Available online at www.sciencedirect.com
ScienceDirect
Journal of Financial Markets 13 (2010) 225–248
Journal of FINANCIAL MARKETS
www.elsevier.com/locate/finmar

How asymmetric is U.S. stock market volatility?

Louis H. Ederington[a,*], Wei Guan[b]

[a]Finance Division, Michael F. Price College of Business, University of Oklahoma, 205A Adams Hall, Norman, OK 73019, USA
[b]College of Business, University of South Florida St. Petersburg, 140 Seventh Avenue South, St. Petersburg, FL 33701, USA

Available online 12 November 2009

Abstract

This paper explores differences in the impact of equally large positive and negative surprise return shocks in the aggregate U.S. stock market on: (1) the volatility predictions of asymmetric time-series models, (2) implied volatility, and (3) realized volatility. Following large negative surprise return shocks, both asymmetric time-series models (such as the EGARCH and GJR models) and implied volatility predict an increase in volatility and, consistent with this, ex post realized volatility normally rises as predicted. Following large positive return shocks, asymmetric time-series models predict an increase in volatility (albeit a much smaller increase than following a negative shock of the same magnitude), but both implied and realized volatilities generally fall sharply. While asymmetric time-series models predict a decline in volatility following near-zero returns, both implied and realized volatility are normally little changed from levels observed prior to the stable market. The reasons for the differences are explored.

JEL classification: G10; G13; C22

Keywords: Asymmetric volatility; Implied volatility; GARCH; EGARCH

1. Introduction

This paper examines how asymmetric the volatility responses to surprise return shocks are for the aggregate U.S. stock market—specifically whether volatility rises or falls following large positive return shocks and how volatility behaves following near-zero

*Corresponding author. Tel.: +1 405 325 5591; fax: +1 405 325 7688.
E-mail addresses: lederington@ou.edu (L.H. Ederington), wguan@mail.usf.edu (W. Guan).

1386-4181/$ - see front matter
doi:10.1016/j.finmar.2009.10.001

returns. We examine and compare the behavior of three volatility measures: (1) the predictions of asymmetric time-series models, such as the EGARCH and GJR models, (2) implied volatility, specifically the Chicago Board Option Exchange's VIX index, and (3) realized volatility.

It is well established that positive and negative return shocks have differing impacts on aggregate U.S. stock market volatility in that volatility tends to be much higher following negative return shocks than following positive return shocks of the same magnitude. This has been documented repeatedly by estimations of asymmetric time-series GARCH-type models, such as Nelson's (1991) EGARCH model and the GJR model of Glosten et al. (1993). Asymmetry has also been found in the behavior of implied volatility. However, the extent of this asymmetry and its exact form for the aggregate U.S. stock market are not settled. While researchers universally find that volatility increases sharply following large negative returns, its behavior following positive and near-zero returns is less clear. Although estimations of asymmetric GARCH-type models consistently predict an increase in volatility following large positive return shocks (albeit much less than following a negative return shock of the same magnitude), numerous studies report that implied volatility declines.[1]

Extant time series and implied volatility studies are not necessarily in conflict since they differ in several dimensions. First, the horizons differ. While the asymmetric GARCH model predictions are generally for a short horizon, such as the next day, the usual implied volatility metric, the CBOE's VIX index, measures implied volatility over the next month. Second, the markets differ. Until September 2003, the VIX measured implied volatility on the S&P 100 index while most asymmetric GARCH estimations are for a broader index. Third, data periods vary.

Hence our first task is to compare changes in asymmetric GARCH model forecasts and implied volatility following positive and negative shocks where the forecast horizons, market index (S&P 500), and data periods are identical for both the GARCH model estimations and implied volatility. Consistent with previous studies, we find clear inconsistencies between the volatility predictions of the asymmetric GARCH models and implied volatility. While the time-series models, specifically Nelson's (1991) EGARCH model, the GJR model of Glosten et al. (1993), and Engle and Ng's (1993) partially nonparametric model, predict a small volatility increase following large positive surprise returns, implied volatility generally falls. No such discrepancy is observed following large negative surprise returns when the asymmetric GARCH models predict a volatility increase and implied volatility rises by roughly the same amount.

These results raise a number of issues that have not been addressed in the literature to date. One, since the asymmetric GARCH and implied volatility forecasts differ following positive but not negative shocks, do they differ (and if so how) following periods without return shocks, i.e., stable markets? Here we find that while the asymmetric time-series models predict a decline in volatility following days with small (or near-zero) surprise returns, implied volatility tends to be little changed. Two, and most important, when the asymmetric GARCH and implied volatility forecasts differ, which is correct? In other words, how does realized volatility behave? On this question, we find that the behavior of realized volatility basically corresponds to that predicted

[1] As discussed more fully in Section 3, these include Fleming et al. (1995), Bates (2000), Poteshman (2001), Pan (2002), Low (2004), Bollerslev and Zhou (2006), and Denis et al. (2006).

by implied volatility, not the asymmetric GARCH models. Specifically, realized volatility generally rises sharply following large negative returns, falls sharply following large positive returns (when the asymmetric GARCH models predict a slight increase), and is roughly unchanged following near-zero returns (when the asymmetric GARCH models predict a decline). Three, why following positive and near-zero returns do the asymmetric GARCH models predict volatility changes opposite to those normally observed? Partial explanations include the high weight that the likelihood function places on extreme observations and the importance of more distant observations.

Our findings imply that a basic tenet of the volatility persistence literature that volatile markets beget volatile markets and stable markets beget stable markets is only partially true. Volatile markets do tend to follow bear markets but implied and realized volatility both tend to fall following bull markets. Following stable markets (near-zero returns), implied and realized volatility are little changed from the levels observed prior to the near-zero return. Our volatility findings are consistent with the price of risk findings of Kane et al. (2000). If a rise (fall) in expected volatility leads investors to require higher (lower) expected returns to compensate, then our findings imply that the price of risk should rise (fall) following large negative (positive) returns as they document. While we find that subsequent implied and realized volatility are monotonic negative functions of the net return over the preceding period, they are both positive functions of intra-period volatility.

The paper is organized as follows. In the next section, we briefly review theoretical explanations for asymmetric volatility. In Section 3, we review previous empirical studies. Our data are described in Section 4 and several asymmetric GARCH models are estimated in Section 5. The impact of surprise return shocks on implied volatility is documented in Section 6 and the behavior of realized or ex post volatility in Section 7. Section 8 explores possible reasons for the differences and Section 9 concludes.

2. Theories of volatility asymmetry

Based on the insight of the early volatility persistence literature that volatile markets tend to follow volatile markets, early time-series models, ARCH and GARCH, incorporated the assumption that volatility increased equally following positive and negative shocks. However, it soon became apparent through the estimation of more flexible models, such as GJR and EGARCH, that stock volatility (especially aggregate U.S. stock market volatility) is asymmetric in that negative return shocks have a greater impact on subsequent volatility than positive return shocks.[2]

The two main theories of this observed asymmetry are the leverage and feedback hypotheses. In the volatility feedback hypothesis of French et al. (1987), Schwert (1989), and Campbell and Hentschel (1992), the asymmetry is due to a risk premium. If a return shock causes expectations of future volatility to rise, this rise in risk causes investors to require a higher expected future return, causing current prices to fall (a negative return). This feedback effect partially offsets a positive return and augments a negative return. In the leverage hypothesis of Black (1976) and Christie (1982), since a large negative return shock means a decline in aggregate equity values, leverage at the average firm is increased.

[2]See the review by Poon and Granger (2003).

Thus for the same volatility of assets, equity volatility should increase. Conversely, a positive return shock implies a decrease in leverage and lower volatility.[3]

Bekaert and Wu (2000) argue that part of the reason for the asymmetry in aggregate market returns, as well as an explanation for why aggregate returns are more asymmetric than firm returns, is that firm returns are more correlated in down markets. Thus if firm level volatility is unchanged, when the market declines, aggregate stock market volatility rises because the covariances rise. Recently, Avramov et al. (2006) have presented a fourth hypothesis. They argue that following a positive return, selling activity is dominated by informed or contrarian investors tending to reduce volatility. Following a negative return, selling is supposedly dominated by uninformed or herding investors leading to an increase in volatility.

Our study clarifies how much asymmetry these and other theories have to explain. For instance, it has been argued by papers, such as Christie (1982) and Schwert (1989), that the leverage effect is too weak to account for the asymmetry predicted by the time-series models. Accounting for a decline in volatility following a positive return is even harder.

3. Evidence on volatility asymmetry

3.1. Time-series models

Numerous studies have found that U.S. stock market volatility is higher following negative returns than following positive returns of the same magnitude. Early work includes Christie (1982), French et al. (1987), and Schwert (1989). Numerous GARCH type models have been proposed to estimate this asymmetry including most prominently Nelson's (1991) EGARCH model, the GJR, or TGARCH, model of Glosten et al. (1993), and Engle and Ng's (1993) partially nonparametric model. For a review of pre-2000 estimations of such models, see Bekaert and Wu (2000). More recent examples include Bekaert and Wu (2000), Wu and Xiao (2002), Li et al. (2005), Caporin and McAleer (2006), and Koulakiotis et al. (2006).

To the authors' knowledge, all estimations of such models for aggregate U.S. stock indices find a reverse-J shape in which future volatility tends to: (1) rise strongly following negative return shocks, (2) rise weakly following positive return shocks, and (3) decline following periods in which surprise returns are close to zero. The same pattern has been observed for non-U.S. stock markets, although the asymmetry is not statistically significant in all (Bekaert and Wu, 2000; Li et al., 2005; Koulakiotis et al., 2006). Employing a nonparametric approach, Kane et al. (2000) find large increases in squared returns following negative price jumps and smaller increases following positive price jumps.

[3]GARCH-in-mean or GARCH-M type models attempt to capture both with a mean equation in which the return is a simultaneous function of expected volatility (intended to capture the feedback effect) and a variance equation in which future volatility is a function of the current return (intended to capture the leverage effect). However, results for the mean equation have been mixed. While some, e.g., Campbell and Hentschel (1992), Wu (2001), and Ghysels et al. (2004) find positive and significant coefficients for expected volatility, others, e.g., Glosten et al. (1993), Lettau and Ludvigson (2009), and Bollerslev et al. (2006), obtain negative or insignificant estimates.

3.2. Implied volatility

Numerous studies have also examined the relation between surprise returns and implied volatility—usually as measured by the CBOE's VIX index. In linear estimations, which do not allow different slopes for positive and negative return shocks, strong negative correlations between U.S. stock market returns and implied volatility changes have been documented by Bates (2000), Poteshman (2001), Pan (2002), and Denis et al. (2006), implying that positive shocks lead to a decline in implied volatility. In estimations that allow differing slopes, results differ. Wu and Xiao (2002) find reverse-J shapes somewhat similar to those generated by the asymmetric GARCH type models, except that the relation between positive shocks and the change in implied volatility is quite weak. However, Fleming et al. (1995), Whaley (2000), and Low (2004) find that implied volatility falls following large positive shocks, instead of the positive relation of the reverse-J relation. Although insignificant, the coefficient estimates in Bollerslev and Zhou (2006) also suggest that positive shocks are associated with declines in implied volatility.

In summary, in the U.S. stock market, implied volatility clearly rises following negative return shocks. However, excepting Wu and Xiao (2002), most have found that implied volatility falls following positive shocks contrary to the predictions of asymmetric GARCH models. As noted above, comparisons between the two sets of studies are difficult because of differing forecast horizons, stock indices, and data periods. The main questions we seek to answer here are: (1) whether this discrepancy following positive return shocks between conditional volatility from time-series models and implied volatility is observed when estimated for the same index, forecast horizon, and data period; (2) how realized volatility behaves; (3) how implied and realized volatility change following near-zero returns; and (4) why the behavior of conditional, implied, and realized volatility differs?

4. Data

Our main data set consists of daily closing price observations from 1/2/1990 to 12/31/2005 on the S&P 500 index (or SPX) and the Chicago Board Option Exchange's (CBOE's) VIX index, the most well known and studied measure of implied stock market volatility. For some calculations, we also make use of intraday observations of the S&P 500 index. The CBOE's calculation of the VIX has changed. Prior to September 2003, the CBOE calculated the VIX as a weighted average of the Black–Scholes implied volatilities for eight at-the-money (ATM) options on the S&P 100 index: the two ATM calls and two ATM puts with expirations close to, but less than, one month and the two ATM calls and puts with expirations close to, but longer than, one month. The eight implied volatilities were weighted so that the average time to expiration was 30 calendar days and the average strike was ATM. In September 2003, the CBOE changed the VIX calculation in three ways: (1) it switched from S&P 100 to S&P 500 options; (2) it changed to a calculation based on most ATM and out-of-the-money options; and (3) it switched from a weighted average of Black–Scholes implied volatilities to a "distribution free" calculation based on the work of Demeterfi et al. (1999) and Britten-Jones and Neuberger (2000).[4] Implied volatilities on

[4]For details of calculation of the new VIX see the CBOE's white paper at http://www.cboe.com/micro/vix/introduction.aspx. Properties of a similar distribution free implied volatility calculation are explored in Jiang and Tian (2005).

Table 1
Descriptive statistics.
Means, standard deviations, and first-order autocorrelations are reported for the S&P 500 index and its implied volatility index, VIX. Statistics are reported for both index levels, daily returns, and monthly returns based on data from 1/2/1990 to 12/31/2005. Return means and standard deviations are annualized.

Series	Mean	Standard deviation	First-order autocorrelation
S&P 500 index-levels	840	366	
Daily returns (annualized)	7.78%	16.09%	−0.002
Monthly returns (annualized)	7.90%	13.61%	
VIX-levels	19.45%	6.40%	
Daily returns (annualized)	−2.23%	88.49%	−0.070
Monthly returns (annualized)	−2.69%	56.04%	

different expiries are still weighted to keep the average expiry at one month. The CBOE has calculated this new VIX back to 1/2/1990 and re-labeled the old calculations as VXO. We use the new VIX. Its main advantage for our purposes is that it measures volatility on a broader market index, the S&P 500. However, the behavior of the VIX and VXO indices is quite similar with a correlation of 0.984 in levels, 0.955 in monthly percentage changes, and 0.818 in daily percentage changes. Virtually all our results are insensitive to this choice.

Descriptive statistics are reported in Table 1. Among the salient statistics in Table 1 are these. One, the oft-noted tendency of implied volatility to over-estimate actual volatility is reflected here in that the mean implied volatility is 19.5% while the annualized standard deviation of returns is 16.1% calculated from daily returns and 13.6% based on monthly returns. Two, daily percentage changes in implied volatility are somewhat negatively serially correlated. Three, the VIX is highly volatile with an annualized standard deviation of daily log percentage changes of 88.5%.

5. Time-series estimates of conditional volatility

We first estimate the asymmetry in U.S. stock market volatility using the GJR (or TGARCH) model of Glosten et al. (1993),[5] and the EGARCH model of Nelson (1991). In both estimations, the mean equation is the GARCH-in-mean (or GARCH-M) (Engle et al., 1987) specification[6]:

$$r_t = \gamma_0 + \gamma_1 \sigma_t + e_t, \tag{1}$$

where r_t is the daily log return and where it is assumed $e_t \sim \mathrm{N}(0, \sigma_t^2)$. In the GJR model:

$$\sigma_{t+1}^2 = \alpha_0 + \alpha_1 \sigma_t^2 + \alpha_2 e_t^2 + \alpha_3 e(-)_t^2, \tag{2}$$

where $e(-)_t = e_t$ if $e_t < 0$ and $= 0$ if $e_t \geq 0$. If $\alpha_3 = 0$, the GJR model collapses to the GARCH(1,1) model and the relation is symmetric. In EGARCH:

$$\ln(\sigma_{t+1}^2) = \beta_0 + \beta_1 \ln(\sigma_t^2) + \beta_2 |e_t/\sigma_t| + \beta_3 (e_t/\sigma_t). \tag{3}$$

[5]Specifically "Model 2" proposed by Glosten et al. (1993). While they present several models, this is the one commonly referred to in the literature as the GJR model.

[6]We have also estimated models in which $r_t = e_t$ and $r_t = \gamma_0 + e_t$ and the results are little changed from those presented here.

If $\beta_3 = 0$, the relation is symmetric. In the GJR model, $\alpha_3 > 0$ implies that negative shocks have a greater impact on conditional volatility than positive shocks. In EGARCH, $\beta_3 < 0$ has the same implication.

Estimations of Eqs. (1)–(3), as well as GARCH(1,1) for comparison, are reported in Table 2. In both the GJR and EGARCH estimations, the null that the relation is symmetric is strongly rejected in that $\alpha_3 > 0$ and $\beta_3 < 0$. Both estimate a reverse-J shape in which both large positive and negative shocks positively impact conditional volatility but the impact of negative shocks is much stronger. Indeed, in the GJR estimation, α_2 is small and insignificant (negative values of α_2 are allowed in our estimation). In the EGARCH estimation, $\beta_2 + \beta_3 = 0.0295$, which is positive and significant at the 0.05 level. Still, the estimated impact of a negative shock on the next day's volatility is more than six times the impact of a positive shock of the same magnitude. The estimated impact of positive shocks is weaker than that in most previous studies. This is at least partially due to the time period. Estimated over the first half of our data period, which more closely coincides with previous studies, α_2 and $\beta_2 + \beta_3$ are somewhat larger.

In columns 2 and 3 of Table 3, we report the implied log percentage changes in conditional volatility, $\ln(\sigma_{t+1}/\sigma_t)$, following various return shocks, e_t, according to the parameter estimates in Table 2 and assuming conditional volatility at time t is at its unconditional level. These news impact curves are graphed in Fig. 1. As reported in Table 3 and Fig. 1, the impact of a negative return shock at time t on conditional volatility for time $t+1$ is considerable. For instance, if the market declines 2.5% (about 3.3% of down days are of this magnitude or greater), conditional volatility for the next day increases about 27.1% according to the GJR model and 21.1% according to the EGARCH model. The impact of positive shocks is much weaker with conditional volatility rising only after very large shocks. Though the GJR model is more convex (as noted by Engle and Ng, 1993, among others), both the GJR and EGARCH estimations suggest that negative shocks

Table 2
Time-series models of asymmetric volatility.

Estimations of the GARCH(1,1) model, the GJR model of Glosten et al. (1993), and the EGARCH model of Nelson (1991) are reported, where the mean equation is $r_t = \gamma_0 + \gamma_1\sigma_t + e_t$ and r_t is the log return on day t. $e(-)_t = e_t$ if $e_t < 0$ and 0 otherwise. The equations are estimated using daily data for the S&P 500 index from 1/2/1990 to 12/31/2005. P-values based on Bollerslev–Woolridge robust standard errors are reported in parentheses for the null that the coefficient equals zero.

Panel A: GARCH(1,1) and GJR models: $\sigma^2_{t+1} = \alpha_0 + \alpha_1\sigma^2_t + \alpha_2 e^2_t + \alpha_3 e(-)^2_t$						
γ_0	γ_1	α_0	α_1	α_2	α_3	Log-likelihood
−9.98E−05	0.0773	5.73E−07	0.9366	0.0586		13 326.7
(0.8011)	(0.1178)	(0.0004)	(0.0000)	(0.0000)		
−1.80E−04	0.0588	1.13E−06	0.9286	0.0076	0.1019	13 373.4
(0.6314)	(0.2220)	(0.0000)	(0.0000)	(0.4100)	(0.0000)	
Panel B: EGARCH: $\ln(\sigma^2_{t+1}) = \beta_0 + \beta_1\ln(\sigma^2_t) + \beta_2\lvert e_t/\sigma_t\rvert + \beta_3(e_t/\sigma_t)$						
γ_0	γ_1	β_0	β_1	β_2	β_3	
6.40E−05	0.0250	−0.2494	0.9829	0.1140	−0.0845	13 382.4
(0.8687)	(0.6129)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	

Table 3
Time-series model estimates of the impact of return shocks on conditional volatility.

Estimates of the impact of a surprise return shock in the stock market on the next day's volatility are presented for the EGARCH and GJR models estimated in Table 2 and the partially nonparametric model of Engle and Ng (1993) with knots at +2%, +1%, +0.5%, 0%, −0.5%, −1%, and −2%. It is assumed that volatility is at its unconditional mean prior to the shock. We also present estimates of the impact on conditional volatility over the next month according to the EGARCH and GJR models.

Surprise market return on day t, e_t (%)	Log percent change in next day conditional volatility, $\ln(\sigma_{t+1}/\sigma_t)$ (%)			Log percent change in next month conditional volatility (%)	
	GJR model	EGARCH model	Partially nonparametric	GJR model	EGARCH model
−3.5	45.15	31.40	36.82	41.49	26.65
−3.0	36.19	26.27	30.53	33.04	22.29
−2.5	27.11	21.13	23.34	24.56	17.92
−2.0	18.22	16.00	14.93	16.37	13.56
−1.5	10.01	10.86	9.08	8.92	9.20
−1.0	3.20	5.72	2.45	2.83	4.85
−0.5	−1.39	0.59	−1.45	−1.22	0.50
0.0	−3.01	−4.55	−3.30	−2.65	−3.85
0.5	−2.90	−3.78	−3.53	−2.55	−3.20
1.0	−2.56	−3.02	−2.39	−2.25	−2.56
1.5	−1.99	−2.26	−2.34	−1.75	−1.91
2.0	−1.22	−1.49	−2.29	−1.07	−1.26
2.5	−0.23	−0.73	−1.49	−0.20	−0.62
3.0	0.94	0.03	−0.69	0.83	0.03
3.5	2.30	0.80	0.09	2.03	0.67

have a significantly greater impact on conditional volatility than positive shocks of the same magnitude.

According to both estimations, the greatest reduction in future volatility occurs when $e_t = 0$, by −3.0% according to the GJR model and −4.6% according to EGARCH. Since both allow a kink only at $e_t = 0$, this is a necessary result given a reverse-J shape. In order to relax this restraint and to test whether the overall relation has the convex shape dictated by the GJR and EGARCH models, we estimated the partially nonparametric model of Engle and Ng (1993) with knots at $e_t = -2\%, -1\%, -0.5\%, 0\%, 0.5\%, 1\%$, and 2%. Since the unconditional standard deviation of daily S&P 500 returns is about 1.01%, these knots correspond to approximately 0, 0.5, 1.0, and 2.0σ. The resulting estimates are shown in column 4 of Table 3 and Fig. 1. While the estimates differ somewhat, the partially nonparametric model's estimation of the impact pattern of a surprise return shock basically matches those of EGARCH and GJR—specifically a reverse-J shape in which: (1) negative shocks have a much greater impact on conditional volatility than positive shocks of the same magnitude, and (2) the greatest decline in conditional volatility occurs near $e_t = 0$.

While the time-series models generate volatility forecasts for the next day, the VIX measures implied volatility over the next month. Hence, for comparison, columns 5 and 6 of Table 3 present GJR and EGARCH estimations of the impact of a return shock on forecast volatility for the next month. These estimates, commonly referred to as

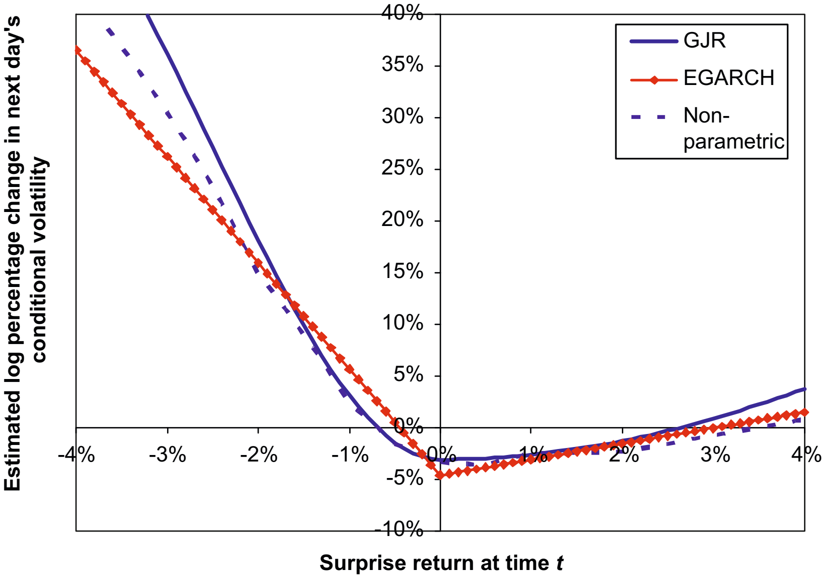

Fig. 1. Asymmetric time-series model estimates of the impact of a surprise return shock on U.S. stock market (S&P 500) volatility.

"integrated volatility," are obtained by the usual recursive procedure in which forecast volatility for day $t+1$ is substituted for e^2 and σ^2 on the right-hand side of Eqs. (2) and (3) to forecast σ_{t+2}, the σ_{t+2} forecast is used to forecast σ_{t+3}. The procedure is repeated through trading day $t+21$, and the 21 forecast volatilities are averaged.[7] To distinguish integrated from single-day volatilities, double subscripts indicate the beginning and ending of the integration period, e.g., $\sigma_{t+1,t+21}$. As presented in Table 3, while both models are characterized by mean reversion, the impact of a single day's negative shock on conditional volatility over the next month is still considerable. A market decline of 2.5% raises the S&P 500 index conditional volatility over the next month by 24.6% according to the GJR model and 17.9% according to the EGARCH model. Following a day with $e_t = 0$, conditional volatility for the next month declines 2.7% or 3.9% according to the GJR and EGARCH models, respectively.

In summary, while forecast magnitudes differ, all three time-series models come to basically the same conclusions. One, negative return shocks have a large positive impact on conditional volatility. Two, large positive return shocks have a much weaker but positive impact. Three, the largest decline in conditional volatility occurs when the market indices are approximately unchanged. These relations hold for both next day and next month volatility forecasts.

[7]Note that such longer-term forecasts cannot be generated with the partially nonparametric model since it generates forecasts of the variance but does not forecast *when* a future observation will be x standard deviations from its mean. While standard procedure in the literature, Ederington and Guan (2009a) point out that this recursive procedure for generating longer-term integrated volatility forecasts assumes that the relative importance of recent and older observations is the same regardless of the forecast horizon while older observations are actually more important in the longer horizons. Ederington and Guan (2009b) find that the resulting integrated volatility forecasts tend to be biased upward.

Table 4
Impact of 1-day return shock on implied volatility and time-series forecasts of conditional volatility partitioned by the day t return.

We report log percent changes in forecast U.S. stock market volatility calculated as $\ln(\sigma_{t+1,t+21}/\sigma_{t,t+20})$, where $\sigma_{t,t+20}$ is the volatility forecast for the period from day t to day $t+20$ based on information through day $t-1$ and $\sigma_{t+1,t+21}$ is the forecast volatility from $t+1$ to $t+21$ based on information through day t. We report means of $\ln(\sigma_{t+1,t+21}/\sigma_{t,t+20})$ partitioned by the day t return, r_t, for: (1) implied volatility as measured by the VIX and (2) conditional volatility according to the GJR and EGARCH models.

Day t return (%)	Obs.	Mean 1 day index return (%)	Mean log percent change in VIX (%)	T-value	Mean log percent change in EGARCH estimate (%)	Mean log percent change in GJR estimate (%)
<-2.5	56	−3.26	12.32	12.99	16.56	22.41
≥ -2.5 and < -2.0	54	−2.24	8.92	9.97	12.70	14.94
≥ -2.0 and < -1.5	135	−1.72	7.92	16.22	9.74	10.01
≥ -1.5 and < -1.0	239	−1.22	5.08	16.58	7.45	6.19
≥ -1.0 and < -0.5	501	−0.72	3.09	16.96	3.81	1.62
≥ -0.5 and < -0.25	391	−0.36	1.38	7.52	0.63	−0.96
≥ -0.25 and <0	533	−0.13	0.67	4.06	−1.82	−1.88
≥ 0 and <0.25	571	0.12	−0.97	−6.63	−3.29	−2.04
≥ 0.25 and <0.5	451	0.37	−2.05	−11.98	−3.02	−2.13
≥ 0.5 and <1.0	594	0.73	−2.89	−18.41	−2.66	−2.11
≥ 1.0 and <1.5	281	1.23	−4.64	−19.00	−2.22	−1.95
≥ 1.5 and <2.0	120	1.73	−5.26	−10.89	−2.20	−2.00
≥ 2.0 and <2.5	52	2.23	−6.10	−7.08	−1.92	−1.76
≥ 2.5	55	3.40	−9.04	−12.86	−1.96	−1.62

6. Implied volatility

6.1. The short-run impact

Next we consider the impact of a return shock on day t on implied volatility. The term VIX_{t+1} designates implied volatility for the month beginning on day $t+1$ and running through day $t+21$, calculated from closing option prices on day t.[8] The log percentage change in the implied volatility is measured as $\%\Delta VIX_t = \ln(VIX_{t+1}/VIX_t)$.[9] To explore how implied volatility responds to the day t return, r_t, we partition r_t into 14 groupings or strata from $r_t < -2.5\%$ to $>2.5\%$, as presented in column 1 of Table 4. For absolute returns of more than 0.5% the partitions are in intervals of 0.5% and for $|r_t|<0.5\%$ in intervals of 0.25%. The number of observations and mean returns in each interval are reported in columns 2 and 3, respectively of Table 4.

In column 4 of Table 4, we report the mean of $\%\Delta VIX_t$, partitioned by the day t return. In Fig. 2, this mean percentage change is graphed against the mean day t return for each of the fourteen strata. As shown in the table and figure, the change in implied volatility is a

[8]Since VIX always covers a 21-day period, we suppress the second subscript used to describe the integrated GARCH-type forecasts.

[9]The VIX is itself normally expressed as a percentage, or more precisely as the standard deviation of log percentage returns. So $\%\Delta VIX_t$ is the percentage change in a percent. For example, if the VIX rises from 20% to 21%, $\%\Delta VIX_t$ is approximately 5%.

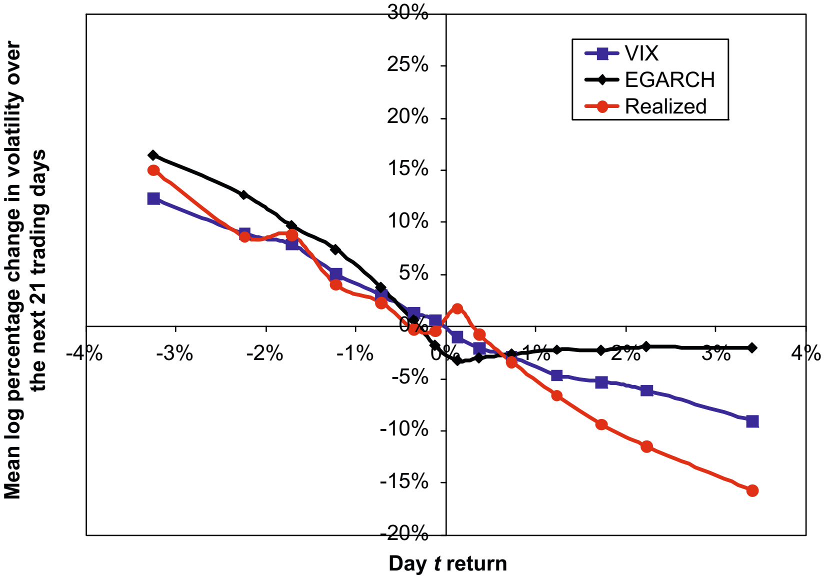

Fig. 2. News impact curves for implied, conditional, and realized volatilities.

strong, monotonic negative function of r_t. For all fourteen return intervals, the mean percentage change is significantly different from zero at the 0.001 level.

The behavior of the VIX following negative returns is in line with the behavior predicted by the asymmetric time-series models in that implied volatility rises and rises more the larger the absolute magnitude of the negative shock. However, the behavior of implied volatility in stable markets and following large positive shocks differs from that predicted by the time-series models. On days when the market is roughly unchanged, implied volatility does not decline as asymmetric GARCH models predict but instead is little changed. On days when the market rises sharply, implied volatility does not rise, as asymmetric GARCH models predict, but falls and falls more the larger the positive shock. For example, when the market rose more than 2.5% in one day, the VIX fell an average of 9.04%.

In terms of implied volatility, the tenet of the volatility persistence literature that stable markets are followed by stable markets and volatile markets by volatile markets appears only partially true. While implied volatility rises following large negative return shocks, it falls following large positive shocks and falls more the greater the positive return. Likewise, there is no indication that near-zero returns lead to a decline in expected volatility. Following a near-zero return, implied volatility is roughly unchanged from that observed prior to the small return.

For comparison, in columns 6 and 7 of Table 4, we report the mean log percentage change in the forecast volatility for the next month according to the GJR and EGARCH model estimates in Table 2. For each day, we calculate the forecast standard deviation over the next 21 trading days, $\hat{\sigma}_{t,t+20}$, based on returns through day $t-1$ and $\hat{\sigma}_{t+1,t+21}$ based on returns through day t and then calculate the log percentage change in this integrated volatility forecast $\ln(\hat{\sigma}_{t+1,t+21}/\hat{\sigma}_{t,t+20})$. These are averaged over all observations in each r_t stratum. Note that this calculation differs from that in Table 3 where $\hat{\sigma}_t$ is assumed to be at

Table 5
Impact of 1-day return shock on implied volatility and time-series forecasts of conditional volatility partitioned by the standardized day t return.

We report log percent changes in forecast U.S. stock market volatility calculated as $\ln(\sigma_{t+1,t+21}/\sigma_{t,t+20})$, where $\sigma_{t,t+20}$ is the volatility forecast for the period from day t to day $t+20$ based on information through day $t-1$ and $\sigma_{t+1,t+21}$ is the forecast volatility from $t+1$ to $t+21$ based on information through day t. We report means of $\ln(\sigma_{t+1,t+21}/\sigma_{t,t+20})$ partitioned: (1) by $r_t/\hat{\sigma}_t$ for implied volatility as measured by the VIX and (2) by $e_t/\hat{\sigma}_t$ for the EGARCH models and GJR models, where r_t is the return on day t, e_t is the surprise return on day t from the mean equation, and $\hat{\sigma}_t$ is the forecast volatility based on information through day $t-1$. * and ** designate means significantly different from zero at the 0.05 and 0.01 levels, respectively.

$r_t/\hat{\sigma}_t$ or $e_t/\hat{\sigma}_t$	VIX		EGARCH		GJR	
	Obs.	Mean of $\ln(\sigma_{t+1,t+21}/\sigma_{t,t+20})$ (%)	Obs.	Mean of $\ln(\sigma_{t+1,t+21}/\sigma_{t,t+20})$ (%)	Obs.	Mean of $\ln(\sigma_{t+1,t+21}/\sigma_{t,t+20})$ (%)
<-2.5	4	23.85^{**}	41	24.13^{**}	38	34.81^{**}
≥-2.5 and <-2.0	9	17.00^{**}	66	14.60^{**}	68	17.40^{**}
≥-2.0 and <-1.5	51	12.13^{**}	163	10.55^{**}	157	10.36^{**}
≥-1.5 and <-1.0	248	7.46^{**}	327	6.47^{**}	340	4.59^{**}
≥-1.0 and <-0.5	558	3.66^{**}	530	2.43^{**}	537	0.21^{**}
≥-0.5 and <-0.25	463	1.70^{**}	406	-0.58^{**}	394	-1.67^{**}
≥-0.25 and <0	583	0.70^{*}	448	-2.65^{**}	453	-2.22^{**}
≥0 and <0.25	614	-1.06^{**}	449	-3.49^{**}	451	-2.26^{**}
≥0.25 and <0.5	500	-2.04^{**}	434	-3.25^{**}	437	-2.30^{**}
≥0.5 and <1.0	618	-3.47^{**}	587	-2.83^{**}	590	-2.19^{**}
≥1.0 and <1.5	265	-4.87^{**}	336	-2.16^{**}	323	-1.84^{**}
≥1.5 and <2.0	100	-5.83^{**}	159	-1.60^{**}	158	-1.39^{**}
≥2.0 and <2.5	18	-9.53^{**}	61	-0.93^{**}	64	-0.62^{**}
≥2.5	8	-9.52^{**}	25	-0.11	22	0.26

its unconditional level while in Table 4 it varies depending on prior market conditions.[10] EGARCH estimates of the news impact curve from Table 4 are graphed along with the VIX news impact curve in Fig. 2.

Comparing the news impact curves for implied volatility and the asymmetric GARCH models, we observe that: (1) both increase with negative return shocks, though the volatility forecasts of the asymmetric GARCH models tend to increase more (particularly the GJR estimates due to its greater convexity); (2) implied volatility falls sharply following large positive returns while the time-series forecasts tend to decline slightly; and (3) implied volatility is little changed when absolute market returns are small while the time-series volatility forecasts fall.

In Table 4, the observations are grouped into partitions or strata on the basis of the return shock r_t without regard to expected volatility at the time. Obviously, a day t return of say 2% is more extreme when volatility is 1% than when it is 2.5%. Consequently, as a robustness check, in Table 5 we present the same statistics where the observations are placed into standardized return strata. For the VIX, the returns are standardized according to $r_t/\hat{\sigma}_{t,t+20}$, where $\hat{\sigma}_{t,t+20}$ is the VIX volatility forecast for the next month based on observations through day $t-1$. For EGARCH and GJR, the returns are standardized

[10] Also the returns in column 1 of Table 4 are raw returns, r_t, while in column 1 of Table 3 they are residuals, e_t, from the mean equation (Eq. (1)). However, this difference is very slight.

Table 6
The short-run VIX reaction to market return shocks.

The log percentage change in the VIX over a one-day period is regressed on: (1) the market (S&P 500) return on the same day, (2) a variable that is equal to the market return if the return is negative and zero otherwise, (3) the market return on the previous day, (4) the log percentage change in the VIX on the previous day, and (5) the same day market return squared. Equations are estimated using daily data from 1/2/1990 to 12/31/2005. Estimated coefficients are reported with t-statistics in parentheses.

	Intercept	Day t market return	Day t market return if negative	Day $t-1$ market return	Lagged percent change in implied volatility	Squared day t market return	Adjusted R^2
Model 1	0.0010	−3.6579					0.442
	(1.57)	(−56.53)					
Model 2	−0.0043	−2.9361	−1.4825	−0.2240	−0.1188		0.460
	(−4.61)	(−26.30)	(−8.08)	(−2.61)	(−7.65)		
Model 3	−0.0048	−2.814	−1.7311	−0.2255	−0.1189	−4.0838	0.460
	(−4.23)	(−14.29)	(−4.59)	(−2.63)	(−7.65)	(−0.75)	

according to $e_t/\hat{\sigma}_t$, where e_t is the surprise return from Eq. (1) and $\hat{\sigma}_t$ is the forecast volatility for day t for that model based on observations through day $t-1$.

While this procedure has the advantage of controlling for differences in expected volatility, a disadvantage is that, since both the numerator and denominator for the standardized returns differ, a given day's observation can wind up in different strata for the three volatility measures. In other words, a given day's observation might be in the $-2<r_t/\hat{\sigma}_{t,t+1}<-1.5$ group for implied volatility, in the $-2.5<e_t/\hat{\sigma}_t<-2$ group for GJR, and in the $e_t/\hat{\sigma}_t<-2.5$ group for EGARCH. Hence, the different volatility forecasts in a row are not as comparable as they were in Table 4. The number of observations in each stratum for each volatility measure are reported in Table 5, along with means of $\ln(\hat{\sigma}_{t+1,t+21}/\hat{\sigma}_{t,t+20})$.

The standardized return results in Table 5 largely mirror the non-standardized results in Table 4. Except for the most extreme negative returns where the GJR model predicts a larger volatility increase, following sizable negative returns, the behavior of the VIX roughly matches that forecast by the two asymmetric time-series models. However, following sizable positive returns, the VIX tends to fall while the time-series models predict little change. Following near-zero return days, the VIX is little changed while the time-series models predict a decline.

We next explore the same issues in a regression format. Results of a univariate regression of the percentage change in implied volatility on the same day return, r_t, are shown in Row 1 of Table 6. Consistent with the results in Tables 4 and 5 and Fig. 2, the percentage change in implied volatility is a strong inverse function of the same day return. According to the results in Row 1, each 1% increase (decrease) in the S&P 500 lowers (raises) the VIX approximately 3.7%.[11] Moreover, with an adjusted R^2 of 0.44, the day t return explains much of the change in implied volatility. The intercept is small, positive, and insignificant, indicating little change in implied volatility following a

[11] Again we would point out that this change is in percentages, not percentage points. For example, if the VIX is initially 20%, and the S&P 500 falls 1%, the VIX tends to rise 20(0.037) = 0.74 percentage points to 20.74%.

zero-return day. Next we estimate

$$\%\Delta\text{VIX}_t = \alpha_0 + \alpha_1 r_t + \alpha_2 r(-)_t + \alpha_3 r_{t-1} + \alpha_4 \%\Delta\text{VIX}_{t-1} + \alpha_5 r_t^2 + u_t. \tag{4}$$

In Eq. (4), $r(-)_t = r_t$ if $r_t < 0$ and 0 otherwise. This variable is included to test whether the relation is kinked at $r_t = 0$. Specifically, $\alpha_2 < 0$ implies that volatility tends to increase more following negative returns than it falls following positive returns. The lagged S&P 500 index return, r_{t-1}, is included to test for the possibility of lags or reversals in the relationship. There are two rationales for including r_{t-1}. First, implied volatility may be mean reverting. Second, there is a possibility of noise or measurement error in the implied volatility index.[12] If there is measurement error, an increase (decrease) one day will tend to be followed by a decrease (increase) the next day, so the successive volatility changes will tend to be negatively correlated, as observed in Table 1. The squared return r_t^2 is included to test for convexities in the relation, as implied by our EGARCH and GJR estimations. Since it complicates the test for a kinked relationship, we first estimate Eq. (4) without r_t^2 and then with.

The results are reported in rows 2 and 3 of Table 6. Again, r_t is highly significant and responsible for most of the regressions' explanatory power. The negative coefficient of $r(-)_t$ indicates that negative shocks raise implied volatility more than positive shocks lower implied volatility. Nonetheless, the relation is continuously downward sloping and, as shown by the negative coefficient of r_t^2, there is no evidence of a convex relation. In other words, the regressions confirm the pattern observed in Tables 4 and 5 and graphed in Fig. 2. Since $\hat{\alpha}_3 < 0$ there is no evidence that the observed contemporaneous relation is temporary or due to non-synchronous trading. The coefficient of lagged changes in the implied volatility index is negative and significant, suggesting that either implied volatility is mean reverting or that there is measurement error or noise in the implied volatility indices.

6.2. Longer-run impact

An obvious question is whether this strong negative correlation between implied volatility and contemporaneous returns holds over longer intervals when changes in the market indices tend to be more substantial. Will they be associated with similarly large changes in implied volatility or, given the evidence of mean reversion in the daily relation, will implied volatility tend to return to its original level? Also, how does implied volatility respond to intra-period volatility? To answer these questions, we estimate regressions for monthly changes in implied volatilities similar to those just estimated for daily changes. Specifically, we relate the change in implied volatility from trading day $t-21$ to day t to the net change in the stock index over the same period and other variables. One advantage of this longer-term perspective is that we can relate the change in implied volatility from day $t-21$ to t to both the net index return over this period and to the volatility of returns within the period. Another advantage is that since VIX measures implied volatility over one month periods, the VIX horizons for trading days t and $t-21$ do not overlap.

[12] If the market rises (falls) and implied call option volatilities are calculated from fresh equity prices but stale option prices, then the calculated implied volatilities will tend to fall (rise). However, for put options, the pattern is reversed. If the market rises, implied volatilities calculated from stale put prices will tend to rise. Since the VIX is an average of implied volatilities from both puts and calls, the two tendencies should approximately offset but we include this variable to test whether this is the case.

The estimated equation is

$$\%\Delta \mathrm{VIX}_{t-21,t} = \alpha_0 + \alpha_1 r_{t-21,t} + \alpha_2 r(-1)_{t-21,t} + \alpha_3 \mathrm{SurpriseVol}_{t-21,t} + \alpha_4 \%\Delta \mathrm{VIX}_{t-42,t-21} + \alpha_5 r^2_{t-21,t} + u_t. \quad (5)$$

The log percentage change in implied volatility over a one-month period is measured as $\%\Delta \mathrm{VIX}_{t-21,t} = \ln(\mathrm{VIX}_t/\mathrm{VIX}_{t-21})$. The log return on the S&P 500 index over the same period is measured as $r_{t-21,t} = \ln(P_t/P_{t-21})$, where P_t is the S&P index on day t. $r(-1)_{t-21,t} = r_{t-21,t}$ if $r_{t-21,t}<0$ and 0 otherwise. $\mathrm{SurpriseVol}_{t-21,t}$ is calculated as $\ln(\mathrm{SD}_{t-20,t}/\mathrm{VIX}_{t-21})$, where $\mathrm{SD}_{t-20,t}$ is the standard deviation of returns from $t-20$ through t. In other words, treating VIX_{t-21} as the market's estimation of volatility from $t-20$ to t, $\mathrm{SurpriseVol}_{t-21,t}$ measures the log percentage difference between actual and expected volatility over the month. The term $r^2_{t-21,t} = (r_{t-21,t})^2$ is included to test for nonlinearities. Since Eq. (5) is estimated using daily data, our observations overlap. Since OLS estimates of standard errors are seriously biased downward due to this overlap, we report t-statistics calculated using Newey–West standard errors with lags out to 21 days.

The results are reported in Table 7. In the univariate regression of $\%\Delta \mathrm{VIX}_{t-21,t}$ on $r_{t-21,t}$ (Model 1), the coefficient of $r_{t-21,t}$ is -2.584, indicating again that the change in implied volatility is strongly related to the direction of change in the S&P 500 index over the month. Combined with the intercept, the implication is that a 10% rise in the market over a month is associated with a fall in implied volatility of 21.8% while a 10% market decline is associated with an increase in implied volatility of 29.8%. The adjusted R^2 is 0.413, indicating that much of the change in implied volatility over one-month horizons can be explained with knowledge of the net market return over that month. Moreover, the intercept is significantly positive, indicating that if the market is unchanged over the period, implied volatility tends to *rise* approximately 4.0%, not fall as the asymmetric GARCH models and most discussions of the volatility persistence hypothesis predict.

Table 7
The longer-run VIX reaction to market return shocks.

The log percentage change in the VIX implied volatility index over a one month period, specifically from trading day $t-21$ to t, is regressed on: (1) the log return for the S&P 500 index over the month, (2) a variable equal to the market return over the month if negative and zero otherwise, (3) the volatility surprise measured as the log percentage difference between the standard deviation of daily returns from $t-20$ to t and the VIX on day $t-21$, (4) the lagged log percentage change in the VIX from $t-42$ to $t-21$, and (5) the index return for the month squared. T-statistics are reported in parentheses using Newey–West standard errors.

	Intercept	S&P 500 index return over month	S&P 500 index return if negative	Volatility surprise	Lagged percent change in implied volatility	Index return squared	Adjusted R^2
Model 1	0.040	−2.5836					0.413
	(2.02)	(−16.00)					
Model 2	0.0619	−1.602	−0.8523	0.1982	−0.2116		0.546
	(5.38)	(−7.46)	(−2.35)	(8.38)	(−7.32)		
Model 3	0.0600	−1.492	−1.0861	0.1980	−0.2110	−1.0427	0.546
	(4.87)	(−4.12)	(−1.53)	(8.39)	(−7.30)	(−0.42)	

Consider next the estimation of Eq. (5) without $r^2_{t-21,t}$—Model 2. Again $\hat{\alpha}_1 < 0$. Since $\hat{\alpha}_2 < 0$, there is also evidence that implied volatility rises more following a negative return than it falls following a positive return. Nonetheless, it is clear from $\hat{\alpha}_1$ that a strong positive market return is associated with a decline in implied volatility. The coefficient of $\text{SurpriseVol}_{t-21,t}$ is positive and significant. The implication is that if over a month actual volatility turns out to be about 10% higher (lower) than that predicted by implied volatility, implied volatility for the next month tends to increase (decrease) by about 2.0%. In other words, as hypothesized, revisions in implied volatility depend on both volatility within the month and the net market movement over the month. Again there is evidence of mean reversion behavior in implied volatility in that $\hat{\alpha}_4 < 0$. As shown by the insignificant coefficient of $r^2_{t-21,t}$ in Model 3, there is no evidence that the relation between $r_{t-21,t}$ and the change in implied volatility is convex.

7. Realized volatility

The finding that the behavior of implied volatility differs considerably from that predicted by the asymmetric GARCH models following large positive and near-zero return shocks raises the question of whether ex post or realized market volatility behaves as predicted by the asymmetric GARCH models or as predicted by implied volatility.

7.1. Short-run

First we examine how realized volatility changes following a single day return shock. For this we measure the percentage change in volatility from the day before the return shock to the day after as $\ln(\sigma_{t+1}/\sigma_{t-1})$, where σ_{t-1} is the standard deviation of 5-min returns the day before the return shock and σ_{t+1} is the standard deviation of 5-min returns the day after. We switch to intraday returns in order to obtain enough return observations each day to calculate a standard deviation.[13] We examine how this realized volatility change varies with the day t return by calculating means of $\ln(\sigma_{t+1}/\sigma_{t-1})$ for the 14 different r_t groups or strata used in Table 4.

The results are presented in column 3 of Table 8 with t-values for the no change null in column 4. Like implied volatility, the change in realized volatility is a strong negative function of r_t. Following days when the market declines 2.5% or more, the standard deviation of 5-min returns increases an average of 29.1%. Following days when the market rises 2.5% or more, the standard deviation *falls* an average of 20.6%.[14] With the exception of the $0.5\% \leq r_t < 1\%$ stratum, the mean percentage change in realized volatility declines monotonically with r_t. The magnitude of the decline in realized volatility following large positive returns is strikingly large since it is even larger than the decline in the VIX, though of course the later is a volatility forecast for the next month, not the next day. When r_t is small in absolute terms, specifically when it is between -0.5% and $+0.25\%$, there is little change in volatility from that observed prior to day t. In summary, realized volatility rises

[13]Since the expected return on a given day is close to zero, we, like Bollerslev and Zhou (2006) and others, measure the standard deviation as the square root of the sum of the squared returns divided by N, the number of 5-min intervals, instead of the sum of squared deviations from the sample mean divided by $N-1$. Results are virtually the same if we use squared deviations from the sample mean and if we use 10-min or 15-min intervals.

[14]In contrast, Kane et al. (2000) find that squared daily returns are higher following very large positive price jumps, albeit not as high as following very large negative jumps.

Table 8
Impact of a 1-day return shock on realized volatility.

Means of the log percentage change in realized volatility stratified by r_t are reported. In column 3, means of $\ln(\sigma_{t+1}/\sigma_{t-1})$ are reported where σ_{t-1} and σ_{t+1} represent standard deviations of 5-min interval returns on days $t-1$ and $t+1$, respectively. Means of $\ln(\sigma_{t+1,t+5}/\sigma_{t-5,t-1})$ are in column 5, where $\sigma_{t-5,t-1}$ and $\sigma_{t+1,t+5}$ represent standard deviations of 5-min interval returns over the week before and after day t, respectively. Means of $\ln(\sigma_{t+1,t+21}/\sigma_{t-21,t-1})$ are reported in column 7, where $\sigma_{t-21,t-1}$ and $\sigma_{t+1,t+21}$ are measured as standard deviations of daily returns over the month prior to day t and after day t, respectively. T-values for the null hypothesis of no change are shown in columns 4, 6, and 8.

Day t return (%)	Obs.	Mean log percentage change in standard deviation of returns					
		From day before to day after: $\ln(\sigma_{t+1}/\sigma_{t-1})$ (%)	T-value	From week before to week after: $\ln(\sigma_{t+1,t+5}/\sigma_{t-5,t-1})$ (%)	T-value	From month before to month after: $\ln(\sigma_{t+1,t+21}/\sigma_{t-21,t-1})$ (%)	T-value
<-2.5	56	29.13	5.57	20.40	5.34	15.05	3.13
≥ -2.5 and <-2.0	54	19.09	4.02	17.07	4.56	8.66	1.86
≥ -2.0 and <-1.5	135	18.68	6.37	13.15	5.72	8.82	2.90
≥ -1.5 and <-1.0	239	14.41	6.81	8.46	5.32	4.05	1.81
≥ -1.0 and <-0.5	501	8.52	5.73	3.71	3.71	2.33	1.57
≥ -0.5 and <-0.25	391	1.92	1.08	2.20	1.85	−0.19	−0.12
≥ -0.25 and <0	533	0.89	0.58	−0.56	−0.53	−0.42	−0.30
≥ 0 and <0.25	571	−2.90	−1.98	−0.07	−0.07	1.77	1.35
≥ 0.25 and <0.5	451	−5.74	−3.52	−1.77	−1.52	−0.74	−0.49
≥ 0.5 and <1.0	594	−9.12	−6.49	−5.10	−5.24	−3.35	−2.44
≥ 1.0 and <1.5	281	−6.18	−3.02	−6.74	−4.87	−6.58	−3.50
≥ 1.5 and <2.0	120	−7.80	−2.45	−9.04	−4.22	−9.34	−2.81
≥ 2.0 and <2.5	52	−12.16	−2.37	−12.53	−3.71	−11.47	−2.72
≥ 2.5	55	−20.62	−4.21	−17.98	−5.17	−15.72	−3.59

following large negative returns, but falls following large positive returns and is little changed following a day when the market changes very little. Next, we examine how the log percentage change in realized volatility from the week before day t to the week after day t, $\ln(\sigma_{t+1,t+5}/\sigma_{t-5,t-1})$, varies with r_t. For this, $\sigma_{t-5,t-1}$ is measured as the standard deviation of 5-min returns over the five days from $t-5$ through $t-1$.[15] Means of $\ln(\sigma_{t+1,t+5}/\sigma_{t-5,t-1})$ are reported in column 5 of Table 8 with t-values in column 6. Results mirror those for the single day horizon in that realized volatility tends to rise following large negative returns, fall following large positive returns, and is little changed following near-zero returns.

Finally, we examine how the change in the standard deviation of returns over the next month is related to the day t return. For this, we switch back to daily return data from 5-min returns. Thus $\sigma_{t+1,t+21}$ represents the standard deviation of daily returns over the 21 days following the return shock and $\sigma_{t-21,t-1}$ over the preceding 21 days. Means of $\ln(\sigma_{t+1,t+21}/\sigma_{t-21,t-1})$ are reported in column 7 of Table 8. A single day's return has a surprisingly strong impact on volatility over the next month. When the stock index falls 2.5% or more on day t, realized volatility over the next month, as measured by the standard deviation of returns, rises by an average of 15.1%. When the stock index rises

[15]Since we are measuring the standard deviation of 5-min returns, overnight returns are excluded.

Table 9
The short-run reaction of realized volatility to market return shocks.
The log percentage change in realized volatility from day $t-1$ to day $t+1$, $\ln(\sigma_{t+1}/\sigma_{t-1})$, where σ_t is the standard deviation of 5-min returns on day t, is regressed on four variables: (1) the day t return, r_t; (2) $r(-)_t$, where $r(-)_t = r_t$ if $r_t < 0$ and zero otherwise; (3) $\ln(\sigma_t/\sigma_{t-1})$; and (4) $\ln(\sigma_{t-1}/\sigma_{t-2})$.

	Coefficient	*P*-value
Intercept	−0.0040	0.5368
r_t	−4.4921	<0.0001
$r(-)_t$	−1.6042	0.2138
$\ln(\sigma_t/\sigma_{t-1})$	0.4351	<0.0001
$\ln(\sigma_{t-1}/\sigma_{t-2})$	−0.2630	<0.0001
Adjusted R^2	0.354	

2.5% or more on day t, realized volatility falls by an average of 15.7%. Again there is little change in realized volatility following near-zero returns.

The percentage change in realized volatility over the month following a day t return shock is graphed along with the percentage change in implied volatility and the EGARCH forecast in Fig. 2. However, it should be pointed out that this series is not completely comparable with the other two since the bases are different. While the percentage changes in the VIX and EGARCH forecasts are from the forecasts on day $t-1$, the base for the percentage change in realized volatility is that observed over the $(t-21, t-1)$ interval. Nonetheless, the news impact curves in Fig. 2 serve to graphically summarize our basic findings. Like implied volatility and asymmetric GARCH model forecasts, realized volatility tends to rise following large negative returns. Like implied volatility and unlike the asymmetric time-series model forecasts, it tends to fall following positive returns and fall more the larger the positive return. Indeed, realized volatility tends to fall more sharply than implied volatility. Like the VIX and unlike the asymmetric GARCH model forecasts, realized volatility is little changed on average following days when the market rises or falls very little.

Having established that the change in realized volatility from day $t-1$ to day $t+1$ is a consistent negative function of the day t return, we are also interested in how it relates to the intraday volatility on day t.[16] For this, we estimate:

$$\ln(\sigma_{t+1}/\sigma_{t-1}) = \alpha_0 + \alpha_1 r_t + \alpha_2 r(-)_t + \alpha_3 \ln(\sigma_t/\sigma_{t-1}) + \alpha_4 \ln(\sigma_{t-1}/\sigma_{t-2}), \quad (6)$$

where σ_t represents the standard deviation of 5-min returns on day t, r_t is the day t return, and $r(-)_t$ is the day t return if negative and zero otherwise. $\mathrm{Ln}(\sigma_t/\sigma_{t-1})$ measures the change in intraday volatility from day $t-1$ to t. Volatility persistence implies $\alpha_3 > 0$. The lagged variable $\ln(\sigma_{t-1}/\sigma_{t-2})$ is included to allow for the possibility of mean reversion.

The results are reported in Table 9. As expected, $\hat{\alpha}_3 > 0$, indicating that, after controlling for the net return on day t, high(low) intraday volatility on day t leads to high (low) intraday volatility on day $t+1$. $\hat{\alpha}_1 < 0$, indicating that, controlling for intraday volatility,

[16]Moreover, days when the market rises or falls a lot are likely to be periods of high volatility and it is possible that days when the market declines are characterized by higher volatility than days when the market rises an equal amount. Hence, we would like to measure how r_t impacts future volatility controlling for intraday volatility.

a positive (negative) return on day t leads to a volatility decline (increase). The sign of $\hat{\alpha}_2$ is negative, which would indicate that the increase in realized volatility following a negative day t return is greater than the volatility decrease following a positive return but the difference is not significant at the 0.05 level.

7.2. Longer-run

Next we explore how realized volatility responds to market movements over a month. As in Table 9, we wish to separate the impact of the net return on future volatility from the impact of current intra-period volatility. For this we estimate monthly realized volatility equations analogous to those estimated in Table 9 for daily realized volatility and in Table 7 for implied volatility. First, we measure realized volatility over the $(t-41, t-21)$ and $(t+1, t+21)$ windows using daily returns and relate the change in volatility between these two periods to both the market return and volatility over the $(t-21, t)$ window. Specifically, we estimate:

$$\%\Delta \mathrm{Vol}_t = \alpha_0 + \alpha_1 r_{t-21,t} + \alpha_2 r(-1)_{t-21,t} + \alpha_3 \mathrm{SurpriseVol}_{t-21,t} + u_t, \tag{7}$$

where $\%\Delta\mathrm{Vol}_t = \ln[\sigma_{t+1,t+21}/\sigma_{t-41,t-21}]$, $\sigma_{t+1,t+21}$ is the standard deviation of daily returns r_t over the 21-day period from day $t+1$ to day $t+21$, and $\sigma_{t-41,t-21}$ is the standard deviation from day $t-41$ to day $t-21$. This is related to the market return over the $(t-21, t)$ interval, $r_{t-21,t}$, and to $\mathrm{SurpriseVol}_{t-21,t}$, which (as defined above and used in Table 7) is the difference between the standard deviation of returns over the $(t-20, t)$ interval and implied volatility for that period as observed on day $t-21$. As before, $r(-)_{t-21,t} = r_{t-21,t}$ if $r_{t-21,t} < 0$ and 0 otherwise.

The results are presented in Table 10. Considering first the univariate results in the first row, we find that the change in realized volatility is strongly negatively related to the market return over the intervening period. Indeed, the coefficient of $r_{t-21,t}$ exceeds that in the implied volatility regression in Table 7. In the multivariate estimation, $\hat{\alpha}_2 < 0$, indicating that realized volatility rises more following a negative return over the period from $t-21$ to t than it falls following a similar size positive return. Nonetheless, $\hat{\alpha}_1$ is significantly (0.05 level) negative, indicating that, like implied volatility, realized volatility falls following a positive return; $\hat{\alpha}_3$ is positive and significant, indicating that (as with implied volatility) higher (lower) than anticipated volatility over the intervening period is associated with an increase (decrease) in volatility.

Table 10
The longer-term reaction of realized volatility to market return shocks.

The log percentage change in realized volatility over a one month period measured as $\ln[\sigma_{t+1,t+21}/\sigma_{t-41,t-21}]$ is regressed on measures of the level and volatility of returns over the intervening month: (1) the S&P 500 index return from $t-21$ to t, (2) the return from $t-21$ to t if negative, and (3) the volatility surprise measured as the log percentage difference between the standard deviation of daily returns from $t-20$ to t and the VIX on day $t-21$. T-statistics are reported in parentheses using Newey–West standard errors.

	Intercept	Index return over previous month	Index return if negative	Volatility surprise	Adjusted R^2
Model 1	0.01798 (1.14)	−3.079 (−8.90)			0.129
Model 2	0.03699 (1.40)	−1.684 (−3.03)	−1.979 (−2.20)	0.162 (3.00)	0.150

In summary, the asymmetry in realized volatility roughly mirrors that in implied volatility and differs considerably from the pattern predicted by the asymmetric GARCH models. All other things equal, realized volatility tends to fall following periods with large positive returns and rise substantially following periods of large negative returns. Also, contrary to the predictions of the asymmetric GARCH models, there is no evidence that realized volatility declines following near-zero returns. This pattern is observed over both daily and monthly horizons. In results not reported here, we have replicated the monthly horizon results in Tables 7, 8, and 10 for two sub-periods: 1990–1997 and 1998–2005. The basic results hold in both sub-periods. To wit, time-series conditional, implied, and realized volatility all tend to rise sharply following large negative return shocks, conditional volatility falls but implied and realized do not following near zero returns, and implied and realized volatility fall sharply following large positive return shocks. The one difference is that the asymmetric time-series models predict an increase in volatility following large positive shocks in the 1990–1997 sub-period and a small decrease in the 1998–2005 sub-period.

The decline we document in implied and realized volatility following large positive return shocks is consistent with the findings of Kane et al. (2000) on the price of risk while the asymmetric GARCH model volatility predictions of a volatility increase are not. According to portfolio theory, an increase (decrease) in expected volatility will lead investors to require an increase (decrease) in expected returns to compensate. Consistent with this and the behavior of implied and realized volatility that we document, Kane et al. (2000) find that returns tend to rise following negative price jumps but fall following positive price jumps.

8. Exploring the differences

The finding that the behavior of realized volatility following positive and near-zero surprise returns differs considerably from that predicted by asymmetric GARCH models leads to the question of why—particularly since the models are estimated from realized volatility observations. One possibility is that the GJR and EGARCH models are mis-specified, i.e., that the true return-generating process does not match either. However, while the GJR and EGARCH models may not be the best specifications, with appropriate parameters, they are capable of predicting the patterns documented here. For example, a GJR model in which $\alpha_2 < 0$ would predict volatility declines following positive return shocks as would an EGARCH model in which $\beta_2 + \beta_3 < 0$. It is possible to specify parameter combinations that would also imply little change after near-zero returns. So why do the GJR and EGARCH estimations not choose these parameters?

Another possibility is that the metrics we have used to measure realized and conditional volatility on day $t+1$ are not identical. While the asymmetric GARCH models predict the volatility of daily returns for the next day, due to the difficulty in measuring daily return volatility from a single observation, our measure of realized volatility on day $t+1$ in Tables 8 and 9 is the standard deviation of 5-min, not daily, returns. It is possible that after a negative shock intraday volatility falls while daily volatility does not.

To more closely match our measure of realized volatility to that forecast by the asymmetric time-series models, we calculate the percentage forecast error in daily volatility as

$$\%\mathrm{FE}_{t+1} = \frac{(e_{t+1}^2 - \hat{\sigma}_{t+1}^2)}{\hat{\sigma}_{t+1}^2}, \tag{8}$$

Table 11
Characteristics of EGARCH volatility forecasts.

We explore how aspects of asymmetric time-series models' volatility forecasts (specifically EGARCH forecasts) vary with the standardized surprise return on day t. We group observations into 14 categories based on $e_t/\hat{\sigma}_t$, where e_t is the surprise return on day t and $\hat{\sigma}_t$ is the EGARCH volatility forecast for day t. The mean volatility forecast error $(e_{t+1}^2-\hat{\sigma}_{t+1}^2)/\hat{\sigma}_{t+1}^2$ is reported in column 3. Excess kurtosis for e_{t+1} is reported in column 4. In columns 5–8, the means of $\partial \text{LL}_{t+i}/\partial\hat{\sigma}_{t+1}$ are reported for $i=1-4$, where LL_{t+i} is the log likelihood for observation $t+i$ and $\hat{\sigma}_{t+1}$ is the EGARCH volatility forecast for day $t+1$. * and ** in columns 3 and 5–8 indicate means significantly different from zero at the 0.05 and 0.01 levels, respectively.

$e_t/\hat{\sigma}_t$	Obs.	Mean, $(e_{t+1}^2-\hat{\sigma}_{t+1}^2)/\hat{\sigma}_{t+1}^2$ (%)	Excess kurtosis, e_{t+1}	Log-likelihood derivative means, $\partial \text{LL}_{t+i}/\partial\hat{\sigma}_{t+1}$			
				$i=1$	$i=2$	$i=3$	$i=4$
<-2.5	41	−13.69	1.81	−388	1429	56	779
≥ -2.5 and <-2.0	66	13.36	0.54	99	5047	106	−364
≥ -2.0 and <-1.5	163	14.09	3.62	2051	495	103	−564
≥ -1.5 and <-1.0	327	1.14	1.56	294	−375	−266	776
≥ -1.0 and <-0.5	530	7.97	3.29	1135	301	581	−391
≥ -0.5 and <-0.25	406	13.52	0.23	1413	−105	896	−599
≥ -0.25 and <0	448	13.13	2.37	1923	541	−767	−324
≥ 0 and <0.25	449	7.90	−0.04	11	1717	1223	1703
≥ 0.25 and <0.5	434	−8.38	0.95	−1104	988	286	1192
≥ 0.5 and <1.0	587	−4.85	0.40	−429	264	−82	41
≥ 1.0 and <1.5	336	−22.39**	1.04	−1904*	−1446*	2411	562
≥ 1.5 and <2.0	159	−28.68*	2.62	−2011	−2827**	−2475**	1783
≥ 2.0 and <2.5	61	−37.13**	0.48	−2343	−2020	205	−370
≥ 2.5	25	−47.04**	−0.46	−5627**	−5066*	−215	2194

where e_{t+1} is the surprise return on day $t+1$ and $\hat{\sigma}_{t+1}^2$ is the EGARCH forecast of the surprise return variance on day $t+1$ based on returns through day t. In column 3 of Table 11, we report the means of $\%\text{FE}_{t+1}$ stratified by the standardized surprise return on day $t, e_t/\hat{\sigma}_t$, as in Table 5. While the high variation in e_{t+1}^2 around $\text{E}(e_{t+1}^2)$ leads to less powerful tests than in Table 8, the results in Table 11 indicate once again that EGARCH tends to substantially over-forecast the surprise return variance following positive surprise returns. Indeed, when $e_t/\hat{\sigma}_t>2.5$, e_{t+1}^2 averages about 47% below the EGARCH forecast. There is also weak evidence that EGARCH under-forecasts volatility following near-zero returns. Similar results are obtained for the GJR forecast. So this measure, like the others, indicates that the asymmetric time-series models tend to over-estimate realized volatility following positive returns.

Another possible explanation lies with the estimation criterion for the GARCH-type models. While our focus has been on the volatility forecast error, these seek to maximize the log-likelihood (LL) of observing the sample. Specifically, the EGARCH and GJR estimations choose parameters that maximize $\text{LL}=\sum_{t=1}^{T}\text{LL}_\text{t}=\sum_{t=1}^{T}\ln(f(e_t))$ where $f()$ is the normal density function that for observation e_t is $[1/2\pi\sigma_t^2]^{0.5}\exp[-0.5(e_t^2/\sigma_t^2)]$. Hence, $\text{LL}_t=-0.5[\ln(2\pi)+\ln(\sigma_t^2)+(e_t^2/\sigma_t^2)]$. Consider the impact on LL of a change in the volatility forecast for day $t+1$:

$$\partial \text{LL}/\partial\hat{\sigma}_{t+1}^2=\sum_{i=1}^{T-t}\partial \text{LL}_{t+i}/\partial\hat{\sigma}_{t+1}^2=\sum_{i=1}^{T-t}(\partial \text{LL}_{t+i}/\partial\hat{\sigma}_{t+i}^2)(\partial\hat{\sigma}_{t+i}^2/\partial\hat{\sigma}_{t+1}^2),$$

where $\partial \mathrm{LL}_{t+i}/\partial\hat{\sigma}^2_{t+i} = -0.5(1/\hat{\sigma}^2_{t+i})[1 - e^2_{t+i}/\hat{\sigma}^2_{t+i}]$. Consider the bracketed term $[1 - e^2_{t+i}/\hat{\sigma}^2_{t+i}]$. If $e^2_{t+1} = \hat{\sigma}^2_{t+1}$, $[1 - e^2_{t+i}/\hat{\sigma}^2_{t+i}] = 0$ and it varies between 1 if $|e_{t+1}|$ is small and ∞ if $|e_{t+1}|$ is very large. So, the log-likelihood estimation criterion puts very high weight on forecasting high volatility when $|e_{t+1}|$ is large and much less weight on correctly forecasting volatility when $|e_{t+1}|$ is small. Hence, if following large positive day t returns, most day $t+1$ returns are small in absolute value but a few are quite large, the log-likelihood estimations will tend to predict increases in volatility when that is rarely the case. As a partial test of this hypothesis, we calculated the excess kurtosis (i.e., minus the normal value of 3) of e_{t+1} for each of our 14 $e_t/\hat{\sigma}_t$ groupings in Table 11. As shown in column 4, there is little pattern across the different groups; so while we cannot rule this cause out, we cannot support it either.

Finally, we consider the relation between $\partial \mathrm{LL}_{t+i}/\partial\hat{\sigma}^2_{t+1}$ and $e_t/\hat{\sigma}_t$ for several i. In column 5 of Table 11, we report the means of $\partial \mathrm{LL}_{t+1}/\partial\hat{\sigma}^2_{t+1}$ for our fourteen $e_t/\hat{\sigma}_t$ groups. As with e^2_{t+1}, $\partial \mathrm{LL}_{t+1}/\partial\hat{\sigma}^2_{t+1}$ is highly variable so the t-statistics are low. Nonetheless, $\partial \mathrm{LL}_{t+1}/\partial\hat{\sigma}^2_{t+1}$ tends to be negative for positive values of $e_t/\hat{\sigma}_t$ and positive for negative and near-zero values of $e_t/\hat{\sigma}_t$, indicating that the $t+1$ part of the log-likelihood function could have been improved by predicting somewhat lower volatility following positive surprise returns. This is consistent with all our previous evidence indicating that the asymmetric time-series models tend to overestimate volatility following large negative returns. Clearly, if it were just a matter of predicting volatility on day $t+1$, the EGARCH estimation would have been different.

However, the objective of log-likelihood estimations of the time-series models is to maximize LL, not LL_{t+1}, and lowering the forecast variance for $t+1$ entails lowering it for future i as well. $\partial \mathrm{LL}_{t+\mathrm{i}}/\partial\hat{\sigma}^2_{t+1} = (\partial \mathrm{LL}_{t+i}/\partial\hat{\sigma}^2_{t+i})(\partial\hat{\sigma}^2_{t+i}/\partial\hat{\sigma}^2_{t+1})$ and for the EGARCH and GJR estimations, $(\partial\hat{\sigma}^2_{t+i}/\partial\hat{\sigma}^2_{t+1})$ declines slowly so observations in the distant future have almost as much impact on $\hat{\sigma}_{t+1}$ as observations at $t+1$. For instance, for our estimations of the EGARCH model, $(\partial\hat{\sigma}^2_{t+i}/\partial\hat{\sigma}^2_{t+1})$ averages 0.986 for $i=2$, 0.947 for $i=5$, 0.885 for $i=10$, and 0.520 for $i=50$. In columns 6–8 of Table 11, we report means of $\partial \mathrm{LL}_{t+i}/\partial\hat{\sigma}^2_{t+1}$ for $i=2-4$. For $i=2$, there is again evidence that the log-likelihood could be improved by predicting lower volatility following negative returns but after that the pattern breaks down. Similar results are observed for $i>4$. In summary, our results suggest that the tendency of the asymmetric GARCH models to over-predict volatility following large positive returns may be due to the heavy weight that the log-likelihood function attaches to extreme return observations and to the high persistence in the volatility forecasts.

9. Conclusions

We find that U.S. stock market volatility is highly asymmetric and that the degree of asymmetry differs sharply depending on whether measured using volatility predictions from asymmetric time-series models or either implied or realized volatility. According to all three volatility measures, volatility rises sharply following large negative returns but the measures differ regarding volatility changes following positive and near-zero returns. While time-series models predict a small increase in volatility following large positive returns, both implied and realized volatilities decline substantially from levels observed prior to the bull market. This result holds for both one-day and one-month horizons. While the time-series models predict that volatility should fall following near-zero returns, implied and realized volatilities are little changed from levels observed prior to the near-zero return. Possible reasons for the differences between asymmetric GARCH estimations of conditional volatility and realized volatility include the heavy weight that the GARCH estimations place on extreme observations and on the likelihood specification. In summary, theories seeking to explain asymmetric volatility in the U.S. stock market face a sterner task than apparent heretofore. Instead of just explaining why volatility rises more following negative shocks than positive shocks, they need to explain why volatility declines following large positive shocks and is roughly unchanged when net returns are close to zero.

Acknowledgments

The authors wish to acknowledge the helpful comments received in presentations of earlier versions of this paper at the University of Oklahoma and the 2008 European meetings of the Financial Management Association. They also wish to thank the referee and the editor, Bruce Lehmann, as well as Guan Jun Wang for their helpful suggestions. Part of this research was completed while Ederington was a Visiting Professor at the University of Queensland and Melbourne University. He thanks them for their assistance.

References

Avramov, D., Chordia, T., Goyal, A., 2006. The impact of trades on daily volatility. Review of Financial Studies 19, 1241–1277.

Bates, D., 2000. Post-87 crash fears in the S&P 500 futures option market. Journal of Econometrics 94, 181–238.

Bekaert, G., Wu, G., 2000. Asymmetric volatility and risk in equity markets. Review of Financial Studies 13, 1–42.

Black, F., 1976. Studies in stock price volatility changes. In: Proceedings of the 1976 Business Meeting of the Business and Economics Section. American Statistical Association, pp. 177–181.

Bollerslev, T., Zhou, H., 2006. Volatility puzzles: a simple framework for gauging return-volatility regressions. Journal of Econometrics 131, 123–150.

Bollerslev, T., Litvinova, J., Tauchen, G., 2006. Leverage and volatility feedback effects in high-frequency data. Journal of Financial Econometrics 4, 353–384.

Britten-Jones, M., Neuberger, A., 2000. Option prices, implied price processes, and stochastic volatility. Journal of Finance 55, 839–866.

Campbell, J., Hentschel, L., 1992. No news is good news: an asymmetric model of changing volatility in stock returns. Journal of Financial Economics 31, 281–318.

Caporin, M., McAleer, M., 2006. Dynamic asymmetric GARCH. Journal of Financial Econometrics 4, 385–412.

Christie, A., 1982. The stochastic behavior of common stock variances—value, leverage, and interest rate effects. Journal of Financial Economics 10, 407–432.

Demeterfi, K., Derman, E., Kamal, M., Zou, J., 1999. More than you ever wanted to know about volatility swaps. Journal of Derivatives 6, 9–32.

Denis, P., Mayhew, S., Stivers, C., 2006. Stock returns, implied volatility innovations, and the asymmetric volatility phenomenon. Journal of Financial and Quantitative Analysis 41, 381–406.

Ederington, L., Guan, W., 2009a. Longer-term time series volatility forecasts. Journal of Financial and Quantitative Analysis, forthcoming.

Ederington, L., Guan, W., 2009b. Bias in time series model volatility forecasts. The Journal of Futures Markets, forthcoming.

Engle, R., Lilien, D., Robins, R., 1987. Estimating time varying risk premia in the term structure: the ARCH-M model. Econometrica 55, 391–407.

Engle, R., Ng, V., 1993. Measuring and testing the impact of news on volatility. Journal of Finance 48 1749–1778

Fleming, J., Ostdiek, B., Whaley, R., 1995. Predicting stock market volatility: a new measure. Journal of Futures Markets 15, 265–302.

French, K., Schwert, G.W., Stambaugh, R., 1987. Expected stock returns and volatility. Journal of Financial Economics 19, 3–30.

Ghysels, E., Santa-Clara, P., Valknov, R., 2004. There is a risk-return tradeoff after all. NBER Working Paper W10913.

Glosten, L., Jagannathan, R., Runkle, D., 1993. On the relation between the expected value and volatility of the nominal excess return on stocks. Journal of Finance 48, 1779–1801.

Jiang, G., Tian, Y., 2005. The model free implied volatility and its information content. Review of Financial Studies 18 (4), 1305–1342.

Kane, A., Lehmann, B., Trippi, R., 2000. Regularities in volatility and the price of risk following large stock market movements in the US and Japan. Journal of International Money and Finance 19, 1–32.

Koulakiotis, A., Papasyriopoulos, N., Molyneux, P., 2006. More evidence on the relationship between stock price returns and volatility: a note. International Research Journal of Finance and Economics 1, 21–28.

Lettau, M., Ludvigson, S., 2009. Measuring and modeling variation in the risk-return tradeoff. In: Ait-Sahalia, Y., Hansen, L. (Eds.), Handbook of Financial Econometrics. Elsevier, Amsterdam.

Li, Q., Yang, J., Hsiao, C., Chang, Y., 2005. The relationship between stock returns, and volatility in international stock markets. Journal of Empirical Finance 12, 650–665.

Low, C., 2004. The fear and exuberance from implied volatility of S&P 100 index. Journal of Business 77, 52–546.

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

Pan, J., 2002. The jump risk premia implicit in options: evidence from an integrated time-series study. Journal of Econometrics 63, 3–50.

Poon, S.H., Granger, C., 2003. Forecasting volatility in financial markets. Journal of Economic Literature 41, 478–539.

Poteshman, A., 2001. Underreaction, overreaction, and increasing misreaction to information in the options market. Journal of Finance 56, 851–876.

Schwert, G.W., 1989. Why does stock market volatility change over time?. Journal of Finance 44, 1115–1153.

Whaley, R., 2000. The investor fear gauge. Journal of Portfolio Management 26, 12–17.

Wu, G., 2001. The determinants of asymmetric volatility. Review of Financial Studies 14, 837–859.

Wu, G., Xiao, Z., 2002. A generalized partially linear model of asymmetric volatility. Journal of Empirical Finance 9, 287–319.