

A generalized partially linear model of asymmetric volatility

Guojun Wu^{a,*}, Zhijie Xiao^{b,1}

^aUniversity of Michigan Business School, 701 Tappan Street, Ann Arbor, MI 48109, USA

^b185 Commerce West Building, University of Illinois at Urbana-Champaign, 1206 South Sixth Street, Champaign, IL 61820, USA

Accepted 23 November 2001

Abstract

In this paper we conduct a close examination of the relationship between return shocks and conditional volatility. We do so in a framework where the impact of return shocks on conditional volatility is specified as a general function and estimated nonparametrically using implied volatility data—the Market Volatility Index (VIX). This setup can provide a good description of the impact of return shocks on conditional volatility, and it appears that the news impact curves implied by the VIX data are useful in selecting ARCH specifications at the weekly frequency. We find that the Exponential ARCH model of Nelson [Econometrica 59 (1991) 347] is capable of capturing most of the asymmetric effect, when return shocks are relatively small. For large negative shocks, our nonparametric function points to larger increases in conditional volatility than those predicted by a standard EGARCH. Our empirical analysis further demonstrates that an EGARCH model with separate coefficients for large and small negative shocks is better able to capture the asymmetric effect. © 2002 Elsevier Science B.V. All rights reserved.

JEL classification: G10; C14

Keywords: Asymmetric volatility; News impact curve

* Corresponding author. Tel.: +1-734-936-3248; fax: +1-734-764-2555.

E-mail addresses: gjwu@umich.edu (G. Wu), zxiao@uiuc.edu (Z. Xiao).

¹ Tel.: +1-217-333-4520; fax: +1-217-244-6678.

1. Introduction

Over the past several decades researchers have documented strong evidence that volatility is asymmetric in equity markets: negative returns are generally associated with upward revisions of the conditional volatility while positive returns are associated with smaller upward or even downward revisions of the conditional volatility (see, for example, Cox and Ross, 1976). Researchers (see Black, 1976; Christie, 1982; Schwert, 1989) believe that the asymmetry could be due to changes in leverage in response to changes in the value of equity. Others have argued that the asymmetry could arise from the feedback from volatility to stock price when changes in volatility induce changes in risk premiums (see Pindyck, 1984; French et al., 1987; Campbell and Hentschel, 1992; Wu, 2001). The presence of asymmetric volatility is most apparent during a market crisis when large declines in stock prices are associated with a significant increase in market volatility. Asymmetric volatility can potentially explain the negative skewness in stock return data, as discussed in Harvey and Siddique (1999).

Formal econometric models have been developed by researchers to capture asymmetric volatility and currently two main classes of time series models allow for asymmetric volatility. The first class is based on continuous-time stochastic volatility. By accommodating a constant correlation between stock price and volatility diffusion, these models generally produce estimates of a negative correlation between return and return volatility (Bates, 1997; Bakshi et al., 1997). The second class of models extends the ARCH models (see Bollerslev, 1986; Pagan and Schwert, 1990; Bollerslev et al., 1992; Engle and Ng, 1993). For example, the exponential ARCH model of Nelson (1991) and the asymmetric GARCH models of Glosten et al. (1993), referred to as GJR below) and of Hentschel (1995) have been found to outperform significantly models that do not accommodate the asymmetry. The threshold ARCH models of Rabemananjara and Zakoian (1993) and Zakoian (1994) are motivated by and designed to capture the asymmetry. Bekaert and Wu (2000) used the asymmetric BEKK model of Baba et al. (1989) and Engle and Kroner (1995) to investigate asymmetric volatility simultaneously at the market level and individual firm level for the Japanese equity market. Asymmetric volatility was found to be significant both at the market level and individual firm level.

These two classes of asymmetric volatility models both assume a parametric form for the asymmetry. The continuous-time stochastic volatility models constrain the negative correlation between the instantaneous stock return and volatility to be a constant. The exponential ARCH model specifies a negative exponential relationship between conditional volatility and return shock. The asymmetric ARCH models essentially allow negative and positive return shocks to have separate positive coefficients, with the former being larger than the latter. All of these models are able to capture the negative correlation between volatility and returns, yet the exact relationship between return and volatility is not well understood. For example, to make sure volatility is always positive, asymmetric ARCH models generally do not allow volatility to decline for a positive return shock. However, positive shocks may not affect or even reduce volatility, rather than increase it. Moreover, it is not clear from the existing literature whether we should model the relationship between negative return shocks and

increases in volatility as linear, exponential, or other functional forms. Since asymmetric volatility is important for time series modeling and derivative pricing, understanding the exact nature of this relationship may have important implications (see Amin and Ng, 1994; Duan, 1995).

In this paper we extend the news impact curve analysis of Engle and Ng (1993) in order to examine the relationship between return shocks and conditional volatility. We do so in a partially linear framework where the impact of return shock on conditional volatility is specified as a general function and estimated nonparametrically by using the implied volatility as a proxy for conditional return volatility. Since we treat volatility as an “observable” in our analysis, we offer a novel approach to the analysis of the news impact curve. Andersen and Bollerslev (1998) showed that ARCH models provide surprisingly good estimates for conditional volatilities. We think that appropriately computed implied volatilities directly incorporate market expectations and hence may provide better estimates of future volatilities. This setup allows us to study this relationship from a different direction and we hope that this analysis will help us better identify the exact impact of return shocks on conditional volatility.

Our generalized partially linear model is estimated with returns and the Black and Scholes (1973) implied volatilities of options on the S&P 100 Index. Specifically, the implied volatility is the Market Volatility Index (VIX) constructed by Whaley (1993, 2000) for the Chicago Board Options Exchange. The VIX series represents the implied volatility of a 30-calendar day (22-trading day) at-the-money option on the S&P 100 Index. The estimation results confirm the existence of asymmetric volatility: negative return shocks are associated with increases in conditional volatility while positive return shocks are generally uncorrelated with changes in conditional volatility. We find that the exponential ARCH model is capable of capturing most of the asymmetric effect when return shocks are relatively small. For larger negative shocks, however, our nonparametric function points to larger increases in conditional volatility than those predicted by a traditional EGARCH model. Our empirical analysis further demonstrates that an EGARCH model with separate coefficients for large and small negative shocks is better able to capture the asymmetric effect.

There are some issues associated with the use of implied volatilities. First of all, since the true data generating process of the S&P 100 Index is likely to be more complicated than the simple geometric Brownian motion assumed in the Black–Scholes model, the implied volatility is at best a proxy for the true conditional volatility. Second, the estimation of the implied volatilities from the option price data introduces further errors in the volatility series. Fortunately, the measurement error problem can be handled properly in our empirical analysis with the use of an instrumental variable. Third, the implied volatility series is unlikely to be identical to the conditional volatility series generated by models based on the historical prices of the underlying index, such as the ARCH-class of models. However, if one assumes that the implied volatility is related to the true conditional volatility process, knowledge of implied volatility will be helpful in the specification analysis of the ARCH models.

The remainder of the article is organized as follows. Section 2 motivates and describes the generalized partially linear asymmetric volatility model. The estimation methodology is then introduced. We also discuss the choice of bandwidths in the nonparametric

analysis. Section 3 provides data descriptions and conducts the estimation. We also provide an analysis of stationarity and conduct specification tests on the generalized partially linear model. Empirical results regarding asymmetric volatility are reported and discussed. Section 4 attempts to utilize the generalized partially linear model to investigate the performance of various EGARCH models. We also conduct a series of sign bias tests on the specification of these models. Section 5 contains the concluding remarks. Some technical details are included in Appendix A.

2. A generalized asymmetric volatility model

2.1. The model

Let's consider a model of stock return r_t ,

$$r_t = \mu_{t-1} + u_t = \mu_{t-1} + \sigma_t \xi_t,$$

where $u_t | \sigma_t \sim N(0, \sigma_t^2)$. The conditional variance σ_t^2 follows a stochastic process specified as

$$\sigma_t^2 = \alpha_0 + \alpha_1 \sigma_{t-1}^2 + g(u_{t-1}) + v_t, \quad (1)$$

where v_t is an error term which may contain errors in estimating σ_t . The relationship between the volatility σ_t^2 and u_{t-1} is given by a general function $g(u_{t-1})$. Setting $g(u_{t-1}) = \alpha_2 u_{t-1}^2$, we see that model (1) becomes the important special case of GARCH(1, 1), which has been widely used in empirical research (see Bollerslev et al., 1992). In the asymmetric GARCH model of Glosten et al. (1993), for example, $g(u_{t-1}) = \alpha_2 u_{t-1}^2 + \alpha_3 \max(0, -u_{t-1})^2$. Many other specifications for $g(u)$ are proposed in the literature.

To ensure positive variance, we focus our partially linear analysis on the log volatility process. The methodology can easily be extended to the variance process itself, but the model needs to be constrained as in a GARCH setting. We may consider, as an alternative to model (1), the following partially linear model:

$$\ln \sigma_t = \alpha_0 + \alpha_1 \ln \sigma_{t-1} + g(u_{t-1}) + \varepsilon_t, \quad (2)$$

where ε_t is an error term. The purpose of the partially linear analysis is to investigate if a general functional form $g(u)$ is implied by the options data. In the following analysis we do not specify any functional form for $g(u)$, but estimate it nonparametrically.

To approach the estimation of $g(u)$, we reparameterize the model (2) and denote $\beta = (\alpha_0, \alpha_1)'$, $x_t = (1, \ln \sigma_{t-1})'$, and $y_t = \ln \sigma_t$, then

$$y_t = \beta' x_t + g(u_{t-1}) + \varepsilon_t, \quad (3)$$

where ε_t is an i.i.d. random variable independent of x_t and u_{t-1} , $g(\cdot)$ is an unknown real function, and β is the vector of unknown parameters that we want to estimate. In this

model, the mean response of log volatility $\ln \sigma_t$ is assumed to be linearly related to $(1, \ln \sigma_{t-1})$ and nonparametrically related to u_{t-1} . We now describe the estimation methodology.

2.2. Estimation methodology

Partially linear models have been an important object of study in econometrics and statistics for a long time. One primary approach is the penalized least squares method employed by Wahba (1984), among others. Estimation of the parameters is obtained by adding one penalized term to the ordinary nonlinear least squares to penalize for roughness in the fitted $g(\cdot)$. Heckman (1986) and Chen (1985) proved that the estimate of β can achieve the \sqrt{n} convergence rate if x and $g(\cdot)$ are not related to each other. Rice (1986) obtained the asymptotic bias of a partial smoothing spline estimator of β due to the dependence between x and $g(\cdot)$ and showed that it is not generally possible to attain the \sqrt{n} convergence rate for β . Green and Yandell (1985) and Speckman (1988) suggested a method of simultaneous equations to estimate both β and $g(\cdot)$. The \sqrt{n} consistency and asymptotic normality are established in Speckman (1988).

Robinson (1988) addressed the efficiency issue when the regression errors are independently and identically distributed (i.i.d.) normal variates. In particular, he used (higher order) Nadaraya–Watson kernel estimates to eliminate the unknown function $g(\cdot)$ and introduced a feasible least squares estimator for β . Under regularity and smoothness conditions, \sqrt{n} consistency and asymptotic normality are obtained. When the errors are normally i.i.d., this estimator achieves the semiparametric information bound. A higher order asymptotic analysis on the estimator is given by Linton (1995). Fan and Li (1999) established the \sqrt{n} -consistent estimation and asymptotic normal distribution results under β -mixing conditions.

Taking expectation of Eq. (3) conditional on u_{t-1} , we have

$$E(y_t | u_{t-1}) = \beta' E(x_t | u_{t-1}) + g(u_{t-1}). \quad (4)$$

Thus, combining Eqs. (3) and (4) gives

$$y_t - E(y_t | u_{t-1}) = \beta' (x_t - E(x_t | u_{t-1})) + \varepsilon_t.$$

Denoting

$$y_t^* = y_t - E(y_t | u_{t-1}),$$

and

$$x_t^* = x_t - E(x_t | u_{t-1}),$$

we get the following parametric regression equation:

$$y_t^* = \beta' x_t^* + \varepsilon_t. \quad (5)$$

This equation shows that if we are able to obtain estimates of the conditional expectations, we may be able to estimate the parameter β using the linear components of the model. With β estimated, we will then be able to estimate $g(\cdot)$ by utilizing Eq. (4). Estimation and inference of partially linear regression models have been widely studied. Among different estimation methods, the semiparametric kernel method is probably one of the most commonly used. This method was originally studied by Robinson (1988) and has been extended in various directions. Under regularity conditions [see, Robinson (1988); Fan and Li (1999), and references therein, for discussions on a wide range of models and regularity conditions], the conditional expectations $E(y_t|u_{t-1})$ and $E(x_t|u_{t-1})$ can be estimated nonparametrically by the standard Nadaraya–Watson kernel method:

$$E(\widehat{y_t | u_{t-1}}) = \frac{1}{nh} \sum_{j=1}^n K\left(\frac{u_{t-1} - u_j}{h}\right) y_j / \hat{f}(u_{t-1}),$$

$$E(\widehat{x_t | u_{t-1}}) = \frac{1}{nh} \sum_{j=1}^n K\left(\frac{u_{t-1} - u_j}{h}\right) x_j / \hat{f}(u_{t-1}),$$

where

$$\hat{f}(u_{t-1}) = \frac{1}{nh} \sum_{j=1}^n K\left(\frac{u_{t-1} - u_j}{h}\right) \tag{6}$$

is a consistent density estimator for $f(u_{t-1})$ under bandwidth conditions, n is the sample size, h is the bandwidth parameter, and $K(\cdot)$ is the kernel function satisfying the properties that it is a real, even function, having support on $[-1, 1]$, $\int K(u)du = 1$, $\int u^l K(u)du = 0$, for $l = 1, \dots, q-1$, and $\int u^q K(u)du \neq 0$. q is defined as the characteristic component of the corresponding kernel $K(\cdot)$. The requirement that K integrates to 1 makes $\hat{f}(u_{t-1})$ an appropriate estimator of $f(u_{t-1})$. Candidate kernel functions can be found in standard econometric texts [also see Robinson (1988) on discussions of kernel functions]. Define

$$\hat{y}_t^* = y_t - E(\widehat{y_t | u_{t-1}}), \text{ and } \hat{x}_t^* = x_t - E(\widehat{x_t | u_{t-1}}); \tag{7}$$

then a least squares estimator for β is given by

$$\hat{\beta} = \left[\sum \hat{x}_t^* \hat{x}_t^{*'} \right]^{-1} \left[\sum \hat{x}_t^* \hat{y}_t^* \right]. \tag{8}$$

and $g(u_{t-1})$ can be estimated based on $\hat{\beta}$:

$$\tilde{g}(u_{t-1}) = E(\widehat{y_t | u_{t-1}}) - \hat{\beta}' E(\widehat{x_t | u_{t-1}}). \tag{9}$$

Given data on y_t , x_t , and u_t , the above procedure delivers \sqrt{n} -consistent estimators of β and $g(u)$. However, notice that in our partially linear model (2), the conditional volatility

σ_t has to be estimated. Specifically we use the VIX, which is a weighted average implied volatility from the S&P 100 Index options.

However, the estimation in σ_t introduces an errors-in-variable problem so that the estimator for β given by Eq. (8) is no longer consistent. As a result, the estimation of $g(u)$ based on Eq. (8) is not appropriate. To obtain a consistent estimator for β , an instrumental variables estimator should be used [see Robinson (1988) and Liviatan (1963) for some related discussion]. Suppose that z_t is an appropriate instrumental variable of x_t^* , then β can be consistently estimated by

$$\hat{\beta}_{IV} = \left[\sum z_t \hat{x}_t^{*'} \right]^{-1} \left[\sum z_t \hat{y}_t^* \right], \quad (10)$$

and the estimation of $g(u)$ can then be constructed based on $\hat{\beta}_{IV}$.

We use the log difference of the total trading volume for the stocks in the S&P 100 Index as an instrument for volatility. Empirical evidence indicates that trading volume and volatility are positively correlated [see, for example, Chen et al. (2001)]. We use the log difference of the total trading volume as an instrument because the volume series is generally not stationary but the log difference series is. In Section 3, we test for stationarity for both series.

2.3. Bandwidth selection

The nonparametric estimates that are used in the partially linear regression estimation entail a choice of bandwidth h . As long as the bandwidth parameter h converges to zero at an appropriate rate, all such semiparametric estimators are (first order) asymptotically equivalent (i.e., the estimators are \sqrt{n} -consistent and have the same asymptotic normal distribution). However, estimates from finite samples can vary considerably with bandwidth choice. The finite sample performance of estimators can be quite satisfactory for some bandwidth choices, but it is also easy to find examples where the estimators have poor performance. Thus, it is useful to have rules that help us select bandwidth in an appropriately defined “optimal” way.

Although the bandwidth parameter does not appear in the first order asymptotics, it plays an important role in higher order asymptotics. Linton (1995) derived a second order expansion for the partially linear regression estimator and proposed a feasible method for optimal bandwidth selection in these models. The intuition of such an optimal bandwidth selection is as follows: the standardized estimator $\Sigma^{1/2} \sqrt{n}(\hat{\beta} - \beta)$, where Σ is the covariance matrix of the corresponding estimator, which is asymptotically standard normal, can be represented by a stochastic Taylor expansion to the second order as

$$\Sigma^{1/2} \sqrt{n}(\hat{\beta} - \beta) = \Phi + \sqrt{nh}^{2q} B + \frac{1}{\sqrt{nh}} V + O_{pp} \left(\sqrt{nh}^{2q} + \frac{1}{\sqrt{nh}} \right),$$

where Φ is distributed as $N(0, I)$, B and V are $O(1)$ and $O_p(1)$ quantities [$O_p(\cdot)$ signifies order in probability], $\sqrt{nh}^{2q} B$ is the leading term of nonparametric estimation bias and

$\frac{1}{\sqrt{nh}}V$ is the leading term of nonparametric estimation variance, q is the characteristic component determined by the kernel as defined in the above section. Notice that since the bandwidth h converges to zero at a rate such that $\sqrt{nh}h^{2q} \rightarrow 0$ and $nh \rightarrow \infty$, both $\sqrt{nh}h^{2q}B$ and $\frac{1}{\sqrt{nh}}V$ converge to zero and thus $\Sigma^{1/2}\sqrt{n}(\hat{\beta} - \beta)$ is asymptotically standard normal. Consequently, it can be shown that the mean squared error of $\hat{\beta}$ (standardized by Σ) has the following moment expansion

$$\text{MSE}(\hat{\beta}) = \frac{1}{n} \left\{ 1 + nh^{4q}\mathcal{B} + \frac{1}{nh}\mathcal{V} + O\left(nh^{4q} + \frac{1}{nh}\right) \right\},$$

where $nh^{4q}\mathcal{B}$ and $(1/nh)\mathcal{V}$ are squared bias and variance, respectively [see Linton (1995) for the exact formulae of these quantities]. Notice that the bandwidth parameter h surfaces in the second order effect $nh^{4q}\mathcal{B} + (1/nh)\mathcal{V}$ and an optimal bandwidth choice may be defined by minimizing the truncated mean squared error

$$\frac{1}{n} \left\{ 1 + nh^{4q}\mathcal{B} + \frac{1}{nh}\mathcal{V} \right\}$$

with respect to h . This bandwidth balances the order of magnitude of variance and squared bias so that nh^{4q} and $(1/nh)$ are of the same order of magnitude. Linton (1995) shows that the optimal bandwidth is

$$h^* = [\mathcal{V}/(4q\mathcal{B})]^{1/[4q+1]} n^{-2/[4q+1]}.$$

The quantities \mathcal{B} and \mathcal{V} are functions of nonparametric bias and variance and are generally estimated by “plug-in” methods. Linton suggested the following two-step bandwidth selection procedure for the partially linear regression models: firstly estimate \mathcal{B} and \mathcal{V} (and thus $V/(4qB)$) by their sample analogs from a preliminary estimation and denote the corresponding estimator by $\hat{\mathcal{V}}/(4q\hat{\mathcal{B}})$. Then $[\hat{\mathcal{V}}/(4q\hat{\mathcal{B}})]^{1/[4q+1]} n^{-2/[4q+1]}$ defines the window used in the partially linear regression. [See Appendix A for formulae of these estimators; also see Linton (1995) (Section 5) for a detailed discussion on the data-dependent bandwidth selection and Hardle et al. (1992) for similar methods].

3. Empirical estimation and results

3.1. Data description

We conduct our estimation using the Chicago Board of Options Exchange’s Market Volatility Index (VIX) dataset and the corresponding returns data on the underlying S&P 100 Index. The sample period is from January 1986 to December 1999, and the analysis is done on both the weekly and daily frequencies. Total trading volume for the stocks in the S&P 100 Index was also collected for the sample period and the log difference series

is used as an instrument for implied volatility. The source of the data is the Chicago Board Options Exchange and the on-line data service Datastream.

The VIX series is constructed by Whaley (1993, 2000) for the Chicago Board Options Exchange as a measure of “investor fear.” The series is an implied volatility computed from the prices of options on the S&P 100 Index (OEX options). Given that option prices are often quoted by the implied volatility and that the OEX options are the most actively traded options contracts, the VIX to a large extent represents investors’ consensus view about expected future stock market volatility. Specifically, the VIX is computed on a minute-by-minute basis from the implied volatilities of the eight near-the-money, nearby, and second nearby OEX option series. These implied volatilities are then weighted in such a manner that the VIX represents the implied volatility of a 30-calendar day (22-trading day) at-the-money OEX option.

Several authors have documented that the Black–Scholes implied volatilities from short-dated, at-the-money option prices approximate conditional volatilities well, even when volatility is stochastic and there is correlation between volatility and stock price (Latane and Rendleman, 1976; Chiras and Manaster, 1978; Beckers, 1981). Although the Black–Scholes model often produces biased prices for out-of-the-money options, the pricing of near-the-money options provides good approximations, even when the true underlying price follows a stochastic volatility process (Sheikh, 1991; Bakshi et al., 1997). This result provides an easy method to back out a time-series of conditional volatilities for empirical analyses and tests. Certainly, this conditional volatility series is unlikely to be identical to the conditional volatility series generated by models based on the historical




Fig. 1. Return and implied volatility (VIX) of the S&P 100 Index. This figure shows the returns and (annualized) implied volatility (VIX) of the S&P 100 stock index at weekly (upper panel) and daily frequencies. The sample is from January 1986 to December 1999. Dashed line plots the volatility series.

prices of the underlying index, such as the ARCH-class of models. Yet we believe that this implied volatility time series contains valuable information that may be useful in guiding our choice of ARCH models.

Fig. 1 shows the weekly and daily returns of the S&P 100 Index and the annualized VIX series. We see that there is much variation in implied volatility and that it mostly moves within the 10% to 30% range. One prominent exception, however, is the 1987 market crash when implied volatility reached a record high of 81.24% for the weekly data and 150.19% for the daily data. Volatility is also higher in 1990 and early 1991, possibly due to the Gulf War, and at the end of 1998 due to financial crises in Asia and Russia. The return series shows a similar pattern, with relatively large negative return movement during these volatile periods. In particular, during the 1987 market crash, we observe a negative return of 15.80% for the weekly data and 23.69% for the daily data. Over the entire weekly sample, there are 11 observations with negative returns larger than 5% in magnitude. There are seven such observations for the daily sample.

The annualized VIX series was divided by the square root of 52 to obtain the weekly series and by the square root of 260 to obtain the daily volatility series. Table 1 reports some summary statistics of the volatility, return and trading volume data at weekly and daily frequencies. $AC(k)$ denotes autocorrelation of order k . $Corr(1)$ is the sample correlation between returns and the VIX series. $Corr(2)$ is the sample correlation between the VIX series and the log difference of the trading volume.

The mean returns for the weekly and daily series are both positive, indicating overall upward trend in the S&P 100 Index. The means of the log difference of volume is also

Table 1
Summary statistics of the data

	Weekly data			Daily data		
	Return	Volatility	Volume	Return	Volatility	Volume
Mean	0.0028	0.0278	0.0032	0.0006	0.0125	0.0004
Std. Dev.	0.0213	0.0099	0.2396	0.0110	0.0049	0.2408
Max	0.0669	0.1127	0.8722	0.0854	0.0931	1.3199
Min	-0.1580	0.0127	-0.8279	-0.2369	0.0056	-1.6650
Skew	-1.0672	2.4797	-0.1375	-3.1003	4.6371	-0.1794
Ex. Kurt	5.5448	13.7134	1.4365	66.8315	51.8265	3.5305
AC(1)	-0.0287	0.9083	-0.3534	-0.0211	0.9417	-0.3554
AC(2)	0.0266	0.8220	-0.1202	-0.0508	0.8921	-0.0869
AC(3)	-0.0190	0.7649	-0.0368	-0.0405	0.8756	-0.0112
AC(4)	-0.0097	0.7243	0.0269	-0.0398	0.8541	-0.0703
AC(5)	-0.0559	0.6937	0.0441	0.0226	0.8320	0.0850
AC(10)	-0.0025	0.6005	0.1232	0.0116	0.7240	0.1127
Corr(1)	-0.2427			-0.1584		
Corr(2)		0.1702			0.1210	

This table reports the summary statistics of returns, conditional volatility (VIX) of the S&P 100 Index, and daily log difference of total trading volume for the stocks in the S&P 100 Index at weekly and daily frequencies. $AC(k)$ denotes autocorrelation of order k . $Corr(1)$ is the sample correlation between returns and the VIX series. $Corr(2)$ is the sample correlation between the VIX series and the log difference of trading volume. The sample period is from January 1986 to December 1999. The source of the data is the the Chicago Board Options Exchange and the on-line data service Datastream.

positive, with a large standard error, suggesting that total trading volume for stock in the index is volatile and upward trending. It is interesting to note the negative skewness of the return and the positive skewness of the conditional variance. The former has been widely documented and volatility asymmetry can potentially account for this skewness. The latter reflects the fact that volatility often increases quickly due to large return shocks and decreases slowly. Such dynamics can be very well captured by the GARCH models which provide an autoregressive mechanism and allow return shocks to impact conditional volatility. The skewness for the S&P 100 Index from 1986 to 1999 is -1.067 for the weekly series and -3.100 for the daily series, which are on the same order as those for the S&P 500 futures prices from 1988 to 1993 reported by Bates (1997). The two return series on the S&P 100 Index also exhibit excess kurtosis, in particular the daily sample (66.832), which suggests a highly non-normal distribution. These sample moment results generally are consistent with documented evidence on stock returns as discussed in Campbell et al. (1997). Although there is little autocorrelation in the stock return, the slow decaying autocorrelation of the conditional variance is characteristic of an autoregressive process. Finally, the negative sample correlation between return and conditional variance demonstrates the existence of volatility asymmetry. The positive correlations between volatility and log difference of volume suggest that the log difference of volume may be an appropriate instrument.

3.2. Stationarity tests of the volume series

To avoid the errors-in-variable problem, we need an instrument in the estimation problem (10) to produce a consistent estimate of the linear parameter β . Given that the log difference of volume is correlated with volatility, the trading volume variable may act as an instrument for volatility. Since the measurement error associated with the estimation of the implied volatility is unlikely to be correlated with the trading volume, the use of the log difference of volume as an instrument can be justified.

We test the stationarity/nonstationarity property of both the level of trading volume and the log difference of the trading volume using the augmented Dickey–Fuller (ADF) unit root test (Dickey and Fuller, 1979) and the KPSS (Kwiatkowski et al., 1992) stationarity test. Recall that the ADF procedure considers the null hypothesis of a unit root, which is unlikely to be rejected unless there is strong evidence against it. The KPSS test considers the null hypothesis of stationarity against the unit root hypothesis. Thus it is useful to perform both tests because the null hypotheses are different.

The parametric ADF t -ratio test is based on a long autoregression where the lag length of the ADF regression should grow at a rate $O(n^{1/3})$. We choose the lag length using the BIC model selection criterion (Schwarz, 1978; Rissanen, 1978). The KPSS test is semiparametric and entails the choice of a bandwidth. According to Monte Carlo evidence in previous research, we consider three choices of bandwidth: $M1=[4(n/100)^{1/4}]$, $M2=[12(n/100)^{1/4}]$, $M3=[4(n/100)^{1/3}]$.

Table 2 reports the testing results for the ADF test. The 5% level critical value is -3.432 , and the unit root hypothesis should be rejected if the calculated statistic is smaller than this value. For the KPSS test, the 5% level critical value is 0.146 and the stationarity hypothesis should be rejected if the calculated statistic is larger. Both tests indicate that the

Table 2
Stationarity test of the volume series

	ADF test	KPSS test		
		M1	M2	M3
Volume	-2.949	0.318	0.297	0.277
Change of volume	-13.4354	0.0988	0.0924	0.0844

This table reports the stationarity test results for the total trading volume series and the log difference of the total trading volume series. The 5% critical value for the ADF test is -3.432, and the null hypothesis of a unit root is rejected if the test statistic is smaller (larger in magnitude). The 5% critical value for the KPSS test is 0.146, and the stationarity hypothesis is rejected if the statistic is larger.

volume series is nonstationary in levels, but stationary in the log differences. Therefore, we use the log difference of the total trading volume as an instrument instead of the volume series.

3.3. Parameter estimates and specification tests

Due to the nonparametric nature of model (2), the intercept term α_0 and the impact function $g(u_{t-1})$ cannot be separately identified. We are mainly interested in the impact function $g(u_{t-1})$. So we estimate $\alpha_0 + g(u_{t-1})$ together and constrain $g(0) = 0$. The following discussion on $g(u_{t-1})$ is presented with this constraint imposed.

Before we move on to examine $g(u_{t-1})$, let's take a look at the AR(1) coefficient α_1 in the volatility Eq. (2). Table 3 reports the estimated coefficient for both the weekly and daily data.

The coefficient is quite stable for different values of the bandwidths and shows that conditional volatility exhibits strong first order autocorrelation. This persistence has been documented by many time series models such as Duan (1995) and Nelson (1991). At the weekly frequency, the 0.7811 coefficient for the optimal bandwidth is consistent with the findings of most empirical studies using S&P 100 Index data. The standard error is fairly small at 0.0307, signifying that the autoregressive effect is strong and the coefficient can be estimated accurately. For the daily data, the AR(1) coefficient is slightly smaller. For the optimal bandwidth, α_1 is estimated to be 0.6191 with a standard error of 0.0401. The estimated coefficient is smaller than those estimated by ARCH models and may indicate the difference between daily volatility dynamics estimated by VIX and ARCH models. The fact that the volatility persistence is partially captured by $g(u_{t-1})$ in the generalized partially linear model could also explain the smaller estimates.² Overall, our stochastic volatility model (2) is able to capture the strong autocorrelation of volatilities.

To ensure that our generalized partially linear model is properly specified, we conduct a series of diagnostics tests. Two groups of tests are performed on the estimated residual $\hat{\varepsilon}_t$ from the partially linear volatility equation in order to check for autocorrelation and conditional heteroskedasticity.

² We thank the associate editor for pointing this out to us.

Table 3
Parameter estimates of the model

Weekly VIX Data			
Bandwidth	$h=0.0317$ (optimal)	$h=0.02$	$h=0.05$
α_1	0.7811	0.7886	0.7856
Standard deviation	0.0307	0.0390	0.0435
Daily VIX Data			
Bandwidth	$h=0.0262$ (optimal)	$h=0.01$	$h=0.04$
α_1	0.6191	0.6394	0.5990
Standard deviation	0.0401	0.0371	0.0426

This table reports the parameter estimates of the partial linear model for both the weekly data and the daily data of VIX index from January 1986 to December 1999. Standard errors are in parentheses.

The Box and Pierce (1970) Q -statistic is a test for autocorrelation in time series data. By summing the squared autocorrelations, the Box–Pierce test is designed to detect departure from zero autocorrelation in either directions and at all lags. Ljung and Box (1978) provide a finite sample correction that yields a better fit to the Q -statistic for small sample size. The selection of the number of autocorrelations (m) requires some care, and generally it is good practice to specify several of them. We test for autocorrelation with $m=3, 6$, and 12 . Under the null hypothesis of no autocorrelation, the test statistics are all distributed as χ_m^2 .

Table 4
Specification tests of the partially linear model

Bandwidth	Ljung–Box test			ARCH test		
	$m=3$	$m=6$	$m=12$	$q=3$	$q=6$	$q=12$
<i>Weekly VIX data</i>						
$h=0.0317$	7.987 (0.046)	9.834 (0.131)	12.325 (0.419)	7.498 (0.057)	9.622 (0.141)	11.700 (0.470)
$h=0.02$	7.467 (0.058)	9.433 (0.150)	11.539 (0.483)	7.119 (0.068)	9.289 (0.157)	10.843 (0.542)
$h=0.05$	9.006 (0.029)	10.134 (0.119)	13.295 (0.347)	8.598 (0.035)	9.529 (0.145)	12.300 (0.421)
<i>Daily VIX data</i>						
$h=0.0262$	7.150 (0.067)	7.529 (0.274)	14.234 (0.285)	6.755 (0.080)	7.270 (0.296)	12.922 (0.374)
$h=0.01$	6.074 (0.108)	6.471 (0.372)	12.496 (0.406)	5.664 (0.129)	6.220 (0.398)	11.402 (0.494)
$h=0.04$	7.487 (0.057)	8.136 (0.228)	15.089 (0.236)	7.075 (0.069)	7.898 (0.245)	13.686 (0.321)

This table reports specification results of the partially linear model for the weekly and daily VIX data. m is the number of autocorrelations for the Ljung–Box test. q is the number of the order of ARCH effects. All tests are distributed as chi-square with the degree of freedom equal to m or q . The p -values are in parentheses. The data are from January 1986 to December 1999.

In Table 4, we report the The Ljung–Box test statistics and their p -values. The upper panel reports the results for the weekly data, where only two of the nine statistics are significant at the 5% level (at 0.046 for $m=3$ and $h=0.0317$, and at 0.029 for $m=3$ and $h=0.05$). Therefore, the null hypothesis of non-autocorrelation can be generally rejected at the 5% level. The null hypothesis, however, cannot be rejected at the 1% level. The results for the daily data are reported in the lower panel of Table 4. All nine statistics are insignificant at the 5% level; hence, the null hypothesis of non-autocorrelation cannot be rejected at that level.

We test for heteroskedasticity using the Lagrange multiplier ARCH test from Engle (1982). The test is designed by Engle to detect the existence of q th order ARCH effects (see also Breusch and Pagan, 1980). We specify several orders q of the ARCH effect.

The results are also presented in Table 4. For the weekly data only one of the nine statistics is significant at the 5% level (at 0.035 for $q=3$ and $h=0.05$). For the daily data, all nine statistics are insignificant at the 5% level. We conclude that the null hypothesis of non-conditional heteroskedasticity in the residual series cannot be rejected. Overall, the specification tests indicate that our partially linear models are properly specified for both the weekly and daily data.

3.4. The nonparametric news impact curves

To interpret the estimates of the impact function $g(u_t)$ meaningfully, we need to derive confidence bands around the estimated function. Notice that although $\hat{\beta}_{IV}$ is \sqrt{n} -consistent, the convergence rate of $E(\widehat{y_t | u_{t-1}})$ (and $E(\widehat{x_t | u_t})$) is slower than the root- n convergence rate of $\hat{\beta}_{IV}$. By standard results of nonparametric estimation, $E(\widehat{y_t | u_{t-1}})$ and $E(\widehat{x_t | u_t})$ converge at rate $n^{1/2}h^{1/2}$, and thus $(\hat{\beta} - \beta)' E(\widehat{x_t | u_{t-1}})$ is of a smaller order of magnitude. If we denote $\zeta_t = (y_t, x_t)'$ and $\gamma = (1, -\beta')'$, then under regularity conditions,

$$\begin{aligned} & \sqrt{nh}[\tilde{g}(u_{t-1}) - g(u_{t-1})] \\ &= \sqrt{nh}[E(\widehat{y_t | u_{t-1}}) - E(y_t | u_{t-1})] \\ &\quad - \sqrt{nh}[\hat{\beta}' E(\widehat{x_t | u_{t-1}}) - \beta' E(x_t | u_{t-1})] \\ &\approx \gamma' \sqrt{nh}[E(\widehat{\zeta_t | u_{t-1}}) - E(\zeta_t | u_{t-1})] \\ &\Rightarrow N\left(0, f(u_{t-1})^{-1} \gamma' V_\zeta \gamma \int K(u)^2\right) \end{aligned}$$

where $V_\zeta = E[\zeta_t \zeta_t' | u_{t-1}] - E(\zeta_t | u_{t-1})E(\zeta_t | u_{t-1})'$ and $K(\cdot)$ is the kernel. Thus, the variance of $\sqrt{nh}[\tilde{g}(u_{t-1}) - g(u_{t-1})]$ can be asymptotically approximated by $\hat{f}(u_{t-1})^{-1} \hat{\gamma}' \hat{V}_\zeta \hat{\gamma} \int K(u)^2$ and a confidence band can be constructed using this approximation. Notice that the convergence rate of $\tilde{g}(u_t)$ is dominated by that of the nonparametric estimate and the confidence band of $\tilde{g}(u_t)$ is wider than those from parametric models.

Figs. 2 and 3 display the estimated function and associated 95% confidence bands for the weekly and daily data. They demonstrate that the function $g(u_t)$ displays a strong asymmetric effect and has a similar shape for the weekly and daily samples. For both the

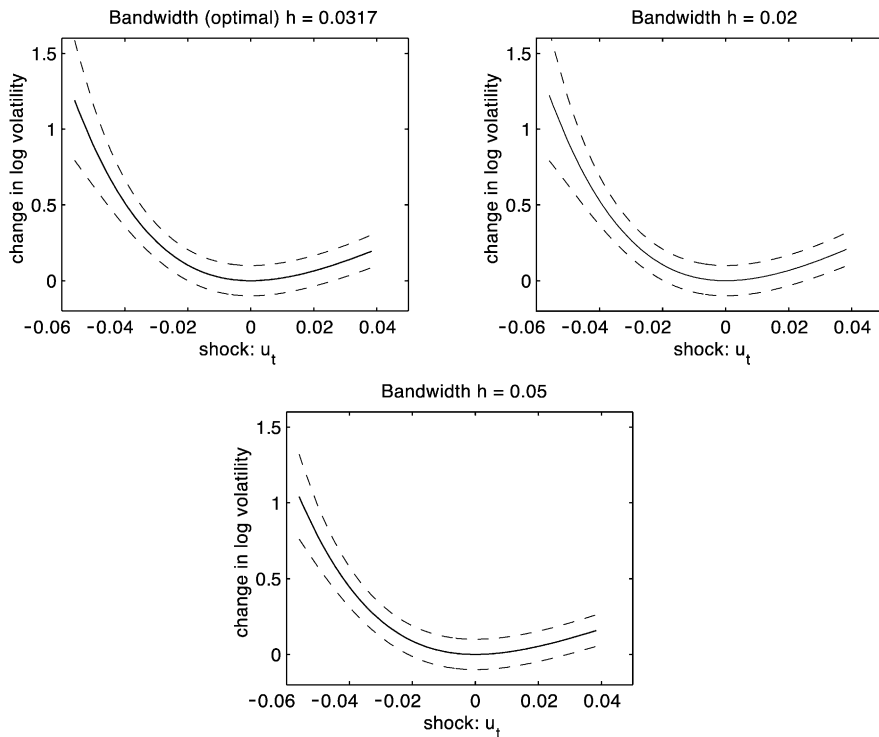


Fig. 2. Log volatility news impact curves for the weekly VIX data. This figure shows the volatility news impact curves estimated using a nonparametric approach. The three panels differ in the bandwidths chosen for the nonparametric estimation. The dashed lines denote the 95% confidence band.

weekly and daily data, we plot the log volatility news impact curves for the optimal bandwidth and two bandwidths around the optimal bandwidth.

Fig. 2 plots the news impact curves from the weekly data, using the optimal bandwidth 0.0317 and two other bandwidths of 0.02 and 0.05. The estimation results are qualitatively similar for the three different bandwidths. Fig. 3 plots the news impact curves from the daily data, using the optimal bandwidth 0.0262 and two other bandwidths of 0.01 and 0.04. For both the weekly and daily data we have used a wide range of different bandwidths and found that the choice of the bandwidth does not change our results qualitatively.

It is important to note that log changes in volatility represent relative (or percentage) changes in volatility. For the weekly data we see that for positive return shocks and negative return shocks with magnitude less than 2%, conditional volatility does not change very much. However, as the magnitude of a negative shock increases, the impact on conditional volatility increases dramatically. For a 3% negative return shock, the impact on log volatility is about 0.25, which translates into about 25% change from the current level of volatility. For example, if the current annualized volatility level is about 20%, the volatility would increase to 25% with the shock. When the size of the negative return

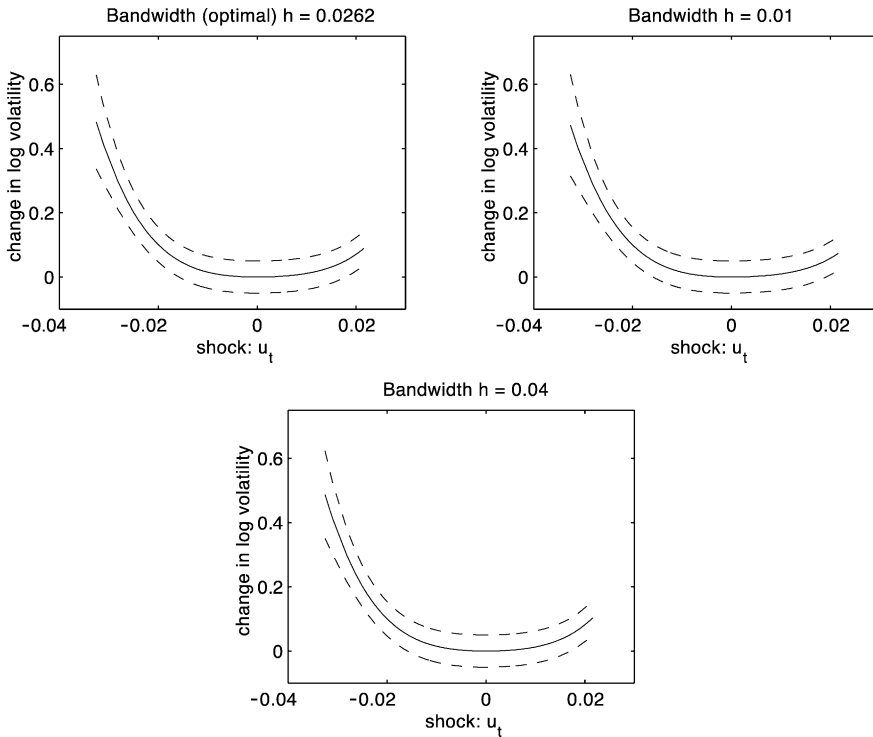


Fig. 3. Log volatility news impact curves for the daily VIX data. This figure shows the volatility news impact curves estimated using a nonparametric approach. The three panels differ in the bandwidths chosen for the nonparametric estimation. The dashed lines denote the 95% confidence band.

shock is 5%, conditional log volatility will increase by about 1, or about double the current level of volatility. A 20% annualized volatility would increase to a level of about 40%. It is likely that a larger negative return shock yields an even bigger impact on conditional volatility.

For the daily data we see that for positive return shocks and negative return shocks with magnitude less than 1%, conditional volatility does not change very much. However, as the magnitude of a negative shock increases, its impact on conditional volatility increases dramatically. For a 2% negative return shock, its impact on log volatility is about 0.1, which translates into about 10% change from the current level of volatility. For example, if the current annualized volatility level is about 20%, the volatility would increase to 22% with the shock. When the size of the negative return shock is 3%, conditional log volatility would increase by about 0.5, or about 50% change from the current level of volatility. A 20% annualized volatility would increase to a level of about 30%. Although a larger negative return shock may yield an even bigger impact on conditional volatility, we do not have enough sample points in the tail to be able to make a prediction at a very high confidence level.

In summary, the news impact curves estimated by the VIX data support the asymmetric volatility findings based on returns data. Large negative return shocks tend to have the most impact on conditional volatility as measured by implied volatility, while smaller negative return shocks tend to increase conditional volatility only slightly. These results are consistent with those from the ARCH models (see Nelson, 1991; Glosten et al., 1993; Hentschel, 1995; Bekaert and Wu, 2000). Positive return shocks, large or small, do not seem to have much impact on conditional volatility.

This result could explain an interesting phenomenon in empirical research on volatility dynamics: while there is much consensus that negative return shocks are associated with increases in conditional volatility, researchers disagree over the role positive return shocks play in volatility dynamics. Whereas a large positive return shock may lead to a higher stock price—hence a lower leverage ratio and lower volatility—either positive or negative large movements may lead to increases in conditional volatility. Whether large positive return shocks increase or decrease volatility depends on the net effect of these two forces. This could explain why $g(u_t)$ does not change much for $u_t > 0$ in Figs. 2 and 3. Furthermore, this result is consistent with high volatility persistence during market turmoils. When the stock market experiences a sharp decline, this large negative return shock usually leads to very high market volatility. When the stock market rebounds, the large positive return does not seem to reduce volatility immediately. Thus high volatility could not be reduced quickly with large positive return shocks, as reflected in the high AR(1) coefficient reported above.

4. Implications for EGARCH models

We now wish to determine if the estimated news impact function $g(u_t)$ can help specify ARCH models. Even though volatilities implied by option prices are often linked to estimated volatilities from ARCH models, it is not clear that the news impact curves implied by the VIX data can be used to specify appropriate ARCH models (see Latane and Rendleman, 1976; Chiras and Manaster, 1978; Beckers, 1981, for properties of implied volatility). In general, the implied volatility series will not be identical to the conditional volatility series generated by models based on historical returns, such as the ARCH-class of models. But if there is a close relationship between the implied volatility process and that of a well-specified ARCH model, then the analysis of the implied volatility will be helpful in the choice of ARCH models.

In this section, we examine various EGARCH models inspired by our partially linear analysis. Unfortunately we cannot include the nonparametrically estimated function $g(u_t)$ in the ARCH model. The best alternative is to specify parametric functions that are close to the nonparametrically estimated $g(u_t)$. A general high order polynomial will not work since some specifications may imply explosive volatility processes. So we use the nonparametrically estimated $g(u_t)$ to specify several plausible EGARCH models and test the performance of these models. We also compare the

implied news impact curves of these models with the nonparametrically estimated $g(u_t)$ function.

4.1. Estimation of EGARCH models

We estimate seven exponential GARCH models using weekly and daily S&P 100 Index returns from January 1986 to December 1999. The models have the same AR(1) mean equation,

$$r_t = \omega_0 + \omega_1 r_{t-1} + u_t = \omega_0 + \omega_1 r_{t-1} + \sigma_t \xi_t \quad \xi_t \sim N(0, 1), \quad (11)$$

but differ in how the return shocks are specified. We consider different specifications of volatility to allow for an appropriate type of shock in various applications, and to contrast the empirical results of the models. For example, some models may allow for infinite variance while others do not. The seven EGARCH models are

$$\text{Model 1 : } \ln \sigma_t = \gamma_0 + \gamma_1 \ln \sigma_{t-1} + \gamma_2 \xi_{t-1} + \gamma_3 |\xi_{t-1}|, \quad (12)$$

$$\text{Model 2 : } \ln \sigma_t = \gamma_0 + \gamma_1 \ln \sigma_{t-1} + \gamma_2 u_{t-1} + \gamma_3 |u_{t-1}|, \quad (13)$$

$$\text{Model 3 : } \ln \sigma_t = \gamma_0 + \gamma_1 \ln \sigma_{t-1} + \gamma_2 \max(0, -\xi_{t-1})^2 + \gamma_3 \max(0, \xi_{t-1})^2, \quad (14)$$

$$\text{Model 4 : } \ln \sigma_t = \gamma_0 + \gamma_1 \ln \sigma_{t-1} + \gamma_2 \max(0, -\xi_{t-1}) + \gamma_3 \max(0, \xi_{t-1}), \quad (15)$$

$$\text{Model 5 : } \ln \sigma_t = \gamma_0 + \gamma_1 \ln \sigma_{t-1} + \gamma_2 \max(0, -u_{t-1})^2 + \gamma_3 \max(0, u_{t-1})^2, \quad (16)$$

$$\text{Model 6 : } \ln \sigma_t = \gamma_0 + \gamma_1 \ln \sigma_{t-1} + \gamma_2 \max(0, -u_{t-1}) + \gamma_3 \max(0, u_{t-1}), \quad (17)$$

$$\begin{aligned} \text{Model 7 : } \ln \sigma_t = & \gamma_0 + \gamma_1 \ln \sigma_{t-1} + \gamma_2 \max(0, -u_{t-1}) + \gamma_3 \max(0, u_{t-1}) \\ & + \gamma_4 \max(0, -u_{t-1} + \theta). \end{aligned} \quad (18)$$

In the estimation of the EGARCH models, we use the quasi-maximum likelihood method (Bollerslev and Wooldridge, 1992), which is robust to nonnormal innovations, although a Gaussian likelihood was maximized. We calculate the robust standard errors (see White, 1982; Bollerslev and Wooldridge, 1992) and conduct specification tests of each model.

Tables 5 and 6 report the parameter estimates of these seven models for the weekly and daily data, respectively. We first consider the parameter estimates for the AR(1) mean Eq. (11). For the weekly data, the estimates of the constant term ω_0 are similar across the seven models and statistically significant. The maximum is 0.0033 and the minimum is 0.0024.

Table 5
Estimated parameters of the weekly exponential ARCH models

	ω_0	ω_1	γ_0	γ_1	γ_2	γ_3	γ_4
Model 1	0.0025 (0.0005)	-0.0736 (0.0293)	-0.6238 (0.1622)	0.9394 (0.0191)	-0.1154 (0.0259)	0.1966 (0.0296)	
Model 2	0.0027 (0.0005)	-0.0646 (0.0296)	-1.3261 (0.2923)	0.8491 (0.0351)	-6.3557 (1.5464)	9.3712 (1.4994)	
Model 3	0.0029 (0.0005)	-0.0814 (0.0299)	-0.7128 (0.2529)	0.9176 (0.0307)	0.1058 (0.0280)	0.0140 (0.0194)	
Model 4	0.0029 (0.0005)	-0.0801 (0.0299)	-0.7785 (0.2278)	0.9088 (0.0281)	0.1136 (0.0240)	-0.0052 (0.0299)	
Model 5	0.0029 (0.0005)	-0.0423 (0.0288)	-0.8633 (0.3156)	0.8936 (0.0399)	126.41 (39.302)	5.3009 (43.862)	
Model 6	0.0033 (0.0005)	-0.0406 (0.0304)	-2.4294 (1.5403)	0.7022 (0.1908)	288.55 (147.94)	4.6351 (3.6713)	
Model 7	0.0024 (0.0005)	-0.0543 (0.0295)	-1.3679 (0.2614)	0.8499 (0.0307)	8.803 (3.6475)	4.3849 (1.7201)	19.250 (5.2716)

This table reports the estimated parameters of the seven exponential ARCH models for the weekly data. Standard errors are in parentheses. The data used are the daily S&P 100 Index returns from January 1986 to December 1999.

The estimates of the AR(1) coefficient ω_1 are negative, either statistically insignificant or weakly significant, indicating overall weak negative autocorrelation in weekly returns. For the daily data, the estimates of the constant term ω_0 are also similar across the seven models and statistically significant. The maximum is 0.0008 and the minimum is 0.0004. The estimates of the AR(1) coefficient ω_1 are negative, all statistically insignificant except for Model 6 (with a t -ratio of 1.99), indicating overall weak negative autocorrelation in daily returns as well. The negative estimates of AR(1) coefficient ω_1 at the weekly and

Table 6
Estimated parameters of the daily exponential ARCH models

	ω_0	ω_1	γ_0	γ_1	γ_2	γ_3	γ_4
Model 1	0.0004 (0.0001)	-0.0173 (0.0130)	-0.4015 (0.0442)	0.9701 (0.0040)	-0.0912 (0.0080)	0.1692 (0.0129)	
Model 2	0.0006 (0.0001)	-0.0150 (0.0127)	-0.9413 (0.0707)	0.9110 (0.0068)	-9.2991 (0.7499)	15.030 (1.0789)	
Model 3	0.0006 (0.0001)	-0.0090 (0.0138)	-0.5980 (0.0669)	0.9429 (0.0067)	0.0777 (0.0061)	0.0494 (0.0065)	
Model 4	0.0007 (0.0001)	-0.0136 (0.0134)	-0.5844 (0.0743)	0.9451 (0.0074)	0.0742 (0.0062)	0.0787 (0.0129)	
Model 5	0.0008 (0.0001)	-0.0245 (0.0130)	-1.8180 (0.2172)	0.8084 (0.0234)	631.67 (51.022)	4.3300 (98.762)	
Model 6	0.0008 (0.0001)	-0.0259 (0.0130)	-1.5534 (0.1438)	0.8400 (0.0151)	572.77 (40.396)	8.2700 (1.4370)	
Model 7	0.0004 (0.0001)	-0.0054 (0.0127)	-0.8067 (0.0627)	0.9228 (0.0063)	23.439 (1.2917)	0.0492 (0.5072)	0.3852 (3.1953)

This table reports the estimated parameters of the seven exponential ARCH models for the daily data. Standard errors are in parentheses. The data used are the daily S&P 100 Index returns from January 1986 to December 1999.

daily level can be attributed to market microstructure effects such as the bid–ask spread and non-synchronous trading (see Roll, 1984; Lo and MacKinlay, 1990; Campbell et al., 1997).

We now examine the seven volatility equations.

4.1.1. Model 1

Model 1 is the exponential ARCH model of Nelson (1991), where the residuals in the volatility equation are standardized with a mean of zero and variance of one. We use this model as a benchmark for our analysis. The term γ_2 , if negative, captures the asymmetric effect of negative and positive return shocks. Indeed, γ_2 is estimated to be -0.1154 with a standard error of (0.0259) for the weekly data, and -0.0912 (0.0080) for the daily data.³ This is not surprising since the model was designed to capture the asymmetric volatility effect. The estimate of γ_1 is large, at 0.9394 (0.0191) for the weekly data and 0.9701 (0.0040) for the daily data, suggesting that volatility persistence could be strong. All estimated coefficients are statistically significant at both the weekly and daily frequency.

4.1.2. Model 2

Model 2 differs from Model 1 in that the general return shock u_t is used in the volatility equation instead of the standardized residual ξ_t . We consider this model since the partially linear model specifies the news impact function with respect to the general residuals, not standardized residuals. In this model, γ_2 is estimated to be -6.3557 (1.5464) for the weekly data and -9.2991 (0.7499) for the daily data, and both estimates are statistically significant. Hence, the EGARCH model with non-standardized residuals can also capture the asymmetry effect. Note that the estimated coefficients are much larger since they are multiplied by smaller, non-standardized residuals. The estimate of the coefficient γ_1 is large and statistically significant, at 0.8491 (0.0351) for the weekly data and 0.9110 (0.0068) for the daily data. Interestingly, the use of non-standardized residuals seems to reduce the size of the estimated coefficient for the lagged log volatility. For example, the weekly γ_1 , estimated at 0.8491 in this model, is much closer to the corresponding coefficient in the partially linear model, estimated at 0.7811 , than γ_1 in Model 1, estimated at 0.9394 . This could be due to the fact that some of the persistence in volatility is captured in $\gamma_2 u_{t-1} + \gamma_3 |u_{t-1}|$. To see this, note that this term can be written as $(\gamma_2 \xi_{t-1} + \gamma_3 |\xi_{t-1}|) \sigma_{t-1}$.

4.1.3. Models 3 and 4

Models 3 and 4 specify a quadratic function for negative standardized shocks which captures the increasing impact of larger return shocks (see Figs. 2 and 3). In Model 3, asymmetric volatility is present when γ_2 is larger than γ_3 . Indeed, for the weekly data γ_2 is estimated to be 0.1058 with a standard error of (0.0280) while γ_3 is estimated to be 0.0140 with a standard error of (0.0194) . For the daily data, γ_2 is estimated to be 0.0777 (0.0061) while γ_3 is estimated to be 0.0494 (0.0065) . Hence, the quadratic EGARCH model with standardized residuals can capture the asymmetry effect. All estimated

³ Standard errors are reported in parentheses.

coefficients are statistically significant at both the weekly and daily frequency, except the estimated γ_3 for the weekly data. The estimate of γ_1 is large, at 0.9176 (0.0307) for the weekly data and 0.9429 (0.0067) for the daily data, again suggesting that volatility persistence is strong.

Model 4 allows the negative return shocks to increase log volatility quadratically and positive return shocks only linearly, so that asymmetric volatility is present when γ_2 is large and γ_3 is close to zero or negative. For the weekly data γ_2 is estimated to be statistically significant at 0.1136 (0.0240) while γ_3 is estimated to be -0.0052 (0.0299). For the daily data, γ_2 is estimated to be 0.0742 (0.0062) while γ_3 is estimated to be 0.0787 (0.0129). Thus, at weekly frequency, Model 4 demonstrates asymmetric volatility, while at daily frequency it does not. Note again that all estimates are statistically significant at both the weekly and daily frequencies except that of γ_3 for the weekly data.

4.1.4. Models 5 and 6

Models 5 and 6 are similar to Models 3 and 4, but general return shocks are used in place of standardized innovations. In Model 5, asymmetric volatility is present when γ_2 is larger than γ_3 , and we estimate γ_2 to be 126.41 (39.302) and γ_3 to be 5.3009 (43.862) in the weekly data. In the daily data, γ_2 is estimated to be 631.67 (51.022) and γ_3 is estimated to be 4.3300 (98.762). Thus, the quadratic EGARCH model with non-standardized residuals (Model 5) can capture the asymmetry effect. The estimated coefficients for γ_2 and γ_3 are much larger than those in Models 3 and 4 since they are multiplied by smaller, non-standardized residuals. The estimate of γ_1 is large, at 0.8936 (0.0399) in the weekly data and 0.8084 (0.0234) in the daily data, suggesting volatility persistence could be strong. However, the use of non-standardized residuals appears to reduce the size of the estimate, compared with the estimate of γ_1 in Model 3. The estimate for γ_1 in Model 3 could be “biased” downward for the same reason as in Model 2 described above. All estimates are statistically significant at both the weekly and daily frequency except that of γ_3 .

In Model 6, asymmetric volatility is present when γ_2 is large and γ_3 is close to zero or negative. We estimate γ_2 to be statistically significant at 288.55 (147.94) while the estimate of γ_3 is insignificant at 4.6351 (3.6713) in the weekly data. In the daily data, γ_2 is estimated to be 572.77 (40.396) while γ_3 is estimated to be 8.2700 (1.4370). All estimates are statistically significant at both the weekly and daily frequency except that of γ_3 for the weekly data. Again, it seems that the use of non-standardized residuals reduces the size of the estimate, compared with Model 4 due to the presence of σ_{t-1} in u_{t-1} . In general, the specification of the shock terms has an impact on the size of the estimate of γ_1 , as we compare the estimates from Models 2, 5 and 6. The weekly estimate of γ_1 in Model 6 is even smaller than 0.7811, the corresponding coefficient in the partially linear model estimated from the VIX data. In general, γ_1 is not invariant to the specification of the shock function since γ_1 by itself is not sufficient to measure the persistence of volatility.

4.1.5. Model 7

Model 7 is a special case of the Engle and Ng (1993) partially nonparametric ARCH model, whereby shocks larger than θ have an additional linear impact on conditional

volatility. Since the quadratic functions in Models 3 through 6 may imply that extreme shocks have an excessive impact on volatility, a piecewise linear function may fit the data better. Based on the estimated function $g(u_t)$ from the partially linear model, we set θ to be -0.03 for the weekly data and -0.02 for the daily data. The coefficient γ_4 measures the additional impact that a negative return shock larger than θ may have in predicting volatility.

For the weekly data, γ_2 is estimated to be 8.803 (3.6475) and γ_3 is estimated to be 4.3849 (1.7201), implying that even for negative returns below θ there is asymmetric volatility effect. γ_4 is estimated to be 19.25 (5.2716), so that negative return shocks above -3% increase volatility further. For the daily data, γ_2 is estimated to be 23.439 (1.2917) and γ_3 is estimated to be 0.0492 (0.5072), indicating that the asymmetric volatility effect is very strong. However, γ_4 is estimated to be statistically insignificant at 0.3852 (3.1953). It seems that at the daily frequency, large negative returns do not increase volatility additionally.

Overall, an important feature of the empirical results is that the estimated γ_1 in the seven models are generally larger than the estimated α_1 in the partial linear model (and other nonlinear models). This is due to the fact that, in the partial linear model, some of the




Fig. 4. Implied news impact curves for weekly EGARCH models. This figure shows the volatility news impact curves for the weekly EGARCH models. The news impact curve from the partially linear model for the weekly VIX series is also plotted for comparison.

persistence in volatility are captured by the nonlinear function $g(\cdot)$. The partial linear model decomposes the volatility into two components and the nonlinear persistent behavior in volatility is contained in the second component. Such a fact is important in the presence of asymmetric persistency. In particular, the second component in the partial linear model provides a more accurate persistence measure that reflects the asymmetric dynamics in the implied volatility.

4.2. Implied news impact curves

To better understand how the EGARCH models capture asymmetric volatility, we plot the implied news impact curves for the EGARCH models. The weekly and daily news impact curves are plotted in Figs. 4 and 5, respectively. We only plot the news impact curves for Models 2, 5, 6, and 7, since the functional form of Models 2, 5, and 6 is identical to that of Models 1, 3 and 4, respectively. We plot the estimated news impact function $g(u_{t-1})$ implied by the VIX data for comparison.

Fig. 4 shows that the impact functions are very similar for positive return shocks at the weekly frequency—positive shocks in all models increase volatility only marginally. For




Fig. 5. Implied news impact curves for daily EGARCH models. This figure shows the volatility news impact curves for the daily EGARCH models. The news impact curve from the partially linear model for the daily VIX series is also plotted for comparison.

negative and positive return shocks, Models 6 and 7 produce results very similar to the VIX model. Model 5, however, under-predicts the increase in volatility for large negative return shocks, compared to the VIX model. Model 2, the linear EGARCH model, over-predicts increases in volatility for small negative return shocks, yet under-predicts increases in volatility for large negative return shocks, compared to the VIX model and Models 6 and 7.

Fig. 5 shows that for positive return shocks at the daily frequency, the impact functions are quite different from each other. Models 5 and 7 do not predict an increase in volatility at all for positive return shocks, whereas Models 2 and 6 predict increases in volatility larger than those from the VIX model. For negative return shocks, all models predict increases in volatility larger than those from the VIX model. Models 5 and 6 are very close to each other, and Model 7 is almost identical to the linear Model 2. These results indicate that the daily EGARCH models are more difficult to estimate. This could be due to the fact that at daily frequency, EGARCH model innovations are further away from conditional normality as shown below in the specification tests.

To provide another perspective on the volatility behavior of the models, we plot the news impact curves as in Engle and Ng (1993) [For more details, see Table 8 and




Fig. 6. Engle–Ng news impact curves for weekly EGARCH models. This figure shows the Engle and Ng (1993) volatility news impact curves for the weekly EGARCH models. The news impact curves from the partially linear model for the weekly VIX series and the GJR–GARCH model are also plotted for comparison. Volatility at time $t-1$ is assumed to be 2.78% for all models.




Fig. 7. Engle–Ng news impact curves for daily EGARCH models. This figure shows the Engle and Ng (1993) volatility news impact curves for the daily EGARCH models. The news impact curves from the partially linear model for the daily VIX series and the GJR–GARCH are also plotted for comparison. Volatility at time $t-1$ is assumed to be 1.25% for all models.

associated discussion in Engle and Ng, pp. 1773]. The conditional volatility at time $t-1$ is assumed to be the same for all models and return shocks are applied to the model to yield the predicted volatility at time t . We hold the time $t-1$ volatility to be the sample average of the VIX index: 2.78% for the weekly data and 1.25% for the daily data, corresponding to 20.43% and 20.16% on an annual basis. Figs. 6 and 7 plot the resulting volatility level for the EGARCH models. The news impact curves from the partially linear model for the VIX series and from the GJR–GARCH model are also plotted for comparison.

For the weekly data, Fig. 6 shows that Model 7 matches the response function of the partially linear model (VIX) most closely, at least for the range of return shocks plotted, which is consistent with our findings above. For negative shocks larger than 6% in size, Models 5 and 6 will likely generate very high increases in volatility, given the quadratic functional form. This may not be appropriate for return shocks larger than 10% since it will imply unreasonably high volatility. Since the partially linear model is based upon non-parametric approximation of the function, we cannot draw conclusions on volatility response for large negative shocks. The GJR–GARCH model is pretty flat, which reflects a large coefficient estimated for the lagged variance (0.9516). Since this number is

unusually high for the weekly data, the GJR-GARCH response function should be interpreted with care.

For the daily data, none of the models shown matches the partially linear VIX model well. The GJR-GARCH model seems to be the only one that is close to matching the VIX model for both the positive and negative shocks. The results in Fig. 7 further indicate that the partially linear model does not seem to help in specifying GARCH models.

4.3. Specification tests

To ensure that the seven EGARCH models are properly specified at the weekly and daily frequencies, we conduct specification tests on the standardized residuals of these models. Recall that we used the quasi-maximum likelihood method in estimating the models (Bollerslev and Wooldridge, 1992). Although the method is robust to nonnormal innovations, we conducted a normality test on the standardized residuals to check to what extent normality was violated. The results in Tables 7 and 8 indicate that there is weak evidence against normality in the weekly data, but stronger evidence in the daily data. The p -values ranges from 0.001 for Model 5 to 0.036 for Model 1 at weekly frequency. For the daily data, there is strong evidence against the null hypothesis of normality since the p -values for all models are around 0.007, smaller than the 0.01 threshold.

Table 7
Specification tests of the weekly exponential ARCH models

	Normality test	Ljung–Box test			ARCH test		
		$m=3$	$m=6$	$m=12$	$q=3$	$q=6$	$q=12$
Model 1	6.672 (0.036)	1.931 (0.586)	2.097 (0.910)	7.932 (0.790)	1.902 (0.592)	2.051 (0.914)	7.719 (0.806)
Model 2	6.9241 (0.031)	1.571 (0.665)	1.848 (0.933)	8.612 (0.735)	1.555 (0.669)	1.760 (0.940)	8.571 (0.739)
Model 3	7.049 (0.030)	2.033 (0.565)	2.094 (0.910)	6.352 (0.897)	2.001 (0.572)	2.026 (0.917)	6.220 (0.904)
Model 4	7.159 (0.028)	1.923 (0.588)	2.015 (0.918)	7.077 (0.852)	1.895 (0.594)	1.935 (0.925)	6.956 (0.860)
Model 5	13.719 (0.001)	1.885 (0.596)	2.386 (0.880)	6.583 (0.883)	1.891 (0.595)	2.399 (0.879)	6.782 (0.871)
Model 6	7.491 (0.024)	1.159 (0.762)	1.682 (0.946)	11.398 (0.495)	1.145 (0.766)	1.606 (0.952)	11.436 (0.491)
Model 7	7.507 (0.023)	1.005 (0.799)	1.258 (0.973)	10.366 (0.583)	1.003 (0.800)	1.215 (0.976)	10.185 (0.599)

This table reports specification results of the seven exponential ARCH models for the weekly data. The test is performed on the standardized residuals. The normality test is a Wald test for the skewness and excess kurtosis coefficients to be jointly equal to zero. The p -values are based on the chi-square distribution with 2 degrees of freedom. m is the number of autocorrelations for the Ljung–Box test. q is the number of the order of ARCH effects. These tests are distributed as chi-square with the degrees of freedom equal to m or q . The p -values are in parentheses. The data used are the weekly S&P 100 Index returns from January 1986 to December 1999.

Table 8
Specification tests of the daily exponential ARCH models

	Normality test	Ljung–Box test			ARCH test		
		$m=3$	$m=6$	$m=12$	$q=3$	$q=6$	$q=12$
Model 1	9.932 (0.007)	1.488 (0.684)	1.874 (0.930)	5.119 (0.953)	1.499 (0.682)	1.830 (0.934)	5.129 (0.953)
Model 2	9.917 (0.007)	1.638 (0.650)	1.812 (0.936)	4.404 (0.974)	1.641 (0.649)	1.745 (0.941)	4.324 (0.976)
Model 3	9.939 (0.007)	1.622 (0.654)	2.001 (0.919)	4.992 (0.958)	1.617 (0.655)	1.996 (0.920)	5.026 (0.957)
Model 4	9.938 (0.007)	3.783 (0.285)	4.052 (0.669)	7.514 (0.821)	3.752 (0.289)	4.041 (0.671)	7.535 (0.820)
Model 5	9.832 (0.007)	7.994 (0.046)	18.556 (0.004)	43.850 (1.61E–5)	7.600 (0.055)	16.824 (0.009)	34.904 (0.0004)
Model 6	9.825 (0.006)	6.943 (0.073)	10.000 (0.124)	21.060 (0.049)	6.635 (0.084)	9.363 (0.154)	17.943 (0.117)
Model 7	9.883 (0.007)	3.246 (0.355)	3.480 (0.746)	6.871 (0.865)	3.255 (0.353)	3.469 (0.747)	6.768 (0.872)

This table reports specification results of the seven exponential ARCH models for the daily data. The test is performed on the standardized residuals. The normality test is a Wald test for the skewness and excess kurtosis coefficients to be jointly equal to zero. The p -values are based on the chi-square distribution with 2 degrees of freedom. m is the number of autocorrelations for the Ljung–Box test. q is the number of the order of ARCH effects. These tests are distributed as chi-square with the degrees of freedom equal to m or q . The p -values are in parentheses. The data used are the daily S&P 100 Index returns from January 1986 to December 1999.

Moreover, if the models are properly specified, these residuals should be non-autocorrelated and free of any conditional heteroskedasticity. As for the partially linear model, we conduct the Ljung and Box (1978) test for autocorrelation and the Lagrange multiplier ARCH test from Engle (1982) for conditional heteroskedasticity. The testing results are reported in Tables 7 and 8 for the weekly and daily data, respectively.

In the weekly data, we never reject non-autocorrelation and conditional homoskedasticity, indicating that all models seem to capture the conditional volatility dynamics well. In the daily data, however, Model 5 is rejected based on both the Ljung–Box test and the ARCH test. There is weak evidence against Model 6, based on the Ljung–Box test with $m=12$.

4.4. Model performance tests

To measure the relative performance of the models, we compute the sum of squared errors (SSE) and the sum of absolute errors (SAE) for the seven models and report them in Table 9. In the upper panel of Table 9, we see that for the weekly data, the SSEs from Model 1 to 4 are slightly larger than those from Models 5 to 7. Model 7 has the smallest SSE at 0.3288. The results based on the SAE are similar. Overall, we find that Model 7 is the preferred model for the weekly data. For the daily data (lower panel), Model 7 is again the preferred model based on the SSE. The sum of absolute errors for Model 7, however, is slightly above those for Models 5 and 6. But recall that these two models are

Table 9
Sum of squared and absolute errors of exponential ARCH models

	Sum of squared errors	Sum of absolute errors
<i>Weekly data</i>		
Model 1	0.3301	11.2884
Model 2	0.3298	11.2822
Model 3	0.3303	11.2698
Model 4	0.3302	11.2712
Model 5	0.3294	11.2745
Model 6	0.3295	11.2633
Model 7	0.3288	11.2624
<i>Daily data</i>		
Model 1	0.4277	25.3851
Model 2	0.4275	25.3720
Model 3	0.4276	25.3786
Model 4	0.4276	25.3710
Model 5	0.4277	25.3615
Model 6	0.4277	25.3620
Model 7	0.4274	25.3694

This table reports the sum of squared errors (SSE) and sum of absolute errors (SAE) of the seven exponential ARCH models. The data are the weekly and daily S&P 100 Index returns from January 1986 to December 1999.

Table 10
Sign and size bias tests of exponential ARCH models

	Sign bias	Positive size bias	Negative size bias
<i>Weekly data</i>			
Model 1	1.0551	-0.3618	0.0474
Model 2	1.2062	-0.1998	0.3248
Model 3	1.5381	-0.0370	0.2576
Model 4	1.8727	0.0877	0.2608
Model 5	0.6610	-0.1175	-0.2713
Model 6	1.6659	0.2027	0.0756
Model 7	0.2041	-0.5959	-0.6233
<i>Daily data</i>			
Model 1	0.8893	-0.4573	0.1830
Model 2	0.8826	-0.4693	0.4444
Model 3	0.5554	-0.5639	-1.8689
Model 4	1.0200	0.0907	-1.8281
Model 5	1.0744	1.7418	-1.8899
Model 6	1.1631	0.5929	-1.9405
Model 7	1.0188	0.2380	0.4278

This table reports the *t*-statistics of the sign and size bias tests of the seven exponential ARCH models. The Engle and Ng (1993) sign bias test, positive size bias test and negative bias test are designed to check volatility model specifications with regard to return shocks. The data used are the weekly and daily S&P 100 Index returns from January 1986 to December 1999.

rejected by our specification tests for the daily data. The slightly lower SAE figures for Models 5 and 6 could be due to the ability of these models to fit certain observations well.

To further examine the specification of these models, we conduct the following three tests described in Engle and Ng (1993): the sign bias test, the positive size bias test, and the negative size bias test. These tests are designed to examine the impact of positive and negative innovations on volatility not predicted by the model. If the models are properly specified, then positive and negative innovations should not have an impact on the squared standardized residuals. Under the null hypothesis, the statistics are distributed under the Student t distribution.

We report the test statistics in Table 10; the weekly statistics are presented in the upper panel. Looking at the weekly data, we reject the null hypothesis only for Model 4 by using the sign bias test. In addition, there is some weak evidence against Model 6. Looking at the daily data, we reject Models 3 to 6 by using the negative size bias test. A close examination of these models reveals the cause of the rejection: although the quadratic function allows large negative return shocks to have a large impact on volatility, the impact is overestimated for very large negative return shocks. Model 7 also allows for a large impact for large shocks, but the increase in impact is linear, which seems to be more appropriate for very large shocks.

To sum up the results of Sections 4.1–4.3 and this section, we note that Models 6 and 7 are closely related to the partially linear model at the weekly frequency. Thus, the news impact curve implied by the weekly VIX data seems useful in selecting ARCH specifications. At the daily frequency, however, the implied news impact curves for EGARCH models are quite different from that for the VIX model. Overall, it seems that the piecewise linear Model 7 is the most appropriate for modeling both the weekly and daily conditional volatilities according to our specification and performance tests.

5. Conclusions

In this paper we conducted a close examination of the relationship between return shocks and conditional volatility. We did so in a framework where the impact of return shocks on conditional volatility is specified as a general function and estimated nonparametrically using implied volatility data—the VIX series. This setup can provide a good description of the impact of return shocks on conditional volatility. The estimation results confirm the existence of asymmetric volatility: negative return shocks are associated with increases in conditional volatility while positive return shocks are generally uncorrelated with changes in conditional volatility. It appears that the news impact curves implied by the VIX data is useful in selecting ARCH specifications at the weekly frequency. We find that the exponential ARCH model of Nelson (1991) is capable of capturing most of the asymmetric effect, when return shocks are relatively small. For large negative shocks, however, our nonparametric function points to larger increases in conditional volatility than those predicted by a standard EGARCH model. Our empirical analysis further demonstrates that an EGARCH model with separate

coefficients for large and small negative shocks is better able to capture the asymmetric effect.

Acknowledgements

We thank the editor, the associate editor, and two referees for extensive and insightful comments on this paper. We have benefited especially from the helpful comments of the associate editor. We thank Kurt Butler, Bruno Gerard and James Wiggins for helpful discussions and comments, and Tom Bishop and Brian Boyer for excellent research assistance. We are responsible for all remaining errors.

Appendix A. Bandwidth selection

The estimation with optimal bandwidth selection can be implemented by a two-step procedure. In the first step, we use a simple bandwidth h_0 to generate some preliminary estimators. Then these estimators are “plugged into” the formula to obtain an estimator for the optimal bandwidth. In order to obtain an estimator for the optimal bandwidth, we need to estimate the quantities \mathcal{B} and \mathcal{V} , which are functions of nonparametric bias and variance. Linton (1995) shows that \mathcal{V} can be estimated by

$$h_0 \sum_i \sum_{j \neq i} \left[\sum_{k \neq i, k \neq j} w_{ki} w_{kj} - 2w_{ij} \right]^2$$

where

$$w_{ij} = K\left(\frac{u_i - u_j}{h_0}\right)(u_i - u_j)^2 \bigg/ \left[\sum_{l=1}^n K\left(\frac{u_i - u_l}{h_0}\right)(u_i - u_l)^2 \right],$$

and \mathcal{B} can be estimated by

$$\hat{\sigma}^{-2} n^{-1} \sum_{t=1}^n \hat{\beta}_{i1} \hat{\beta}_{i2}$$

where

$$\hat{\sigma}^2 = \frac{1}{n} \sum_{t=1}^n \left[y_t - E(\widehat{y_t | u_{t-1}}) - \tilde{\beta}(x_t - E(\widehat{x_t | u_{t-1}})) \right]^2,$$

and

$$\hat{B}_{i1} = h_0^{-q} \left[\sum_{t \neq i} w_{it} E(\widehat{x_t | u_{t-1}}) - E(\widehat{x_i | u_{i-1}}) \right],$$

$$\hat{B}_{i2} = h_0^{-q} \left[\sum_{t \neq i} w_{it} \widehat{g}(u_{t-1}) - \widehat{g}(u_{i-1}) \right].$$

After substituting the estimates $\hat{\mathcal{B}}$ and $\hat{\mathcal{V}}$ into the following formula

$$h^* = [\mathcal{V}/(4q\mathcal{B})]^{1/[4q+1]} n^{-2/[4q+1]},$$

we obtain an estimator for the optimal bandwidth.

References

- Amin, K., Ng, V., 1994. A Comparison of Predictable Volatility Models Using Option Data, Working Paper, International Monetary Fund.
- Andersen, T., Bollerslev, T., 1998. Answering the skeptics: yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39, 885–905.
- Baba, Y., Engle, R., Kraft, D., Kroner, K.F., 1989. Multivariate Simultaneous Generalized ARCH, Discussion Paper 89-57, University of California, San Diego.
- Bakshi, G.S., Cao, C., Chen, Z., 1997. Empirical performance of alternative option pricing models. *Journal of Finance* 52, 2003–2049.
- Bates, D.S., 1997. Post-'87 Crash Fears in S&P 500 Futures Options, Working Paper No. 5894, NBER.
- Beckers, S., 1981. Standard deviations implied in option prices as predictors of future stock price volatility. *Journal of Banking and Finance* 5, 363–381.
- Bekaert, G., Wu, G., 2000. Asymmetric volatility and risk in equity markets. *Review of Financial Studies* 13, 1–42.
- Black, F., 1976. Studies of stock price volatility changes. Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economical Statistics Section, pp. 177–181.
- Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637–659.
- Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31, 307–327.
- Bollerslev, T., Wooldridge, J., 1992. Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. *Econometric Reviews* 11, 142–172.
- Bollerslev, T., Chou, R.Y., Kroner, K.F., 1992. ARCH modeling in finance. *Journal of Econometrics* 52, 5–59.
- Box, G., Pierce, D., 1970. Distribution of residual autocorrelations in autoregressive-integrated moving average time series models. *Journal of the American Statistical Association* 65, 1509–1526.
- Breusch, T.S., Pagan, A.R., 1980. The Lagrange multiplier test and its applications to model specification in econometrics. *Review of Economic Studies* 47, 239–254.
- Campbell, J.Y., Hentschel, L., 1992. No news is good news: an asymmetric model of changing volatility in stock returns. *Journal of Financial Economics* 31, 281–318.
- Campbell, J.Y., Lo, A.W., MacKinlay, A.C., 1997. *The Econometrics of Financial Markets*. Princeton Univ. Press, Princeton, NJ.
- Chen, H., 1985. Data Smoothing in Analysis of Covariance, Technical Report AMS-85-65, State University of New York, Stony Brook.
- Chen, J., Hong, H., Stein, J.C., 2001. Forecasting crashes: trading volume, past returns and conditional skewness in stock prices. *Journal of Financial Economics* 61 (3), 345–381.
- Chiras, D.P., Manaster, S., 1978. The information content of option prices and a test of market efficiency. *Journal of Financial Economics* 6, 213–234.
- Christie, A.A., 1982. The stochastic behavior of common stock variances—value, leverage and interest rate effects. *Journal of Financial Economics* 10, 407–432.

- Cox, J.C., Ross, S.A., 1976. The valuation of options for alternative stochastic processes. *Journal of Financial Economics* 3, 145–166.
- Dickey, D.A., Fuller, W.A., 1979. Distribution of estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association* 74, 427–431.
- Duan, J., 1995. Fitting and Smile Family—A GARCH Approach, Working Paper, McGill University.
- Engle, R., 1982. Autoregression conditional heteroskedasticity with estimates of the variance of the UK inflation. *Econometrica* 50, 987–1007.
- Engle, R.F., Kroner, K., 1995. Multivariate simultaneous generalized ARCH. *Econometric Theory* 11, 122–150.
- Engle, R.F., Ng, V.K., 1993. Measuring and testing the impact of news on volatility. *Journal of Finance* 48, 1749–1778.
- Fan, Y., Li, Q., 1999. Root-n-consistent estimation of partially linear time series models. *Journal of Nonparametric Statistics* 11, 251–269.
- French, K.R., Schwert, G.W., Stambaugh, R., 1987. Expected stock returns and volatility. *Journal of Financial Economics* 19, 3–29.
- Glosten, L.R., Jagannathan, R., Runkle, D.E., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance* 48, 1779–1801.
- Green, P., Yandell, B.S., 1985. Semiparametric generalized linear models. In: Gilchrist, R. (Ed.), *Proceedings of the International Conference on Generalized Linear Models*. Springer-Verlag.
- Hardle, W., Hart, J., Marron, S., Tsybakov, A.B., 1992. Bandwidth choice for average derivative estimation. *Journal of American Statistical Association* 87, 218–226.
- Harvey, C., Siddique, A., 1999. Autoregressive conditional skewness. *Journal of Financial and Quantitative Analysis* 34, 465–487.
- Heckman, N., 1986. Spline smoothing in a partial linear model. *Journal of Royal Statistics Society B*, 244–248.
- Hentschel, L., 1995. All in the family: nesting symmetric and asymmetric GARCH models. *Journal of Financial Economics* 39, 71–104.
- Kwiatkowski, D., Phillips, P.C.B., Schmidt, P., Shin, Y., 1992. Testing the null hypothesis of stationarity against the alternative of a unit root: how sure are we that economic time series have a unit root? *Journal of Econometrics* 54, 159–178.
- Latane, H., Rendleman, R.J., 1976. Standard deviations of stock price ratios implied by option premia. *Journal of Finance* 31, 369–382.
- Linton, O.B., 1995. Second order approximation in the partially linear regression model. *Econometrica* 63, 1079–1112.
- Liviatan, N., 1963. Consistent estimation of distributed lags. *International Economic Review* 4, 44–52.
- Ljung, G., Box, G., 1978. On a measure of lack of fit in time series models. *Biometrika* 66, 67–72.
- Lo, A., MacKinlay, A.C., 1990. An econometric analysis of nonsynchronous-trading. *Journal of Econometrics* 45, 181–212.
- Nelson, D.B., 1991. Conditional heteroskedasticity in asset returns: a new approach. *Econometrica* 59, 347–370.
- Pagan, A.R., Schwert, G.W., 1990. Alternative models for conditional stock volatility. *Journal of Econometrics* 45, 267–290.
- Pindyck, R.S., 1984. Risk, inflation, and the stock market. *American Economic Review* 74, 334–351.
- Rabemananjara, R., Zakoian, J.M., 1993. Threshold ARCH models and asymmetries in volatility. *Journal of Applied Econometrics* 8, 31–49.
- Rice, J., 1986. Convergence rate for partially linear splined models. *Statistics and Probability Letters*, 203–208.
- Rissanen, J., 1978. Modeling by shortest data description. *Automatica* 14, 465–471.
- Robinson, P.M., 1988. Root-N-consistent semiparametric regression. *Econometrica* 56, 931–954.
- Roll, R., 1984. A simple implicit measure of the effective bid–ask spread in an efficient market. *Journal of Finance* 39, 1127–1140.
- Schwarz, G., 1978. Estimating the dimension of a model. *Annals of Statistics* 6, 461–464.
- Schwert, G.W., 1989. Why does stock market volatility change over time? *Journal of Finance* 44, 1115–1153.
- Sheikh, A., 1991. Transaction data tests of S&P 100 call option pricing. *Journal of Financial and Quantitative Analysis* 26, 459–475.
- Speckman, P., 1988. Kernel smoothing in partial linear models. *Journal of Royal Statistics Society B*, 413–436.

- Wahba, G., 1984. Partial spline models for the semiparametric estimation of functions of several variables. *Statistical Analysis of Time Series*. Institute of Statistical Mathematics, Tokyo, pp. 319–329.
- Whaley, R.E., 1993. Derivatives on market volatility: hedging tools long overdue. *Journal of Derivatives* 1, 71–84.
- Whaley, R.E., 2000. The Investor Fear Gauge, Working Paper, Duke University.
- White, H., 1982. Maximum likelihood estimation of misspecified models. *Econometrica* 50, 1–26.
- Wu, G., 2001. The determinants of asymmetric volatility. *Review of Financial Studies* 14, 837–859.
- Zakoian, J.M., 1994. Threshold heteroskedastic models. *Journal of Economic Dynamics and Control* 18, 931–955.