Is ARCH Useful in High Frequency Foreign Exchange Applications?

Brad Jones

Applied Finance Centre
Macquarie University
North Ryde, NSW, 2109, Australia
Telephone: +61 2 9181 1197
Email: brad.jones@mafc.mq.edu.au

Abstract

One of the many challenges posed by the study of high frequency financial market data is to develop models capable of explaining asset price behaviour at a range of frequencies. At the same time as presenting researchers with new opportunities, it also calls into question whether standard time series models are useful in high frequency applications. This paper addresses this issue from two perspectives. First, a Monte Carlo procedure is used to investigate whether the unconditional distribution of high frequency foreign exchange returns can be approximated by the unconditional distribution of returns simulated by a range of popular stochastic processes. Second, high frequency data is used to generate and appraise forecasts of daily variance. Forecasts are evaluated using statistical criteria as well as a profitability measure based on a trading game in a pseudo options market.

The simulation exercise demonstrates that the autoregressive conditional heteroskedasticity (ARCH) family of models is unable to reproduce the unconditional distribution of foreign exchange returns at frequencies higher than 24 hours. This is largely a legacy of the heavy-tailed feature of intraday returns. However, results from the forecasting analysis extend those in Andersen and Bollerslev (1998) by showing that a range of standard volatility models can in fact produce accurate forecasts of realized daily variance. In other words, it is possible for ARCH models to predict variability in the conditional second moment of daily foreign exchange returns. This is attributed to the use of frequently sampled data in the construction of estimates of realized variance (against which forecasts are measured). In addition, the inclusion of the sum of squared intraday returns in the Generalized ARCH(1,1) model yields improvements in the modelling, and most notably forecasting, of realized daily variance. This appears to be an artifact of the noise inherent in using the daily squared return as an estimator of realized daily variance.

This paper demonstrates that whilst standard econometric models do not capture the intraday foreign exchange return generating process, this should not immediately preclude these models from high frequency applications. Instead, the forecasting exercise demonstrates practical benefits are easily attainable from using high frequency data to develop and evaluate existing asset pricing models.

1. Introduction

The application of misspecified asset pricing models has the potential to induce the mispricing of financial assets, in turn leading to serious implications for portfolio selection and risk management. As such, exercises attempting to establish an accurate representation of asset price movements are non-trivial. Indeed, the past four decades have spawned a voluminous literature attempting to develop models consistent with the behaviour of speculative prices. Arguably the most celebrated of these is Engle's (1982) autoregressive conditional heteroskedasticity (ARCH) model, designed primarily to capture the positive serial correlation in financial market volatility[1].

In recent times, the asset pricing literature has become increasingly interested in using high frequency data to address a range of issues in financial markets. This interest has been fueled by the development of intraday databases spanning a host of financial instruments and markets[2]. Much of the high frequency analysis has centered on the largest market of all – the foreign exchange market[3]. This research has resulted in a richer understanding of the intraday foreign exchange return generating process (RGP), revealing some behaviours that were not observed at lower frequencies. Discoveries of this nature have stimulated the development of a new breed of models designed to account for characteristics observed only in intraday markets[4].

At the same time as presenting researchers with new opportunities, these findings have also called into question whether established time series models are useful in high frequency applications. For instance, the ARCH family of models were originally developed to capture features of financial time series measured at daily and lower frequencies. Recent studies have argued that standard ARCH processes are unable to replicate the autocorrelation structure of intraday (squared and absolute) returns[5]. In addition, ARCH type models provide seemingly poor forecasts of daily volatility when judged by standard forecast evaluation criteria (see Andersen and Bollerslev, 1998). Predictably, these findings have led to the perception that ARCH models are of limited practical use in high frequency studies. The primary contribution of the current exercise is to address this issue from two perspectives.

---

[1] Bollerslev, Chou and Kroner (1992) and Bera and Higgins (1993) review a number of univariate and multivariate applications of ARCH models involving equity, debt, foreign exchange and derivative pricing.
[2] Sources of intraday data sets are discussed by Goodhart and O'Hara (1997) and Bollerslev and Zhou (2001). Goodhart and O'Hara (1997) also provide an excellent survey of the applications and issues associated with the use of high frequency financial market data.
[3] Note that the operations of intraday foreign exchange traders account for more than ninety percent of foreign exchange market volume (Dacorogna, Gencay, Müller, Olsen and Pictet, 2001).
[4] For instance, the Heterogeneous ARCH model of Müller, Dacorogna, Dave, Olsen, Pictet and von Weizsacker (1997) and the Fractionally Integrated Generalized ARCH model of Baillie, Bollerslev and Mikkelsen (1996) have been designed to capture the long memory of volatility observed in intraday returns.
[5] For instance, Dacorogna et al. (2001) and Zumbach (2002) argue that the Generalized ARCH(1,1) model assumes an exponential decay in volatility autocorrelation, whereas the autocorrelation of intraday foreign exchange volatility typically decays at a hyperbolic rate.

First, a Monte Carlo procedure is used to test whether high frequency foreign exchange RGP's can be approximated by the distributional characteristics of popular stochastic processes. The simulation exercise uses parameter values obtained by fitting these models to foreign exchange returns measured at different intraday frequencies. Using a simple test, the unconditional distribution of returns simulated by a range of ARCH models is compared to the unconditional distribution of high frequency foreign exchange returns.

Second, one-step-ahead daily volatility forecasts generated by the ARCH class of models are assessed. A statistical evaluation of model forecasts is supplemented with a profitability measure based on trading option contracts. Following Engle, Hong, Kane and Noh (1993) and Maheu and McCurdy (2001), model estimates of ex-ante volatility are used to price European-style options combined into a straddle position. A trading game in a pseudo options market is devised, where the criterion for success is profits associated with trading foreign exchange straddle contracts. In the forecasting exercise, realized daily volatility (referred to as realized daily variance henceforth) is defined as the summation of intraday squared returns. The paper demonstrates that this measure of realized variance substantially reduces the noise that plagues the use of the daily squared return for the same purpose.

The main findings of the paper are as follows. The simulation exercise demonstrates that the ARCH family of models is unable to reproduce the unconditional distribution of foreign exchange returns at frequencies higher than 24 hours. This is largely a legacy of the heavy-tailed feature of intraday returns. However, results from the forecasting analysis extend those in Andersen and Bollerslev (1998) by showing that a range of standard volatility models can in fact produce accurate forecasts of realized daily variance. In other words, it is possible for ARCH models to predict variability in the conditional second moment of daily foreign exchange returns. This is attributed to the use of frequently sampled data in the construction of estimates of realized variance (against which forecasts are measured). In addition, the inclusion of the sum of squared intraday returns in the Generalized ARCH(1,1) model yields improvements in the modelling, and most notably forecasting, of realized daily variance. This appears to be an artifact of the noise inherent in using the daily squared return as an estimator of realized daily variance.

This paper demonstrates that whilst established econometric models do not capture the intraday foreign exchange RGP, this should not immediately preclude these models from high frequency applications. Instead, the forecasting exercise demonstrates practical benefits are readily attainable from using high frequency data to develop and evaluate existing asset pricing models.

The remainder of the paper is organized as follows. The features of the data are presented in section 2. Section 3 presents the models that constitute high frequency foreign exchange RGP's under the null hypothesis. Section 4 presents the simulation exercise investigating whether these models are able to reproduce the moments of intraday and daily exchange rate returns. Model forecasts of realized daily variance are evaluated in section 5. Conclusions are drawn and suggestions for future research offered in section 6.

2. Data

The data set consists of spot foreign exchange prices for the following four currencies expressed against the US Dollar: the Australian Dollar (USD/AUD), the British Pound (USD/GBP), the Deutschemark (DEM/USD) and the Japanese Yen (JPY/USD)[6]. The spot rates were obtained from Olsen and Associates, a Zurich based institute specializing in the collection and analysis of high frequency foreign exchange data. The sample consists of continuously recorded five-minute bid and ask prices from January 1, 1997, 00:00 GMT through December 31, 1998, 00:00 GMT, for a total of 210,240 observations. Five-minute prices were defined as the midpoint of the bid and ask. Prices measured at thirty-minute, hourly, eight-hourly and daily frequencies were obtained by sampling from this initial grid of five-minute prices (ie. hourly prices were obtained by observing the midpoint of the bid and ask on the stroke of each hour). Continuously compounded returns were formed by taking the first difference of logarithmic prices.

Whilst it is possible to trade currencies 24 hours per day, 7 days a week, there are clearly periods such as weekends and holidays during which there will be very low trading activity, a phenomenon that has the potential to generate seasonal effects. As such, weekend prices covering the period from Friday 21:05 GMT to Sunday 21:00 GMT were removed from the sample following Maheu and McCurdy (2001) among others. Quiet trading days falling on fixed holidays (December 24-26,31, January 1-2) as well as moving holidays (Good Friday, Easter Monday, Memorial Day, July Fourth, Labor Day and Thanksgiving) were also eliminated following Maheu and McCurdy (2001). In sum, deseasonalizing by removing days with abnormally low trading activity resulted in a final sample of 144,864 five-minute returns spanning 503 days.

3. Possible Representations of High Frequency Exchange Rate Behaviour

This section presents the models consistent with high frequency foreign exchange RGP's under the null hypothesis. Prior to simulating unconditional return distributions for these processes, parameter values must first be determined. This is achieved by estimating the parameters (using Maximum Likelihood) from models fitted to foreign exchange returns measured at daily and intraday frequencies.

Many statistical processes proposed in the asset pricing literature, and consequently those investigated in this study, are assumed to follow the general formula below describing equally spaced continuously compounded returns r_t,

$$r_t = \ln p_t - \ln p_{t-1} = \mu + \sigma_t \cdot \varepsilon_t \tag{1}$$

[6] The DEM/USD was the most actively traded foreign currency during the sample. The JPY remains the second most heavily traded currency against the USD, the GBP the third, and the AUD the ninth; see Dacorogna et al. (2001) for average daily tick volumes. During this period, the DEM/USD and JPY/USD constituted the main axes of the international financial system, and thus spanned the majority of the systematic currency risk faced by most large institutional investors and international corporations.

where ε_t is an identically and independently distributed (i.i.d.) random variable with zero mean and unit variance, σ_t is the square root of the variance of the return, μ is the constant drift and $\ln p_t$ denotes the log price, with all variables defined at time t (besides the time-invariant drift parameter, μ). To model the conditional variance of high frequency exchange rate returns, the following parametric specifications are employed[7].

3.1 Generalized ARCH (GARCH)

Many studies have found Bollerslev's (1986) GARCH(1,1) model provides a reasonable first approximation to the temporal dependencies observed in financial asset returns[8]. As such, the GARCH(1,1) formulation is the first process considered as a possible representation of high frequency exchange rate behaviour. This model specifies the conditional variance of the current period return as a function of the conditional variance of the last period's return, σ_{t-1}^2, updated by the news revealed by last period's return ε_{t-1}^2,

$$\sigma_t^2 = \gamma + \alpha \cdot \varepsilon_{t-1}^2 + \beta \cdot \sigma_{t-1}^2 \tag{2}$$

where the stationarity condition imposes $\alpha + \beta < 1$, and to ensure a well-defined process, $\gamma > 0$, $\alpha \geq 0$, and $\beta \geq 0$. Some authors have suggested the assumption of normally distributed error terms in the GARCH(1,1) model may not be sufficient to capture the excess kurtosis commonly observed in high frequency returns (see McCurdy and Morgan, 1988, Hsieh, 1989, and Baillie and Bollerslev, 1989). Consequently, this model is estimated assuming the error term conforms to either a normal or Student-*t* distribution.

3.2 GARCH-I: GARCH(1,1) extended with Intraday Information

Andersen and Bollerslev (1998) show that whilst the daily squared return is an unbiased estimator of ex-post daily variance, it is also an extremely noisy one. In response to this, Martens (2001) augments the standard GARCH(1,1) model with a variable, I_{t-1}, comprising the sum of thirty-minute squared returns (this is referred to as the GARCH-I model henceforth). The conditional variance is defined by,

$$\sigma_t^2 = \gamma + \alpha \cdot \varepsilon_{t-1}^2 + \beta \cdot \sigma_{t-1}^2 + \kappa \cdot I_{t-1} \tag{3}$$

where the stationarity condition imposes $\alpha + \beta + \kappa < 1$. To ensure the process is well-defined, $\gamma > 0$, $\alpha \geq 0$, $\beta \geq 0$ and $\kappa \geq 0$. The intuition behind this model is simple: if the standard GARCH(1,1) model is fitted to daily returns, and on the previous trading day the return was zero but prices fluctuated heavily during the day, the lagged squared return (equal to zero) provides misleading information. By extending the standard GARCH model with intraday returns, this specification is able to capture the information that the previous trading day was actually quite volatile[9]. In this case, the I_{t-1} variable would be defined as the sum of 48 thirty-minute squared returns over the day.

---

[7] A similar treatment of alternative models, including those which simultaneously incorporate jumps and stochastic volatility, is beyond the scope of the present analysis and is left for future research.

[8] For instance, see Engle and Bollerslev (1986) and Bollerslev (1987) for early evidence.

[9] This intuition provides a partial explanation for the success of including the daily high and low price (see Parkinson, 1980, Garman and Klass, 1980, Beckers, 1983, and Taylor, 1987, among others), daily volume

3.3 Exponential GARCH (EGARCH)

A possible limitation of the GARCH formulations described above is that the conditional variance is assumed symmetric in the lagged error terms. In response to this, Nelson (1991) proposed the EGARCH model specifically to capture the (negative) correlation between variance and past returns (a phenomenon most frequently observed in equity markets)[10]. This model links the conditional variance to the deviation of the absolute magnitude of the lagged scaled residual from its mean, as well as the sign of this innovation. In another distinction from the GARCH model, there are no restrictions placed on the parameters of the EGARCH process to ensure non-negativity of the conditional variance, as a logarithmic form is used. The conditional variance of the EGARCH(1,1) model is given by,

$$\log(\sigma_t^2) = \gamma + \alpha \cdot g_{t-1} + \beta \cdot \sigma_{t-1}^2 \tag{4}$$

$$\text{where } g_t = \left[\left| \frac{\varepsilon_t}{\sigma_t} \right| - \sqrt{\frac{2}{\pi}} - \eta \left(\frac{\varepsilon_t}{\sigma_t} \right) \right] \tag{5}$$

where the stationarity condition imposes $|\beta| < 1$.

3.4 Heterogeneous ARCH (HARCH)

In attempting to capture the long memory of squared intraday returns described by Ding, Granger and Engle (1993), Müller et al. (1997) proposed the HARCH model[11]. The conditional variance of the HARCH process appears as a linear combination of the squares of aggregated returns measured at different intraday frequencies,

$$\sigma_t^2 = \gamma + \sum_{j=1}^{n} \alpha_j \left(\sum_{i=1}^{j} r_{t-i} \right)^2 \tag{6}$$

where $\gamma > 0$, $\alpha_n > 0$ and $\alpha_j \geq 0$ to ensure a well defined process. The conditional variance equation of a HARCH(3) model may help illustrate the properties of the model,

$$\sigma_t^2 = \gamma + \alpha_1 \cdot (r_{t-1})^2 + \alpha_2 \cdot (r_{t-1} + r_{t-2})^2 + \alpha_3 \cdot (r_{t-1} + r_{t-2} + r_{t-3})^2$$

(Bessembinder and Seguin, 1993), number of price changes (Laux and Ng, 1993), and the standard deviation of intraday returns (Taylor and Xu, 1997) in studies attempting to model asset return variability.

[10] The GARCH model links the conditional variance to past squared innovations, implying the conditional variance is insensitive to the sign of past returns. The GJR GARCH model of Glosten, Jagannathan and Runkle (1993), the Nonlinear GARCH and Nonlinear Asymmetric GARCH models of Engle and Ng (1993) are alternative specifications to the EGARCH model that are also designed to capture the (negative) correlation between variance and the sign of past returns.

[11] Whilst the FIGARCH model of Baillie et al. (1996) has also been designed for a similar purpose, it cannot reproduce the lead-lag correlation feature of the HARCH model reported in Dacorogna et al. (2001).

3.5 Interpretation of Parameter Estimates

The parameter values (estimated by maximum likelihood) for the GARCH(1,1) model with normal and Student-*t* distributed errors, the GARCH(1,1) model extended with cumulative intraday squared returns, the EGARCH(1,1) model and the HARCH model, are presented in Tables 1(a) and 1(b)[12]. Two features of the GARCH coefficient estimates are worthy of a brief discussion.

Referring to Table 1(b), it is interesting to find the α parameter is dominated by the additional κ parameter (representing intraday squared returns) in the GARCH-I model. Recall that the α parameter represents hourly, eight hourly or daily squared returns respectively, while the κ term represents the sum of 2, 16 or 48 thirty-minute squared returns respectively[13]. This suggests the sum of squared intraday returns provides a relatively less noisy update of the conditional variance, as noted in Martens (2001).

Theoretical aggregation results are available for the GARCH(1,1) model (see Drost and Nijman, 1993 and Nelson and Foster, 1994)[14]. From this it can be inferred that if GARCH constitutes the RGP at one particular frequency, the behaviour of the data sampled at any other frequency can be determined by temporal aggregation or disaggregation of the original process. These implied processes can then be compared to the empirically estimated processes at the same frequencies. The figures presented in the 'implied' columns in Table 1(a) use the daily estimations as a reference basis to imply out the parameter values for the higher frequencies (ie. via disaggregation)[15]. The estimated coefficients appear quite reasonable, moving in the direction suggested by disaggregation in the majority of cases. For instance, as the data is sampled more frequently, the α parameter slowly approaches zero from above while the β parameter approaches unity from below[16]. This results in the autoregressive root of the conditional variance process, $\alpha + \beta$, also approaching unity, implying shocks to the conditional variance become more persistent as the sampling frequency increases[17]. Given the GARCH(1,1) coefficients generally appear well behaved (according to disaggregation theory), they should facilitate a reliable analysis in the simulation exercise that follows[18].

---

[12] Diagnostic results for the standardized residuals are presented in the Appendix.

[13] Note at the half hour frequency, the α and κ parameters both represent thirty-minute squared returns.

[14] Of the models presented in section 3, aggregation results are only available for the GARCH model.

[15] The daily frequency is used as the reference basis, given that at this time horizon, the GARCH(1,1) model has served as a reasonable approximation of the RGP for a range of asset classes.

[16] In this context, the GARCH model can be interpreted as either a jump process according to Drost and Nijman (1993), or a diffusion process based on the results presented in Nelson and Foster (1994).

[17] The two most flagrant exceptions occur at the hourly frequency for the JPY/USD for the GARCH model assuming normally distributed errors, and the DEM/USD for the GARCH model with Student-*t* errors.

[18] A number of studies have documented that explicitly accounting for intraday seasonal volatility patterns does not have a discernible impact on the modelling ability (see Andersen and Bollerslev, 1997b, and Dacorogna et al., 2001) or forecasting performance (see Martens, 2001) of ARCH processes. However, to circumvent the potential problem of coefficient estimates becoming biased if these intraday patterns are not treated, the simulations for the GARCH(1,1) model (with normal and Student-*t* errors) were also performed using the intraday parameter values implied from the coefficient estimates of the GARCH(1,1) model fitted to daily returns (whose parameters are unaffected by intraday seasonality). The qualitative results from this method and the method using maximum likelihood coefficients from models fitted directly to intraday returns were indistinguishable, so for consistency only those based on the latter approach are reported.

[INSERT TABLE 1(a) & (b)]

4. Comparing the Unconditional Distribution of Empirical and Simulated Returns

This section investigates whether the unconditional distribution of high frequency foreign exchange returns can be approximated by the unconditional distribution of returns simulated by the ARCH models presented in the previous section.

4.1 Simulating Discrete Time Volatility Models

In the Monte Carlo procedure, simulated residuals are generated by sampling from the i.i.d. standard normal (0,1) random number generator[19]. The simulated residuals are combined with the conditional mean to form a series of simulated returns following equation (1). The conditional variance of simulated returns is determined by a GARCH(1,1), GARCH-I, EGARCH(1,1), or HARCH process (as outlined in section 3), using the parameter values in Table 1(a) and 1(b).

4.2 Simulating the GARCH(1,1) model in Continuous Time: GARCH Diffusion

A continuous time GARCH diffusion process may also be capable of reproducing the distributional characteristics of intraday exchange rate returns. In line with many theoretical asset pricing models and derivative pricing theories, it is assumed that instantaneous returns are generated by the continuous time martingale,

$$d \ln p_t = \sigma_t \cdot dW_{p,t} \tag{7}$$

where $d\ln p_t$ is the instantaneous change in the log price and σ_t is a stochastic process independent of the instantaneous change in the standard Weiner process, $dW_{p,t}$[20]. A natural continuous time model for the conditional variance process is given by the diffusion limit of the GARCH(1,1) process as developed by Nelson (1990),

$$d\sigma_t^2 = \theta(\omega - \sigma_t^2) \cdot dt + (2\lambda\theta)^{0.5} \sigma_t^2 \cdot dW_{\sigma,t} \tag{8}$$

where $\omega > 0$, $\theta > 0$, $0 < \lambda < 1$, and the Weiner processes, $W_{p,t}$ and $W_{\sigma,t}$, are independent. The exact relationship between the discrete time weak GARCH(1,1) parameters and the continuous time stochastic volatility parameters in equation (8) can be expressed by,

---

Furthermore, the GARCH diffusion coefficients (presented in section 4) are unaffected by intraday seasonal volatility patterns as the intraday parameters are implied from the GARCH(1,1) model fitted to daily returns. Note also that even after conducting the simulation exercise for the GARCH-I, EGARCH(1,1) and HARCH models using a range of different coefficient estimates, the qualitative results remained unchanged. This suggests that even if the parameter estimates were slightly affected by intraday seasonal volatility patterns, for the purposes of this study, the impact would have been negligible.

[19] Simulated returns for the GARCH(1,1) model are also generated by using residuals sampled from the Student-*t* random number generator. Note that the seed was kept fixed while generating the random numbers for each of the simulated models.

[20] Any mean predictability could easily be incorporated into this model, but the assumption of mean-zero returns in (7) is consistent with the empirical evidence for the four exchange rates in this analysis.

$$\theta = -m \cdot \log(\alpha + \beta), \tag{9}$$

$$\omega = m \cdot \gamma \cdot (1 - \alpha + \beta)^{-1}, \tag{10}$$

$$\lambda = 2 \cdot \alpha \cdot \log^2(\alpha + \beta) \cdot [1 - \beta \cdot (\alpha + \beta)] \cdot \left\{ \begin{matrix} [1 - (\alpha + \beta)^2] \cdot (1 - \beta)^2 - \alpha \cdot [1 - \beta \cdot (\alpha + \beta)] \cdot \\ [6 \cdot \log(\alpha + \beta) + 2 \cdot \log^2(\alpha + \beta) + 4 \cdot (1 - \alpha - \beta)] \end{matrix} \right\}^{-1} \tag{11}$$

with a sampling frequency m times per day (ie. m = 24 for hourly modelling when the stochastic volatility parameters are implied by the α and β parameters fitted to daily returns). Equation (9) implies that $\lim_{m\to\infty} (\alpha + \beta) = 1$, so the weak GARCH(1,1) model converges to the IGARCH case of Engle and Bollerslev (1986) as the sampling frequency increases. Therefore, this diffusion approximation provides a possible explanation for the empirical findings of IGARCH behaviour in high frequency asset returns, as noted by Nelson (1990). The numerical simulations of the model in equations (7) and (8) are performed using a standard Euler discretization scheme as follows,

$$\ln p_{t+\Delta} = \ln p_t + \sigma_t \cdot \Delta^{0.5} \cdot w_{p,t} \tag{12}$$

$$\sigma^2_{t+\Delta} = \theta \cdot \omega \cdot \Delta + \sigma^2_t \cdot (1 - \theta \cdot \Delta + [2 \cdot \lambda \cdot \theta \cdot \Delta]^{0.5} \cdot w_{\sigma,t}) \tag{13}$$

where $w_{p,t}$ and $w_{\sigma,t}$ denote independent standard normal variables. Following Andersen and Bollerslev (1998), in the implementation Δ = 1/2880, corresponding to ten observations per five-minute interval.

4.3 Tests for Comparing the Distribution of Empirical and Simulated Returns

To test if the simulated and empirical unconditional return distributions are significantly different, this study tests whether the unconditional moments of the empirical distribution lie inside the 95th or 99th percentiles of the unconditional simulated moments.

Each sample path of simulated returns produces an associated mean, standard deviation, skewness and kurtosis. Given that 10,000 return paths are simulated for each model, this results in the generation of 10,000 values for each of the first four moments[21]. Rather than taking a single number (such as the average, median etc.) as being representative of these 10,000 values, the test uses confidence intervals based on the entire distribution of

[21] 1,000 warm-up replications were performed prior to basing the results on a further 10,000 replications. Also note that the number of observations used in the simulation exercise corresponded to the number in the empirical sample. This resulted in 24,114 observations for the thirty-minute series, 12,057 for the hourly series, 1507 for the eight-hour series and 503 for the daily series. All simulations were performed in RATS programming language.

values. To illustrate, assume that the empirical moment lies inside the 95th percentile of the simulated moments. In this instance, the null hypothesis that the simulated and empirical moments are insignificantly different cannot be rejected at the 5 percent level. If the empirical moment lies outside the 95th percentile of the simulated moments but inside the 99th percentile, the null is rejected at the 5 percent level, but not at the 1 percent level. If the empirical moment lies outside the 99th percentile of the simulated moments, the null is rejected at the 1 percent level. A model can be considered representative of the true RGP only when all four empirical (unconditional) moments are insignificantly different to the unconditional moments of the simulated process.

4.4 Statistical Properties of High Frequency Foreign Exchange Returns

The existence of volatility clustering in squared returns measured at different frequencies has been extensively documented, dating back to Mandelbrot (1963) and Fama (1965). This feature is ubiquitous in Figure 1, with DEM/USD squared returns (measured at the thirty-minute, hourly, eight-hourly and daily frequencies) exhibiting well-defined periods of tranquility and turbulence[22]. The visual impression of volatility clustering is confirmed by the Ljung-Box portmanteau tests for serial correlation in squared returns presented in Table 2. The Q-statistics based on squared returns over different frequencies for each of the exchange rates are highly significant. The Lagrange multiplier test applied to squared returns confirms the existence of ARCH disturbances.

[INSERT FIGURE 1]

The figures presented in Table 2 also highlight the extreme departures from normality exhibited by intraday foreign exchange returns. Whilst mean returns are indistinguishable from zero, at each frequency almost all exchange rates are significantly skewed away from zero. In addition, excess kurtosis is rife in that the values exceed 3, which is the theoretical value for a Gaussian distribution. All of the rates display the same general behaviour. For instance, a decreasing kurtosis is associated with a coarser sampling frequency, and at the shortest time intervals, the kurtosis values are extremely large. The standard deviation of returns increases dramatically as the frequency increases, though is quite stable when expressed in annualized terms. Figure 1 and Table 2 are entirely consistent with the statistical features of financial asset returns reported in the literature.

[INSERT TABLE 2]

4.5 Results of the Tests Comparing Empirical and Simulated Return Distributions

Table 3 presents the results of the Monte Carlo exercise examining whether the unconditional moments of the empirical distribution lie inside the 95th and 99th percentiles of the unconditional simulated moments. In this table, "accept" refers to occasions where the null hypothesis that the simulated and empirical moments are the same cannot be rejected at the 5 percent level; "5%" refers to instances where the null is rejected at the 5 (but not the 1) percent level; and "1%" refers to instances where the null is rejected at the 1 percent level.

[22] Similar patterns of volatility clustering were exhibited by each currency across different frequencies.

The GARCH(1,1) model with normal errors does a reasonable job at characterizing the unconditional mean and standard deviation of exchange rate returns at different frequencies. However it cannot be considered representative of the true RGP as it fails to reproduce the behaviour of the third and fourth moments of the empirical return distribution at frequencies higher than a calendar day (24 hours). As expected, the GARCH(1,1) specification following a Student-*t* distribution performs slightly better at capturing the behaviour of the unconditional fourth moment than its normal-error counterpart. The results for the GARCH-I model are similar to those reported for the GARCH(1,1) models. As anticipated, it appears the EGARCH(1,1) model is fairly successful in tracking the skewness associated with exchange rate returns at the trading day (eight hour) and calendar day frequencies. The unconditional fourth moment of returns simulated by the HARCH model and the continuous time GARCH diffusion - in fact all the models examined here – are unable to approximate the leptokurtic unconditional fourth moment of high frequency foreign exchange returns.

In reconciling these results with related research, Müller et al. (1997) suggest ARCH models are incapable of replicating intraday foreign exchange RGP's due to a range of independent volatility components inherent in high frequency data. They imply the GARCH model is unable to capture the heterogeneity of traders acting under different time horizons and objectives. Andersen and Bollerslev (1997a) claim standard ARCH models cannot accommodate the regular cyclical patterns in intraday volatility associated with the opening and closing times of financial centers around the world. At the highest frequencies, institutional and behavioural features of the trading process such as non-synchronous trading, the bid-ask bounce and other microstructure effects may also preclude ARCH models from constituting the intraday RGP.

[INSERT TABLE 3]

The results presented in Table 3 show the ARCH class of models are reasonably successful in approximating high frequency foreign exchange RGP's only at the calendar day frequency. This is not surprising given the ARCH family of models were originally developed to capture features of financial time series measured at daily and lower frequencies. The analysis presented thus far appears to support the findings in related studies, suggesting standard ARCH processes may be of limited use in high frequency applications. However, arguably the most rigorous test of the veracity of an asset price model is its ability to forecast future movements in a state variable. The most critical feature of conditional financial asset return distributions is the structure of the second moment - the dominant time-varying moment. Given that volatility permeates modern financial theories and the functioning of markets, it is non-trivial to investigate the performance of ARCH models in forecasting the conditional second moment of daily foreign exchange returns.

5. Forecasting Daily Variance with ARCH Models

In this section, one-step-ahead daily variance forecasts are assessed via a host of statistical procedures, given no universally accepted loss function exists for the ex-post evaluation and comparison of model forecasts. In addition to the statistical evaluation, forecasts are assessed using a profitability measure based on a trading game in a pseudo options market. Model parameters are estimated from daily returns from January 3, 1997 through June 30, 1998, leaving the period July 1, 1998 through December 30 (125 days) in which to assess the forecasts[23].

Note that the HARCH model is not used in the daily forecasting exercise given the parameter values from the model fitted to daily returns violated the associated stationarity conditions (see the notes under Table 1(b)). In addition, the theoretical GARCH diffusion framework is not designed for empirical implementations, so it is also excluded from the forecasting analysis. These two models are replaced with a nonparametric autoregressive (AR) process, and J.P. Morgan's Riskmetrics™ model. The AR estimator simply uses the daily realized variance (defined as the sum of thirty-minute squared returns) as the variance forecast for the following day. In the Riskmetrics™ model, the conditional variance of the current period return appears as a function of the conditional variance of yesterday's return σ_{t-1}^2, updated by the news revealed by yesterday's return ε_{t-1}^2,

$$\sigma_t^2 = \psi \cdot \sigma_{t-1}^2 + (1-\psi) \cdot \varepsilon_{t-1}^2 \tag{14}$$

where returns are sampled once per day. Whilst (14) bears some resemblance to the GARCH(1,1) process, Riskmetrics™ differs from the GARCH model in a number of ways. For instance, Riskmetrics™ does not allow for mean reversion in variance forecasts, and rather than estimating parameters on each data set using an optimization method, it fixes the only parameter, ψ, equal to 0.94. According to J.P. Morgan (1996), this value has been found to optimize the daily forecasting quality of (14) over a range of financial assets and test periods[24].

5.1 Constructing Estimates of Realized Variance

To evaluate the forecasting performance of competing models, forecasts must be compared to some measure of realized daily variance. It is common to see the daily squared return used for this purpose. Whilst the daily squared return is an unbiased estimator of ex-post daily variance, Andersen and Bollerslev (1998) point out it is also an extremely noisy one as the idiosyncratic component of daily returns is large. Prior to this, Merton (1980) demonstrated that latent (true but unobservable) variance can be approximated to an arbitrary precision by using the sum of intraday squared returns[25].

[23] As only 503 daily observations (from January 3, 1997, through to December 30, 1998) were available for the study, parameters used in the simulation exercise were based on the full sample period so as to utilize all the available data. To facilitate an out-of-sample forecasting analysis, parameters were re-estimated over the first 18 months of this period, leaving the remaining 125 days in which to assess the forecasts.

[24] As the Riskmetrics™ model has a parameter with a fixed value based on the use of daily data, it was not appropriate to use the model in the intraday simulation exercise.

[25] However, market microstructure effects may make sampling at the very highest frequencies problematic.

The figures in Table 4 highlight the advantage of using the sum of squared intraday returns, rather than the squared daily return, as the estimator of realized daily variance. In particular, the range and standard deviation of the ex-post daily realized USD/AUD and USD/GBP variance fall by more than 50% upon replacing the daily squared return with the sum of five-minute intraday squared returns. Consequently, realized daily variance is defined as the sum of intraday squared returns in the forecasting exercise that follows. This quantity is fully observable and provides a nonparametric estimate of latent variance over the same time interval.

[INSERT TABLE 4]

5.2 Statistical Analysis

The first metric used to evaluate daily variance forecasts is the coefficient of determination (R^2) resulting from the regression of realized daily variance, $\sigma^2_{real,t+k}$, on its forecast, $\hat{\sigma}^2_{t+k}$,

$$\sigma^2_{real,t+k} = \varphi_0 + \varphi_1 \cdot \hat{\sigma}^2_{t+k} + e_{t+k} \tag{15}$$

where $k=1$ (corresponding to one step ahead daily forecasts)[26]. If the conditional mean of returns is zero, and the conditional variance is specified correctly, φ_0 and φ_1 should equal zero and one respectively. However this regression is sensitive to extreme values, so the coefficient of determination from the log form of (15) may also be useful,

$$\log(\sigma^2_{real,t+k}) = \varphi_0 + \varphi_1 \cdot \log(\hat{\sigma}^2_{t+k}) + e_{t+k} \tag{16}$$

which is less sensitive to outliers than (15) as severe mispredictions are given relatively less weight. The remaining metrics used in the forecasting analysis are as follows,

$$\text{Root Mean Squared Error (RMSE)} = \sqrt{\frac{1}{T}\sum_{t=0}^{T-1}\left(\hat{\sigma}^2_{t+k} - \sigma^2_{real,t+k}\right)^2} \tag{17}$$

$$\text{Mean Absolute Error (MAE)} = \frac{1}{T}\sum_{t=0}^{T-1}\left|\hat{\sigma}^2_{t+k} - \sigma^2_{real,t+k}\right| \tag{18}$$

$$\text{Mean Absolute Percent Error (MAPE)} = \frac{1}{T}\sum_{t=0}^{T-1}\frac{\left|\hat{\sigma}^2_{t+k} - \sigma^2_{real,t+k}\right|}{\sigma^2_{real,t+k}} \tag{19}$$

$$\text{Median Squared Error (MSE)} = Median\left(\hat{\sigma}^2_{t+k} - \sigma^2_{real,t+k}\right)^2 \tag{20}$$

where forecasts of realized daily variance are assessed over a period of T days.

[26] Note that the R^2 of this regression is the squared correlation between the realized and forecasted variance.

5.3 Results from the Statistical Evaluation of Daily Variance Forecasts

Whilst significant in-sample intertemporal volatility persistence has been extensively documented in ARCH related research, a number of studies have found, albeit rather surprisingly, that ARCH models explain little of the variability in ex-post realized variance. These findings have been based on the measurement of ex-post realized variance as the squared or absolute return over the relevant forecast horizon in the regression of realized variance on the model forecast, as in equation (15)[27]. Examining the figures in Table 5 from left to right in the '24-hr' rows appears to offer further support to the perception that daily ARCH forecasts are of limited practical use. The low R^2's for the one day ahead ARCH forecasts indicates these models perform very poorly, consistently explaining less than one percent of the ex-post variability in the USD/AUD, USD/GBP and DEM/USD realized daily variance, and less than four percent in JPY/USD realized daily variance.

However, upon increasing the sampling frequency of squared returns in the measurement of ex-post realized daily variance, the fallacy of this conclusion becomes evident. For instance, when realized daily variance is defined as the sum of thirty-minute squared returns, the ARCH models produce R^2's as high as thirty-two percent for the USD/AUD, and fifty-six percent for the JPY/USD. Increasing the sampling frequency of realized variance to five-minutes results in even higher correlations, with R^2's as high as thirty-eight percent for the USD/AUD and sixty-one percent for the JPY/USD. These statistics signify an enormous increase in explanatory power relative to the inference based on R^2's documented previously[28]. Combined with the figures in Table 4, these results support Andersen and Bollerslev's (1998) argument that the poor predictive power of ARCH models, when judged by standard forecast evaluation criteria (ie. daily squared returns for daily forecasts), is a consequence of the noise inherent in the RGP. Furthermore, the trend of increasing R^2's associated with the use of squared intraday returns as the proxy for ex-post daily realized variance is also evident for the non-ARCH predictors - namely J.P. Morgan's Riskmetrics™ and the nonparametric AR estimator. Recall this AR estimator simply uses the realized daily variance (defined as the sum of thirty-minute squared returns) observed at the completion of the current day as the variance forecast for the following day. As such, the relatively high R^2's produced by the AR estimator is consistent with the volatility clustering phenomena displayed in Figure 1.

[INSERT TABLE 5]

In assessing the relative forecasting performance of the models in Table 5, recall that equation (15) regresses realized daily variance on the forecast of daily variance. Equation (16) mitigates the influence of outliers in equation (15) as it regresses the logarithm of realized variance on the logarithmic variance forecast. Using these metrics, Table 5 shows the nonparametric AR estimator to be superior in the case of the USD/GBP and

[27] See for example Cumby, Figlewski and Hasbrouck (1993), Figlewski (1994) and Jorion (1995, 1996).

[28] In the regression of realized daily variance on the model forecast in equations (15) and (16), the values of φ_0 and φ_1 were generally indistinguishable from zero and one respectively when the regression produced a relatively high R^2 (as expected). In instances where the regression in (15) and (16) produced a relatively low R^2, the parameter values of φ_0 and φ_1 invariably deviated from zero and one respectively.

DEM/USD using equation (15), and the AUD/USD and GBP/USD following equation (16). On one occasion each, the GARCH(1,1) models assuming a normal and Student-*t* distribution perform best following (15), and both models are consistently second best, or close to it, following both (15) and (16). However, the GARCH-I model using the sum of intraday squared returns (from the previous day) is perhaps the most consistent of all models. For the DEM/USD and JPY/USD it is the superior model following (16), and on two occasions it is the second best model using both (15) and (16). It is also interesting to see that whilst honours are evenly split between the GARCH(1,1) with normal errors and the GARCH-I model following (15) (in that they outperform each other twice), the GARCH-I model produces relatively higher R^2's in every instance when the influence of extreme mispredictions is reduced following (16). The EGARCH model and J.P. Morgan's Riskmetrics™ are relatively inferior at forecasting realized daily variance.

Table 6 presents the relative performance rankings of model forecasts, based on the loss functions described in section 5.2. Using the RMSE or Theil's U statistic as the criterion, the GARCH(1,1) with Gaussian errors is the superior model, and is arguably the best performed when examining the MAE rankings also. The GARCH-I model appears to perform best under the MAPE and MSE criterion, while the AR estimator seems to be particularly useful in forecasting the variability in DEM/USD realized variance using any criterion. These results, coupled with those presented in Table 5, make it quite difficult to distinguish between the superiority of the GARCH(1,1) model with normal errors, the GARCH-I model augmented with intraday returns, and the simple AR estimator. The pseudo options market trading game presented next is used to establish which of these estimators generate superior forecasts of realized daily variance.

[INSERT TABLE 6]

5.4 Profitability Assessment of Daily Variance Forecasts: An 'Economic' Interpretation

This section follows Engle, Hong, Kane and Noh (1993) and Maheu and McCurdy (2001) by using a profitability measure to determine the relative ranking of competing variance forecasts. Ex-ante forecasts of daily variance are used to price at-the-money European put and call options (with the same maturity) on a spot foreign currency position following the Garman and Kohlhagen (1983) model[29]. A short time-to-expiration (one day) is used to better approximate the constant volatility assumption in the Black and Scholes (1973) formula.

To make the pseudo options market operational, only two investors, A and B, are assumed to exist. Each investor is assigned their own forecasting model. Option contracts are combined into a straddle position with both investors computing a straddle price based on variance forecasts from their respective models[30]. Trades are initiated based on the investors comparing their straddle prices. For instance, if investor A's straddle price

[29] This is simply the Black and Scholes (1973) formula adapted to spot foreign currency positions.

[30] Straddles are a natural choice to gauge the relative merits of variance forecasting models as they can be used essentially as a bet on future volatility. A long (short) straddle position involves the simultaneous purchase (sale) of a call and a put option on the same currency. A long (short) straddle position will become profitable if realized volatility is substantially higher (lower) than the market expectation.

is greater than investor B's, investor A buys one straddle from investor B (at the mid point of the two prices)[31]. Otherwise investor A takes a short position and sells one straddle to investor B. Straddles are repriced each day using updated one-day-ahead variance forecasts[32]. Once the next day's spot price has been realized, payouts from the zero-sum game are calculated.

The figures presented in Table 7 show that no clear winner emerges when the investors trade against one another using the AR and GARCH(1,1) estimators respectively. For instance, these predictors outperform each other twice, and on one of these two occasions, the profits generated (at the expense of the other) are statistically significantly different from zero. However, it is possible to infer from the remaining figures in Table 7 that the GARCH-I model is the superior of the three, as it makes statistically significant profits at the expense of the AR estimator and GARCH(1,1) model for all but the USD/AUD contract[33]. This result can also be reconciled with the work in Nelson (1992) and Nelson and Foster (1995), who demonstrated theoretically that variance forecasts can be made as accurate as required for many diffusion models by using ARCH estimates and sufficiently frequent price measurements. This suggests that using the sum of intraday squared returns in both the ex-post measurement of realized variance, as well as in the forecasting of variance directly, has valuable practical applications.

[INSERT TABLE 7]

6. Conclusion

One of the many challenges posed by the study of high frequency financial market data is to develop models capable of explaining asset price behaviour at a range of frequencies. This presents researchers with new opportunities. However at the same time, it calls into question whether established time series models are useful in high frequency applications. This paper addresses this issue from two perspectives. A Monte Carlo procedure is used to investigate whether the unconditional distribution of high frequency foreign exchange returns can be approximated by the unconditional distribution of returns simulated by widely used ARCH models. The study then uses high frequency data to generate and evaluate forecasts of daily variance.

---

[31] To illustrate this idea further, if investor A's straddle price is say \$10, then they will be willing to buy (sell) a straddle at any price below (above) \$10, assuming no transaction costs. Now, if investor B's straddle price is \$9, it is reasonable to assume the trade would take place at the mid point of these two prices, here \$9.50, given that investor A wishes to buy at the lowest possible price below \$10, and investor B wishes to sell at the highest possible price above \$9.

[32] These contracts were set at AUD\$50,000, GBP£31,250, DM62,500 and ¥6,250,000 respectively, as is the case for standardized European currency option contracts traded on the Philadelphia Stock Exchange. In the pricing of the straddle contracts, the annualized domestic and foreign interest rates were arbitrarily assumed to equal 4 and 5 percent respectively. Different combinations of domestic and foreign interest rates yielded similar results to those reported in Table 7.

[33] Note that realized daily variance, against which the competing forecasts are measured, is defined as the sum of 48 thirty-minute squared returns in the trading game.

The simulation exercise demonstrates that the ARCH family of models is unable to reproduce the unconditional distribution of foreign exchange returns at frequencies higher than 24 hours. This is largely a legacy of the heavy-tailed feature of intraday returns. However, results from the forecasting analysis extend those in Andersen and Bollerslev (1998) by showing that a range of standard volatility models (not just the GARCH(1,1) model) can in fact produce accurate forecasts of daily variance. In other words, it is possible for ARCH models to predict variability in the conditional second moment of daily foreign exchange returns. This is attributed to the use of frequently sampled data in the construction of estimates of realized variance (against which forecasts are measured). In addition, the inclusion of the sum of squared intraday returns in the GARCH(1,1) model yields improvements in the modelling, and most notably forecasting, of realized daily variance. This appears to be an artifact of the noise inherent in using the daily squared return as an estimator of realized daily variance.

This paper demonstrates that whilst established econometric models do not constitute the intraday foreign exchange RGP, this should not immediately preclude these models from high frequency applications. Instead, the forecasting exercise demonstrates practical benefits are easily attainable from using high frequency data to develop and evaluate existing asset pricing models.

Emerging interest in the field of high frequency finance has seen a new set of stylized facts specific to intraday foreign exchange (and other) markets begin to surface. For example, the identification of a hyperbolic long memory decay in the autocorrelation of volatility, and the associated implications for scaling laws and fractal structures are promising contemporary developments. No doubt a number of similar discoveries are waiting to be made. Given that the evolution of asset pricing models has largely been motivated by empirical findings and economic interpretations, future research along these lines may pave the way for the development of models consistent with the behaviour of asset prices across different frequencies. In the case of any successful achievement, benefits will extend across a range of financial applications, not limited to derivative pricing, portfolio selection and risk management.

References

Andersen, T.G. and T. Bollerslev (1997a). "Intraday Periodicity and Volatility Persistence in Financial Markets," Journal of Empirical Finance, 4, 115-158.

Andersen, T.G. and T. Bollerslev (1997b). "Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns," Journal of Finance, 52, 975-1005.

Andersen, T.G. and T. Bollerslev (1998). "Answering the Critics: Yes, ARCH Models do Provide Good Volatility Forecasts," International Economic Review, 39, 885-905.

Baillie, R.T. and T. Bollerslev (1989). "The Message in Daily Exchange Rates: A Conditional Variance Tale," Journal of Business and Economic Statistics, 7, 297-305.

Baillie, R.T., T. Bollerslev and H.O. Mikkelsen (1996). "Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity," Journal of Econometrics, 74, 3-30.

Beckers, S. (1983). "Variances of Security Price Returns Based on High, Low and Closing Prices," Journal of Business, 56, 97-112.

Bera, A.K. and M.L. Higgins (1993). "A Survey of ARCH Models: Properties, Estimation and Testing," Journal of Economic Surveys, 7 (4), 1-58.

Besseminder, H. and P.J. Seguin (1993). "Price Volatility, Trading Volume and Market Depth: Evidence from Futures Markets," Journal of Financial and Quantitative Analysis, 28, 21-39.

Black, F. and M. Scholes (1973). "The Pricing of Options and Corporate Liabilities," Journal of Political Economy, 81, 637-654.

Bollerslev, T. (1986). "Generalized Autoregressive Heteroskedasticity," Journal of Econometrics, 31, 307-327.

Bollerslev, T. (1987). "A Conditional Heteroskedastic Time Series Model for Speculative Prices and Rates of Return," Review of Economics and Statistics, 69, 542-547.

Bollerslev, T., R.T. Chou and K.F. Kroner (1992). "ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence," Journal of Econometrics, 52, 5-59.

Bollerslev, T. and H. Zhou (2002). "Estimating Stochastic Volatility Diffusion Using Conditional Moments of Integrated Volatility," Research Paper, National Bureau of Economic Research.

Cumby, R.S., S. Figlewski and J. Hasbrouck (1993). "Forecasting Volatility and Correlations with EGARCH Models," Journal of Derivatives, Winter, 51-63.

Dacorogna, M.M., R. Gencay, U.A. Müller, R.B. Olsen and O.V. Pictet (2001), "An Introduction to High Frequency Finance," San Diego: Academic Press.

Ding, Z., C.W.J. Granger, and R.F. Engle (1993). "A Long Memory Property of Stock Market Returns and a New Model," Journal of Empirical Finance, 1, 83-106.

Drost, F.C. and T.E. Nijman (1993). "Temporal Aggregation of GARCH Processes," Econometrica, 61, 909-927.

Engle, R.F. (1982). "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation," Econometrica, 50 (4), 987-1007.

Engle, R.F. and T. Bollerslev (1986). "Modeling the Persistence of Conditional Variances," Econometric Reviews, 5, 1-50.

Engle, R.F. and V.K. Ng (1993). "Measuring and Testing the Impact of News on Volatility," Journal of Finance, 48 (5), 1749-1779.

Engle, R.F., C. Hong, A. Kane and J. Noh (1993), "Arbitrage Valuation of Variance Forecasts," Advances in Futures and Options Research, 6, 393-415.

Fama, E.F. (1965). "The Behaviour of Stock Market Prices," Journal of Business, 38, 34-105.

Figlewski, S. (1994). "Forecasting Volatility Using Historical Data," Manuscript, Department of Finance, New York University.

Garman, M.B. and M.J. Klass (1980). "On the Estimation of Security Price Volatilities from Historical Data," Journal of Business, 53, 67-78.

Garman, M.B. and S. Kohlhagen (1983). "Foreign Currency Option Values," Journal of International Money and Finance, 2, 231-237.

Glosten, L.R., R. Jagannathan and D. Runkle (1993). "Relationship Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks," Journal of Finance, 48 (5), 1779-1801.

Goodhart, C.A.E. and M. O'Hara (1997). "High frequency Data in Financial Markets: Issues and Applications," Journal of Empirical Finance, 4, 73-114.

Hsieh, D.A. (1989). "Modeling Heteroskedasticity in Daily Foreign Exchange Rates," Journal of Business and Economic Statistics, 7, 307-317.

J.P. Morgan (1996). "Riskmetrics," Technical Document, 4th ed. Reuters.

Jorion, P. (1995). "Predicting Volatility in the Foreign Exchange Market," Journal of Finance, 50, 507-528.

Jorion, P. (1996). "Risk and Turnover in the Foreign Exchange Market," in The Microstructure of Foreign Exchange Markets (J.A. Frankel, G. Galli and A. Giovannini eds.). Chicago: The University of Chicago Press.

Laux, P.A. and Ng, L.K. (1993). "The Sources of GARCH: Empirical Evidence from an Intraday Returns Model Incorporating Systematic and Unique Risks," Journal of International Money and Finance, 12, 543-560.

Maheu, J.M. and T.H McCurdy (2001). "Nonlinear Features of Realized FX Volatility," Working Paper, J.L Rotman School of Management, University of Toronto.

Mandelbrot, B. (1963). "The Variation of Certain Speculative Prices," Journal of Business, 36, 394-419.

Martens, M. (2001). "Forecasting Daily Exchange Rate Volatility Using Intraday Returns," Journal of International Money and Finance, 20, 1-23.

McCurdy, T.H. and I. Morgan (1988). "Tests of the Martingale Hypothesis for Foreign Currency Futures with Time Varying Volatility," International Journal of Forecasting, 3, 131-148.

Merton, R.C. (1980). "On Estimating the Expected Return on the Market: An Exploratory Investigation," Journal of Financial Economics, 8, 323-361.

Müller, U.A., M.M. Dacorogna, R.D. Dave, R.B. Olsen, O.V. Pictet and J.E. von Weizsacker (1997). "Volatilities of Different Time Resolutions – Analyzing the Dynamics of Market Components," Journal of Empirical Finance, 4, 213-239.

Nelson, D.B. (1990). "ARCH Models as Diffusion Approximations," Journal of Econometrics, 45, 7-38.

Nelson, D.B. (1991). "Conditional Heteroskedasticity in Asset Returns: A New Approach," Econometrica, 59, 347-370.

Nelson, D.B. (1992). "Filtering and Forecasting with Misspecified ARCH Models I: Getting the Right Variance with the Wrong Model," Journal of Econometrics, 67, 303-335.

Nelson, D.B. and D.P. Foster (1994). "Asymptotic Filtering Theory for Univariate ARCH Models," Econometrica, 62 (1), 1-41.

Nelson D.B. and D.P. Foster (1995). "Filtering and Forecasting with Misspecified ARCH Models II: Making the Right Forecast with the Wrong Model," Journal of Econometrics, 67, 303-335.

Parkinson, M. (1980). "The Extreme Value Method for Estimating the Variance of the Rate of Return," Journal of Business, 53, 61-65.

Taylor, S.J. (1987). "Forecasting the Volatility of Currency Exchange Rates," International Journal of Forecasting, 3, 159-170.

Taylor, S.J. and X. Xu (1997). "The Incremental Volatility Information in One Million Foreign Exchange Quotations," Journal of Empirical Finance, 4, 317-340.

Zumbach, G. (2002). "Volatility Processes and Volatility Forecasts with Long Memory," Manuscript, Charles Baudouin, Switzerland.

Figure 1. Annualized Squared DEM/USD Returns Across Different Frequencies

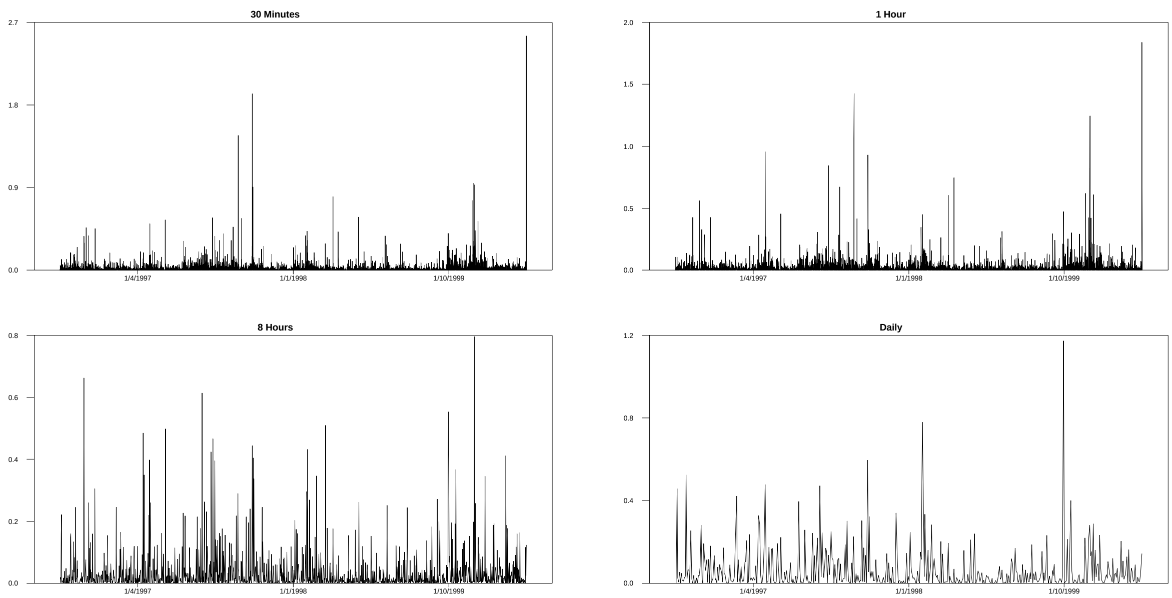

Figure 1 displays the time series of squared DEM/USD returns measured at different frequencies (converted to an annualized figure and expressed in percentage terms) from the period commencing January 3, 1997 through December 30, 1998 inclusive.

Table 1(a). Empirical and Implied GARCH(1,1) Parameter Estimates

	GARCH(1,1) (normal)					GARCH(1,1) (Student-*t*)					
				Implied						Implied	
	$\gamma \cdot 10^8$	α	β	$\hat{\alpha}$	$\hat{\beta}$	$\gamma \cdot 10^8$	α	β	υ	$\hat{\alpha}$	$\hat{\beta}$
30 Minutes											
USD/AUD	0.32	0.013	0.984	0.004	0.993	0.94	0.014	0.985	3.19	0.003	0.996
	(0.01)	(0.000)	(0.000)			(0.09)	(0.001)	(0.001)	(0.05)		
USD/GBP	0.43	0.012	0.981	0.006	0.987	0.65	0.015	0.985	3.27	0.005	0.991
	(0.02)	(0.000)	(0.001)			(0.08)	(0.001)	(0.001)	(0.05)		
DEM/USD	0.29	0.009	0.987	0.036	0.911	0.10	0.015	0.985	3.01	0.033	0.938
	(0.01)	(0.000)	(0.000)			(0.04)	(0.001)	(0.001)	(0.06)		
YEN/USD	0.69	0.014	0.982			0.75	0.014	0.984	3.02		
	(0.03)	(0.000)	(0.000)			(0.04)	(0.001)	(0.001)	(0.05)		
1 Hour											
USD/AUD	0.71	0.017	0.980	0.002	0.995	2.73	0.014	0.980	3.31	0.001	0.996
	(0.05)	(0.001)	(0.001)			(1.26)	(0.002)	(0.002)	(0.06)		
USD/GBP	3.05	0.033	0.944	0.004	0.991	1.16	0.014	0.981	3.44	0.002	0.992
	(0.15)	(0.001)	(0.002)			(0.20)	(0.002)	(0.003)	(0.08)		
DEM/USD	2.02	0.025	0.961	0.018	0.936	6.22	0.018	0.825	3.17	0.009	0.943
	(0.88)	(0.001)	(0.001)			(1.03)	(0.003)	(0.025)	(0.07)		
YEN/USD	28.45	0.119	0.808			7.35	0.013	0.980	3.19		
	(0.78)	(0.004)	(0.005)			(1.52)	(0.001)	(0.002)	(0.07)		
8 Hours											
USD/AUD	32.28	0.060	0.922	0.003	0.993	16.70	0.038	0.927	5.61	0.007	0.990
	(6.21)	(0.007)	(0.009)			(7.02)	(0.008)	(0.015)	(0.83)		
USD/GBP	44.79	0.023	0.923	0.007	0.988	7.41	0.009	0.958	4.17	0.013	0.982
	(14.26)	(0.005)	(0.008)			(3.61)	(0.004)	(0.018)	(0.36)		
DEM/USD	27.01	0.021	0.957	0.030	0.927	14.17	0.011	0.956	3.84	0.046	0.914
	(10.33)	(0.005)	(0.012)			(6.17)	(0.004)	(0.020)	(0.50)		
YEN/USD	140.93	0.054	0.900			14.98	0.015	0.956	3.46		
	(32.30)	(0.006)	(0.012)			(7.01)	(0.004)	(0.010)	(0.33)		
24 Hours											
USD/AUD	360.14	0.065	0.785	0.013	0.983	182.62	0.064	0.851	6.96	0.002	0.996
	(175.84)	(0.030)	(0.056)			(81.46)	(0.028)	(0.053)	(2.27)		
USD/GBP	314.06	0.027	0.843	0.024	0.967	153.21	0.014	0.850	6.36	0.004	0.992
	(33.13)	(0.011)	(0.019)			(16.89)	(0.006)	(0.112)	(2.11)		
DEM/USD	834.24	0.043	0.833	0.100	0.830	285.03	0.062	0.822	7.03	0.027	0.940
	(333.85)	(0.021)	(0.198)			(27.72)	(0.016)	(0.117)	(2.31)		
YEN/USD	1461.35	0.141	0.663			350.27	0.057	0.847	5.27		
	(387.67)	(0.027)	(0.061)			(106.20)	(0.024)	(0.060)	(1.23)		

Table 1(a) presents the parameter estimates for the GARCH(1,1) model with the error term following a normal and Student-*t* distribution. The conditional variance equation is defined in equation (2) for the GARCH(1,1) model. The υ parameter in the GARCH(1,1) model with Student-*t* errors refers to the corresponding degrees of freedom. Standard errors are reported in parenthesis. Table 1(a) also presents the disaggregation results where the GARCH(1,1) parameters estimated on daily data serve as the reference basis for implying out the higher frequency parameters. The figures in the implied columns (underneath the $\hat{\alpha}$ and $\hat{\beta}$ coefficients) are implied from Drost and Nijman's (1993) disaggregation formulas.

Table 1(b). GARCH-I, EGARCH and HARCH Parameter Estimates

	GARCH-I				EGARCH(1,1)				HARCH				
	$\gamma \cdot 10^7$	α	β	κ	γ	α	β	η	$\gamma \cdot 10^6$	IMP_1	IMP_2	IMP_3	IMP_4
30 Minutes													
USD/AUD	0.03	0.01	0.98	0.01	0.03	0.03	1.00	0.04	0.72	0.16	0.40	0.30	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.02)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
USD/GBP	0.04	0.01	0.98	0.01	0.36	0.35	0.83	0.03	0.50	0.14	0.21	0.36	0.17
	(0.00)	(0.00)	(0.00)	(0.00)	(0.04)	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)
DEM/USD	0.03	0.00	0.99	0.00	0.85	0.27	0.09	0.05	0.56	0.14	0.28	0.51	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.03)	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)
YEN/USD	0.07	0.01	0.98	0.01	0.05	0.02	1.00	0.24	1.56	0.14	0.24	0.21	0.28
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.02)	(0.01)	(0.00)	(0.01)	(0.00)	(0.00)
1 Hour													
USD/AUD	0.07	0.00	0.98	0.02	0.04	0.04	1.00	0.04	1.95	0.15	0.27	0.56	0.00
	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)	(0.03)	(0.01)	(0.00)	(0.00)	(0.00)
USD/GBP	0.18	0.00	0.95	0.03	0.59	0.36	0.73	0.12	0.98	0.15	0.20	0.57	0.00
	(0.01)	(0.00)	(0.00)	(0.00)	(0.10)	(0.01)	(0.01)	(0.02)	(0.02)	(0.01)	(0.00)	(0.00)	(0.00)
DEM/USD	0.27	0.01	0.60	0.24	0.95	0.28	0.83	0.05	1.03	0.15	0.38	0.07	0.01
	(0.01)	(0.00)	(0.01)	(0.01)	(0.06)	(0.01)	(0.00)	(0.02)	(0.02)	(0.01)	(0.00)	(0.01)	(0.00)
YEN/USD	3.03	0.07	0.78	0.17	0.09	0.04	0.99	0.10	3.33	0.14	0.25	0.48	0.00
	(0.08)	(0.01)	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)	(0.03)	(0.03)	(0.01)	(0.00)	(0.00)	(0.00)
8 Hours													
USD/AUD	3.18	0.01	0.89	0.08	0.17	0.13	0.98	0.16					
	(0.81)	(0.01)	(0.02)	(0.01)	(0.04)	(0.01)	(0.00)	(0.08)					
USD/GBP	1.76	0.02	0.94	0.03	0.00	0.00	1.00	0.22					
	(0.63)	(0.02)	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)	(0.09)					
DEM/USD	16.02	0.02	0.80	0.07	0.90	0.07	0.92	0.87					
	(5.78)	(0.02)	(0.06)	(0.02)	(0.27)	(0.02)	(0.02)	(0.28)					
YEN/USD	14.41	0.09	0.84	0.12	0.55	0.10	0.95	0.77					
	(2.67)	(0.03)	(0.02)	(0.01)	(0.07)	(0.01)	(0.01)	(0.13)					
24 Hours													
USD/AUD	29.82	0.03	0.65	0.24	0.63	0.27	0.94	0.20					
	(20.53)	(0.04)	(0.11)	(0.09)	(0.14)	(0.05)	(0.03)	(0.01)					
USD/GBP	202.58	0.07	0.80	0.14	0.01	(0.00)	1.00	(0.02)					
	(73.09)	(0.06)	(0.13)	(0.02)	(0.00)	(0.00)	(0.00)	(0.01)					
DEM/USD	123.89	0.04	0.50	0.20	0.55	0.10	0.65	0.33					
	(15.83)	(0.07)	(0.12)	(0.02)	(0.19)	(0.04)	(0.18)	(0.13)					
YEN/USD	174.40	0.15	0.45	0.38	0.54	0.29	0.73	0.64					
	(62.43)	(0.03)	(0.10)	(0.05)	(0.26)	(0.05)	(0.06)	(0.17)					

Table 1(b) presents the parameter estimates for the GARCH(1,1) model extended with the sum of thirty-minute intraday squared returns (referred to as GARCH-I model), along with parameter estimates for the EGARCH(1,1) and HARCH models. The conditional variance equation is defined in equation (3) for the GARCH-I model, equations (4) and (5) for the EGARCH model, and equation (6) for the HARCH model. Standard errors are reported in parenthesis. When reporting parameter values for the HARCH model in equation (6), Muller et al. (1997) suggest that the coefficients α_j reflect the relative impact of different market components, i, via different relevant time intervals. Therefore, the impact, IMP_i,, of the i^{th} component can be defined as, $IMP_i = (q-1)q^{i-2}\frac{q^{i-1}+q^{i-2}+1}{2}\alpha_i$ where the j ranges are separated by powers of a natural number q, so the typical time interval size of a component differs from that of the neighbor component by a factor q (which is set equal to 4). Components of the HARCH model estimated for frequencies lower than 1 hour are not reported due to violation of the stationarity condition, which imposes that the sum of these impacts lie inside the unit circle. See Muller et al. (1997) for further discussion of the properties of the HARCH process

Table 2. Statistical Properties of High Frequency Foreign Exchange Returns

Time Interval	Mean * 10^4		StDev * 10^4	Ann (%)	Skewness		Kurtosis		LB Q-Stat	LM test
USD/AUD										
30 Minutes	-0.11	(0.07)	11.62	12.78	0.35	[0.00]	15.52	[0.00]	[0.00]	[0.00]
1 Hour	-0.22	(0.15)	16.51	12.84	0.54	[0.00]	14.42	[0.00]	[0.00]	[0.00]
8 Hours	-1.73	(1.20)	42.87	11.78	0.29	[0.00]	8.11	[0.00]	[0.00]	[0.00]
24 Hours	-5.13	(3.32)	74.49	11.83	0.77	[0.00]	6.75	[0.00]	[0.00]	[0.00]
USD/GBP										
30 Minutes	-0.01	(0.05)	8.07	8.87	-0.17	[0.00]	11.46	[0.00]	[0.00]	[0.00]
1 Hour	-0.03	(0.10)	11.33	8.81	-0.25	[0.00]	12.81	[0.00]	[0.00]	[0.00]
8 Hours	-0.21	(0.76)	29.50	8.81	-0.12	[0.03]	6.23	[0.00]	[0.00]	[0.00]
24 Hours	-0.49	(2.29)	51.36	8.15	0.11	[0.16]	4.01	[0.00]	[0.00]	[0.01]
DEM/USD										
30 Minutes	0.03	(0.06)	8.91	9.79	-0.39	[0.00]	19.68	[0.00]	[0.00]	[0.00]
1 Hour	0.06	(0.11)	12.54	9.75	-0.43	[0.00]	13.63	[0.00]	[0.00]	[0.00]
8 Hours	0.52	(0.91)	35.19	9.68	-0.16	[0.01]	5.31	[0.00]	[0.00]	[0.00]
24 Hours	1.62	(2.80)	62.80	9.97	-0.27	[0.01]	4.00	[0.00]	[0.00]	[0.01]
JPY/USD										
30 Minutes	-0.01	(0.09)	13.98	15.37	-0.69	[0.00]	49.73	[0.00]	[0.00]	[0.00]
1 Hour	-0.02	(0.18)	20.06	15.60	-0.83	[0.00]	38.76	[0.00]	[0.00]	[0.00]
8 Hours	-0.16	(1.42)	55.07	15.14	-0.67	[0.00]	9.80	[0.00]	[0.00]	[0.00]
24 Hours	-0.35	(4.28)	95.79	15.21	-1.26	[0.00]	8.78	[0.00]	[0.00]	[0.00]

Table 2 presents the first four moments of the unconditional return distribution at different time intervals for the USD/AUD, USD/GBP, DEM/USD and JPY/USD exchange rates. The numbers in parenthesis to the right of the mean values are the associated standard errors, also raised by 10^4. The annualized standard deviation, expressed in percentage terms, is found by multiplying the standard deviation by the square root of the number of times the returns are sampled during a year. The numbers in rectangular brackets to the right of the skewness and kurtosis values are the marginal significance levels of the test of the null hypothesis that the skewness and excess kurtosis values are equal to zero and three respectively. The numbers in rectangular brackets reported in the second last column are the marginal significance levels of the Ljung-Box portmanteau tests for autocorrelation in squared returns with 48, 24, 15 and 10 degrees of freedom for the thirty-minute, hourly, eight hourly and daily series respectively. All Ljung-Box Q-statistics are significant at any conventional level of significance of a χ^2 with the corresponding degrees of freedom. The numbers in rectangular brackets reported in the final column are the marginal significance levels of the Lagrange multiplier tests for the presence of heteroskedasticity, and are obtained by regressing squared returns on a constant and the five most recent lagged squared returns. These statistics follow a χ^2 with 5 degrees of freedom, and again all are significant at any conventional significance level.

Table 3. Comparison of Simulated and Empirical Return Distributions

	GARCH(1,1) (normal)				GARCH(1,1) (Student-*t*)				GARCH -I			
	Mean	St Dev	Skew	Kurt	Mean	St Dev	Skew	Kurt	Mean	St Dev	Skew	Kurt
30 Minutes												
USD/AUD	accept	5%	1%	1%	accept	5%	1%	1%	accept	5%	1%	1%
USD/GBP	accept	5%	1%	1%	accept	accept	1%	1%	accept	5%	1%	1%
DEM/USD	accept	5%	1%	1%	accept	accept	1%	1%	accept	5%	1%	1%
YEN/USD	accept	5%	1%	1%	accept	5%	1%	1%	accept	5%	1%	1%
1 Hour												
USD/AUD	accept	accept	1%	1%	accept	accept	1%	1%	accept	accept	1%	1%
USD/GBP	accept	accept	1%	1%	accept	accept	1%	1%	accept	accept	1%	1%
DEM/USD	accept	accept	1%	1%	accept	accept	1%	1%	accept	accept	1%	1%
YEN/USD	accept	5%	1%	1%	accept	5%	1%	1%	accept	5%	1%	1%
8 Hours												
USD/AUD	accept	accept	1%	1%	accept	accept	1%	5%	accept	accept	1%	1%
USD/GBP	accept	accept	accept	1%	accept	accept	accept	accept	accept	accept	accept	1%
DEM/USD	accept	accept	5%	1%	accept	accept	accept	accept	accept	accept	5%	1%
YEN/USD	accept	accept	1%	1%	accept	accept	1%	5%	accept	accept	1%	1%
24 Hours												
USD/AUD	accept	accept	1%	5%	accept	accept	1%	accept	accept	accept	1%	5%
USD/GBP	accept	accept	accept	accept	accept	accept	accept	accept	accept	accept	accept	accept
DEM/USD	accept	accept	5%	accept	accept	accept	accept	accept	accept	accept	5%	accept
YEN/USD	accept	accept	1%	5%	accept	accept	1%	accept	accept	accept	1%	5%
	EGARCH(1,1)				**HARCH**				**GARCH DIFFUSION**			
	Mean	St Dev	Skew	Kurt	Mean	St Dev	Skew	Kurt	Mean	St Dev	Skew	Kurt
30 Minutes												
USD/AUD	accept	5%	1%	1%	accept	accept	1%	1%	accept	5%	1%	1%
USD/GBP	accept	accept	1%	1%	accept	accept	1%	1%	accept	5%	1%	1%
DEM/USD	accept	accept	1%	1%	accept	accept	1%	1%	accept	5%	1%	1%
YEN/USD	accept	5%	1%	1%	accept	accept	1%	1%	accept	5%	1%	1%
1 Hour												
USD/AUD	accept	accept	1%	1%	accept	accept	1%	1%	accept	1%	1%	1%
USD/GBP	accept	accept	1%	1%	accept	accept	1%	5%	accept	accept	1%	1%
DEM/USD	accept	accept	1%	1%	accept	accept	1%	1%	accept	accept	1%	1%
YEN/USD	accept	5%	1%	1%	accept	accept	1%	1%	accept	5%	1%	1%
8 Hours												
USD/AUD	accept	accept	accept	1%					accept	accept	1%	1%
USD/GBP	accept	accept	accept	1%					accept	accept	accept	1%
DEM/USD	accept	accept	accept	1%					accept	accept	5%	1%
YEN/USD	accept	accept	1%	1%					accept	accept	1%	1%
24 Hours												
USD/AUD	accept	accept	5%	5%					accept	accept	1%	5%
USD/GBP	accept	accept	accept	accept					accept	accept	accept	accept
DEM/USD	accept	accept	accept	accept					accept	accept	accept	accept
YEN/USD	accept	accept	1%	1%					accept	accept	1%	5%

Table 3 presents the results of the Monte Carlo exercise examining whether the unconditional moments of the empirical distribution lie inside the 95^{th} and 99^{th} percentiles of the unconditional simulated moments. In this table, "accept" refers to occasions where the null hypothesis that the simulated and empirical moments are the same cannot be rejected at the 5 percent level; "5%" refers to instances where the null is rejected at the 5 (but not the 1) percent level; and "1%" refers to instances where the null is rejected at the 1 percent level. Blank spaces for the HARCH model refer to instances where the model's parameter values violated the stationarity condition (see the notes under Table 1(b) for a discussion).

Table 4. Descriptive Statistics for Estimates of Daily Realized Variance

	USD/AUD			USD/GBP		
	5-min	30-min	24-hr	5-min	30-min	24-hr
Std Dev	0.064	0.072	0.130	0.019	0.024	0.046
Range	0.601	0.784	1.609	0.193	0.283	0.412
	DEM/USD			JPY/USD		
	5-min	30-min	24-hr	5-min	30-min	24-hr
Std Dev	0.041	0.042	0.068	0.170	0.188	0.255
Range	0.412	0.458	0.746	3.154	3.522	4.039

In Table 4, the column headings "5-min" and "30-min" refer to the construction of daily realized variance from the sum of 288 five-minute and 48 thirty-minute squared intraday returns respectively. "24-hr" refers to the construction of daily realized variance by simply squaring daily returns. All figures are raised by 10^3.

Table 5. Daily Variance Forecasting Performance: Predictive R^2's

		R^2 from equation:	GARCH(1,1) (normal)	GARCH(1,1) (Student-t)	GARCH-I	EGARCH(1,1)	AR	Riskmetrics
USD/AUD	5-min	15	0.340	0.376	0.270	0.196	0.235	0.123
		16	*0.176*	*0.222*	*0.237*	*0.121*	*0.350*	*0.066*
	30-min	15	0.293	0.315	0.254	0.196	0.195	0.137
		16	*0.180*	*0.206*	*0.229*	*0.137*	*0.279*	*0.112*
	24-hr	15	0.002	0.004	0.006	0.003	0.007	0.006
		16	*0.003*	*0.004*	*0.006*	*0.005*	*0.013*	*0.002*
USD/GBP	5-min	15	0.110	0.112	0.108	0.086	0.288	0.080
		16	*0.193*	*0.195*	*0.195*	*0.183*	*0.229*	*0.168*
	30-min	15	0.084	0.087	0.101	0.063	0.257	0.058
		16	*0.165*	*0.166*	*0.179*	*0.145*	*0.249*	*0.139*
	24-hr	15	0.001	0.000	0.020	0.001	0.020	0.002
		16	*0.011*	*0.013*	*0.019*	*0.005*	*0.008*	*0.003*
DEM/USD	5-min	15	0.107	0.117	0.200	0.094	0.517	0.077
		16	*0.255*	*0.243*	*0.420*	*0.178*	*0.254*	*0.296*
	30-min	15	0.065	0.070	0.154	0.058	0.441	0.051
		16	*0.208*	*0.186*	*0.372*	*0.144*	*0.194*	*0.267*
	24-hr	15	0.008	0.008	0.006	0.008	0.008	0.008
		16	*0.008*	*0.008*	*0.009*	*0.008*	*0.007*	*0.007*
JPY/USD	5-min	15	0.612	0.475	0.498	0.238	0.407	0.207
		16	*0.460*	*0.429*	*0.499*	*0.275*	*0.216*	*0.334*
	30-min	15	0.560	0.426	0.437	0.205	0.356	0.175
		16	*0.412*	*0.386*	*0.439*	*0.247*	*0.177*	*0.272*
	24-hr	15	0.035	0.014	0.024	0.008	0.099	0.007
		16	*0.009*	*0.004*	*0.010*	*0.002*	*0.008*	*0.007*

Table 5 reports the R^2 from equations (15) and (16). The R^2 values in the equation (15) rows are based on the regression of realized daily variance on the forecast of daily variance. The R^2 values below in the equation (16) rows (in italics) are based on the regression of the log of realized daily variance on the log variance forecast. The figures in the "5-min" and "30-min" rows are the R^2 where daily realized volatility is defined as the sum of 288 five-minute and 48 thirty-minute squared intraday returns respectively. The figures in the "24-hr" rows are the R^2 where daily realized volatility is defined as the squared daily return.

Table 6. Daily Variance Forecasting Performance: Alternative Criterion

	USD/AUD		USD/GBP		DEM/USD		JPY/USD	
	5-min	30-min	5-min	30-min	5-min	30-min	5-min	30-min
RMSE								
GARCH(1,1) normal	1	1	3	4	3	4	1	1
GARCH(1,1) Student-*t*	4	3	1	1	5	5	3	3
GARCH-I	5	5	6	6	4	3	2	2
EGARCH(1,1)	6	6	5	3	6	6	6	5
AR	2	4	2	5	1	1	5	6
Riskmetrics	3	2	4	2	2	2	4	4
MAE								
GARCH(1,1) normal	3	1	2	3	4	6	1	1
GARCH(1,1) Student-*t*	4	3	1	1	5	5	3	3
GARCH-I	6	6	4	2	3	2	2	2
EGARCH(1,1)	5	4	6	5	6	4	6	5
AR	1	5	5	6	2	3	5	6
Riskmetrics	2	2	3	4	1	1	4	4
MAPE								
GARCH(1,1) normal	2	4	2	5	6	6	3	4
GARCH(1,1) Student-*t*	1	1	1	4	5	5	2	3
GARCH-I	5	3	3	1	2	4	1	1
EGARCH(1,1)	4	2	6	3	4	2	4	2
AR	6	6	5	6	3	1	5	5
Riskmetrics	3	5	4	2	1	3	6	6
MSE								
GARCH(1,1) normal	2	2	1	2	5	6	4	2
GARCH(1,1) Student-*t*	4	1	2	3	6	5	2	3
GARCH-I	6	5	3	1	3	3	1	1
EGARCH(1,1)	5	3	6	6	4	1	6	5
AR	1	6	5	4	2	2	3	4
Riskmetrics	3	4	4	5	1	4	5	6
Theil's U								
GARCH(1,1) normal	1	1	3	4	3	4	1	1
GARCH(1,1) Student-*t*	4	3	2	1	5	5	3	3
GARCH-I	5	5	6	6	4	3	2	2
EGARCH(1,1)	6	6	5	3	6	6	6	5
AR	2	4	1	5	1	1	5	6
Riskmetrics	3	2	4	2	2	2	4	4

Table 6 presents the relative ranking of model forecasting performance for each loss function, with 1 (6) referring to the superior (worst) model. These criteria are the root mean squared error (RMSE), mean absolute error (MAE), mean absolute percent error (MAPE), median squared error (MSE) and Theil's U statistic, and are defined by equations (17), (18), (19) and (20) respectively. The numbers reported under the columns "5-min" and "30-min" refer to instances where realized daily variance is given by the sum of 288 five-minute and 48 thirty-minute squared intraday returns respectively.

Table 7. Profitability Assessment of Variance Forecasts

	Model Used by Investor A	Model Used by Investor B	**Superior Model**	Average Daily Profit of Superior Model	t-stat	% of Days Superior Model Profited
USD/AUD	GARCH(1,1)	AR	**AR**	AUD $ 5.55	(0.64)	52.4
	AR	GARCH-I	**AR**	AUD $ 14.73	(1.62)	54.8
	GARCH(1,1)	GARCH-I	**GARCH(1,1)**	AUD $ 37.80	(4.20)	69.8
USD/GBP	GARCH(1,1)	AR	**GARCH(1,1)**	£ 2.46	(0.87)	54.0
	AR	GARCH-I	**GARCH-I**	£ 7.88	(2.61)	57.9
	GARCH(1,1)	GARCH-I	**GARCH-I**	£ 6.61	(2.04)	57.1
DEM/USD	GARCH(1,1)	AR	**AR**	DM 20.54	(2.52)	59.5
	AR	GARCH-I	**GARCH-I**	DM 21.84	(2.36)	52.4
	GARCH(1,1)	GARCH-I	**GARCH-I**	DM 27.28	(2.90)	65.1
JPY/USD	GARCH(1,1)	AR	**GARCH(1,1)**	¥ 3,636	(2.14)	55.6
	AR	GARCH-I	**GARCH-I**	¥ 7,152	(3.75)	58.7
	GARCH(1,1)	GARCH-I	**GARCH-I**	¥ 4,673	(2.11)	57.1

Table 7 presents the results from the trading game where investors use the AR, GARCH(1,1) or GARCH-I estimator to generate variance forecasts one-day-ahead. The column entitled "Superior Model" reports which of the two competing models generated a positive average daily profit.

Appendix

	GARCH(1,1) (normal)			GARCH(1,1) (Student-*t*)			GARCH-I		
	Skewness	Kurtosis	Q-St	Skewness	Kurtosis	Q-St	Skewness	Kurtosis	Q-St
30 Minutes									
USD/AUD	-0.14 [0.00]	7.72 [0.00]	[0.00]	-0.33 [0.00]	6.87 [0.00]	[0.00]	-0.14 [0.00]	7.72 [0.00]	[0.00]
USD/GBP	-0.14 [0.00]	7.44 [0.00]	[0.00]	-0.17 [0.00]	6.27 [0.00]	[0.00]	-0.14 [0.00]	7.44 [0.00]	[0.00]
DEM/USD	-0.32 [0.00]	12.83 [0.00]	[0.00]	-0.14 [0.00]	11.22 [0.00]	[0.00]	-0.32 [0.00]	12.83 [0.00]	[0.00]
YEN/USD	0.57 [0.00]	41.20 [0.00]	[0.99]	0.63 [0.00]	36.25 [0.00]	[0.99]	0.57 [0.00]	41.20 [0.00]	[0.99]
1 Hour									
USD/AUD	-0.07 [0.00]	6.67 [0.00]	[0.03]	-0.14 [0.00]	6.63 [0.00]	[0.03]	-0.08 [0.00]	6.77 [0.00]	[0.03]
USD/GBP	-0.25 [0.00]	9.31 [0.00]	[0.00]	-0.19 [0.00]	8.51 [0.00]	[0.00]	-0.17 [0.00]	9.30 [0.00]	[0.00]
DEM/USD	-0.28 [0.00]	8.61 [0.00]	[0.00]	-0.26 [0.00]	7.61 [0.00]	[0.00]	-0.06 [0.00]	7.37 [0.00]	[0.00]
YEN/USD	0.61 [0.00]	31.14 [0.00]	[0.99]	0.38 [0.00]	24.81 [0.00]	[0.41]	0.69 [0.00]	34.14 [0.00]	[0.99]
8 Hours									
USD/AUD	0.23 [0.00]	2.16 [0.00]	[0.96]	0.24 [0.00]	2.12 [0.00]	[0.97]	0.18 [0.00]	2.35 [0.00]	[0.93]
USD/GBP	-0.06 [0.17]	3.23 [0.00]	[0.03]	0.00 [0.48]	3.20 [0.00]	[0.03]	-0.05 [0.21]	3.14 [0.00]	[0.03]
DEM/USD	-0.16 [0.01]	2.31 [0.00]	[0.00]	-0.16 [0.01]	2.30 [0.00]	[0.00]	-0.14 [0.01]	2.16 [0.00]	[0.00]
YEN/USD	-0.65 [0.00]	7.85 [0.00]	[0.98]	-0.64 [0.00]	6.96 [0.00]	[0.99]	-0.64 [0.00]	8.04 [0.00]	[0.93]
24 Hours									
USD/AUD	0.18 [0.05]	1.28 [0.00]	[0.75]	0.23 [0.02]	1.18 [0.00]	[0.48]	0.04 [0.36]	1.18 [0.00]	[0.26]
USD/GBP	0.10 [0.18]	1.02 [0.00]	[0.26]	0.10 [0.18]	1.01 [0.00]	[0.22]	0.10 [0.18]	1.01 [0.00]	[0.17]
DEM/USD	-0.20 [0.03]	0.98 [0.00]	[0.82]	-0.24 [0.01]	0.92 [0.00]	[0.74]	-0.25 [0.01]	0.99 [0.00]	[0.82]
YEN/USD	-0.74 [0.00]	2.11 [0.00]	[0.39]	-0.91 [0.00]	1.74 [0.00]	[0.17]	-0.69 [0.00]	2.15 [0.00]	[0.27]

	EGARCH(1,1)			HARCH		
	Skewness	Kurtosis	Q-St	Skewness	Kurtosis	Q-St
30 Minutes						
USD/AUD	-0.19 [0.00]	9.11 [0.00]	[0.00]	0.01 [0.35]	6.99 [0.00]	[0.00]
USD/GBP	-0.14 [0.00]	7.97 [0.00]	[0.00]	-0.12 [0.00]	7.48 [0.00]	[0.00]
DEM/USD	0.19 [0.00]	14.08 [0.00]	[0.01]	-0.33 [0.00]	14.66 [0.00]	[0.00]
YEN/USD	0.49 [0.00]	39.87 [0.00]	[0.58]	-0.07 [0.00]	24.03 [0.00]	[0.42]
1 Hour						
USD/AUD	-0.07 [0.00]	6.98 [0.00]	[0.00]	0.17 [0.00]	7.21 [0.00]	[0.00]
USD/GBP	-0.23 [0.00]	9.29 [0.00]	[0.00]	-0.23 [0.00]	9.11 [0.00]	[0.00]
DEM/USD	-0.02 [0.19]	8.27 [0.00]	[0.00]	-0.22 [0.00]	8.57 [0.00]	[0.00]
YEN/USD	0.10 [0.00]	24.60 [0.00]	[0.04]	-0.76 [0.00]	33.82 [0.00]	[0.26]
8 Hours						
USD/AUD	0.25 [0.00]	2.41 [0.00]	[0.95]			
USD/GBP	-0.08 [0.10]	3.19 [0.00]	[0.02]			
DEM/USD	-0.15 [0.01]	2.15 [0.00]	[0.00]			
YEN/USD	-0.54 [0.00]	7.63 [0.00]	[0.99]			
24 Hours						
USD/AUD	0.21 [0.03]	1.31 [0.00]	[0.85]			
USD/GBP	0.06 [0.29]	0.85 [0.00]	[0.15]			
DEM/USD	-0.23 [0.02]	0.89 [0.00]	[0.76]			
YEN/USD	-0.68 [0.00]	1.83 [0.00]	[0.38]			

The table in the Appendix presents the skewness, excess kurtosis and Ljung Box Q-Statistics for the standardized residuals after fitting the various models to returns measured at different frequencies. The numbers in rectangular brackets to the right of the skewness and kurtosis values are the marginal significance levels of the test of the null hypothesis that the skewness and excess kurtosis values are equal to zero. The numbers in rectangular brackets under the Q-St columns are the marginal significance levels of the Ljung-Box portmanteau tests for autocorrelation in squared standardized residuals with 48, 24, 15 and 10 degrees of freedom for the thirty-minute, hourly, eight hourly and daily series respectively.