

Longer-term Time Series Volatility Forecasts

Louis H. Ederington
University of Oklahoma

Wei Guan
University of South Florida St. Petersburg

November 2008
Initial Draft: July 2005

Keywords: GARCH, EGARCH, Derivatives
JEL Classification: G13, C22, C53

Contact Info: Louis Ederington: Finance Division, Michael F. Price College of Business, University of Oklahoma, 205A Adams Hall, Norman, OK 73019, USA. lederington@ou.edu Phone: (405) 325-5591

Wei Guan: College of Business, University of South Florida St. Petersburg, 140 Seventh Avenue South, St. Petersburg, FL 33701, USA. wguan@stpt.usf.edu. Phone: (727)873-4945

We have benefitted from the comments of workshop participants at the University of Oklahoma, Oklahoma State University, Auburn University, and the University of Queensland as well as participants at the European meetings of the FMA (2006), Southern Finance Association (2007) and American Finance Association (2008). In particular, we thank Binh Do, Steve Swidler, Mark Bertus, Duane Stock, Scott Linn, Tim Krehbiel, Dean Leistkow, and George Jiang for helpful suggestions. Any errors are our responsibility of course.

Abstract

Longer-term Time Series Volatility Forecasts

Option pricing models and longer-term VaR models typically require volatility forecasts over horizons considerably longer than the data frequency. These are generally generated from short-horizon forecasts by successive forward substitution. We document deficiencies with the resulting long-horizon volatility predictions generated by GARCH type models, such as GARCH(1,1), EGARCH, and GJR. Since volatility forecasts for forward periods are functions of forecast volatility for the next period, this recursive procedure keeps the *relative* weights of recent and older observations the same whether forecasting volatility in the near or distant future. In contrast, we find that older observations are relatively more important in forecasting at longer horizons, e.g., more important in forecasting volatility next month than in forecasting volatility tomorrow. We find that the ARLS model of Ederington and Guan (2005) and a modified version of the standard EGARCH model in which the parameter values vary with the forecast horizon forecast better out-of-sample than the GARCH(1,1), EGARCH, and GJR models across a wide variety of markets and forecast horizons.

Longer-term Time Series Volatility Forecasts

I. Introduction

This paper explores the problems that arise when GARCH type time-series models, such as GARCH(1,1), EGARCH, and the GJR (or TGARCH) model, estimated from daily (or higher frequency) data are used to forecast volatility (either the standard deviation, variance, or log of the standard deviation of returns)¹ over the longer horizons common to option valuations and longer term value-at-risk measures. While the GARCH type models generate volatility forecasts for the very next period or observation (normally the next day)², option pricing models and longer horizon value-at-risk measures commonly require volatility forecasts for much longer periods of weeks, months, or even years. These are generally obtained by successive forward substitution in which the volatility forecast for period $t+1$ is used together with the model parameters to forecast volatility for period $t+2$, the forecast for $t+2$ is used to forecast volatility for period $t+3$, etc. These are then combined to obtain the “integrated volatility” forecast for the interval from $t+1$ through $t+N$. Most evaluations of volatility forecasting models in the econometrics literature have focused on their ability to forecast volatility at $t+1$. As Christoffersen and Diebold (2000) have observed, “much less is known about volatility forecastability at longer horizons.”³ We seek to fill this gap.

We argue that a problem with time series volatility forecasts over multi-period horizons is that since the forecast volatility for day (or period) $t+1$ is used to forecast volatility for any future day $t+k$, the *relative*⁴ importance of observed volatility today (t) versus volatility yesterday ($t-1$) or last week ($t-5$) is forced to be the same whether forecasting volatility for tomorrow, next week, or next month. In other words, in the usual procedure, today’s volatility receives the same weighting *relative to* volatility a week ago in forecasting volatility a month from now as it does in forecasting volatility tomorrow. We show that for the GARCH(1,1), EGARCH, and GJR models, the parameters which best forecast volatility for the next day are not those which forecast best over longer horizons - specifically older observations are relatively more important in forecasting volatility at longer horizons. One model in which the relative importance of older and more recent observations varies with the forecast horizon is the Absolute Restricted Least Squares (ARLS) model of Ederington and

Guan (2005). We find that in this model as well the relative importance of older observations increases with the forecast horizon.

We compare the out-of-sample forecasting ability of the GARCH(1,1), EGARCH, GJR, models, regression based modifications of those three models in which the parameter values vary with the forecast horizon, and the ARLS model. Of these seven models, we find that the ARLS and modified EGARCH models forecast best across a wide variety of markets, forecast horizons, and volatility measures. ARLS tends to have a somewhat lower root mean squared forecast error than the modified EGARCH model while modified EGARCH generally has a lower mean absolute error.

Our data set consists of daily data on one equity index market (the S&P 500), two interest rates (3 month T-Bills and 1 year T-Notes), two exchange rates (Yen/Dollar and Dollar/Pound), two commodities (crude oil and gold), and five individual equities from the Dow-Jones index (Caterpillar, Disney, Dupont, GE, and Walmart). Volatility forecasts are examined for horizons of 10, 20, 40, and 80 trading days.

While many volatility forecasting models have been proposed in the econometrics literature, GARCH type models (including Riskmetric's variant of GARCH(1,1)) dominate in practice along with the simple historical standard deviation. We focus on the three which appear to be the most popular: GARCH(1,1), EGARCH, and the GJR (or TGARCH) model of Glosten et al (1993). While not as popular as the other three, we also examine the ARLS model of Ederington and Guan (2005) since it allows older observations to be relatively more important in forecasting at longer horizons. Implied volatilities provide another source of volatility forecasts. While theoretically these should reflect all available information, including time-series information, much evidence indicates that this is not the case.⁵ Moreover implied volatilities cannot simultaneously be used to price the derivatives from whose prices they are calculated and are only available for some assets and some time horizons. Consequently, time-series models remain a major source of volatility forecasts.

The paper is organized as follows. The next section explores how the relative weights attached to recent and older observations in forecasting volatility depend on the forecast horizon in

the context of the GARCH(1,1) model. The same issue is explored in Section III for the GJR, EGARCH and ARLS models. Out-of-sample forecasting ability is compared in IV. Section V concludes the paper.

II. The Forecast Horizon and the Relative Importance of Past Observations in GARCH(1,1)

II.1. Multi-period Horizons Volatility Forecasts

As pointed out by Figlewski (1997) and Christoffersen and Diebold (2000), many uses of volatility forecasts, such as option pricing and longer-term value-at-risk models, require volatility estimates over a much longer horizon than the data frequency used to estimate the model. Typically, models estimated using daily data are used to generate forecasts of volatility for day $t+1$ but for option valuation purposes, what is required is a volatility estimate over the life of the option which may expire months in the future. Sometimes it is simply assumed that the forecast volatility for day $t+1$ will continue through the end of the period. This ignores volatility's mean reverting tendency and as Christoffersen et al (1998) show can lead to serious estimation error. More appropriately and typically, volatility over the longer period is forecast through a recursive procedure in which the volatility forecast for day $t+1$ is used together with the model parameters to forecast volatility on day $t+2$, the forecast for day $t+2$ is used to forecast volatility for day $t+3$, etc. These are then combined to obtain forecast volatility for the period from $t+1$ through $t+N$, a measure which Andersen et al (2006) term "integrated volatility."

To illustrate, consider the GARCH(1,1) model:

$$v_{t+1} = \alpha_0 + \alpha_1 r_t^2 + \beta v_t \quad (1)$$

where r_t is the surprise log return (i.e., $r_t = R_t - E_{t-1}(R_t)$, where $R_t = \ln(P_t/P_{t-1})$ and P_t is the asset price at time t) and v_t is the variance of r_t .⁶ Since $E_t(r_{t+1}^2) = v_{t+1}$, successive forward substitution yields the expression for the expected variance at time $t+k$ based on the forecast for $t+1$:

$$v_{t+k} = \alpha_0 \sum_{j=0}^{k-2} (\alpha_1 + \beta)^j + (\alpha_1 + \beta)^{k-1} v_{t+1} = \alpha_0 \sum_{j=0}^{k-1} (\alpha_1 + \beta)^j + (\alpha_1 + \beta)^{k-1} [\alpha_1 r_t^2 + \beta v_t] \quad (2)$$

While v_{t+1} and v_{t+k} are point volatility estimates, option valuation and longer-term value-at-risk measures require volatility forecasts for a multi-day period, not a single day.⁷ To generate such forecasts, it is normally assumed that returns are independent so the forecast variance for the period (integrated volatility) is obtained by summing or averaging the forecast daily variances. Summing equation 2 from $k=1$ to s and dividing by s yields the integrated volatility forecast V_{t+s} :

$$V_{t+s} = (1/s) \sum_{k=1}^s v_{t+k} = \alpha_s + [\alpha_1 r_t^2 + \beta v_t] \sum_{k=1}^s (\alpha_1 + \beta)^{k-1} \quad (3)$$

where $\alpha_s = (\alpha_0/s) \sum_{k=1}^s \sum_{i=0}^{k-1} (\alpha_1 + \beta)^i$.

II.2. The Relative Importance of Past Observations in GARCH(1,1)

The main issue of this paper is how the relative importance of recent versus older observations in predicting future volatility depends on the forecast horizon. Suppose at the end of trading on a Tuesday, you are forecasting volatility for (1) tomorrow (Wednesday), and (2) Wednesday a week or month forward. Given evidence on volatility persistence, Tuesday's volatility should be more important than Monday's in predicting tomorrow's volatility. But is it much more important than Monday's volatility in predicting volatility a week or month forward? The successive substitution procedure preserves the relative importance of recent and older observations regardless of the forecast horizon, while we hypothesize that differences in relative importance between recent and past observations should decline as the forecast horizon lengthens.

Consider the relative importance of past observations in the GARCH(1,1) model. Since $v_t = \alpha_0 + \alpha_1 r_{t-1}^2 + \beta v_{t-1}$, $v_{t+1} = (\alpha_0 + \beta \alpha_0) + \alpha_1 r_t^2 + \beta \alpha_1 r_{t-1}^2 + \beta^2 v_{t-1}$, and successive substitution back to time $t-J$ yields the alternative expression of the GARCH(1,1) model in equation 1:

$$v_{t+1} = \alpha_0' + \alpha_1 \sum_{j=0}^J \beta^j r_{t-j}^2 \quad (4)$$

where $\alpha_0' = \alpha_0 \sum_{j=0}^J \beta^j + \beta^{J+1} v_{t-J}$. Substituting equation 4 into equation 2, yields:

$$v_{t+k} = \alpha_0 \sum_{j=0}^{k-2} (\alpha_1 + \beta)^j + (\alpha_1 + \beta)^{k-1} \alpha_0' + (\alpha_1 + \beta)^{k-1} \alpha_1 \sum_{j=0}^J \beta^j r_{t-j}^2 \quad (5)$$

As equation 5 makes clear, while the absolute weights decline with the horizon k , assuming $\alpha_1 + \beta < 1$, the relative weights on past squared surprise returns decline at the same exponential rate whether forecasting volatility for tomorrow or for the distant future. Since $\partial v_{t+k} / \partial r_{t-j}^2 = \alpha_1 (\alpha_1 + \beta)^{k-1} \beta^j$, relative partials for past observations m days apart are:

$$\frac{\partial v_{t+k} / \partial r_{t-j-m}^2}{\partial v_{t+k} / \partial r_{t-j}^2} = \beta^m \quad (6)$$

So in GARCH(1,1) forecasts, the relative weights assigned to past observations m days or periods apart are in the ratio β^m regardless of the forecast horizon, k , and how far in the past, j .

While related to the criticism of Engle and Bollerslev (1986), Ding and Granger (1996), Baillie et al (1996), and Bollerslev and Mikkelsen (1996) that the GARCH(1,1) model's memory is too short, it is not the same. There the main issue is how rapidly the impact of a return shock decays in the future. Consider the impact of a squared return shock, r_{t-j}^2 , on volatility forecasts n days apart:

$$\frac{\partial v_{t+k+n} / \partial r_{t-j}^2}{\partial v_{t+k} / \partial r_{t-j}^2} = (\alpha_1 + \beta)^n \quad (7)$$

So the impact of a shock on volatility at $t+k$ declines at the exponential rate $\alpha_1 + \beta$ while there is evidence, e.g., Ding, Granger, and Engle (1993), Ding and Granger (1996), Andersen and Bollerslev (1997) that its impact lasts longer. Note that the weights in equation 6 decline much faster than those in equation 7. For instance, for daily observations on the S&P 500 index from 1/3/1968 to

12/31/2002, GARCH estimates of α_1 and β are .0690 and .9228 respectively so $(\alpha_1 + \beta)^{10} = .9210$ and the reversion of forward forecasts to the mean is fairly slow. At the same time, $\beta^{10} = .4478$ so r_{t-10}^2 (that is an observation 10 days or two weeks ago) receives a weight only 44.8% of that attached to r_t^2 in forecasting any future volatility. A number of long-memory models employ a slower decay rate than the GARCH model's exponential, β^m . However, in any model in which the forecasts for future days are linear functions of the forecast for $t+1$, the relative weights attached to recent and older observations will be the same regardless of the forecast horizon while we would argue that, regardless of the model, older observations should be relatively more important in forecasting volatility in the more distant future.

The relation is basically the same for integrated volatility. Summing equation 5 from $k=1$ to s and dividing by s yields the integrated volatility forecast from $t+1$ through $t+s$, V_{t+s} :

$$V_{t+s} = (1/s) \sum_{k=1}^s v_{t+k} = \alpha_s + \lambda_s \sum_{j=0}^J \beta^j r_{t-j}^2 \quad (8)$$

where

$$\alpha_s = (1/s) \sum_{k=1}^s \left| \alpha_0 \sum_{i=0}^{k-2} (\alpha_1 + \beta)^i + \alpha_0' (\alpha_1 + \beta)^{k-1} \right| \text{ and } \lambda_s = (\alpha_1/s) \sum_{k=1}^s (\alpha_1 + \beta)^{k-1} \quad (9)$$

It is easily seen from equation 8, that the relative impact of r_{t-j}^2 and r_{t-j-m}^2 on the integrated volatility from t to $t+s$ is in the same ratio β^m . That is:

$$\frac{\partial V_{t+s} / \partial r_{t-j-m}^2}{\partial V_{t+s} / \partial r_{t-j}^2} = \beta^m \quad \text{for all } s \quad (10)$$

We hypothesize that the GARCH(1,1) parameter estimates which maximize the likelihood of generating observed returns for $t+1$, e.g., the next day if the model is estimated from daily data, do not forecast volatility very well for $t+k$ when k is large. Specifically, we hypothesize that better long horizon forecasts are obtained if equation 8 is altered to allow β to vary with the forecast horizon s :

$$V_{t+s} = (1/s) \sum_{k=1}^s v_{t+k} = \alpha_s + \lambda_s \sum_{j=0}^J \beta_s^j r_{t-j}^2 \quad (11)$$

We further hypothesize that the β_s which yield the best volatility forecasts increase with the horizon s so that $\partial\beta_s/\partial s > 0$.⁸ The basic idea is simple and intuitive. Suppose at the end of trading on Tuesday we want to forecast volatility for both tomorrow (Wednesday) and Wednesday a week hence. Among other information, we know volatility today (Tuesday) and Tuesday a week ago. For forecasting tomorrow's volatility, today's volatility is likely much more important than volatility a week ago. For forecasting volatility next Wednesday, we expect the difference in the importance of these two past volatilities to be smaller. Therefore, we hypothesize that as the forecast horizon lengthens, better forecasts are obtained by increasing the relative weights on older observations.

One way to avoid this problem is to match the data frequency to the forecast horizon. For example, if the goal is to forecast volatility over the next month, one could use monthly data to estimate the GARCH model and forecast volatility for month $t+1$. But if the forecast horizon is long, the number of observations is sharply reduced and, as Figlewski (1997) points out, convergence often requires a long time series. Moreover, as Andersen et al (1999) show, the resulting volatility forecasts are considerably less accurate than those generated from higher frequency data through the usual forward substitution strategy.

Our argument could explain a surprising finding in past studies. A common alternative to GARCH for option valuation purposes (and the model presented most often in textbooks) is the historical variance or standard deviation. Conceptually, the historical variance would appear decidedly inferior to GARCH(1,1) forecasts in that it weights all included past surprise returns, r_{t-j}^2 , equally (i.e., $\beta=1$). Also it imposes an arbitrary cutoff, date, J . It is hard to believe that the informativeness of the day $t-J$ return is equal to that of day t while there is no information in knowing the return on day $t-J-1$. Yet many studies find that historical volatility forecasts better than GARCH(1,1). In a survey of 39 such studies, Poon and Granger (2003) report that 22 find that historical volatility (including some weighted measures) forecasts actual volatility better. Our hypothesis provides a possible explanation. If the GARCH(1,1) estimate of β is too low when forecasting beyond $t+1$, then the historical variance imposition of $\beta=1$ may be better at long horizons.

II.3. Testing Procedure

To explore whether volatility forecasts can be improved by varying β_s with the forecast horizon, we use nonlinear least squares regressions (NLS) to obtain estimates of β_s (also α_s and λ_s) in equation 11. Letting $AV(s)_t$ represent the actual (or ex post) realized variance over the period from $t+1$ through $t+s$. i.e., $AV(s)_t = (1/s) \sum_{i=1}^s r_{t+i}^2$, equation 11 is estimated by applying least squares estimation to the equation pair:⁹

$$AV(s)_t = \alpha_s + \lambda_s Z_t + \varepsilon_t \quad \text{where} \quad Z_t = \sum_{j=0}^J \beta_s^j r_{t-j}^2 \quad (12)$$

This procedure finds the equation 11 parameter values which minimize the in-sample root mean squared variance forecast errors. Thus we can test whether the parameters which minimize the sum of squared forecast errors of the variance differ from the GARCH(1,1) parameter estimates of equation 1 and vary with the forecast horizon. Criteria other than this quadratic loss function will be considered in Section IV when we compare out-of-sample forecasts.

II.4. Data

Our data set consists of daily log returns on one equity index market (the S&P 500), two interest rates (3 month T-Bills and 1 year T-Notes), two exchange rates (Yen/Dollar and Dollar/Pound), two commodities (crude oil and gold), and five individual equities chosen from the Dow-Jones index (Caterpillar, Disney, Dupont, GE, and Walmart). The series start in January 1968 or the first available date after January 1968 and run through December 2002. Data sources, time periods, and descriptive statistics are reported in Table 1. We consider horizons, s , of 10, 20, 40, and 80 trading days. Ten days is a popular horizon for value-at-risk measures. The other horizons cover expiries of the more heavily traded options. Surprise log returns are estimated as $r_t = R_t - E_{t-1}(R_t)$ where, reflecting the possibility of serially correlated returns, $E_{t-1}(R_t) = \eta_0 + \eta_1 R_{t-1}$. The coefficients η_0 and η_1 in this “mean equation” are estimated simultaneously with the variance equation (e.g. equation 11).¹⁰

II.5. Results

GARCH(1,1) and NLS estimates of β_s in equation 11 are reported in Table 2. The results confirm our argument that better volatility forecasts are obtained by increasing the relative weight on older observations as the forecast horizon lengthens. With the single exception of the S&P500 index at the 10 day horizon, the beta values which minimize the sum of the squared forecast errors are all greater than the GARCH(1,1) estimates. Moreover, in 31 of the 36 adjacent pairs in Table 2, the regression estimates of β_s increase as the horizon s lengthens. Two sets of statistical tests are reported. A $(**)$ on $\hat{\beta}_{40}$ indicates that it is significantly greater than the GARCH(1,1) estimate at at least the .05(.01) level.¹¹ In all markets β_{40} is significantly greater than the GARCH estimate at at least the .05 level. A $(**)$ on $\hat{\beta}_{80}$ indicates that it is significantly greater than the estimated $\hat{\beta}_{10}$ at the .05(.01) level.¹² The difference is significant at at least the .05 level in nine of the twelve markets.

The differences in the implied importance of older versus recent observations is substantial. Consider, for instance, the weight attached to r^2_{t-20} (the observation approximately one month ago) relative to that attached to today's volatility, r^2_t , in generating volatility forecasts. For the GARCH(1,1) model, the average estimated β over our twelve markets is .9091 which translates into a relative weight on r^2_{t-20} only 14.9% of the weight on r^2_t . This is the same regardless of the forecast horizon. By comparison, at the 20-day forecast horizon, the average NLS estimate of β_s over our twelve markets is .968 which implies a relative weight on r^2_{t-20} 52.2% of that on r^2_t . For older observations, the weighting differences are even more stark. Again using the average β estimates across the twelve markets, the implied weight on r^2_{t-41} (approximately two months ago) relative to r^2_t is only 2.01% for GARCH, versus 26.4% for NLS at the 20-day horizon and 36.7% at the 80-day horizon. So while volatility two months or more ago has virtually no impact on the GARCH(1,1) forecasts, there is evidence that it contains considerable information about likely volatility at the longer horizons. Estimates of the other parameters, α_s and λ_s , in equation 11 are available from the authors on request. Since the relative weights on older observations (as represented by β_s), increase

as the forecast horizon s lengthens, we expect the absolute weight on the most recent observation, λ_s , to decline with s and that is generally the case.

III. The Forecast Horizon and the Relative Importance of Past Observations in the GJR and EGARCH Models

III.1 The GJR Model

Since regardless of the model, multi-period volatility forecasts are normally generated by first forecasting volatility for period $t+1$ and then generating forecasts for later periods from the $t+1$ forecast, most time series models keep the relative importance of older and recent observations the same at all forecast horizons. We present evidence for two additional models: the GJR (or TGARCH) model of Glosten et al (1993) and Nelson's (1991) EGARCH model starting with the GJR model. The GJR model adds to the GARCH(1,1) model of equation 1 a term to capture asymmetric volatility: $\alpha_2 D_t r_t^2$ where $D_t=1$ if $r_t < 0$ and $=0$ if $r_t \geq 0$:

$$v_{t+1} = \alpha_0 + \alpha_1 r_t^2 + \alpha_2 D_t r_t^2 + \beta v_t \quad (13)$$

Backward and forward substitution leads to the following expression for volatility at time $t+k$:

$$v_{t+k} = \alpha_k + \alpha_1 \delta_k \sum_{j=0}^J \beta^j r_{t-j}^2 + \alpha_2 \delta_k \sum_{j=0}^J \beta^j D_{t-j} r_{t-j}^2 \quad (14)$$

where $\delta_k = \prod_{j=0}^{k-2} \zeta_{t+j}$ and $\zeta_t = \alpha_1 + \alpha_2 D_t + \beta$.

Consequently, the partials for past observations m periods apart are in the ratio:

$$\frac{\partial v_{t+k} / \partial r_{t-j-m}^2}{\partial v_{t+k} / \partial r_{t-j}^2} = \beta^m \left[\frac{(\alpha_1 + \alpha_2 D_{t-j-m})}{(\alpha_1 + \alpha_2 D_{t-j})} \right] \quad (15)$$

Thus if the shocks at times $t-j-m$ and $t-j$ have the same sign, the relative weights in forecasting future volatility are in the ratio β^m as in the GARCH(1,1) model and in any case do not depend on the forecast horizon k .

III.2. GJR Results

Equation 14 leads to the following expression for integrated volatility:

$$V_{t+s} = (1/s) \sum_{k=1}^s v_{t+k} = \alpha_s + \lambda_{1s} \sum_{j=0}^J \beta^j r_{t-j}^2 + \lambda_{2s} \sum_{j=0}^J \beta^j D_{t-j} r_{t-j}^2 \quad (16)$$

where $\lambda_{1s} = (1/s) \sum_{k=1}^s \alpha_1 \delta_k$ and $\lambda_{2s} = (1/s) \sum_{k=1}^s \alpha_2 \delta_k$. As with the GARCH(1,1) model we hypothesize that volatility forecasts with smaller root mean squared forecast errors may be obtained by allowing the weighting parameter β to vary with the forecast horizon, s , and that the optimal parameter β_s increases with the forecast horizon. To test this, we use NLS to estimate the model:

$$AV(s)_t = \alpha_s + \lambda_{1s} Z_{1t} + \lambda_{2s} Z_{2t} + \varepsilon_t \quad \text{where} \quad Z_{1t} = \sum_{j=0}^J \beta_s^j r_{t-j}^2 \quad \text{and} \quad Z_{2t} = \sum_{j=0}^J \beta_s^j D_{t-j} r_{t-j}^2 \quad (17)$$

for the same twelve markets where again $AV(s)_t$ represents the actual (or ex post) realized variance over the period from $t+1$ through $t+s$. The resulting estimates of β_s are reported in Table 3 along with the usual GJR estimates of β obtained through maximum likelihood estimation of equation 13.

In Table 3, we again report tests of the null that $\beta_{80} \leq \beta_{10}$ and that β_{40} is less than or equal to the usual GJR estimate of equation 13. As one would expect given the similar structures of the GJR and GARCH(1,1) models, the β estimates in Table 3 are close to those in Table 2. In 33 of the 36 pairs in Table 3, the estimated β_s increases as s increases.

III.3. The EGARCH Model

Consider next the EGARCH model, which is of particular interest since it tends to have a longer memory than GARCH(1,1). It has the form:

$$\ln(v_{t+1}) = \alpha_0 + \beta \ln(v_t) + \gamma_1 |r_t/\sigma_t| + \gamma_2 (r_t/\sigma_t) \quad (18)$$

Assuming a normal distribution (as is normally assumed for maximum likelihood estimation of the EGARCH model and in option pricing applications), forward substitution yields:

$$\ln(v_{t+k}) = [\alpha + \gamma_1\sqrt{(2/\pi)}] \sum_{i=0}^{k-2} \beta^i + \beta^{k-1}\ln(v_{t+1}) \quad (19)$$

while backward substitution yields:

$$\ln(v_{t+1}) = \alpha' + \gamma_1 \sum_{j=0}^J \beta^j |r_{t-j}/\sigma_{t-j}| + \gamma_2 \sum_{j=0}^J \beta^j (r_{t-j}/\sigma_{t-j}) \quad (20)$$

Substituting equation 20 into equation 19 yields the following expression for the relative impact of returns (not squared as before) at times t-j and t-j-m on $\ln(v_{t+k})$:

$$\frac{\partial \ln(v_{t+k})/\partial r_{t-j-m}}{\partial \ln(v_{t+k})/\partial r_{t-j}} = \beta^m \left[\frac{(D_{t-j-m}/\sigma_{t-j-m}) + (1/\sigma_{t-j-m})}{(D_{t-j}/\sigma_{t-j}) + (1/\sigma_{t-j})} \right] \quad (21)$$

Where $D_t=1$ if $r_t>0$ and $D_t=-1$ if $r_t<0$.

So if 1) the returns at t-j and t-j-m have the same sign and 2) the conditional volatilities σ_{t-j} and σ_{t-j-m} are equal, then we have a result analogous to that for the GARCH(1,1) and GJR models in that the two partials are in the ratio β^m . More generally, in the EGARCH model the relative impact of past returns on the volatility forecast does not depend on the forecast horizon k.

Defining integrated volatility, V_{t+s} as the geometric average¹³ of volatilities from t+1 through t+s, yields the integrated volatility expression:¹⁴

$$\ln(V_{t+s}) = (1/s) \sum_{k=1}^s \ln(v_{t+k}) = \lambda_{1s} + \lambda_{2s} \sum_{j=0}^J \beta^j |r_{t-j}/\sigma_{t-j}| + \lambda_{3s} \sum_{j=0}^J \beta^j (r_{t-j}/\sigma_{t-j}) \quad (22)$$

III.4 EGARCH Results

As with the other two models, we hypothesize that better volatility forecasts are obtained by allowing the weighting parameter β to vary with the forecast horizon, s, and that the optimal β_s increases with the forecast horizon. To test this we use NLS to estimate the model:

$$\ln(AV(s)_t) = \lambda_{1s} + \lambda_{2s} Z_{2t} + \lambda_{3s} Z_{3t} + \varepsilon_t \text{ where } Z_{2t} = \sum_{j=0}^J \beta_s^j |r_{t-j}/\hat{\sigma}_{t-j}| \text{ and } Z_{3t} = \sum_{j=0}^J \beta_s^j (r_{t-j}/\hat{\sigma}_{t-j}) \quad (23)$$

where $\hat{\sigma}_{t-j}$ is the volatility estimate for day $t-j$.¹⁵ The resulting estimates of β_s are reported in Table 4 along with the usual EGARCH estimates of β obtained through maximum likelihood estimation of equation 18.

As shown in Table 4, estimates of β for the EGARCH model are much higher than those for the GARCH(1,1) and GJR models. Across the twelve markets, $\hat{\beta}$ averages .9795 for EGARCH versus .9091 for GARCH(1,1) implying a much longer memory and also implying that older observations receive a much higher weight (relative to more recent observations) in the EGARCH model. For instance, the .9091 GARCH(1,1) figure implies that the weight on r_{t-20}^2 is only 14.9% of that on r_t^2 while the .9795 EGARCH figure implies a weight on r_{t-20} 66.1% of that on r_t .

Nonetheless, again there is strong evidence that older observations are relatively more important in forecasting volatility at longer horizons than at shorter horizons. In every market except crude oil, the EGARCH estimate of β using the usual maximum likelihood procedure for the next day is less than the regression estimates for longer horizons. In 30 of 36 pairwise cases (the crude oil market is the major exception), the regression estimates of β_s increase as s increases. Across our twelve markets the average $\hat{\beta}_s$ is .9885 for $s=10$ implying a weight on r_{t-40} 63.0% of the weight on r_t . For $s=80$, the average is .9917 implying that r_{t-40} 's weight is 71.8% of r_t 's. In contrast, the usual maximum likelihood EGARCH estimate of .9795 implies that the weight on r_{t-40} is only 43.7% of r_t 's.

III. 5. The ARLS Model

One existing model which allows the relative weights on recent versus older observations to vary with the forecast horizon is the ARLS (Absolute Restricted Least Squares) model suggested by Ederington and Guan (2005). As in GARCH(1,1), in this model the weights on past volatilities decline exponentially and it incorporates mean reversion. However, it models the standard deviation instead of the variance and is based on absolute, not squared, surprise returns.¹⁶

Specifically, the model is:

$$\text{ASD}(s)_t = \alpha_s + \lambda_s \sum_{j=0}^J \sqrt{\pi/2} \beta_s^j |r_{t-j}| \quad (24)$$

where $\text{ASD}(s)_t$ is the standard deviation of returns from $t+1$ to $t+s$. This is structurally identical to the GARCH expression in equation 11 except: (1) the standard deviation replaces the variance on the left-hand side of the equation, (2) the absolute return, $|r_{t-j}|$, replaces r_{t-j}^2 on the right-hand side, (3) the coefficients β_s , α_s and λ_s are allowed to vary with the horizon s and (4) the addition of the term $\sqrt{\pi/2}$. While $E(r^2)=\sigma^2$, under quite general conditions, $E(|r|)$ is distribution dependent. If log returns r_t are normally distributed with mean μ , then $E(|r_t|) = \sigma \sqrt{2/\pi}$ where $r_t = R_t - \mu$, and $E\left(\sqrt{\pi/2} \sum_{j=0}^{n-1} W_j |r_{t-j}|\right) = \sigma$ where $\sum W_j = 1$. Hence the term $\sqrt{\pi/2}$ in equation 24. Obviously, this means that this model presumes that log returns are approximately normally distributed. However this assumption also underlies maximum likelihood estimations of GARCH type models. Moreover most option pricing models (e.g., Black-Scholes and Barone-Adesi-Whaley) and value-at-risk measures assume normality so no new assumption is imposed in these applications. In any case, we test below how well the resulting model forecasts over a broad range of markets which could deviate from normality.

We estimate equation 24 by regressing the ex post standard deviation from $t+1$ to $t+s$ on $|r_{t-j}|$ from $j=0$ to $J=250$.¹⁷ Estimates of α_s , λ_s , and β_s for $s=40$ days are reported in Table 5. As shown there, as compared with the GARCH(1,1) estimates in Table 2, the least squares estimates of β_s for the ARLS model tend to be considerably higher. Estimates of λ_s are considerably lower. This means that (as compared with GARCH(1,1)) older (recent) observations receive considerably more (less) weight in forecasting future volatility over longer horizons. For instance, the average β across our twelve markets is .9637 for the ARLS model with a 40 day horizon versus .9091 for GARCH. Thus in forecasting the standard deviation from $t+1$ to $t+40$, the weight attached to $|r_{t-10}|$ (i.e., 10 days ago) is about 69.1% of the weight attached to $|r_t|$ while in the GARCH model the weight attached to r_{t-10}^2 is only 38.6% of that attached to r_t^2 .

Also reported in Table 5 are estimates of β_s for horizons s of 10, 20, and 80 days. As hypothesized, the estimated β_s generally rises with the forecast horizon s with the exception of the crude oil market. In Figure 1, we graph the coefficients of $|r_{t,j}|$ as a function of the lag j for $s=10$ and 80 days based on the average values of β and λ across our twelve markets. As reflected in Figure 1, the longer the forecast horizon, the greater the importance of older observations.

IV. Out-of-Sample Forecast Accuracy

Next we compare how well the models forecast out-of-sample. For the GARCH(1,1), GJR, and EGARCH models, we estimate both the standard models and the modified versions in equations 12, 17, and 23 respectively which allow the parameter values to vary with the forecast horizon and which are estimated using non-linear least squares. For instance, we estimate separately both the GARCH model in equation 1 and the modified version in equation 12 and use both to forecast the future variance. So in total we compare seven models: GARCH(1,1), GJR, EGARCH, modified least squares versions of each with parameters specific to the forecast horizon and ARLS.

Following tradition in this literature, our primary measure of forecast accuracy is the out-of-sample root mean squared volatility forecast error defined as:

$$\text{RMSE}(s,j,k) = \left[(1/T) \sum_{t=1}^T \text{FE}(s,j,k)_t^2 \right]^{1/2} \quad (25)$$

where $\text{FE}(s, j, k)_t$ is model k 's volatility forecast error for an s -day horizon in market j on day t . One issue is whether to measure volatility as the standard deviation of surprise returns, the variance or some other measure. Poon and Granger (2003) argue that the standard deviation is better than the variance since the variance is more susceptible to outliers and deviations from normality. In addition in Black-Scholes (and most other option pricing models), the option price is an approximately linear function of the standard deviation for near-the-money options and of course value-at-risk measures are linear functions of the standard deviation. As we have seen,

GARCH(1,1) and GJR model the variance, ARLS models the standard deviation, and EGARCH models the log of the standard deviation (or variance).¹⁸ Hence we consider these three volatility measures starting with the standard deviation. The standard deviation forecast error is calculated as $FE(s, j, k)_t = \text{Fore}(s, j, k)_t - \text{Act}(s, j)_t$, where $\text{Fore}(s, j, k)_t$ is the forecast standard deviation (annualized) for the s day horizon from $t + 1$ to $t + s$ in market j using model k , and $\text{Act}(s, j)_t$ is the actual standard deviation observed ex-post.¹⁹ To generate out-of-sample forecasts, the models are estimated using 1260 daily return observations - approximately five years of daily data. To limit the computational burden, parameter values are re-estimated every 40 days.²⁰ In generating multi-day forecasts for the GJR and EGARCH models we assume that positive and negative surprise returns on any future day are equally likely. In generating multi-day forecasts for the standard GARCH, GJR, and EGARCH model, we employ the standard recursive substitution procedure described above. As in the in-sample estimations in Tables 2-5, in the modified models which allow beta to vary with the forecast horizon, the estimated betas generally exceed the estimated betas for the corresponding standard models and rise with the forecast horizon.

The resulting root mean squared forecast errors, $RMSE(s, j, k)$, for forecasts of the standard deviation of returns are reported in Table 6 for all seven forecasting models k , and forecast horizons $s = 10, 20, 40$, and 80 trading days in all twelve markets, j , where we shade the cell of forecasting model with the lowest RMSE for each market and forecast horizon. To provide an indication of how much of the variation in volatility is forecast by the models, we also report the standard deviation of the realized standard deviation. This corresponds to the RMSE for a “naive” model which assumes an unchanging volatility equal to mean volatility for the entire sample.²¹

As shown in Table 6, no one model predicts best in all markets at all forecast horizons but the ARLS and modified EGARCH models clearly dominate. We have twelve markets and four forecast horizons. Of these 48 market/horizon combinations, the ARLS model has the lowest out-of-sample RMSE in 26. The modified EGARCH model, which allows the beta parameter to vary with

the forecast horizon, has the lowest RMSE in 15 market/horizon combinations. No other forecasting model has the lowest RMSE in more than three market/horizon combinations.

Note that while the modified version of the GJR model dominates the standard GJR model in in-sample forecasts, this ranking is generally reversed out-of-sample. The reason for this is readily apparent. The independent variables, Z_{1t} and Z_{2t} in equation 17 are highly correlated by construction. For the 40-day horizon, their average correlation across our markets and samples is .972 and the minimum average correlation is .934 in the gold market. This high multicollinearity results in high standard errors for the two coefficients leading to high forecast errors. In contrast, for the 40-day horizon, the average absolute correlation between the Z_{2t} and Z_{3t} terms for the modified EGARCH model in equation 23 is only .138. As a result, its standard errors and forecasting errors are small and the modified EGARCH model has lower RMSEs than the standard EGARCH model in most markets at most horizons.

In addition to forecasts of the standard deviation we generate forecasts and root mean squared forecast errors for forecasts of the variance of surprise returns and of the natural log of the standard deviation of surprise returns. Note that, as compared to RMSEs for forecasts of the standard deviation, outliers have more impact on the RMSE for forecasts of the variance and less on the RMSE for forecasts of the log of the standard deviation. Summary measures of the relative forecast accuracy for forecasts of the variance and log standard deviation are reported in Table 7 along with results for the standard deviation forecasts. To keep the presentation manageable, in table 7 we report root mean squared forecast errors averaged across the twelve markets. Since $RMSE(s, j, k)$ varies considerably for different markets, j , to weight the markets equally, we first standardize by calculating for each market, horizon, and model, a relative root mean squared forecast error $RRMSE(s, j, k)$ where each model's RMSE is divided by the RMSE for a naive model in which the forecast is a constant equal to the mean of actual volatility over the entire sample. Specifically, $RRMSE(s, j, k) = RMSE(s, j, k)/RMSE(s, j)^*$ where

$$RMSE(s, j)^* = \left\{ \left[\frac{1}{T} \sum_{t=1}^T [Act(s, j)_t - \overline{Act(s, j)}]^2 \right]^{.5} \right\} \text{ and } \overline{Act(s, j)} = \frac{1}{T} \sum_{t=1}^T Act(s, j)_t.$$

For example, for the standard deviation forecasts, $RRMSE(s, j, k)$ is calculated by dividing the figures in columns 3-9 of Table 6 by the figures in column 2. Note that since $\overline{Act(s, j)}$ is an average over the entire sample, it is not out-of-sample like the seven forecasting models. We then average the relative root mean squared forecast errors over the twelve markets:

$$RRMSE(s, k) = (1/12) \sum_{j=1}^{12} RRMSE(s, j, k).$$

Values of $RRMSE(s, k)$ are reported in Table 7 for each model k and horizon s for forecasts of the standard deviation, variance, and natural log of the standard deviation. For each horizon we shade the cell with the lowest $RRMSE(s, k)$ and cross-hatch the second lowest. In general, the ARSL model has the lowest $RRMSE(s, k)$ followed by the modified EGARCH model (in which the values of β are allowed to vary with the forecast horizon). This pattern is weakest for forecasts of the variance.

Of course, the relative root mean squared forecast errors in Table 7 are averages over twelve markets so the figures in Table 7 don't mean that the ARLS and modified EGARCH models are the best in all markets. For forecasts of the variance, the ARLS model has the lowest RMSE in 23 of our 48 market/horizon combinations; the modified EGARCH model in 13, the standard EGARCH model in 6, the standard GJR model in 3, the modified GARCH in 2 and the standard GARCH in 1. For forecasts of the log of the standard deviation, the number of minimum RMSEs for the respective models are: Modified EGARCH: 24, ARLS: 19, standard GARCH: 3, standard EGARCH: 1, and standard GJR: 1.

Finally, We compare the seven models using a different loss criterion, the mean absolute forecast error, $MAE(s, j, k)$, defined as: $MAE(s, j, k) = [(1/T) \sum_{t=1}^T |FE(s, j, k)_t|]$. Summary

measures are presented in Table 8. Since the MAE scale differs across the different markets, we again first calculate each model k 's MAE relative to that of a naive model in which the forecast volatility is a constant equal to mean volatility over the period, i.e., $RMAE(s, j, k) = MAE(s, j,$

$k)/MAE(s, j)^*$ where $MAE(s, j)^*$ is the naive volatility forecast. RMAE's are then averaged over the twelve markets, j , yielding the average RMAE(s, k) reported in Table 8.

The basic result is the same as for the root mean squared forecast error criterion in that the ARLS and modified EGARCH models clearly dominate the other five forecasting models. However, under this criterion the relative ranking of the ARLS and modified EGARCH models is reversed. Whereas the ARLS model tended to have a lower root mean squared forecast error in more markets and horizons, the modified EGARCH models generally has the lower MAE. For forecasts of the standard deviation, of our 48 market/horizon combinations, the modified EGARCH has the lowest MAE in 30, ARLS the lowest in 16, and the standard GARCH model in 2 (both at the 80-day horizon where all models predict poorly). For forecasts of the variance, the modified EGARCH model has the lowest MAE in 32 market/horizon combinations, ARLS in 13 and standard GARCH in 3 (again all at the 80-day horizon). For forecasts of the log of the standard deviation, the modified EGARCH model has the lowest MAE in 26 market/horizon combinations, ARLS in 20 and standard GARCH in 2 (80-day horizon). The finding that the modified EGARCH model tends to have a lower mean absolute forecast error than ARLS but a higher root mean squared forecast error, indicates that a majority of the time the modified EGARCH's forecasts are closer to actual volatility than ARLS's forecast but that it also has more large forecast errors.

In summary, the ARLS model of Ederington and Guan (2005) and the modified EGARCH model developed here clearly dominate the other five models in terms of out-of-sample forecasting ability over multi-period horizons. The latter model is structurally identical to the standard EGARCH model except that the parameters vary with the forecast horizon and it is estimated using non-linear least squares. The ARLS model tends to have lower root mean squared forecast errors in more markets at more horizons while the modified EGARCH model tends to have lower mean absolute forecast errors. However, by both criteria these two forecasting models clearly forecast better over multi-day horizons than the other five models including traditional GARCH, EGARCH, and GJR.

V. Conclusions

Option valuation and value-at-risk applications typically require volatility forecasts over longer horizons than the frequency of the data used to generate the volatility forecast. We have shown that, in GARCH type models, the standard practice of generating these longer horizon forecasts by recursive substitution forces the importance of older observations relative to more recent observations to be the same whether forecasting volatility for the near or distant future. In contrast, we have also shown that older observations are relatively more important in forecasting volatility in the more distant future than in forecasting volatility in the near future.

We estimate modified versions of the GARCH(1,1), EGARCH, and GJR models which allow the relative importance of the weighting parameter to vary with the forecast horizon as well as the ARLS model of Ederington and Guan (2005) which allows the relative importance of older and recent observations to vary with the forecast horizon. Out-of-sample forecast ability is compared over a variety of markets and forecast horizons for forecasts of the standard deviation, variance, and log of the standard deviation of returns. While no one model forecasts best over all markets and horizons, the ARLS and modified EGARCH models (both of which allow relative weights to vary with the forecast horizon) clearly dominate in most markets at most horizons with the ARLS model generally having the lower root mean squared forecast error and modified EGARCH the lower mean absolute forecast error.

Our main points are: 1) that generating longer term volatility forecasts by the standard recursive procedure is an inappropriate procedure, and 2) that volatility forecasting procedures need to allow older observations to be relatively more important in forecasting volatility at more distant horizons.

REFERENCES

- Andersen, T.G., and T. Bollerslev, 1997, Heterogeneous information arrivals and return volatility dynamics: uncovering the long-run in high frequency returns, *Journal of Finance* 52, 975-1005.
- Andersen, T.G., T. Bollerslev, P.F. Christoffersen, and F.X. Diebold, 2006, Volatility and correlation forecasting, *Handbook of Economic Forecasting*, Elliott, Granger, and Timmerman (eds.), Amsterdam, North Holland.
- Andersen, T.G., T. Bollerslev, and S. Lange, 1999, Forecasting financial market volatility: sample frequency vis-a-vis forecast horizon, *Journal of Empirical Finance* 6, 457-477.
- Baillie, R. T., T. Bollerslev, and H. O. Mikkelsen, 1996, Fractionally integrated generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 74, 3-30.
- Bollerslev, T., and H. O. Mikkelsen, 1996, Modeling and pricing long memory in stock market volatility, *Journal of Econometrics* 73, 151-184.
- Christoffersen, P. F., and F. X. Diebold, 2000, How relevant is volatility forecasting for financial risk management?, *Review of Economics and Statistics* 82(1), 12-22.
- Christoffersen, P. F., F. X. Diebold, and T. Schuermann, 1998, Horizon problems and extreme events in financial risk management, *Federal Reserve Bank of New York Economic Policy Review* 4(3), 109-118.
- Ding, Z., and C.W. Granger, 1996, Modeling volatility persistence of speculative returns: A new approach, *Journal of Econometrics* 73, 185-215.
- Ding, Z., C.W. 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.
- Ederington, L. and W. Guan, 2005, Forecasting volatility, *Journal of Futures Markets* 25(5), 465-490.
- Engle, R.F., and T. Bollerslev, 1986, Modeling the persistence of conditional variances, *Econometric Reviews* 5, 1-50.

- Figlewski, S., 1997, *Forecasting volatility*, Financial Markets, Institutions, and Instruments 6, Stern School of Business, (Blackwell Publishers: Boston).
- Glosten, L., R. Jagannathan, and D. Runkle. 1993. On the relation between the expected value and volatility of the nominal excess return on stocks, *Journal of Finance* 48, 1779-1801.
- Hull, J., and A. White, 1987, The pricing of options on assets with stochastic volatilities, *Journal of Finance* 42, 281-300.
- Nelson, D.B., 1991, Conditional heteroskedasticity in asset returns: A new approach, *Econometrica* 59, 347-70.
- Poon, S. H. and C. Granger, 2003, Forecasting volatility in financial markets: A review, *Journal of Economic Literature* 41, 478-539.

ENDNOTES

1. Specifically GARCH(1,1) and GJR forecast the variance of returns and EGARCH the natural log of the variance. These are often converted to standard deviation forecasts. Since the exact volatility measure varies between models, we will use the term “volatility” in a general sense to refer to either the variance, standard deviation, or log of the variance when the implications are the same for all. When a statement applies to a particular volatility measure, such as the variance or standard deviation, we use the more specific term.
2. For expositional ease we will henceforth assume daily data and use the terms “today,” “yesterday,” and “tomorrow” for observations t , $t-1$, and $t+1$. However, our arguments are not frequency specific.
3. An exception to this statement are explorations of the long memory property of many financial series, e.g., Ding, Granger, and Engle (1993), Ding and Granger (1996) and Andersen and Bollerslev (1997) among others. These studies find that absolute and squared returns are more correlated at long lags than GARCH(1,1) predicts they should be - leading to the development of alternatives, such as FIGARCH.
4. As seen below, “relative” is the operative word here. Most models are mean reverting so that the unconditional volatility becomes more important and past volatilities less important as the forecast horizon lengthens. Our issue is with the importance of older observations relative to more recent observations, not the absolute importance of both.
5. For a review of evidence on this issue, see Poon and Granger (2003).
6. Our measure of $E_{t-1}(R_t)$ is explained in section II.4. Of course in the case of dividends or other payments, $R_t = \ln((P_t + D_t)/P_{t-1})$.
7. While the Black-Scholes model assumes constant volatility over the life of the option, this is generally ignored based on the Hull and White (1987) observation that if volatility varies but the volatility of volatility is not priced then the option price is equal to the expected Black-Scholes price based on average volatility.
8. Andersen et al (1999) estimate a continuous time GARCH model in which the values of β vary with the forecast horizon s . However, in their model β varies inversely with the horizon s while we hypothesize and find a positive relation.
9. For these estimations we set $J=250$. Since $\beta < 1$, β^J is nil for high values of J so our estimations are not sensitive to this choice as long as J is fairly high.
10. Estimates of η_0 and η_1 are not reported here but are available on request. In all except the two commodity markets, there is small but positive autocorrelation and in nine of the ten, the estimate of η_1 is significantly greater than zero at the .01 level in the GARCH(1,1) estimation. In the gold market, η_1 is significantly less than zero at the .01 level. However, the explanatory power of this equation is low and the results below are not sensitive to this specification. We have run estimations without the lagged return term R_{t-1} in the mean equation and the results are virtually identical to those presented here.

11. These are based on likelihood ratio tests. First we calculated the log-likelihood for the least squares estimates of equation 12 and again using the GARCH (equation 1) estimates of β allowing α_s and λ_s to take log-likelihood minimizing values. We calculate the difference λ between the two log-likelihoods. Since we utilize daily data but forecast volatility over multi-day periods, s , volatility forecasts on days less than s days apart overlap. Hence the log-likelihoods for observations less than s days apart are not independent. To correct for this we calculate the likelihood ratio as $2\lambda/s$ instead of the usual 2λ . This is a very conservative approach which treats any two observations with any days in common as perfectly correlated. Thus the significance levels in Table 2 are understated, i.e., if the two pairs of beta estimates are significant at the .05 level, they are significant at at least the .05 level.

12. These are based on volatility forecasts for the shorter horizon to minimize the data overlap. In other words the test of the difference in β_{80} and β_{10} is based on forecasts for the 10 day horizon. We estimate the log-likelihood for the 10-day horizon using the parameter estimates for the 10 day horizon and again using β from the 80-day horizon estimation (allowing α_0 and α_1 to take likelihood maximizing values). As explained in footnote 11, due to the data overlap, we calculate the likelihood ratio as $2\lambda/s$ where λ is the difference in the two log-likelihoods and $s=10$.

13. This is different from the usual arithmetic average expression for integrated volatility as utilized above for the GARCH(1,1) and GJR models but this should be a better measure since the return over the forecast horizon period is a compound return.

14. The implied λ parameters are:

$$\lambda_{1s} = (1/s) \sum_{k=1}^s \left[(\alpha + \gamma_1 \sqrt{2/\pi}) \sum_{j=0}^{k-2} \beta^j + \alpha' \beta^{k-1} \right], \quad \lambda_{2s} = (\gamma_1/s) \sum_{k=1}^s \beta^{k-1}, \quad \text{and} \quad \lambda_{3s} = (\gamma_2/s) \sum_{k=1}^s \beta^{k-1}$$

15. Note that this minimizes the root mean squared forecast error of the log of the variance - not the variance itself as with GARCH(1,1) and GJR.

16. In an earlier version of the paper, we included the AGARCH model among those analyzed since it is also based on absolute, not squared returns. However, since its forecasting ability proved weaker than GARCH(1,1), GJR, EGARCH, and ARLS so we have dropped it to save space.

17. Procedurally we first generate the series $W(\beta)_t = (\pi/2)^{-.5} \sum_{j=0}^{\infty} \beta^j |r_{t-j}|$ for $j=0$ to 250 using values of β from .500 through 1.000 in increments of .0001, then regress $ASD(s)_t$ on $W(\beta)_t$ using OLS, repeat the regression for all values of β , and choose the values of β , α , and λ for the regression resulting in the lowest residual sum of squares.

18. Of course results for the log of the standard deviation also hold for the log of the variance since the latter is 2 times the former.

19. In a few instances the two GJR models (and very rarely the GARCH models) forecast a negative variance. When this occurred it was replaced by the minimum variance over the last five years. Also if the forecast standard deviation was more than double the maximum standard deviation over the last five years, it was replaced with two times the maximum standard deviation over the last five years.

20. For instance, the models are first estimated using observations 1 through 1260 and these parameter estimates are used to generate volatility forecasts for the period $t+1$ (1261) through $t+s$ ($1260+s$) where s is the forecast horizon. The same parameters but return observations through day $t+1$ are used to generate volatility forecasts for the period $t+2$ through $t+s+1$, and so forth so that the forecasts for the period $t+40$ through $t+s+39$ are generated using the same parameters but return observations through day $t+39$. Then the models are re-estimated using data from day 41 through day 1300.

21. This model assumes that the mean is known so is not purely “out-of-sample” like the other seven models.

Figure 1 - Relative Weights on Past Return Observations in the ARLS Model For 10 and 80 Day Forecast Horizons. Weights Based on Average Parameters for Twelve Markets




Table 1 - Markets and Descriptive Statistics					
Market or Index	Data Source	Data Period	Daily Obs.	Daily Returns (x 1000)	
				Mean	Standard Deviation
S&P 500	CRSP	1/3/68 - 12/31/02	8,809	0.2514	9.9407
TBill	Federal Reserve	1/3/68 - 12/31/02	8,809	-0.1629	14.2690
TNote	Federal Reserve	1/3/68 - 12/31/02	8,809	-0.1663	12.7353
Yen/Dollar	Federal Reserve	1/5/71 - 12/31/02	8,078	-0.1303	6.4371
Dollar/Pound	Federal Reserve	1/5/71 - 12/31/02	8,078	-0.0499	5.8547
Crude Oil	COMEX	3/31/83 - 12/31/02	4,954	0.0108	24.3518
Gold	NYMEX	1/3/75 - 12/31/02	7,026	0.0964	12.6908
Caterpillar	CRSP	4/2/71 - 12/31/02	8,017	0.3182	18.4326
Disney	CRSP	1/3/68 - 12/31/02	8,809	0.5134	21.3474
DuPont	CRSP	1/3/68 - 12/31/02	8,809	0.3410	16.3096
GE	CRSP	1/3/68 - 12/31/02	8,809	0.4865	15.9671
Walmart	CRSP	11/21/72 - 12/31/02	7,602	0.8981	21.5836

Table 2 - The Impact of the Forecast Horizon on Estimates of the GARCH(1,1) model.

GARCH and regression estimates of β_s in the equation:

$$V_{t+s} = (1/s) \sum_{k=1}^s v_{t+k} = \alpha_s + \lambda_s \sum_{i=0}^J \beta_s^i r_{t-j}^2$$

are reported where V_{t+s} is the variance of surprise returns from $t+1$ to $t+s$ and r_{t-j} is the surprise return on day $t-j$. The reported "GARCH" parameters are those estimated using the normal procedure which chooses parameter values to maximize the likelihood of observing r_{t+1} .

Regression estimates are obtained by estimating the above equation using non-linear least squares. Separate regressions are estimated for horizons, s , of 10, 20, 40, and 80 trading days.

*(**) on the 40-day β estimate indicates that it is significantly greater than the GARCH estimate at the .05 (.01) level based on likelihood ratio tests adjusted for the data overlap. *(**) on the 80-day β estimate indicates that it is significantly greater than the 10-day estimate at the .05 (.01) level.

Market	GARCH	Least Squares Regression Estimates by Horizon			
		10 day	20 day	40 day	80 day
S&P 500	.923	.852	.964	.968	.984**
TBill	.827	.943	.976	.985**	.989**
TNote	.913	.981	.984	.984**	.988*
Yen/Dollar	.863	.969	.972	.980**	.980
Dollar/Pound	.902	.941	.945	.962**	.972**
Crude Oil	.918	.962	.965	.964*	.957
Gold	.905	.903	.939	.953**	.948**
Caterpillar	.952	.969	.980	.984*	.989**
Disney	.886	.964	.979	.982**	.990**
Dupont	.953	.972	.982	.988**	.992**
GE	.946	.953	.972	.978**	.986**
Walmart	.921	.955	.955	.950*	.950

Table 3 - The Impact of the Forecast Horizon on Estimates of the GJR model

GJR and regression estimates of β_s in the equation :

$$V_{t+s} = (1/s) \sum_{k=1}^s v_{t+k} = \alpha_s + \lambda_s \sum_{j=0}^J \beta_s^j r_{t-j}^2 + \delta_s \sum_{j=0}^J \beta_s^j D_{t-j} r_{t-j}^2$$

are reported where V_{t+s} is the variance of surprise returns from $t+1$ to $t+s$, r_{t-j} is the surprise return on day $t-j$ and $D_{t-j}=1$ if $r_{t-j}<0$ and 0 otherwise. The reported "GJR" parameters are those estimated using the normal procedure which chooses parameter values to maximize the likelihood of observing r_{t+1} . Regression estimates are obtained by estimating the above equation using non-linear least squares. Separate regressions are estimated for horizons, s , of 10, 20, 40, and 80 trading days. (**) on the 40-day β estimate indicates that it is significantly greater than the GJR estimate at the .05(.01) level based on likelihood ratios adjusted for the data overlap. (**) on the 80-day estimate indicates that it is significantly greater than the 10-day estimate at the .05(.01) level.

Market	GJR	Least Squares Regression Estimates by Horizon			
		10 day	20 day	40 day	80 day
S&P 500	.924	.863	.961	.966	.987**
TBill	.822	.944	.973	.985**	.988**
TNote	.914	.983	.987	.987**	.990*
Yen/Dollar	.856	.965	.966	.973**	.973
Dollar/Pound	.899	.941	.946	.962**	.973**
Crude Oil	.900	.964	.969	.970**	.968
Gold	.904	.936	.961	.969**	.970**
Caterpillar	.942	.968	.981	.985**	.988**
Disney	.892	.965	.979	.982**	.992**
Dupont	.946	.972	.982	.988**	.994**
GE	.943	.951	.971	.978**	.987**
Walmart	.922	.957	.958	.952*	.949

Table 4 - The Impact of the Forecast Horizon on Estimates of the EGARCH model

EGARCH and regression estimates of β_s in the equation :

$$\ln(V_{t+s}) = \lambda_{1s} + \lambda_{2s} \sum_{j=0}^J \beta_s^j |r_{t-j}/\sigma_{t-j}| + \lambda_{3s} \sum_{j=0}^J \beta_s^j (r_{t-j}/\sigma_{t-j})$$

are reported where V_{t+s} is the variance of surprise returns from $t+1$ to $t+s$ and r_{t-j} is the surprise return on day $t-j$. The reported "EGARCH" parameters are those estimated using the normal procedure which chooses parameter values to maximize the likelihood of observing r_{t+1} . Regression estimates are obtained by estimating the above equation using non-linear least squares. Separate regressions are estimated for horizons, s , of 10, 20, 40, and 80 trading days. (**) on the 40-day β estimate indicates that it is significantly greater than the EGARCH estimate at the .05(.01) level. (***) on the 80-day estimate indicates that it is significantly greater than the 10-day estimate at the .05(.01) level.

Market	EGARCH	Least Squares Regression Estimates by Horizon			
		10 day	20 day	40 day	80 day
S&P 500	.9853	.9902	.9908	.9917**	.9937**
TBill	.9678	.9828	.9857	.9885**	.9897**
TNote	.9828	.9897	.9903	.9915**	.9922*
Yen/Dollar	.9495	.9798	.9817	.9843**	.9953**
Dollar/Pound	.9592	.9782	.9798	.9805**	.9779
Crude Oil	.9895	.9913	.9908	.9899	.9887**
Gold	.9829	.9895	.9906	.9922**	.9939**
Caterpillar	.9890	.9936	.9939	.9939**	.9944
Disney	.9859	.9915	.9919	.9916**	.9927
Dupont	.9919	.9958	.9966	.9965**	.9969
GE	.9893	.9922	.9926	.9935**	.9953**
Walmart	.9809	.9877	.9886	.9894**	.9901

Table 5 - ARLS Model Parameter Estimates

Regression estimates for the ARLS volatility forecasting model:

$$ASD(s)_t = \alpha_s + \lambda_s \sum_{j=0}^J \sqrt{\pi/2} \beta_s^j |r_{t-j}|$$

are reported where $ASD(s)_t$ is the standard deviation of returns over the period from $t+1$ to $t+s$ and r_{t-j} is the return on day $t-j$ and $s=40$ trading days. We report estimates of all three parameters for a 40-day horizon and estimates of β_s for horizons of 10, 20, 40, and 80 days.

Market	Parameter Estimates for a 40-day forecast horizon			Estimates of β_s for different forecasting horizons		
	α_s	λ_s	β_s	10-day	20-day	80-day
S&P 500	0.00325	0.0371	0.9443	0.9280	0.9397	0.9809
TBill	0.00381	0.0211	0.9729	0.9140	0.9360	0.9822
TNote	0.00259	0.0259	0.9695	0.9542	0.9603	0.9794
Yen/Dollar	0.00214	0.0214	0.9676	0.9465	0.9521	0.9768
Dollar/Pound	0.00283	0.0283	0.9578	0.9376	0.9479	0.9655
Crude Oil	0.00457	0.0457	0.9422	0.9452	0.9442	0.9336
Gold	0.00306	0.0306	0.9623	0.9301	0.9522	0.9636
Caterpillar	0.00181	0.0181	0.9763	0.9727	0.9755	0.9793
Disney	0.00245	0.0245	0.9663	0.9632	0.9687	0.9766
Dupont	0.00145	0.0145	0.9824	0.9684	0.9771	0.9875
GE	0.00257	0.0257	0.9668	0.9526	0.9615	0.9797
Walmart	0.00283	0.0283	0.9557	0.9510	0.9555	0.9647

Table 6 - Forecast Accuracy - Standard Deviation

Root mean squared forecast errors (RMSE) for out-of-sample forecasts of the standard deviation of returns are reported for seven volatility forecasting models: ARLS, GARCH(1,1), EGARCH, GJR and modifications of the last three that allow the parameters to vary with the forecast horizon and are estimated using least squares regression. For comparison, the standard deviation of the realized standard deviation over the entire sample is reported in column 2. Forecast standard deviations are generated for horizons of 10, 20, 40, and 80 trading days. For each market and forecast horizon, the cell with the lowest RMSE is shaded.

Market	Std. Dev.	ARLS	GARCH		EGARCH		GJR	
			standard	modified	standard	modified	standard	modified
Panel A: 10 Day forecast horizon								
S&P 500	0.0775	0.0665	0.0684	0.0727	0.0670	0.0657	0.0681	0.0739
TBill	0.1277	0.1037	0.1117	0.1130	0.1132	0.1083	0.1142	0.1195
TNote	0.1044	0.0834	0.0884	0.0895	0.0870	0.0855	0.0901	0.0939
Yen/Dollar	0.0452	0.0415	0.0440	0.0427	0.0429	0.0440	0.0442	0.0437
Dollar/Pound	0.0417	0.0346	0.0357	0.0356	0.0356	0.0362	0.0365	0.0372
Crude Oil	0.2178	0.1785	0.1914	0.2056	0.1844	0.1865	0.1931	0.2355
Gold	0.0896	0.0703	0.0785	0.0781	0.0754	0.0698	0.0768	0.0809
Caterpillar	0.1292	0.1159	0.1221	0.1198	0.1196	0.1186	0.1224	0.1215
Disney	0.1493	0.1340	0.1391	0.1429	0.1367	0.1343	0.1436	0.1473
Dupont	0.1051	0.0908	0.0970	0.0939	0.0962	0.0918	0.0958	0.0948
GE	0.1086	0.0889	0.0887	0.0925	0.0880	0.0872	0.0882	0.0935
Walmart	0.1229	0.1059	0.1069	0.1065	0.1068	0.1101	0.1069	0.1120

Table 6 continued - Panel B: 20 Day forecast horizon

Market	Std. Dev.	ARLS	GARCH		EGARCH		GJR	
			standard	modified	standard	modified	standard	modified
S&P500	0.0722	0.0620	0.0639	0.0688	0.0629	0.0614	0.0640	0.0688
TBill	0.1183	0.0975	0.1061	0.1059	0.1040	0.1000	0.1081	0.1125
TNote	0.0964	0.0767	0.0804	0.0821	0.0794	0.0773	0.0821	0.0832
Yen/Dollar	0.0389	0.0359	0.0380	0.0370	0.0366	0.0392	0.0383	0.0384
Dollar/Pound	0.0373	0.0308	0.0317	0.0317	0.0314	0.0324	0.0327	0.0342
Crude Oil	0.2008	0.1610	0.1718	0.1945	0.1627	0.1661	0.1726	0.2312
Gold	0.0830	0.0632	0.0732	0.0720	0.0690	0.0632	0.0715	0.0740
Caterpillar	0.1130	0.0978	0.1060	0.1025	0.1032	0.1009	0.1059	0.1107
Disney	0.1343	0.1199	0.1270	0.1290	0.1239	0.1196	0.1315	0.1325
Dupont	0.0921	0.0767	0.0838	0.0801	0.0833	0.0772	0.0823	0.0820
GE	0.0990	0.0799	0.0783	0.0856	0.0777	0.0770	0.0778	0.0900
Walmart	0.1084	0.0924	0.0931	0.0919	0.0931	0.0964	0.0931	0.0997
Panel C: 40 day forecast horizon								
S&P 500	0.0674	0.0599	0.0613	0.0671	0.0606	0.0585	0.0614	0.0722
TBill	0.1091	0.0937	0.1027	0.1019	0.0996	0.0924	0.1045	0.1084
TNote	0.0897	0.0745	0.0737	0.0774	0.0760	0.0730	0.0754	0.0786
Yen/Dollar	0.0338	0.0334	0.0345	0.0336	0.0328	0.0360	0.0349	0.0355
Dollar/Pound	0.0333	0.0278	0.0283	0.0296	0.0280	0.0300	0.0293	0.0340
Crude Oil	0.1881	0.1496	0.1627	0.1844	0.1503	0.1566	0.1589	0.2185
Gold	0.0768	0.0583	0.0705	0.0666	0.0656	0.0587	0.0692	0.0751

Table 6 continued - Panel C - 40 Day horizon continued

Market	Std. Dev.	ARLS	GARCH		EGARCH		GJR	
			standard	modified	standard	modified	standard	modified
Caterpillar	0.1002	0.0850	0.0943	0.0898	0.0907	0.0884	0.0932	0.0978
Disney	0.1205	0.1095	0.1171	0.1193	0.1121	0.1081	0.1203	0.1217
Dupont	0.0822	0.0659	0.0751	0.0696	0.0744	0.0674	0.0728	0.0717
GE	0.0916	0.0756	0.0716	0.0821	0.0723	0.0708	0.0710	0.0883
Walmart	0.0972	0.0860	0.0850	0.0855	0.0852	0.0857	0.0850	0.0943
Panel D: 80 Day forecast horizon								
S&P 500	0.0634	0.0595	0.0601	0.0695	0.0602	0.0575	0.0599	0.0848
TBill	0.0995	0.0912	0.0994	0.1002	0.0981	0.0864	0.1017	0.1047
TNote	0.0844	0.0749	0.0713	0.0762	0.0777	0.0767	0.0739	0.0782
Yen/Dollar	0.0291	0.0321	0.0304	0.0312	0.0298	0.0353	0.0308	0.0343
Dollar/Pound	0.0299	0.0283	0.0269	0.0316	0.0266	0.0298	0.0278	0.0372
Crude Oil	0.1768	0.1438	0.1705	0.1720	0.1549	0.1558	0.1620	0.2290
Gold	0.0709	0.0558	0.0693	0.0651	0.0655	0.0572	0.0689	0.0763
Caterpillar	0.0904	0.0796	0.0862	0.0846	0.0826	0.0824	0.0846	0.0935
Disney	0.1087	0.1041	0.1143	0.1132	0.1066	0.1034	0.1208	0.1163
Dupont	0.0752	0.0604	0.0699	0.0640	0.0697	0.0631	0.0674	0.0704
GE	0.0860	0.0804	0.0685	0.0884	0.0712	0.0684	0.0675	0.0905
Walmart	0.0878	0.0823	0.0826	0.0844	0.0820	0.0781	0.0818	0.0901

Table 7 - Relative Root Mean Squared Forecast Errors

Relative root mean squared forecast errors (RRMSE) for seven volatility forecasting models are reported. For each market and forecast horizon, relative root mean squared forecast errors are calculated as the ratio of the root mean squared forecast error for each model relative to the root mean squared error for a naive model in which forecast volatility is a constant equal to mean volatility over the entire period. RRMSEs are then averaged across the twelve markets listed in previous tables and reported for forecasts of the standard deviation, variance, and log standard deviation of surprise returns in panels A, B, and C respectively. For each horizon, the lowest mean RRMSE is shaded and the next lowest cross-hatched.

Forecast Horizon	ARLS	GARCH		EGARCH		GJR	
		standard	modified	standard	modified	standard	modified
Panel A: Forecasts of the standard deviation of surprise returns							
10 day horizon	0.8465	0.8895	0.8993	0.8759	0.8652	0.8948	0.9359
20 day horizon	0.8367	0.8853	0.8989	0.8652	0.8542	0.8908	0.9504
40 day horizon	0.8520	0.9008	0.9194	0.8779	0.8622	0.9026	0.9956
80 day horizon	0.9088	0.9457	0.9870	0.9293	0.9141	0.9478	1.1025
All horizons	0.8610	0.9053	0.9262	0.8871	0.8739	0.9090	0.9961
Panel B: Forecasts of the variance of surprise returns							
10 day horizon	.9238	.9662	.9408	.9500	.9554	.9958	.9671
20 day horizon	.9084	.9734	.9428	.9300	.9310	1.0008	.9787
40 day horizon	.9097	.9781	.9556	.9177	.9119	.9940	1.0161
80 day horizon	.9453	1.0034	1.0138	.9431	.9397	1.0103	1.1442
All horizons	.9218	.9803	.9633	.9352	.9345	1.0002	1.0265
Panel C: Forecasts of the log of the standard deviation of surprise returns							
10 day horizon	.8285	.8890	.9268	.8842	.8231	.8859	.9981
20 day horizon	.8104	.8690	.9063	.8649	.8174	.8678	1.0082
40 day horizon	.8302	.8846	.9200	.8837	.8442	.8856	1.0445
80 day horizon	.9041	.9343	1.0039	.9425	.9130	.9405	1.1705
All horizons	.8433	.8942	.9392	.8938	.8494	.8949	1.0553

Table 8 - Relative Mean Absolute Forecast Errors

Relative mean absolute forecast errors (RMAE) for seven volatility forecasting models are reported. For each market and horizon, relative mean absolute forecast errors are calculated as the ratio of the mean absolute forecast error for each model relative to the mean absolute error for a naive model in which forecast volatility is a constant equal to the mean volatility over the entire period. RMAEs are then averaged across the twelve markets listed in previous tables and reported for forecasts of the standard deviation, variance, and log standard deviation of surprise returns in panels A, B, and C respectively. For each horizon, the lowest mean RMAE is shaded and the next lowest cross-hatched.

Forecast Horizon	ARLS	GARCH		EGARCH		GJR	
		standard	modified	standard	modified	standard	modified
Panel A: Forecasts of the standard deviation of surprise returns							
10 day horizon	0.8089	0.8788	0.9159	0.8692	0.7954	0.8761	0.9484
20 day horizon	0.7882	0.8579	0.8966	0.8497	0.7773	0.8562	0.9370
40 day horizon	0.8105	0.8760	0.9132	0.8736	0.7970	0.8750	0.9743
80 day horizon	0.8730	0.9245	0.9753	0.9359	0.8635	0.9235	1.0618
All horizons	0.8201	0.8843	0.9252	0.8821	0.8083	0.8827	0.9804
Panel B: Forecasts of the variance of surprise returns							
10 day horizon	.7656	.8394	.8599	.8201	.7522	.8451	.8841
20 day horizon	.7557	.8318	.8582	.8096	.7416	.8365	.8892
40 day horizon	.7857	.8544	.8858	.8362	.7597	.8566	.9419
80 day horizon	.8482	.9060	.9541	.9015	.8207	.9045	1.0444
All horizons	.7891	.8579	.8895	.8419	.7686	.8607	.9399
Panel C: Forecasts of the log of the standard deviation of returns							
10 day horizon	.8288	.8918	.9327	.8890	.8194	.8873	.9896
20 day horizon	.8058	.8700	.9105	.8705	.8021	.8671	.9783
40 day horizon	.8243	.8859	.9215	.8931	.8249	.8849	1.0067
80 day horizon	.8905	.9325	.9860	.9529	.8984	.9340	1.0986
All horizons	.8374	.8950	.9377	.9014	.8362	.8933	1.0183