA

Received February 1992, revised version received September 1992

When estimates of variances are used to make asset allocation decisions, underestimates of population variances lead to lower expected utility than equivalent overestimates: a utility-based criterion is asymmetric, unlike standard criteria such as mean squared error. To illustrate how to estimate a utility-based criterion, we use five bilateral weekly dollar exchange rates, 1973–1989, and the corresponding pair of Eurodeposit rates. Of homoskedastic, GARCH, autoregressive and non-parametric models for the conditional variance of each exchange rate, GARCH models tend to produce the highest utility, on average. A mean squared error criterion also favors GARCH, but not as sharply.

1. Introduction

This paper evaluates the out-of-sample performance of some univariate models for exchange rate volatility, using bilateral weekly data for the dollar versus the currencies of Canada, France, Germany, Japan and the United Kingdom, 1973–1989, and the corresponding pair of Eurodeposits, 1981–1989. The models considered include homoskedastic, GARCH, and non-parametric ones, as well as autoregressions in both the absolute value and square

Correspondence to: Prof. K.D. West, Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, USA.

*We thank an anonymous referee, Buz Brock, Frank Diebold, Takatoshi Ito, Blake LeBaron, Robin Lumsdaine and participants in various seminars for helpful comments and discussions, Karen Lewis for providing data, and John Hulbert for research assistance. West thanks the National Science Foundation, the Sloan Foundation, and the University of Wisconsin Graduate School for financial support. This paper represents the views of the authors and not necessarily those of the Board of Governors of the Federal Reserve System or other members of its staff.

of exchange rate changes. The metric we use to compare the models is a utility-based one: How much would an investor with a mean–variance utility function, who uses the estimates of one of these models to divide her wealth between a pair of Eurodeposits, be willing to pay to use one model rather than another?

Recent research on conditional volatility has established that for many financial variables, including exchange rates, squared changes that are large tend to be followed by squared changes that are also large [Bollerslev et al. (1990)]. This empirical fact has stimulated a variety of formal statistical models. Since the relative merits of many of these models are as yet not well established, there is a need for systematic evaluation and comparison.

Some previous authors have compared the out-of-sample performance of univariate models applied to stock price data. Using a mean squared error criterion, Pagan and Schwert (1990) found that GARCH and ARMA models are preferred to non-parametric and Markov switching ones, and Akgiray (1989) that GARCH dominates naive and ARMA models. Using a criterion based on performance in a simulated market, Engle et al. (1990) also found GARCH preferable to naive and ARMA models. Finally, Friedman and Kuttner (1988) compared multivariate GARCH and AR models, using stock and bond data. Among other statistics, they examined mean squared errors, but did not seem to find strong grounds for preferring one model to another.

One inessential sense in which the present paper differs from any of these is in its use of exchange rate data, which we study largely because such data apparently have yet to be used in a systematic comparison of volatility models. More importantly, we also depart from earlier work in how we measure performance. An appropriate measure of performance depends on the use to which one puts the estimates of volatility, and our measure is probably not the best if one wants to, say, study the links between observable macro variables and volatility [e.g. Schwert (1989a)]. But insofar as models for volatility are motivated by references to investment by risk-averse utility maximizers – as, indeed, they often are [e.g. Engle and Bollerslev (1986); Friedman and Kuttner (1988)] – a utility-based measure seems quite appropriate.¹

Our measure is based on the following presumption: at a given point in time, one estimate of a conditional variance is better than another if investment decisions based on it lead to higher (population) expected utility. Similarly, an estimator or model of a conditional variance is preferred if, on average, over many time periods, it leads to higher expected utility. We show that under the assumption that utility is either (a) exponential, and asset

¹For a model of the conditional mean of stock returns, McCulloch and Rossi (1990) also use a utility approach, and Breen et al. (1989) aim, as do we, to estimate how much an investor would pay to use a model.

returns are jointly normal, or (b) quadratic, such a utility-based criterion is fundamentally different from statistical ones based on mean squared and mean absolute error: the utility criterion is asymmetric, with underestimates of the population conditional variance-covariance matrix leading to lower expected utility than equivalent overestimates.

To illustrate the use of our measure empirically, we assume that an investor divides her wealth between two assets, weekly or quarterly Euro-deposits denominated in (1) dollars and (2) the currency of one other country (Canada, France, Germany, Japan, or the United Kingdom). We consider an investor who knows the population conditional variance of exchange rate changes, but is forced to make a wealth allocation using not the population value but one of a set of estimates. Different estimates will lead to different wealth allocations and, therefore, different levels of expected utility. We envision the investor using estimates from each of our models to produce a sequence of hypothetical wealth allocations over a number of successive periods, and ask the following: Which model's implied allocations produce the highest expected utility, on average, and how much would such an investor pay for the right to allocate wealth according to that model rather than another?

For quadratic utility, we show that one can estimate the average expected utility produced using a given volatility model, even when one does not have our hypothetical investor's knowledge of the time series of population conditional variances. If, for a given level of beginning-of-period wealth, one model produces higher expected utility, on average, than does another, then the better model will produce equal average utility with a lower beginning level. We interpret the difference in beginning wealth as the average per period fee that our hypothetical investor would be willing to pay to use the higher rather than lower utility model.

Although there was some variation across data sets, we find that GARCH models tend to do best. Depending on the dataset, an investor would typically be willing to pay about 0.05 to 2 percent, or 5 to 200 basis points of her wealth, annually, to switch to GARCH from another model. Confidence intervals around these point estimates, however, tend to be large. The *t*-statistics indicate that the fee is statistically significantly different from zero at conventional levels only about one-fourth of the time; *F*-tests of the null that all six models yield the same utility are significant a little less than half the time. Under an out-of-sample mean squared error criterion, however, the statistical significance of differences across models is even less pronounced.

One way to gauge the economic significance of the utility-based figures is to interpret them as a transactions fee that a professional money manager could charge an investor capable of estimating, say, homoskedastic but not GARCH models of exchange rate risk. As such, the 5 to 200 range seems to bracket what Wall Street mutual funds charge for their services [Ippolito

(1989), *New York Times*, 14 May, 1991, p. F14], which seems to us a substantial figure.

While the immediate motivation for our research is the relatively recent literature on conditional volatility, our results are relevant for evaluation of any models for second moments of asset returns. Eun and Resnick (1984), for example, use a mean squared error criterion in evaluating models for correlations across share prices, motivating their study with reference to mean-variance portfolio analysis. An implication of this paper is that such a criterion is probably not the best.

Before turning to the body of the paper, two introductory cautions seem advisable, to set the reader's expectations straight. First, while we have tried to make a sensible choice of models to study; we do not claim to be comprehensive, and some readers may feel that we have unwisely excluded some important models. For such readers we emphasize that we consider one of our contributions to be the technique used to produce the rankings of the models. Second, we abstract from a number of potentially important complications involved in real-world investments. We ignore, for example, default risk, transactions costs such as bid-asked spreads, and issues about the timing of settlement of transactions, including that our exchange and interest rate series, which we obtained from two different sources, are sampled at slightly different times (of the same day); we also acknowledge that the very simple portfolios that we consider are not well diversified. Our aim is simply to get a rough idea of the magnitude of the potential benefits of better volatility models, not to quantify these benefits to many decimal places.

Section 2 briefly outlines the motivation for our utility-based measure rather than a standard statistical one, in a more general framework than is required for our empirical work. Section 3 describes how we apply our measure to exchange rate data, section 4 our data and models, and section 5 our empirical results. Section 6 concludes. An appendix contains some technical details, and an additional appendix, available from the authors upon request, contains some results omitted from the paper to save space.

2. Utility versus statistical measures of estimator quality

A basic message of our paper is that when comparing estimators of conditional variances, rankings from a utility-based criterion might differ from those from a statistical mean squared or mean absolute error criterion, because of a certain asymmetry in utility evaluation of estimators. A general statement is given in the proposition below.

We begin, however, with a simple numerical example. This example does *not* illustrate the asymmetry, but it does point out that utility and statistical measures may be dramatically different, and thus motivates our desire to

estimate a utility based measure. Suppose one has an exponential utility function, $U(W_{t+1}) = -\exp(-\theta W_{t+1})$, where $\theta > 0$ and W_{t+1} is period $t+1$ wealth. Suppose that there are three assets, one riskless. Let $\mu \equiv (\mu_1, \mu_2)'$, H denote the mean and covariance matrix of the (2×1) vector of excess returns, and $f \equiv (f_1, f_2)'$ the (2×1) vector of fractions of period t wealth put in the two risky assets. As is well known, maximization of expected utility leads to $f \equiv (1/\theta W)H^{-1}\mu$, where W is period t wealth. Suppose further that H is the identity matrix, and $\mu_1 = \mu_2 > 0$. Then the optimal fraction satisfies $f_1 = f_2 = (\mu_1/\theta W) = (\mu_2/\theta W)$.

Assume that an investment decision must be made using the true μ and one of two noisy estimates of H ,

$$\hat{H}_1 = \begin{vmatrix} 2 & -1 \\ -1 & 2 \end{vmatrix}, \quad \hat{H}_2 = \begin{vmatrix} 1 & 0 \\ 0 & 1.001 \end{vmatrix}.$$

Which is the better estimate of the true H (which equals the identity matrix)? By common statistical measures such as the average of the squared differences between the non-redundant elements of H and the \hat{H}_i 's, \hat{H}_2 is 'closer' to H and therefore is better. But a routine calculation indicates that \hat{H}_1 leads to exactly the optimal (expected utility maximizing) fraction. The basic presumption of this paper is that \hat{H}_1 is therefore a preferable estimate.

This numerical example obviously is special. We now state a proposition that illustrates that, in a very general sense, utility and standard statistical criteria are different. Let W_{t+1} be wealth in period $t+1$. Assume: (1) The utility function is either (a) exponential, $U(W_{t+1}) = -\exp(-\theta W_{t+1})$, $\theta > 0$, and asset returns are normally distributed, or (b) quadratic, $U(W_{t+1}) = W_{t+1} - 0.5\gamma W_{t+1}^2$, $\gamma > 0$, and asset returns have finite means and variances. (2) There are $k \geq 1$ risky assets, with positive definite variance-covariance matrix H . There may or may not be a $(k+1)$ st riskless asset. If not, $k \geq 2$; if so, expected returns on the risky assets are greater than those on the riskless asset. (3) There are no constraints on short sales; the fraction of wealth put in a given asset may be less than zero or greater than one. (4) The population conditional mean of returns is used in making investment decisions.

Assumptions (1)–(3) are used to get a convenient closed-form solution. Assumption (4) is used to focus on the effects of errors in estimation of H . Note that this assumption rules out a comparison of various parameterizations of GARCH-M models, for example.

Suppose we wish to compare two estimates of H , \hat{H}_1 and \hat{H}_2 . Let $E_t U_{it+1}$, $i=1, 2$, denote expected utility that results when model i is used to make an investment decision, where the true variance-covariance matrix H is used in computing expected utility. Our basic result is that there is an asymmetry in the utility loss from estimation error, with estimates of H that are too large being preferred to those that are too small.

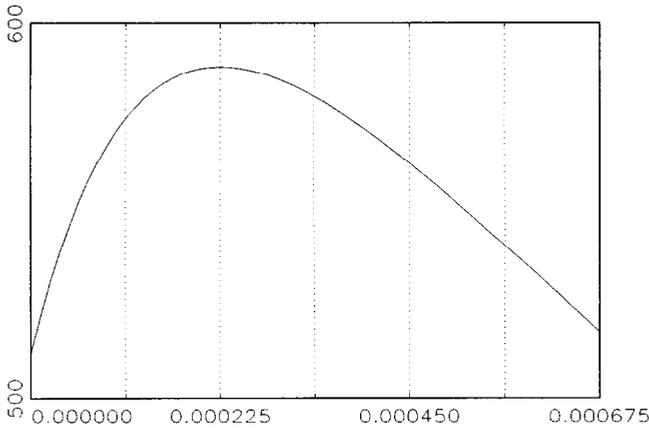


Fig. 1. Graph of $u(h_t, \hat{h}_{mt})$ as a function of \hat{h}_{mt} , $h_t = 0.000225$.

Proposition. Suppose that $\hat{H}_1 = H + V$ and $\hat{H}_2 = H - V$, where V is a positive semidefinite symmetric matrix and \hat{H}_2 is a positive definite matrix. Then $E_t U_{1t+1} \geq E_t U_{2t+1}$; equality holds if and only if use of \hat{H}_1 and of \hat{H}_2 result in the allocation implied by use of the population variance-covariance matrix H .

Algebra to derive the proposition is in the additional appendix available on request.

To illustrate the proposition, consider fig. 1, which plots expected utility as a function of \hat{h}_{mt} , the estimate of H , when H is a scalar and utility is quadratic with parameters matching those in our empirical work. By assumption, highest expected utility occurs when $\hat{h}_{mt} = H = (0.015)^2 = 0.000225$. Expected utility declines the farther away is \hat{h}_{mt} from h_t .² What is to be noted is that, in contrast to the usual mean squared or mean absolute error criterion, this objective function is asymmetric around h_t , penalizing estimates that are too small more sharply than those that are too large. As we shall see, this asymmetry plays a role in the empirical results.

3. Estimation of average utility

It is helpful to begin by defining some notation. Let

$$e_t = \log \text{ difference of weekly exchange rate (dollars per unit of foreign currency);} \tag{1a}$$

²This scalar result does not generalize in an obvious way to higher dimensions. If H is a matrix, it is possible to have $\hat{H}_1 > \hat{H}_2 > H$ (where for matrices A and B , $A > B$ means $A - B$ is positive definite) with $E_t U_{1t+1} > E_t U_{2t+1}$.

$$h_{t,j} = \text{var}_t(e_{t+j}) = E_t e_{t+j}^2 = (\text{population}) \text{ variance of } e_{t+j}, \text{ conditional on information generated by past } e_s, s \leq t; \quad (1b)$$

$$\hat{h}_{mt,j} = \text{fitted conditional variance of } e_{t+j}, \text{ according to model } m \text{ (e.g. model } m \text{ is GARCH(1,1) or homoskedastic), estimated using data on past } e_s, s \leq t; \quad (1c)$$

$$h_t \equiv h_{t,1}, \hat{h}_{mt} \equiv \hat{h}_{mt,1}; \quad (1d)$$

$$N = \text{endpoint of first sample used in estimation}; \quad (1e)$$

$$T = \text{endpoint of last sample used in estimation.} \quad (1f)$$

Note that in (1b) the conditional variance is equated with the raw (as opposed to central) conditional second moment of e_t . This is in accord with Meese and Rogoff (1983), Diebold and Nason (1990), Meese and Rose (1991) and the findings summarized below that the conditional first moment of e_t is zero. For concreteness in interpreting (1b) and (1c), it may help to note that in the tables below we report results for weekly and quarterly horizons for investment decisions, which require estimates of $h_{t,1}$ (weekly) and $h_{t,1}, h_{t,2}, \dots$, and $h_{t,13}$ (quarterly). To do so for, say, weekly horizons, we obtain for each model $T - N + 1$ fitted values $\hat{h}_{mt}, t = N, \dots, T$, for models $m = 1, \dots, M$, where the number of models M in the tables below is 6. Note, finally, our dating convention: what we denote h_t corresponds to what is often called h_{t+1} or σ_{t+1} [e.g. Engle (1982)].

We specialize the general environment described in the previous section to one in which utility is quadratic. We assume a two-country world with two assets, one sold domestically, the other sold abroad. To focus on the question at hand, we assume that apart from the conditional variance of exchange rate changes, all relevant moments of the return distribution are known.

Let the utility function and wealth constraint be

$$\begin{aligned} \text{utility in period } t+1 &= W_{t+1} - 0.5\gamma W_{t+1}^2, \\ W_{t+1} &= W_t [f_t(R_{t+1}^* + e_{t+1}) + (1-f_t)R_{t+1}], \end{aligned} \quad (2)$$

where W_{t+1} and γ are, as in the previous section, wealth and taste for risk, f_t is the fraction of wealth put in the foreign asset (possibly negative, possibly greater than one), R_{t+1}^* is the gross return on the foreign asset in terms of foreign currency, and R_{t+1} is the gross return on the domestic asset.

For each period, use each model, one by one, to choose the fraction of wealth that maximizes expected utility, taking each model's point estimate for the conditional variance as the correct expectation. Given the assumption that the mean return on the asset is known, in a given period the optimal fraction will vary across competing models only insofar as the estimates of the conditional variance vary. Let f_{mt} be the fraction that results when model m is used [the exact formula is given in eq. (A.1) in the appendix],

$W_{mt+1} = W_t[f_{mt}(R_{t+1}^* + e_{t+1}) + (1 - f_{mt})R_{t+1}]$ and U_{mt+1} the resulting wealth and utility. If f_{mt} depends only on information known at time t – as it will in an out-of-sample study such as ours – it is straightforward to show that mathematically expected utility may be written

$$E_t U_{mt+1} \equiv E_t[W_{mt+1} - 0.5\gamma W_{mt+1}^2] = W_t[c_t + d_t u(h_t, \hat{h}_{mt})], \quad (3)$$

where E_t is mathematical expectations, c_t and d_t are certain variables that depend on γ and $R_{t+1} - R_{t+1}^*$ but not on h_t , and $u(h_t, \hat{h}_{mt})$ is a certain function that is linear in h_t . Explicit formulas are given in the appendix. Fig. 1, which was already discussed in the preceding section, plots $u(h_t, \hat{h}_{mt})$ as a function of \hat{h}_{mt} , for $h_t = (0.015)^2$ (approximately the sample variance of e_t in our data).

We cannot use (3) directly to determine which model yields the highest mathematically expected utility, since the whole problem is that we do not know the population (mathematical) expectation h_t . But since $u(h_t, \hat{h}_{mt})$ is linear in h_t , we can get an estimate that is right on average by replacing h_t with the ex-post realized value e_{t+1}^2 . We therefore compute *average utility* for a given model by replacing h_t with the ex-post realized value e_{t+1}^2 and averaging the result for $t = N, \dots, T$, with W_t held fixed at a constant level W :

$$(T - N + 1)^{-1} \sum_{t=N}^T W[c_t + d_t u(e_{t+1}^2, \hat{h}_{mt})] \equiv \bar{U}_m. \quad (4)$$

In a large sample, this will be close to the average of the conditional mathematical expectation, $(T - N + 1)^{-1} \sum_{t=N}^T E_t U_{mt+1}$. Of course the asymmetry in fig. 1 now revolves around e_{t+1}^2 rather than h_t .

Average utility depends on taste for risk. Consider fixing the coefficient of relative risk aversion (CRRA), which for quadratic utility is $\gamma W / (1 - \gamma W)$. In this case, the variables c_t and d_t in (3) do not depend on W and expected utility is linearly homogeneous in wealth: double wealth (holding the CRRA constant) and expected utility doubles. (Of course, by fixing relative risk aversion rather than γ , we are implicitly interpreting quadratic utility as an approximation to a non-quadratic utility function, with the approximating choice of γ dependent on wealth.)

By definition, the optimal model requires less wealth to yield any specified level of average utility than does any given suboptimal model. We interpret the difference in required wealth as the average per period fee that the investor would be willing to pay to switch from a suboptimal to the optimal model. For a given suboptimal model, we report the ratio of this fee to an initial level of wealth (the exact level is arbitrary, since the linear homogeneity noted in the previous paragraph means that ratio is independent of the initial level). For convenience of interpretation, we express this in annual basis points.

Example. Suppose that with a horizon of one week, an optimal GARCH model with initial wealth of \$9,999 yields the same average utility as does a suboptimal homoskedastic model with initial wealth of \$10,000. Then we report an annualized fee of $52 \text{ weeks/year} \times [(\$10,000 - \$9,999) / \$10,000] \times 100 \times 100 = 52$, where the first 100 converts to percentage and the second to basis points.

The appendix shows that if, say, model 1 is the optimal model, model m an arbitrary suboptimal model, this fee may be computed as

$$(52/j) \times 10,000 [1 - (\bar{U}_m / \bar{U}_1)], \quad (5)$$

where j is the horizon, $j=1$ or 13.

How does variation in risk aversion affect this fee? In the general framework of the previous section, the effects are ambiguous. But when there is a risk-free asset, as in our empirical work, it can be shown that the expected utility benefits of a better model are lower for more risk-averse investors: if, say, $E_t U_{1t+1} - E_t U_{mt+1} > 0$ (i.e. model 1 is better than model m), then $\partial(E_t U_{1t+1} - E_t U_{mt+1}) / \partial(\text{CRRA}) < 0$. The intuition is that greater risk aversion leads to larger fractions of wealth in the safe asset and less variation in expected outcomes across models. A likely empirical implication is that for a given time series of returns and volatility estimates, the higher is the CRRA $\gamma W / (1 - \gamma W)$, the lower will be the estimated value of (5).³

4. Data and models

4.1. Data

Our exchange rates are measured as dollars per unit of foreign currency, between the United States and Canada, France, Germany, Japan and the United Kingdom.⁴ The data are Wednesday, New York noon bid rates, as published in *The Federal Reserve Bulletin*.

The returns R_{t+1} and R_{t+1}^* are Eurodeposit rates. For one week maturities, the data are from the London market. Wednesday closing rates (which we believe are at noon New York time, apart from variation induced by daylight savings), average of bid and asked, were available for France, Germany, Japan, the United Kingdom and United States (but not Canada). These were kindly supplied by Karen Lewis; the ultimate source is *The*

³It is not absolutely certain that in any given sample increased risk aversion will lead to a lower fee; a sufficient condition is $U_{1t+1} - U_{mt+1} > 0$ for all t .

⁴We also obtained Italian data. But in-sample statistics suggest a *non-zero* unconditional mean. We dropped Italy rather than fit means as well as variances.

London Financial Times. We cleaned up some obvious recording errors before using these data (details available on request). For one quarter maturities, the data are generally from the Zurich market, occasionally (when Zurich data were not available) from the London market. Wednesday bid rates, 10:00 AM Swiss time (4:00 AM New York time, again apart from variation induced by daylight savings) were available for all six countries. The sources is the Bank of International Settlements. For both exchange and interest rate data, when Wednesday was a holiday we used Thursday data; when Thursday was a holiday as well we used Tuesday data.

After an initial observation was lost due to differencing the exchange rate data, the exchange rate sample for each country included the 863 observations from 14 March 1973, to 20 September 1989. Plots and summary statistics on the exchange rates are presented in West and Cho (1992). To conserve space, we limit ourselves here to a brief summary of the familiar pattern; exchange rate changes appear to have zero unconditional means, be serially uncorrelated, have zero skewness and very fat tails; the squares of exchange rate changes appear to be highly serially correlated.

We arbitrarily began our out-of-sample exercise at the midpoint of the exchange rate data, and the first sample for which we fit any volatility models included the 432 observations from 14 March 1973 to 17 June 1981 [$N=432$ in the notation of eq. (1e)]; accordingly, the first interest rate observations that we used were those for 17 June 1981. As we added additional observations, we rolled the sample, fixing the sample size at 432, and dropping what had been the initial observation as each additional observation was added on. The final week used in estimation was 5 April 1989 [$T=839$ in eq. (1f)], which means that our final sample spanned the 432 observations from 17 December 1980 to 5 April 1989 and the number of forecasts, as well as the size of our sample of interest rate observations, was 408. An earlier version of this paper tried not only 1 and 13 week but 4 and 24 week horizons as well, and this accounts for our withholding the final 24 (instead of 12) weeks of data [i.e. accounts for $T=839$ instead of $T=851$ in (1f)]. Results for 4 and 24 weeks are not reported since they are similar to those for 1 and 13 weeks.

Table 1 contains some basic statistics on the foreign–U.S. differential. For ease of interpretation, these are expressed at annualized rates; the corresponding weekly or quarterly rates were used in the empirical work. The standard errors here and in subsequent tables were computed by (1) applying the asymptotic theory in Hansen (1982) to the moment conditions used to produce the estimates, and (2) using the Newey and West (1987) technique to estimate a certain spectral density that this theory requires.

According to table 1, the differentials have broadly similar patterns. Lines (7) and (9) indicate that they tended to stay within a band about 3 percentage points wide; line (11), which in columns (1) to (5) gives the

Table 1
Summary statistics on annualized interest rate differentials.

| | (1) Canada | (2) France | (3) Germany | (4) Japan | (5) U.K. | (6) U.S. |
|--------------------------------|-----------------|-----------------|-----------------|-----------------|-----------------|----------------|
| (a) Weekly | | | | | | |
| (1) Mean | | 3.76 (1.21) | -3.59 (0.22) | -3.61 (0.43) | 1.65 (0.43) | 9.32 (0.52) |
| (2) Standard deviation | | 17.94 (7.90) | 1.32 (0.14) | 2.47 (0.39) | 2.53 (0.27) | 2.89 (0.45) |
| (3) ρ_1 | | 0.50 (0.08) | 0.90 (0.036) | 0.95 (0.02) | 0.94 (0.03) | 0.96 (0.02) |
| (4) ρ_2 | | 0.07 (0.03) | 0.86 (0.04) | 0.92 (0.02) | 0.89 (0.04) | 0.94 (0.03) |
| (5) Correlation with U.S. rate | | 0.08 (2.78) | -0.70 (0.32) | -0.91 (0.58) | -0.76 (0.45) | 1.00 |
| (6) Minimum | | -2.51 | -7.56 | -12.87 | -7.37 | 5.75 |
| (7) Q1 | | 0.63 | -4.13 | -4.50 | -0.06 | 7.13 |
| (8) Median | | 1.63 | -3.44 | -3.00 | 1.94 | 8.63 |
| (9) Q3 | | 2.75 | -2.75 | -2.06 | 3.57 | 9.88 |
| (10) Maximum | | 306.00 | -0.65 | 0.50 | 6.25 | 19.63 |
| (11) No. of obs > 0 | | 341 | 0 | 7 | 300 | 408 |
| (b) Quarterly | | | | | | |
| (1) Mean | 1.19 (0.18) | 2.87 (0.54) | -3.75 (0.23) | -3.82 (0.43) | 1.50 (0.43) | 9.57 (0.55) |
| (2) Standard deviation | 0.99 (0.10) | 3.53 (0.63) | 1.27 (0.15) | 2.38 (0.38) | 2.42 (0.22) | 2.98 (0.45) |
| (3) ρ_1 | 0.95 (0.02) | 0.89 (0.02) | 0.97 (0.01) | 0.98 (0.01) | 0.98 (0.01) | 0.98 (0.01) |
| (4) ρ_2 | 0.90 (0.03) | 0.77 (0.04) | 0.93 (0.02) | 0.96 (0.03) | 0.95 (0.02) | 0.97 (0.02) |
| (5) Correlation with U.S. rate | -0.11 (0.14) | 0.21 (0.25) | -0.77 (0.32) | -0.93 (0.55) | -0.79 (0.40) | 1.00 |
| (6) Minimum | -1.06 | -1.69 | -7.06 | -12.13 | -5.62 | 5.63 |
| (7) Q1 | 0.56 | 0.75 | -4.31 | -4.75 | -0.43 | 7.44 |
| (8) Median | 1.25 | 1.82 | -3.56 | -3.19 | 1.75 | 8.94 |
| (9) Q3 | 1.75 | 4.00 | -3.00 | -2.25 | 3.50 | 10.31 |
| (10) Maximum | 4.19 | 25.75 | -1.12 | 0.00 | 5.37 | 19.38 |
| (11) No. of obs > 0 | 344 | 346 | 0 | 0 | 281 | 408 |

Notes:

1. The sample includes 408 observations from 17 June 1981 to 5 April 1989; both quarterly and weekly rates are sampled weekly. Non-U.S. interest rates are expressed as an excess over the U.S. rate. Data are described in the text.

2. In rows (3) and (4), ρ_1 and ρ_2 are the first and second autocorrelation coefficients.

3. Asymptotic standard errors are in parentheses.

number of weeks in which the foreign rate is higher than the U.S. rate, reveals that the differential rarely changed sign during the sample period. With the exception of the French weekly rate, there is considerable serial correlation in the interest rate differentials [lines (2) and (3)]; nonetheless, computation of a statistic not reported in the table, $T(\hat{\rho}_1 - 1)$ [where $T = 408$ is the sample size and $\hat{\rho}_1$ is the first-order autocorrelation reported in row (2)], rejects the null of a unit root at the 5 percent levels in all four weekly differentials and in the Canadian and French quarterly differentials as well.

For some brief periods in the early part of the sample, French interest rates were rather high, at times extraordinarily so.⁵ These temporary spikes account for the large standard deviation [line (2)] and the relatively little serial correlation in French differential [lines (3), (4)]. Apart from France, the other interest rate differentials followed more stable patterns.

In computing our utility-based measure, we treat each currency in isolation, and produce nine sets of estimates, four for weekly and five for quarterly rates. Under our assumptions, a U.S. resident will invest a positive amount in a bond denominated in foreign currency only if the expected return on the bond denominated in foreign currency is higher than that on the dollar bond; since we also assume that the expected change in exchange rates is zero, this happens only if the foreign nominal return is higher. The converse is true for a foreign resident dividing her portfolio between bonds denominated in her own and in U.S. currency.

It is evident from line (11) of table 1 that there would be little grounds for comparing volatility estimates for German and Japanese exchange rates if our hypothetical investor were a U.S. resident: since weekly and quarterly deutschmark rates, and quarterly yen rates, were lower than dollar rates for every single week in the sample, a U.S. resident dividing her wealth between deutschmark or yen bonds on the one hand and dollar bonds on the other would never put any money in the former. We would have a similar though less dramatic problem for Canadian and French data if the investor were a resident of one of those two countries. So that our utility-based measure could use all 408 estimates of conditional variances, for each model and exchange rate we elected to make the hypothetical investor in a given week a U.S. resident if the U.S. interest rate is lower, and a foreign resident if the foreign interest rate is lower. For a given exchange rate, the fee that an

⁵On at least one occasion the high rate preceded an EMS realignment that depreciated the franc: the interest rate differential of 306 percent [table 1, part (a) column (2), line (10); corresponding weekly rate is about 2.7 percent], occurred on 16 March 1983, and the following week there was a realignment that depreciated the franc against the deutschmark by about 8 percent [Edison and Kaminsky (1991)]. This suggests that our assumption that the change in exchange rates is never predictable is a little extreme, at least in the first part of the sample; in the empirical work we therefore make sure that our results hold when we exclude the earlier parts of the sample.

Table 2
Models.

| (1) Model | (2) Formula for h_t | (3) Acronym |
|------------------------------|---|----------------|
| <i>Homoskedastic model</i> | | |
| 1. Homoskedastic | $h_t = \omega$ | homo |
| <i>GARCH models</i> | | |
| 2. GARCH(1, 1) | $h_t = \omega + \alpha e_t^2 + \beta h_{t-1}$ | (1, 1) |
| 3. IGARCH(1, 1) | $h_t = \alpha e_t^2 + (1 - \alpha)h_{t-1}$ | ig |
| <i>Autoregressive models</i> | | |
| 4. AR(12) in e_t^2 | $h_t = \omega + \sum_{i=1}^{12} \alpha_i e_{t-i+1}^2$ | e2AR |
| 5. AR(12) in $ e_t $ | $h_t = (\pi/2)(E_t e_{t+1})^2$; $E_t e_{t+1} = \omega + \sum_{i=1}^{12} \alpha_i e_{t-i+1} $ | e AR |
| <i>Non-parametric model</i> | | |
| 6. Gaussian kernel | $h_{t,j} = E(e_{t+j}^2 e_t)$; $\hat{h}_{ij} = \sum_{s=1}^{N-j} w_{iN,j} e_{s+j}^2$; $w_{iN,j} = c_{iN,j} / \sum_{s=1}^{N-j} c_{sN,j}$; $c_{iN,j} = \exp[-0.5(e_N - e_i)^2 / b^2]$; $b = \text{bandwidth defined in text}$ | nonp |

investor would pay to switch to the best model is then interpreted as the sum of the fees paid by investors in the two countries.

4.2. Models and estimation techniques

Column (1) of table 2 lists the models we estimated, and column (3) the acronyms used in some subsequent tables. Column (2) gives the formula for the one-period-ahead conditional variance, except for the non-parametric estimator for which the formula for the arbitrary j -period-ahead forecast is given. Since all the other models are linear, multiperiod forecasts can be obtained by the usual recursive prediction formulas.

The homoskedastic model [line (1)] simply set the conditional variance at all horizons equal to the sample mean of lagged e_t^2 's.

Two GARCH models were used [lines (2) and (3)]. Both were estimated by maximum likelihood assuming conditional normality, using analytical derivatives, with presample values of h and e^2 set to sample means. Lee and Hansen (1991) and Lumsdaine (1989) show that the conditional normality assumption is not necessary for the consistency and asymptotic normality of the estimators. We chose GARCH(1, 1) and IGARCH from a larger set of possible GARCH models after some preliminary in- and out-of-sample analysis suggested that these were the best GARCH models.

We also studied two autoregressive models, both of which were estimated

by OLS. One autoregression used e_t^2 [line (4)]. It is included because GARCH models imply ARMA processes for e_t^2 [see Bollerslev (1986)]; OLS estimation of such autoregressions therefore might perform comparably to more complicated GARCH estimation (although under the GARCH null, such OLS estimation is asymptotically inefficient). (In practice, this model occasionally produced negative point estimates of the conditional variance, in which case we used the homoskedastic estimate.) As in Schwert (1989a, b), the other autoregression used $|e_t|$ [line (4)]. Schwert suggests the factor of $(\pi/2)$ because the variance of a zero mean normally distributed random variables is $(\pi/2)$ times the square of the expected value of its absolute value. For both autoregressions, the lag length of 12 was chosen because for all countries in-sample results indicated that such a lag length was more than sufficient to produce a Q -statistic that implied white noise residuals.

Finally, we also tried a non-parametric estimator [line (6)]. It can be interpreted as working off the basic definition $E(e_{t+j}^2 | e_t) = \int_0^\infty e_{t+j}^2 f(e_{t+j}^2 | e_t) de_{t+j}^2$, where $f(e_{t+j}^2 | e_t)$ is the density of e_{t+j}^2 conditional on e_t . See Pagan and Ullah (1990a, b) for an excellent exposition. As in Pagan and Schwert (1990) we used a Gaussian kernel, defined in column (2), with the bandwidth $b = \hat{\sigma}(N-j)^{-1/5}$ and $\hat{\sigma}$ the sample standard deviation of e_t , $t=1, \dots, N-j$. We did not try any other kernel. We did a little experimentation with some alternative fixed bandwidths and information sets, comparing out-of-sample mean squared errors, but found that these yielded similar results.

5. Empirical results

For our utility-based measure we report in detail results with a CRRA of one [i.e. $\gamma W/(1-\gamma W)=1$, in the notation of section 3]; below we summarize results with a CRRA of 10. Table 3 has estimates and asymptotic standard errors of (5), with Eurodeposits of one week maturity in part (a), one quarter maturity in part (b).

One's eyes are drawn to the '0.0' entries for GARCH(1, 1), which appear for five of the nine rows. IGARCH yielded the highest average utility in two other data sets (Germany, both horizons), and was second to GARCH(1, 1) in four others. The non-parametric model was best for France (weekly) and Canada (quarterly), but otherwise did not perform very impressively. The remaining three models generally did poorly. Note that the fine performance of the GARCH models as a class is not an artifact of the presence of two such models: had we not estimated IGARCH models, GARCH(1, 1) would have been best in 6 rather than 5 datasets; had we not estimated GARCH(1, 1) models, IGARCH would have been best in 6 rather than 2 datasets.

The statistical significance of differences across models is weak, however.

Table 3
Wealth sacrifice for right to highest utility model.

| Country | Model | | | | | | |
|------------------------------|--------------------|------------------|-------------------|-------------------|-------------------|-------------------|---------------------|
| | homo | (1, 1) | ig | e2AR | e AR | nonp | $\chi^2(5)$ |
| <i>(a) Weekly horizon</i> | | | | | | | |
| France | 862.6 (973.0) | 78.7 (267.0) | 144.5 (394.8) | 245.2 (441.7) | 108.5 (187.4) | 0.0 | 3.174 [0.673] |
| Germany | 42.5 (31.5) | 47.1 (38.7) | 0.0 | 319.1* (177.1) | 208.9** (76.5) | 306.7 (230.7) | 12.246** [0.032] |
| Japan | 37.3 (171.7) | 0.0 | 40.2 (39.1) | 62.6 (181.8) | 169.0* (90.3) | 687.7 (473.3) | 9.912* [0.078] |
| U.K. | 204.7** (102.9) | 0.0 | 40.2 (30.9) | 412.3* (229.1) | 234.6 (153.2) | 482.2 (366.0) | 8.517 [0.130] |
| <i>(b) Quarterly horizon</i> | | | | | | | |
| Canada | 1.7 (9.0) | 56.8** (28.8) | 1031.3 (736.8) | 28.9** (11.1) | 94.4 (62.2) | 0.0 | 9.480* [0.091] |
| France | 181.3 (116.0) | 0.0 | 8.4 (20.9) | 148.5 (95.1) | 210.4* (120.4) | 180.2* (102.6) | 7.893 [0.162] |
| Germany | 29.4 (46.6) | 4.9 (25.1) | 0.0 | 29.7 (46.4) | 48.5 (48.9) | 44.7 (52.7) | 5.128 [0.400] |
| Japan | 120.1* (71.9) | 0.0 | 25.2 (34.8) | 40.8 (34.6) | 75.1 (51.4) | 153.6* (80.4) | 7.838 [0.165] |
| U.K. | 33.0 (71.3) | 0.0 | 81.8* (46.8) | 0.1 (29.3) | 8.5 (22.8) | 32.5 (68.7) | 5.643 [0.343] |

Notes:

1. An investor is assumed to divide her wealth between Eurodeposits in dollars and those in the currency of the indicated country, 17 June 1981–5 April 1989. Each row reports estimates of (5), the wealth that the investor would give up to use the model that yielded the highest average utility (the model with the '0.0' entry). The units are annual basis points. Smaller numbers mean better performance. Table 2 describes the acronyms for the models. Relative risk aversion is set to 1.

2. Asymptotic standard errors are in parentheses; ** indicates significance at 10 percent level, *** at 5 percent level (two-tailed test).

3. The $\chi^2(5)$ column reports a test of the null that all five non-zero entries in a given row are equal to zero, with asymptotic p -values in parentheses.

Only five of the twenty entries in part (a) of table 3, and seven of the 25 entries in part (b) of table 3 are significantly different from zero at the 10 percent level (two-tailed test). As indicated in the last column of each panel, the nine tests for the equality of utility levels across all models is rejected at the 5 percent level once and at the 10 percent but not 5 percent level twice.

The economic significance of differences across models appears to be more pronounced. In the weekly data in table 3, part (a), the four non-GARCH models had three-digit estimates more often than not, indicating that an investor would be willing to give up more than 100 basis points of her wealth, annually, to switch from using one of these models to the optimal one. At the longer horizon, performance is more similar across models: the median figure in part (b) is 45, in part (a) is 187. This is consistent with the

well-known fact that conditional heteroskedasticity in exchange rates tends to die out rapidly [Diebold (1988)].

One way to gauge these figures is to interpret them as a fee that a professional money manager could charge an investor not capable of estimating GARCH models. As such, these figures seem to be above what Wall Street mutual funds typically charge for their services [Ippolito (1989), *New York Times*, 14 May 1991, p. F14], which suggests to us that they are substantial.

Table 4 summarizes some experiments we performed to see whether these results are sensitive to the sample used and to the choice of relative risk aversion. In table 4, specification A is the one used in table 3, and is repeated for convenience. Specification B recalculated the entries in table 3 using each of the two halves of the sample rather than the whole sample, and specification C recalculated using each of the four quarters of the sample. Specification D recalculated using the whole sample, and a higher assumed level of risk aversion.

As one can see in column (2) of part (b), GARCH models tended to perform best in all these additional experiments. Columns (3) and (4) indicate that the statistical significance of differences between models about as strong as was suggested by table 3. Column (1) indicates that the estimates of the wealth one would sacrifice to use the best model are a little lower in the later parts of the sample, and that increased risk aversion (specification D), which, as we noted above, will lead to a narrowing of differences across models, happens to do so rather sharply. We therefore slightly amend our summary of results to state that our estimates imply that an investor would be willing to give up 5 to 200 basis points of her wealth, annually, to switch to GARCH from another model; even the lower bound of this range strikes us as substantial.

How do these results compare with those of the usual mean squared error criterion? Table 5 presents rankings by this criterion, for a one week horizon. (The mean squared errors underlying the rankings are available on request.) While in each country either GARCH(1, 1) or IGARCH has the lowest mean squared error, the GARCH(1, 1) model overall does not perform as well as it did by the utility-based criterion (see the entries for France and Germany). Moreover, the $\chi^2(5)$ statistics in the next to last column suggest that there is little to recommend one model over another, in the sense that for no country can one reject the null that all six mean squared errors are the same at conventional significance levels. Even more striking is that there is precious little evidence that whichever GARCH model had the lowest mean squared error is substantially better than the homoskedastic model. The last column indicates that one cannot reject the null that the mean squared error for the homoskedastic model is the same as the best GARCH model at anything close to conventional significance levels. We conclude that the mean squared

Table 4
Effects of alternative specifications.

| Sample period | | CRRA | Description | |
|-----------------------------------|--|--|--|--|
| (a) Description of specifications | | | | |
| A | 17 June 1981 to 5 April 1989 | 1 | Table 3 specification | |
| B1 | 17 June 1981 to 8 May 1985 | 1 | First half of table 3 sample | |
| B2 | 15 May 1985 to 5 April 1989 | 1 | Last half of table 3 sample | |
| C1 | 17 June 1981 to 25 May 1983 | 1 | First quarter of table 3 sample | |
| C2 | 1 June 1983 to 8 May 1985 | 1 | Second quarter of table 3 sample | |
| C3 | 15 May 1985 to 22 April 1987 | 1 | Third quarter of table 3 sample | |
| C4 | 29 April 1987 to 5 April 1989 | 1 | Fourth quarter of table 3 sample | |
| D | 17 June 1981 to 5 April 1989 | 10 | Table 3 sample, with higher CRRA | |
| | (1) | (2) | (3) | (4) |
| | Median estimate of (5), wealth sacrifice | Number of countries for which best model is: | Number of <i>t</i> -statistics significant at: | Number of $\chi^2(5)$ statistics significant at: |
| | | (1, 1) ig | 0.10 0.05 | 0.10 0.05 |
| (b) Summary of empirical results | | | | |
| <i>Weekly:</i> | | | | |
| A | 186.9 | 2 1 | 5 2 | 2 1 |
| B1 | 297.0 | 2 1 | 4 1 | 2 0 |
| B2 | 54.5 | 1 2 | 2 0 | 0 0 |
| C1 | 321.4 | 2 1 | 3 1 | 1 1 |
| C2 | 74.2 | 1 1 | 3 2 | 2 2 |
| C3 | 63.7 | 0 2 | 5 2 | 0 0 |
| C4 | 10.0 | 1 3 | 4 2 | 2 0 |
| D | 4.6 | 2 1 | 5 2 | 2 0 |
| <i>Quarterly:</i> | | | | |
| A | 44.7 | 3 1 | 7 2 | 1 0 |
| B1 | 100.4 | 2 2 | 6 1 | 2 1 |
| B2 | 23.2 | 1 0 | 6 3 | 2 2 |
| C1 | 55.4 | 1 3 | 5 2 | 2 1 |
| C2 | 49.0 | 1 3 | 6 1 | 3 2 |
| C3 | 53.4 | 1 1 | 8 6 | 3 3 |
| C4 | 17.4 | 0 2 | 11 5 | 3 2 |
| D | 0.7 | 2 1 | 5 1 | 1 1 |

Notes:

1. Specification A is the one reported in detail in table 3, and is repeated here for convenience.
2. Part B is based on estimates for the 4 (weekly) or 5 (quarterly) countries and 6 models listed in table 3 part (a). Since, for a given country, the estimates of the best model (the '0.0' model) do not figure into the computation of the number of the part (b) values, the total number of values underlying each weekly row is 20 for columns (1) and (3), 4 for columns (2) and (4); the corresponding quarterly figures are 25 and 5.

Table 5
 Rankings by out-of-sample mean squared error, weekly horizon.

| Country | Model | | | | | | $\chi^2(5)$ | $\chi^2(1)$ |
|---------|-------|--------|----|------|------|------|------------------|------------------|
| | homo | (1, 1) | ig | e2AR | e AR | nonp | | |
| Canada | 5 | 1 | 3 | 4 | 2 | 6 | 7.244 [0.203] | 1.243 [0.265] |
| France | 2 | 6 | 1 | 5 | 3 | 4 | 8.911 [0.113] | 0.011 [0.918] |
| Germany | 2 | 5 | 1 | 6 | 4 | 3 | 8.147 [0.148] | 0.012 [0.912] |
| Japan | 3 | 1 | 2 | 5 | 4 | 6 | 6.414 [0.268] | 0.770 [0.380] |
| U.K. | 4 | 2 | 1 | 5 | 3 | 6 | 3.779 [0.582] | 1.521 [0.217] |

Notes:

1. In each row, '1' indicates best (smallest) mean squared error for the indicated country, '2' second best, ..., '6' worst.

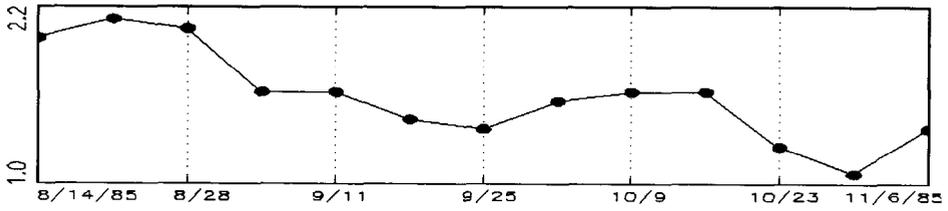
2. The $\chi^2(5)$ column reports a test of the null of the equality of the six mean squared errors underlying the ranking in a given row, with asymptotic p -value in brackets.

3. The $\chi^2(1)$ column reports a test of the null of the equality of mean squared error for the homoskedastic and best model (either GARCH(1, 1), or IGARCH, as indicated), with asymptotic p -value in brackets.

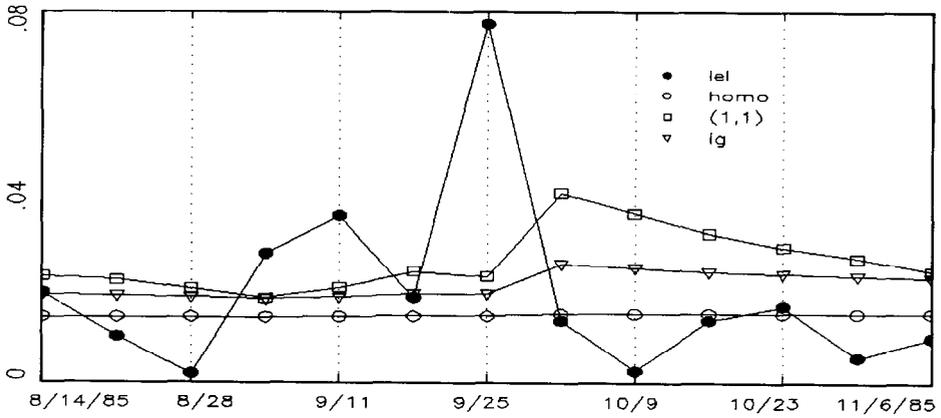
error criterion also favors GARCH as a class, but not as sharply as does our utility-based criterion.

We close this section with a closer look at a particular period, which suggests that it is the asymmetry in our utility-based criterion that accounts for the differences between mean squared error and utility rankings. A comparison of tables 3 and 5 indicated to us that a detailed examination of French data might be revealing in this connection, because for such data GARCH(1,1) does poorly by the mean squared error, well by the utility based criterion. Consider the 13 weeks from 14 August 1985 to 6 November 1985. The length of this interval was chosen because the relevant figures could be graphed clearly; the dating of this interval was chosen because it is centered around the Plaza Accord, which was announced on 22 September 1985, and which caused the largest weekly change in the dollar/franc exchange rate in our sample (7.7 percent).

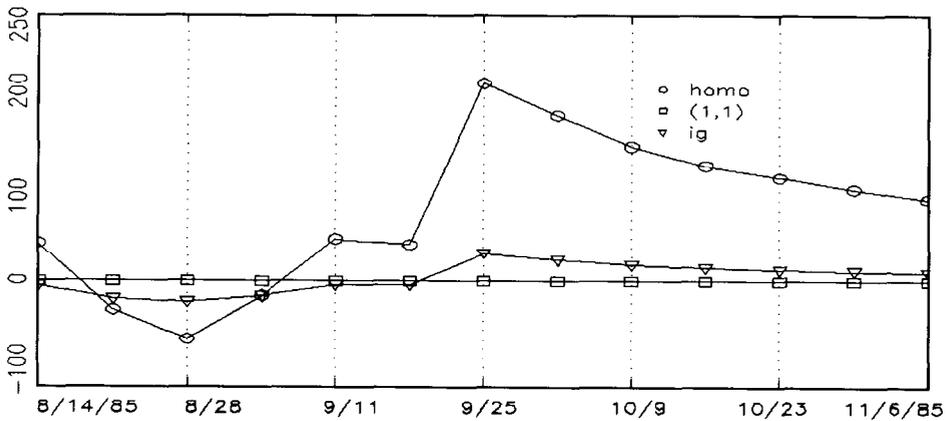
Fig. 2(a) plots the annualized interest rate differential, which we present simply to reassure the reader that the estimates of our utility-based measure that we are about to present are not based on unusual interest rates. Fig. 2(b) plots the absolute value of the exchange rate together with the square root of the corresponding conditional variance for the GARCH(1,1), IGARCH, and homoskedastic model. Only three models, and square roots rather than squares, were plotted to make the figure more legible. Fig. 2(c) plots the evolution over time of estimates of the wealth an investor would



(a)



(b)



(c)

Fig. 2. \hat{h}_{m-1} and wealth sacrifice for France, weekly. (a) Interest rate differential (percent). (b) $|e_t|$ and square root of \hat{h}_{m-1} . (c) Wealth sacrifice for right to use (1,1) model.

Table 6
Results for weekly horizon, France, 14 August 1985 to 6 November 1985.

| | Model | | | | | |
|---|------------|-----|------|------|------|-------|
| | homo (1,1) | ig | e2AR | e AR | nonp | |
| Estimates of (5), wealth sacrifice to use highest utility model | 77.6 | 0.0 | 8.4 | 48.2 | 16.5 | 105.5 |
| Rankings by out of sample mean squared error | 3 | 6 | 1 | 5 | 2 | 4 |

Notes:

1. For interpretation of the estimates of (5), see the notes to table 3.
2. For interpretation of the rankings by mean squared error, see the notes to table 4.

sacrifice to use GARCH(1,1); the first estimate, for 14 August 1985, is based on one observation, the final estimate, for 6 November 1985, is based on 13.

In the first three weeks of this period, fig. 2(b) suggests that GARCH(1,1) did a poorer job of fitting the realized square of the exchange rate than did the other two models, and fig. 2(c) bears out this impression. During the next four weeks, from 3 to 25 September, it is hard to tell from fig. 2(b) which models are tracking the exchange rate best. But fig. 2(c) indicates that by 25 September, the GARCH(1,1) model delivered the highest average utility, a ranking that was maintained until the end of the 13-week period that is graphed. In comparing figs. 2(b) and 2(c), what is particularly striking is (1) the degeneration of the homoskedastic relative to the GARCH(1,1) model during the week ending 25 September (the week of the Plaza Accord), and (2) the continued domination of GARCH(1,1) after 25 September despite its substantial overestimates of the conditional variance.⁶ This illustrates the asymmetry of the utility-based criterion: it may be seen in fig. 2(b) that the homoskedastic model underestimated only slightly relative to the GARCH(1,1) model for the week ending 25 September, and that the GARCH(1,1) model overestimated dramatically relative to the homoskedastic model in some subsequent weeks. But the underestimate has a much stronger effect on utility than do the overestimates.

Fig. 2(b) suggests to the eye that the GARCH(1,1) model does poorly by a mean squared error criterion. This impression is borne out by a formal calculation. Table 6 contains wealth sacrifices and rankings by mean squared errors for all six models, for this 13-week period. GARCH(1,1) was the best by the utility-based criterion, worst by the mean squared error criterion.

⁶Here, we are identifying the square of the exchange rate with the population conditional variance, although these in fact differ by a zero mean expectational error; note that the fact that the sample contains only 13 observations means that this expectational error may contribute substantially to our estimates of average utility.

6. Conclusions

We conclude with some suggestions for future research. An obvious possibility is to see if other models, such as those surveyed in Bollerslev et al. (1990), dominate GARCH. Another is to apply our analysis to a portfolio of assets that is better diversified, such as one that includes equities. A third is to permit flexible use of a variety of models by allowing for weighted combinations of fitted conditional means and variances and/or implied fractions, possibly with time varying weights. Finally, it would be very desirable to compare volatility models in an environment of dynamic rather than static utility maximization.

Appendix

Average utility is estimated as follows. For a one-period (one week) horizon, use models $m=1, \dots, M$ to solve

$$\begin{aligned} & \max_{\{f_{mt}\}} [E_t U_{m+1} \equiv E_t(W_{m+1} - 0.5\gamma W_{m+1}^2) \mid \text{model } m \text{ used to} \\ & \text{compute } E_t e_{t+1}^2] \\ & \text{s.t. } W_{m+1} = W_t [f_{mt}(R_{t+1}^* + e_{t+1}) + (1 - f_{mt})R_{t+1}]. \end{aligned}$$

Assume that the interest rate differential is uncorrelated with the change in exchange rates, $E_t(R_{t+1}^* - R_{t+1})e_{t+1} = 0$. Let $\mu_{t+1} \equiv R_{t+1}^* - R_{t+1} > 0$, $W \equiv W_t$, and assume $1 - \gamma WR_{t+1} > 0$ (otherwise the investor can reach satiation with certainty). Elementary calculus yields

$$\begin{aligned} f_{mt} &= [\mu_{t+1}(1 - \gamma WR_{t+1})/(\gamma W)] [1/(\mu_{t+1}^2 + \hat{h}_{mt})], \\ &\Rightarrow E_t U_{m+1} = [c_t + d_t u(h_t, \hat{h}_{mt})] W, \\ c_t &\equiv (R_{t+1} - 0.5\gamma WR_{t+1}), \\ d_t &\equiv (\gamma W)^{-1} \mu_{t+1}^2 (1 - \gamma WR_{t+1})^2, \\ u(h_t, \hat{h}_{mt}) &\equiv [(\mu_{t+1}^2 + \hat{h}_{mt})^{-1} - 0.5(\mu_{t+1}^2 + \hat{h}_{mt})^{-2}(\mu_{t+1}^2 + h_t)]. \end{aligned} \tag{A.1}$$

Let $\delta \equiv \gamma W / (1 - \gamma W)$ be the coefficient of relative risk aversion. Substitute the ex-post realized exchange rate square e_{t+1}^2 for its conditional expectation h_t and average over many time periods to get average utility,

$$\begin{aligned} \bar{U}_m &= [\bar{c} + \bar{u}_m] W, \\ \bar{c} &\equiv (T - N + 1)^{-1} \sum_{t=N}^T c_t, \\ \bar{u}_m &\equiv (T - N + 1)^{-1} \sum_{t=N}^T d_t u(e_{t+1}^2, \hat{h}_{mt}). \end{aligned} \tag{A.2}$$

Suppose that model 1 turns out to be the best. Let m be an arbitrary suboptimal model. We see from (A.2) that when model 1 with wealth $W - \Delta W$ yields average utility equivalent to model m with wealth W , ΔW satisfies $(\bar{c} + \bar{u}_1)(W - \Delta W) = (\bar{c} + \bar{u}_m)W$. The corresponding fraction of wealth is $\Delta W/W = [(\bar{c} + \bar{u}_1) - (\bar{c} + \bar{u}_m)]/(\bar{c} + \bar{u}_1) = [1 - (\bar{U}_m/\bar{U}_1)]$. As indicated in eq. (5), we express this ratio in basis points.

For a 13-week horizon: in (A.2) replace e_{t+1}^2 with $(e_{t+1} + \dots + e_{t+13})^2$ and \hat{h}_{mt} with $\hat{h}_{mt,1} + \dots + \hat{h}_{mt,13}$. The implicit timing assumption is that investors are using weekly data to make investment decisions every quarter (every 13 weeks). One-thirteenth are investing the first week in the quarters, ..., one-thirteenth the last week of the quarter. The figure for average utility that we compute is the average of average utility for each of the 13 groups of investors.

References

- Akgiray, Vedat, 1989, Conditional heteroskedasticity in time series of stock returns: Evidence and forecasts, *Journal of Business* 62, 55–80.
- Bollerslev, Tim, 1986, Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 31, 307–327.
- Bollerslev, Tim, Ray Y. Chou, Narayanan Jayaraman and Kenneth F. Kroner, 1990, ARCH modelling in finance: A selective review of the theory and empirical evidence, with suggestions for future research, manuscript, Northwestern University.
- Breen, William, Lawrence R. Glosten and Ravi Jagannathan, 1989, Economic significance of predictable variations in stock index returns, *Journal of Finance* 44, 1177–1191.
- Diebold, Francis X., 1988, Empirical modelling of exchange rate dynamics (Springer-Verlag, Berlin).
- Diebold, Francis X. and James Nason, 1990, Nonparametric exchange rate prediction?, *Journal of International Economics* 28, 315–332.
- Edison, Hali J. and Garciela Laura Kaminsky, 1991, Target zones, intervention and exchange rate volatility, France, 1979–1990, manuscript, Federal Reserve Board of Governors.
- Engle, Robert F., 1982, Autoregressive conditional heteroskedasticity, with estimates of the variance of United Kingdom inflation, *Econometrica* 50, 987–1007.
- Engle, Robert F. and Tim Bollerslev, 1986, Modelling the persistence of conditional variances, *Econometric Reviews* 5, 1–50.
- Engle, Robert F., Che-Hsiung Hong and Alex Kane, 1990, Valuation of variance forecasts with simulated options markets, manuscript, UCSD.
- Eun, Cheol S. and Bruce G. Resnick, 1984, Estimating the correlation structure of international share prices, *Journal of Finance* 39, 1311–1324.
- Friedman, Benjamin M. and Kenneth N. Kuttner, 1988, Time-varying risk perceptions and the pricing of risky assets, NBER Working Paper No. 2694.
- Hansen, Lars Peter, 1982, Asymptotic properties of generalized method of moments estimators, *Econometrica* 50, 1029–1054.
- Ippolito, Richard A., 1989, Efficiency with costly information: A study of mutual fund performance, *Quarterly Journal of Economics* 1–24.
- Lee, Sang-Won and Bruce E. Hansen, 1991, Asymptotic properties of the maximum likelihood estimator and test of the stability of parameters of the GARCH and IGARCH models, manuscript.
- Lumsdaine, Robin L., 1989, Asymptotic properties of the maximum likelihood estimator in GARCH(1, 1) and IGARCH(1, 1) models, manuscript, Harvard University.
- McCulloch, Robert and Peter E. Rossi, 1990, Posterior, predictive and utility-based approaches to testing the arbitrage pricing theory, *Journal of Financial Economics* 28, 7–38.

- Meese, Richard A. and Kenneth S. Rogoff, 1983, Empirical exchange rate models of the seventies: Do they fit out of sample?, *Journal of International Economics* 14, 3–24.
- Meese, Richard A. and Andrew K. Rose, 1991, An empirical assessment of non-linearities in models of exchange rate determination, *Review of Economic Studies* 58, 603–618.
- Newey, Whitney K. and Kenneth D. West, 1987, A simple, positive semidefinite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.
- Pagan, Adrian R. and G. William Schwert, 1990, Alternative models for conditional stock volatility, *Journal of Econometrics* 45, 267–290.
- Pagan, Adrian R. and Aman Ullah, 1990a, Chapter 3: Methods of density estimation, manuscript, University of Rochester.
- Pagan, Adrian R. and Aman Ullah, 1990b, Chapter 4: Non-parametric estimation of conditional moments, manuscript, University of Rochester.
- Schwert, G. William, 1989a, Business cycles, financial crises and stock volatility, in: K. Brunner and A. H. Meltzer, eds., *IMF policy advice, market volatility, commodity price rules and other essays*, Carnegie Rochester Series on Public Policy No. 31 (Elsevier Science Publishers, Amsterdam) 983–126.
- Schwert, G. William, 1989b, Why does stock market volatility change over time?, *Journal of Finance* 44, 1115–1154.
- West, Kenneth D. and Dongchul Cho, 1992, The predictive ability of several models of exchange rate volatility, manuscript, University of Wisconsin.