

# Modeling and Forecasting Stock Return Volatility Using a Random Level Shift Model

Yang K. Lu<sup>†</sup>

Pierre Perron<sup>‡</sup>

Boston University

Boston University

July 28, 2008; Revised April 4, 2009.

## Abstract

We consider the estimation of a random level shift model for which the series of interest is the sum of a short memory process and a jump or level shift component. For the latter component, we specify the commonly used simple mixture model such that the component is the cumulative sum of a process which is 0 with some probability  $(1 - \alpha)$  and is a random variable with probability  $\alpha$ . Our estimation method transforms such a model into a linear state space with mixture of normal innovations, so that an extension of Kalman filter algorithm can be applied. We apply this random level shift model to the logarithm of absolute returns for the S&P 500, AMEX, Dow Jones and NASDAQ stock market return indices. Our point estimates imply few level shifts for all series. But once these are taken into account, there is little evidence of serial correlation in the remaining noise and, hence, no evidence of long-memory. Once the estimated shifts are introduced to a standard GARCH model applied to the returns series, any evidence of GARCH effects disappears. We also produce rolling out-of-sample forecasts of squared returns. In most cases, our simple random level shifts model clearly outperforms a standard GARCH(1,1) model and, in many cases, it also provides better forecasts than a fractionally integrated GARCH model.

**JEL Classification Number:** C22.

**Keywords:** structural change, forecasting, GARCH models, long-memory.

---

\*Perron acknowledges financial support for this work from the National Science Foundation under Grant SES-0649350. We are grateful to Zhongjun Qu and Adrien Verdelhan for helpful comments, and to Adam McCloskey for detailed comments on a previous draft that improved the presentation.

<sup>†</sup>Department of Economics, Boston University, 270 Bay State Rd., Boston, MA, 02215 (yanglu@bu.edu).

<sup>‡</sup>Department of Economics, Boston University, 270 Bay State Rd., Boston, MA, 02215 (perron@bu.edu).

## 1 Introduction

In the economics and finance literature, there has long existed an interest in long-memory models, especially regarding applications to financial times series. For reviews of this literature, see Robinson (1994), Beran (1994) and Baillie (1996). Of particular interest to econometrics is the fractionally integrated process (see Granger, 1981, Granger and Joyeux, 1980, and Hosking, 1981) whose difference of order  $d$  is a short-memory process, often modeled as an ARMA process for which autocorrelations decay exponentially. The parameter  $d$  can be non-integer valued and when  $0 < d < 1$ , the autocorrelations decay very slowly, a characteristic of long-memory processes. Recently, there has been an upsurge of interest in the possibility of confusing long-memory with structural change in levels. This idea extends that expositied by Perron (1989, 1990) who showed that structural change and unit roots ( $d = 1$ ) are easily confused: when a stationary process is contaminated by structural change, the estimate of the sum of its autoregressive coefficients is biased towards 1 and tests of the null hypothesis of a unit root are biased toward non-rejection. This phenomenon has been shown to apply to the long-memory context as well. That is, when a stationary short-memory process is contaminated by structural change in levels, the estimate of  $d$  is biased away from 0 and the autocovariance function of the process exhibits a slow rate of decay. Relevant references on this issue include Diebold and Inoue (2001), Engle and Smith (1999), Gouriéroux and Jasiak (2001), Granger and Ding (1996), Granger and Hyung (2004), Lobato and Savin (1998), Mikosch and Stărică (2004), Parke (1999) and Teverosovky and Taqqu (1997).

The literature on modeling and forecasting stock return volatility is voluminous. Two approaches that have proven useful are the GARCH and stochastic volatility (SV) models<sup>1</sup>. In their standard forms, the ensuing volatility processes are stationary and weakly dependent with autocorrelations that decrease exponentially. This contrasts with the empirical findings obtained using various proxies for volatility (e.g., daily absolute returns) which indicate autocorrelations that decay very slowly at long lags. In light of this, several long-memory models have been proposed. For example, Baillie, Bollerslev and Mikkelsen (1996) and Bollerslev and Mikkelsen (1996) considered fractionally integrated GARCH and EGARCH models, while Breidt, Crato and De Lima (1998) and Harvey (1998) proposed long memory SV (LSV) models where the log of volatility is modelled as a fractionally integrated process. A related literature explores semiparametric estimation of the memory parameter;

---

<sup>1</sup>For extensive reviews and collected works, see Engle (1995) and Shephard (2005).

e.g., Arteche (2004) and Deo and Hurvich (2001).

More recently, attempts have been made to distinguish between stationary noise plus level shift and long-memory models; see, in particular, Granger and Hyung (2004). They documented the fact that, when breaks determined via some pre tests are accounted for, the evidence for long-memory is weaker. This evidence is, however, inconclusive since structural change tests are severely biased in the presence of long-memory and log periodogram estimates of the memory parameter are biased downward when sample-selected breaks are introduced. This is an overfitting problem that Granger and Hyung (2004, p. 416) clearly recognized. Starić and Granger (2005) presented evidence that log-absolute returns of the S&P 500 index is an i.i.d. series affected by occasional shift in the unconditional variance and show that this specification has better forecasting performance than the more traditional GARCH(1,1) model and its fractionally integrated counterpart. Mikosch and Starica (2005) considered the autocorrelation function of the absolute returns of the S&P 500 index for the period 1953-1977. They documented the fact that for the full period, it resembles that of a long-memory process. But, interestingly, if one omits the last four years of data, the autocorrelation function is very different and looks like one associated with a short-memory process. They explain this finding by arguing that the volatility of the S&P 500 returns has increased over the period 1973-1977. Morana and Beltratti (2004) also argue that breaks in the level of volatility partially explain the long-memory features of some exchange rate series. Perron and Qu (2007) analyzed the time and spectral domain properties of a stationary short memory process affected by random level shifts. They showed that the shape of the autocorrelation function is different from that obtained from a long-memory process when a wide range of lags is considered and that the limiting distribution of the log periodogram estimate changes as the number of frequencies used varies. Perron and Qu (2008) showed that, when applied to daily S&P 500 absolute returns, their square roots and log absolute returns over the period 1928-2002, the level shift model explains both the shape of the autocorrelations and the path of log periodogram estimates as a function of the number of frequency ordinates used. Qu and Perron (2008) estimated a stochastic volatility model with level shifts using a Bayesian approach. They applied their method to daily data on returns from the S&P 500 and NASDAQ indices over the period 1980.1-2005.12. They showed that the level shifts account for most of the variation in volatility, that their model provides a better in-sample fit than alternative models and that its forecasting performance is better for the NASDAQ and just as good for the S&P 500 as standard short or long-memory models without level shifts.

Our approach extends the work of Starica and Granger (2005) by directly estimating a structural model. We adopt a specification for which the series of interest is the sum of a short-memory process and a jump or level shift component. For the latter component, we specify the commonly used simple mixture model such that the component is the cumulative sum of a process which is 0 with some probability  $(1 - \alpha)$  and is a random variable with probability  $\alpha$ . To estimate such a model, we transform it into a linear state space form with innovations having a mixture of two normal distributions. This transformation allows a wide range of estimation methods. In particular, for this paper, we adopt an algorithm similar to the one by Wada and Perron (2006), which is based on an extension of the Kalman filter applied to four possible states. We restrict the variance of one of the two normal distributions to be zero, allowing us to achieve a simple but efficient algorithm.

We apply this random level shift model to the logarithm of absolute returns. We do so for the following stock market return indices: S&P 500 (1962/07/03 to 2004/03/25; 10504 observations), AMEX (1962/07/03 to 2006/12/31; 11201 observations), Dow Jones (1957/03/04 to 2002/10/30; 11534 observations) and NASDAQ (1972/12/15 to 2006/12/31; 8592 observations). Our point estimates imply few level shifts for all series. But once these are taken into account, there is little evidence of serial correlation in the remaining noise and, hence, no evidence of long-memory. Furthermore, once the estimated shifts are introduced to a standard GARCH model, any evidence of GARCH effects disappears. We also produce recursive out-of-sample forecasts of squared returns. In most cases, our simple random level shifts model clearly outperforms a standard GARCH(1,1) model and, in many cases, it also provides better forecasts than a fractionally integrated GARCH model.

A few comments on our modeling and forecasting strategy are in order. A main goal is to provide forecasts of volatility proxied by daily squared returns. We however, apply our level shift model to log-absolute returns. The main reason is that they do not suffer from a non-negativity constraint as do, say, absolute or squared returns. This is important given that our level shift model does not have a built-in mechanism to bound the modeled series above zero. Also, the evidence for long-memory is just as strong for log-absolute returns as it is for any other transformation of returns, including squared returns (see Granger and Ding, 2005). There is also no loss relative to using squared returns in identifying level shift since log-absolute returns are a monotonic transformation. Since we wish to identify the probability of shifts and their locations, the fact that log-absolute returns are quite noisy is not problematic since our methods are robust to the presence of noise. Another reason is the fact that for many asset returns, a log-absolute transformation yields a series that

is closer to being normally distributed (see, e.g., Andersen, Bollerslev, Diebold and Labys, 2001). A second comment of interest is the fact we use daily returns as opposed to realized volatility series constructed from intra-daily high frequency data which has recently become popular; see, e.g., Andersen and Bollerslev (1997) for the DM-US\$ exchange rate, Bollerslev and Wright (2000) for the Yen-US\$ exchange rate, and Andersen, Bollerslev, Diebold and Labys (2003) for currency spot rates, among many others. While realized volatility series are a less noisy measure of volatility than daily squared returns, their use in our context is problematic. First and foremost, such series are typically available for a short time span. Given the fact that level shifts will turn out to be relatively rare, it is imperative to have a long span of data available to obtain reliable estimates of the probability of occurrence of level shifts. Secondly, such series are available only for specific assets, as opposed to market indices. In our framework, the intent of the level shift model is to have a framework which allows for special events affecting overall markets. Using data on a specific asset would confound such market-wide events with idiosyncratic ones associated with the particular asset used. Third, while the theory about the higher efficiency of realized volatility is well established, the practical applications are subject to several problems. Indeed, the results are very sensitive to a variety of factors: the window used to construct the daily realized volatility series (e.g., 5 or 10 minutes intervals), the treatment of seasonal effects, the precise method to handle missing observations, etc. The use of daily returns avoids any elements of arbitrariness in the exact construction of realized volatility series. Finally, we wish to re-evaluate the adequacy of GARCH models applied to daily returns when taking into account the possibility of level shifts. Hence, it is important to have estimates of these level shifts for squared daily returns which are equivalent to those obtained using log-absolute returns.

The structure of the paper is as follows. Section 2 presents the model and the specifications adopted. Section 3 discusses the estimation procedure while the estimation results and various diagnostics are presented in Section 4. The forecasting comparisons between our model and GARCH and FIGARCH models are reported in Section 5. Section 6 offers brief conclusions.

## 2 The model

The model we apply to log-absolute returns is given by

$$y_t = a + \epsilon_t + \zeta_t, \quad (1)$$

where  $a$  is a constant,  $\tau_t$  is the random level shift components and  $c_t$  is a short-memory process, included to model the remaining noise. The level shift component is specified by

$$\tau_t = \tau_{t-1} + \epsilon_t,$$

where

$$\tau_t = \tau_t \epsilon_t.$$

Here,  $\tau_t$  is a binomial variable that takes value 1 with probability  $\alpha$  and value 0 with probability  $(1 - \alpha)$ . If it takes value 1, then a random level shift  $\epsilon_t$  occurs, specified by  $\epsilon_t \sim \text{i.i.d. } N(0, \sigma^2)$ .

In its most general form, the short-memory process  $c_t$  is defined by  $c_t = C(L) \epsilon_t$ , with  $\epsilon_t \sim \text{i.i.d. } N(0, \sigma^2)$  and  $E|\epsilon_t|^r < \infty$  for some  $r > 2$ . The polynomial  $C(L)$  satisfies  $C(L) = \sum_{i=0}^{\infty} c_i L^i$ ,  $\sum_{i=0}^{\infty} i |c_i| < \infty$  and  $C(1) \neq 0$ . These conditions allow us to approximate the process by a finite order autoregression, provided its order is chosen suitably. As we shall see, once the level shifts are accounted for, very little autocorrelation remains in the process so that an AR(1) specification of the short-memory component is sufficient. Even then, the autoregressive parameter estimate will be small and insignificant. Hence, in what follows, we shall for simplicity assume that

$$c_t = c_{t-1} + \theta_t,$$

understanding that our method of estimation can easily be extended to cover more general processes with straightforward modifications. We also assume that the components  $\tau_t$ ,  $\epsilon_t$  and  $c_t$  are mutually independent. The normality assumption for  $\epsilon_t$  is needed to construct the likelihood function, though the parameter estimates remain consistent without it.

It is easy to confuse this model with the popular Markov regime switching model (e.g., see Hamilton 1989) since  $\tau_t$  can be interpreted as a latent regime variable. The two models share the feature of allowing the time series to follow different processes over different subsamples. But there is a key distinction that allows our model to have more flexibility. In the standard Markov regime switching model, only a finite number of possible regimes, usually two or some other small value, are allowed (see, for example, Hamilton 1989, Gray 1996 and Filardo and Gordon 1998). In contrast, our model does not restrict the magnitude of the level shifts since they are draws from a normal distribution. Hence, any number of regimes is possible. It is also important to note that the probability that  $\tau_t$  be 0 or 1 is independent of past realizations, unlike the Markov switching type models. Here, the different regimes are associated with different magnitudes of the shifts. Our goal is to develop a framework which

allows for special events. Hence, the probability that a shift occurs should be independent of past shifts. Accordingly, it is appropriate to make the probability of being in a given regime as independent of past realizations.

In order to estimate the model, we shall embed it in a state space framework involving errors that have a mixture of two normal distributions. The level shift component  $\mu_t$  can be specified as a random walk process with innovations distributed according to a mixture of two normally distributed processes:

$$\begin{aligned}\mu_t &= \mu_{t-1} + \nu_t \\ \nu_t &= \nu_{1t} + (1 - \nu_{1t}) \nu_{2t},\end{aligned}$$

where  $\nu_{it}$  i.i.d.  $N(0, \sigma_i^2)$  and  $\nu_{1t}$  is a Bernoulli random variable that takes value one with probability  $\pi$  and value 0 with probability  $1 - \pi$ . By specifying  $\sigma_1^2 = \sigma^2$  and  $\sigma_2^2 = 0$ , we recover our level shift model.

We next specify the model in terms of first-differences of the data:

$$\begin{aligned}\Delta y_t &= \mu_t - \mu_{t-1} + C_t - C_{t-1} \\ &= \nu_t + C_t - C_{t-1},\end{aligned}$$

where

$$\nu_t = \nu_{1t} + (1 - \nu_{1t}) \nu_{2t}.$$

We then have the following state space form:

$$\begin{aligned}\Delta y_t &= C_t - C_{t-1} + \nu_t \\ C_t &= C_{t-1} + \epsilon_t,\end{aligned}$$

or more generally,

$$\begin{aligned}\Delta y_t &= H X_t + \nu_t \\ X_t &= F X_{t-1} + U_t,\end{aligned}$$

where, in the case of an AR(p) process,

$$X_t = [C_t, C_{t-1}, \dots, C_{t-p+1}]',$$

$$F = \begin{bmatrix} 1 & 2 & \cdots & p \\ 1 & & & \\ & 1 & & \\ & & \ddots & \\ & & & 1 \end{bmatrix},$$

$H = [1, -1, 0, \dots, 0]$  and  $U_t$  is a  $p$ -dimensional normally distributed random vector with mean zero and covariance matrix

$$Q = \begin{pmatrix} \sigma_e^2 & 0_{1 \times (p-1)} \\ 0_{(p-1) \times 1} & 0_{(p-1) \times (p-1)} \end{pmatrix}.$$

The important difference between this model and a standard state space model is that the distribution of  $\epsilon_t$  is a mixture of two normal distributions with variances  $\sigma_e^2$  and 0, occurring with probabilities  $\alpha$  and  $1 - \alpha$ , respectively. Note that it is also easy to extend our model to have the short-memory component  $c_t$  follow an ARMA process as long as  $c_t$  can be represented in a state space form.

### 3 The estimation method

The model described in the previous section is a special case of the class of models considered by Wada and Perron (2006). They modeled the trend-cycle decomposition of a macroeconomic time series allowing for mixtures of normal distributions for the shocks affecting the level, slope and cyclical components. Here, we only have shocks affecting the level of the series and we simply need to impose the restriction that the variance of one component of the mixture of normal distributions is zero. In what follows, we provide a brief description of the method, more details can be found in Wada and Perron (2006).

The first thing to note is that despite their fundamental differences, our model and the Markov regime switching models share similarities in estimation methodology. The basic ingredient for estimation is the augmentation of the states by the realizations of the mixture at time  $t$  so that the Kalman filter can be used to generate the likelihood function, conditional on the realizations of the states. The latent states are eliminated from the final likelihood expression by summing over all possible state realizations. As we shall show, our model takes a structure that is similar to a version of the Markov regime switching model.

Let  $Y_t = (\Delta y_1, \dots, \Delta y_t)$  be the vector of data available up to time  $t$  and denote the vector of parameters by  $\theta = [\sigma_e^2, \alpha, \sigma_e^2, \beta_1, \dots, \beta_q]$ . To illustrate the similarities we adopt

the notations in Hamilton (1994), where  $\mathbf{1}$  represents a  $(4 \times 1)$  vector of ones, the symbol  $\odot$  denotes element-by-element multiplication,  $\hat{\omega}_{t|t-1} = \text{vec}(\tilde{\omega}_{t|t-1})$  with the  $(i, j)^{\text{th}}$  element of  $\tilde{\omega}_{t|t-1}$  being  $\Pr(S_{t-1} = i, S_t = j | Y_{t-1}; \cdot)$  and  $\omega_t = \text{vec}(\tilde{\omega}_t)$  with the  $(i, j)^{\text{th}}$  element of  $\tilde{\omega}_t$  being  $f(\Delta y_t | S_{t-1} = i, S_t = j, Y_{t-1}; \cdot)$  for  $i, j \in \{1, 2\}$ . Here  $s_t = 1$  (resp., 0) when  $\Delta y_t = 1$  (resp., 0), i.e., a level shift occurs (resp., does not occur). The log likelihood function is

$$\ln(L) = \sum_{t=1}^T \ln f(\Delta y_t | Y_{t-1}; \cdot), \quad (2)$$

where

$$f(\Delta y_t | Y_{t-1}; \cdot) = \sum_{i=1}^2 \sum_{j=1}^2 f(\Delta y_t | S_{t-1} = i, S_t = j, Y_{t-1}; \cdot) \Pr(S_{t-1} = i, S_t = j | Y_{t-1}; \cdot) \mathbf{1}' \left( \hat{\omega}_{t|t-1} \odot \omega_t \right). \quad (3)$$

We first focus on the evolution of  $\hat{\omega}_{t|t-1}$ . Applying rules for conditional probabilities, Bayes' rule and the independence of  $s_t$  with past realizations, we have

$$\begin{aligned} \tilde{\omega}_{t|t-1}^{ij} &= \Pr(S_{t-1} = i, S_t = j | Y_{t-1}; \cdot) \\ &= \Pr(S_t = j) \sum_{k=1}^2 \Pr(S_{t-2} = k, S_{t-1} = i | Y_{t-1}; \cdot), \end{aligned}$$

and

$$\begin{aligned} \tilde{\omega}_{t-1|t-1}^{ki} &= \Pr(S_{t-2} = k, S_{t-1} = i | Y_{t-1}; \cdot) \\ &= \frac{f(\Delta y_{t-1} | S_{t-2} = k, S_{t-1} = i, Y_{t-2}; \cdot) \Pr(S_{t-2} = k, S_{t-1} = i | Y_{t-2}; \cdot)}{f(\Delta y_{t-1} | Y_{t-2}; \cdot)}. \end{aligned}$$

Therefore, the evolution of  $\hat{\omega}_{t|t-1}$  is given by:

$$\begin{bmatrix} \tilde{\omega}_{t+1|t}^{11} \\ \tilde{\omega}_{t+1|t}^{21} \\ \tilde{\omega}_{t+1|t}^{12} \\ \tilde{\omega}_{t+1|t}^{22} \end{bmatrix} = \begin{bmatrix} \left( \tilde{\omega}_{t|t}^{11} + \tilde{\omega}_{t|t}^{21} \right) \\ \left( \tilde{\omega}_{t|t}^{12} + \tilde{\omega}_{t|t}^{22} \right) \\ (1 - \cdot) \left( \tilde{\omega}_{t|t}^{11} + \tilde{\omega}_{t|t}^{21} \right) \\ (1 - \cdot) \left( \tilde{\omega}_{t|t}^{12} + \tilde{\omega}_{t|t}^{22} \right) \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ (1 - \cdot) & (1 - \cdot) \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \tilde{\omega}_{t|t}^{11} \\ \tilde{\omega}_{t|t}^{21} \\ \tilde{\omega}_{t|t}^{12} \\ \tilde{\omega}_{t|t}^{22} \end{bmatrix}.$$

or more compactly by:

$$\hat{\omega}_{t+1|t} = \Pi \hat{\omega}_{t|t}, \quad (4)$$

with

$$\hat{\omega}_{t|t} = \frac{\hat{\omega}_{t|t-1} \odot \omega_t}{\mathbf{1}'(\hat{\omega}_{t|t-1} \odot \omega_t)}.$$

The conditional likelihood for  $\Delta y_t$  is the following Normal density:

$$\tilde{\omega}_t^{ij} = f(\Delta y_t | S_{t-1} = i, S_t = j, Y_{t-1}; ) = \frac{1}{2} |f_t^{ij}|^{-\frac{1}{2}} \exp \left\{ -\frac{ij' (f_t^{ij})^{-1} ij}{2} \right\}, \quad (5)$$

where  $ij = \Delta y_t - \Delta y_{t|t-1}^i$  is the prediction error and  $f_t^{ij} = E ( ij' ij )$  is the prediction error variance. Note that  $\Delta y_{t|t-1}^i = E [\Delta y_t | S_{t-1} = i, Y_{t-1}; ]$  does not depend on the state  $j$  at time  $t$  because we are conditioning on time  $t-1$  information. However,  $\Delta y_t$  does depend on  $s_t = j$  so that the prediction error and its variance depend on both  $i$  and  $j$ . The best forecast for the state variable and its associated variance conditional on past information and  $s_{t-1} = i$  are

$$\begin{aligned} X_{t|t-1}^i &= F X_{t-1|t-1}^i, \\ P_{t|t-1}^i &= F P_{t-1|t-1}^i F' + Q. \end{aligned} \quad (6)$$

We have the measurement equation  $\Delta y_t = H X_t + \epsilon_t$ , where the measurement error  $\epsilon_t$  has mean zero and a variance which can take two possible values:  $R_1 = \sigma^2$ , with probability  $\alpha$ , or  $R_2 = 0$ , with probability  $1 - \alpha$ . Hence, the prediction error is  $ij = \Delta y_t - H X_{t|t-1}^i$  with associated variance  $f_t^{ij} = H P_{t|t-1}^i H' + R_j$ . Applying standard updating formulae, we have given  $s_t = j$  and  $s_{t-1} = i$ ,

$$\begin{aligned} X_{t|t}^{ij} &= X_{t|t-1}^i + P_{t|t-1}^i H' (H P_{t|t-1}^i H' + R_j)^{-1} (\Delta y_t - H X_{t|t-1}^i), \\ P_{t|t}^{ij} &= P_{t|t-1}^i - P_{t|t-1}^i H' (H P_{t|t-1}^i H' + R_j)^{-1} H P_{t|t-1}^i. \end{aligned} \quad (7)$$

To reduce the dimension of the estimation problem, we adopt the re-collapsing procedure suggested by Harrison and Stevens (1976), given by

$$\begin{aligned} X_{t|t}^j &= \frac{\sum_{i=1}^2 \Pr(S_{t-1} = i, S_t = j | Y_t; ) X_{t|t}^{ij}}{\Pr(S_t = j | Y_t; )} = \frac{\sum_{i=1}^2 \tilde{ij} X_{t|t}^{ij}}{\sum_{i=1}^2 \tilde{ij}}, \\ P_{t|t}^j &= \frac{\sum_{i=1}^2 \Pr(S_{t-1} = i, S_t = j | Y_t; ) \left[ P_{t|t}^{ij} + (X_{t|t}^j - X_{t|t}^{ij}) (X_{t|t}^j - X_{t|t}^{ij})' \right]}{\Pr(S_t = j | Y_t; )} \\ &= \frac{\sum_{i=1}^2 \tilde{ij} \left[ P_{t|t}^{ij} + (X_{t|t}^j - X_{t|t}^{ij}) (X_{t|t}^j - X_{t|t}^{ij})' \right]}{\sum_{i=1}^2 \tilde{ij}}. \end{aligned} \quad (8)$$

By doing so, we make  $\omega_t^{ij}$  unaffected by the history of states before time  $t - 1$ .

If we define  $S_t = (s_t, s_{t-1})$ , we then have four possible states corresponding to  $S_t = 1$  when  $(s_t = 1, s_{t-1} = 1)$ ,  $S_t = 2$  when  $(s_t = 1, s_{t-1} = 2)$ ,  $S_t = 3$  when  $(s_t = 2, s_{t-1} = 1)$  and  $S_t = 4$  when  $(s_t = 2, s_{t-1} = 2)$  with the transition matrix  $\Pi$  as defined in (4). The vector of conditional densities  $\omega_t = (\omega_t^1, \dots, \omega_t^4)$  thus have a more compact representation given by:

$$\omega_t^t = f(\Delta y_t | S_t = t, Y_{t-1}; \theta) = \frac{1}{2} |f_t^t|^{-\frac{1}{2}} \exp \left\{ -\frac{t' (f_t^t)^{-1} t}{2} \right\} \quad (9)$$

where  $t$  and  $f_t^t$  are as defined in (5) with the values of  $s_t$  and  $s_{t-1}$  corresponding to  $S_t = t$ . This definition of  $\omega_t$ , together with  $\hat{\pi}_{t|t}$  and  $\hat{\pi}_{t+1|t}$ , the collection of conditional probabilities  $\Pr(S_t = t | Y_t; \theta)$  for  $t = 1, \dots, 4$  and its the one-period ahead forecasts, evolving as in (4), give us the same structure as a version of the Markov regime switching model used by Hamilton (1994). However, there are two extra complexities here. Firstly, the mean and variance in the conditional density function are nonlinear functions of the fundamental parameters and past realizations  $\{\Delta y_{t-j}; j = 1\}$ . This non-linearity and time dependence complicate the maximization of the log-likelihood function since we cannot separate out some elements of  $\theta$  in the first order conditions. Accordingly, the standard EM algorithm does not apply. Secondly, the conditional probability of being in a given regime  $\hat{\pi}_{t|t}$  is not separable from the conditional densities  $\omega_t$  since  $\hat{\pi}_{t|t}^{ij}$  enters in its construction (see (8)).

#### 4 Empirical results for returns on stock market indices

In this section, we apply our model and estimation method to the returns of four major market indices: the S&P 500, AMEX, Dow Jones and NASDAQ. The daily returns are computed by first differencing the logarithm of the index price series  $r_t = \ln(P_t) - \ln(P_{t-1})$ . The data coverage is from 1962/07/03 to 2004/03/25 for the S&P 500 (10504 observations), from 1962/07/03 to 2006/12/31 for the AMEX (11201 observations), from 1957/03/04 to 2002/10/30 for the Dow Jones (11534 observations) and from 1972/12/15 to 2006/12/31 for the NASDAQ (8592 observations). For reasons stated in the introduction, we model log absolute returns. When returns are zero or close to it, the log absolute value transformation implies extreme negative values. Using our method, these outliers would be attributed to the level shift component and thus bias the probability of shifts upward. To avoid this problem, we bound absolute returns away from zero by adding a constant, i.e., we use  $\log(|r_t| + 0.001)$ .

To specify the short-memory component  $c_t$ , we start by setting  $c_t = e_t$ , following Stărica and Granger (2005) who report that the short-memory component in such series is just white

noise. As a robustness check, we also report the estimates for the specification that  $c_t$  follows an AR(1) process,  $c_t = \rho c_{t-1} + e_t$ . Since all components of the state vector are stationary, we can initialize the state vector and its covariance matrix by their unconditional expected values, i.e.,  $X_{0|0} = (0, 0)'$  and

$$P_{0|0} = \begin{bmatrix} \frac{2}{e} & 0 \\ 0 & 0 \end{bmatrix}.$$

We obtain estimates by directly maximizing the likelihood function (2). In order to avoid the problem of local maxima, we re-estimate with different initial values of  $\rho_0$  and pick the estimates associated with the largest likelihood value upon convergence.

#### 4.1 Estimation results

The parameter estimates are presented in Table 1 for both the cases in which the short-memory component  $c_t$  is specified to be white noise and an AR(1) process. To assess the relative importance of the two components to the total variation of the series, we also report the standard deviation of the original series  $\log|r_t|$ . The estimates reported are the standard deviation of level shift component  $\sigma_\mu$ , the probability of a shift  $\rho_\mu$ , the standard deviation of the stationary component  $\sigma_e$  and the autoregressive coefficient  $\rho$  when considering the AR(1) specification for  $c_t$ .

These results exhibit noteworthy features. First, when considering the AR(1) specification for  $c_t$ , the estimate of  $\rho$  is very small and close to zero, except for the AMEX series for which it is the largest but still only 0.063. In all cases, adopting either an AR(1) or white noise specification for  $c_t$  yields very similar results (there are some slight differences for the AMEX series but the subsequent results to be discussed here are unaffected by the specification of  $c_t$ ). Thus in subsequent sections we shall only consider results based on the white noise specification for the short-memory component. These results are in agreement with those of Stărică and Granger (2005).

The second noteworthy feature is the fact that the probability of level shifts is very small in all cases considered. Indeed, the point estimates of  $\rho_\mu$  imply the following number of shifts for each index: 15 for the S&P 500, 28 for the AMEX, 12 for the Dow Jones and 7 for the NASDAQ. However, the shifts are important given that their standard deviation is of the same order as the standard deviation of the series. Nevertheless, given that the shifts occur so infrequently, the noise component accounts for the bulk of total variation.

## 4.2 The effect of level shifts on long-memory and conditional heteroskedasticity

Given the features documented above, it would be interesting to determine whether the level shifts are important to dictating the overall behavior of the series. We shall address this issue by investigating whether the shifts can explain a) the well-documented feature of long-memory and b) the presence of conditional heteroskedasticity.

To do so we need estimates of the dates at which the level shifts occur as well as the means within each segment. Ideally, one would use a smoothed estimate of the level shift component  $\mu_t$ . However, in this context of multiple abrupt changes, conventional smoothers perform very poorly (at least we are not aware of one such smoother that would yield reliable results). Hence, we resort to the use of the following strategy. We use the point estimate of  $\mu_t$  to obtain a point estimate of the number of changes, as reported above. We then apply the method of Bai and Perron (2003) to obtain the estimates of the break dates that globally minimize the following sum of squared residuals:

$$\sum_{i=1}^{m+1} \sum_{t=T_{i-1}+1}^{T_i} [y_t - \mu_i]^2,$$

where  $m$  is the number of breaks,  $T_i$  ( $i = 1, \dots, m$ ) are the break dates with  $T_0 = 0$  and  $T_{m+1} = T$  and  $\mu_i$  ( $i = 1, \dots, m + 1$ ) are the means within each regime which can easily be estimated once the break dates are. This method is efficient and can easily handle the large number of observations we are using, even with as many as 28 breaks (see Bai and Perron, 2003 for more details). Note that since our model allows for consecutive level shifts, we set the minimal length of a segment to just one observation. The fitted level shift component is presented in Figure 1 along with the original series  $\log|r_t|$ . A look at the graph indicates that the shifts agree with the major changes in the scale of  $\log|r_t|$  and, as expected, there is a brief but important shift in October 1987 for all series.

The first issue we address is whether the level shift can account for the long-memory feature of the series  $\log|r_t|$ . To do so, we plot the autocorrelation functions (up to 300 lags) of the original series  $\log|r_t|$  and its short-memory component  $c_t$ , obtained by subtracting the fitted trend from  $\log|r_t|$ . The results are presented in Figure 2. One can see that for all cases, the log-absolute returns clearly display an autocorrelation function that resembles that of a long-memory process: it decays very slowly and the values remain important even at lag 300. On the other hand, once the level shifts are accounted for, the picture is completely different and there is barely any evidence of remaining correlation, consistent with the estimate of  $d$  obtained above. Given that we have ten thousands or so observations

for each series, the autocorrelations at small lags may be deemed significant but they have no practical importance given their small values. Hence, for all practical purposes, we can view the short-memory component as being nearly white noise. These results are quite impressive. Despite the fact that the shifts are rare and account for only a small portion of the total variation of the series, they account for (nearly) all of the autocorrelation present in log-absolute returns.

It is almost universally accepted that stock returns exhibit conditional heteroskedasticity. For that reason, the GARCH(1,1) model, introduced by Bollerslev (1986) following the work of Engle (1982), has been used widely to model returns and is generally perceived as one of the leading candidates of forecasting models. Though Lamoureux and Lastrapes (1990) documented that structural changes in the level of variance can amplify the evidence of conditional heteroskedasticity, no strong evidence has been previously presented on assessing whether accounting for such changes can completely eliminate all evidence of conditional heteroskedasticity. We shall present evidence to this effect.

To account for the fat-tailed distribution of returns, the most popular model is the GARCH(1,1) model with Student-t errors given by, for the demeaned returns process  $\tilde{r}_t$ ,

$$\begin{aligned}\tilde{r}_t &= \epsilon_t \nu_t, \\ \nu_t^2 &= \mu + \alpha \tilde{r}_{t-1}^2 + \beta \nu_{t-1}^2,\end{aligned}\tag{10}$$

where  $\epsilon_t$  is i.i.d Student-t distributed with mean 0 and variance 1. The parameters of interest are  $\alpha$  and  $\beta$ , which measure the extent of conditional heteroskedasticity present in the data. As is standard, we let the degrees of freedom of the Student-t distribution be an unknown parameter, estimated jointly with the others. We estimate such a model for the four returns indices using maximum likelihood. We also consider the case that accounts for level shifts at the dates documented in the fitted level shift component, as presented of Figure 1. This implies the following extended specification, sometimes referred to as a components GARCH model:

$$\begin{aligned}\tilde{r}_t &= \epsilon_t \nu_t, \\ \nu_t^2 &= \mu + \alpha (\nu_{t-1}^2 - \mu) + (\tilde{r}_{t-1}^2 - \nu_{t-1}^2) + \sum_{i=2}^{m+1} D_{i,t} \nu_i, \\ (\nu_t^2 - \nu_t) &= \alpha (\nu_{t-1}^2 - \nu_{t-1}) + (\tilde{r}_{t-1}^2 - \nu_{t-1}),\end{aligned}\tag{11}$$

where  $D_{i,t} = 1$  if  $t$  is in regime  $i$ , i.e.,  $t \in \{T_{i-1} + 1, \dots, T_i\}$ , and 0 otherwise, with  $T_i$  ( $i = 1, \dots, m$ ) being the break dates documented in Figure 1 (again  $T_0 = 0$  and  $T_{m+1} = T$ ).

The coefficients  $\beta_i$ , which index the magnitude of the shifts, are treated as unknown and are estimated with the remaining parameters, while the number of breaks is obtained from the point estimate of  $\beta$ . The estimates of  $\alpha_r$  and  $\beta$  obtained from the standard GARCH(1,1) model (10) and the extended model (11) that allows for level shifts are presented in Table 2. The results are quite informative. Using the standard GARCH(1,1) model, both estimates are highly significant for all series. In particular, the value of  $\beta$  is quite high, as often documented, with values ranging between 0.866 and 0.935. We also estimated the components GARCH model without level shifts and the results are similar for the estimates of  $\alpha_r$  and  $\beta$ . The estimates of  $\beta$  are very close to one. This is due to the fact that the model attempts to capture features akin to long memory so that a value of  $\beta$  close to one implies highly persistent shocks.

The picture is very different when we allow for level shifts. None of the estimates of  $\beta$  are significant, with small values ranging from -0.045 to 0.151. The estimate of  $\alpha_r$  is insignificant for the S&P 500 and Dow Jones series. It is significant with p-values close to 0.01 for the AMEX and NASDAQ series, though the point estimates are much smaller (0.05 for both) so that their economic importance is minor. Also of interest is the fact that with level shifts the estimates of  $\beta$  are now well below one, between 0.85 and 0.95, implying that shocks have a half life between between 4 and 13 days. Hence, introducing few level shifts imply a markedly different interpretation of the data. These results indicate that level shifts in log-absolute or squared returns account for nearly all of the documented evidence of conditional heteroskedasticity in these stock returns series. They also imply that shock to volatility have very little persistence.

To summarize the results so far, the level shifts model with white noise errors appears to provide an accurate description of the data. The level shift component is an important feature that explains both the long-memory and conditional heteroskedasticity features generally perceived as stylized facts. Our analysis so far is retrospective and one may rightfully raise the issue of data overfitting as an explanation of the results, especially when multiple structural changes are present. It remains therefore to see whether our level shift model provides reasonable forecasts in comparison to traditional models.

## 5 Forecasting

We now consider the performance of the level shift model with white noise errors for log-absolute returns in forecasting volatility proxied by squared returns relative to two popular models: the GARCH(1,1) model and its fractionally integrated counterpart, the FI-

GARCH(1,1) model. The design of the forecasting experiment follows Stărică and Granger (2005). We start forecasting at observations 2,000. We re-estimate the models every 20 days, at which point forecasts of up to 200 days are constructed. The proxy for volatility being the realized squared returns are quite noisy, and to reduce the effect of sampling variability, we follow Stărică and Granger (2005) in the construction of a metric to gauge relative performance. Let  $\hat{\sigma}_{t,p}^2$  be a p-step ahead forecast of  $\sigma_{t+p}^2$ , the variance of returns  $r_t$  at time  $t+p$ , proxied by the squared demeaned returns. Let  $n$  be the number of forecasts produced, then the estimated MSE is constructed as

$$\text{MSE}(p) = \frac{1}{n} \sum_{t=1}^n (r_{t,p}^2 - \hat{\sigma}_{t,p}^2)^2, \quad (12)$$

where  $\hat{\sigma}_{t,p}^2 = \sum_{k=1}^p \hat{\sigma}_{t+k}^2$  and  $r_{t,p}^2 = \sum_{k=1}^p r_{t+k}^2$  is the realized volatility over the interval  $[t+1, t+p]$ . This estimate of the MSE is preferable to the simpler version  $\sum_{t=1}^n (r_{t+p}^2 - \hat{\sigma}_{t+p}^2)^2$ , because the latter uses a poor measure of realized return volatility (see Anderson and Bollerslev, 1998). Through averaging, some of the idiosyncratic noise in the high-frequency component of squared returns is canceled out. Throughout, the relative forecasting performance of the two models is evaluated by the ratio of their MSEs as defined above.

## 5.1 Construction of the forecasts

We now describe how the forecasts are constructed for each method. We start with the level shift model applied to log-absolute returns for which we assume, following the evidence documented above, a short-memory component that is simply white noise. Since our model and estimates pertain to log-absolute returns, we need an appropriate transformation that yields forecasts of the variance of returns. Recall that our model specifies the following process for log absolute returns <sup>2</sup>:

$$\log(|r_t| + C) = \mu_t + \epsilon_t,$$

where  $C$  is a positive constant used to bound returns away from zero. This yields the following model for returns:

$$|r_t| + C = h_t^{1/2} \tilde{\epsilon}_t,$$

where  $h_t = e^{2\mu_t} E(e^{2\epsilon_t})$  and  $\tilde{\epsilon}_t = e^{\epsilon_t} / [E(e^{2\epsilon_t})]^{1/2}$  so that  $E(\tilde{\epsilon}_t^2) = 1$  and  $\tilde{\epsilon}_t$  is independent of  $h_t$ . The sequence  $\tilde{\epsilon}_t$  is also i.i.d., given that  $\epsilon_t = e_t$ . Hence,  $2\epsilon_t \sim$  i.i.d.  $N(0, 4\frac{\sigma^2}{e})$  and

---

<sup>2</sup>Note that the constant term present in (1) does not enter in the estimation procedure. The unconditional mean of  $\log(|r_t| + C)$  can be captured by constructing  $\mu_t$  appropriately.

$E(e^{2c_t}) = e^{0+4\sigma^2/2} = e^{2\sigma^2}$ . Since the level shifts are rare and there is considerable uncertainty about their timing and magnitudes, we ignore them when forecasting. We then have,

$$E_t(|r_{t+k}| + C)^2 = E_t h_{t+k} = \exp(2\sigma^2 + 2C),$$

so that the k-period ahead forecast of the squared returns  $r_{t+k}^2$  is

$$E_t r_{t+k}^2 = \exp(2\sigma^2 + 2C) - 2CE_t|r_{t+k}| - C^2,$$

where

$$E_t|r_{t+k}| = E_t \exp(\sigma^2 + C_{t+k}) - C = \exp(\sigma^2 + 0.5\sigma^2) - C.$$

We now consider forecasting with the Student-t GARCH(1,1) model, which can be written as:

$$\tilde{r}_t = r_t - \mu = \epsilon_t \epsilon_t, \quad (13)$$

$$\epsilon_t^2 = \alpha_1 + \alpha_2 \tilde{r}_{t-1}^2 + \alpha_3 \epsilon_{t-1}^2, \quad (14)$$

where the innovation  $\epsilon_t$  is i.i.d Student-t distributed with mean 0 and variance 1 and  $\tilde{r}_t$  are demeaned returns. We start with the following transformation:

$$\tilde{r}_t^2 = \alpha_1 + (\alpha_2 + \alpha_3) \tilde{r}_{t-1}^2 + \epsilon_t - \alpha_3 \epsilon_{t-1}, \quad (15)$$

where  $\epsilon_t = \tilde{r}_t^2 - \alpha_1$  is the forecast error associated with the forecast of  $\tilde{r}_t^2$  based on its lagged values (to see this, note that  $E_{t-1}(\tilde{r}_t^2) = E_{t-1}(\alpha_1 + \alpha_2 \tilde{r}_{t-1}^2 + \alpha_3 \epsilon_{t-1}^2) = \alpha_1 + \alpha_2 \tilde{r}_{t-1}^2 + \alpha_3 \epsilon_{t-1}^2$ ). Thus,  $\epsilon_t$  is a white noise process that is fundamental to  $\tilde{r}_t^2$ . The unconditional mean of  $\tilde{r}_t^2$  can then be computed easily from (15), assuming  $\alpha_2 + \alpha_3 < 1$ , by

$$E(\tilde{r}_t^2) = \alpha_1 / (1 - \alpha_2 - \alpha_3).$$

The recursive form for the squared demeaned returns is

$$\tilde{r}_t^2 - \alpha_1 = (\alpha_2 + \alpha_3)(\tilde{r}_{t-1}^2 - \alpha_1) + \epsilon_t - \alpha_3 \epsilon_{t-1}$$

so that, for  $k > 1$ , the k-period ahead forecast of the squared demeaned returns is

$$E_t \tilde{r}_{t+k}^2 = \alpha_1 + (\alpha_2 + \alpha_3)^{k-1} (E_t \tilde{r}_{t+1}^2 - \alpha_1) = \alpha_1 + (\alpha_2 + \alpha_3)^{k-1} (\tilde{r}_{t+1}^2 - \alpha_1),$$

where  $\tilde{r}_{t+1}^2 = \alpha_1 + \alpha_2 \tilde{r}_t^2 + \alpha_3 \epsilon_t^2$ . To recover the forecasts for the squared returns, we make the following adjustment, using the fact that the time variation in the conditional mean of returns is quantitatively negligible:

$$E_t r_{t+k}^2 = E_t \tilde{r}_{t+k}^2 + (E_t r_{t+k})^2 = E_t \tilde{r}_{t+k}^2 + \mu^2.$$

We finally consider the FIGARCH(1,1) model, given by

$$\begin{aligned}\tilde{r}_t &= r_t - \mu = \varepsilon_t \sigma_t, \\ (1 - \alpha_3 L) \sigma_t^2 &= \alpha_1 + [1 - \alpha_3 L - (1 - \alpha_2 L)(1 - L)^d] \tilde{r}_t^2,\end{aligned}$$

where the innovation  $\varepsilon_t$  is i.i.d standard normal and  $(1 - L)^d = \sum_{j=0}^{\infty} \psi_j L^j$  with

$$\psi_j = \frac{\Gamma(j + d)}{\Gamma(j + 1)\Gamma(d)} = \prod_{0 < k \leq j} \frac{k - 1 - d}{k}.$$

To facilitate forecasting, we transform the conditional variance equation into an infinite order ARCH representation:

$$\sigma_t^2 = \alpha_1 / (1 - \alpha_3) + (L) \tilde{r}_t^2,$$

where

$$\alpha_1 = \alpha_2 - \alpha_3 + d,$$

and

$$\alpha_k = \alpha_3 \alpha_{k-1} + \left[ \frac{k - 1 - d}{k} - \alpha_2 \right] \alpha_{k-1}, \text{ for } k \geq 2.$$

Then, the recursive forecasts can be constructed from

$$E_t \tilde{r}_{t+k}^2 = E_t \sigma_{t+k}^2 = \alpha_1 / (1 - \alpha_3) + \sum_{i=1}^{t+k-1} \alpha_i E_t \tilde{r}_{t+k-i}^2,$$

where  $E_t \tilde{r}_{t+k-i}^2 = \tilde{r}_{t+k-i}^2$  for  $i \geq k$ . To recover the forecasts of the squared returns, we adjust  $E_t \tilde{r}_{t+k}^2$  by adding back the squared mean, so that  $E_t r_{t+k}^2 = E_t \tilde{r}_{t+k}^2 + \mu^2$ .

## 5.2 Forecasting Comparisons

We start with an unfair but instructive comparison. For the GARCH(1,1) and FIGARCH(1,1) models, we re-estimate every 20 observations. However, for the random level shifts model, we use the fitted means obtained from the full sample, i.e., the fitted level shift process depicted in Figure 1 obtained using the Bai and Perron (2003) algorithm and the number of breaks implied by the full sample estimate of  $\lambda$ , the probability of a level shift at each period. The results are presented in Figure 3 (level shifts versus GARCH) and Figure 4 (level shifts versus FIGARCH). They show that the random level shift model has much better forecasting performance at all horizons and for all series, except for very short horizons when considering the AMEX and NASDAQ indices, in which case the GARCH and FIGARCH have a slight advantage. As stated, this is an unfair comparison but it indicates that if we have a precise

estimate of the mean of log-absolute returns at a given date, we can obtain much better forecasts from the level shift model than the other models.

The issue then becomes how to obtain a good estimate of the current mean of a regime at a given date, at which the forecasts are made, without using information after that date. This turns out to be a delicate issue. The use of the filtered estimates of  $\mu_t$ , the level shift component, obtained via the Kalman filter algorithm are too volatile to be useful. An obvious approach then is to repeat what we did for the full sample every 20 observations and forecast using the estimate of the mean for the last regime. The problem here is that if one allows for the possibility of a change at each date, as in the theoretical model, the fitted values often indicate that a change occurred at the end of the sample when, ex-post, no such long lasting change have occurred. In light of these issues, we resort to using a backwards CUSUM procedure, as in Pesaran and Timmerman (1999). At each forecasting period, we use the CUSUM test of Brown, Durbin and Evans (1975). We then estimate the date of the last break as the first time the CUSUM statistic crosses one of the critical lines, determined by the criterion that the probability of at least one of the last 1,000 cumulative sums of standardized recursive residuals crossing a line is 10%.

The results are presented in Figure 5 (level shifts versus GARCH) and Figure 6 (level shifts versus FIGARCH). For the S&P 500 and Dow Jones indices, the level shift model forecasts better than either the GARCH(1,1) or the FIGARCH(1,1) model at nearly all horizons (except at a few isolated short horizons). For the AMEX and NASDAQ indices, the GARCH(1,1) model performs better only at very short horizons, while the random level shift model is better for longer forecasts. Compared to the FIGARCH(1,1) model, the random level shift model performs better in the case of the AMEX index, except at a few isolated short horizons. For the NASDAQ index, the level shift model performs better at horizons greater than 100 days, but the FIGARCH model performs better at short horizons.

Following Anderson et al. (2003) and Stărică and Granger (2005), in the tradition of Mincer and Zarnowitz (1969), we also evaluate the models' relative forecasting accuracy by regressing proxies for volatility on a constant and the forecast values provided by the various models. We consider regressions of the form

$$r_{t+p}^2 - r_t^2 = b_0 + b_1(f_{t,p}^{LS} - r_t^2) + b_2(f_{t,p}^i - r_t^2) + u_t, \quad (16)$$

where  $f_{t,p}^{LS}$  denotes the p-step ahead forecast of  $r_{t+p}^2$  from the level shift model and  $f_{t,p}^i$  denotes the p-step ahead forecast from the GARCH model ( $i = GA$ ) or the FIGARCH model ( $i = FI$ ). See the discussion in Stărică and Granger (2005) for the choice of this

particular regression. The goal is to see first if the forecasts from the level shift model are uncorrelated with the forecast errors from either of the two competitors. This is done by testing the null hypothesis  $H_0^A : (b_0, b_1, b_2) = (0, 0, 1)$  using a standard Wald test with an asymptotic chi-square distribution. The results favor the level shift model if the null hypothesis can be rejected. Conversely, one can test if the forecast errors from the level shift model are uncorrelated with the forecasts from either of the two competitors. This is done by testing the null hypothesis  $H_0^B : (b_0, b_1, b_2) = (0, 1, 0)$  using again a standard Wald test. The results favor the level shift model if the null hypothesis cannot be rejected. An outcome for which the p-value for the test pertaining to  $H_0^A$  is less (resp., greater) than 5% and that for the test pertaining to  $H_0^B$  is greater (resp., less) than 5% is viewed as evidence in favor (resp., against) the level shift model. The cases for which both p-values are either greater or smaller than 5% are viewed as providing no evidence in favor of either model.

The results are presented in Table 3 for various horizons between 20 and 200 days. Consider first the comparison between the level shift and GARCH models. For the S&P 500 and AMEX indices, the evidence is in favor of the level shift model at all horizons. For the Dow Jones index, the evidence is also in favor of the level shift model for horizons up to 100 days while for longer horizons the tests are inconclusive (both tests have p-values less than 5%). For the NASDAQ index, the level shift model is deemed superior for horizons 20, 40, 60, 80, 140 and 160 while the GARCH model performs better at horizons 100, 120 and 180 (for a horizon of 200 days, the evidence is inconclusive). Hence, the overall results strongly indicate that the forecasting performance of the level shift model is superior to that of the GARCH(1,1) model. When compared to the FIGARCH model, the level shift model generally performs better at shorter horizons: up to 100 days for the S&P 500 and Dow Jones indices, up to 60 days for the NASDAQ index and at horizon 20 for the AMEX index. The FIGARCH occasionally performs better: at horizon 200 for the S&P 500 and AMEX indices and at horizons 80, 100, 120 and 180 for the NASDAQ index. For all other cases, the evidence does not favor one model over the other, using the criterion adopted.

Overall, the results are encouraging and show that the random level shift model is a serious contender as a forecasting model and that gains in forecast accuracy can be substantial. These results are important since not only does the level shift model provide a better description of the data in sample, it does so without sacrificing forecasting performance, often increasing it. Hence, level shifts appear to be genuinely present in the data and not simply a modeling convenience that allows for a better in-sample fit. This contrasts with the common perception that structural change models are not useful for forecasting.

## 6 Conclusion

In this paper, we first proposed a simple estimation method for a model of a series composed of a random level shift and a short-memory component. The model, applied to log-absolute returns of the S&P 500, AMEX, Dow Jones and NASDAQ indices, yields impressive results. The level shift model with white noise errors appears to provide an accurate description of the data. The level shift component is an important feature that explains the presence of both the long-memory and conditional heteroskedasticity features that are generally perceived as stylized facts. When accounting for level shifts, the evidence in favor of long-memory and conditional heteroskedasticity disappears. The model can also provide important improvements in forecasting volatility when using squared returns as a proxy. These are especially noticeable when the mean of the regime in effect at the time of forecasting is well estimated. Precise estimates of this mean are, however, difficult to obtain in real time. We have nevertheless shown that in many cases, forecasting improvements can still be obtained from our model by using a backward recursive CUSUM test to determine the length of the last regime.

## References

- Andersen, T.G., and T. Bollerslev (1997), "Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-run in High Frequency Data," *Journal of Finance* 52, 975-1005.
- Andersen, T.G., and T. Bollerslev (1998), "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts," *International Economic Review* 39, 885-905.
- Anderson, T.G., T. Bollerslev, F.X. Diebold, and P. Labys (2001): "The Distribution of Realized Exchange Rate Volatility," *Journal of the American Statistical Association* 96, 42-55.
- Anderson, T.G., T. Bollerslev, F.X. Diebold, and P. Labys (2003): "Modeling and Forecasting Realized Volatility," *Econometrica*, 71, 579-626.
- Arteche, J. (2004), "Gaussian Semiparametric Estimation in Long Memory in Stochastic Volatility and Signal Plus Noise Models," *Journal of Econometrics* 119, 131-154.
- Baillie, R.T. (1996), "Long Memory Processes and Fractional Integration in Econometrics," *Journal of Econometrics* 73, 5-59.
- Baillie, R.T., T. Bollerslev, and H.O. Mikkelsen (1996), "Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics* 73, 3-30.
- Beran, J. (1994), *Statistics for Long Memory Processes*, New York: Chapman & Hall.
- Bollerslev, T. (1986), "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics* 31, 307-327.
- Bollerslev, T., and H.O. Mikkelsen (1996), "Modeling and Pricing Long Memory in Stock Market Volatility," *Journal of Econometrics* 73, 151-184.
- Bollerslev, T., and J.H. Wright (2000), "Semiparametric Estimation of Long-memory Volatility Dependencies: The Role of High Frequency Data," *Journal of Econometrics*, 98, 81-106.
- Breidt, J.F., N. Crato, and P.J.F. de Lima (1998), "On the Detection and Estimation of Long-memory in Stochastic Volatility," *Journal of Econometrics* 83, 325-348.
- Brown, R.L., J. Durbin, and J.M. Evans (1975), "Techniques for Testing the Constancy of Regression Relationships Over Time," *Journal of the Royal Statistical Society B* 37, 149-163.
- Deo, R.S., and C.M. Hurvich (2001), "On the Log Periodogram Regression Estimator of the Memory Parameter in the Long Memory Stochastic Volatility Models," *Econometric Theory* 17, 686-710.
- Diebold, F., and A. Inoue (2001), "Long Memory and Regime Switching," *Journal of Econometrics* 105, 131-159.
- Engle, R. F. (1982), "Autoregressive Conditional Heteroscedasticity with Estimates of Variance of United Kingdom Inflation," *Econometrica* 50, 987-1008.
- Engle, R.F. (1995), *ARCH: Selected Readings*. Oxford University Press.

- Engle, R.F., and A.D. Smith (1999), "Stochastic Permanent Breaks," *Review of Economics and Statistics* 81, 533-574.
- Filardo, A. J., and S. F. Gordon (1998), "Business Cycle Durations," *Journal of Econometrics* 85, 99-123.
- Gourieroux, C., and J. Jasiak (2001), "Memory and Infrequent Breaks," *Economics Letters* 70, 29-41.
- Granger, C.W.J. (1981), "Some Properties of Time Series Data and Their Use in Econometric Model Specification," *Journal of Econometrics* 16, 121-130.
- Granger, C.W.J., and Z. Ding (1995), "Some Properties of Absolute Return: an Alternative Measure of Risk," *Annales d'Économie et de Statistique*, 40, 67-91.
- Granger, C.W.J., and Z. Ding (1996), "Varieties of Long Memory Models," *Journal of Econometrics* 73, 61-77.
- Granger, C.W.J., and N. Hyung (2004), "Occasional Structural Breaks and Long Memory with an Application to the S&P 500 Absolute Stock Returns," *Journal of Empirical Finance* 11, 399-421.
- Granger, C.W. and R. Joyeux (1980), "An Introduction to Long-Memory Time Series Models and Fractional Differencing," *Journal of Time Series Analysis* 1, 15-29.
- Gray, S. (1996), "Modeling the Conditional Distribution of Interest Rates as a Regime Switching Process," *Journal of Financial Economics* 42, 27-62.
- Hamilton, J. (1989), "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business-Cycle," *Econometrica* 57, 357-384.
- Harrison, P. J., and C. F. Stevens (1976), "Bayesian Forecasting," *Journal of the Royal Statistical Society, Series B* 38, 205-247.
- Harvey, A.C. (1998), "Long Memory in Stochastic Volatility," in *Forecasting Volatility in Financial Markets*, ed. by J. Knight and S. Satchell. Oxford: Butterworth-Heinemann, 307-320.
- Hosking, J. R. (1981), "Fractional Differencing," *Biometrika* 68, 165-176.
- Lamoureux, C.G., and W.D. Lastrapes (1990), "Persistence in Variance, Structural Change, and the GARCH Model," *Journal of Business & Economic Statistics* 8, 225-234.
- Lobato, I.N., and N.E. Savin (1998), "Real and Spurious Long-memory Properties of Stock-market Data," *Journal of Business and Economic Statistics* 16, 261-268.
- Mikosch, T., and C. Stărică (2005), "Changes of Structure in Financial Time Series and the GARCH Model," Unpublished manuscript.
- Mincer, J. and V. Zarnowitz (1969): "The Evaluation of Economic Forecasts," in *Economic Forecasts and Expectations*, ed. by J. Mincer. New York: National Bureau of Economic Research.

- Morana, C., and A. Beltratti (2004), "Structural Change and Long-Range Dependence in Volatility of Exchange Rates: Either, Neither, or Both," *Journal of Empirical Finance* 11, 629-658.
- Parke, W.R. (1999), "What is fractional integration?" *Review of Economics and Statistics* 81, 632-638.
- Perron, P. (1989), "The Great Crash, the Oil Price Shock and the Unit Root Hypothesis," *Econometrica* 57, 1361-1401.
- Perron, P. (1990), "Testing for a Unit Root in a Time Series Regression with a Changing Mean," *Journal of Business and Economic Statistics* 8, 153-162.
- Perron, P. and Z. Qu (2007), "An Analytical Evaluation of the Log-periodogram Estimate in the Presence of Level Shifts," Unpublished Manuscript, Department of Economics, Boston University.
- Perron, P., and Z. Qu (2008), "Long-Memory and Level Shifts in the Volatility of Stock Market Return Indices", Unpublished Manuscript, Department of Economics, Boston University.
- Pesaran, M.H., and A. Timmerman (1999), " Model Instability and Choice of Observation Window," Discussion Paper 99-19, Department of Economics, University of California, San Diego.
- Qu, Z., and P. Perron (2008), "A Stochastic Volatility Model with Random Level Shifts: Theory and Applications to S&P 500 and NASDAQ Indices," Unpublished Manuscript, Department of Economics, Boston University.
- Robinson, P.M. (1994), "Time Series with Strong Dependence," in: C. Sims, ed., *Advances in Econometrics*, 6th World Congress (Cambridge University Press, Cambridge), 47-95.
- Shephard, N. (2005): *Stochastic Volatility: Selected Readings*, Edited volume, Oxford University Press.
- Stărică, C., and C. Granger (2005), "Nonstationarities in Stock Returns," *The Review of Economics and Statistics* 87, 503-522.
- Teverovsky, V., and Taqqu, M. (1997), "Testing for Long-range Dependence in the Presence of Shifting Means or a Slowly Declining Trend, Using a Variance-type Estimator," *Journal of Time Series Analysis* 18, 279-304.
- Wada, T., and P. Perron (2006), "Let's Take a Break: Trends and Cycles in U.S. Real GDP," mimeo, Boston University.
- Wada, T., and P. Perron (2007), "An Alternative Trend-Cycle Decomposition using a State Space Model with Mixtures of Normals: Specifications and Applications to International Data," mimeo, Boston University.

**Table 1: Maximum Likelihood Estimates.**

			e	
<b>S&amp;P 500</b>	0.75123	0.00145	0.73995	
(SD: 0.8021)	0.74290	0.00150	0.73950	-0.01040
<b>AMEX</b>	0.98696	0.00246	0.70346	
(SD: 0.7773)	1.08830	0.00167	0.70607	0.06321
<b>Dow Jones</b>	0.95720	0.00100	0.73947	
(SD: 0.7888)	0.95860	0.00103	0.73909	-0.00984
<b>NASDAQ</b>	1.45396	0.00077	0.74255	
(SD: 0.8528)	1.45589	0.00076	0.74258	0.00070

**Table 2: Parameter Estimates; GARCH and CGARCH models**

	<b>S&amp;P500</b>					<b>Dow Jones</b>				
	Coefficient	Estimate	Std. Error	t-Statistic	p-value	Coefficient	Estimate	Std. Error	t-Statistic	p-value
No level shifts in GARCH	$\beta_r$	0.066	0.005	14.38	0	$\beta_r$	0.059	0.004	14	0
	$\beta_\sigma$	0.932	0.005	209.35	0	$\beta_\sigma$	0.935	0.004	212.39	0
No level shifts in CGARCH	$\beta_r$	0.047	0.007	7.14	0	$\beta_r$	0.041	0.007	5.889	0
	$\beta_\sigma$	0.926	0.010	89.07	0	$\beta_\sigma$	0.931	0.010	91.23	0
	$\rho$	0.9998	0.0003	3149.13	0	$\rho$	0.998	0.001	1044.74	0
	$\varphi$	0.027	0.005	5.36	0	$\varphi$	0.023	0.006	3.709	0
Level shifts in CGARCH	$\beta_r$	0.009	0.011	0.786	0.432	$\beta_r$	0.015	0.010	1.48	0.139
	$\beta_\sigma$	0.151	1.186	0.127	0.899	$\beta_\sigma$	0.034	0.629	0.053	0.958
	$\rho$	0.942	0.007	141.92	0	$\rho$	0.948	0.007	140.6	0
	$\varphi$	0.062	0.006	9.65	0	$\varphi$	0.053	0.006	9.47	0
	<b>AMEX</b>					<b>NASDAQ</b>				
	Coefficient	Estimate	Std. Error	t-Statistic	p-value	Coefficient	Estimate	Std. Error	t-Statistic	p-value
No level shifts in GARCH	$\beta_r$	0.127	0.007	17.15	0	$\beta_r$	0.094	0.007	14.40	0
	$\beta_\sigma$	0.867	0.007	124.77	0	$\beta_\sigma$	0.904	0.006	145.36	0
No level shifts in CGARCH	$\beta_r$	0.123	0.010	12.01	0	$\beta_r$	0.082	0.010	7.83	0
	$\beta_\sigma$	0.797	0.018	43.49	0	$\beta_\sigma$	0.839	0.023	36.16	0
	$\rho$	0.999	0.001	973.16	0	$\rho$	0.9999	0.000	4964.46	0
	$\varphi$	0.035	0.006	5.82	0	$\varphi$	0.039	0.005	7.39	0
Level shifts in CGARCH	$\beta_r$	0.050	0.020	2.57	0.010	$\beta_r$	0.049	0.017	2.91	0.004
	$\beta_\sigma$	-0.045	0.291	-0.155	0.877	$\beta_\sigma$	0.068	0.304	0.222	0.824
	$\rho$	0.854	0.018	48.39	0	$\rho$	0.952	0.008	122.69	0
	$\varphi$	0.114	0.015	7.54	0	$\varphi$	0.086	0.01	8.85	0

Table 3: Comparison of forecasting performance between the level shift, GARCH and FIGRACH models.

Horizon p (days)	p value of Wald statistics for				p value of Wald statistics for			
	$H_0^A$		$H_0^B$		$H_0^A$		$H_0^B$	
	$f_{t,p}^{LS}$	$e_{t,p}^{GR}$	$e_{t,p}^{LS}$	$f_{t,p}^{GR}$	$f_{t,p}^{LS}$	$e_{t,p}^{FI}$	$e_{t,p}^{LS}$	$f_{t,p}^{FI}$
	S&P 500				Dow Jones			
20	0.00	0.64	0.00	0.57	0.00	0.13	0.00	0.09
40	0.00	0.85	0.01	0.75	0.00	0.55	0.00	0.55
60	0.00	0.79	0.00	0.45	0.00	0.43	0.00	0.25
80	0.00	0.99	0.08	0.58	0.00	0.24	0.00	0.18
100	0.00	0.67	0.05	0.36	0.00	0.26	0.00	0.17
120	0.00	0.54	0.13	0.09	0.00	0.00	0.00	0.00
140	0.00	0.50	0.27	0.07	0.00	0.02	0.00	0.06
160	0.00	0.96	0.32	0.19	0.00	0.01	0.00	0.02
180	0.00	0.97	0.30	0.27	0.00	0.00	0.00	0.04
200	0.00	0.17	0.11	0.02	0.02	0.00	0.04	0.00
	AMEX				NASDAQ			
20	0.00	0.68	0.04	0.51	0.00	0.48	0.00	0.53
40	0.00	0.58	0.26	0.15	0.00	0.80	0.00	0.94
60	0.00	0.67	0.42	0.14	0.00	0.76	0.08	0.44
80	0.00	0.69	0.51	0.35	0.00	0.12	0.07	0.01
100	0.00	0.64	0.53	0.18	0.35	0.00	0.16	0.00
120	0.00	0.52	0.39	0.31	0.22	0.00	0.45	0.00
140	0.00	0.86	0.79	0.12	0.00	0.40	0.66	0.10
160	0.00	0.78	0.71	0.11	0.00	0.49	0.53	0.11
180	0.00	0.85	0.79	0.11	0.12	0.00	0.35	0.00
200	0.00	0.93	0.67	0.03	0.00	0.00	0.01	0.00




Figure 1.a: S&P 500; fitted level shift component (right axis) and series (left axis).




Figure 1.b: AMEX; fitted level shift component (right axis) and series (left axis).




Figure 1.c: Dow Jones; fitted level shift component (right axis) and series (left axis).






Figure 1.d: NASDAQ; fitted level shift component (right axis) and series (left axis).


**ACF of S&P500**



**ACF of DJ**



**ACF of S&P500 Residual**



**ACF of DJ Residual**






Figure 2.a: S&P 500; Autocorrelations.

Figure 2.b: Dow Jones; Autocorrelations.


**ACF of AMEX**



**ACF of NASDAQ**



**ACF of AMEX Residual**



**ACF of NASDAQ Residual**




Figure 2.c: AMEX; Autocorrelations.

Figure 2.d: NASDAQ; Autocorrelations.


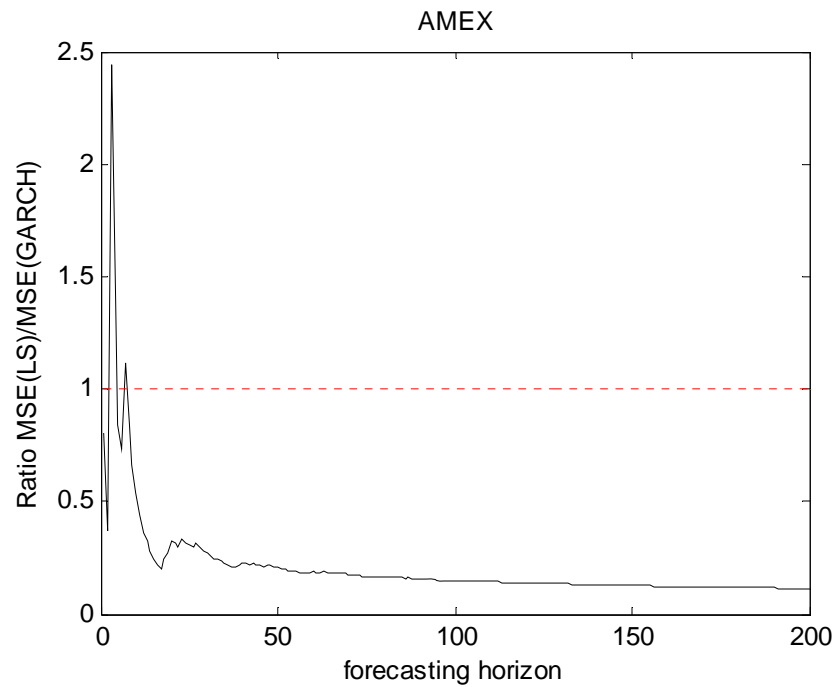




Figure 3: In-sample level shift versus out-of-sample GARCH forecasts



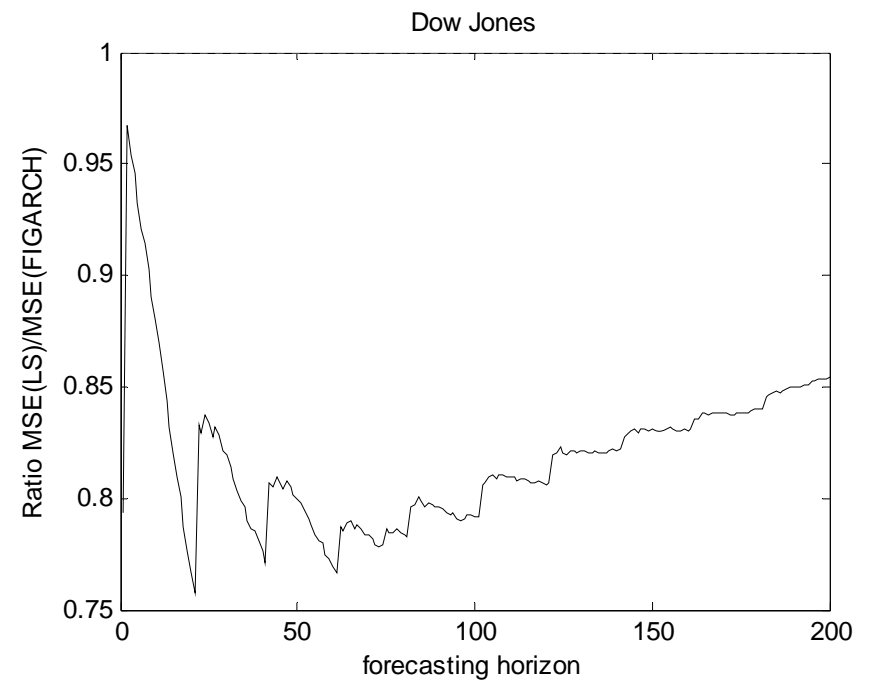
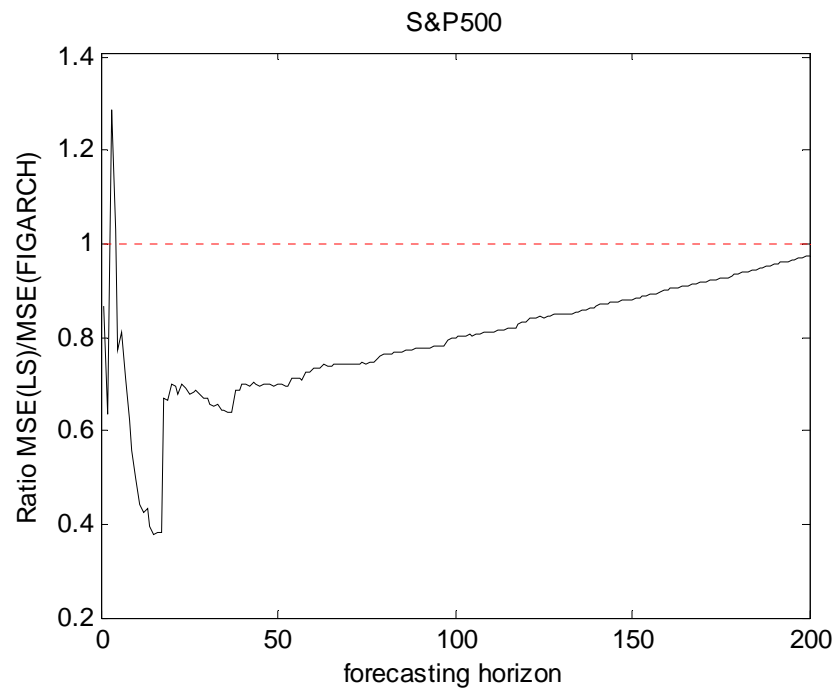


Figure 4: In-sample level shift versus out-of-sample FIGARCH forecasts

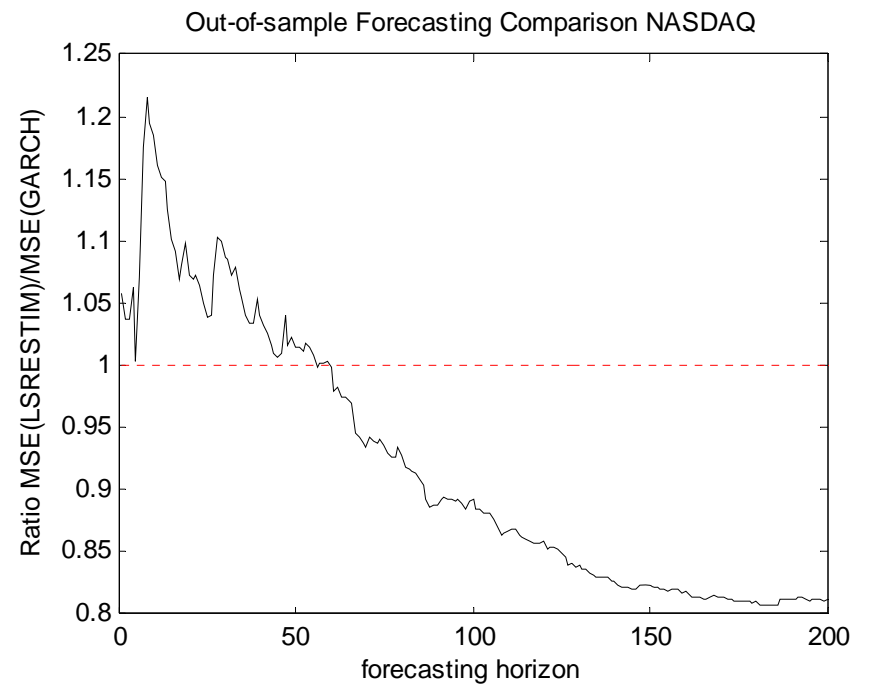
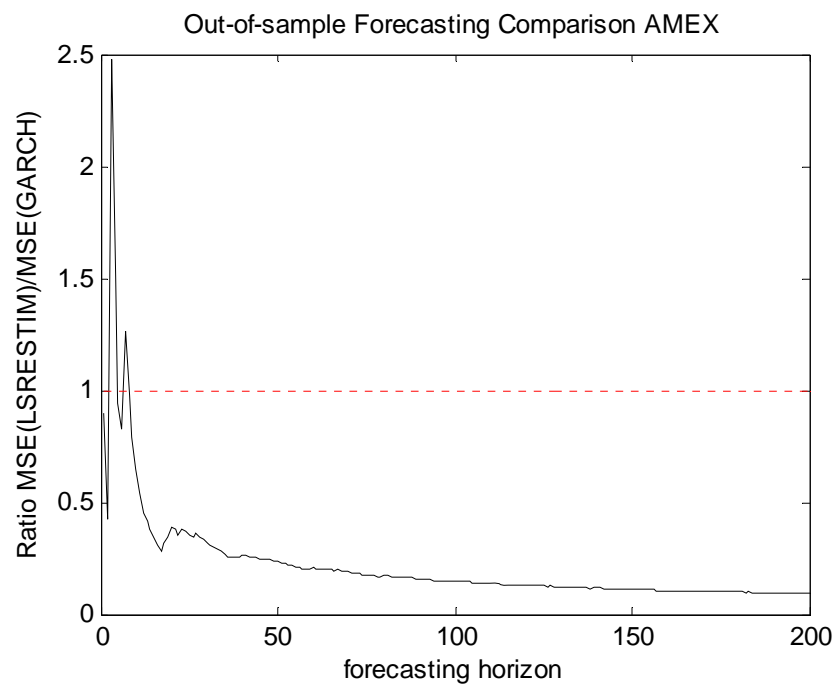
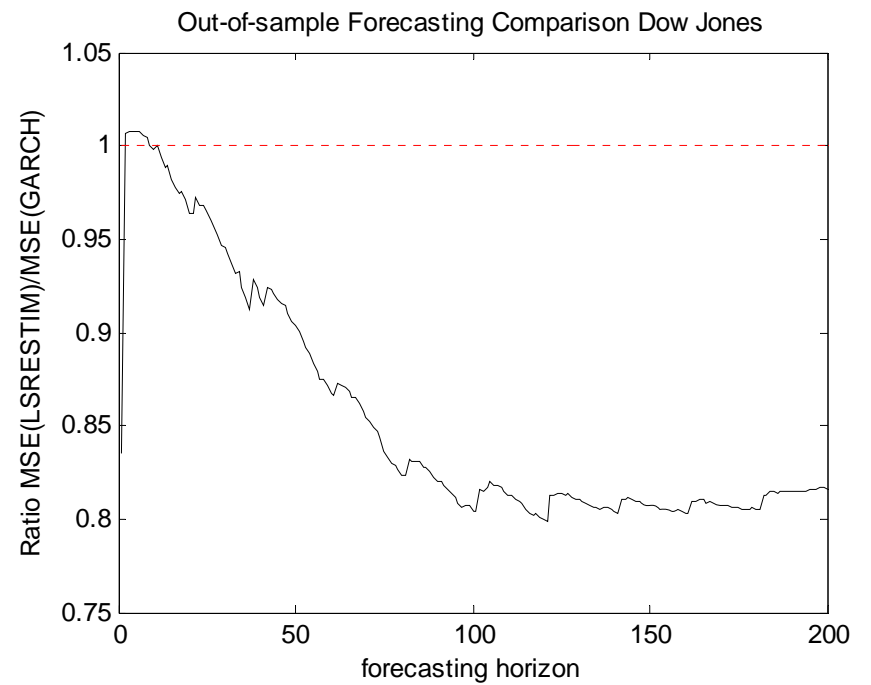
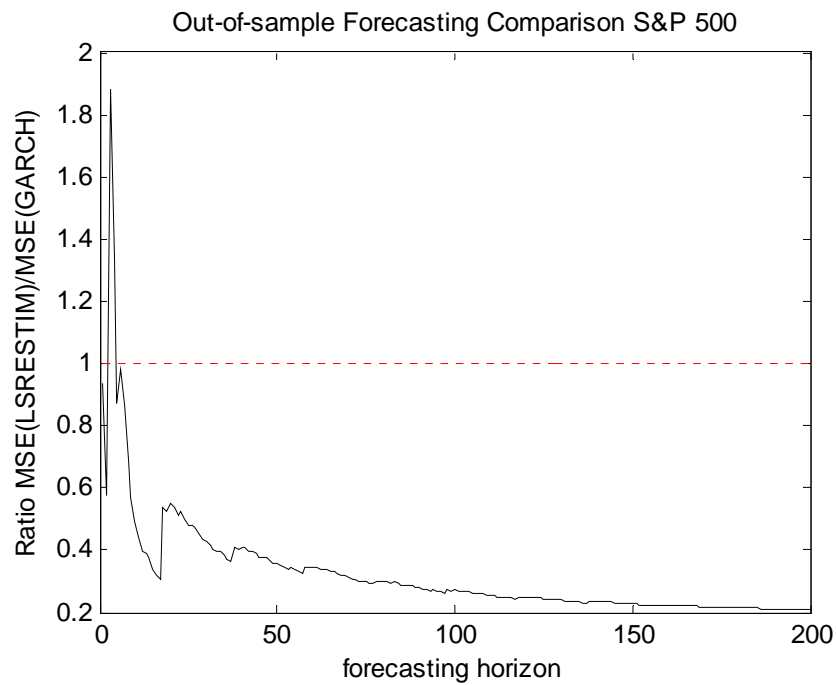


Figure 5: Out-of-sample level shift versus out-of-sample GARCH forecasts

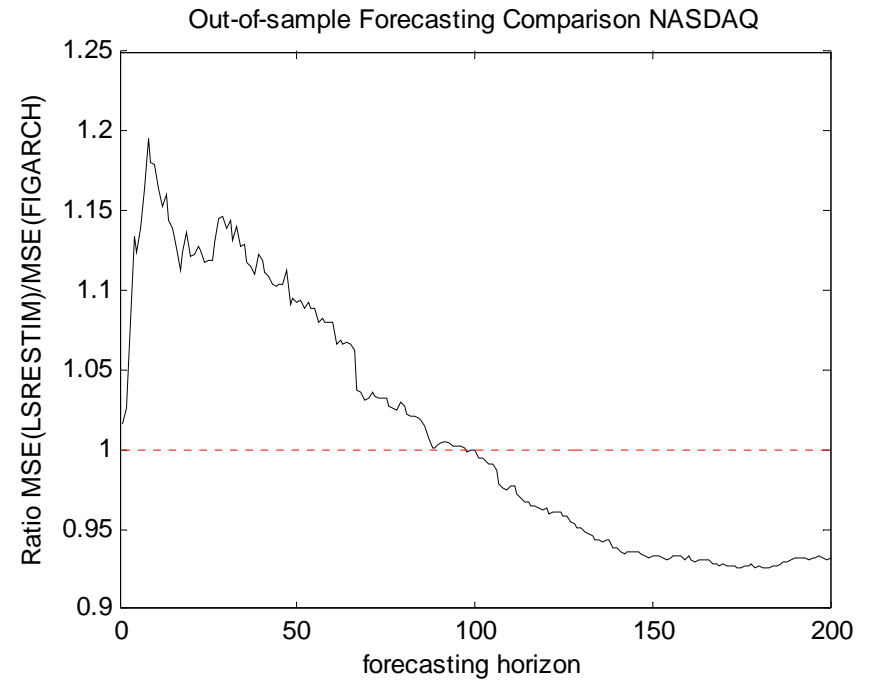


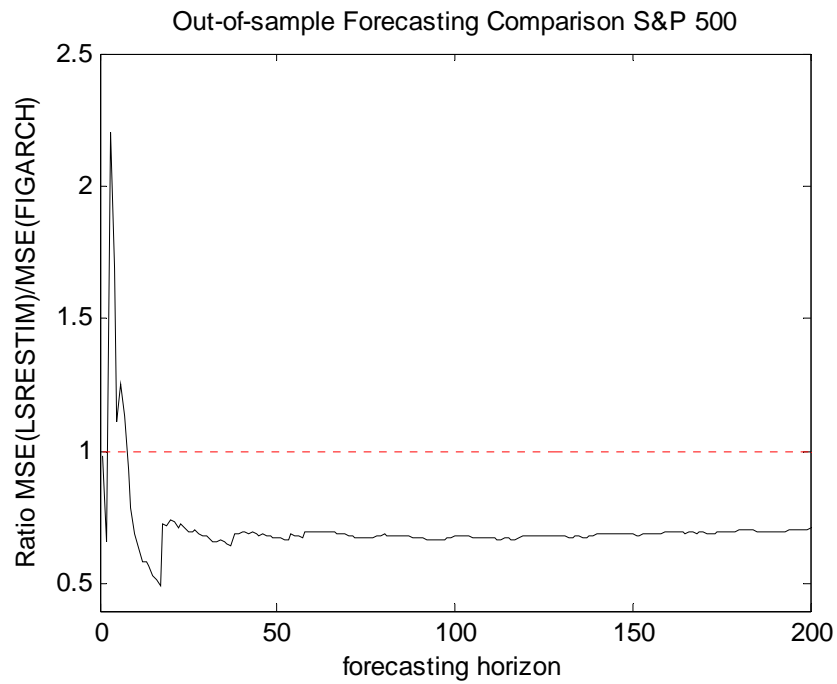


Figure 6: Out-of-sample level shift versus out-of-sample FIGARCH forecasts