



Lancaster University
MANAGEMENT SCHOOL

Lancaster University Management School
Working Paper
2001/017

**Consequences for Option Pricing of a Long Memory in
Volatility**

Stephen John Taylor

The Department of Accounting and Finance
Lancaster University Management School
Lancaster LA1 4YX
UK

©Stephen John Taylor

All rights reserved. Short sections of text, not to exceed
two paragraphs, may be quoted without explicit permission,
provided that full acknowledgement is given.

The LUMS Working Papers series can be accessed at <http://www.lums.co.uk/publications>
LUMS home page: <http://www.lums.lancs.ac.uk/>

Consequences for option pricing of a long memory in volatility

Stephen J Taylor*

Department of Accounting and Finance, Lancaster University, England LA1 4YX

December 2000

Abstract

The economic consequences of a long memory assumption about volatility are documented, by comparing implied volatilities for option prices obtained from short and long memory volatility processes. Numerical results are given for options on the S & P 100 index from 1984 to 1998, with lives up to two years. The long memory assumption is found to have a significant impact upon the term structure of implied volatilities and a relatively minor impact upon smile effects. These conclusions are important because evidence for long memory in volatility has been found in the prices of many assets.

* S.Taylor@lancaster.ac.uk, www.lancs.ac.uk/staff/afasjt. The author acknowledges the support of Inquire Europe, with particular thanks to Bernard Dumas and Ton Vorst for their helpful advice.

1. Introduction

Long memory effects in a stochastic process are effects that decay too slowly to be explained by stationary processes defined by a finite number of autoregressive and moving-average terms. Long memory is often represented by fractional integration of shocks to the process, which produces autocorrelations that decrease at a hyperbolic rate compared with the faster asymptotic rate of stationary ARMA processes.

Long memory in volatility occurs when the effects of volatility shocks decay slowly. This phenomenon can be identified from the autocorrelations of measures of realised volatility. Two influential examples are the study of absolute daily returns from stock indices by Ding, Granger and Engle (1993) and the recent investigation of daily sums of squared five-minutes returns from exchange rates by Andersen, Bollerslev, Diebold and Labys (2000).

Stochastic volatility causes option prices to display both smile and term structure effects. An implied volatility obtained from the Black-Scholes formula then depends on both the exercise price and the time until the option expires. Exact calculation of smile and term effects is only possible for special, short memory volatility processes, with the results of Heston (1993) being a notable example. Monte Carlo methods are necessary, however, when the volatility process has a long memory. These have been evaluated by Bollerslev and Mikkelsen (1996, 1999).

The objective of this paper is to document the economic consequences of a long memory assumption about volatility. This is achieved by comparing implied volatilities for option prices obtained from short and long memory specifications. It is necessary to

use a long history of asset prices when applying a long memory model and this paper uses levels of the S & P 100 index from 1984 to 1998. For this data it is found that a long memory assumption has a significant economic impact upon the term structure of implied volatilities and a relatively minor impact upon smile effects.

Market makers in options have to make assumptions about the volatility process. The effects of some assumptions are revealed by the prices of options with long lives. Bollerslev and Mikkelsen (1999) find that the market prices of exchange traded options on the S & P 500 index, with lives between nine months and three years, are described more accurately by a long memory pricing model than by the short memory alternatives. Thus some option prices already reflect the long memory phenomenon in volatility, although Bollerslev and Mikkelsen (1999) find that significant biases remain to be explained. The same LEAPS contracts are also investigated by Bakshi, Cao and Chen (2000) who show that options with long lives can differentiate between competing pricing models, but their analysis is restricted to a short memory context.

Three explanatory sections precede the illustrative option pricing results in Section 5, so that this paper provides a self-contained description of how to price options with a long memory assumption. A general introduction to the relevant literature is provided by Granger (1980) and Baillie (1996) on long memory, Andersen, Bollerslev, Diebold and Labys (2000) on evidence for long memory in volatility, Duan (1995) on option pricing for ARCH processes, and Bollerslev and Mikkelsen (1999) on applying these pricing methods with long memory specifications.

Section 2 defines and characterises long memory precisely. These characteristics are illustrated theoretically for the fractionally integrated white noise process and then the empirical evidence for these characteristics in volatility is surveyed. The empirical evidence for the world's major markets is compelling and explanations for the source of long memory effects in volatility are summarised.

Section 3 describes parsimonious volatility models that incorporate long memory either within an ARCH or a stochastic volatility framework. The former framework is easier to use and we focus on using the fractionally integrated extension of the exponential GARCH model, that is known by the acronym FIEGARCH. An important feature of applications is the unavoidable truncation of an autoregressive component of infinite order. Empirical results are provided for ten years of S & P 100 returns.

Section 4 provides the necessary option pricing methodology. Contingent claim prices are obtained by simulating terminal payoffs using an appropriate risk-neutral measure. A risk-premium term then appears in the simulated process and has a non-trivial effect upon the results that is typically seen in term structures that slope upwards on average. Numerical methods that enhance the accuracy of the simulations are also described.

Section 5 compares implied volatilities for European option prices obtained from short and long memory volatility specifications, for hypothetical S & P 100 options whose lives range from one month to two years. Options are valued on ten dates, one per annum from 1989 to 1998. The major impact of the long memory assumption is seen to be the very slow convergence of implied volatilities to a limit as

the option life increases. This convergence is so slow that the limit can not be estimated precisely. Section 6 contains conclusions.

2. Long memory

2.1. Definitions

There are several definitions that categorise stochastic processes as having either a short memory or a long memory; examples can be found in McLeod and Hipel (1978), Brockwell and Davis (1991), Baillie (1996) and Granger and Ding (1996). The fundamental characteristic of a long-memory process is that dependence between variables separated by t time units does not decrease rapidly as t increases.

Consider a covariance stationary stochastic process $\{x_t\}$ that has variance \mathbf{s}^2 and autocorrelations \mathbf{r}_t , spectral density $f(\mathbf{w})$ and n -period variance-ratios V_n defined by

$$\mathbf{r}_t = \text{cor}(x_t, x_{t+t}), \quad (1)$$

$$f(\mathbf{w}) = \frac{\mathbf{s}^2}{2\mathbf{p}} \sum_{t=-\infty}^{\infty} \mathbf{r}_t \cos(t\mathbf{w}), \quad \mathbf{w} > 0, \quad (2)$$

$$V_n = \frac{\text{var}(x_{t+1} + \dots + x_{t+n})}{n\mathbf{s}^2} = 1 + 2 \sum_{t=1}^{n-1} \frac{n-t}{n} \mathbf{r}_t. \quad (3)$$

Then a covariance stationary process is here said to have a short memory if $\sum_{t=1}^n \mathbf{r}_t$

converges as $n \rightarrow \infty$, otherwise it is said to have a long memory. A short memory process then has

$$\sum_{t=1}^n \mathbf{r}_t \rightarrow C_1, \quad f(\mathbf{w}) \rightarrow C_2, \quad V_n \rightarrow C_3, \quad \text{as } n \rightarrow \infty, \mathbf{w} \rightarrow 0, \quad (4)$$

for constants C_1, C_2, C_3 . Examples are provided by stationary ARMA processes. These processes have geometrically bounded autocorrelations, so that $|\mathbf{r}_t| \leq C\mathbf{f}^t$ for some $C > 0$ and $1 > \mathbf{f} > 0$, and hence (4) is applicable.

In contrast to the above results, all the limits given by (4) do not exist for a typical covariance stationary long memory process. Instead, it is typical that the autocorrelations have a hyperbolic decay, the spectral density is unbounded for low frequencies and the variance-ratio increases without limit. Appropriate limits are then provided for some positive $d < \frac{1}{2}$ by

$$\frac{\mathbf{r}_t}{t^{2d-1}} \rightarrow D_1, \quad \frac{f(\mathbf{w})}{\mathbf{w}^{-2d}} \rightarrow D_2, \quad \frac{V_n}{n^{2d}} \rightarrow D_3, \quad \text{as } n \rightarrow \infty, \mathbf{w} \rightarrow 0, \quad (5)$$

for positive constants D_1, D_2, D_3 . The limits given by (5) characterise the stationary long memory processes that are commonly used to represent long memory in volatility. The fundamental parameter d can be estimated from data using a regression, either of $\ln(\hat{f}(\mathbf{w}))$ on \mathbf{w} or of $\ln(\hat{V}_n)$ on n as in, for example, Andersen, Bollerslev, Diebold and Ebens (2000).

2.2. Fractionally integrated white noise

An important example of a long memory process is a stochastic process $\{y_t\}$ that requires fractional differencing to obtain a set of independent and identically distributed residuals $\{\mathbf{e}_t\}$. Following Granger and Joyeux (1980) and Hosking (1981), such a process is defined using the filter

$$(1-L)^d = 1 - dL + \frac{d(d-1)}{2!}L^2 - \frac{d(d-1)(d-2)}{3!}L^3 + \dots \quad (6)$$

where L is the usual lag operator, so that $Ly_t = y_{t-1}$. Then a fractionally integrated white noise (FIWN) process $\{y_t\}$ is defined by

$$(1-L)^d y_t = \mathbf{e}_t \quad (7)$$

with the \mathbf{e}_t assumed to have zero mean and variance \mathbf{s}_e^2 . Throughout this paper it is assumed that the differencing parameter d is constrained by $0 \leq d < 1$.

The mathematical properties of FIWN are summarised in Baillie (1996) and were first derived by Granger and Joyeux (1980) and Hosking (1981). The process is covariance stationary if $d < \frac{1}{2}$ and then the following results apply. First, the autocorrelations are given by

$$\mathbf{r}_1 = \frac{d}{1-d}, \quad \mathbf{r}_2 = \frac{d(d+1)}{(1-d)(2-d)}, \quad \mathbf{r}_3 = \frac{d(d+1)(d+2)}{(1-d)(2-d)(3-d)}, \dots \quad (8)$$

or, in terms of the Gamma function,

$$\mathbf{r}_t = \frac{\Gamma(1-d)\Gamma(t+d)}{\Gamma(d)\Gamma(t+1-d)}, \quad (9)$$

with

$$\frac{\mathbf{r}_t}{t^{2d-1}} \rightarrow \frac{\Gamma(1-d)}{\Gamma(d)} \quad \text{as } t \rightarrow \infty. \quad (10)$$

Second, the spectral density is

$$f(\mathbf{w}) = \frac{\mathbf{s}_e^2}{2\mathbf{p}} |1 - e^{-i\mathbf{w}}|^{-2d} = \frac{\mathbf{s}_e^2}{2\mathbf{p}} \left[2 \sin\left(\frac{\mathbf{w}}{2}\right) \right]^{-2d}, \quad \mathbf{w} > 0, \quad (11)$$

so that

$$f(\mathbf{w}) \cong \frac{\mathbf{s}_e^2}{2\mathbf{p}} \mathbf{w}^{-2d} \quad \text{for } \mathbf{w} \text{ near } 0. \quad (12)$$

Also,

$$\frac{V_n}{n^{2d}} \rightarrow \frac{\Gamma(1-d)}{(1+2d)\Gamma(1+d)} \quad \text{as } n \rightarrow \infty. \quad (13)$$

When $d \geq \frac{1}{2}$ the FIWN process has infinite variance and thus the autocorrelations are not defined, although the process has some stationarity properties for $\frac{1}{2} \leq d < 1$.

2.3. Evidence for long memory in volatility

When returns r_t can be represented as $r_t = \mathbf{m} + \mathbf{s}_t u_t$, with \mathbf{s}_t representing volatility and independent of an i.i.d. standardised return u_t , it is often possible to make inferences about the autocorrelations of volatility from the autocorrelations of either $|r_t - \mathbf{m}|$ or $(r_t - \mathbf{m})^2$, see Taylor (1986). In particular, evidence for long memory in powers of daily absolute returns is also evidence for long memory in volatility. Ding, Granger and Engle (1993) observe hyperbolic decay in the autocorrelations of powers

of daily absolute returns obtained from U.S. stock indices. Dacorogona, Muller, Nagler, Olsen and Pictet (1993) observe a similar hyperbolic decay in 20-minute absolute exchange rate returns. Breidt, Crato and de Lima (1998) find that spectral densities estimated from the logarithms of squared index returns have the shape expected from a long memory process at low frequencies. Further evidence for long memory in volatility has been obtained by fitting appropriate fractionally integrated ARCH models and then testing the null hypothesis $d = 0$ against the alternative $d > 0$. Bollerslev and Mikkelsen (1996) use this methodology to confirm the existence of long memory in U.S. stock index volatility.

The recent evidence for long memory in volatility uses high-frequency data to construct accurate estimates of the volatility process. The quadratic variation $\hat{\mathbf{S}}_t^2$ of the logarithm of the price process during a 24-hour period denoted by t can be estimated by calculating intraday returns $r_{t,j}$ and then

$$\hat{\mathbf{S}}_t^2 = \sum_{j=1}^N r_{t,j}^2. \quad (14)$$

The estimates are illustrated by Taylor and Xu (1997) for a year of DM/\$ rates with $N = 288$ so that the $r_{t,j}$ are 5-minute returns. As emphasised by Andersen, Bollerslev, Diebold and Labys (2000), the estimate $\hat{\mathbf{S}}_t^2$ will be very close to the integral of the unobservable volatility during the same 24-hour period providing N is large but not so large that the bid-ask spread introduces bias into the estimate. Using five-minute returns provides conclusive evidence for long memory effects in the estimates $\hat{\mathbf{S}}_t^2$ in four studies: Andersen, Bollerslev, Diebold and Labys (2000) for ten years of DM/\$

and Yen/\$ rates, Andersen, Bollerslev, Diebold and Ebens (2000) for five years of stock prices for the 30 components of the Dow-Jones index, Ebens (1999) for fifteen years of the same index and Areal and Taylor (2000) for eight years of FTSE-100 stock index futures prices. These papers provide striking evidence that time series of estimates $\hat{\mathbf{S}}_t^2$ display all three properties of a long memory process: hyperbolic decay in the autocorrelations, spectral densities at low frequencies that are proportional to ω^{-2d} and variance-ratios whose logarithms are very close to linear functions of the aggregation period n . It is also seen from these papers that estimates of d are between 0.3 and 0.5, with most estimates close to 0.4.

2.4. Explanations of long memory in volatility

Granger (1980) shows that long memory can be a consequence of aggregating short memory processes; specifically if AR(1) components are aggregated and if the AR(1) parameters are drawn from a Beta distribution then the aggregated process converges to a long memory process as the number of components increases. Andersen and Bollerslev (1997) develop Granger's theoretical results in more detail for the context of aggregating volatility components and also provide supporting empirical evidence obtained from only one year of 5-minute returns. It is plausible to assert that volatility reflects several sources of news, that the persistence of shocks from these sources depends on the source and hence that total volatility follows a long memory process. Scheduled macroeconomic news announcements are known to create additional volatility that is very short-lived (Ederington and Lee, 1993), whilst other sources of

news that have a longer impact on volatility are required to explain volatility clustering effects that last several weeks.

Gallant, Hsu and Tauchen (1999) estimate a volatility process for daily IBM returns that is the sum of only two short memory components yet the sum is able to mimic long memory. They also show that the sum of a particular pair of AR(1) processes has a spectral density function very close to that of fractionally integrated white noise with $d = 0.4$ for frequencies $\omega \geq 0.01\pi$. Consequently, evidence for long memory may be consistent with a short memory process that is the sum of a small number of components whose spectral density happens to resemble that of a long memory process except at extremely low frequencies. The long memory specification may then provide a much more parsimonious model. Barndorff-Nielsen and Shephard (2001) model volatility in continuous-time as the sum of a few short memory components. Their analysis of ten years of 5-minute DM/\$ returns, adjusted for intraday volatility periodicity, shows that the sum of four volatility processes is able to provide an excellent match to the autocorrelations of squared 5-minute returns, which exhibit the long memory property of hyperbolic decay.

3. Long memory volatility models

A general set of long memory stochastic processes can be defined by first applying the filter $(1 - L)^d$ and then assuming that the filtered process is a stationary ARMA(p, q) process. This defines the ARFIMA(p, d, q) models of Granger (1980), Granger and Joyeux (1980) and Hosking (1981). This approach can be used to obtain long memory

models for volatility, by extending various specifications of short memory volatility processes. We consider both ARCH and stochastic volatility specifications.

3.1. ARCH specifications

The conditional distributions of returns r_t are defined for ARCH models using information sets I_{t-1} that are here assumed to be previous returns $\{r_{t-i}, i \geq 1\}$, conditional mean functions $\mathbf{m}_t(I_{t-1})$, conditional variance functions $h_t(I_{t-1})$ and a probability distribution D for standardised returns z_t . Then the terms

$$z_t = \frac{r_t - \mathbf{m}_t}{\sqrt{h_t}} \quad (15)$$

are independently and identically distributed with distribution D and have zero mean and unit variance.

Baillie (1996) and Bollerslev and Mikkelsen (1996) both show how to define a long memory process for h_t by extending either the GARCH models of Bollerslev (1986) or the exponential ARCH models of Nelson (1991). The GARCH extension can not be recommended because the returns process then has infinite variance for all positive values of d , which is incompatible with the stylized facts for asset returns. For the exponential extension, however, $\ln(h_t)$ is covariance stationary for $d < \frac{1}{2}$; it may then be conjectured that the returns process has finite variance for particular specifications of h_t .

Like Bollerslev and Mikkelsen (1996, 1999), this paper applies the FIEGARCH(1, d , 1) specification :

$$\ln(h_t) = \mathbf{a} + (1 - \mathbf{f}L)^{-1}(1 - L)^{-d}(1 + \mathbf{y}L)g(z_{t-1}), \quad (16)$$

$$g(z_t) = \mathbf{q}z_t + \mathbf{g}(|z_t| - C), \quad (17)$$

with \mathbf{a} , \mathbf{f} , d , \mathbf{y} respectively denoting the location, autoregressive, differencing and moving average parameters of $\ln(h_t)$. The i.i.d. residuals $g(z_t)$ depend on a symmetric response parameter \mathbf{g} and an asymmetric response parameter \mathbf{q} that enables the conditional variances to depend on the signs of the terms z_t ; these residuals have zero mean because C is defined to be the expectation of $|z_t|$. The EGARCH(1,1) model of Nelson (1991) is given by $d = 0$. If $\mathbf{f} = \mathbf{y} = 0$ and $d > 0$, then $\ln(h_t) - \mathbf{a}$ is a fractionally integrated white noise process. In general, $\ln(h_t)$ is an ARFIMA(1, d , 1) process.

Calculations using equation (16) require series expansions in the lag operator L .

We note here the results :

$$(1 - L)^d = 1 - \sum_{j=1}^{\infty} a_j L^j, \quad a_1 = d, \quad a_j = \frac{j-d-1}{j} a_{j-1}, \quad j \geq 2, \quad (18)$$

$$(1 - \mathbf{f}L)(1 - L)^d = 1 - \sum_{j=1}^{\infty} b_j L^j, \quad b_1 = d + \mathbf{f}, \quad b_j = a_j - \mathbf{f}a_{j-1}, \quad j \geq 2, \quad (19)$$

$$(1 - \mathbf{f}L)(1 - L)^d(1 + \mathbf{y}L)^{-1} = 1 - \sum_{j=1}^{\infty} \mathbf{f}_j L^j,$$

$$\mathbf{f}_1 = d + \mathbf{f} + \mathbf{y}, \quad \mathbf{f}_j = b_j - (-\mathbf{y})^j + \sum_{k=1}^{j-1} (-\mathbf{y})^{j-k} b_k, \quad j \geq 2, \quad (20)$$

and that the autoregressive weights in (20) can be denoted as $\mathbf{f}_j(d, \mathbf{f}, \mathbf{y})$. Also,

$$(1 - \mathbf{f}L)^{-1}(1 - L)^{-d}(1 + \mathbf{y}L) = 1 + \sum_{j=1}^{\infty} \mathbf{y}_j L^j,$$

$$\mathbf{y}_1 = d + \mathbf{f} + \mathbf{y}, \quad \mathbf{y}_j = -\mathbf{f}_j(-d, -\mathbf{y}, -\mathbf{f}). \quad (21)$$

It is necessary to truncate the infinite summations when evaluating empirical conditional variances. Truncation after N terms of the summations in (21), (20) and (19) respectively give the MA(N), AR(N) and ARMA(N , 1) approximations :

$$\ln(h_t) = \mathbf{a} + g(z_{t-1}) + \sum_{j=1}^N \mathbf{y}_j g(z_{t-j-1}), \quad (22)$$

$$\ln(h_t) = \mathbf{a} + \sum_{j=1}^N \mathbf{f}_j [\ln(h_{t-j}) - \mathbf{a}] + g(z_{t-1}), \quad (23)$$

$$\ln(h_t) = \mathbf{a} + \sum_{j=1}^N b_j [\ln(h_{t-j}) - \mathbf{a}] + g(z_{t-1}) + \mathbf{y}g(z_{t-2}). \quad (24)$$

As $j \rightarrow \infty$, the coefficients b_j and \mathbf{f}_j converge much more rapidly to zero than the coefficients \mathbf{y}_j . Consequently it is best to use either the AR or the ARMA approximation.

3.2. Estimates for the S & P 100 index

Representative parameters are required in Section 5 to illustrate the consequences of long memory in volatility for pricing options. As ARCH specifications are preferred for these illustrations, a discussion is presented here of parameter values for the FIEGARCH(1, d , 1) specification. These parameters are estimated from daily returns

r_t for the S & P 100 index, excluding dividends, calculated from index levels p_t as

$$r_t = \ln(p_t / p_{t-1}).$$

Evaluation of the conditional variances requires truncation of the infinite series defined by the fractional differencing filter. Here the variances are evaluated for $t \geq 1$ by setting $N = 1000$ in equation (24), with $\ln(h_{t-j})$ replaced by \mathbf{a} and $g(z_{t-j})$ replaced by zero whenever $t - j \leq 0$. The log-likelihood function is calculated for the 2,528 trading days during the ten-year estimation period from 3 January 1989 to 31 December 1998, which corresponds to the times $1,221 \leq t \leq 3,748$ for our dataset; thus the first 1,220 returns are reserved for the calculation of conditional variances before 1989 which are needed to evaluate the subsequent conditional variances.

Results are first discussed and are tabulated when returns have a constant conditional mean which is estimated by the sample mean. The conditional variances are obtained recursively from equations (15), (17), (19) and (24). The conditional distributions are assumed to be Normal when defining the likelihood function. This assumption is known to be false but it is made to obtain consistent parameter estimates (Bollerslev and Wooldridge, 1992). Preliminary maximisations of the likelihood showed that a suitable value for $C = E[|z_t|]$ is 0.737, compared with $\sqrt{2/\pi} \cong 0.798$ for the standard Normal distribution. They also showed that an appropriate value of the location parameter \mathbf{a} of $\ln(h_t)$ is -9.56; the log-likelihood is not sensitive to minor deviations from this level because \mathbf{a} is multiplied by a term $1 - \sum_{j=1}^N b_j$ in equation (24) that is small for large N . Consequently, the results summarised in Table 1 are given by

maximising the log-likelihood function over some or all of the parameters

$\mathbf{q}, \mathbf{g}, \mathbf{f}, \mathbf{y}, d$.

The estimates of \mathbf{q} and \mathbf{g} provide the usual result for a series of U.S. stock index returns that changes in volatility are far more sensitive to the values of negative returns than those of positive returns, as first reported by Nelson (1991). When z_t is negative, $g(z_t) = (\mathbf{g} - \mathbf{q})(-z_t) - \mathbf{g}C$, otherwise $g(z_t) = (\mathbf{g} + \mathbf{q})z_t - \mathbf{g}C$. The ratio $\frac{\mathbf{g} - \mathbf{q}}{\mathbf{g} + \mathbf{q}}$ is at least 4 and hence is substantial for the estimates presented in Table 1.

The first two rows of Table 1 report estimates for short memory specifications of the conditional variance. The AR(1) specification has a persistence of 0.982 that is typical for this volatility model. The ARMA(1,1) specification has an additional parameter and increases the log-likelihood by 3.0. The third row shows that the fractional differencing filter alone ($d > 0$, $\mathbf{f} = \mathbf{y} = 0$) provides a better description of the volatility process than the ARMA(1,1) specification; with $d = 0.66$ the log-likelihood increases by 10.9. A further increase of 7.8 is then possible by optimising over all three volatility parameters, d , \mathbf{f} and \mathbf{y} , to give the parameter estimates¹ in the fifth row of Table 1.

The estimates for the most general specification identify two issues of concern. First, d equals 0.57 for our daily data which is more than the typical estimate of 0.4

¹ The log-likelihood function is maximised using a complete enumeration algorithm and hence standard errors are not immediately available. A conservative robust standard error for our estimate of d is 0.12, using information provided by Bollerslev and Mikkelsen (1996).

produced by the studies of higher frequency data mentioned in Section 2.3. The same issue arises in Bollerslev and Mikkelsen (1996) with d estimated as 0.63 (standard error 0.06) from 9,559 daily returns of the S & P 500 index, from 1953 to 1990; there are similar results in Bollerslev and Mikkelsen (1999). Second, the sum $d + \mathbf{f} + \mathbf{y}$ equals 1.39. As this sum equals \mathbf{y}_1 in equations (21) and (22), more weight is then given to the volatility shock at time $t - 2$ than to the shock at time $t - 1$ when calculating $\ln(h_t)$. This is counterintuitive. To avoid this outcome, the constraint $d + \mathbf{f} + \mathbf{y} \leq 1$ is applied and the results given in the penultimate row of Table 1 are obtained. The log-likelihood is then reduced by 2.0. Finally, if d is constrained to be 0.4 then the log-likelihood is reduced by an additional 8.3.

The estimates obtained here for \mathbf{f} and \mathbf{y} , namely -0.27 and 0.68 for the most general specification, are rather different to the 0.78 and -0.68 given by Bollerslev and Mikkelsen (1999, Table 1), although the estimates of d are similar, namely 0.59 and 0.65. However, the moving-average representations obtained from these sets of parameters estimates are qualitatively similar. This is shown on Figure 1 that compares the moving-average coefficients \mathbf{y}_j defined by equation (21). The coefficients are positive and monotonic decreasing for the four sets of parameter values used to produce Figure 1. They show the expected hyperbolic decay when $d > 0$ and a geometric decay when $d = 0$. The values of b_j in equations (19) and (24) that are used to calculate the conditional variances decay much faster. For each of the four curves shown on Figure 1, $\mathbf{y}_{10} > 0.33$ and $\mathbf{y}_{100} > 0.07$ whilst $0 < b_{10} < 0.02$ and $0 < b_{100} < 0.0003$.

The results reported in Table 1 are for a constant conditional mean, $\mathbf{m}_t = \mathbf{m}$. Alternative specifications such as $\mathbf{m}_t = \mathbf{m} + \mathbf{b}r_{t-1}$, $\mathbf{m}_t = \mathbf{m} - \frac{1}{2}h_t$ and $\mathbf{m}_t = \mathbf{m} + \mathbf{I}\sqrt{h_t}$ give similar values of the log-likelihood when the volatility parameters are set to the values in the final row of Table 1. First, including the lagged return r_{t-1} is not necessary because the first-lag autocorrelation of the S&P 100 returns equals -0.022 and is statistically insignificant. Second, including the adjustment $-\frac{1}{2}h_t$ makes the conditional expectation of $\frac{P_t - P_{t-1}}{P_{t-1}}$ constant when the conditional distribution is Normal. The adjustment reduces the log-likelihood by an unimportant 0.3. Third, incorporating the ARCH-M parameter \mathbf{I} gives an optimal value of 0.10 and an increase in the log-likelihood of 1.5. This increase is not significant using a non-robust likelihood-ratio test at the 5% level.

3.3. Stochastic volatility specifications

Two shocks per unit time characterise stochastic volatility (SV) models, in contrast to the single shock z_t that appears in ARCH models. A general framework for long memory stochastic volatility models is given for returns r_t by

$$r_t = \mathbf{m} + \mathbf{S}_t u_t \quad (25)$$

with $\ln(\mathbf{S}_t)$ following an ARFIMA(p, d, q) process. For example, with $p = q = 1$,

$$\ln(\mathbf{S}_t) = \mathbf{a} + (1 - \mathbf{f}L)^{-1}(1 - L)^{-d}(1 + \mathbf{y}L)v_t. \quad (26)$$

This framework has been investigated by Breidt, Crato and de Lima (1998), Harvey (1998) and Bollerslev and Wright (2000), all of whom provide results for the simplifying assumption that the two i.i.d. processes $\{u_t\}$ and $\{v_t\}$ are independent. This assumption can be relaxed and has been for short memory applications (Taylor, 1994, Shephard, 1996).

Parameter estimation is difficult for SV models, compared with ARCH models, because SV models have twice as many random innovations as observable variables. Breidt, Crato and de Lima (1998) describe a spectral-likelihood estimator and provide results for a CRSP index from 1962 to 1989. For the ARFIMA(1, d , 0) specification they estimate $d = 0.44$ and $f = 0.93$. Bollerslev and Wright (2000) provide detailed simulation evidence about semiparametric estimates of d , related to the frequency of the observations.

It is apparent that the ARCH specification (15)-(17) has a similar structure to the SV specification (25)-(26). Short memory special cases of these specifications, given by $d = q = 0$, have similar moments (Taylor, 1994). This is a consequence of the special cases having the same bivariate diffusion limit when appropriate parameter values are defined for increasingly frequent observations (Nelson, 1990, Duan, 1997). It seems reasonable to conjecture that the multivariate distributions for returns defined by (15)-(17) and (25)-(26) are similar, with the special case of independent shocks $\{u_t\}$ and $\{v_t\}$ corresponding to the symmetric ARCH model that has $\mathbf{q} = 0$ in equation (17).

4. Option pricing methodology

4.1. A review of SV and ARCH methods

The pricing of options when volatility is stochastic and has a short memory has been studied by several researchers using a variety of methods. The most popular methods commence with separate diffusion specifications for the asset price and its volatility. These are called stochastic volatility (SV) methods. Option prices then depend on several parameters including a volatility risk premium and the correlation between the differentials of the Wiener processes in the separate diffusions. Hull and White (1987) provide solutions that include a simple formula when volatility risk is not priced and the correlation between the differentials is zero. The closed form solution of Heston (1993) assumes that volatility follows a square-root process and permits a general correlation and a non-zero volatility risk premium; for applications see, for example, Bakshi, Cao and Chen (1997, 2000) and for extensions see Duffie, Pan and Singleton (2000).

There is much less research into option pricing for short memory ARCH models. Duan (1995) provides a valuation framework and explicit results for the GARCH(1,1) process that can be extended to other ARCH specifications. Ritchken and Trevor (1999) provide an efficient lattice algorithm for GARCH(1,1) processes and extensions for which the conditional variance depends on the previous value and the latest return innovation.

Methods for pricing options when volatility has a long memory have been described by Comte and Renault (1998) and Bollerslev and Mikkelsen (1996, 1999). The former authors provide analysis within a bivariate diffusion framework. They replace the usual Wiener process in the volatility equation by fractional Brownian motion. However, their option pricing formula appears to require independence between the Wiener process in the price equation and the volatility process. This assumption is not consistent with the empirical evidence for stock returns. The assumption is refuted, for example, by finding that \mathbf{q} is not zero in the function $g(z_t)$ that appears in an exponential ARCH model.

The most practical way to price options with long memory in volatility is probably based upon ARCH models, as demonstrated by Bollerslev and Mikkelsen (1999). We follow the same strategy. From the asymptotic results in Duan (1997), also discussed in Ritchken and Trevor (1999), it is anticipated that insights about options priced from a long memory ARCH model will be similar to the insights that can be obtained from a related long memory SV model.

4.2. The ARCH pricing framework

When pricing options it will be assumed that returns are calculated from prices (or index levels) as $r_t = \ln(p_t/p_{t-1})$ and hence exclude dividends. A constant risk-free interest rate and a constant dividend yield will also be assumed and, to simplify the notation and calculations, it will be assumed that interest and dividends are paid once

per trading period. Conditional expectations are defined with respect to current and prior price information represented by $I_t = \{p_{t-i}, i \geq 0\}$.

To obtain fair option prices in an ARCH framework it is necessary to make additional assumptions in order to obtain a risk-neutral measure Q . Duan (1995) and Bollerslev and Mikkelsen (1999) provide sufficient conditions to apply a risk-neutral valuation methodology. For example, it is sufficient that a representative agent has constant relative risk aversion and that returns and aggregate growth rates in consumption have conditional normal distributions. Kallsen and Taqqu (1998) derive the same solution as Duan (1995) without making assumptions about utility functions and consumption. Instead, they assume that intraday prices are determined by geometric Brownian motion with volatility determined once a day from a discrete-time ARCH model.

At time t' , measured in trading periods, the fair price of an European contingent claim that has value $y_{t'+n}(p_{t'+n})$ at the terminal time $t'+n$ is given by

$$y_{t'} = E^Q[e^{-rn} y_{t'+n}(p_{t'+n}) | I_{t'}] \quad (27)$$

with r the risk-free interest rate for one trading period. Our objective is now to specify an appropriate way to simulate $p_{t'+n}$ under a risk-neutral measure Q and thereby to evaluate the above conditional expectation using Monte Carlo methods.

Following Duan (1995), it is assumed that observed returns are obtained under a probability measure P from

$$r_t | I_{t-1} \sim^P N(\mathbf{m}_t, h_t), \quad (28)$$

with

$$z_t = \frac{r_t - \mathbf{m}_t}{\sqrt{h_t}} \sim^P i.i.d. N(0,1) \quad (29)$$

and that in a risk-neutral framework returns are obtained under measure Q from

$$r_t | I_{t-1} \sim^Q N\left(\mathbf{r} - \mathbf{d} - \frac{1}{2}h_t, h_t\right), \quad (30)$$

with

$$z_t^* = \frac{r_t - \left(\mathbf{r} - \mathbf{d} - \frac{1}{2}h_t\right)}{\sqrt{h_t}} \sim^Q i.i.d. N(0,1). \quad (31)$$

Here \mathbf{d} is the dividend yield, that corresponds to a dividend payment of $d_t = (e^{\mathbf{d}} - 1)p_t$ per share at time t . Then $E^Q[p_t | I_{t-1}] = e^{\mathbf{r} - \mathbf{d}} p_{t-1}$ and the expected value at time t of one share and the dividend payment is $E^Q[p_t + d_t | I_{t-1}] = e^{\mathbf{r}} p_{t-1}$, as required in a risk-neutral framework.

Note that the conditional means are different for measures P and Q but the functions $h_t(p_{t-1}, p_{t-2}, \dots)$ that define the conditional variances for the two measures are identical. Duan (1995) proves that this is a consequence of the sufficient assumptions that he states about risk preferences and distributions. The same conclusion applies for the less restrictive assumptions of Kallsen and Taquq (1998).

Option prices depend on the specifications for \mathbf{m}_t and h_t . We again follow Duan (1995) and assume that

$$\mathbf{m}_t = \mathbf{r} - \mathbf{d} - \frac{1}{2}h_t + \mathbf{I}\sqrt{h_t} \quad (32)$$

with \mathbf{I} representing a risk-premium parameter. Then the conditional expectations of r_t for measures P and Q differ by $\mathbf{I}\sqrt{h_t}$ and

$$z_t - z_t^* = -\mathbf{1}. \quad (33)$$

4.3. Long memory ARCH equations

Option prices are evaluated in this paper when the conditional variances are given by the ARMA($N, 1$) approximation to the FIEGARCH($1, d, 1$) specification. From (17), (24) and (33),

$$\left(1 - \sum_{j=1}^N b_j L^j\right) (\ln(h_t) - \mathbf{a}) = (1 + \mathbf{y}L)g(z_{t-1}) = (1 + \mathbf{y}L)g(z_{t-1}^* - \mathbf{1}), \quad (34)$$

and

$$g(z_t) = \mathbf{q}z_t + \mathbf{g}(|z_t| - C) \quad (35)$$

with the autoregressive coefficients b_j defined by (19) as functions of \mathbf{f} and d ; also $C = \sqrt{2/\mathbf{p}}$ for conditional normal distributions². Suppose there are returns observed at times $1 \leq t \leq t'$, whose distributions are given by measure P , and that we then want to simulate returns for times $t > t'$ using measure Q . Then $\ln(h_t)$ is calculated for $1 \leq t \leq t'+1$ using the observed returns, with $\ln(h_t) = \mathbf{a}$ and $g(z_t) = 0$ for $t < 1$, followed by simulating $z_t^* \sim^Q N(0,1)$ and hence obtaining r_t and $\ln(h_{t+1})$ for $t > t'$.

² When $z \sim N(0,1)$, $E[|z - \mathbf{1}|] = \sqrt{2/\mathbf{p}} \exp(-\frac{1}{2}\mathbf{1}^2) + \mathbf{1}(2\Phi(\mathbf{1}) - 1)$ with Φ the cumulative distribution function of z .

The expectation of $\ln(h_t)$ depends on the measure³ when $\mathbf{I} \neq 0$. It equals \mathbf{a} for measure P . It is different for measure Q because

$$E^Q \left[g(z_t^* - \mathbf{I}) \right] = -\mathbf{I}\mathbf{q} + \mathbf{g} \left(E^Q \left[|z_t^* - \mathbf{I}| \right] - \sqrt{2/\mathbf{p}} \right) \cong -\mathbf{I}\mathbf{q} + \frac{\mathbf{I}^2 \mathbf{g}}{\sqrt{2\mathbf{p}}} \quad (36)$$

when \mathbf{I} is small, and this expectation is in general not zero. For a fixed t' , as $t \rightarrow \infty$,

$$E^Q \left[\ln(h_t) | I_{t'} \right] \rightarrow \mathbf{a} + \left(1 - \sum_{j=1}^N b_j \right)^{-1} (1 + \mathbf{y}) E^Q \left[g(z_t^* - \mathbf{I}) \right]. \quad (37)$$

The difference between the P and Q expectations of $\ln(h_t)$ could be interpreted as a volatility risk premium. This "premium" is typically negative, because typically $\mathbf{I} > 0$, $\mathbf{q} \leq 0$ and $\mathbf{g} > 0$. Furthermore, when \mathbf{q} is negative the major term in (36) is $-\mathbf{I}\mathbf{q}$, because \mathbf{I} is always small, and then the "premium" reflects the degree of asymmetry in the volatility shocks $g(z_t)$.

The magnitude of the volatility "risk premium" can be important and, indeed, the quantity defined by the limit in (37) becomes infinite⁴ as $N \rightarrow \infty$ when d is positive. A realistic value of \mathbf{I} for the S & P 100 index is 0.028, obtained by assuming that the equity risk premium is 6% per annum⁵. For the short memory parameter values in the first row of Table 1, when $d = 0$ and $N = 1000$, the limit of $E^Q \left[\ln(h_t) | I_{t'} \right] - \mathbf{a}$

³ The dependence of moments of h_t on the measure is shown by Duan (1995, p. 19) for the GARCH(1,1) model.

⁴ As $(1-L)^d 1 = 0$ for $d > 0$, it follows from (18) and (19) that $\sum_{j=1}^{\infty} a_j = \sum_{j=1}^{\infty} b_j = 1$.

⁵ The conditional expectations of r_t for measures P and Q differ by $\mathbf{I}\sqrt{h_t}$ and a typical average value of $\sqrt{h_t}$ is 0.00858. Assuming 253 trading days in one year gives the stated value of \mathbf{I} .

equals 0.10. This limit increases to 0.20 for the parameter values in the final row of Table 1, when $d = 0.4$ and $N = 1000$. The typical effect of adding 0.2 to $\ln(h_t)$ is to multiply standard deviations $\sqrt{h_t}$ by 1.1 so that far-horizon expected volatilities, under Q , are slightly higher than might be expected from historical standard deviations. Consequently, on average the term structure of implied volatilities will slope upwards.

4.4. Numerical methods

The preceding equations can be used to value an European contingent claim at time t' by simulating prices under the risk-neutral measure Q , followed by estimating the expected discounted terminal payoff at time $t'+n$ as stated in equation (27). Two variations on these equations are used when obtaining representative results in Section 5. First, the specification of \mathbf{m}_t can be different for times on either side of t' to allow for changes through time in risk-free interest rates and risk premia. Equation (32) is then replaced by

$$\begin{aligned}\mathbf{m}_t &= m - \frac{1}{2}h_t + \mathbf{I}'\sqrt{h_t}, & t \leq t', \\ &= \mathbf{r} - \mathbf{d} - \frac{1}{2}h_t + \mathbf{I}'\sqrt{h_t}, & t > t'.\end{aligned}\tag{38}$$

Second, because the observed conditional distributions are not Normal whilst the simulations assume that they are, it is necessary to define

$$\begin{aligned}C &= C', & t \leq t', \\ &= \sqrt{2/\mathbf{p}}, & t > t'.\end{aligned}\tag{39}$$

for a constant C' estimated from observed returns. An alternative method, described by Bollerslev and Mikkelsen (1999), is to simulate from the sample distribution of standardised observed returns.

Standard variance reduction techniques can be applied to increase the accuracy of Monte Carlo estimates of contingent claim prices. A suggested antithetic method uses one i.i.d. $N(0,1)$ sequence $\{z_t^*\}$ to define the further i.i.d. $N(0,1)$ sequences, $\{-z_t^*\}$, $\{z_t^\bullet\}$ and $\{-z_t^\bullet\}$ with the terms z_t^\bullet chosen so that there is negative correlation between $|z_t^*|$ and $|z_t^\bullet|$; this is achieved by defining $\Phi(z_t^\bullet) + \Phi(z_t^*) = 1 + \frac{1}{2} \text{sign}(z_t^*)$. The four sequences provide claim prices whose average, \bar{y} say, is much less variable than the claim price from a single sequence. An overall average \hat{y} is then obtained from a set of K values $\{\bar{y}_k, 1 \leq k \leq K\}$.

The control variate method makes use of an unbiased estimate \hat{y}_{CV} of a known parameter y_{CV} , such that \hat{y} is positively correlated with \hat{y}_{CV} . A suitable parameter, when pricing a call option in an ARCH framework, is the price of a call option when volatility is deterministic. The deterministic volatility process is defined by replacing all terms $\ln(h_t)$, $t > t'+1$, by their expectations under P conditional on the history $I_{t'}$. Then y_{CV} is given by a simple modification of the Black-Scholes formula, whilst \hat{y}_{CV} is obtained by using the same $4K$ sequences of i.i.d. variables that define \hat{y} . Finally, a more accurate estimate of the option price is then given by $\tilde{y} = \hat{y} - \mathbf{b}(\hat{y}_{CV} - y_{CV})$ with \mathbf{b} chosen to minimise the variance of \tilde{y} .

5. Illustrative long memory option prices

5.1. Inputs

The Black-Scholes formula has six parameters, when an asset pays dividends at a constant rate, namely the current asset price S , the time until the exercise decision T , the exercise price X , the risk-free rate R , the dividend yield D and the volatility σ . There are many more parameters and additional inputs for the FIEGARCH option pricing methodology described in Section 4. To apply that methodology to value European options it is necessary to specify eighteen numbers, a price history and a random number generator, as follows :

- Contractual parameters - time until exercise T measured in years, the exercise price X and whether a call or a put option.
- The current asset price $S = p_t$, and a set of previous prices $\{p_t, 1 \leq t < t'\}$.
- Trading periods per annum M , such that consecutive observed prices are separated by $1/M$ years and likewise for simulated prices $\{p_t, t' < t \leq t'+n\}$ with $n = MT$.
- Risk-free annual interest rate R , from which the trading period rate $r = R/M$ is obtained.
- Annual dividend yield D giving a constant trading period payout rate of $d = D/M$; both R and D are continuously compounded and applicable for the life of the option contract.

- The risk premium \mathbf{I} for investment in the asset during the life of the option, such that one-period conditional expected returns are $\mathbf{m}_t = \mathbf{r} - \mathbf{d} - \frac{1}{2}h_t + \mathbf{I}\sqrt{h_t}$.
- Parameters m and \mathbf{I}' that define conditional expected returns during the time period of the observed prices by $\mathbf{m}_t = m - \frac{1}{2}h_t + \mathbf{I}'\sqrt{h_t}$.
- Eight parameters that define the one-period conditional variances h_t . The integration level d , the autoregressive parameter \mathbf{f} and the truncation level N determine the parameters b_j (given by equation (19)) of the AR(N) filter in equation (34). The mean \mathbf{a} and the moving-average parameter \mathbf{y} complete the ARMA($N, 1$) specification for $\ln(h_t)$ in equation (34). The values of the shocks to the ARMA($N, 1$) process depend on \mathbf{g} and \mathbf{q} , that respectively appear in the symmetric function $\mathbf{g}(z_t|C)$ and the asymmetric function $\mathbf{q}z_t$ whose total determines the shock term $g(z_t)$; the constant C is a parameter C' for observed prices but equals $\sqrt{2/\mathbf{p}}$ when returns are simulated.
- K , the number of independent simulations of the terminal asset price $S_T = p_{t'+n}$.
- A set of Kn pseudo-random numbers distributed uniformly between 0 and 1, from which pseudo-random standard normal variates can be obtained. These numbers typically depend on a seed value and a deterministic algorithm.

5.2. Parameter selections

Option values are tabulated for hypothetical European options on the S & P 100 index. Options are valued for ten dates defined by the last trading days of the ten years from 1989 to 1998 inclusive. For valuation dates from 1992 onwards the size of the price history is set at $t' = 2000$; for previous years the price history commences on 6 March 1984 and $t' < 2000$. It is assumed that there are $M = 252$ trading days in one year and hence exactly 21 trading days in one simulated month. Option values are tabulated when T is 1, 2, 3, 6, 12, 18 and 24 months.

Table 2 lists the parameter values used to obtain the main results. The annualised risk-free rate and dividend yield are set at 5% and 2% respectively. The risk parameter I is set at 0.028 to give an annual equity risk premium of 6% (see footnote 5). The mean return parameter m is set to the historic mean of the complete set of S & P 100 returns from March 1984 to December 1998 and I' is set to zero.

There are two sets of values for the conditional variance process because the primary objective here is to compare option values when volatility is assumed to have either a short or a long memory. The long memory parameter set takes the integration level to be $d = 0.4$ because this is an appropriate level based upon the recent evidence from high-frequency data, reviewed in Section 2.3. The remaining variance parameters are then based on Table 1; as the moving-average parameter is small it is set to zero and the autoregressive parameter is adjusted to retain the unit total, $d + \mathbf{f} + \mathbf{y} = 1$. The AR

filter⁶ is truncated at lag 1000, although the results obtained will nevertheless be referred to as long memory results. The short memory parameters are similar to those for the AR(1) estimates provided in Table 1. The parameters \mathbf{g} and \mathbf{q} are both 6% less in Table 2 than in Table 1 to ensure that selected moments are matched for the short and long memory specifications; the unconditional mean and variance⁷ of $\ln(h_t)$ are then matched for the historic measure P , although the unconditional means differ by approximately 0.10 for the risk-neutral measure Q as noted in Section 4.3.

Option prices are estimated from $K = 10,000$ independent simulations of prices $\{p_t, t' < t \leq t'+n\}$ with $n = 504$. Applying the antithetic and control variate methods described in Section 4.4 then produces results for a long memory process in about 50 minutes, using a PC running at 466 MHz. Most of the time is spent evaluating the high-order AR filter; the computation time is less than 5 minutes for the short memory

⁶ This filter is

$$\begin{aligned} (1-0.6L)(1-L)^{0.4} &= 1 - (L - 0.12L^2 - 0.008L^3 + 0.0032L^4 + 0.0050L^5 + 0.0050L^6 + \dots) \\ &= 1 - \sum_{j=1}^{\infty} b_j L^j. \end{aligned}$$

Also, $b_6 > b_j > b_{j+1} > 0$ for $j > 6$, $b_{100} = 0.00017$, $b_{1000} = 7 \times 10^{-6}$

and $\sum_{j=1}^{1000} b_j$ equals 0.983.

⁷ I thank Granville Tunncliffe-Wilson for calculating the variance of the AR(1000) process.

process. Separate seed values are used to commence the "random number" calculations⁸ for the ten valuation dates; these seed values are always the same for calculations that have the same valuation date.

5.3. Comparisons of implied volatility term structures

The values of all options are reported using annualised implied volatilities rather than prices. Each implied volatility (IV) is calculated from the Black-Scholes formula, adjusted for continuous dividends. The complete set of IV outputs for one set of inputs forms a matrix with rows labelled by the exercise prices X and columns labelled by the times to expiry T ; examples are given in Tables 5 and 6 and are discussed later.

Initially we only consider at-the-money options, for which the exercise price equals the forward price $F = Se^{(R-D)T}$, with IV values obtained by linear interpolation across two adjacent values of X . As T varies, the IV values represent the term structure of implied volatility. Tables 3 and 4 respectively summarise these term structures for the short and long memory specifications. The same information is plotted on Figures 2 and 3 respectively. The IV values for $T = 0$ are obtained from the conditional variances on the valuation dates. The standard errors of the tabulated implied volatilities increase with T . The maximum standard errors for at-the-money options are respectively 0.0003 and 0.0004 for the short and long memory specifications.

⁸ The Excel VBA pseudo-random number generator was used. This generator has cycle length 2^{24} . Use is made of 30% of the complete cycle when $K = 10,000$ and $n = 504$.

The ten IV term structures for the short memory specification commence between 9.5% (1993) and 18.8% (1997) and converge towards the limiting value of 14.3%. The initial IV values are near the median level from 1989 to 1991, are low from 1992 to 1995 and are high from 1996 to 1998. Six of the term structures slope upwards, two are almost flat and two slope downwards. The shapes of these term structures are completely determined by the initial IV values because the volatility process is Markovian.

There are three clear differences between the term structures for the short and long memory specifications that can be seen by comparing Figures 2 and 3. First, the long memory term structures can and do intersect because the volatility process is not Markovian. Second, some of the term structures have sharp kinks for the first month. This is particularly noteworthy for 1990 and 1996 when the term structures are not monotonic. For 1990, the initial value of 14.1% is followed by 15.6% at one month and a gradual rise to 16.2% at six months and a subsequent slow decline. For 1996, the term structure commences at 15.6%, falls to 13.6% after one month and reaches a minimum of 12.8% after six months followed by a slow incline. The eight other term structures are monotonic and only those for 1997 and 1998 slope downwards. Third, the term structures approach their limiting value very slowly⁹. The two-year IVs range from 12.1% to 16.1% and it is not possible to deduce the limiting value, although 15.0% to

⁹ The results support the conjecture that $IV(T) \cong a_1 + a_2 T^{2d-1}$ for large T with a_2 determined by the history of observed returns.

16.0% is a plausible range¹⁰. It is notable that the dispersion between the ten IV values for each T decreases slowly as T increases, from 2.2% for one-month options to 1.4% for two-year options.

There are substantial differences between the two IV values that are calculated for each valuation date and each option lifetime. Figure 4 shows the differences between the at-the-money IVs for the long memory specification minus the number for the short memory specification. When $T = 0$ these differences range from -1.9% (1997) to 1.5% (1992), for three-month options from -1.5% (1995, 1996) to 2.1% (1990) and for two-year options from -1.9% (1995) to 1.7% (1990). The standard deviation of the ten differences is between 1.1% and 1.4% for all values of T considered so it is common for the short and long memory option prices to have IVs that differ by more than 1%.

5.4. Comparisons of smile effects

The columns of the IV matrix provide information about the strength of the so-called smile effect for options prices. These effects seem to be remarkably robust to the choice of valuation date and they are not very sensitive to the choice between the short and

¹⁰ An estimate of the constant a_1 (defined in the previous footnote) is 16.0%. An estimate of 15.0% follows by supposing the long memory limit is 105% of the short memory limit, based on the limit of $\ln(h_t)$ being higher by 0.1 for the long memory process as noted in Section 4.3. The difference in the limits is a consequence of the risk premium obtained by owning the asset; its magnitude is mainly determined by the pronounced asymmetry in the volatility shock function $g(z_t)$.

long memory specifications. This can be seen by considering the ten values of $\Delta IV = IV(T, X_1) - IV(T, X_2)$ obtained for the ten valuation dates, for various values of T , various pairs of exercise prices X_1, X_2 and a choice of volatility process. First, for one-month options with $S = 100$, $X_1 = 92$ and $X_2 = 108$, the values of ΔIV range from 3.0% to 3.3% for the short memory specification and from 3.7% to 4.0% for the long memory specification. Second, for two-year options with $X_1 = 80$ and $X_2 = 120$, the values of ΔIV range from 1.8% to 2.0% and from 1.8% to 1.9%, respectively for the short and long memory specifications.

Figure 5 shows the smiles for three-month options valued using the short memory model, separately for the ten valuation dates. As may be expected from the above remarks the ten curves are approximately parallel to each other. They are almost all monotonic decreasing for the range of exercise prices considered, so that a U-shaped function (from which the idea of a smile is derived) can not be seen. The near monotonic decline is a standard theoretical result when volatility shocks are negatively correlated with price shocks (Hull, 2000). It is also a stylized empirical fact for U.S. equity index options, see, for example, Rubinstein (1994) and Dumas, Fleming and Whaley (1998).

Figure 6 shows the three-month smiles for the long memory specification. The shapes on Figures 5 and 6 are similar, as all the curves are for the same expiry time, but they are more dispersed on Figure 6 because the long memory effect induces more dispersion in at-the-money IVs. The minima of the smiles are generally near an exercise price of 116. Figure 7 shows further long memory smiles, for two-year options

when the forward price is 106.2. The parallel shapes are clear; the two highest curves are almost identical, and the third, fourth and fifth highest curves are almost the same.

Tables 5 and 6 provide matrices of implied volatilities for options valued on 31 December 1998. When either the call or the put option is deep out-of-the-money it is difficult to estimate the option price accurately because the risk-neutral probability $q(X)$ of the out-of-the-money option expiring in-the-money is small. Consequently, the IV information has not been presented when the corresponding standard errors exceed 0.2%; estimates of $q(X)$ are less than 3%. The standard errors of the IVs are least for options that are near to at-the-money and most of them are less than 0.05% for the IVs listed in Tables 5 and 6. All the sections of the smiles summarised by Tables 5 and 6 are monotonic decreasing functions of the exercise price. The IV decreases by approximately 4% to 5% for each tabulated section.

5.5. Sensitivity analysis

The sensitivity of the IV matrices to three of the inputs has been assessed for options valued on 31 December 1998. First, consider a change to the risk parameter I that corresponds to an annual risk premium of 6% for the tabulated results. From Section 4.3, option prices should be lower for large T when I is reduced to zero. Changing I to zero reduces the at-the-money IV for two-year options from 16.0% to 15.4% for the long memory inputs, with a similar reduction for the short memory inputs. Second, consider reducing the truncation level N in the AR(N) filter from 1000 to 100. Although this has the advantage of a substantial reduction in the computational time it changes

the IV numbers by appreciable amounts and can not be recommended; for example, the two-year at-the-money IV then changes from 16.0% to 14.7%.

The smile shapes on Figures 5, 6 and 7 are heavily influenced by the negative asymmetric shock parameter \mathbf{q} , that is substantial relative to the symmetric shock parameter \mathbf{g} . The asymmetry in the smile shapes can be expected to disappear when \mathbf{q} is zero, which is realistic for some assets including exchange rates. Figures 8 and 9 compare smile shapes when \mathbf{q} is changed from the values used previously to zero, with \mathbf{g} scaled to ensure the variance of $\ln(h_t)$ is unchanged for measure P . Figure 8 shows that the one-month smile shapes become U-shaped when \mathbf{q} is zero, whilst Figure 9 shows that the IV are then almost constant for one-year options.

6. Conclusions

The empirical evidence for long memory in volatility is strong, for both equity (Andersen, Bollerslev, Diebold and Ebens, 2000, Areal and Taylor, 2000) and foreign exchange markets (Andersen, Bollerslev, Diebold and Labys, 2000). This evidence may more precisely be interpreted as evidence for long memory effects, because there are short memory processes that have similar autocorrelations and spectral densities, except at very low frequencies (Gallant, Hsu and Tauchen, 1999, Barndorff-Nielsen and Shephard, 2001). There is also evidence that people trade at option prices that are more compatible with a long memory process for volatility than with a parsimonious short memory process (Bollerslev and Mikkelsen, 1999).

The theory of option pricing when volatility follows a discrete-time ARCH process relies on weak assumptions about the continuous-time process followed by prices (Kallsen and Taqqu, 1998) and the numerical implementation of the theory is straightforward. Application of the theory when the volatility process is fractionally integrated does, however, require pragmatic approximations because the fundamental filter $(1-L)^d$ is an infinite order polynomial that must be truncated at some power N . Option prices are sensitive to the truncation point N , so that large values and long price histories from an assumed stationary process are required.

The term structure of implied volatility for at-the-money options can be notably different for short and long memory ARCH specifications applied to the same price history. Long memory term structures have more variety in their shapes. They may have kinks for short maturity options and they may not have a monotonic shape. Also, term structures on different valuation dates sometimes intersect each other. None of these possibilities occurs for a Markovian short memory specification. Long memory term structures do not converge rapidly to a limit as the lifetime of options increases. It is difficult to estimate the limit for the typical value $d = 0.4$.

Implied volatilities as functions of exercise prices have similar shapes for short and long memory specifications. The differences in these shapes are minor in comparison to the differences in the term structure shapes. It is common for the short and long memory implied volatilities to differ by more than 1% for options on the S & P 100 index, regardless of the option lifetime and the exercise price; if the short memory implied is at its average level of 14% then the long memory implied is often

below 13% or above 15%. Consequently, the economic consequences of a long memory assumption are important.

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, F.X. Diebold and H. Ebens, 2000, The distribution of stock return volatility, *Journal of Financial Economics*, forthcoming.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys, 2000, The distribution of exchange rate volatility, *Journal of the American Statistical Association*, forthcoming.
- Areal, N.M.P.C. and S.J. Taylor, 2000, The realized volatility of FTSE-100 futures prices, working paper, Lancaster University.
- Baillie, R.T., 1996, Long memory processes and fractional integration in econometrics, *Journal of Econometrics* 73, 5-59.
- Bakshi, G., C. Cao and Z. Chen, 1997, Empirical performance of alternative option pricing models, *Journal of Finance* 52, 2003-2049.
- Bakshi, G., C. Cao and Z. Chen, 2000, Pricing and hedging long-term options, *Journal of Econometrics* 94, 277-318.
- Barndorff-Nielsen, O.E. and N. Shephard, 2001, Non-Gaussian Ornstein-Uhlenbeck based models and some of their uses in financial economics, *Journal of the Royal Statistical Society B*, forthcoming.
- Bollerslev, T., 1986, Generalized autoregressive conditional heteroscedasticity, *Journal of Econometrics* 31, 307-327.
- Bollerslev, T., 1987, A conditionally heteroskedastic time series model for security prices and rates of return data, *Review of Economics and Statistics* 59, 542-547.
- Bollerslev, T. and H.O. Mikkelsen, 1996, Modeling and pricing long memory in stock market volatility, *Journal of Econometrics* 73, 151-184.
- Bollerslev, T. and H.O. Mikkelsen, 1999, Long-term equity anticipation securities and stock market volatility dynamics, *Journal of Econometrics* 92, 75-99.

- Bollerslev, T. and J.M. Wooldridge, 1992, Quasi-maximum likelihood estimation and inference in dynamic models with time varying covariances, *Econometric Reviews* 11, 143-172.
- Bollerslev, T. and J.H. Wright, 2000, Semiparametric estimation of long-memory volatility dependencies: the role of high-frequency data, *Journal of Econometrics* 98, 81-106.
- Breidt, F.J., N. Crato and P. de Lima, 1998, The detection and estimation of long memory in stochastic volatility, *Journal of Econometrics* 83, 325-348.
- Brockwell, P.J. and R.A. Davis, 1991, *Time Series: Theory and Methods*, second edition (Springer, New York).
- Comte, F. and E. Renault, 1998, Long memory in continuous-time stochastic volatility models, *Mathematical Finance* 8, 291-323.
- Dacorogna, M.M., U.A. Muller, R.J. Nagler, R.B. Olsen and O.V. Pictet, 1993, A geographical model for the daily and weekly seasonal volatility in the foreign exchange market, *Journal of International Money and Finance* 12, 413-438.
- Ding, Z., C.W.J. Granger and R.F. Engle, 1993, A long memory property of stock market returns and a new model, *Journal of Empirical Finance* 1, 83-106.
- Duan, J.-C., 1995, The GARCH option pricing model, *Mathematical Finance* 5, 13-32.
- Duan, J.-C., 1997, Augmented GARCH(p,q) process and its diffusion limit, *Journal of Econometrics* 79, 97-127.
- Duffie, D., J. Pan and K.J. Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica*, forthcoming.
- Dumas, B., J. Fleming and R.E. Whaley, 1998, Implied volatility functions : empirical tests, *Journal of Finance* 53, 2059-2106.
- Ebens, H., 1999, Realized stock volatility, working paper, Department of Economics, John Hopkins University.
- Ederington, L.H. and J.H. Lee, 1993, How markets process information : news releases and volatility, *Journal of Finance* 49, 1161-1191.
- Gallant, A.R., C.T. Hsu and G.E. Tauchen, 1999, Using high/low data to calibrate volatility diffusions and extract the forward integrated variance, working paper, Duke University.
- Granger, C.W.J., 1980, Long memory relationships and the aggregation of dynamic models, *Journal of Econometrics* 14, 227-238.
- Granger, C.W.J. and Z. Ding, 1996, Varieties of long memory models, *Journal of Econometrics* 73, 61-77.

- Granger, C.W.J. and R. Joyeux, 1980, An introduction to long-memory time series models and fractional differencing, *Journal of Time Series Analysis* 1, 15-29.
- Harvey, A.C., 1998, Long-memory in stochastic volatility, in J. Knight and S. Satchell (eds.), *Forecasting Volatility in Financial Markets* (Butterworth-Heinemann, London).
- Heston, S., 1993, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* 6, 327-343.
- Hosking, J., 1981, Fractional differencing, *Biometrika* 68, 165-176.
- Hull, J., 2000, *Options, Futures and Other Derivative Securities*, fourth edition (Prentice-Hall, New Jersey).
- Hull, J. and A. White, 1987, The pricing of options on assets with stochastic volatilities, *Journal of Finance* 42, 281-300.
- Kallsen, J. and M. Taqqu, 1998, Option pricing in ARCH-type models, *Mathematical Finance* 8, 13-26.
- McLeod, A.I. and K.W. Hipel, 1978, Preservation of the rescaled adjusted range, 1 : a reassessment of the Hurst phenomenon, *Water Resources Research* 14, 491-508.
- Nelson, D.B., 1990, ARCH models as diffusion approximations, *Journal of Econometrics* 45, 7-38.
- Nelson, D.B., 1991, Conditional heteroskedasticity in asset returns : a new approach, *Econometrica* 59, 347-370.
- Ritchken, P. and R. Trevor, 1999, Pricing options under generalized GARCH and stochastic volatility processes, *Journal of Finance* 54, 377-402.
- Rubinstein, M., 1994, Implied binomial trees, *Journal of Finance* 49, 771-818.
- Shephard, N., 1996, Statistical aspects of ARCH and stochastic volatility, in : D.R. Cox, D.V. Hinkley and O.E. Barndorff-Nielsen, eds., *Likelihood, Time Series with Econometric and Other Applications* (Chapman and Hall, London), 1-67.
- Taylor, S.J., 1986, *Modelling Financial Time Series* (John Wiley, Chichester, UK).
- Taylor, S.J., 1994, Modeling stochastic volatility : a review and comparative study, *Mathematical Finance* 4, 183-204.
- Taylor, S.J. and X. Xu, 1997, The incremental volatility information in one million foreign exchange quotations, *Journal of Empirical Finance* 4, 317-340.

Table 1. Parameter estimates for short and long memory ARCH models.

Parameters are estimated by maximising the log-likelihood of daily returns from the S & P 100 index, from 3 January 1989 to 31 December 1998. Returns are modelled as:

$$r_t = \mathbf{m} + \sqrt{h_t} z_t,$$

$$\ln(h_t) = \mathbf{a} + (1 - \mathbf{fL})^{-1} (1 - L)^{-d} (1 + \mathbf{yL})g(z_{t-1}),$$

$$g(z_t) = \mathbf{q}z_t + \mathbf{g}(|z_t| - C).$$

The z_t are assumed to be i.i.d., standard Normal variables when defining the likelihood function. The values $C = 0.737$ and $\mathbf{a} = -9.56$ are used for all likelihood calculations. The relative log-likelihood (LL) for a model equals the maximum LL for that model minus the maximum LL for the AR(1) model. The maximum log-likelihood for the AR(1) model is 8561.6.

<u>Model type</u>	<u>Constraints</u>	<u>\mathbf{q}</u>	<u>\mathbf{g}</u>	<u>\mathbf{f}</u>	<u>\mathbf{y}</u>	<u>d</u>	<u>Relative log-likelihood</u>
AR(1)	$\mathbf{y} = d = 0$	-0.06	0.10	0.982			0.0
ARMA(1,1)	$d = 0$	-0.09	0.15	0.988	-0.440		3.0
FI(d)	$\mathbf{f} = \mathbf{y} = 0$	-0.12	0.19			0.66	13.9
ARFI(1, d)	$\mathbf{y} = 0$	-0.11	0.17	0.30		0.59	17.2
ARFIMA(1, d , 1)	None	-0.10	0.15	-0.16	0.98	0.57	21.7
ARFIMA(1, d , 1)	$\mathbf{f} + \mathbf{y} + d \leq 1$	-0.12	0.18	-0.27	0.68	0.59	19.7
ARFIMA(1, d , 1)	$d = 0.4, \mathbf{f} + \mathbf{y} \leq 0.6$	-0.11	0.18	0.64	-0.04	0.4	11.4

Table 2. Parameter values for option price calculations.

Options are valued on the final trading day of ten consecutive years, from 1989 to 1998. The returns history is drawn from the set of daily returns from the S & P 100 index from 6 March 1984 to 31 December 1998. A set of $t' = 2000$ historical returns is used from 1992 onwards, and as many as possible before then. The current level of the index is reset to $S = 100$ when option values are determined.

Trading periods per annum M	252
Risk-free interest rate r	$0.05 / M$

Conditional mean

Historic intercept m	$0.161 / M$
Historic equity risk premium term $I'0$	
Dividend yield d	$0.02 / M$
Future equity risk premium term I	0.028

Conditional variance

	<u>Short memory</u>	<u>Long memory</u>
Integration level d	0	0.4
Truncation limit N		1000
Mean \mathbf{a} of $E^P[\ln(h_t)]$	-9.56	-9.56
Autoregressive parameter \mathbf{f}	0.982	0.6
Moving-average parameter \mathbf{y}	0	0
Asymmetric shock parameter \mathbf{q}	-0.056	-0.11
Symmetric shock parameter \mathbf{g}	0.094	0.18
Historic value of $E^P[z]$	0.737	0.737

Table 3. At-the-money implied volatilities for a short memory volatility process.

The parameters of the EGARCH price process are listed in Table 2. The half-life of a volatility shock is 1.8 months.

European option prices are estimated from 10,000 simulations of the asset prices on the dates that the options expire. The implied volatilities are for options whose exercise prices equal the forward rates for the expiry dates. The standard errors of the implied volatilities are between 0.0001 and 0.0003.

Options are valued on the final trading day of ten consecutive years. The returns history is drawn from the set of daily returns from the S & P 100 index from 6 March 1984 to 31 December 1998. A set of $t' = 2000$ historical returns is used from 1992 onwards, and as many as possible before then. The column for $T = 0$ provides the annualised volatilities on the valuation dates.

Implied volatilities for at-the-money options that
expire after 1, 2, 3, 6, 12, 18 and 24 months

	0	1	2	3	6	12	18	24
<u>Year</u>								
1989	0.1211	0.1238	0.1267	0.1291	0.1340	0.1384	0.1404	0.1414
1990	0.1378	0.1380	0.1388	0.1395	0.1409	0.1421	0.1427	0.1431
1991	0.1213	0.1240	0.1271	0.1296	0.1339	0.1383	0.1403	0.1413
1992	0.1085	0.1129	0.1175	0.1210	0.1288	0.1356	0.1385	0.1401
1993	0.0945	0.1008	0.1069	0.1119	0.1226	0.1321	0.1363	0.1382
1994	0.1165	0.1201	0.1236	0.1263	0.1321	0.1375	0.1398	0.1411
1995	0.1118	0.1158	0.1200	0.1234	0.1304	0.1367	0.1391	0.1405
1996	0.1462	0.1449	0.1446	0.1445	0.1442	0.1438	0.1440	0.1441
1997	0.1883	0.1791	0.1730	0.1683	0.1599	0.1526	0.1500	0.1486
1998	0.1694	0.1640	0.1607	0.1581	0.1531	0.1488	0.1476	0.1469
Mean	0.1315	0.1323	0.1339	0.1352	0.1380	0.1406	0.1419	0.1425
St. dev.	0.0291	0.0243	0.0205	0.0175	0.0116	0.0063	0.0043	0.0032

Table 4. At-the-money implied volatilities for a long memory volatility process.

The parameters of the FIEGARCH price process are listed in Table 2 and include an integration level of $d = 0.4$.

European option prices are estimated from 10,000 simulations of the asset prices on the dates that the options expire. The implied volatilities are for options whose exercise prices equal the forward rates for the expiry dates. The standard errors of the implied volatilities are all less than 0.0004.

Options are valued on the final trading day of ten consecutive years. The returns history is drawn from the set of daily returns from the S & P 100 index from 6 March 1984 to 31 December 1998. A set of $t' = 2000$ historical returns is used from 1992 onwards, and as many as possible before then. The column for $T = 0$ provides the annualised volatilities on the valuation dates.

Implied volatilities for at-the-money options that
expire after 1, 2, 3, 6, 12, 18 and 24 months

	0	1	2	3	6	12	18	24
<u>Year</u>								
1989	0.1194	0.1356	0.1403	0.1429	0.1467	0.1496	0.1507	0.1515
1990	0.1413	0.1556	0.1592	0.1609	0.1624	0.1624	0.1614	0.1607
1991	0.1215	0.1368	0.1416	0.1441	0.1478	0.1502	0.1510	0.1516
1992	0.1239	0.1261	0.1283	0.1301	0.1338	0.1375	0.1395	0.1409
1993	0.1080	0.1101	0.1128	0.1150	0.1195	0.1245	0.1272	0.1292
1994	0.1114	0.1189	0.1213	0.1228	0.1256	0.1284	0.1300	0.1314
1995	0.0965	0.1041	0.1067	0.1085	0.1119	0.1160	0.1187	0.1210
1996	0.1564	0.1357	0.1311	0.1295	0.1283	0.1288	0.1297	0.1308
1997	0.1697	0.1650	0.1626	0.1609	0.1574	0.1540	0.1523	0.1515
1998	0.1734	0.1693	0.1682	0.1672	0.1650	0.1623	0.1610	0.1602
Mean	0.1322	0.1357	0.1372	0.1382	0.1399	0.1414	0.1422	0.1429
St. dev.	0.0267	0.0221	0.0211	0.0204	0.0186	0.0165	0.0151	0.0141


Figure 1. Moving-average coefficients for four ARFIMA(1, d, 1) processes.


Figure 2. Ten volatility term structures for a short memory process.

Figure 3. Ten volatility term structures for a long memory process with $d = 0.4$.





Figure 4. Differences between ten pairs of term structures.

Figure 5. Ten smile shapes for three-month options and a short memory process.




Figure 6. Ten smile shapes for three-month options and a long memory process with $d = 0.4$.




Figure 7. Ten smile shapes for two-year options and a long memory process with $d = 0.4$.





Figure 8. Impact of asymmetric volatility shocks on one-month options.

Figure 9. Impact of asymmetric volatility shocks on one-year options.