Long-term equity anticipation securities and stock market volatility dynamics

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

[a] *Department of Economics, Duke University Durham, NC 27708-0097, USA and NBER*
[b] *Marshall School of Business, University of Southern California, Los Angeles, CA 90089-1421, USA*

Abstract

Recent empirical findings suggest that the long-run dependence in U.S. stock market volatility is best described by a slowly mean-reverting fractionally integrated process. The present study complements this existing time-series-based evidence by comparing the risk-neutralized option pricing distributions from various ARCH-type formulations. Utilizing a panel data set consisting of newly created exchange traded long-term equity anticipation securities, or leaps, on the Standard and Poor's 500 stock market index with maturity times ranging up to three years, we find that the degree of mean reversion in the volatility process implicit in these prices is best described by a Fractionally Integrated EGARCH (FIEGARCH) model. © 1999 Elsevier Science S.A. All rights reserved.

JEL classification: C13; C15; C22; C52

Keywords: Option pricing; Leaps; Stock market volatility; Long memory; Fractionally integrated EGARCH

1. Introduction

Exchange traded long-term equity anticipation securities (Leaps) were first introduced by the Chicago Board Options Exchange (CBOE) in the Fall of 1990. Since then organized trading in leaps on the CBOE have grown substantially both in terms of contract scope and trading volume, and other options

* Corresponding author. Tel.: 919 660 1846; fax: 919 684 8974; E-mail: boller@econ.duke.edu.

exchanges now list similar long-lived contingent claims. Like a conventional equity call or put option, a leap gives the owner the right to buy or sell an individual stock or an underlying basket of stocks at or within a given time at a pre-specified price. As such, the price of a leap depends intimately on the forecasted return volatility of the underlying asset over the life of the contract. Whereas the maturity times for exchange traded options do not exceed nine months, the expiration cycles for leaps range up to 3 yr. Leap prices therefore contain potentially new valuable information about the markets' expected degree-of-mean reversion in the underlying volatility processes over such multi-year horizons,[1] the extent of which has received much recent attention in the empirical finance and time series literatures.

In particular, on fitting conventional ARCH and stochastic volatility type models to long time series of daily or weekly returns, several studies have documented evidence in favor of what appears to be a unit root in the autoregressive part of the conditional variance equation; see e.g. Bollerslev et al. (1992), Ghysels et al. (1996), and Shephard (1995) for surveys of this literature. More recently, however, Baillie et al. (1996), Bollerslev and Mikkelsen (1996), Breidt et al. (1998), Comte and Renault (1998), and Harvey (1998) have argued that these findings of a unit root in the variance may simply be an artefact of the restrictive time series specifications employed in previous studies. By adopting more flexible ARCH and stochastic volatility parameterizations, that allow for a slowly mean-reverting process, it appears that the temporal dependencies in aggregate U.S. stock market and exchange rate volatilities are better characterized by fractionally integrated time series models.

The present paper complements this existing time-series-based evidence on the long-run dependence in the aggregate market volatility by inferring the degree of mean-reversion implicit in a panel data set of leaps transactions prices on the Standard and Poor's 500 (S&P500) composite stock price index. As shown by Amin and Ng (1993a), no closed-form, preference-free option valuation formula is generally available when the volatility is time varying. The results reported on below relies on the specific equilibrium-based valuation schemes put forward by Amin and Ng (1993b) and Duan (1995a,b), which allow for the determination of contingent claims prices under fairly weak assumptions about the joint distribution of the underlying price process and the marginal rate of intertemporal substitution of a representative investor. By comparing the observed leaps prices and their simulated theoretically counterparts, we find that among the different hypothesized ARCH models for the underlying index returns, the Fractionally Integrated version of Nelson's (1991) EGARCH model (FIEGARCH) developed by Bollerslev and Mikkelsen (1996) results in the

---

[1] In a closely related context, Bakshi et al. (1998) argue that long-term options, or leaps, are also better able to differentiate among alternative pricing models.

lowest average absolute and relative pricing errors. While the discussion in the paper is focused on the model estimates with daily returns, the corresponding results for temporally aggregated weekly returns also support the FIEGARCH formulation.

Instead of informally comparing the actual leaps prices to their theoretical counterparts as a method for differentiating between the competing volatility models, the parameters of the models could be estimated by indirect inference procedures that simultaneously make use of the underlying asset returns and the leaps prices; see e.g. Renault (1997) for a more detailed discussion of these ideas. Such an approach would allow for a direct test of the significance of the long-memory volatility component implicit in the observed leaps prices. Furthermore, as argued by Pastorello et al. (1996) the use of the leaps prices for inference purposes may also result in non-trivial efficiency gains. However, with several thousand multi-year contracts, the practical implementation of such indirect inference procedures would be extremely expensive from a computational point of view. We therefore did not pursue this estimation strategy in the present context.

The plan for the rest of the paper is as follows. The theoretical option pricing framework employed in valuing the leaps is reviewed in Section 2. Section 3 discusses the EGARCH, IEGARCH and FIEGARCH time-series formulations used in modeling the S & P500 volatility. The characteristics of the leaps data set is detailed in Section 4. This section also provides an illustration of the valuation methodology for a few specific leaps. Section 5 presents our general findings related to the average pricing errors across the different volatility models and sampling frequencies, along with a simple regression-based characterization of the most important biases in the preferred FIEGARCH leaps valuations. Section 6 concludes.

2. Option pricing framework

The seminal Black and Scholes (1973) option pricing formula is derived under the assumption that the price of the underlying asset follows a continuous time random walk with a constant instantaneous variance. Numerous subsequent studies have sought to relax this implausible empirical assumption. In particular, option pricing formulas that permit the latent instantaneous volatility to be determined by a separate continuous time stochastic process have been developed by Hull and White (1987) among others. Unfortunately, the volatility processes employed in most of these studies appear to be dictated more by their theoretical tractability than by their empirical plausibility. Excellent surveys of this extensive literature along with a discussion of the complications that arise in the estimation of continuous time models are provided by Bates (1996), Melino (1994), and Renault (1997).

It is not the purpose of the present paper to develop another option pricing model. Instead we seek to uncover the importance of long-memory features in characterizing volatility dynamics more generally by exploiting the pertinent information in long-term options contracts. As noted by Renault (1997), the apparent persistence of the 'smile', or the 'sneer', in the plot of implied Black–Scholes volatilities as a function of the underlying exercise price is difficult to explain within the context of standard option pricing models that only allow for short memory in the return volatility process. This is also one of the primary motivations behind the continuous time long-memory stochastic volatility model proposed by Comte and Renault (1998), in which the volatility process evolves according to a fractional Brownian motion. Under the simplifying assumption that the volatility risk is not priced, it is possible to show that the Hull and White (1987) option pricing formula involving the quadratic variation of the future latent volatility over the life of the option remains intact. While this allows for an elegant theoretical treatment of option pricing, the implementation of the continuous time long-memory volatility model with discretely sampled data raises a number of practical questions. Generalizations to allow for more realistic short-run dynamics are also difficult. More importantly, the vast majority of empirical asset pricing studies tend to rely on the ARCH stochastic difference equation representations when modeling temporal dependencies in financial market volatility. We therefore choose to frame our assessment of the importance of long-memory volatility dependencies in terms of this popular class of models.

The corresponding theoretical option prices are based on the partial equilibrium formula in Amin and Ng (1993b) and Duan (1995a,b). The derivation of these formula follow the standard risk-neutral valuation approach for determining the price of a European option as the discounted expected payout at the maturity date, where the expectation is evaluated in the risk-adjusted martingale measure; see e.g. Harrison and Kreps (1979). As such, the empirical results necessarily depends upon the particular risk-neutralization scheme. However, the results in Bollerslev and Mikkelsen (1996) highlight, that for the time-series models investigated here, the prices of long-term options contracts are much more sensitive to the assumed underlying price dynamics than they are to assumptions about the specific option pricing model.

To set out the notation, let P_t, $t = 1, 2, \ldots$, denote the discretely observed price process for the underlying index value. Also, suppose that there exists a state pricing density process, M_t, such that the time $t - 1$ price of a single period, risk free, pure discount bond equals $E_{t-1}(M_t/M_{t-1}) = \exp(-r_{t-1})$, where r_t refers to the risk free one-period interest rate, and the conditional expectation is evaluated with respect to the agents full information set. In a representative agent economy with a time-separable utility function, M_t/M_{t-1} corresponds to the intertemporal marginal rate of substitution. By the law of iterated expectations, the time t price of a derivative risky claim with a single,

random, time $t + T$ payout, say $\pi_{t+T}(P_{t+T})$, may then be expressed as $\pi_t = E_t[\pi_{t+T}(P_{t+T}) \cdot M_{t+T}/M_t]$. Specifically, for a European call with an exercise price of X_t, $\pi_{t+T}(P_{t+T}) = \max\{0, P_{t+T} - X_t\}$, whereas for a put $\pi_{t+T}(P_{t+T}) = \max\{0, X_t - P_{t+T}\}$.

Of course, the state pricing density is generally unobservable. However, suppose that the corresponding latent logarithmic pricing kernel, $m_t \equiv \ln(M_t/M_{t-1})$, and the index return, $R_t \equiv \ln(P_t/P_{t-1})$, are conditionally jointly normally distributed. Consider the following representation of the return:

$$R_t \equiv \ln(P_t/P_{t-1}) = \mu_t - 1/2\sigma_t^2 + \varepsilon_t, \tag{1}$$

where μ_t is measurable with respect to the time $t - 1$ information set, $E_{t-1}(\varepsilon_t) = 0$, and the conditional variance, $\sigma_t^2 \equiv \mathrm{Var}_{t-1}(\varepsilon_t) = \mathrm{Var}_{t-1}(R_t)$, follows an ARCH process which depends on the past returns and innovations only.[2] Then by the assumption of conditional normality $E_{t-1}(P_t/P_{t-1}) = E_{t-1}[\exp(R_t)] = \exp(\mu_t)$. Also, by the fundamental pricing relationship $E_{t-1}(M_t/M_{t-1}) = E_{t-1}[\exp(m_t)] = \exp(-r_{t-1})$, and $E_{t-1}[(P_t/P_{t-1}) \cdot (M_t/M_{t-1})] = 1$. Utilizing the joint normality assumption, this implies that $E_{t-1}[\exp(R_t + m_t)] = \exp[\mu_t - r_{t-1} + \mathrm{Cov}_{t-1}(R_t, m_t)] = 1$, or $\mathrm{Cov}_{t-1}(R_t, m_t) = r_{t-1} - \mu_t$. It is this covariance between R_t and m_t that renders preference free option valuation infeasible.[3] However, assuming that μ_t is a function of the past observable returns only, Amin and Ng (1993b) show that the covariance term may be eliminated from the risk-neutral distribution implied by the equivalent martingale measure.

Let P_t^* denote the price process under this transformed distribution. The theoretical price of any contingent claim is then conveniently expressed as the

---

[2] The representation of the conditional mean return as $\mu_t - 1/2\sigma_t^2$ entails no loss of generality. However, the assumption of no other state variables in the information set is restrictive.

[3] The covariance between R_t and m_t corresponds to the negative of the risk premium, or the expected return in excess of the risk free rate. In general investors will accept (require) a low (high) expected return for an asset that has a low (high) correlation with the pricing kernel, m_t. In the representative agent economy discussed above, a high (low) marginal rate of substitution, or high (low) m_t, corresponding to a low (high) level of future consumption, coupled with a high (low) return would help to smooth consumption, calling for a small, or even negative, risk premium. Conversely, a negative covariance between m_t and R_t implies that the asset 'adds' variance to consumption thus requiring a positive risk premium. Since the value of the pricing kernel, and therefore $\mathrm{Cov}_{t-1}(R_t, m_t)$, depends on preferences through the marginal utility of future consumption, μ_t necessarily depends on preferences also. Preference free option valuation is therefore infeasible in this context. Only by making explicit assumptions about m_t may the covariance term be properly accounted for. For instance, under the assumption of risk neutrality $\mathrm{Cov}_{t-1}(R_t, m_t) = 0$, whereby the expected return simply equals the risk free rate, thus eliminating any preference parameters from the risk-neutral distribution.

expected discounted future value of the terminal payoff,

$$\pi_t^* = E_t^*[\pi_{t+T}(P_{t+T}^*)\exp(-\Sigma_{\tau=1,\dots,T}r_{t+\tau-1})], \tag{2}$$

where the expectation $E_t^*(\cdot)$ is evaluated under the transformed measure. Furthermore, this risk-neutralized stock price distribution takes the simple form

$$P_t^* = P_{t-1}^* \exp(r_{t-1} - 1/2\sigma_t^2 + \varepsilon_t^*), \tag{3}$$

where ε_t^* is conditionally normally distributed with mean zero and variance σ_t^2 under the risk-neutralized measure.[4] Under this transformed measure, $\varepsilon_t^* = \varepsilon_t + \mu_t - r_{t-1}$. Assumptions about μ_t therefore only affect the risk-neutralized distribution through the ε_t error term in the ARCH equation for σ_t^2. Since the logarithmic price process, $\ln(P_t)$, is closely approximated by a martingale, any reasonable specification for μ_t, will be dominated empirically by the ε_t^* innovation term in determining the adjusted price process for P_{t+T}^*. Meanwhile, the importance of the risk-neutralization scheme is underscored by the results for individual stock options in Amin and Ng (1993b). Of course, in evaluating the expectations in Eq. (2), a functional form for σ_t^2 is still called for. This is the theme of the paper.

The derivation of the risk-neutral distribution for P_{t+T}^* outlined above relied on the conditional joint normality of R_t and m_t, along with the assumption that μ_t only depends on past returns. Duan (1995a) has shown how this same pricing formula obtains in a representative agent economy with a constant relative risk aversion utility function and conditionally normally distributed returns and aggregate growth rates in consumption. Furthermore, on utilizing the generalized local risk-neutral valuation principle, this same basic formula has recently been extended by Duan (1995b) to allow for conditionally leptokurtic returns. In the subsequent analysis, we shall therefore allow the ε_t^* error term to be conditionally non-normally distributed.

Given the ARCH model estimates for μ_t and σ_t^2, the theoretical leaps prices in Eq. (2) are easily computed through Monte Carlo simulation techniques. Averaging over a large number of simulated price paths, N, the theoretical price in Eq. (2) is consistently estimated by

$$\hat{\pi}_t^* = \exp(-\Sigma_{\tau=1,\dots,T}r_{t+\tau-1})N^{-1}\Sigma_{n=1,\dots,N}\pi_{t+T}(P_{t+T,n}^*), \tag{4}$$

---

[4] The present value ARCH option prices in Engle and Mustafa (1992) are obtained by evaluating the expectation in Eq. (2) in the sampling distribution for P_{t+T} as opposed to the risk-neutralized distribution for P_{t+T}^*.

where $P^*_{t+T,n}$ denotes the nth simulated terminal price from Eq. (3). The accuracy of the numerical approximation in Eq. (4) for any finite N will depend upon the particular leap contract and the characteristics of the simulated price process. In particular, the further out-of-the-money the option, the longer the maturity time, and the more persistent the underlying price process, the higher the value of N required to achieve a given level of accuracy. The results reported below are all based on a total of $N = 5000$ replications.

The T-period risk free rate, $\Sigma_{\tau=1,\ldots,T} r_{t+\tau-1}$, is interpolated linearly from the available 1, 2, 3, 6, 12, 24, and 36 month Treasury yields for day t, as the option pricing framework does not accommodate stochastic interest rates. The relevant Treasury yields were collected from subsequent issues of the London Financial Times. Although other term structure assumptions may lead to slightly different leaps prices, the orders of magnitude are inconsequential for the relative model comparisons reported below.[5]

Explicitly modeling the dividend process for the index may also help to further improve the valuations of the leaps. Specifically, let $d_{t,T}$ denote the average expected dividend payout rate over the maturity time of the leap for the 500 stocks making up the index. The effect of dividends may then be approximately accounted for through a reduction in the risk-neutral drift; i.e., $P^*_t = P^*_{t-1} \exp(r_{t-1} - 1/2\sigma_t^2 - d_{t,T} + \varepsilon^*_t)$. In light of the steady decreasing importance of dividends over the sample, we chose to ignore this additional complication.

3. ARCH modeling and long-run volatility dependence

Implementation of the option pricing model discussed in the previous section is predicated on a specification for the conditional mean and variance processes for the underlying returns, μ_t and σ_t^2. In this section, we first present a detailed discussion of the conditional mean and volatility functions based on the daily S&P500 returns, followed by a brief summary of the corresponding estimation results obtained with weekly data.

Numerous studies have found that daily U.S. stock market index returns exhibit positive low-order serial correlation, the existence of which is often attributed to a non-synchronous trading effect; see e.g. Campbell et al. (1997). In order to allow for this dependence, a simple AR(2) specification was adopted for μ_t,

$$\mu_t = \mu + (\rho_1 L + \rho_2 L^2) R_t, \tag{5}$$

---

[5] This is consistent with the recent results in Bakshi et al. (1998) who find no significant improvements from modeling stochastic interest rates when pricing S&P500 leaps.

where L denotes the lag operator; i.e., $LR_t \equiv R_{t-1}$. During the 141 week period from 21 January 1991 to 30 September 1993 for which options data are available we re-estimated all of the models on a weekly basis. In rolling through the sample, every week we added the five most recent observations while discarding the five most distant observations, thus keeping the sample size at a constant 7500 daily returns. The maximum and minimum estimates for ρ_1 obtained across these 141 rolling samples equaled 0.120 and 0.113, respectively. Also, the 141 estimates for ρ_2 range from a low of -0.057 to a high of -0.049. Even though these estimates are all statistically significant at the conventional five percent level, the degree of return predictability is virtually zero beyond a few days.

In contrast, the conditional volatility for the daily S&P500 returns is clearly time varying and predictable. To illustrate, let $Q_{\ell,k}(|R_t|) \equiv T(T+2) \Sigma_{i=\ell,\ldots k}(T-i)^{-1} \text{Corr}(|R_t|, |R_{t-i}|)$ denote the Ljung and Box (1978) portmanteau tests for the joint significance of the absolute return autocorrelations from lag ℓ through k. Although no formal justification is available, the simulations in Bollerslev and Mikkelsen (1996) suggest that the conventional asymptotic chi-square distribution with $k - \ell + 1$ degrees of freedom works well in testing the null hypothesis of no volatility clustering in the present context. As such, it is noteworthy that for the 7500 daily returns from 24 March 1961 through 18 January 1991, that underlies the model estimates for the first leap contracts in our sample, $Q_{1,21}(|R_t|) = 4543.1$. Furthermore, this overwhelming significance remains at high lags. For instance, the portmanteau test for the lack of any serial correlation from year three through five equals $Q_{781,1301}(|R_t|) = 1519.5$.

Following the seminal work of Engle (1982), numerous ARCH and Generalized ARCH (GARCH)-type formulations have been proposed for modeling this volatility clustering phenomenon; for recent surveys of this literature see e.g. Bollerslev et al. (1992), and Bollerslev et al. (1994). Since future stock return volatility tends to respond asymmetrically to negative and positive return innovations, we shall here rely on the Exponential GARCH (EGARCH) model developed by Nelson (1991). Conventional model selection criteria suggest that an EGARCH(2, 1) specification conveniently captures the short-run serial dependence in the daily S&P500 volatility process,

$$\ln(\sigma_t^2) = \omega_t + (1 + \psi_1 L)(1 - \phi_1 L)^{-1}(1 - \phi_2 L)^{-1}\{\theta z_{t-1}$$
$$+ \gamma[|z_{t-1}| - \text{E}(|z_{t-1}|)]\}, \tag{6}$$

where $z_t \equiv \varepsilon_t \sigma_t^{-1}$. Since $\text{E}_{t-1}\{\theta z_{t-1} + \gamma[|z_{t-1}| - \text{E}(|z_{t-1}|)]\} = 0$, the EGARCH model is readily interpreted as an ARMA model for the logarithm of the conditional variance. If $\theta < 0$ future conditional variances will increase proportionally more following a negative return shock, than for a positive return shock of the same absolute magnitude. This effect is sometimes referred to as the

'leverage effect'. The importance of allowing for the asymmetry in the EGARCH formulations in the implied risk-neutral distributions when pricing shorter-lived individual U.S. stock options has previously been documented by Amin and Ng (1993b). Similarly, in their analysis of Dutch equity index options and Philips stock options, Heynen et al. (1994) find that an EGARCH model does a better job of replicating the implied term structure of volatilities, than does a symmetric GARCH formulation. Kuwahara and Marsh (1992) also find that the EGARCH model valuations of Japanese equity warrants are vastly superior to their CEV based counterparts.

To complete the specification, the intercept in the conditional variance equation is parameterized as

$$\omega_t = \omega + \ln(1 + \delta N_t), \tag{7}$$

where N_t gives the number of non-trading days between day $t - 1$ and t. For $0 < \delta < 1$ this captures the tendency for a less than proportional increase in the daily variability following weekend and holiday non-trading periods.

The resulting quasi maximum-likelihood estimates (QMLE) for the AR(2)-EGARCH(2, 1) model in Eq. (1) and (5)–(7), obtained under the auxiliary assumption of conditionally normal errors, are directly in line with previous evidence reported in the literature.[6] In particular, the inverse of the largest autoregressive root for $\ln(\sigma_t^2)$ is always very close to unity, ranging from a low of $\hat{\phi}_1 = 0.9988$ to a high of $\hat{\phi}_1 = 1.0002$ across the 141 rolling samples. The corresponding confidence intervals generally include unity. For instance, for the first of the 141 rolling samples $\hat{\phi}_1 = 0.99989$, with an asymptotic standard error of 0.00081.[7] As such, these estimates are highly suggestive of a unit root in the conditional variance equation, as in the Integrated GARCH (IGARCH) class of models proposed by Engle and Bollerslev (1986).

Even though the assumption of $\phi_1 = 1.0$ in the Integrated EGARCH (IEGARCH) model may adequately capture the short-run volatility clustering, recent evidence suggest that the long-term volatility persistence may be too high. To illustrate, consider the portmanteau tests for the first difference of the

---

[6] Apart from initial conditions, the quasi log-likelihood function for a sample of T observations takes the form, $Q \log L = -1/2 \cdot T \ln(2\pi) - 1/2 \cdot \Sigma_{t=1,...,T} [\ln(\sigma_t^2) + \varepsilon_t^2 \sigma_t^{-2}]$. Even though the standardized innovations, $z_t \equiv \varepsilon_t \sigma_t^{-1}$, may be non-normally distributed, this approach remains consistent under quite general conditions, provided that the first two conditional moments are correctly specified; see Bollerslev and Wooldridge (1992). All of the subsequent model estimates are based on this same approach.

[7] Based on daily CRSP value weighted index returns from July 1962 through December 1987, Nelson (1991) estimates the largest root in the EGARCH model to be 0.9996 with an asymptotic standard error of 0.0009.

absolute returns from lag three through five years, $Q_{781,1301}[(1 - L)|R_t|] = 701.0$. Although the test statistic is below the corresponding value for the absolute returns given above, the p-value based on the asymptotic chi-squared distribution is less than 0.0001. Meanwhile, the long-run portmanteau statistic for the fractionally differenced absolute returns equals $Q_{781,1301}[(1 - L)^{0.5}|R_t|] = 563.5$, which has a chi-square p-value of 0.091.[8] These results are also consistent with the slow hyperbolic decay of the sample autocorrelation function for the daily absolute S&P500 returns depicted in Ding, Granger and Engle (1993). Furthermore, on applying more formal log-periodogram regression procedures for estimating the value of d so that $(1 - L)^d|R_t|$ is stationary, Breidt et al. (1998) report strong evidence in favor of fractional integration for a time series of daily value-weighted CRSP index returns.

To accommodate these empirical regularities within the context of an ARCH model, Baillie et al. (1996) proposed the new class of fractionally integrated GARCH (FIGARCH) models.[9] In order to also account for the aforementioned asymmetries in equity index volatility, we shall here rely on the fractionally integrated EGARCH (FIEGARCH) formulation developed by Bollerslev and Mikkelsen (1996). Specifically, consider the following FIEGARCH$(1, d, 1)$ model:

$$\ln(\sigma_t^2) = \omega_t + (1 + \psi_1 L)(1 - \phi_1 L)^{-1}(1 - L)^{-d}\{\theta z_{t-1} + \gamma[|z_{t-1}|$$
$$- \mathrm{E}(|z_{t-1}|)]\}. \tag{8}$$

By replacing the apparent unit root in the estimated EGARCH model with the fractional differencing operator, the FIEGARCH model nests the conventional EGARCH model for $d = 0$, and the IEGARCH model for $d = 1$. However, for $0 < d < 1$ the effect of a shock to the forecast of $\ln(\sigma_{t+T}^2)$ dissipates at a slow hyperbolic rate of decay.

On actually estimating the FIEGARCH model in Eqs. (1), (5), (7) and (8), the quasi-maximum-likelihood estimates for d range from a low of 0.648 to a high of 0.708, across the 141 different rolling samples. The overall mean for $\hat{d}$ equals 0.675. All of the 141 estimates for d are also significantly different from zero and unity at conventional significance levels. For instance, for the first rolling sample

---

[8] The fractional differencing operator is formally defined by its infinite Maclaurin series expansion, $(1 - L)^d = 1 - d \cdot \Sigma_{k=1,...,\infty} \Gamma(k - d)\Gamma(1 - d)^{-1}\Gamma(k + 1)^{-1}L^k$, where $\Gamma(\cdot)$ denotes the gamma function. In all of the results reported on here the infinite summations were truncated at $k = 1000$, with all pre-sample values fixed at their unconditional sample analogues.

[9] Similar ideas within the context of stochastic volatility models have been put forth by Breidt et al. (1998), Comte and Renault (1995), and Harvey (1998).

$\hat{d} = 0.654$, with an asymptotic standard error of 0.052.[10] Thus, from a purely statistical perspective the estimates clearly suggest that the long-run dynamics in the S&P500 volatility process is best characterized by a fractionally integrated time series model.

These findings form the basis for the five different daily time series specifications for σ_t^2 investigated in the leaps valuations. For comparison purposes, all five simulation designs are based on individually estimated AR(2) models for μ_t. All of the model parameters are estimated under the assumption of conditional normality using rolling samples of the most recent 7500 daily returns until the Friday immediately preceding each of the leaps. The first two models take the variance to be higher following non-trading periods as in Eq. (7), but ignore any ARCH effects. In the first of these AR(2) benchmark models, the future risk-neutralized prices, $P_{t+T,n}^*$, are simulated under the assumption that the innovations for the future returns are normally distributed. This is consistent with the assumptions underlying the Black–Scholes formula. In the subsequent discussion we shall refer to this model as NORMAL. To accommodate the non-normality in the conditional return distributions, the out-of-sample errors for the FATTAIL model is generated by sampling with replacement from the corresponding 7500 in-sample standardized residuals; i.e., $\hat{z}_{t+\tau,n}^* \equiv \hat{\varepsilon}_{t+\tau,n}^* / \hat{\sigma}_{t+\tau}$ for $\tau = 1, 2, \ldots, 7500$. This same bootstrap procedure is employed in the simulations from the AR(2)-EGARCH(2, 1), AR(2)-IEGARCH(2, 1), and AR(2)-FIEGARCH(1, d, 1) specifications, referred to below as EGARCH, IEGARCH and FIEGARCH, respectively. Table 1 provides the QMLEs for the first of the 141 rolling sub-samples consisting of the 7500 daily returns from 24 March 1961 to 18 January 1991.[11]

The choice of a daily sampling frequency for the underlying models employed in evaluating the leaps is arbitrary. In order to investigate the robustness of the long-memory dependencies we therefore also calculated the leaps prices based on a weekly sampling frequency. The discrete time EGARCH class of models is not closed under temporal aggregation, and the weekly counterparts to the daily models in Table 1 would involve very complicated non-linear dynamics.[12]

---

[10] These estimates are in close accordance with the results in Bollerslev and Mikkelsen (1996), who on using a longer sample of 9559 daily S&P500 returns report $\hat{d} = 0.633$, with an asymptotic standard error of 0.063.

[11] Because of the high correlation among the parameter estimates, the apparent cancellation of common roots in the ARMA representations for $\ln(\sigma_t^2)$ is misleading. For instance, the conservative Schwartz information criteria for the EGARCH(2, 1) model reported in Table 1 equals 52460.8, compared to 52391.6 for the more parsimonious EGARCH(1, 0) model, providing strong evidence against any parameter redundancy.

[12] The long-memory stochastic volatility model in Comte and Renault (1995) provides an elegant theoretical approach for analyzing issues related to temporal aggregation in a continuous time setting.

Table 1
Daily S&P500 time-series models

NORMAL and FATTAIL:
$R_t = 0.00038 + 0.120 \cdot R_{t-1} - 0.057 \cdot R_{t-2} - 0.5 \cdot \sigma_t^2 + \varepsilon_t,$
$\quad\quad (0.00011) \quad (0.037) \quad\quad\quad (0.033)$
$\ln(\sigma_t^2) = -9.60 + \ln(1 + 0.373 \cdot N_t)$
$\quad\quad\quad\quad (0.04) \quad\quad\quad (0.227)$
$b_3 = -1.10 \quad b_4 = 32.9 \quad Q_{1,20}(\hat{\varepsilon}_t\hat{\sigma}_t^{-1}) = 43.9 \quad Q_{1,20}(|\hat{\varepsilon}_t\hat{\sigma}_t^{-1}|) = 5255.9$

EGARCH:
$R_t = 0.00028 + 0.180 \cdot R_{t-1} - 0.021 \cdot R_{t-2} - 0.5 \cdot \sigma_t^2 + \varepsilon_t$
$\quad\quad (0.00007) \quad (0.013) \quad\quad\quad (0.013)$
$\ln(\sigma_t^2) = -10.55 + \ln(1 + 0.163 \cdot N_t) + (1 - 0.986L)(1 - 0.99989L)^{-1}(1 - 0.940L)^{-1} \cdot g(z_{t-1})$
$\quad\quad\quad\quad (0.41) \quad\quad\quad\quad (0.040) \quad\quad\quad\quad (0.016) \quad\quad (0.00081) \quad\quad\quad (0.037)$
$g(z_t) = -0.109 \cdot z_t + 0.155 \cdot [|z_t| - E(|z_t|)]$
$\quad\quad\quad (0.018) \quad\quad (0.027)$
$b_3 = -0.33 \quad b_4 = 6.07 \quad Q_{1,20}(\hat{\varepsilon}_t\hat{\sigma}_t^{-1}) = 16.8 \quad Q_{1,20}(|\hat{\varepsilon}_t\hat{\sigma}_t^{-1}|) = 28.6$

IEGARCH:
$R_t = 0.00026 + 0.180 \cdot R_{t-1} - 0.021 \cdot R_{t-2} - 0.5 \cdot \sigma_t^2 + \varepsilon_t$
$\quad\quad (0.00006) \quad (0.024) \quad\quad\quad (0.024)$
$\ln(\sigma_t^2) = -10.58 + \ln(1 + 0.163 \cdot N_t) + (1 - 0.987L)(1 - L)^{-1}(1 - 0.943L)^{-1} \cdot g(z_{t-1})$
$\quad\quad\quad\quad (0.39) \quad\quad\quad\quad (0.040) \quad\quad\quad\quad (0.012) \quad\quad\quad\quad\quad (0.030)$
$g(z_t) = -0.108 \cdot z_t + 0.154 \cdot [|z_t| - E(|z_t|)]$
$\quad\quad\quad (0.018) \quad\quad (0.027)$
$b_3 = -0.32 \quad b_4 = 6.04 \quad Q_{1,20}(\hat{\varepsilon}_t\hat{\sigma}_t^{-1}) = 17.0 \quad Q_{1,20}(|\hat{\varepsilon}_t\hat{\sigma}_t^{-1}|) = 28.2$

FIEGARCH:
$R_t = 0.00032 + 0.179 \cdot R_{t-1} - 0.025 \cdot R_{t-2} - 0.5 \cdot \sigma_t^2 + \varepsilon_t$
$\quad\quad (0.00007) \quad (0.013) \quad\quad\quad (0.013)$
$\ln(\sigma_t^2) = -10.34 + \ln(1 + 0.162 \cdot N_t) + (1 - 0.684L)(1 - L)_{(0.052)}^{-0.654}(1 - 0.777L)^{-1} \cdot g(z_{t-1})$
$\quad\quad\quad\quad (0.35) \quad\quad\quad\quad (0.039) \quad\quad\quad\quad (0.112) \quad\quad\quad\quad\quad (0.087)$
$g(z_t) = -0.134 \cdot z_t + 0.188 \cdot [|z_t| - E(|z_t|)]$
$\quad\quad\quad (0.024) \quad\quad (0.030)$
$b_3 = -0.32 \quad b_4 = 5.80 \quad Q_{1,20}(\hat{\varepsilon}_t\hat{\sigma}_t^{-1}) = 16.2 \quad Q_{1,20}(|\hat{\varepsilon}_t\hat{\sigma}_t^{-1}|) = 30.0$

Note: The model estimates are based on 7500 daily return observations for the S&P500 composite index excluding dividends; i.e., $R_t \equiv \ln(P_t/P_{t-1})$. The sample period extends from 24 March 1961 to 18 January 1991. All the parameter estimates are obtained under the assumption of conditional normality; i.e., $z_t \equiv \varepsilon_t\sigma_t^{-1}$ i.i.d. N(0, 1). Robust standard errors, as in Bollerslev and Wooldridge (1992), are reported in parentheses. The skewness and kurtosis of the standardized residuals, $\hat{z}_t = \hat{\varepsilon}_t\hat{\sigma}_t^{-1}$, are denoted by b$_3$ and b$_4$, respectively. $Q_{1,20}(\cdot)$ refer to the Ljung and Box (1978) portmanteau test for up to twentieth order serial correlation.

Meanwhile, the EGARCH models employed throughout the empirical finance literature are best viewed as convenient approximations to the true data generating process. Since the purpose of the present paper is to learn about the quality of these competing discrete time approximations, and the role of

long-memory in particular, we choose to rely on the same five popular models in characterizing the weekly returns.

However, the dynamic dependencies are somewhat simpler at the weekly level. The returns are approximately serially uncorrelated. The N_t no-trade dummy variable is no longer important. Furthermore, conventional testing procedures suggest a cancellation of the roots in the conditional variance equations, so that the orders of the three weekly ARCH models are EGARCH(1, 0), IEGARCH(1, 0) and FIEGARCH(0, d, 0), respectively. Table 2

Table 2
Weekly S&P500 time series models

NORMAL and FATTAIL:
$R_t = 0.0013 - 0.5 \cdot \sigma_t^2 + \varepsilon_t$
       (0.0005)
$\ln(\sigma_t^2) = -7.78$
            (0.06)
$b_3 = -0.31 \quad b_4 = 6.49 \quad Q_{1,20}(\hat{\varepsilon}_t\hat{\sigma}_t^{-1}) = 32.5 \quad Q_{1,20}(|\hat{\varepsilon}_t\hat{\sigma}_t^{-1}|) = 620.7$

EGARCH:
$R_t = 0.0012 - 0.5 \cdot \sigma_t^2 + \varepsilon_t$
       (0.0004)
$\ln(\sigma_t^2) = -7.99 + (1 - 0.946)^{-1} \cdot g(z_{t-1})$
            (0.14)   (0.018)
$g(z_t) = -0.117 \cdot z_t + 0.243 \cdot [|z_t| - E(|z_t|)]$
          (0.025)   (0.040)
$b_3 = -0.34 \quad b_4 = 3.98 \quad Q_{1,20}(\hat{\varepsilon}_t\hat{\sigma}_t^{-1}) = 23.4 \quad Q_{1,20}(|\hat{\varepsilon}_t\hat{\sigma}_t^{-1}|) = 14.7$

IEGARCH:
$R_t = 0.0013 - 0.5 \cdot \sigma_t^2 + \varepsilon_t$
       (0.0004)
$\ln(\sigma_t^2) = -8.85 + (1 - L)^{-1} \cdot g(z_{t-1})$
            (0.34)
$g(z_t) = -0.076 \cdot z_t + 0.229 \cdot [|z_t| - E(|z_t|)]$
          (0.018)   (0.042)
$b_3 = -0.50 \quad b_4 = 4.22 \quad Q_{1,20}(\hat{\varepsilon}_t\hat{\sigma}_t^{-1}) = 25.9 \quad Q_{1,20}(|\hat{\varepsilon}_t\hat{\sigma}_t^{-1}|) = 17.7$

FIEGARCH:
$R_t = 0.0019 - 0.5 \cdot \sigma_t^2 + \varepsilon_t$
       (0.0003)
$\ln(\sigma_t^2) = -8.68 + (1 - L)_{(0.038)}^{-0.634} \cdot g(z_{t-1})$
            (0.35)
$g(z_t) = -0.198 \cdot z_t + 0.265 \cdot [|z_t| - E(|z_t|)]$
          (0.037)   (0.045)
$b_3 = -0.54 \quad b_4 = 4.15 \quad Q_{1,20}(\hat{\varepsilon}_t\hat{\sigma}_t^{-1}) = 22.4 \quad Q_{1,20}(|\hat{\varepsilon}_t\hat{\sigma}_t^{-1}|) = 15.5$

Note: See Table 1. The estimates are based on 1560 weekly returns from 24 March 1961 to 18 January 1991.

summarizes these weekly QMLE's for the 1560 returns in the first of the 141 rolling sub-samples.

The estimate of $\hat{d} = 0.634$ with an asymptotic standard error of 0.038 for the FIEGARCH model is particularly noteworthy. This value of d is almost identical to the value obtained for the corresponding daily FIEGARCH model reported in Table 1, and highlights the robustness of the estimated long-memory dependencies. This commonality also implies that for longer horizons the simulated risk-neutral distributions for P^*_{t+T} should be very similar for the daily and weekly FIEGARCH models. In contrast, the estimated EGARCH parameters differ markedly across Tables 1 and 2. The half-life of a volatility shock for the weekly EGARCH model is only $-\ln(2)/\ln(0.94626) \approx 12.55$ weeks, or 63 days, compared to an estimated half-life of $-\ln(2)/\ln(0.99989) \approx 6301$ days for the largest root in the daily EGARCH model. The EGARCH model prices may therefore depend importantly on the sampling frequency in this case because the model is misspecified.

4. Leaps data and illustrative examples

Leaps were first created and marketed by the CBOE in October, 1990. Aside from the longer maturity times, the margin requirements and other institutional regulations mirror those for conventional, shorter-lived, exchange traded equity options. The data set analyzed here consist of the daily closing prices for all the CBOE traded S&P500 leaps contracts over the 141 week period from 21 January 1991 to 30 September 1993, for a total of 3256 observations.[13] The leaps expire on the third Saturday in December within one of the following 3 yr. The average maturity time is 489 calendar days, or 340 trading days. Since the S&P500 leaps are European-type contracts, there are no complications associated with early exercise premia. In order to minimize the problems associated with thin trading, we excluded all leaps with daily trading volume of four or less contracts. Of the resulting 2912 price observations, 286 and 2626 are calls and puts, respectively. The corresponding aggregate put trading volume over the sample period is 375,305 contracts compared to only 46,989 call contracts. This

---

[13] With multiple contracts written on the same date the option pricing errors obey a panel data structure, and an assumption of independent stochastic pricing errors is logically inconsistent with a model based on a single state variable. Empirical non-degenerate error distributions could be rationalized by the introduction of additional state variables, transactions costs, non-synchronous prices, or other market frictions. Renault (1997) provides a lucid discussion of these and related issues that arise in the actual estimation of stochastic volatility models based on option pricing errors.

reflects the predominant use of the S&P500 leaps in portfolio insurance type strategies, the scope of which has increased enormously since the October 1987 stock market crash. In the subsequent analysis we shall therefore concentrate exclusively on the empirical findings for the puts.[14]

Our focus on the put prices is also motivated by the observation, that the corresponding left tails of the simulated risk-neutral distributions afford the greatest room for differentiating among the alternative time series models in Tables 1 and 2. Numerous studies have also documented a tendency for the implied risk neutral distributions calculated from regular options to exhibit increasing negative skewness following the October 1987 stock market crash; see e.g. the survey by Bates (1996). The non-parametric density estimates for the risk neutral distributions implicit in regular S&P500 index options in Aït-Sahalia and Lo (1998) suggest that this negative skewness may in fact be increasing with the maturity time, and as such may be very informative about longer-run dependencies. Furthermore, Bates (1995) has recently argued that this long-lived 'skewness premium' seems particularly difficult to explain within the context of standard finite variance jump diffusion processes which tend to converge too quickly to their log-normal counterparts.

Before proceeding to an overall summary of the results, it is informative to consider the findings for a few specific contracts. To this end, Table 3 lists the characteristics for four sets of contracts traded at the beginning, middle and end of the sample period. Since the leaps are written on 1/10 of the S&P500 index, the closing prices listed under P_t have been scaled accordingly. In addition to the actual contract specifications and closing prices, π_t, the column labeled $\hat{\sigma}_t$ reports the one-step ahead 'annualized' percentage volatilities from the estimated daily AR(2)-FIEGARCH(1, d, 1) models. The simulated theoretical risk-neutralized leaps prices from the daily NORMAL, FATTAIL, EGARCH, IEGARCH, and FIEGARCH models are reported in the last five columns.

The first two leaps in our sample are written on 21 January 1991 with expiration 485 trading days later on 18 December 1992. The exercise prices for these two puts are 30 and 35, respectively. The annualized volatility for 21 January 1991 from the FIEGARCH model is 0.160. In comparison, the median S&P500 FIEGARCH volatility estimates from 21 January 1991 through 30 September 1993 equals 0.111, ranging from a low of 0.067 on 31 August 1993 to a high of 0.281 on 15 November 1991. These one-step ahead conditional variances are quite close to the results for the EGARCH and IEGARCH models. Meanwhile, the in-sample annualized volatility estimates from the

---

[14] The findings for the 286 calls are available upon request.

Table 3
S&P500 put leaps pricing examples

t	$t+T$	T	$\hat{\sigma}_t$	P_t	X_t	π_t	NORMAL	FATTAIL	EGARCH	IEGARCH	FIEGARCH
1/21/91	12/18/92	485	0.160	33.109	30.0	1.875	0.348	0.476	1.728	1.642	1.697
—	—	—	—	—	35.0	3.625	1.484	1.795	3.359	3.210	3.429
1/21/91	12/17/93	737	0.160	33.109	35.0	3.625	1.156	1.530	3.659	3.487	3.615
5/15/92	12/17/93	403	0.140	41.007	40.0	2.438	1.277	1.550	1.922	1.779	2.166
—	—	—	—	—	42.5	3.500	2.158	2.575	2.838	2.642	3.147
9/1/93	12/16/94	328	0.068	46.316	40.0	0.813	0.391	0.511	0.366	0.297	0.635
—	—	—	—	—	42.5	1.125	0.819	1.003	0.705	0.591	1.109
—	—	—	—	—	45.0	1.938	1.520	1.765	1.284	1.100	1.877

Note: The S&P500 put leaps are traded on date t with expiration T trading days later on day $t + T$. The exercise prices are denoted by X_t. The annualized FIEGARCH volatility estimate for day t is denoted by $\hat{\sigma}_t$, while P_t gives 1/10 times the S&P500 index value for day t. The actual CBOE leaps transaction prices are listed in the π_t column. The simulated theoretical model prices are given in the last five columns.

NORMAL and FATTAIL models ranged from a low of 0.130 to a high of only 0.173.

Turning to the simulated prices in the last five columns, it is noteworthy that all three ARCH models are quite close to the actual transaction prices of $1\frac{7}{8}$ and $3\frac{5}{8}$. In contrast, the NORMAL and FATTAIL model both severely underprices these two leaps. The reason is immediately evident from the first panel of Fig. 1, which graphs the corresponding simulated risk-neutralized probability densities for P^*_{t+T}. Although the upper tails of all five 485 steps-ahead conditional densities are fairly similar, the asymmetric response to negative and positive shocks in the EGARCH formulations give rise to very differently shaped downside risks. Since the $\hat{\pi}^*_t$ put prices in Eq. (4) are evaluated as the discounted integral under each of these densities below the exercise price X_t, it is obvious
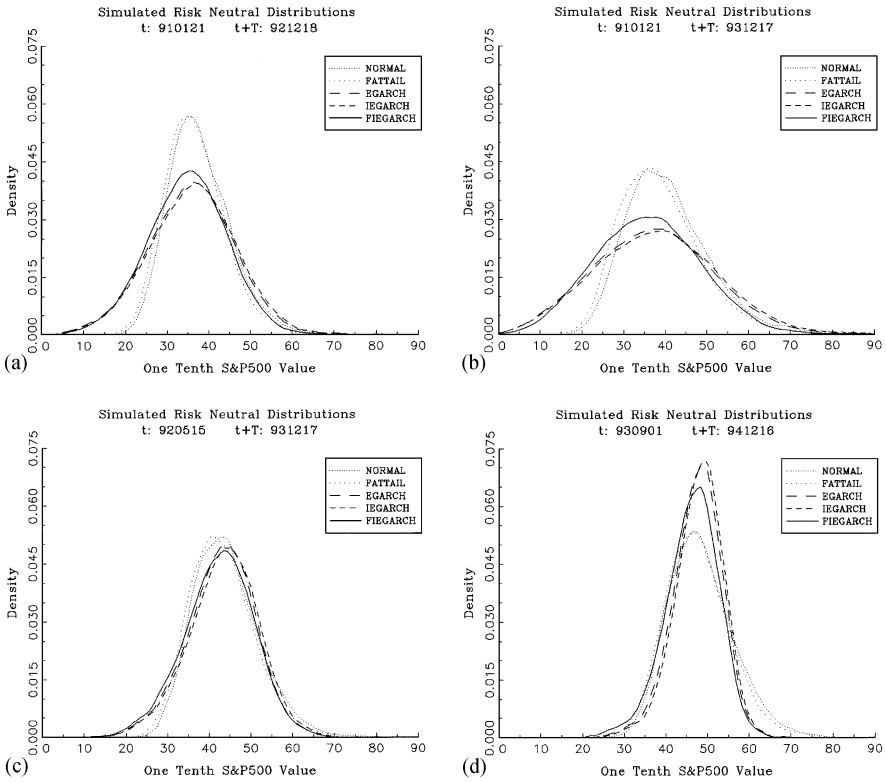

Fig. 1. The figure graphs the simulated risk-neutral distributions for P^*_{t+T} corresponding to the eleven leaps listed in Table 3 for each of the five different daily model specifications. The leaps are written on day t and expire on day $t + T$.

that the prices for the NORMAL and FATTAIL models will differ substantially from the three ARCH prices in this situation.[15]

While the differences between the three ARCH formulations for the first two leaps are relatively minor, the valuations for the longer-lived put with 737 days until expiration that were traded on the same date clearly favor the FIEGARCH model. The corresponding risk-neutral distributions in the second panel of Fig. 1 highlight the larger discrepancies across the models when forecasting almost three years into the future. Of course, the resulting distributions for P^*_{t+737} are also more dispersed than the distributions for P^*_{t+485} in the first panel of the figure.

The effect of the reduced market volatility for the two puts traded on 15 May 1992 is immediately evident. The discrepancies across the simulated risk-neutral distributions graphed in the third panel of Fig. 1 and the corresponding model prices are markedly smaller in this situation. Yet, it is noteworthy that for both puts the FIEGARCH prices are the most accurate.

The last three leaps reported in Table 3 were written on 1 September 1993 towards the end of our sample period. The annualized volatility at that time was only 0.068, which gives rise to even more peaked 328 days-ahead risk-neutralized distributions. Note that for these three contracts, the FATTAIL model actually yields better market valuations than do the EGARCH and IEGARCH formulations. The FIEGARCH model prices are again the closest to the actual π_t transactions prices.

5. Leaps pricing errors

The specific leap valuations discussed in Table 3 above leave the impression that the FIEGARCH model tends to produce the most accurate prices, followed by the other EGARCH formulations. This is confirmed by a more detailed summary of the pricing errors for the full data set. In particular, on comparing the accuracy of all the 2626 daily FIEGARCH model prices to the other daily model valuations by a simple non-parametric sign test, the values of the test statistics for the NORMAL, FATTAIL, EGARCH and IEGARCH models are

---

[15] As discussed above, our simulation based procedure for evaluating options will likely be less accurate for out-of-the-money options with long maturities. The first option in Table 3 is a deep-out-of-the-money put with maturity well above average. In order to gauge the accuracy of the procedure in a worst-case-scenario, we therefore simulated one-hundred independent leaps prices, each based on $N = 5000$ replications, for this particular contract specification. The standard deviations for the five different model prices ranged from 0.015 to 0.056, or roughly three percent for the preferred FIEGARCH model. Meanwhile, most of the comparative statistics reported below are based on averages across several thousand different leaps, thus further reducing any numerical inaccuracy.

40.3, 30.1, 26.2, and 23.3, respectively.[16] Under the null hypotheses that the daily FIEGARCH model and the comparison models are equally accurate, these test statistics should be asymptotically standard normally distributed. This hypothesis is overwhelmingly rejected in favor of the FIEGARCH based valuations for all the pairwise comparisons. On comparing the accuracy of the EGARCH model and the FATTAIL model, this test statistic equals 2.888. While this value is less striking, the corresponding one-sided hypothesis test for no difference between these two valuation schemes still has a *p*-value of only 0.002.

This same ranking of the daily models is borne out by standard summary measures for the average pricing errors. Specifically, define Bias $\equiv K^{-1}\Sigma_{k=1,...,K}$ $(\hat{\pi}_k^* - \pi_k)$, where the subscript k indexes the K leaps available for the particular comparison. Also, let the mean absolute error (MAE) and the mean-squared error (MSE) be defined by MAE $\equiv K^{-1}\Sigma_{k=1,...,K}|\hat{\pi}_k^* - \pi_k|$ and MSE $\equiv K^{-1}\Sigma_{k=1,...,K}$ $(\hat{\pi}_k^* - \pi_k)^2$, respectively. Of course, none of these three measures correspond directly to the loss function that would be of immediate interest to market participants. However, the numbers do provide a simple yardstick for the long-run volatility dynamics implicit in the actual leaps transactions prices.

The first set of figures reported in Table 4 give the averages calculated over all the 2626 leaps. These measures clearly favor the FIEGARCH model. The results for the EGARCH and IEGARCH specifications are fairly close and by all three measures substantially lower than the average errors for the homoskedastic FATTAIL model. This is directly in line with previous studies that have documented the importance of allowing for time-varying volatility in the pricing of regular options. Not surprisingly, the NORMAL model performs the worst by all three criteria.[17] The next two groups of numbers in the table decompose the performance into conventional S&P500 index like options, or leaps with maturity times less than nine months, and genuine leaps with maturities ranging from nine months to three years. Whereas the results for the shorter lived contracts are less clear-cut, the findings for the 2049 puts with maturities exceeding nine months are unambiguous. When judged by the MAE or MSE error criteria, the long-term volatility dynamic implied by these most active leaps is best described by a slowly mean-reverting fractionally integrated process.

---

[16] The test statistic is calculated as $Z = (\hat{\kappa} - 0.5 \cdot K)/(0.5 \cdot K^{1/2})$, where $\hat{\kappa}$ denotes the number of times that the absolute differences between the baseline model prices and the actual transactions prices are less than the absolute pricing errors for the comparison model across all $K = 2626$ leaps.

[17] We also experimented with the use of shorter rolling samples in the estimation of the daily NORMAL model. The results from these estimates were generally worse. For instance, on estimating the volatility from the most recent sixty days, as it is often done by investment analysts in pricing regular short-lived options, the Bias, MAE and MSE measures were -1.194, 1.194, and 1.814, respectively. Further details are available upon request.

Table 4
Average put leaps pricing errors: daily S&P500 time series models

	NORMAL	FATTAIL	EGARCH	IEGARCH	FIEGARCH
All leaps ($K = 2626$)					
Bias	-0.815	-0.623	-0.270	-0.423	0.070
MAE	0.817	0.630	0.484	0.478	0.379
MSE	0.947	0.591	0.375	0.339	0.265
$T \leqslant$ Nine months ($K = 577$)					
Bias	-0.225	-0.147	-0.267	-0.296	-0.101
MAE	0.237	0.178	0.267	0.296	0.173
MSE	0.086	0.050	0.099	0.115	0.052
$T >$ Nine months ($K = 2049$)					
Bias	-0.981	-0.757	-0.271	-0.459	0.117
MAE	0.981	0.758	0.545	0.530	0.437
MSE	1.189	0.743	0.453	0.402	0.324

Note: The average pricing errors are calculated from all the CBOE traded S&P500 put leaps from 21 January 1991 to 30 September 1993, for which the daily trading volume exceeded four contracts. The time-series models for the S&P500 index are estimated with 7500 daily returns. The Bias is defined by $K^{-1}\Sigma_{k=1,...,K}(\hat{\pi}_k^* - \pi_k)$, where K denotes the number of observations for the particular group of leaps and $\hat{\pi}_k^*$ indicate the theoretical simulated model prices. The mean absolute error (MAE) and the mean-squared error (MSE) are averages of $|\hat{\pi}_k^* - \pi_k|$ and $(\hat{\pi}_k^* - \pi_k)^2$, respectively.

Table 5 contains a similar breakdown of the relative pricing errors; i.e., RBias $\equiv K^{-1}\Sigma_{k=1,...,K}(\hat{\pi}_k^* - \pi_k)/\pi_k$, RMAE $\equiv K^{-1}\Sigma_{k=1,...,K}|\hat{\pi}_k^* - \pi_k|/\pi_k$, and RMSE $\equiv K^{-1}\Sigma_{k=1,...,K}(\hat{\pi}_k^* - \pi_k)^2/\pi_k^2$. The exact same model rankings obtain for the longest lived contracts. The FIEGARCH model is superior by all three criteria.

To further explore the robustness of these findings to the choice of a daily sampling frequency, we report the same summary measures for the weekly models in Tables 6 and 7. To conserve space, the tables only give the results for the 2049 genuine leaps with maturities in excess of nine months. The average pricing errors for the NORMAL model are almost identical across Tables 4 and 6, and Tables 5 and 7. The performance of the weekly FATTAIL model is very close to that of the NORMAL model, and somewhat worse than the daily FATTAIL model. This reflects the faster convergence to unconditional normality of the temporally aggregated weekly in-sample residuals. For the shorter maturity options (not reported in the tables) the differences between the two models are larger. Whereas the daily EGARCH and IEGARCH models clearly dominate the daily NORMAL and FATTAIL models, the average pricing errors for the weekly models in Tables 6 and 7 are close, and markedly worse than the daily EGARCH summary measures in Tables 4 and 5. Meanwhile, the superiority of the FIEGARCH model prices is ubiquitous. The commonality in

Table 5
Average relative put leaps pricing errors: daily S&P500 time-series models

	NORMAL	FATTAIL	EGARCH	IEGARCH	FIEGARCH
All leaps ($K = 2626$)					
RBias	− 0.494	− 0.404	− 0.220	− 0.299	− 0.050
RMAE	0.497	0.410	0.331	0.326	0.236
RMSE	0.313	0.237	0.164	0.153	0.090
$T \leqslant$ Nine months ($K = 577$)					
RBias	− 0.376	− 0.299	− 0.404	− 0.429	− 0.242
RMAE	0.387	0.325	0.404	0.429	0.280
RMSE	0.243	0.190	0.227	0.247	0.141
$T >$ Nine months ($K = 2049$)					
RBias	− 0.528	− 0.434	− 0.168	− 0.262	0.004
RMAE	0.528	0.434	0.311	0.297	0.224
RMSE	0.333	0.250	0.146	0.127	0.076

Note: The average relative pricing errors are based on all of the CBOE traded S&P500 put leaps from 21 January, 1991 to 30 September, 1993, for which the daily trading volume exceeded four contracts. The time series models for the S&P500 index are estimated with 7,500 daily returns. The Relative Bias (RBias) is given by $K^{-1}\Sigma_{k=1,...,K}(\hat{\pi}_k^* - \pi_k)/\pi_k$, where K refers to the number of observations for the particular group of leaps and $\hat{\pi}_k^*$ denotes the theoretical simulated model prices. The relative mean absolute error (RMAE) and the Relative Mean-Squared Error (RMSE) are defined by $K^{-1}\Sigma_{k=1,...,K}|\hat{\pi}_k^* - \pi_k|/\pi_k$ and $K^{-1}\Sigma_{k=1,...,K}(\hat{\pi}_k^* - \pi_k)^2/\pi_k^2$, respectively.

Table 6
Average put leaps pricing errors: weekly S&P500 time-series models

	NORMAL	FATTAIL	EGARCH	IEGARCH	FIEGARCH
$T >$ Nine months ($K = 2049$)					
Bias	− 0.948	− 0.948	− 1.082	− 0.859	− 0.219
MAE	0.948	0.948	1.083	0.921	0.315
MSE	1.133	1.136	1.428	1.196	0.193

Note: See Table 4. The time-series models for the S&P500 index are estimated with 1560 weekly returns.

the QMLE estimates of d for the daily and weekly FIEGARCH models also suggest that the long-run forecasts from this model should be similar across the two frequencies. This is indeed the case. The close correspondence between the pricing errors in Tables 4 and 6, and Tables 5 and 7 indicate a remarkable robustness of the FIEGARCH model prices with respect to the sampling frequency. In fact, the MAE and MSE for the weekly model are both less than the corresponding daily numbers.

Table 7
Average relative put leaps pricing errors: weekly S&P500 time-series models

	NORMAL	FATTAIL	EGARCH	IEGARCH	FIEGARCH
$T >$ Nine months ($K = 2049$)					
RBias	− 0.508	− 0.506	− 0.529	− 0.384	− 0.129
RMAE	0.508	0.506	0.535	0.460	0.206
RMSE	0.317	0.314	0.313	0.296	0.125

Note: See Table 5. The time-series models for the S&P500 index are estimated with 1560 weekly returns.

While the FIEGARCH valuations give rise to the most accurate prices, important biases remain. To illustrate, consider the linear regression involving the daily FIEGARCH model relative pricing errors for the 2049 long-lived puts,

$$
(\hat{\pi}_k^* - \pi_k)/\pi_k = \underset{(0.020)}{-0.425} \quad \underset{(0.049)}{-0.832} \quad \max\{0, (P_k - X_k)/X_k\}
$$
$$
+ \underset{(0.00005)}{0.00075} \quad T_k + \underset{(0.228)}{1.126} \quad \hat{\sigma}_k + \underset{(0.006)}{0.015} \quad r_k + \varepsilon_k,
\tag{9}
$$

where the numbers in parentheses give the Eicker–White heteroskedasticity robust standard errors. The significant relationship between the relative pricing errors and the one-step-ahead FIEGARCH volatility estimates is evident. The results also indicate a serious underpricing of deep out-of-the-money puts. These are often quoted at a quarter of a dollar or less, so that bid-ask spreads and transaction costs may be important factors. The residual sum of squares from the ad hoc regression in Eq. (9) is only 0.047, compared to the MSE for the 'structural' daily FIEGARCH model prices of 0.078 reported in Table 5. Furthermore, by including all the 10 possible interaction terms among the explanatory variables on the right-hand side, the residual sum of squares drops to only 0.037. Thus, while the theoretical FIEGARCH model based valuations give rise to the lowest average error rates across the five models, the magnitude of the pricing errors could easily be reduced through the estimation of a non-structural statistical relationship.

6. Conclusion

The specification and comparison of competing time-series models for financial market volatility has been the topic of much recent research. Most of the forecast comparisons reported in the literature have invariably focused on

relatively short-lived volatility dynamics. In contrast, the primary goal of the present study was to shed new light on the longer-run dependencies in aggregate stock market volatility. To this end we employed a new data set consisting of S&P500 index leaps. The results from comparing the actual transactions prices for these leaps to various simulated theoretical risk-neutralized ARCH prices were striking. The summary measures for the average pricing errors from the FIEGARCH model for the longest-lived, most actively traded leap puts uniformly dominated the valuations obtained with the more traditional EGARCH and IEGARCH specifications. As such, our findings are consistent with the recent time-series-based evidence for the statistical significance of fractional integration in financial market volatility. Of course, allowing for fractional integration can not in a general sense result in a worse in-sample fit. However, the present findings add an important new dimension to the existing statistical evidence for in-sample significance, by highlighting the economic importance of incorporating long-memory features in out-of-sample forecasting. Yet in conclusion, it is important to recognize, that significant biases remained in the preferred FIEGARCH leap valuations. We leave further investigations and volatility modeling refinements along these lines for future research.

Acknowledgements

References

Aït-Sahalia, Y., Lo, A.W., 1998. Nonparametric estimation of state-price densities implicit in financial asset prices. Journal of Finance 53, 499–547.

Amin, K.I., Ng, V., 1993a. Option valuation with systematic stochastic volatility. Journal of Finance 48, 881–909.

Amin, K.I., Ng, V., 1993b. ARCH processes and option valuation. Working paper, Department of Finance, University of Michigan.

Bakshi, G., Cao, C., Chen, Z., 1998. Pricing and Hedging long-term options. Journal of Econometrics, forthcoming.

Baillie, R.T., Bollerslev, T., Mikkelsen, H.O., 1996. Fractionally integrated generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 74, 3–30.

Bates, D.S., 1996. Testing option pricing models. In: Maddala, G.S., Rao, C.R. (Eds.), Handbook of Statistics, vol. 14, Statistical Methods in Finance. North-Holland, Amsterdam, pp. 567–611.

Bates, D.S., 1995. Post-'87 Crash fears in S and P 500 futures options. Working paper, Department of Finance, The Wharton School, University of Pennsylvania.

Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. Journal of Political Economy 81, 637–659.

Bollerslev, T., Chou, R.Y., Kroner, K.F., 1992. ARCH modeling in finance: a review of the theory and empirical evidence. Journal of Econometrics 52, 5–59.

Bollerslev, T., Engle, R.F., Nelson, D.B., 1994. ARCH models. In: Engle, R.F., McFadden, D. (Eds.), Handbook of Econometrics, vol. 4, North-Holland, Amsterdam, pp. 2959–3038.

Bollerslev, T., Mikkelsen, H.O., 1996. Modeling and pricing long-memory in stock market volatility. Journal of Econometrics 73, 151–184.

Bollerslev, T., Wooldridge, J.M., 1992. quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. Econometric Reviews 11, 143–172.

Breidt, F.J., Crato, N., de Lima, P., 1998. On the detection and estimation of long memory in stochastic volatility. Journal of Econometrics 83, 325–348.

Campbell, J.Y., Lo, A.W., MacKinlay, A.C., 1997. The Econometrics of Financial Markets. Princeton University Press, Princeton.

Comte, F., Renault, E., 1998. Long-memory in continuous time stochastic volatility models. Mathematical Finance, forthcoming.

Ding, Z., Granger, C.W.J., Engle, R.F., 1993. A long memory property of stock market returns and a new model. Journal of Empirical Finance 1, 83–106.

Duan, J.-C., 1995a. The GARCH option pricing model. Mathematical Finance 5, 13–32.

Duan, J.-C., 1995b. Option pricing under stochastic volatility and conditional fat-tailed distributions. Working paper, Faculty of Management, McGill University.

Engle, R.F., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of U. K. inflation. Econometrica 50, 987–1008.

Engle, R.F., Bollerslev, T., 1986. Modeling the persistence of conditional variances. Econometric Reviews 5, 1–50.

Engle, R.F., Mustafa, C., 1992. Implied ARCH models from options prices. Journal of Econometrics 52, 289–311.

Ghysels, E., Harvey, A., Renault, E., 1996. Stochastic volatility. In: Maddala, G.S., Rao, C.R. (Eds.), Handbook of Statistics, vol. 14, Statistical Methods in Finance. North-Holland, Amsterdam, pp. 119–191.

Harrison, G.M., Kreps, D., 1979. Martingales and arbitrage in multi-period securities markets. Journal of Economic Theory 20, 381–408.

Harvey, A.C., 1998. Long-memory in stochastic volatility. In: Knight, J., Satchell, S. (Eds.), Forecasting Volatility in Financial Markets. Butterworth-Heineman, London, forthcoming.

Heynen, R., Kemna, A., Vorst, T., 1994. Analysis of the term structure of implied volatilities. Journal of Financial and Quantitative Analysis 29, 31–56.

Hull, J., White, A., 1987. The pricing of options on assets with stochastic volatilities. Journal of Finance 42, 381–400.

Kuwahara, H., Marsh, T.A., 1992. The pricing of Japanese equity warrants. Management Science 38, 1610–1641.

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

Melino, A., 1994. The estimation of continuous-time models in finance. In: Sims, C.A. (Ed.), Advances in Econometrics, vol. 2, Cambridge University Press, Cambridge, pp. 313–351.

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

Pastorello, S., Renault, E., Touzi, N., 1996. Statistical inference for random variance option pricing. Working paper, GREMAQ-IDEI, University of Toulouse.

Renault, E., 1997. Econometric models of option pricing errors. In: Kreps, D.M., Wallis, K.F. (Eds.), Advances in Econometrics, vol. 3, Cambridge University Press, Cambridge, pp. 223–278.

Shephard, N., 1995. Statistical aspects of ARCH and stochastic volatility. In: Cox, D.R., Hinkley, D.V., Barndorff-Nielsen, O.E. (Eds.), Likelihood, Time Series with Econometric and Other Applications. Chapman & Hall, London, pp. 1–67.