

What good is a volatility model?

Robert F Engle and Andrew J Patton

Department of Finance, NYU Stern School of Business and Department of Economics, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0508, USA

Received 15 October 2000

Abstract

A volatility model must be able to forecast volatility; this is the central requirement in almost all financial applications. In this paper we outline some stylized facts about volatility that should be incorporated in a model: pronounced persistence and mean-reversion, asymmetry such that the sign of an innovation also affects volatility and the possibility of exogenous or pre-determined variables influencing volatility. We use data on the Dow Jones Industrial Index to illustrate these stylized facts, and the ability of GARCH-type models to capture these features. We conclude with some challenges for future research in this area.

1. Introduction

A volatility model should be able to forecast volatility. Virtually all the financial uses of volatility models entail forecasting aspects of future returns. Typically a volatility model is used to forecast the absolute magnitude of returns, but it may also be used to predict quantiles or, in fact, the entire density. Such forecasts are used in risk management, derivative pricing and hedging, market making, market timing, portfolio selection and many other financial activities. In each, it is the predictability of volatility that is required. A risk manager must know today the likelihood that his portfolio will decline in the future. An option trader will want to know the volatility that can be expected over the future life of the contract. To hedge this contract he will also want to know how volatile is this forecast volatility. A portfolio manager may want to sell a stock or a portfolio before it becomes too volatile. A market maker may want to set the bid–ask spread wider when the future is believed to be more volatile.

There is now an enormous body of research on volatility models. This has been surveyed in several articles and continues to be a fruitful line of research for both practitioners and academics. As new approaches are proposed and tested, it is helpful to formulate the properties that these models should satisfy. At the same time, it is useful to discuss properties that standard volatility models do not appear to satisfy.

We will concern ourselves in this paper only with the volatility of univariate series. Many of the same issues will arise in multivariate models. We will focus on the volatility of

asset returns and consequently will pay very little attention to expected returns.

1.1. Notation

First we will establish notation. Let P_t be the asset price at time t and $r_t = \ln(P_t) - \ln(P_{t-1})$ be the continuously compounded return on the asset over the period $t - 1$ to t . We define the conditional mean and conditional variance as:

$$m_t = E_{t-1}[r_t] \quad (1)$$

$$h_t = E_{t-1}[(r_t - m_t)^2] \quad (2)$$

where $E_{t-1}[u]$ is the expectation of some variable u given the information set at time $t - 1$ which is often denoted $E[u|\mathfrak{S}_{t-1}]$. Without loss of generality this implies that R_t is generated according to the following process:

$$R_t = m_t + \sqrt{h_t}\varepsilon_t, \quad \text{where } E_{t-1}[\varepsilon_t] = 0 \text{ and } V_{t-1}[\varepsilon_t] = 1.$$

In this paper we are often concerned with the conditional variance of the process and the distribution of returns. Clearly the distribution of ε is central in this definition. Sometimes a model will assume:

$$\{\varepsilon_t\} \sim \text{i.i.d.} F() \quad (3)$$

where F is the cdf of ε .

We can also define the unconditional moments of the process. The mean and variance are naturally defined as

$$\mu = E[r_t], \quad \sigma^2 = E[(r_t - \mu)^2] \quad (4)$$

and the unconditional distribution is defined as

$$(r_t - \mu)/\sigma \sim G \quad (5)$$

where G is the cdf of the normalized returns.

A model specified as in equations (1), (2) and (2) will imply properties of (4) and (5) although often with considerable computation. A complete specification of (4) and (5) however, does not imply conditional distributions since the degree of dependence is not formulated. Consequently, this does not deliver forecasting relations. Various models for returns and volatilities have been proposed and employed. Some, such as the GARCH type of models, are formulated in terms of the conditional moments. Others, such as stochastic volatility models, are formulated in terms of latent variables which make it easy to evaluate unconditional moments and distributions but relatively difficult to evaluate conditional moments. Still others, such as multifractals or stochastic structural break models, are formulated in terms of the unconditional distributions. These models often require reformulation to give forecasting relations.

Higher moments of the process often figure prominently in volatility models. The unconditional skewness and kurtosis are defined as usual by

$$\xi = \frac{E[(r_t - \mu)^3]}{\sigma^3}, \quad \zeta = \frac{E[(r_t - \mu)^4]}{\sigma^4}. \quad (6)$$

The conditional skewness and kurtosis are similarly defined

$$s_t = \frac{E_{t-1}[(r_t - m_t)^3]}{h_{t-1}^{3/2}}, \quad k_t = \frac{E_{t-1}[(r_t - m_t)^4]}{h_{t-1}^2}. \quad (7)$$

Furthermore, we can define the proportional change in conditional variance as

$$\text{variance return} = \frac{h_t - h_{t-1}}{h_{t-1}}. \quad (8)$$

Some of the *variance return* is predictable and some is an innovation. The volatility of the variance (VoV) is therefore the standard deviation of this innovation. This definition is directly analogous to price volatility:

$$\text{VoV} = \sqrt{V(\text{variance return})}. \quad (9)$$

A model will also generate a term structure of volatility. Defining $h_{t+k|t} \equiv E_t[r_{t+k}^2]$, the term structure of volatility is the forecast standard deviation of returns of various maturities, all starting at date t . Thus for an asset with maturity at time $t+k$, this is defined as

$$v_{t+k|t} \equiv \sqrt{V_t \left(\sum_{j=1}^k r_{t+j} \right)} \cong \sqrt{\sum_{j=1}^k E_t \left(r_{t+j}^2 \right)}. \quad (10)$$

The term structure of volatility summarizes all the forecasting properties of second moments. From such forecasts, several specific features of volatility processes are easily defined.

1.2. Types of volatility models

There are two general classes of volatility models in widespread use. The first type formulates the conditional variance directly as a function of observables. The simplest examples here are the ARCH and GARCH models which will be discussed in some detail in section 3.

The second general class formulates models of volatility that are not functions purely of observables. These might be called latent volatility or (misleadingly) stochastic volatility models. For example, a simple stochastic volatility specification is:

$$\begin{aligned} r_t &= m_t + \sqrt{v_t} \varepsilon_t \\ v_t &= \omega v_{t-1}^\beta \exp(\kappa \eta_t) \\ \varepsilon_t, \eta_t &\sim \text{in}(0, 1). \end{aligned}$$

Notice that v is not simply a function of elements of the information set of past returns, and therefore it cannot be the conditional variance or the one step variance forecast. Intuitively, this happens because there are two shocks and only one observable so that current and past v are never observed precisely. The conditional variance in this model is well defined but difficult to compute since it depends on a nonlinear filtering problem defined as (2).

Latent volatility models can be arbitrarily elaborate with structural breaks at random times and with random amplitudes, multiple factors, jumps and fat-tailed shocks, fractals and multifractals, and general types of nonlinearities. Such models can typically be simulated but are difficult to estimate and forecast. A general first-order representation could be expressed in terms of a latent vector \vec{v} and a vector of shocks $\vec{\eta}$.

$$\begin{aligned} r_t &= m_t + \sqrt{v_{1,t}} \varepsilon_t \\ \vec{v}_t &= f(\vec{v}_{t-1}, \vec{\eta}_t) \\ \begin{pmatrix} \varepsilon \\ \vec{\eta} \end{pmatrix} &\sim G. \end{aligned}$$

This system can be simulated if all the functions and distributions are known. Yet the forecasts and conditional variances must still be computed. Many of the stylized facts about volatility are properties of the volatility forecasts so a model like this is only a starting point in checking consistency with the data.

2. Stylized facts about asset price volatility

A number of stylized facts about the volatility of financial asset prices have emerged over the years, and been confirmed in numerous studies. A good volatility model, then, must be able to capture and reflect these stylized facts. In this section we document some of the common features of asset price volatility processes.

2.1. Volatility exhibits persistence

The clustering of large moves and small moves (of either sign) in the price process was one of the first documented features of the volatility process of asset prices. Mandelbrot (1963) and Fama (1965) both reported evidence that large changes in the price of an asset are often followed by other large changes, and small changes are often followed by small changes. This behaviour has been reported by numerous other studies, such as Baillie *et al* (1996), Chou (1988) and Schwert (1989). The implication of such volatility clustering is that volatility shocks today will influence the expectation of volatility many periods in the future. Figure 2, which will be described in the following section, displays the daily returns on the Dow Jones Industrial Index over a twelve-year period and shows evidence that the volatility of returns varies over time.

To make a precise definition of volatility persistence, let the expected value of the variance of returns k periods in the future be defined as

$$h_{t+k|t} \equiv E_t[(r_{t+k} - m_{t+k})^2]. \tag{11}$$

The forecast of future volatility then will depend upon information in today's information set such as today's returns. Volatility is said to be persistent if today's return has a large effect on the forecast variance many periods in the future. Taking partial derivatives, the forward persistence is:

$$\theta_{t+k|t} = \frac{\partial h_{t+k|t}}{\partial r_t^2}. \tag{12}$$

This is a dimensionless number as squared returns and conditional variance are in the same units.

For many volatility models this declines geometrically but may be important even a year in the future. A closely related measure is the cumulative persistence, which is the impact of a return shock on the average variance of the asset return over the period from t to $t+k$. It is defined as

$$\begin{aligned} \phi_{t+k|t} &= \frac{\partial \left(\frac{1}{k} (h_{t+k|t} + h_{t+k-1|t} + \dots + h_{t+1|t}) \right)}{\partial r_t^2} \\ &= \frac{1}{k} (\theta_{t+k|t} + \theta_{t+k-1|t} + \dots + \theta_{t+1|t}). \end{aligned} \tag{13}$$

The response of long-term option prices to volatility shocks suggests that volatility models should have significant cumulative persistence a year in the future.

A further measure of the persistence in a volatility model is the 'half-life' of volatility. This is defined as the time taken for the volatility to move halfway back towards its unconditional mean following a deviation from it:

$$\tau = k : |h_{t+k|t} - \sigma^2| = \frac{1}{2} |h_{t+1|t} - \sigma^2|. \tag{14}$$

2.2. Volatility is mean reverting

Volatility clustering implies that volatility comes and goes. Thus a period of high volatility will eventually give way to more normal volatility and similarly, a period of low volatility will be followed by a rise. Mean reversion in volatility is

generally interpreted as meaning that there is a normal level of volatility to which volatility will eventually return. Very long run forecasts of volatility should all converge to this same normal level of volatility, no matter when they are made. While most practitioners believe this is a characteristic of volatility, they might differ on the normal level of volatility and whether it is constant over all time and institutional changes.

More precisely, mean reversion in volatility implies that current information has no effect on the long run forecast. Hence

$$\text{plim}_{k \rightarrow \infty} \theta_{t+k|t} = 0, \quad \text{for all } t \tag{15}$$

which is more commonly expressed as

$$\text{plim}_{k \rightarrow \infty} h_{t+k|t} = \sigma_t^2 < \infty, \quad \text{for all } t \tag{16}$$

even though they are not quite equivalent.

It is possible to generalize the concept of mean reversion to cover processes without finite variance. Consider some other statistic such as the interquartile range or the 5% quantile and call it q_t . The same definitions in (12), (15) and (16) can be used to describe persistence and mean reversion. The cumulative versions however typically do not have the same simple form as (13), see for example the CAViAR model of Engle and Manganelli (1999).

Options prices are generally viewed as consistent with mean reversion. Under simple assumptions on option pricing, the implied volatilities of long maturity options are less volatile than those of short maturity options. They usually are closer to the long run average volatility of the asset than short maturity options.

2.3. Innovations may have an asymmetric impact on volatility

Many proposed volatility models impose the assumption that the conditional volatility of the asset is affected symmetrically by positive and negative innovations. The GARCH(1,1) model, for example, allows the variance to be affected only by the square of the lagged innovation, completely disregarding the sign of that innovation.

For equity returns it is particularly unlikely that positive and negative shocks have the same impact on the volatility. This asymmetry is sometimes ascribed to a *leverage effect* and sometimes to a *risk premium* effect. In the former theory, as the price of a stock falls, its debt-to-equity ratio rises, increasing the volatility of returns to equity holders. In the latter story, news of increasing volatility reduces the demand for a stock because of risk aversion. The consequent decline in stock value is followed by the increased volatility as forecast by the news.

Black (1976), Christie (1982), Nelson (1991), Glosten *et al* (1993) and Engle and Ng (1993) all find evidence of volatility being negatively related to equity returns. In general, such evidence has not been found for exchange rates. For interest rates a similar asymmetry arises from the boundary of zero interest rates. When rates fall (prices increase), they become less volatile in many models and in most empirical estimates, see Engle *et al* (1990b), Chan *et al* (1992) and

Brenner *et al* (1996). In diffusion models with stochastic volatility, this phenomenon is associated with correlation between the shock to returns and the shock to volatility.

The asymmetric structure of volatility generates skewed distributions of forecast prices and under simple derivative pricing assumptions, this gives option-implied volatility surfaces which have a skew. That is, the implied volatilities of out-of-the-money put options are higher than those of at-the-money options, which in turn are higher than the implieds of in-the-money puts.

2.4. Exogenous variables may influence volatility

Most of the volatility characteristics outlined above have been univariate, relating the volatility of the series to only information contained in that series' history. Of course, no-one believes that financial asset prices evolve independently of the market around them, and so we expect that other variables may contain relevant information for the volatility of a series. Such evidence has been found by, *inter alia*, Bollerslev and Melvin (1994), Engle and Mezrich (1996) and Engle *et al* (1990a,b).

In addition to other assets having an impact on the volatility series, it is possible that deterministic events also have an impact. Such things as scheduled company announcements, macroeconomic announcements and even deterministic time-of-day effects may all have an influence on the volatility process. Andersen and Bollerslev (1998a), for example, find that the volatility of the Deutsche Mark–Dollar exchange rate increases markedly around the time of the announcement of US macroeconomic data, such as the Employment Report, the Producer Price Index or the quarterly GDP. Glosten *et al* (1993) find that indicator variables for October and January assist in explaining some of the dynamics of the conditional volatility of equity returns.

2.5. Tail probabilities

It is well established that the unconditional distribution of asset returns has heavy tails. Typical kurtosis estimates range from 4 to 50 indicating very extreme non-normality. This is a feature that should be incorporated in any volatility model. The relation between the conditional density of returns and the unconditional density partially reveals the source of the heavy tails. If the conditional density is Gaussian, then the unconditional density will have excess kurtosis due simply to the mixture of Gaussian densities with different volatilities. However there is no reason to assume that the conditional density itself is Gaussian, and many volatility models assume that the conditional density is itself fat-tailed, generating still greater kurtosis in the unconditional density. Depending on the dependence structure of the volatility process, the returns may still satisfy standard extreme value theorems.

2.6. Forecast evaluation

Establishing the effectiveness of a volatility forecast is not straightforward since volatility itself is not observed. The method most consistent with the estimated models is simply to take each return divided by its one-step-ahead forecast standard

deviation and then apply any type of test to see if the square of this variable is predictable.

An alternative type of test is to examine the forecast accuracy of the model in predicting 'realized volatility', future values of sample variances. For a one-period problem, this amounts to regressing squared returns on a constant and the conditional variance. The test is whether the intercept is zero and the slope is one. Various forecasts can be entered into this equation to determine which is the best:

$$r_t^2 = a + bh_t + u_t. \quad (17)$$

This approach is not recommended for several reasons. Because r is heteroskedastic, r^2 will be much more heteroskedastic; hence this regression will be very inefficient and will have misleading standard errors. Robust standard errors should be used, however these may not make an adequate adjustment. Correcting for the heteroskedasticity would involve dividing both sides by h , leading simply to the original approach.

A second drawback is that r^2 is a noisy estimate of the volatility to be estimated. Hence the maximum R^2 that can be achieved by this regression, if all is perfectly correct, is very low. To improve this, investigators may use volatility measured over longer periods such as weekly or monthly realized volatilities. When non-overlapping periods are used, the sample becomes much smaller, and when overlapping data are used, the standard errors become far more problematic. See for example Stock and Richardson (1989). Andersen and Bollerslev (1998b) proposed using a measure of realized volatility based on observations within the period. For forecasting daily volatility, they used 5 minute data to construct a daily volatility. This improves the efficiency of this regression greatly. There is however a limit as high frequency data have lots of potential pitfalls due to bid–ask bounce and irregular spacing of the price quotes.

A third drawback to this approach is that it measures the level of variance errors rather than the more realistic proportional errors. This criterion will assess primarily the performance for high volatilities. A solution might be to take logs of the realized volatility and its forecast. For more discussion see Bollerslev *et al* (1994).

3. An empirical example

To illustrate the above points, we now present a concrete example. We use daily close price data on the Dow Jones Industrial Index, over the period 23 August 1988 to 22 August 2000, representing 3131 observations¹. The Dow Jones Industrial Index is comprised of 30 industrial companies' stocks, and represent about a fifth of the total value of the US stock market. We take the log-difference of the value of the index, so as to convert the data into continuously compounded returns. Figures 1 and 2 plot the price level and the returns on the index over the sample period.

¹ These data in ASCII format are available from the second author's web site at <http://www.econ.ucsd.edu/~apatton/dowjones.txt>.

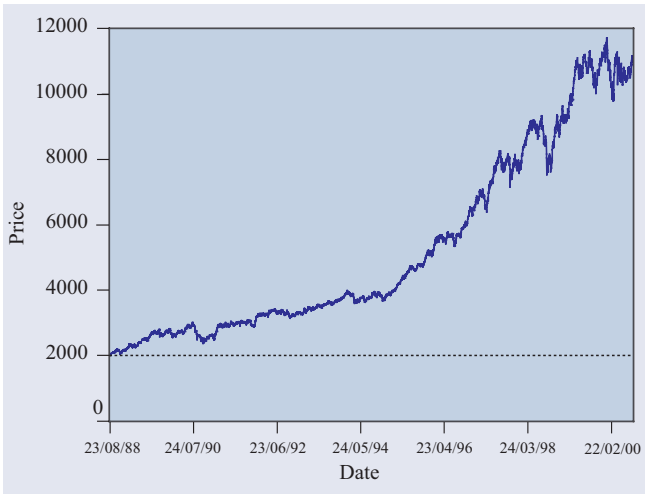


Figure 1. The Dow Jones Industrial Index, 23 August 1988 to 22 August 2000.




Figure 2. Returns on the Dow Jones Industrial Index.

3.1. Summary of the data

Some summary statistics on the data are presented in table 1. As this table shows, the index had a small positive average return of about one-twentieth of a per cent per day. The daily variance was 0.8254, implying an average annualized volatility of 14.42%. The skewness coefficient indicates that the returns distribution is substantially negatively skewed; a common feature of equity returns. Finally, the kurtosis coefficient, which is a measure of the thickness of the tails of the distribution, is very high. A Gaussian distribution has kurtosis of 3, implying that the assumption of Gaussianity for the distribution of returns is dubious for this series².

An analysis of the correlogram of the returns, presented in figure 3, indicates only weak dependence in the mean of the series, and so for the remainder of the paper we will

² The Jarque–Bera test for normality of the returns distribution yields a statistic of 4914.116, much greater than any critical value at conventional confidence levels, thus rejecting the null hypothesis of normally distributed returns.




Figure 3. Correlograms of returns and squared returns.

assume a constant conditional mean. The correlogram of the squared returns, however, indicates substantial dependence in the volatility of returns.

3.2. A volatility model

A widely used class of models for the conditional volatility is the autoregressive conditionally heteroskedastic class of models introduced by Engle (1982), and extended by Bollerslev (1986), Engle *et al* (1987), Nelson (1991), Glosten *et al* (1993), amongst many others. See Bollerslev *et al* (1992) or (1994) for summaries of this family of models.

A popular member of the ARCH class of models is the GARCH(*p, q*) model:

$$h_t = \omega + \sum_{i=1}^p \alpha_i (R_{t-i} - \mu)^2 + \sum_{j=1}^q \beta_j h_{t-j}. \quad (18)$$

Table 1. Dow Jones Industrial Index returns summary statistics.

Mean	0.0550
Variance	0.8254
Skewness	-0.5266
Kurtosis	9.0474

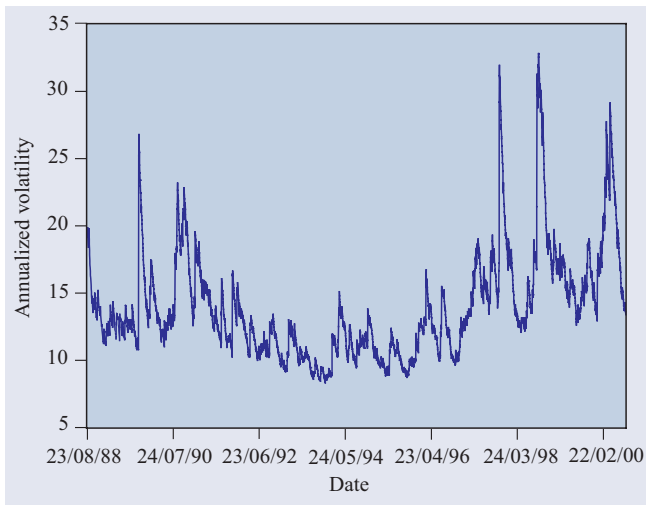


Figure 4. Estimated conditional volatility using a GARCH(1,1) model.

This model can be estimated via maximum likelihood once a distribution for the innovations, ε_t , has been specified. A commonly employed assumption is that the innovations are Gaussian³.

Using the Schwarz Information Criterion we found that the best model in the GARCH(p, q) class for $p \in [1, 5]$ and $q \in [1, 2]$ was a GARCH(1,1). The results for this model are presented in table 2.

A test for whether this volatility model has adequately captured all of the persistence in the variance of returns is to look at the correlogram of the standardized squared residuals. If the model is adequate, then the standardized squared residuals should be serially uncorrelated. The Ljung-Box Q-statistic at the twentieth lag of the standardized squared residuals was 8.9545, indicating that the standardized squared residuals are indeed serially uncorrelated.

3.3. Mean reversion and persistence in volatility

The results above indicate that the volatility of returns is quite persistent, with the sum of α and β being 0.9904, implying a volatility half-life of about 73 days. Although the returns volatility appears to have quite long memory, it is still mean reverting: the sum of α and β is significantly less than one⁴, implying that, although it takes a long time, the volatility process does return to its mean. The unconditional mean of the GARCH(1,1) process is calculated as the ratio of ω to the difference between 1 and the sum of α and β . For the Dow Jones over the sample period this turns out to be 0.8542, which implies that the mean annualized volatility over the sample was 14.67%, very close to the sample estimate of the unconditional

³ Bollerslev and Wooldridge (1992) showed that the maximum likelihood estimates of the parameters of the GARCH model assuming Gaussian errors are consistent even if the true distribution of the innovations is not Gaussian. The usual standard errors of the estimators are not consistent when the assumption of Gaussianity of the errors is violated, so Bollerslev and Wooldridge supply a method for obtaining consistent estimates of these.

⁴ A one-sided t -test that the sum of alpha and beta is greater than or equal to one yields a test statistic of -2.54 , which is greater (in absolute value) than the 5% critical value of -1.96 .




Figure 5. θ and ϕ for k ranging from 1 to 100.

volatility given in table 1. A plot of the annualized conditional volatility estimates over the sample period is given in figure 4.

As described in section 2.1, a measure of the persistence in volatility is the partial derivative of the overnight return volatility at time $t+k$ with respect to the squared return at time t , denoted $\theta_{t+k,t}$. A plot of $\theta_{t+k,t}$ for k ranging from 1 to 100 is given in figure 5. This plot shows that the impact decays geometrically, and is essentially zero beyond one hundred days. The limit of this sequence is zero, confirming that this volatility process is mean reverting. The equivalent measure for the volatility of k -period returns, denoted $\phi_{t+k,t}$ in section 2.1, also declines toward zero, though at a slower rate, as equation (13) suggests that it should.

An alternative way to observe the mean-reverting behaviour in h_t is in the structure of long-term forecasts of volatility. Figure 6 presents forecasts at 23 August 1995 and 23 August 1997 of the annualized daily return volatility out to a year from each of those dates. The first of these forecasts was made at a date with unusually high volatility, and so the forecasts of volatility decline gradually to the unconditional variance level. The second of these forecasts was made during a tranquil period, and so the sequence of forecasts is increasing toward the unconditional volatility level.

One way to examine the volatility of volatility, VoV, is to plot the one-period-ahead volatility and the k -periods-ahead forecast volatility. In figure 7 we present these forecasts for the one-day, one-quarter, one-year and two-year cumulative forecasts. It is immediately apparent that the movements in the one-day horizon are larger than the movements in the two-year

Table 2. Results from the GARCH(1,1) model.

	Coefficient	Robust standard error
Constant	0.0603	0.0143
ω	0.0082	0.0025
α	0.0399	0.0104
β	0.9505	0.0105

Table 3. Volatility of volatility for various forecast horizons from GARCH(1,1).

	One day	One quarter	One year	Two years
Standard deviation	51.19845	39.45779	22.52775	13.77907

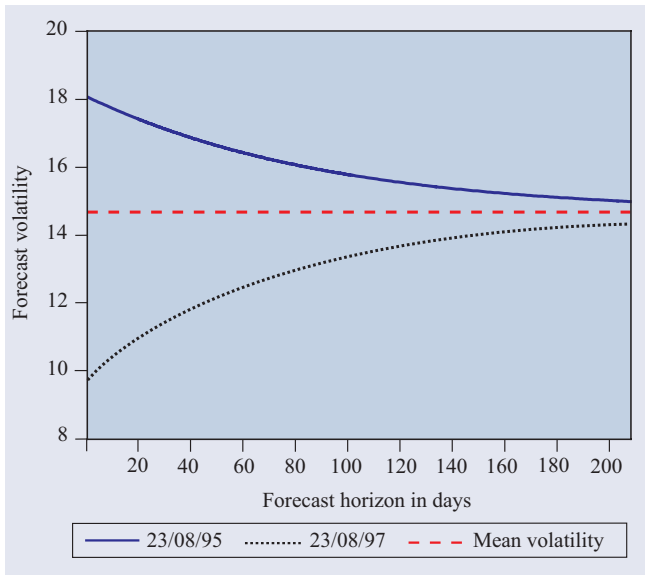


Figure 6. Forecasts of daily return volatility using the GARCH(1,1) model.

horizon. The intermediate horizons lie between. This is an implication of the mean reversion in volatility. The annualized estimates of the volatility of volatility for these forecasts are given in table 3.

3.4. An asymmetric volatility model

As mentioned in the previous section, the sign of the innovation may influence the volatility in addition to its magnitude. There are a number of ways of parametrizing this idea, one of which is the Threshold GARCH (or TAR) model. This model was proposed by Glosten *et al* (1993) and Zakoian (1994) and was motivated by the EGARCH model of Nelson (1991).

$$h_t = \omega + \sum_{i=1}^p \alpha_i (R_{t-i} - \mu)^2 + \sum_{j=1}^q \beta_j h_{t-j} + \sum_{k=1}^r \delta_{t-k} \gamma_k (R_{t-k} - \mu)^2 \tag{19}$$

where δ_{t-k} is an indicator variable, taking the value one if the residual at time $t - k$ was negative, and zero elsewhere.

Table 4. Results from the TAR(1,1,1) model.

	Coefficient	Robust standard error
Constant	0.0509	0.0151
ω	0.0184	0.0024
α	0.0151	0.0070
γ	0.0654	0.0083
β	0.9282	0.0073

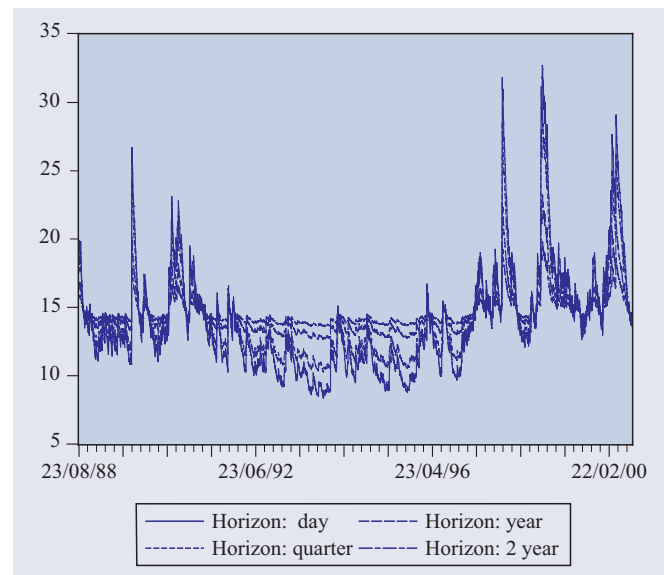


Figure 7. Volatilities at different horizons from GARCH(1,1).

The TAR model implies that a positive innovation at time t has an impact on the volatility at time $t+1$ equal to α times the residual squared, while a negative innovation has impact equal to $(\alpha + \gamma)$ times the residual squared. The presence of the leverage effect would imply that the coefficient γ is positive, that is, that a negative innovation has a greater impact than a positive innovation.

We estimated the TAR(1,1,1) model, and present the results in table 4. These results indicate that the sign of the innovation has a significant influence on the volatility of returns. The coefficient on negative residuals squared is large and significant, and implies that a negative innovation at time t increases the volatility at time $t + 1$ by over four times as much as a positive innovation of the same magnitude.

3.5. A model with exogenous volatility regressors

It may also be of interest to gauge the impact of exogenous variables on the volatility process. This type of model could offer a structural or economic explanation for volatility. Such

Table 5. Results from the GARCH(1,1)-X model.

	Coefficient	Robust standard error
Constant	0.0608	0.0145
ω	-0.0010	0.0016
α	0.0464	0.0040
β	0.9350	0.0065
φ	0.0031	0.0005

Table 6. GARCH(1,1) parameter estimates for data of differing frequencies.

	Daily data	2-day period	3-day period	4-day period	Weekly data
Constant	0.0603	0.1145	0.1715	0.2148	0.2730
ω	0.0082	0.0138	0.0304	0.0238	0.0577
α	0.0399	0.0419	0.0528	0.0416	0.0496
β	0.9505	0.9498	0.9358	0.9529	0.9408

a model may be written as:

$$h_t = \omega + \sum_{i=1}^p \alpha_i (R_{t-i} - \mu)^2 + \sum_{j=1}^q \beta_j h_{t-j} + \varphi X_{t-1}. \quad (20)$$

As an example, we used the lagged level of the three-month US Treasury bill rate as an exogenous regressor in our model of Dow Jones Industrial Index returns volatility. The T-bill rate is correlated with the cost of borrowing to firms, and thus may carry information that is relevant to the volatility of the Dow Jones Industrial Index.

As the reader can see, the impact of the T-bill rate on the volatility process of the Dow Jones Industrials is small, but quite significant. The positive sign on this coefficient indicates that high interest rates are generally associated with higher levels of equity return volatility. This result confirms that of Glosten *et al* (1993) who also find that the Treasury bill rate is positively related to equity return volatility.

3.6. Aggregation of volatility models

Despite the success of GARCH models in capturing the salient features of conditional volatility, they have some undesirable characteristics. Most notably, the theoretical observation that if a GARCH model is correctly specified for one frequency of data, then it will be misspecified for data with different time scales, makes a researcher uneasy. Similarly, if assets follow a GARCH model, then portfolios do not exactly do so. Below, we present some evidence of this for our example data set. We consider the estimation of the simple GARCH(1,1) model on the data, sampled at various frequencies. The results are presented in table 6.

These results indicate that the sampling frequency does indeed affect the results obtained. As an example, the implied half-life of volatility implied by each of the models (in days) is 73, 168, 183, 508 and 365. Clearly these are substantial differences although the statistical and forecast significance of these differences should be assessed. To some extent, the interpretation of these models with aggregate data is slightly different.

Ideas such as the weak GARCH specification of Drost and Nijman (1993) may represent an alternative solution. However, the empirical estimates on different time scales or portfolios are typically reasonable, suggesting that GARCH can be interpreted as an approximation or filter rather than a full statistical specification. Steps in this direction are developed by Nelson and Foster (1994).

4. Conclusions and challenges for future research

The goal of this paper has been to characterize a good volatility model by its ability to forecast and capture the commonly held stylized facts about conditional volatility. The stylized facts include such things as the persistence in volatility, its mean-reverting behaviour, the asymmetric impact of negative versus positive return innovations and the possibility that exogenous or pre-determined variables may have a significant influence on volatility.

We used twelve years of daily data on the Dow Jones Industrial Index to illustrate these stylized facts, and the ability of models from the GARCH family to capture these characteristics. The conditional volatility of the Dow Jones Industrial Index was found to be quite persistent, with a volatility half-life of about 73 days, yet tests for non-stationarity indicated that it is mean reverting. A negative lagged return innovation was found to have an impact on conditional variance roughly four times as large as a positive return innovation, and the three-month US Treasury bill rate was found to be positively correlated with volatility, implying that higher interest rates lead to higher equity return volatility. Finally, we found evidence consistent with the theoretical result that the empirical results obtained are dependent on the sampling frequency—a drawback of the GARCH specification.

Various aspects of the volatility process are important topics of research. The need for a model to forecast 100 or even 1000 steps into the future has suggested long memory or fractionally integrated processes. In spite of substantial research input, the value for these forecast situations has not yet been established. Shifts in the volatility process are sometimes thought to be discrete events; only the Hamilton and Susmel (1994) model and its extension by Gray (1996) have been developed for this task. Time-varying higher conditional moments are clearly of interest but have proven difficult to estimate. Hansen (1994) and more recently Harvey and Sidiqi (1999) have had some success.

References

- Andersen T G and Bollerslev T 1998a Deutsche Mark–Dollar volatility: intraday activity patterns, macroeconomic announcements, and longer run dependencies *J. Finance* **53** 219–65
- Andersen T G and Bollerslev T 1998b Answering the skeptics: yes, standard volatility models do provide accurate forecasts *Int.*

- Economic Rev.* **39** 885–905
- Andersen T G, Bollerslev T, Diebold F X and Labys P 2001 The distribution of exchange rate volatility *J. Am. Statist. Ass.* **96** 42–55
- Baillie R T, Bollerslev T and Mikkelsen H O 1996 Fractionally integrated generalized autoregressive conditional heteroskedasticity *J. Econometrics* **74** 3–30
- Black F 1976 Studies of stock market volatility changes *Proc. 1976 Meetings of the American Statistical Association, Business and Economic Statistics Section* pp 177–81
- Bollerslev T 1986 Generalized autoregressive conditional heteroskedasticity *J. Econometrics* **31** 307–27
- Bollerslev T, Chou R Y and Kroner K F 1992 ARCH modeling in finance: a review of the theory and empirical evidence *J. Econometrics* **52** 5–59
- Bollerslev T, Engle R F and Nelson D B 1994 ARCH models *Handbook of Econometrics* vol IV, ed R F Engle and D McFadden (Amsterdam: North-Holland) pp 2959–3038
- Bollerslev T and Melvin M 1994 Bid–ask spreads and the volatility in the foreign exchange market: an empirical analysis *J. Int. Economics* **36** 355–72
- Bollerslev T and Wooldridge J M 1992 Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances *Econometric Rev.* **11** 143–72
- Brenner R J, Harjes R H and Kroner K F 1996 Another look at models of the short-term interest rate *J. Financial Quantitative Analysis* **31** 85–107
- Chan K C, Karolyi G A, Longstaff F A and Sanders A B 1992 An empirical comparison of alternative models of the short-term interest rate *J. Finance* **47** 1209–27
- Chou R Y 1988 Volatility persistence and stock valuations: some empirical evidence using GARCH *J. Appl. Econometrics* **3** 279–94
- Christie A A 1982 The stochastic behavior of common stock variances: value, leverage and interest rate effects *J. Financial Economics* **10** 407–32
- Drost F C and Nijman T E 1993 Temporal aggregation of GARCH processes *Econometrica* **61** 909–27
- Engle R F 1982 Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation *Econometrica* **50** 987–1007
- Engle R F, Ito T and Lin W-L 1990a Meteor showers or heat waves? Heteroscedastic intra-daily volatility in the foreign exchange market *Econometrica* **58** 525–42
- Engle R F, Lilien D and Robins R 1987 Estimation of time varying risk premia in the term structure: the ARCH-M model *Econometrica* **55** 391–407
- Engle R F and Manganelli S 1999 CAViaR: conditional autoregressive value at risk by regression quantiles *University of California, San Diego, Department of Economics Working Paper 99-20*
- Engle R F and Mezrich J 1996 GARCH for groups *Risk* **9** 36–40
- Engle R F and Ng V 1993 Measuring and testing the impact of news on volatility *J. Finance* **48** 1749–78
- Engle R F, Ng V K and Rothschild M 1990b Asset pricing with a factor-ARCH covariance structure *J. Econometrics* **45** 235–7
- Fama E F 1965 The behavior of stock-market prices *J. Business* **38** 34–105
- Glosten L R, Jagannathan R and Runkle D E 1993 On the relation between the expected value and the volatility of the nominal excess returns on stocks *J. Finance* **48** 1779–801
- Gray S F 1996 Modeling the conditional distribution of interest rates as a regime-switching process *J. Financial Economics* **42** 27–62
- Hamilton J D and Susmel R 1994 Autoregressive conditional heteroskedasticity and changes in regime *J. Econometrics* **64** 307–33
- Hansen B E 1994 Autoregressive conditional density estimation *Int. Economic Rev.* **35** 705–30
- Harvey C R and Siddique A 1999 Autoregressive conditional skewness *J. Financial Quantitative Anal.* **34** 465–87
- Mandelbrot B 1963 The variation of certain speculative prices *J. Business* **36** 394–419
- Nelson D B 1991 Conditional heteroscedasticity in asset returns: a new approach *Econometrica* **59** 347–70
- Nelson D B and Foster D P 1994 Asymptotic filtering theory for univariate arch models *Econometrica* **62** 1–41
- Schwert G W 1989 Why does stock market volatility change over time? *J. Finance* **44** 1115–53
- Richardson M P and Stock J H 1989 Drawing inferences from statistics based on multi-year asset returns *J. Financial Economics* **25** 323–48
- Zakoian J -M 1994 Threshold heteroskedastic models *J. Economic Dynamics Control* **18** 931–55