ELSEVIER Journal of Econometrics 73 (1996) 151–184

JOURNAL OF Econometrics

Modeling and pricing long memory in stock market volatility

Tim Bollerslev[a,*], Hans Ole Mikkelsen[b]

[a]Department of Economics, University of Virginia, Charlottesville, VA 22901, USA
[b]School of Business Administration, University of Southern California, Los Angeles, CA 90089, USA

Abstract

A new class of fractionally integrated GARCH and EGARCH models for characterizing financial market volatility is discussed. Monte Carlo simulations illustrate the reliability of quasi maximum likelihood estimation methods, standard model selection criteria, and residual-based portmanteau diagnostic tests in this context. New empirical evidence suggests that the apparent long-run dependence in U.S. stock market volatility is best described by a mean-reverting fractionally integrated process, so that a shock to the optimal forecast of the future conditional variance dissipate at a slow hyperbolic rate. The asset pricing implications of this finding is illustrated via the implementation of various option pricing formula.

Key words: Fractional integrated EGARCH; Mean reversion; Model selection; Stock market volatility; Option pricing
JEL classification: C15; C22; C52; C13

1. Introduction

Much recent interest in econometrics and empirical finance has centered on modeling the temporal variation in financial market volatility. Particularly

* Corresponding author.

Tim Bollerslev is also affiliated with the NBER. Part of this paper was completed while Hans Ole Mikkelsen was visiting Northwestern University. We would like to thank two anonymous referees, Dick Baillie, Miguel Delgado, Francis X. Diebold, Andrew Harvey, Campbell Harvey, Daniel B. Nelson, Peter M. Robinson, and participants at the 1993 workshop on 'Modern Time Series Analysis in Finance' at the University of Aarhus, Denmark, the 1994 conference on 'Asymmetries and Non-Linearities in Dynamic Economic Models' in Madrid, the 1994 NBER Summer Institute, as well as seminar audiences at Amsterdam University, Tilburg University, University of Arizona, University of California at Santa Barbara, University of Minnesota, University of Montreal, University of Texas at Austin, and University of Virginia, for helpful comments and suggestions.

SSDI 0304-4076(95)01736-W

instrumental in these developments has been the Autoregressive Conditional Heteroskedastic (ARCH) class of models introduced by Engle (1982). In its most general form the ARCH model simply postulates the conditional variance to be a nontrivial function of the current information set. Unfortunately, economic theory offers little guidance as to which variables should be important in determining the observed time variation in the conditional variances. In light of this numerous ad hoc parametric formulations have been suggested in the literature. Two of the most successful such parameterizations for characterizing high-frequency financial market volatility have been the Generalized ARCH (GARCH) model introduced by Bollerslev (1986) and the Exponential GARCH (EGARCH) model proposed by Nelson (1991). The GARCH and the EGARCH models are readily interpreted as ARMA-type models for the conditional second-order moments and the logarithm of the conditional variance, respectively. A common finding in many empirical applications with both of these models concerns the apparent persistence of the estimated conditional variance processes; see Bollerslev, Chou, and Kroner (1992). The so-called Integrated GARCH (IGARCH) class of models was introduced by Engle and Bollerslev (1986) to capture this effect. In the IGARCH model a shock to the conditional variance remains important for the optimal forecasts of the variance for all future horizons. Thus, from a forecasting perspective the difference between the covariance-stationary GARCH formulation and the IGARCH model provides a natural analog to the difference between I(0) and I(1) type processes for the conditional mean.[1]

The distinction between I(0) or I(1) time series for the conditional mean may be far too narrow, however. An added flexibility is obtained by allowing for fractional orders of integration, as in the I(d) class of models introduced by Granger (1980), Granger and Joyeux (1980), Hosking (1981), and Mandelbrot and Van Ness (1968). In contrast to an I(0) time series in which shocks die out at an exponential rate, or an I(1) series in which there is no mean reversion, shocks to an I(d) time series with $0 < d < 1$ dissipate at a slow hyperbolic rate. The importance of this generalization in modeling long-run economic phenomena has recently been illustrated by a number of studies including Baillie and Bollerslev (1994), Baillie, Chung, and Tieslau (1996), Cheung and Lai (1993), Diebold, Husted, and Rush (1991), Lo (1991), and Sowell (1992). Just as the generalization of the standard ARIMA class of models to the fractionally integrated ARFIMA models have proven empirically important, a corresponding result may hold true when modeling long-term dependence in conditional variances. The new class of Fractionally Integrated GARCH (FIGARCH)

[1] As discussed further below, considerable care should be exercised in interpreting persistence in nonlinear models. The analogy to I(1) processes for the conditional mean is far from complete.

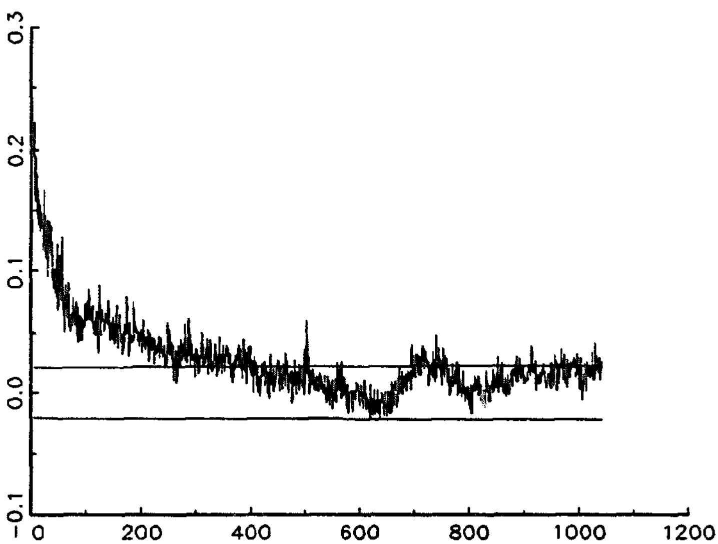

Fig. 1. Autocorrelations for absolute returns.

The figure graphs the lag 5 through 1040 sample autocorrelations for the absolute daily returns, $|r_t|$, on the Standard and Poor's 500 composite index from January 2, 1953 through December 31, 1990. The 95% confidence bands for no serial dependence are also indicated in the figure.

models proposed by Baillie, Bollerslev, and Mikkelsen (1996) allows for such an added flexibility. Analogous to the ARFIMA class of models for the conditional mean, a shock to the conditional variance in the FIGARCH model is transitory, in the sense that the influence on the forecast of the future conditional variance dies out at a slow hyperbolic rate of decay. In this paper we present some new results on the theoretical properties and the importance of allowing for such fractional unit roots in the conditional variance process.

In order to motivate the empirical relevance of these ideas, Fig. 1 plots the lag 5 through 1040 sample autocorrelations of the daily absolute returns, $|r_t|$, on the Standard and Poor's 500 composite index from January 2, 1953 through December 31, 1990. A more detailed description and analysis of the data is contained in Section 4 below. The volatility clustering phenomenon is immediately evident from this figure. The absolute return correlations for very long lags frequently exceed the two 95% Bartlett (1946) confidence bands for no serial dependence. Also, the Ljung and Box (1978) portmanteau test for the joint significance of lags 781 through 1040, corresponding roughly to a dependence between three to four years, equals 580.7, which is highly significant when tested in a chi-square distribution with 260 degrees of

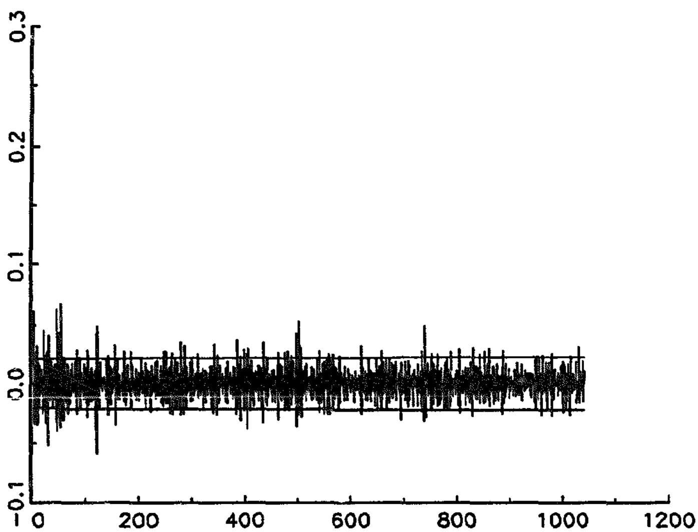

Fig. 2. Autocorrelations for the first difference of absolute returns.

The figure graphs the lag 5 through 1040 sample autocorrelations for the first difference of the absolute daily returns, $(1-L)|r_t|$, on the Standard and Poor's 500 composite index from January 2, 1953 through December 31, 1990. The 95% confidence bands for no serial dependence are also indicated in the figure.

freedom.[2] This apparent persistence in the autocorrelation function is substantially reduced in Fig. 2, which plots the sample autocorrelations for the first difference of the absolute returns; i.e., $(1-L)|r_t| \equiv |r_t| - |r_{t-1}|$. The portmanteau test for the joint significance of the lag 781 through 1040 autocorrelations is reduced to 340.1. Judged by the chi-square distribution the p-value is only 0.001, however. In contrast, on applying the fractional differencing operator $(1-L)^{0.5}$ to the absolute returns, where $(1-L)^d$ is formally defined in Section 2 below, the autocorrelations for the filtered series, $(1-L)^{0.5}|r_t|$, show much less long-term dependence. In particular, the same portmanteau test for no serial dependence beyond the three-year lag in the filtered series now equals 268.6, corresponding to a conventional p-value of 0.344. Of course, this particular test statistic may be even further reduced by judiciously choosing the value of d. This preliminary analysis therefore suggests that the important long-run features of financial

[2] The Bartlett standard errors and the Ljung–Box test statistic both assume that the variance of $|r_t|$ is constant, so that the conventional p-values reported below are merely indicative. We shall return to a more formal discussion of the distributional properties of the Ljung–Box statistic in the presence of ARCH in Section 3 below.

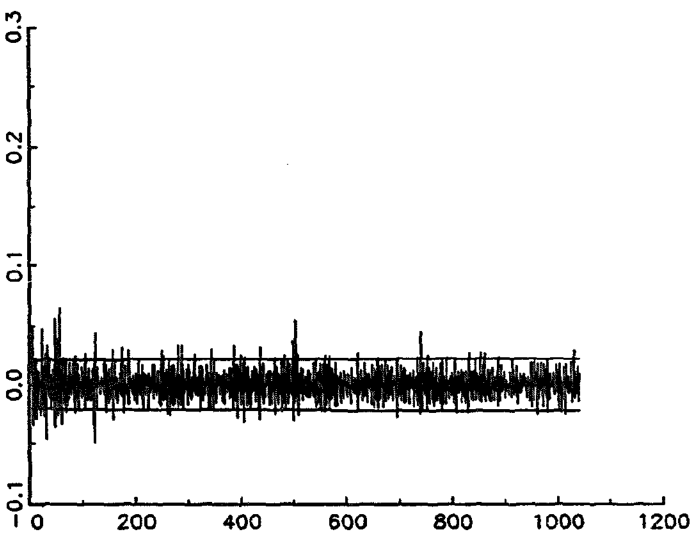

Fig. 3. Autocorrelations for the fractionally differenced absolute returns.
The figure graphs the lag 5 through 1040 sample autocorrelations for the fractionally differenced absolute daily returns, $(1-L)^{0.5}|r_t|$, on the Standard and Poor's 500 composite index from January 2, 1953 through December 31, 1990. The 95% confidence bands for no serial dependence are also indicated in the figure.

market volatility may be conveniently modeled by a fractionally integrated process.[3]

The plan for the rest of the paper is as follows. The basic definitions and theoretical properties of the new FIGARCH and FIEGARCH models are discussed in the next section. An approximate maximum likelihood estimation strategy is outlined in Section 3, which also reports the results from a small-scale Monte Carlo analysis of this estimator along with various model selection criteria and diagnostics in an ARCH context. Section 4 contains an empirical analysis of the daily Standard and Poor's 500 composite stock price index over the 1953 through 1990 period. The estimated degree of fractional integration across the different model formulations are all highly suggestive for the existence of long-memory features in the conditional variance of aggregate U.S. stock market volatility. The practical importance of this findings is further explored in

[3] Dacorogna, Müller, Nagler, Olsen, and Pictet (1993) also note that, when measured in terms of a market activity scale, the autocorrelation function for twenty-minute absolute returns on the U.S. dollar/Deutschmark exchange rate show a clear hyperbolic rate of decay up until a lag length of ten days.

Section 5 through the simulation of synthetic option prices for the estimated EGARCH, IEGARCH, and FIEGARCH data-generating mechanisms. Section 6 concludes.

2. Fractionally integrated ARCH

To set out the notation, let $\{\varepsilon_t\}$ denote a discrete-time real-valued stochastic process. Also, let $\mathrm{E}_{t-1}(\cdot)$ refer to the mathematical expectation conditional on the information set available at time $t-1$, including the past of the process $\{\varepsilon_\tau\}_{\tau=t-1,t-2,\ldots}$. Following Engle (1982), the process $\{\varepsilon_t\}$ is then said to follow an ARCH model if there exist a representation such that

$$\varepsilon_t = z_t \sigma_t, \tag{1}$$

where

$$\mathrm{E}_{t-1}(z_t) = \mathrm{E}_{t-1}(z_t^2 - 1) = 0, \tag{2}$$

and σ_t is measurable with respect to the time $t-1$ information set. In most applications, ε_t will correspond to the innovations for the conditional mean from some other process $\{y_t\}$; i.e., $\varepsilon_t \equiv y_t - \mathrm{E}_{t-1}(y_t)$ so that $\mathrm{var}_{t-1}(y_t) = \mathrm{E}_{t-1}(\varepsilon_t^2) = \sigma_t^2$. The setup in Eqs. (1) and (2) is extremely general and does not lend itself directly to empirical implementation without first imposing some simplifying assumptions regarding the temporal dependencies in the conditional variance function. Arguably the two most successful such parameterizations to date have been the Generalized ARCH, or GARCH(p,q), model of Bollerslev (1986), and the Exponential GARCH, or EGARCH(p,q), model proposed by Nelson (1991).

In the GARCH(p,q) model, the conditional variance is parameterized as a distributed lag of past squared innovations,

$$\sigma_t^2 = \omega + \sum_{i=1,q} \alpha_i \varepsilon_{t-i}^2 + \sum_{i=1,p} \beta_i \sigma_{t-i}^2 \equiv \omega + \alpha(L)\varepsilon_t^2 + \beta(L)\sigma_t^2, \tag{3}$$

where L denotes the lag or backshift operator; i.e., $L^i x_t \equiv x_{t-i}$. For the conditional variance in Eq. (3) to be well-defined, all the coefficients in the lag polynomial in the corresponding infinite ARCH representation, $[1-\beta(L)]^{-1}\alpha(L)$, must be positive; see Nelson and Cao (1992). On rearranging the terms in Eq. (3) it follows that

$$[1 - \alpha(L) - \beta(L)]\varepsilon_t^2 = \omega + (1 - \beta(L)]v_t, \tag{4}$$

where $v_t \equiv \varepsilon_t^2 - \sigma_t^2$. Thus, $\mathrm{E}_{t-1}(v_t) = 0$, and following Bollerslev (1988), the GARCH(p,q) formulation in Eq. (3) is readily interpreted as an ARMA$(\max\{p,q\}, p)$ model for the squared process, $\{\varepsilon_t^2\}$. In many applications of the GARCH(p, q) model the estimated lag polynomial $1 - \hat{\alpha}(x) - \hat{\beta}(x) = 0$ has a root which is statistically indistinguishable from unity.[4] Motivated by this empirical regularity, Engle and Bollerslev (1986) proposed the so-called Integrated GARCH, or IGARCH(p, q), process, in which the autoregressive polynomial in Eq. (4) has one unit root. Thus, factorizing this polynomial as $1 - \alpha(L) - \beta(L) \equiv (1 - L)\phi(L)$, where $\phi(z) = 0$ has all the roots outside the unit circle, the IGARCH(p, q) model may be written as

$$\phi(L)(1 - L)\varepsilon_t^2 = \omega + [1 - \beta(L)]v_t. \tag{5}$$

Of course, the notion of a unit root is intrinsically a linear concept and considerable care should be exercised in interpreting notions of persistence in nonlinear models. For instance, as shown by Nelson (1990), even though the IGARCH(1, 1), model is not covariance-stationary, the process is strictly stationary and ergodic. Specifically, Bollerslev and Engle (1993) define a process to be integrated, or persistent, in variance if $\lim_{j \to \infty} \mathrm{E}_t[\mathrm{var}_{t+j-1}(y_{t+j})]$ does not converge to a constant with probability one for all t; that is, the long-term forecasts for the conditional variance remain sensitive to the initial conditions for forecasts of all horizons. This same notion of persistence in nonlinear structures is also the motivation behind the conditional volatility profiles discussed in Gallant, Rossi, and Tauchen (1993).[5] Whereas the conditional volatility profile for a covariance-stationary GARCH model will decay at a geometric rate, the conditional volatility profile for an IGARCH model exhibits infinite dependence on initial conditions.

Although the empirical estimates for the parametric GARCH(p, q) model in Eq. (3) may not reject the null hypothesis of an IGARCH(p, q) model, intuition suggest that in most situations the volatility process is mean-reverting. A possible reconciliation of this conflicting evidence may be available by allowing for fractional orders of integration. To that end Baillie, Bollerslev, and Mikkelsen (1996), recently introduced the so-called Fractionally Integrated GARCH, or FIGARCH(p, d, q), class of models,

$$\phi(L)(1 - L)^d\varepsilon_t^2 = \omega + [1 - \beta(L)]v_t, \tag{6}$$

[4] For instance, Chou (1988) on estimating a GARCH(1, 1) model for weekly data on the value-weighted returns on the NYSE from July 1962 through December 1985 finds $\hat{\alpha}_1 + \hat{\beta}_1 = 0.986$, with the corresponding likelihood ratio test statistic for a unit root equal to 0.980. Lumsdaine (1996) and Lee and Hansen (1994) have shown that this likelihood ratio test and the t-test for $\alpha_1 + \beta_1 = 1$ in the GARCH(1, 1) model have standard asymptotic chi-square and normal distributions, respectively.

[5] A conditional volatility profile is defined as the conditional expectation of σ_{t+j}, $j = 1, 2, \ldots$, evaluated at time t. The difference between a particular conditional volatility profile and a baseline profile provides a natural generalization of linear impulse-response analysis to nonlinear models..

where the roots of $\phi(z) = 0$ lie outside the unit circle. The fractional differencing operator is defined by its Maclaurin series expansion,

$$\begin{aligned}(1-L)^d &= 1 - d \cdot \sum_{k=1,\infty} \Gamma(k-d)\Gamma(1-d)^{-1}\Gamma(k+1)^{-1}L^k \\ &\equiv 1 - \delta_d(L),\end{aligned} \tag{7}$$

in which $\Gamma(\cdot)$ denotes the gamma function. Also, by definition $(1-L)^0 \equiv 1$. Thus, the FIGARCH(p,d,q) model nests the covariance-stationary GARCH(p,q) model for $d = 0$ and the IGARCH(p,q) model in Eq. (5) for $d = 1$. Allowing for values of d in the interval between zero and unity gives an added flexibility that may be important when modeling long-term dependence in the conditional variance.

This new class of conditional variance models provides a direct analogy to the Fractionally Integrated ARMA, or ARFIMA, class of models for the conditional mean; see, e.g., Diebold and Rudebusch (1989), Diebold, Husted, and Rush (1991), Lo (1991), and Sowell (1992) for a discussion of the importance of allowing for noninteger values of integration when modeling long-run dependencies in the conditional mean of economic time series. In the ARFIMA class of models, the short-run behavior of the time series is captured by the conventional ARMA parameters, while the long-run dependence is conveniently modelled through the fractional differencing parameter. A similar result may well hold true when modeling conditional variances. While a shock to the optimal forecast of the future conditional variance decays at an exponential rate for the covariance-stationary GARCH(p,q) model, and remains important for forecasts of all horizons for the IGARCH(p,q) model, in the FIGARCH(p,d,q) model the effect of a shock to the forecast of the future conditional variance will die out at a slow hyperbolic rate. The fractional differencing parameter is therefore identified by the decay rate of a shock to the conditional variance, and not by the ultimate impact on the forecast for the long-run conditional variance.

In order to better understand the statistical properties and the estimation strategy proposed in Section 3 below, it is convenient to rewrite the FIGARCH(p,d,q) model in Eq. (6) in terms of the observationally equivalent infinite ARCH representation,

$$\begin{aligned}\sigma_t^2 &= [1-\beta(1)]^{-1}\omega + \{1 - [1-\beta(L)]^{-1}\phi(L)(1-L)^d\}\varepsilon_t^2 \\ &\equiv [1-\beta(1)]^{-1}\omega + \lambda(L)\varepsilon_t^2.\end{aligned} \tag{8}$$

Since $\delta_d(1) = 1$ for all $d > 0$, it follows immediately from the representation in Eq. (8) that the FIGARCH(p,d,q) model is not covariance-stationary.[6]

[6] However, since the coefficients in the infinite lag polynomial, $\lambda(L)$, are dominated by the coefficients in the infinite ARCH representation of an appropriately defined high-order IGARCH model, it follows from Bougerol and Picard (1992) and Nelson (1990) that the FIGARCH(p,d,q) model is strictly stationary and ergodic.

Of course, for the FIGARCH(p, d, q) model to be well-defined and the conditional variance positive almost surely for all t, all the coefficients in the infinite ARCH representation in Eq. (8) must be nonnegative. General conditions for this to hold are difficult to establish, although the requirements are relatively straightforward to verify on a case by case basis following Nelson and Cao (1992).[7]

While the GARCH(p, q) model conveniently captures the own short-run temporal dependencies for a wide variety of speculative assets, the formulation in Eq. (3) leaves no room for the so-called leverage effect in stock returns. As first noted by Black (1976), stock return volatility tend to be negatively correlated with past returns, possibly due to the increased leverage following a drop in the stock price. In the GARCH models discussed above, the conditional variances are functions only of the absolute magnitudes of the lagged residuals and not their signs. To circumvent this shortcoming, several recent studies have suggested the inclusion of different asymmetric terms in the conditional variance equation; see, e.g., Ding, Granger, and Engle (1993), Engle and Ng (1993), and Glosten, Jagannathan, and Runkle (1993). Alternatively, in the Exponential GARCH, or EGARCH(p, q), model developed by Nelson (1991), the logarithm of the conditional variance is parameterized as an ARMA(p, q) model,

$$\begin{aligned}\ln(\sigma_t^2) &= \omega + \left(1 - \sum_{i=1,p} \varphi_i L^i\right)^{-1} \left(1 + \sum_{i=1,q} \psi_i L^i\right) g(z_{t-1}) \\ &\equiv \omega + [1 - \varphi(L)]^{-1}[1 + \psi(L)] g(z_{t-1}),\end{aligned} \tag{9}$$

where

$$g(z_t) = \theta z_t + \gamma[|z_t| - \mathrm{E}(|z_t|)]. \tag{10}$$

For $\theta < 0$ the future conditional variances will therefore increase proportionally more as a result of a negative shock than for a positive shock of the same absolute magnitude. Note, by definition, the news impact function, $g(\cdot)$, satisfies $\mathrm{E}_{t-1}[g(z_t)] = 0$. A richer parametrization for this function, which downweighs the influence of large absolute innovations has recently been investigated by Bollerslev, Engle, and Nelson (1994).

[7] For instance, for the FIGARCH$(1, d, 1)$ model estimated below, the ARCH parameters in the lag polynomial $\lambda(L)$ may be written as

$$\lambda_1 = \phi_1 - \beta_1 + d, \qquad \lambda_k = \beta_1 \lambda_{k-1} + [(k - 1 - d)k^{-1} - \phi_1]\delta_{d,k-1} \quad \text{for} \quad k \geqslant 2,$$

where $\delta_{d,k} \equiv \delta_{d,k-1}(k - 1 - d)k^{-1}$, $k = 2, 3, \ldots$, refer to the coefficients in the Maclaurin series expansion of $(1 - L)^d$; i.e., $\delta_d(L) = \sum_{k=1,\infty} \delta_{d,k} L^k$. From these recursions, it follows fairly easily that the inequality constraints,

$$\beta_1 - d \leqslant \phi_1 \leqslant (2 - d)/3, \qquad d[\phi_1 - (1 - d)/2] \leqslant \beta_1(\phi_1 - \beta_1 + d),$$

are sufficient to ensure that the corresponding ARCH parameters are all nonnegative.

Just as the estimates for the standard GARCH(p, q) model often indicate an approximate unit root in the autoregressive polynomial in Eq. (4), when estimating the EGARCH(p, q) model in Eqs. (9) and (10), the largest root of the estimated polynomial $\hat{\varphi}(x) = 1$ is often very close to unity.[8] However, as noted by Nelson (1991), the EGARCH(p, q) model could be extended to allow for fractional orders of integration also. Specifically, by factorizing the autoregressive polynomial $[1 - \varphi(L)] = \phi(L)(1 - L)^d$ where all the roots of $\phi(z) = 0$ lie outside the unit circle, the model may be written as

$$\ln(\sigma_t^2) = \omega + \phi(L)^{-1}(1 - L)^{-d}[1 + \psi(L)]g(z_{t-1}). \tag{11}$$

This FIEGARCH(p, d, q) formulation obviously nests the conventional EGARCH model for $d = 0$ and the Integrated EGARCH (IEGARCH) model for $d = 1$. By analogy to the ARFIMA class of models for the conditional mean, $\{\ln(\sigma_t^2)\}$ is covariance-stationary and invertible for d in the interval between -0.5 and 0.5; see, e.g., Hosking (1981). However, shocks to the optimal forecasts for future values of $\ln(\sigma_t^2)$ will dissipate for all values of $d < 1$.[9] It is worth noting that in contrast to the FIGARCH formulation, the parameters for the FIEGARCH model do not have to satisfy any nonnegativity constraints in order for the model to be well-defined.

3. Model specification and estimation

The most common approach for estimating ARCH models relies on the maximization of a conditional likelihood function. In particular, assuming that the one-step-ahead prediction errors are conditionally normally distributed, the likelihood function for the sample $\{y_1, y_2, \ldots, y_T\}$ equals

$$\begin{aligned}\log L(\theta; y_1, y_2, \ldots, y_T | I_0) = &-0.5 \cdot T \cdot \ln(2\pi) \\ &-0.5 \cdot \sum_{t=1,T} [\ln(\sigma_t^2) + \varepsilon_t^2 \sigma_t^{-2}],\end{aligned} \tag{12}$$

where the initial conditions, I_0, are used to start up the recursions for the conditional mean and variance functions.

In many applications with high-frequency financial data, the assumption of conditional normally distributed standardized innovations, $z_t \equiv \varepsilon_t \sigma_t^{-1}$, is violated. However, following Weiss (1986) and Bollerslev and Wooldridge (1992)

[8] For instance, Nelson (1991) on modelling the daily return on the value weighted CRSP index from July 1962 through December 1987 estimates the largest autoregressive parameter to be 0.9996 with an asymptotic t-statistic for a unit root equal to 0.442.

[9] It follows also from Theorem 2.1 in Nelson (1991) that $\{\ln(\sigma_t^2)\}$ is strictly stationary and ergodic for $d < 1/2$.

asymptotically valid inference regarding the normal Quasi Maximum Likelihood Estimates (QMLE) resulting from Eq. (12), say $\hat{\theta}$, may be based on robustified versions of the standard test statistics. In particular, an asymptotic robust covariance matrix for the parameter estimates is consistently estimated by $A(\hat{\theta})^{-1}B(\hat{\theta})A(\hat{\theta})^{-1}$, where $A(\hat{\theta})$ and $B(\hat{\theta})$ denote the Hessian and the outer product of the gradients, respectively, evaluated at $\hat{\theta}$.

Another complication that arises in the estimation of ARCH type models concerns the proper treatment of the initial conditions, I_0. The approach taken here for the fractionally integrated models is based on the infinite ARCH-type representation for the FIGARCH(p, d, q) model in Eq. (8) and the corresponding expansion for the FIEGARCH(p, d, q) model in Eq. (11), with $\varepsilon_t \equiv 0$ for $t = 0, -1, -2, \ldots$, and the pre-sample values of ε_t^2 and σ_t^2 set equal to the unconditional sample variance. Of course, for the FIGARCH model with $d > 0$ and the FIEGARCH model with $d \geqslant 0.5$ the population variance does not exist. This approach directly mirrors the conventional way of estimating stationary GARCH as well as IGARCH models, however.[10] Unlike the finite-lag representations for the standard GARCH(p, q) and EGARCH(p, q) models in Eqs. (3) and (9), the approximate maximum likelihood techniques for the fractionally integrated models also necessitates the truncation of the infinite distributed lags in Eqs. (8) and (11). Since the fractional differencing operator is designed to capture the long-memory features of the process, truncating at too low a lag may destroy important long-run dependencies. In the simulations and the actual estimation results reported on here we fixed the truncation lag at $J = 1{,}000$.[11]

To gauge the accuracy of this approximate maximum likelihood method, Tables 1, 2, and 3 report the results from a detailed simulation study. Some of the findings complement earlier evidence in Baillie, Bollerslev, and Mikkelsen (1996) for different parameter settings. Tables 1, 2, and 3 also contain new important results on model specification and diagnostic checking of ARCH-type models, however. The true Data Generating Process (DGP) for the results in Table 1 is the FIGARCH(1, d, 0) model. The results in Tables 2 and 3 are for a covariance-stationary GARCH(1, 1) and IGARCH(1, 1) DGP, respectively.

[10] As an alternative to fixing the pre-sample values at their unconditional sample analogues backcasting procedures could be employed. The pre-sample values could also be simulated, and the parameter estimates then averaged over a number of such replications until convergence is achieved. Interestingly, however, Diebold and Schuermann (1996) on using nonparametric density estimation techniques in evaluating the exact likelihood function for low-order ARCH models find that for sample sizes of fifty of larger, the exact results are almost identical to the estimates based on the conditional likelihood function used here.

[11] For $0 < d < 1$ and $J = \infty$, the true value of $\delta_d(1) = 1$. For a fixed value of J, the magnitude of the truncation bias is decreasing in d. With $J = 1000$ and $d = 0.5$ as studied in the simulations below, $\delta_{0.5}(1) = 0.982$, whereas for the estimate of d from the FIEGARCH model reported in Section 5, $\delta_{0.633}(1) = 0.995$.

The length of all of the simulated time series is $T = 3{,}000$. The other parameter values are given in the first column of the tables. A total of 500 replications were generated for each simulation.[12]

Turning to the results, it is immediately clear from Table 1, that the approximate maximum likelihood method for the FIGARCH model outlined above works reasonably well for the sample sizes typically encountered with high-frequency financial data. In particular, let $\bar{d}_N$ denote the sample mean of $\hat{d}$ across the $N = 500$ replications; i.e., $\bar{d}_N \equiv N^{-1}\sum_{i=1,N}\hat{d}_i$, where $\hat{d}_i$ refers to the estimate of d from the ith replication. By a standard law of large numbers, the bias, $\mathrm{E}(\hat{d}) - d$, is then consistently estimated by $\bar{d}_N - d$. For the FIGARCH$(1, d, 0)$ model in Table 1, this bias equals $0.513 - 0.500 = 0.013$. Also, by a central limit theorem argument, $N^{1/2}[\bar{d}_N - \mathrm{E}(\hat{d})] \rightarrow \mathrm{N}(0, \sigma^2(\hat{d}))$, where $\sigma^2(\hat{d})$ denotes the variance of $\hat{d}$. Following Cheung and Diebold (1994), the Monte Carlo standard error associated with $\bar{d}_N$ as an estimator for $\mathrm{E}(\hat{d})$ may therefore be consistently estimated by $N^{-1/2}\bar{\sigma}_N$, where $e\bar{\sigma}_N^2 \equiv N^{-1}\sum_{i=1,N}(\hat{d}_i - \bar{d}_N)^2$. For the results reported here $\bar{d}_N = 0.513$ and $\bar{\sigma}_N = 0.075$, so that a symmetric 95% confidence interval for $\mathrm{E}(\hat{d}_i)$, that takes into account the Monte Carlo sampling error, equals [0.506, 0.520]. Similarly, the 95% confidence intervals for $\mathrm{E}(\hat{\mu})$, $\mathrm{E}(\hat{\omega})$, and $\mathrm{E}(\hat{\beta}_1)$ for the FIGARCH$(1, d, 0)$ model in Table 1 are $[-0.002, 0.004]$, [0.111, 0.119], and [0.454, 0.468], respectively. Although the individual 95% confidence intervals only include the true value of ω, it is evident that the estimation procedure tend to produce fairly reliable estimates. Note also, that the Monte Carlo standard errors for the parameter estimates reported in parentheses are generally close to the mean values of the asymptotic robust standard error estimates given in square brackets. Interestingly, the estimates for $d = 1$ for the FIGARCH$(1, d, 0)$ model reported in Table 3 appear equally precise but slightly downward biased, with a 95% confidence band for $\mathrm{E}(\hat{d})$ of [0.991, 0.999]. Similar results are reported in Baillie, Bollerslev, and Mikkelsen (1996) for a FIGARCH model with $d = 0.75$. Table 1 also illustrates that the estimation of a misspecified GARCH(1, 1) model for the FIGARCH$(1, d, 0)$ DGP tend to produce IGARCH type estimates. The mean value of $\hat{\alpha}_1 + \hat{\beta}_1 = \hat{\phi}_1$ for the estimated GARCH(1, 1) models is 0.983. This is in direct accordance with the results for the conditional mean reported in Diebold and Rudebusch (1991).

The results in Tables 2 and 3 indicate that the bias in the estimated parameters for the correctly specified covariance-stationary GARCH(1, 1) model and the IGARCH(1, 1) model are of the same order of magnitude as for the FIGARCH$(1, d, 0)$ estimates reported in Tables 1 and 3. Thus, the truncation

[12] A total of 10,000 observations were generated for each replication, discarding the first 7,000 realizations of each to avoid start-up problems. The normal random variables were generated by the RNDNS subroutine in GAUSS.

Table 1
Finite-sample distribution for the FIGARCH(1, d, 0) DGP

$y_t = \mu + \varepsilon_t$, $\varepsilon_t\sigma_t^{-1}$ i.i.d. N(0, 1)

$\sigma_t^2 = \omega(1-\beta_1)^{-1} + \{1 - [1-\beta_1 L]^{-1}(1-\phi_1 L)(1-L)^d\}\varepsilon_t^2$, $t = 1, 2, \ldots, 3000$

	DGP	OLS	GARCH	IGARCH	FIGARCH
μ	0.00	0.002	0.001	0.001	0.001
		(0.053)	(0.039)	(0.039)	(0.039)
		[0.053]	[0.040]	[0.040]	[0.039]
ω	0.10	9.448	0.141	0.085	0.115
		(7.675)	(0.072)	(0.046)	(0.044)
		[0.384]	[0.039]	[0.024]	[0.036]
β_1	0.45	—	0.867	0.875	0.461
			(0.018)	(0.016)	(0.077)
			[0.017]	[0.017]	[0.068]
ϕ_1	—	—	0.983	1.000	—
			(0.011)	(—)	
			[0.007]	[—]	
d	0.50	—	—	—	0.513
					(0.075)
					[0.065]
AIC	—	0.000	0.018	0.000	0.982
SIC	—	0.000	0.012	0.024	0.964
Q_{10}	—	0.458	0.056	0.050	0.048
Q_{100}	—	0.634	0.066	0.066	0.070
Q_{10}^A	—	1.000	0.732	0.670	0.070
Q_{100}^A	—	1.000	0.342	0.374	0.066
Q_{10}^2	—	1.000	0.792	0.702	0.054
Q_{100}^2	—	1.000	0.354	0.378	0.070

The table reports the means across the 500 Monte Carlo replications for the Quasi Maximum Likelihood Estimates (QMLE) under the FIGARCH(1, d, 0) Data Generating Process (DGP). The Monte Carlo root mean square error of the parameter estimates are given in parentheses, with the mean of the QMLE-based standard error estimates in square brackets. The rows labelled AIC and SIC report the proportion of times that the different models were favored by the Akaike and the Schwarz Information Criterion. Simulated rejection frequencies based on the adjusted nominal 5% level for the Ljung–Box portmanteau tests for up to Kth-order serial correlation in the standardized residuals, $\hat{\varepsilon}_t\hat{\sigma}_t^{-1}$, the absolute standardized residuals, $|\hat{\varepsilon}_t\hat{\sigma}_t^{-1}|$, and the squared standardized residuals, $\hat{\varepsilon}_t^2\hat{\sigma}_t^{-2}$, are denoted by Q_K, Q_K^A, and Q_K^2, respectively.

bias for the FIGARCH model with $J = 1{,}000$ seems rather inconsequential. For instance, the 95% confidence bands for $E(\hat{\mu})$, $E(\hat{\omega})$, $E(\hat{\beta}_1)$, and $E(\hat{\phi}_1)$ for the correctly specified covariance-stationary GARCH(1, 1) model in Table 2 are [– 0.002, 0.002], [0.105, 0.109], [0.847, 0.849], and [0.972, 0.974],

Table 2
Finite-sample distribution for the GARCH(1, 1) DGP

$$y_t = \mu + \varepsilon_t, \quad \varepsilon_t\sigma_t^{-1} \text{ i.i.d. N(0, 1)}$$

$$\sigma_t^2 = \omega(1-\beta_1)^{-1} + \{1 - [1-\beta_1 L]^{-1}(1-\phi_1 L)(1-L)^d\}\varepsilon_t^2, \quad t = 1, 2, \ldots, 3000$$

	DGP	OLS	GARCH	IGARCH	FIGARCH
μ	0.000	0.001	−0.000	−0.000	−0.000
		(0.036)	(0.027)	(0.027)	(0.027)
		[0.035]	[0.029]	[0.029]	[0.029]
ω	0.100	4.004	0.107	0.055	0.138
		(0.888)	(0.024)	(0.013)	(0.076)
		[0.156]	[0.024]	[0.013]	[0.039]
β_1	0.850	—	0.848	0.860	0.552
			(0.017)	(0.017)	(0.265)
			[0.016]	[0.015]	[0.084]
ϕ_1	0.975	—	0.973	1.000	—
			(0.009)	(—)	
			[0.009]	[—]	
d	—	—	—	—	0.675
					(0.273)
					[0.085]
AIC	—	0.000	0.980	0.000	0.020
SIC	—	0.000	0.790	0.202	0.008
Q_{10}	—	0.480	0.068	0.068	0.070
Q_{100}	—	0.552	0.070	0.060	0.060
Q_{10}^A	—	1.000	0.088	0.088	0.108
Q_{100}^A	—	1.000	0.058	0.098	0.086
Q_{10}^2	—	1.000	0.088	0.088	0.108
Q_{100}^2	—	1.000	0.062	0.090	0.084

See footnote to Table 1. The true data-generating process is GARCH(1, 1).

respectively.[13] Note also that, in contrast to the fairly tight confidence bands for the parameter estimates for the correctly specified FIGARCH models in Tables 1 and 3, the estimates for d and β_1 for the misspecified FIGARCH(1, d, 0) model in Table 2 are very imprecisely determined.

Even though standard statistical model selection criteria such as the Akaike (1973) (AIC) or the Schwarz (1978) (SIC) information criteria are commonly

[13] A finite-sample correction for the asymptotic $O_p(T^{-1})$ bias in the GARCH(1, 1) parameter estimates has recently been developed by Linton (1994).

Table 3
Finite-sample distributions for the IGARCH(1, 1) DGP

$y_t = \mu + \varepsilon_t, \quad \varepsilon_t\sigma_t^{-1}$ i.i.d. N(0, 1)

$\sigma_t^2 = \omega(1-\beta_1)^{-1} + \{1 - [1-\beta_1 L]^{-1}(1-\phi_1 L)(1-L)^d\}\varepsilon_t^2, \quad t = 1, 2, \ldots, 3000$

	DGP	OLS	GARCH	IGARCH	FIGARCH
μ	0.00	−0.003	−0.003	−0.003	−0.003
		(0.101)	(0.040)	(0.040)	(0.040)
		[0.096]	[0.042]	[0.042]	[0.042]
ω	0.10	43.7	0.113	0.107	0.108
		(191.3)	(0.031)	(0.026)	(0.030)
		[4.2]	[0.029]	[0.024]	[0.027]
β_1	0.85	—	0.847	0.848	0.843
			(0.013)	(0.013)	(0.037)
			[0.013]	[0.013]	[0.029]
ϕ_1	1.00	—	0.998	1.000	—
			(0.006)	(—)	
			[0.006]	[—]	
d	—	—	—	—	0.995
					(0.051)
					[0.043]
AIC	—	0.000	0.128	0.730	0.142
SIC	—	0.000	0.008	0.984	0.008
Q_{10}	—	0.804	0.042	0.042	0.044
Q_{100}	—	0.960	0.060	0.062	0.058
Q_{10}^4	—	1.000	0.090	0.060	0.080
Q_{100}^4	—	1.000	0.070	0.070	0.064
Q_{10}^2	—	1.000	0.082	0.052	0.070
Q_{100}^2	—	1.000	0.074	0.062	0.068

See footnote to Table 1. The true data generating process is IGARCH(1, 1) or, equivalently, FIGARCH(1, 1, 0).

used in the specification of ARCH models, little is known about their statistical properties in this context.[14] Tables 1, 2, and 3 demonstrate how these criteria may be effectively used in discriminating between the GARCH(1, 1), IGARCH(1, 1), and FIGARCH(1, d, 0) alternatives analyzed here. From Table 1, the correct FIGARCH(1, d, 0) model is chosen 98.2% and 96.4% of the

[14] The AIC criteria to be maximized is defined by $2\log L(\hat{\theta}) - 2k$ where k denotes the number of estimated model parameters. The SIC criteria to be maximized is given by $2\log L(\hat{\theta}) - k\cdot\log T$.

times by the AIC and SIC criteria, respectively. For the covariance-stationary GARCH(1, 1) model in Table 2, the AIC criteria correctly identifies this model 98.0% of the times. Interestingly, for the IGARCH(1, 1) DGP in Table 3 the AIC criteria selects the GARCH(1, 1) and FIGARCH(1, d, 0) formulations, both of which nest the true IGARCH(1, 1) model, 12.8% and 14.2% of the times, respectively. The more parsimonious SIC criteria correctly identifies the IGARCH(1, 1) model in Table 3 in 492 out of the 500 replications, however.

Residual autocorrelations are also commonly employed in the specification and diagnostic checking of ARCH-type models. As for the model selection criteria discussed above, relatively little is known about the sampling distributions of these test statistics in an ARCH context. The last six rows of Tables 1, 2, and 3 therefore report the simulated rejection frequencies for the Ljung and Box (1978) portmanteau tests for up to Kth-order serial correlation in the standardized residuals, $\hat{\varepsilon}_t\hat{\sigma}_t^{-1}$, the absolute standardized residuals, $|\hat{\varepsilon}_t\hat{\sigma}_t^{-1}|$, and the squared standardized residuals, $\hat{\varepsilon}_t^2\hat{\sigma}_t^{-2}$, for the different estimated models. The three tests statistics are denoted by Q_K, Q_K^A, and Q_K^2, respectively. The simulated rejection frequencies for the Q_K test are based on the nominal 5% critical value in the chi-square distribution with K degrees of freedom. McLeod and Li (1983) have shown, that the asymptotic distribution of the squared-residual autocorrelations does not depend on the parameter estimates for the conditional mean under the null hypothesis of homoskedasticity. The usual $T^{-1/2}$ asymptotic standard errors do not apply to the squared standardized residuals from estimated ARCH models, however. In the results reported on here the nominal 5% critical values for the Q_K^A and Q_K^2 test statistics were approximated by the corresponding fractiles in the chi-square distribution with $K - k$ degrees of freedom, where k denotes the number of estimated ARCH parameters; i.e., $k = 2$ for the GARCH(1, 1) and FIGARCH(1, 1) models and $k = 1$ for the IGARCH(1, 1) model.[15]

It is well-known that the presence of heteroskedasticity invalidates the usual asymptotic distributions for the sample autocorrelations for the mean.[16]

[15] The residual autocorrelations from estimated ARMA-type models for the conditional mean may be approximated by a singular linear transformation of the true disturbance autocorrelations; see Durbin (1970). In particular, as shown by Box and Pierce (1970), when testing the residuals from an estimated ARMA model, the portmanteau test is asymptotically chi-square distributed with $K - k$ degrees of freedom, where k denotes the number of estimated ARMA parameters. Following this logic, we simply adjust the degrees of freedom for the Q_K^A and Q_K^2 tests by the number of estimated ARCH parameters. A more complicated exact adjustment procedure that takes account of the ARCH parameter estimation error uncertainty has recently been developed by Li and Mak (1993).

[16] For instance, Stambaugh (1993) shows that with symmetrically distributed GARCH(1, 1) errors, the asymptotic variance for the first-order autocorrelation for the mean, ρ, equals $T^{-1}(1-\rho^2)[1+\phi_1(1-\rho^2)(\kappa+2)(1-\rho^2\phi_1)^{-1}]$, where κ denotes the unconditional excess kurtosis. In particular, with conditionally normal errors $\kappa = 6(\phi_1-\beta_1)^2[1-\phi_1^2-2(\phi_1-\beta_1)^2]^{-1}$. Note that for $\phi_1 = 0$ the expression reduces to the conventional $T^{-1}(1-\rho^2)$ Bartlett (1946) standard error; see also Milhøj (1985).

This result is immediately evident from the second row of all three tables. The empirical rejection frequencies for the Q_{10} and Q_{100} tests for the mean-adjusted series dramatically exceed the 5% nominal significance levels, thus illustrating how the presence of ARCH may give rise to spurious significance of portmanteau type tests for serial correlation in the mean.[17] This effect is especially pronounced under the IGARCH(1, 1) DGP, for which the nominal 5% Q_{100} test statistic for no serial correlation in the mean falsely rejects 96.0% of the times. Interestingly, the actual significance levels for the Q_K tests based on the standardized residuals from the estimated ARCH models are all close to the 5% nominal size. Note, by a standard binomial type argument a 95% confidence interval for the true significance level may be estimated by $[\bar{p}_N - 1.96\,(\bar{p}_N(1-\bar{p}_N)/N)^{1/2}, \bar{p}_N + 1.96(\bar{p}_N(1-\bar{p}_N)/N)^{1/2}]$, where p_i equals one or zero depending upon whether the null hypothesis is rejected or not for the ith simulation, and $\bar{p}_N \equiv N^{-1}\sum_{i=1,N} p_i$ denotes the rejection frequency across all N replications. For $\bar{p}_N = 0.05$ and $N = 500$ this 95% confidence interval covers [0.031, 0.069].

When testing for serial correlation in the absolute or the squared OLS residuals, the null hypothesis of homoskedasticity is uniformly rejected for all three different DGP's. The actual size of the Q_K^A and Q_K^2 tests based on the adjusted degrees of freedom is fairly close to the nominal size for the correctly specified models; i.e., the FIGARCH(1, d, 0) model in Table 1, the GARCH(1, 1) model in Table 2, and all three models in Table 3. At the same time, the results also indicate that these portmanteau tests do have some power in detecting fractionally integrated variance processes. For instance, under the FIGARCH(1, d, 0) DGP in Table 1, the Q_{10}^2 test for the estimated GARCH(1, 1) model rejects the null hypothesis of no remaining heteroskedasticity 79.2% of the times. The distribution of this Q_{10}^2 test statistic is further illustrated in Fig. 4, which graphs the simulated distributions for the estimated FIGARCH(1, d, 0), GARCH(1, 1), and IGARCH(1, 1) models, together with the density for the χ_8^2 distribution.[18] It is obvious that the χ_8^2 distribution provides a very close approximation to the distribution of the test statistic under the true null hypothesis of FIGARCH(1, d, 0). The simple *ad hoc* adjustment procedure for the degrees of freedom in the asymptotic chi-square distribution, obtained by

---

[17] A simple correction to take account of heteroskedasticity within the context of a particular ARCH model has been suggested by Diebold (1988). The Cumby and Huizinga (1992) l-test for models estimated by instrumental variables procedures also has the correct size in the presence of ARCH.

[18] The smooth densities were calculated by an Epanechnikov kernel,

$$\hat{f}(Q) = 0.75\cdot(N\cdot h)^{-1} \sum_{i=1,N} [1 - ((Q-\hat{Q}_i)h^{-1})^2]\cdot I(|(Q-\hat{Q}_i)h^{-1}| \leq 1),$$

where $\hat{Q}_i$ denotes the value of the test statistic from the ith Monte Carlo replication, and the bandwidth, h, was chosen by formula (3.31) in Silverman (1986). We thank Bo Honoré for sharing his GAUSS computer program for carrying out this kernel estimation.

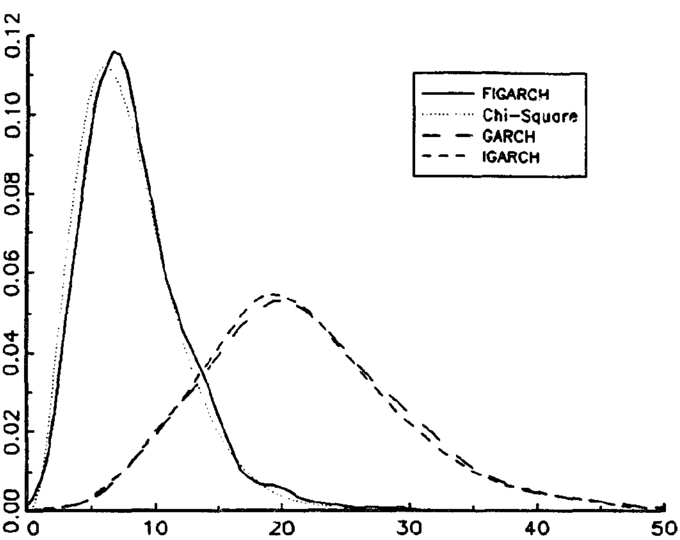

Fig. 4. Finite-sample distributions of the Q_{10}^2 portmanteau tests.

The figure graphs the kernel estimates of the simulated small sample distributions of the Ljung–Box portmanteau tests for up to tenth-order serial correlation in the squared standardized residuals $\hat{\varepsilon}_t^2 \hat{\sigma}_t^{-2}$, from the estimated FIGARCH(1, d, 0), GARCH(1, 1), and IGARCH(1, 1) models. The data generating process is FIGARCH(1, d, 0) with $\mu = 0.0$, $\omega = 0.1$, $d = 0.5$, $\beta = 0.45$, and $T = 3{,}000$. The dotted line gives the density for a chi-square distribution with eight degrees of freedom.

subtracting the number of estimated ARCH parameters from the number of autocorrelations being tested equal to zero, apparently gives rise to fairly accurate nominal significance levels. Note also that the power functions for the GARCH(1, 1) and IGARCH(1, 1) alternatives are almost identical.

In summary, the results discussed in this section illustrate how the approximate maximum likelihood procedure may be used in estimating fractional integrated variance models. Our findings also show how the AIC and SIC model selection criteria and the portmanteau tests for residual autocorrelation may be used effectively in deciding on the correct conditional variance specification. In the next section we shall rely extensively on these tools in the formulation and estimation of conditional variance models for the return on the aggregate U.S. stock market.

4. Long-memory models of stock market volatility

Numerous recent studies have been directed at modeling the temporal variation in stock market volatility, the characteristics of which have very important implications for most modern asset pricing paradigms. Most of these studies using the ARCH methodology have found the volatility process to be highly persistent and possibly not covariance-stationary; see, for instance, Baillie and DeGennero (1990), Bollerslev, Engle, and Nelson (1994), Chou (1988), Engle and Lee (1992), French, Schwert, and Stambaugh (1987), and Nelson (1989, 1991). However, the new class of fractionally integrated variance models discussed above provides an added flexibility that may be important in properly understanding the long-run dependencies in the volatility process.

The data set analyzed here consists of daily prices on the Standard and Poor's 500 composite stock index from January 2, 1953 through December 31, 1990, for a total of $T = 9{,}559$ observations.[19] Following standard practice, we transformed the price index into a continuously compounded capital gains series, $y_t \equiv \log(P_t/P_{t-1})$, $t = 1, 2, \ldots, T$. The time t subscript refers to trading days.

As argued by Scholes and Williams (1977) and Lo and MacKinlay (1990) discontinuous trading in the stocks that make up the index may result in significant serial dependence in the index returns. The exact structure of this autocorrelation will depend on the specific features of the nonsynchroneity. In order to take account of such serial dependence, we here parameterized the mean for all the estimated models as an unrestricted AR(3) model. More complicated parametric formulations for the conditional mean in which the magnitude of the serial correlation depends on the level of the volatility has recently been investigated by LeBaron (1992) and Bollerslev, Engle, and Nelson (1994). The degree of predictability in the mean is very minor and inconsequential for the conditional variance formulations. We shall therefore not pursue any of these more complicated structures any further here.

The variance of returns tend to be higher following weekend and holiday nontrading periods, although the increase is proportionally less than the length of the nontrading period; see, e.g., French and Roll (1986). To capture this effect we included an N_t indicator variable, which gives the number of nontrading days between day t and $t-1$, in all the conditional variance equations.

The estimates from the AR(3) model with the no-trade dummy included in the variance are reported in the first column of Table 4. The evidence for conditional heteroskedasticity is overwhelming. The Q_{10}^2 portmanteau test for up to tenth-order serial correlation in the squared residuals, $\hat{\varepsilon}_t^2$, equals 1697.0. Interestingly, the Q_{100} portmanteau test for no remaining serial correlation in $\hat{\varepsilon}_t$ is also highly significant when judged by the nominal significance level in the corresponding

[19] The Standard 90 index was replaced by the broader Standard and Poor's 500 index in January 1953, corresponding to the start of our data. The properties of these indexes are discussed further in Schwert (1990), who kindly provided the Standard and Poor's 500 data from 1953 through 1962. For a more detailed description of the data set see Bollerslev, Engle, and Nelson (1994), where the same data are analyzed from a different perspective.

Table 4
AR(3)-FIGARCH(1, d, 2) models for daily stock returns

$$y_t \equiv \log(P_t/P_{t-1}) = \mu_0 + \mu_1 y_{t-1} + \mu_2 y_{t-2} + \mu_3 y_{t-3} + \varepsilon_t$$
$$\sigma_t^2 = \omega(1-\beta_1)^{-1} + [1-(1-\beta_1 L)^{-1}(1-\phi_1 L)(1-\phi_2 L)(1-L)^d]\{\varepsilon_t^2 - \delta N_t\} + \delta N_t$$
$$z_t \equiv \varepsilon_t \sigma_t^{-1}, \quad \mathrm{E}_{t-1}(z_t) = 0, \quad \mathrm{E}_{t-1}(z_t^2) = 1$$

	AR	AR–GARCH	AR–IGARCH	AR–FIGARCH
μ_0	$4.04 \cdot 10^{-4}$	$4.62 \cdot 10^{-4}$	$4.57 \cdot 10^{-4}$	$4.80 \cdot 10^{-4}$
	$(0.92 \cdot 10^{-4})$	$(0.70 \cdot 10^{-4})$	$(0.70 \cdot 10^{-4})$	$(0.69 \cdot 10^{-4})$
μ_1	0.120	0.179	0.179	0.182
	(0.032)	(0.012)	(0.011)	(0.012)
μ_2	−0.068	−0.060	−0.060	−0.061
	(0.028)	(0.011)	(0.011)	(0.012)
μ_3	0.002	0.028	0.028	0.026
	(0.024)	(0.012)	(0.011)	(0.012)
ω	$6.15 \cdot 10^{-5}$	$4.49 \cdot 10^{-7}$	$3.40 \cdot 10^{-7}$	$1.27 \cdot 10^{-6}$
	$(0.20 \cdot 10^{-5})$	$(1.26 \cdot 10^{-7})$	$(0.88 \cdot 10^{-7})$	$(0.39 \cdot 10^{-6})$
δ	$2.28 \cdot 10^{-5}$	$5.58 \cdot 10^{-6}$	$5.22 \cdot 10^{-6}$	$5.63 \cdot 10^{-6}$
	$(1.22 \cdot 10^{-5})$	$(1.59 \cdot 10^{-6})$	$(1.45 \cdot 10^{-6})$	$(1.59 \cdot 10^{-6})$
β_1	—	0.933	0.934	0.669
		(0.012)	(0.011)	(0.058)
ϕ_1	—	0.995	1.000	0.365
		(0.003)		(0.052)
ϕ_2	—	0.083	0.089	—
		(0.034)	(0.032)	
d	—	—	—	0.447
				(0.071)
AIC	64351.0	67168.4	67164.8	67212.0
SIC	64308.0	67103.9	67107.5	67147.5
Q_{10}	15.8	14.9	15.7	15.4
Q_{100}	176.4	100.1	100.7	102.7
Q_{10}^2	1697.0	7.1	6.5	10.5
Q_{100}^2	2281.1	73.4	72.3	78.3

The table reports Quasi Maximum Likelihood Estimates (QMLE) for the daily returns on the Standard and Poor's 500 composite index from January 2, 1953 through December 31, 1990, for a total of 9,558 observations. Robust standard errors are reported in parentheses. N_t gives the number of nontrading days between day t and $t-1$. AIC and SIC refer to the Akaike and the Schwarz Information Criterion, respectively. The values of the Ljung–Box portmanteau tests for up to Kth-order serial correlation in the standardized residuals, $\hat{\varepsilon}_t \hat{\sigma}_t^{-1}$, and the squared standardized residuals, $\hat{\varepsilon}_t^2 \hat{\sigma}_t^{-2}$, are denoted by Q_K and Q_K^2, respectively.

χ^2_{97} distribution. However, as the results in Tables 1, 2, and 3 demonstrate, the nominal size of the Q_{10} and Q_{100} tests are grossly misleading in the presence of persistent conditional heteroskedasticity.

Following French, Schwert, and Stambaugh (1987) and Pagan and Schwert (1990), the second row in Table 4 presents the results from estimating a GARCH(2, 1) model for the conditional variance. In all of the ARCH models analyzed below the lag polynomials were factorized to allow for the direct estimation of the inverse real roots. The no-trade dummy was entered in the conditional variance to allow for an impulse effect of the market closures, $\sigma_t^2 = [1 - \beta(1)]^{-1}\omega + \lambda(L)(\varepsilon_t^2 - \delta N_t) + \delta N_t$; see Baillie and Bollerslev (1989). The estimated AR, GARCH, and no-trade dummy parameters are all statistically significant at conventional levels. Interestingly, the Q_{100} test for any remaining serial correlation in the conditional mean beyond the estimated AR(3) model is no longer significant. The largest root in the AR-type representation for the conditional variance is very close to unity, with an estimated half-line in excess of half a year. Indeed, the t-test for $\phi_1 = 1$ does not reject the null hypothesis of an IGARCH model at the usual 5% significance level.

The estimates from the AR(3)–IGARCH(2, 1) model, with $\phi_1 \equiv 1$ are reported in the third column of the table. Not surprisingly, the results are very close to the estimates from the unrestricted model. Whereas the AIC criteria selects the unrestricted AR(3)–GARCH(2, 1) model, the more parsimonious SIC criteria comes out in favor of the AR(3)–IGARCH(2, 1) formulation. Motivated by the classical Beveridge and Nelson (1981) decomposition for the conditional mean of a time series, Engle and Lee (1992) recently proposed a permanent–transitory components model for stock market volatility. The reduced form of this components model for the conditional variance is an unrestricted IGARCH(2, 2) model. It is interesting to note, that on estimating this AR(3)–IGARCH(2, 2) components model for the present data set, the IGARCH(2, 1) model in Table 4 is not rejected against this more general specification.[20]

The more conventional ARCH type models which imply either exponential decay or infinite persistence may be overly restrictive, however. Thus, the estimates for the new AR(3)–FIGARCH(1, d, 1) model are reported in the final row of Table 4. The estimate for the fractional differencing parameter is striking. Judged by standard significance levels, $\hat{d}$ is statistically very different from both zero and one.[21] Also, the AIC and the SIC model selection criteria strongly favor the FIGARCH(1, d, 1) formulation over the GARCH(2, 1) and

[20] The rubust t-test for $\beta_2 = 0$ equals -1.452. In their preferred representation, Engle and Lee (1992) also include asymmetric terms to capture the leverage effect.

[21] Note, that the parameter estimates, $\hat{\beta}_1 = 0.669$, $\hat{\phi}_1 = 0.365$, and $\hat{d} = 0.447$, satisfy the sufficient conditions in Footnote 7 for the conditional variances for the FIGARCH(1, d, 1) model to be positive almost surely for all t.

the IGARCH(2, 1) models. Even though a shock to the volatility process eventually does die out in a forecasting sense, the decay occurs at a slow hyperbolic rate. This is consistent with the conditional volatility profiles in Gallant, Rossi, and Tauchen (1993), which suggest that shocks to the variance are very slowly damped, but do die out. Similarly, Ding, Granger, and Engle (1993) find that the empirical autocorrelations of absolute returns exhibit fairly rapid decay at short lags, while the rate of decay for longer lags is much slower.

As noted above, stock market volatility tend to increase proportionally more for negative than for positive innovations. Several recent studies have proposed various extensions of the GARCH formulations estimated in Table 4 to take account of such asymmetries. We shall here rely on the EGARCH model of Nelson (1991) discussed in Section 2 above. The estimates for this AR(3)–EGARCH(2, 1) model reported in the second column of Table 5 correspond very closely to the findings in Nelson (1991).[22] Although the individual asymptotic standard errors suggest a possible cancellation of $(1 + \psi_1 L)$ and $(1 - \phi_1 L)$, it is worth noting that the portmanteau tests for the squared standardized residuals for this AR(3)–EGARCH(1, 0) model clearly reject the null hypothesis of no remaining serial correlation. Also, the conservative SIC criteria for this restricted model is 67265.3, much lower than the value of 67355.5 obtained for the EGARCH(2, 1) formulation. Note also, that the asymmetric relation between past returns and changes in volatility, as represented by θ, is highly significant. The importance of allowing for asymmetries is further underscored by the large increase in the value of the two model selection criteria from the symmetric GARCH formulations Table 4. The estimated half-life of the largest autoregressive root in the logarithmic conditional variance equation is close to one year.[23] Consequently, the results for the AR(3)–IEGARCH(2, 1) model in the third column correspond very closely to the estimates for the unrestricted AR(3)–EGARCH(2, 1) model. This finding of an approximate unit root in the EGARCH formulation is in direct accordance with the result for the GARCH model.

Similarly, a much better fit for the EGARCH model is obtained by replacing the polynomial corresponding to the largest root, $(1 - \phi_1 L)$, with the fractional differencing operator, $(1 - L)^d$. The values of both model selection criteria are clearly maximized for this fractional integrated AR(3)–FIEGARCH(1, d, 1)

[22] When maximizing the normal quasi-likelihood function in Eq. (12), the $E(|z_t|)$ term in the news impact function, $g(z_t)$, in Eq. (10) was replaced by the sample mean of the absolute standardized residuals; i.e., $T^{-1}\sum_{t=1,T}|\hat{z}_t|$.

[23] It is worth noting, that when estimating the same AR(3)–EGARCH(2, 1) model for the 8,651 daily observations up until June 1, 1987, i.e., excluding the stock market crashes of October 1987 and October 1989, the estimate for the largest root is virtually unaltered compared to the full sample; $\hat{\phi}_1 = 0.996$ versus $\hat{\phi}_1 = 0.997$.

Table 5
AR(3)-FIEGARCH(2, d, 1) models for daily stock returns

$$y_t \equiv \log(P_t/P_{t-1}) = \mu_0 + \mu_1 y_{t-1} + \mu_2 y_{t-2} + \mu_3 y_{t-3} + \varepsilon_t$$
$$\ln(\sigma_t^2) = \omega + \ln(1 + \delta N_t)^{-1} + (1 + \psi_1 L)(1 - \phi_1 L)^{-1}(1 - \phi_2 L)^{-1}(1 - L)^{-d} g(z_{t-1})$$
$$g(z_t) = \theta z_t + \gamma[|z_t| - \mathrm{E}(|z_t|)], \qquad z_t \equiv \varepsilon_t \sigma_t^{-1}, \quad \mathrm{E}_{t-1}(z_t) = 0, \quad \mathrm{E}_{t-1}(z_t^2) = 1$$

	AR	AR-EGARCH	AR-IEGARCH	AR-FIEGARCH
μ_0	$4.04 \cdot 10^{-4}$	$3.42 \cdot 10^{-4}$	$2.98 \cdot 10^{-4}$	$3.48 \cdot 10^{-4}$
	$(0.92 \cdot 10^{-4})$	$(0.73 \cdot 10^{-4})$	$(0.82 \cdot 10^{-4})$	$(0.72 \cdot 10^{-4})$
μ_1	0.120	0.186	0.185	0.184
	(0.032)	(0.011)	(0.012)	(0.011)
μ_2	−0.068	−0.056	−0.053	−0.057
	(0.028)	(0.012)	(0.010)	(0.012)
μ_3	0.002	0.019	0.020	0.021
	(0.024)	(0.011)	(0.009)	(0.011)
ω	−9.696	−10.117	−10.997	−10.273
	(0.032)	(0.312)	(0.413)	(0.414)
δ	0.372	0.212	0.214	0.217
	(0.199)	(0.057)	(0.057)	(0.058)
θ	—	−0.103	−0.098	−0.118
		(0.021)	(0.020)	(0.023)
γ	—	0.203	0.197	0.231
		(0.033)	(0.033)	(0.030)
ψ_2	—	−0.948	−0.974	−0.717
		(0.036)	(0.016)	(0.142)
ϕ_1	—	0.997	1.000	0.774
		(0.002)		(0.124)
ϕ_2	—	0.838	0.894	—
		(0.092)	(0.053)	
d	—	—		0.633
				(0.063)
AIC	64351.0	67434.3	67426.0	67459.9
SIC	64308.0	67355.5	67354.3	67381.1
Q_{10}	15.8	9.4	9.9	10.0
Q_{100}	176.4	99.6	101.1	100.5
Q_{10}^A	3877.0	18.7	18.5	15.1
Q_{100}^A	12621.8	133.2	132.7	122.7

See footnote to Table 4. The Q_K^A row gives the Ljung–Box portmanteau test for up to Kth-order serial correlation in the absolute standardized residuals, $|\hat{\varepsilon}_t \hat{\sigma}_t^{-1}|$.

model. The estimated value of $\hat{d} = 0.633$ is almost six asymptotic standard errors away from unity. As discussed in Section 2, the EGARCH model is readily interpreted as an ARMA model for $\{\ln\{\sigma_t^2\}\}$. It is therefore interesting to note that the p-value for the corresponding Q_{100}^A portmanteau test for up to 100th-order autocorrelation in the absolute standardized residuals, $|\hat{\varepsilon}_t\hat{\sigma}_t^{-1}|$, reported in Table 5 is also the lowest for the FIEGARCH model.

In concluding our discussion of the GARCH and EGARCH estimation results, we note that Engle and Lee (1992) and Gallant, Rossi, and Tauchen (1993) have recently argued that the so-called leverage effect is primarily a short-run phenomenon. This might explain why the results for the EGARCH formulations in Table 5 pertaining to the long-run features and the finding of highly significant fractional integration in the stock market volatility process correspond so closely to the results for the symmetric GARCH models reported in Table 4. The practical importance of modeling these long-run volatility characteristics is illustrated in the next section through the pricing of long-term stock index options.

5. Simulated option prices

A call option gives the owner the right, but not the obligation, to buy a particular security at a pre-specified price within a pre-specified time period. The value of such an option will therefore be intimately related to the distribution of the price of the underlying instrument at the time of maturity. Specifically, the more volatile the underlying price process, the more valuable the option. Organized trading in long-term options, or leaps, with maturity times of one year or longer has increased dramatically in recent years. In this section, we present various simulation based results for the pricing of such hypothetical long-term options on the Standard and Poor's 500 composite stock index.

The standard approach for pricing options rely on risk-neutral valuation methods; see, e.g., Brennan (1979), Lo and Wang (1995), and Rubinstein (1976). In this risk-neutralized probability measure, the price of a call option, that does not allow for early exercise and pays no dividends, will be equal to the discounted expected value of the payoffs at the maturity date. Harrison and Kreps (1979) derive sufficient conditions for the existence of such an equivalent martingale measure. Unfortunately, as shown recently by Amin and Ng (1993), preference-free option valuation is not available under general ARCH-type volatility processes. Thus, rather than actual minimizing the pricing errors from a particular preference-dependent pricing formula, our analysis below is meant primarily to illustrate the practical importance of correctly modeling long-run volatility dependencies when calculating option prices. To that end we will compare the price paths of options with different maturity times for three alternative pricing schemes. Based on the in-sample analysis in the previous section we shall restrict our analysis to the forecasts from the four different EGARCH data-generating mechanisms for the underlying Standard and Poor's index estimated in Table 5.[24]

To set up the notation, let T refer to the time that the option is written; i.e., December 31, 1990 in the experiments reported on below. The maturity time of the option in days is denoted by τ. Thus, the results for $\tau = 70$ and $\tau = 260$ correspond to roughly three-month and one-year maturity times, respectively. An option is said to be at-the-money if the exercise price, K, equals the current value of the underlying security; i.e., here $K = P_T$ where $P_T = 330.2$ refers to the value of the index on December 31, 1990. Similarly, the index option is in-the-money if $K < P_T$ and out-of-the money if $K > P_T$. We shall here concentrate on the results for $K = P_T$ and $K = 1.25 \cdot P_T$. A more detailed investigation of the pricing biases would be interesting but beyond the scope of the present exploratory analysis. The risk-free interest rate over the life of the option is denoted by r. In all of the experiments we took $r = 0.07$ per year. Some informal analysis revealed little sensitivity to this choice.

The celebrated Black and Scholes (1973) option pricing formula for the price of a call option, $C(\sigma, \tau, K, P_T, r)$, is derived under the assumption that the underlying price process, $\{P_{T+t}\}$, $0 \leqslant t \leqslant T$, follows a continuous-time random walk with instantaneous variance σ^2. As the results in the previous section illustrates, this assumption is clearly at odds with the actual time series properties of the Standard and Poor's 500 composite index. The three alternative pricing schemes analyzed here have all been used elsewhere in the literature to value options in the presence of time-varying volatility. For a more thorough discussion of option pricing issues in the presence of stochastic volatility we refer to Amin and Ng (1993), Day and Lewis (1992), Engle and Mustafa (1992), Engle, Hong, Kane, and Noh (1993), Hull and White (1987), Lamoreux and Lastrapes (1993), Melino and Turnbull (1990), Scott (1987), and Wiggins (1987).

Our first pricing equation simply evaluates the Black–Scholes formula with the average time-varying conditional variance of the index over the life of the option in place of σ^2. Denote this expected average per period volatility by $\sigma^{BS}(\tau)^2$. The resulting option prices are then given by

$$C^{BS}(\tau, K) = C(\sigma^{BS}(\tau), \tau, K, P_T, r). \tag{13}$$

[24] Using U.S. stock return data up until 1937 only, Pagan and Schwert (1990) found the EGARCH(2, 1) model to be superior in a simple out-of-sample forecast comparison with other parametric and nonparametric models. In their analysis of actual option pricing errors, Amin and Ng (1993) also find that allowing for asymmetries in the volatility process, as in an EGARCH(1, 1) formulation, tend to produce the most accurate prices.

An estimate of the τ-period volatility was obtained by calculating the sample variance over $N = 1000$ simulations of the future price paths,

$$\hat{\sigma}^{BS}(\tau)^2 = [\tau P_T^2(N-1)]^{-1} \sum_{n=1,N} \left(P_{n,T+\tau} - P_T - N^{-1} \sum_{i=1,N} [P_{i,T+\tau} - P_T] \right)^2$$
$$= [\tau P_T^2(N-1)]^{-1} \sum_{n=1,N} \left(P_{n,T+\tau} - N^{-1} \sum_{i=1,N} P_{i,T+\tau} \right)^2, \qquad (14)$$

where $P_{n,T+t}$ refers to the simulated value of the index at time $T + t$ for the nth replication.

Each of the N replications were generated by sampling randomly from the T in-sample standardized residuals for the particular model under study; i.e., $\hat{z}_t \equiv \hat{\varepsilon}_t \hat{\sigma}_t^{-1}$, $t = 1, 2, \ldots, T$. This bootstrap procedure takes account of the conditional nonnormality in the returns distribution, a feature which may be especially important for far out-of-the-money or in-the-money options; see Baillie and Bollerslev (1992) for a general discussion of prediction error distributions in ARCH models.[25] As would be the case with any other simulation approach based on an ARCH model for the conditional mean and variance, this bootstrap procedure necessarily destroys any higher-order conditional moment dependencies in the standardized residuals. The existence of such higher-order moment dependencies in the returns is at best very weak, however; see, e.g., Hansen (1994).

As noted above, no closed form option pricing formula is generally available in the presence of time-varying volatility. However, Hull and White (1987) show that if the continuous-time volatility process is instantaneously uncorrelated with the aggregate consumption in the economy, the theoretical price of a call option is equal to the expected Black–Scholes price integrated over the average instantaneous variance during the life of the option. Our second pricing formula is motivated by this idea. Specifically, we calculate the prices,

$$C^{HW}(\tau, K) = N^{-1} \sum_{n=1,N} C(\sigma^{HW}(\tau)_n, \tau, K, P_T, r), \qquad (15)$$

where $\sigma^{HW}(\tau)_n$ denotes the volatility per period for the nth simulation. In practice, $\hat{\sigma}^{HW}(\tau)_n$ is estimated by

$$\hat{\sigma}^{HW}(\tau)_n^2 = [P_T^2(\tau-1)]^{-1} \sum_{t=1,\tau} \left(\Delta P_{n,T+t} - \tau^{-1} \sum_{i=1,\tau} \Delta P_{n,T+i} \right)^2$$
$$= [P_T^2(\tau-1)]^{-1} \sum_{t=1,\tau} (\Delta P_{n,T+t} - \tau^{-1}[P_{n,T+\tau} - P_T])^2, \qquad (16)$$

[25] For the AR, EGARCH, IEGARCH, and FIEGARCH models in Table 5, the sample kurtosis for the standardized residuals, $\hat{z}_t$, equal 31.98, 7.62, 7.77, and 7.69, respectively. The sample skewness coefficients are -1.05, -0.43, -0.43, and -0.44, respectively.

where $\Delta P_{n,T+t} \equiv P_{n,T+t} - P_{n,T+t-1}$. This estimator of the per-period variance is based directly on the discretized version of the stochastic differential equation in Hull and White (1987). This volatility estimate takes into account the drift but not the autocovariances of the underlying price process.[26]

Under the equivalent martingale measure discussed above, the price of an option is simply equal to the discounted present value of the payoffs at expiration. The final set of option prices that we consider approximates this theoretical price by replacing the risk-neutralized probability measure with the simulated sampling distribution for $P_{T+\tau}$. The resulting Present Value prices are denoted by

$$C^{PV}(\tau, K) = e^{-r\tau} N^{-1} \sum_{n=1,N} \max\{0, P_{n,T+\tau} - K\}. \tag{17}$$

Options prices based on this approach have previously been studied within an ARCH context by Engle and Mustafa (1992).

Before discussing the actual prices obtained by the three alternative option valuation methods, Table 6 briefly summarizes the simulated predictive distributions for P_{T+70} and P_{T+260} for each of the four different DGP's. The confidence bands from the homoskedastic AR(3) model are the most narrow. Among the three EGARCH models, the AR(3)–EGARCH(2, 1) model always results in the most conservative confidence bands, while the AR(3)–IEGARCH(2, 1) model yields the most dispersed predictions. This is especially pronounced for the one-year-ahead predictions.[27] As the owner of an option has the right, but not the obligation to exercise the option, the simulated options prices should reflect this ranking. It is worth noting, that the conditional variances at the origin of the forecasts are above average by historical standards. Judged by the in-sample distribution of the 9558 conditional variance estimates from the EGARCH, IEGARCH, and FIEGARCH models, the conditional variances on December 31, 1990, correspond to the 0.872, 0.864, and 0.853 fractiles, respectively. It would be interesting, but beyond the scope of the present analysis, to investigate the pricing behavior in less as well as even more volatile periods, also.

The first set of results reported in Table 7 give the simulated prices for at-the-money options with 70 days until maturity. As expected, the IEGARCH model results in the highest prices for all three different pricing schemes, whereas the homoskedastic AR model uniformly produces the lowest valuations. The prices

[26] We also experimented with an alternative estimator which takes into account both the drift and the autocovariances; i.e., $[P_T^2(\tau - 1)]^{-1} (P_{n,T+t} - N^{-1}\sum_{i=1,N} P_{i,T+t})^2$. The simulated options prices obtained with this estimator were generally very close to the results from the $C^{HW}(\tau, K)$ prices using $\hat{\sigma}^{HW}(\tau, K)_n$ in Eq. (16).

[27] The Standard and Poor's 500 index was equal to 417.1 on December 31, 1992; well within the 90% confidence bands for all four models.

Table 6
Simulated predictive price distributions

	Fractile	AR	AR–EGARCH	AR–IEGARCH	AR–FIEGARCH
	0.025	290.6	268.1	260.0	266.4
	0.050	298.0	284.0	279.6	283.0
P_{T+70}	0.500	339.5	341.6	342.3	341.5
	0.950	382.3	387.3	393.1	386.2
	0.975	390.4	398.5	406.3	400.1
	0.025	270.4	225.1	199.7	220.5
	0.050	288.4	255.9	240.8	251.4
P_{T+200}	0.500	361.1	362.7	365.5	362.5
	0.950	458.8	464.3	487.1	471.9
	0.975	477.7	482.7	516.0	492.6

The table reports the simulated fractiles of the predictive distribution for the Standard and Poor's 500 composite index based on the four EGARCH models estimated in Table 5. The origin of the forecasts is December 31, 1990, corresponding to $P_T = 330.2$. All the simulated forecast distributions are based on $N = 1{,}000$ replications. The first group of numbers give the fractiles for predictions 70 days ahead, while the second group of numbers are for predictions 260 days ahead.

Table 7
Simulated options prices

	AR	AR–EGARCH	AR–IEGARCH	AR–FIEGARCH
$C^{BS}(70, P_T)$	13.6	16.2	17.5	16.2
$C^{HW}(70, P_T)$	12.6	14.0	15.1	14.1
$C^{PV}(70, P_T)$	15.6	17.5	19.1	17.5
$C^{BS}(260, P_T)$	33.2	37.5	41.8	38.6
$C^{HW}(260, P_T)$	31.4	33.2	36.8	34.2
$C^{PV}(260, P_T)$	39.8	40.9	46.1	41.6
$C^{BS}(260, 1.25 \cdot P_T)$	4.9	8.5	12.4	9.5
$C^{HW}(260, 1.25 \cdot P_T)$	3.8	5.3	8.2	6.0
$C^{PV}(260, 1.25 \cdot P_T)$	5.8	6.4	10.3	7.3

The table reports the simulated call option prices for the Standard and Poor's 500 composite index based on the four EGARCH models estimated in Table 5. The simulations are based on $N = 1{,}000$ replications. The prices labelled $C^{BS}(\tau, K)$ refer to the Black–Scholes prices in Eq. (15), $C^{HW}(\tau, K)$ gives the Hull–White prices computed from Eq. (17), while $C^{PV}(\tau, K)$ gives the discounted Present Value prices in Eq. (19). The options are all written at time T corresponding to December 31, 1990 and expire τ days later. The value of the index at time T is denoted by P_T. The second argument in the option pricing formulas, K, refers to the exercise price. The other parameter settings are as discussed in the text.

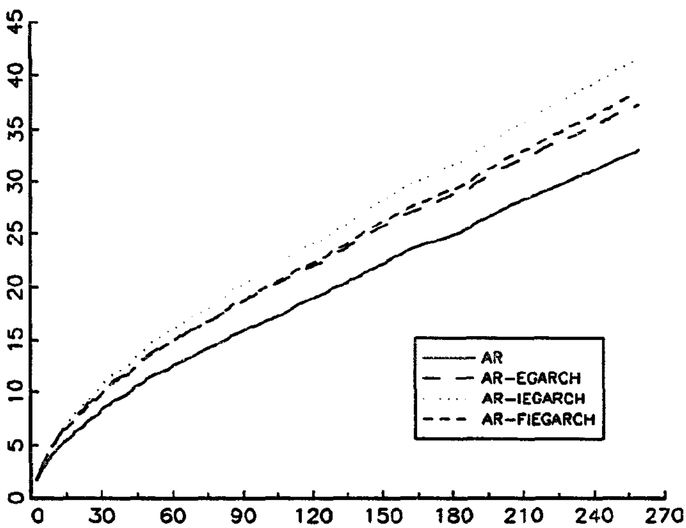

Fig. 5. At-the-money options prices.

The figure graphs the Black–Scholes prices, $C^{BS}(\tau, P_T)$, for at-the-money call options on the Standard and Poor's 500 composite index as a function of the number of days to maturity. All the options are written on December 31, 1990. The four different price paths are calculated from the AR(3), AR(3)–EGARCH(2, 1), AR(3)–IEGARCH(2, 1), and AR(3)–FIEGARCH(1, d, 1) estimates reported in Table 5.

for the EGARCH and the FIEGARCH models are very close. As argued above, the practical importance of the fractional integrated variance formulations stems primarily from the added flexibility when modeling long-run volatility features. Indeed, for the one-year at-the-money options, reported in the second group of numbers in Table 7, the FIGARCH prices are very clearly between the EGARCH and the IEGARCH valuations. This clearcut ranking as the maturity time increases is also evident from Fig. 5, which graphs the Black–Scholes prices $C^{BS}(\tau, P_T)$, $\tau = 1, 2, \ldots, 260$, for all four DGP's.

Of course, the tail behavior of the predicted price distibutions will be more important the further apart the current price of the underlying asset, P_T, and the exercise price, K. To illustrate this effect, the last set of numbers in Table 7 reports the results for one-year out-of-the-money options with an exercise price 25% above the current value of the index; i.e., $K = 1.25 \cdot P_T$. Trading in long-term options that are much further out-of-the-money is quite common. The results for these options are even more illuminating. The prices based on the IEGARCH model are roughly double the homoskedastic AR prices. The FIEGARCH model again yield valuations that are between the EGARCH and the IEGARCH models, but the relative differences are much greater than for the

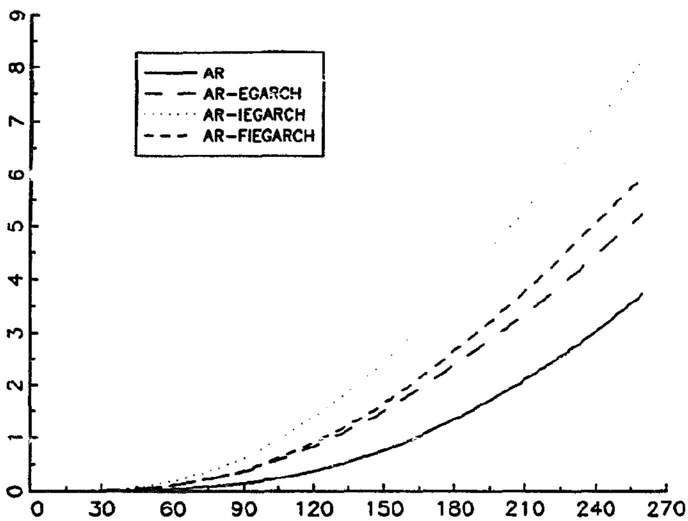

Fig. 6. Out-of-the-money options prices.

The figure graphs the Hull–White prices, $C^{HW}(\tau, 1.25 \cdot P_T)$, for out-of-the-money call options on the Standard and Poor's 500 composite index as a function of the number of days to maturity. All the options are written on December 31, 1990. The four different price paths are calculated from the AR(3), AR(3)–EGARCH(2, 1), AR(3)–IEGARCH(2, 1), and AR(3)–FIEGARCH(1, d, 1) estimates reported in Table 5.

at-the-money options. This is also illustrated in Fig. 6, which graphs the Hull–White prices $C^{HW}(\tau, 1.25 \cdot P_T)$, $\tau = 2, 3, \ldots, 260$.[28] It would be interesting, but beyond the scope of the present paper, to further explore the economic implicatioins of these fairly large price differentials. One possible route would be to formulate trading strategies based on the various pricing formulas as in Engle, Hong, Kane, and Noh (1993). We are currently pursuing these ideas using actual leaps, or long-term options, prices for a number of individual stocks as well as the Standard and Poor's 500 stock index. Along these lines, it is interesting to note that Amin and Ng (1993) on simulating an EGARCH(1, 1) model for the returns on individual stocks, find that the corresponding longer maturity options tend to be underpriced relative to short term options. The results presented here suggest that correctly modeling the long-run dependencies in the volatility process of the underlying asset may be as important as the choice of approximate option valuation method when pricing long maturity contracts.

[28] Note, the per period volatility in Eq. (16) is not defined for $\tau = 1$.

6. Conclusion

A new class of more flexible fractionally integrated EGARCH models for characterizing the long-run dependencies in U.S. stock market volatility was proposed. Strong evidence was uncovered that the conditional variance for the Standard and Poor's 500 composite index is best modeled as a mean-reverting fractionally integrated process. The practical importance of this finding was illustrated via the simulation of hypothetical prices for long-term options on the index. New simulation results on the finite-sample distributions of maximum likelihood estimation procedures, model selectioin criteria, and portmanteau diagnostic checking for the fractional integrated and general ARCH-type models for the conditional variance were also presented.

It would be interesting to extend the empirical analysis for the Standard and Poor's 500 composite index presented here to individual stock and other more broadly defined asset categories. Along these lines, we note that Baillie, Bollerslev, and Mikkelsen (1996) also report significant evidence for the presence of fractional integrated behavior in the conditional variance of nominal U.S. dollar–Deutschmark exchange rates. These findings of long-memory components in the volatility processes of asset returns have important implications for many paradigms in modern financial economics. In addition to the pricing of long-term options as discussed above, optimal portfolio allocations may become extremely sensitive to the investment horizon if the volatility of the returns are long-range dependent. Similarly, optimal hedging decisions must take into account any such long-run dependencies. A more formal and detailed empirical investigation of these issues would be an important task for future research.

References

Akaike, H., 1973, Information theory and an extension of the maximum likelihood principle, in: B.N. Petrov and F. Csáki, eds., Second International symposium on information theory (Akadémiai Kiadó, Budapest).

Amin, K.I. and V. Ng, 1993, Option valuation with systematic stochastic volatility, Journal of Finance 48, 881–910.

Baillie, R.T. and T. Bollerslev, 1989, The message in daily exchange rates: A conditional variance tale, Journal of Business and Economic Statistics 7, 297–305.

Baillie, R.T. and T. Bollerslev, 1992, Prediction in dynamic models with time-dependent conditional variances, Journal of Ecdonometrics 52, 91–113.

Baillie, R.T. and T. Bollerslev, 1994, Cointegration, fractional cointegration, and exchange rate dynamics, Journal of Finance 49, 737–745.

Baillie, R.T. and R.P. DeGennaro, 1990, Stock return and volatility, Journal of Financial and Quantitative Analysis 25, 203–214.

Baillie, R.T., T. Bollerslev, and H.O. Mikkelsen, 1996, Fractionally integrated generalized autoregressive conditional heteroskedasticity, Journal of Econometrics, forthcoming.

Baillie, R.T., C.-F. Chung, and M.A. Tieslau, 1996, Analyzing inflation by the fractionally integrated ARFIMA–GARCH model, Journal of Applied Econometrics, forthcoming.
Bartlett, M.S., 1946, On the theoretical specification of sampling properties of autocorrelated time series, Journal of the Royal Statistical Society B 8, 27–41.
Beveridge, S. and C.R. Nelson, 1981, A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the business cycle, Journal of Monetary Economics 7, 151–174.
Black, F., 1976, Studies of stock price volatility changes, Proceedings from the American Statistical Association, Business and Economic Statistics Section, 177–181.
Black, F. and M. Scholes, 1973, The pricing of options and corporate liabilities, Journal of Political Economy 81, 637–659.
Bollerslev, T., 1986, Generalized autoregressive conditional heteroskedasticity, Journal of Econometrics 31, 307–327.
Bollerslev, T., 1988, On the correlation structure for the generalized autoregressive conditional heteroskedastic process, Journal of Time Series Analysis 9, 121–131.
Bollerslev, T. and R.F. Engle, 1993, Common persistence in conditional variances, Econometrica 61, 167–186.
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., R.Y. Chou, and K.F. Kroner, 1992, ARCH modeling in finance: A review of the theory and empirical evidence, Journal of Econometrics 52, 5–59.
Bollerslev, T., R.F. Engle, and D.B. Nelson, 1994, ARCH models, in: R.F. Engle and D. McFadden, eds., Handbook of econometrics, Vol. 4 (North-Holland, Amsterdam).
Bougerol, P. and N. Picard, 1992, Stationarity of GARCH processes and of some nonnegative time series, Journal of Econometrics 52, 115–128.
Box, G.E.P. and D.A. Pierce, 1970, Distribution of residual autocorrelations in autoregressive integrated moving average time series models, Journal of the American Statistical Association 65, 1509–1526.
Brennan, M., 1979, The pricing of contingent claims in discrete time models, Journal of Finance 34, 53–68.
Cheung, Y-W. and F.X. Diebold, 1994, On maximum likelihood estimation of the differencing parameter of fractionally integrated noise with unknown mean, Journal of Econometrics 62, 301–316.
Cheng, Y.W. and K.S. Lai, 1993, A fractional cointegration analysis of purchasing power parity, Journal of Business and Economic Statistics 11, 103–112.
Chou, R.Y., 1988, Volatility persistence and stock valuations: Some empirical evidence using GARCH, Journal of Applied Econometrics 3, 279–294.
Cumby, R.E. and J. Huizinga, 1992, Testing the autocorrelation structure of disturbances in ordinary least squares and instrumental variables regressions, Econometrica 60, 185–195.
Dacorogna, M.M., U.A. Müller, 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.
Day, T.E. and C.M. Lewis, 1992, Stock market volatility and the information content of stock index options, Journal of Econometrics 52, 267–287.
Diebold, F.X., 1988, Empirical modeling of exchange rate dynamics (Springer-Verlag, New York, NY).
Diebold, F.X. and G.D. Rudebusch, 1989, Long memory and persistence in aggregate output, Journal of Monetary Economics 24, 189–209.
Diebold, F.X. and G.D. Rudebusch, 1991, On the power of Dickey–Fuller tests against fractional alternatives, Economics Letters 35, 155–160.

Diebold, F.X. and T. Schuermann, 1996, Exact maximum likelihood estimation of observation driven econometric models, in: R.S. Mariano, M. Wrecks, and T. Schuermann, eds., Simulation based inference in econometrics: Methods and applications (Cambridge University, Cambridge).
Diebold, F.X., S. Husted, and M. Rush, 1991, Real exchange rates under the gold standard, Journal of Political Economy 99, 1252–1271.
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.
Durbin, J., 1970, Testing for serial correlation in least squares regression when some of the regressors are lagged dependent variables, Econometrica 38, 410–421.
Engle, R.F., 1982, Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation, Econometrica 50, 987–1008.
Engle, R.F. and T. Bollerslev, 1986, Modeling the persistence of conditional variances, Econometric Reviews 5, 1–50.
Engle, R.F. and G.G.J. Lee, 1992, A permanent and transitory component model of stock return volatility, Manuscript (Department of Economics, University of California, San Diego, CA).
Engle, R.F. and C. Mustafa, 1992, Implied ARCH models from options prices, Journal of Econometrics 52, 289–311.
Engle, R.F. and V. Ng, 1993, Measuring and testing the impact of news on volatility, Journal of Finance 48, 1749–1778.
Engle, R.F., C-H. Hong, A. Kane, and J. Noh, 1993, Arbitrage valuation of variance forecasts with simulated options, in: D.M. Chance and R.R. Trippi, eds., Advances in futures and options research (JAI Press, Greenwich, CT).
French, K.R. and R. Roll, 1986, Stock return variances: The arrival of information and the reaction of traders, Journal of Financial Economics 17, 5–26.
French K.R., G.W. Schwert, and R.F. Stambaugh, 1987, Expected stock returns and volatility, Journal of Financial Economics 19, 3–29.
Gallant, A.R., P.E. Rossi, and G. Tauchen, 1993, Nonlinear dynamic structures, Econometrica 61, 871–907.
Glosten, L.R., R. Jagannathan, and D.E. Runkle, 1993, On the relation between the expected value and the volatility of the nominal excess return on stocks, Journal of Finance 48, 1779–1801.
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 R. Joyeux, 1980, An introduction to long memory time series models and fractional differencing, Journal of Time Series Analysis 1, 15–39.
Hansen, B.E., 1994, Autoregressive conditional density estimation, International Economic Review 35, 705–730.
Harrison, G.M. and D. Kreps, 1979, Martingales and arbitrage in multiperiod securities markets, Journal of Economic Theory 20, 381–408.
Hosking, J.R.M., 1981, Fractional differencing, Biometrika 68, 165–176.
Hull, J. and A. White, 1987, The pricing of options on assets with stochastic volatilities, Journal of Finance 42, 381–400.
Lamoureux, C.G. and W.D. Lastrapes, 1993, Forecasting stock return variance: Toward an understanding of stochastic implied volatilities, Review of Financial Studies 6, 293–326.
LeBaron, B., 1992, Some relations between volatility and serial correlation in stock market returns, Journal of Business 65, 199–220.
Lee, S.-W. and B.E. Hansen, 1994, Asymptotic theory for the GARCH(1, 1) quasi-maximum likelihood estimator, Econometric Theory 10, 29–52.
Li, W.K. and T.K. Mak, 1993, On the squared residual autocorrelations in conditional heteroskedastic time series modeling, Manuscript (Department of Statistics, University of Hong Kong, Hong Kong).
Linton, O., 1994, The asymptotic bias of the QMLE in the GARCH(1, 1) model, Manuscript (Department of Economics, Yale University, New Haven, CT).

Ljung, G.M. and G.E.P. Box, 1978, On a measure of lack of fit in time series models, Biometrika 65, 297–303.

Lo, A.W., 1991, Long term memory in stock market prices, Econometrica 59, 1279–1313.

Lo, A.W. and A.C. MacKinlay, 1990, An econometric analysis of nonsynchronous trading, Journal of Econometrics 45, 181–211.

Lo, A.W. and J. Wang, 1995, Implementing option pricing models when asset returns are predictable, Journal of Finance 50, 87–129.

Lumsdaine, R.L., 1996, Consistency and asymptotic normality of the quasi-maximum likelihood estimator in IGARCH (1, 1) and covariance stationary GARCH (1, 1) models, Econometrica, forthcoming.

Mandelbrot, B.B. and J.W. Van Ness, 1968, Fractional Brownian motions, fractional noises and applications, SIAM Review 10, 422–437.

McLeod, A.I. and W.K. Li, 1983, Diagnostic checking ARMA time series models using squared-residual autocorrelations, Journal of Time Series Analysis 4, 269–273.

Melino, A. and S.M. Turnbull, 1990, Pricing foreign currency options with stochastic volatility, Journal of Econometrics 45, 239–265.

Milhøj, A., 1985, The moment structure of ARCH processes, Scandinavian Journal of Statistics 12, 281–292.

Nelson, D.B., 1989, Modeling stock market volatility changes, Proceedings of the American Statistical Association, Business and Economic Statistics Section, 93–98.

Nelson, D.B., 1990, Stationarity and persistence in the GARCH(1, 1) model, Econometric Theory 6, 318–334.

Nelson, D.B. 1991, Conditional heteroskedasticity in asset returns: A new approach, Econometrica 59, 347–370.

Nelson, D.B. and C.Q. Cao, 1992, Inequality constraints in the univariate GARCH model, Journal of Business and Economic Statistics 10, 229–235.

Pagan, A.R. and G.W. Schwert, 1990, Alternative models for conditional stock volatility, Journal of Econometrics 45, 267–290.

Rubinstein, M., 1976, The valuation of uncertain income streams and the pricing of options, Bell Journal of Economics and Management Science 2, 407–425.

Scholes, M. and J. Williams, 1977, Estimating betas from non-synchronous data, Journal of Financial Economics 5, 309–327.

Schwarz, G., 1978, Estimating the dimension of a model, Annals of Statistics 6, 461–464.

Schwert, G.W., 1990, Indexes of U.S. stock prices from 1802 to 1987, Journal of Business 63, 399–426.

Scott, L., 1987, Option pricing when the variance changes randomly: Theory, estimation, and an application, Journal of Financial and Quantitative Analysis 22, 419–436.

Silverman, B.W., 1986, Density estimation for statistics and data analysis (Chapman and Hall, New York, NY).

Sowell, F.B., 1992, Modeling long-run behavior with the fractional ARIMA model, Journal of Monetary Economics 29, 277–302.

Stambaugh, R.F., 1993, Estimating conditional expectations when volatility fluctuates, Manuscript (Department of Finance, Wharton School, University of Pennsylvania, Philadelphia, PA).

Weiss, A.A., 1986, Asymptotic theory for ARCH models: Estimation and testing, Econometric Theory 2, 107–131.

Wiggins, M., 1987, Option values under stochastic volatility: Theory and empirical estimates, Journal of Financial Economics 19, 351–372.