
THE SPECIFICATION OF GARCH MODELS WITH STOCHASTIC COVARIATES

JEFF FLEMING
CHRIS KIRBY*
BARBARA OSTDIEK

A number of studies investigate whether various stochastic variables explain changes in return volatility by specifying the variables as covariates in a GARCH(1, 1) or EGARCH(1, 1) model. The authors show that these models impose an implicit constraint that can obscure the true role of the covariates in the analysis. They illustrate the problem by reconsidering the role of contemporaneous trading volume in explaining ARCH effects in daily stock returns. Once the constraint imposed in earlier research is relaxed, it is found that specifying volume as a covariate does little to diminish the importance of lagged squared returns in capturing the dynamics of volatility. © 2008 Wiley Periodicals, Inc. *Jrl Fut Mark* 28:911–934, 2008

An earlier version of this article was circulated under the title “ARCH Effects and Trading Volume.” The authors thank Scott Baggett for his help in constructing the price and volume data sets from Trade and Quote (TAQ) data and Joel Hasbrouck and James Weston for providing useful insights and advice regarding TAQ data. The authors also thank the seminar participants at the Rice University Econometrics Workshop, the Bank of Canada, and the University of Texas at San Antonio.

*Correspondence author, John E. Walker Department of Economics, Clemson University, Box 341309, Clemson, South Carolina 29634-1309, e-mail: cmkirby@clemson.edu

Received August 2007; Accepted November 2007

-
- *Jeff Fleming is a Professor of Finance and Barbara Ostdiek is an Associate Professor of Finance at the Jones Graduate School of Management, Rice University, Houston, Texas.*
 - *Chris Kirby is an Associate Professor of Economics at the John E. Walker Department of Economics, Clemson University, Clemson, South Carolina.*

INTRODUCTION

The extensive use of generalized autoregressive conditional heteroscedasticity (GARCH) models in financial economics is testimony to their success in capturing volatility dynamics. As low-order GARCH and exponential GARCH (EGARCH) models typically perform well relative to more complex specifications (see, e.g., Hansen & Lunde, 2005a), researchers often use these models to investigate the relation between changes in return volatility and various stochastic variables. In particular, they assess whether the stochastic variables explain changes in volatility by including the variables as covariates in a GARCH(1, 1) or EGARCH(1, 1) model. The variables considered in the literature include interest rate levels (Engle & Patton, 2001; Glosten, Jagannathan, & Runkle, 1993), interest rate spreads (Dominguez, 1998; Hagiwara & Herce, 1999), forward-spot spreads (Hodrick, 1989), implied volatilities (Blair, Poon, & Taylor, 2001; Day & Lewis, 1992; Lamoureux & Lastrapes, 1993), futures open interest (Girma & Mougoue, 2002), a proxy for the information flow during the overnight market closure (Gallo & Pacini, 2000), and contemporaneous trading volume (Fujihara & Mougoue, 1997; Lamoureux & Lastrapes, 1990; Marsh & Wagner, 2005).

Here a closer look is taken at the specification of GARCH models with stochastic covariates, highlighting a specification issue that makes it difficult to draw reliable inferences from many of the models considered in the literature. These models impose an implicit constraint that requires the coefficients on the lagged squared returns and the lagged stochastic variables to decline with the lag length at the same rate. This constraint is problematic for cases in which the stochastic variables provide little information about future return volatility beyond that contained in lagged squared returns. Obtaining precise fitted volatilities in such cases requires giving little weight to the lagged stochastic variables, but owing to the implicit constraint, this also requires giving little weight to lagged squared returns. Hence, if there is a strong contemporaneous relation between the stochastic variables and return volatility, the covariates can drive ARCH effects out of the fitted models regardless of whether they capture volatility persistence.

The problem is illustrated by reconsidering the role of contemporaneous trading volume in explaining ARCH effects in daily stock returns. Lamoureux and Lastrapes (1990) examined this issue by fitting a stochastic-covariate GARCH (SC-GARCH) model for 20 firms. They found that the coefficient on contemporaneous trading volume is highly significant and concluded that lagged squared returns provide little if any information about return volatility beyond that contained in volume. The most common criticism of this approach is that it requires treating volume as exogenous to establish consistency of the

SC-GARCH estimator (see, e.g., Fleming, Kirby, & Ostdiek, 2006). However, regardless of whether the estimator is consistent, using a SC-GARCH(1, 1) model to investigate the extent to which volume explains ARCH effects may be inappropriate in light of the implicit constraint. To investigate this issue, the performance of an SC-EGARCH(1, 1) model is compared with that of a higher-order SC-GARCH model that has the flexibility to simultaneously capture both transitory and persistent volatility shocks. As the model is a version of a standard SC-EGARCH(2, 2) process that nests the SC-EGARCH(1, 1) model as a special case, it allows one to directly test the impact of relaxing the constraint.

The SC-EGARCH(2, 2) model is fitted to daily returns for the 20 stocks in the major market index (MMI). The results of the analysis are consistent with those reported by Fleming et al. (2006). Specifically, it is found that volume is strongly correlated with contemporaneous return volatility, but the correlation is driven by transitory shocks to the volatility process, that are largely unrelated to the persistent component of volatility captured by standard volatility models. No support is found for the hypothesis that inserting volume into the conditional variance function of the model reduces the importance of lagged squared returns in capturing volatility dynamics. Similar findings were obtained by Liesenfeld (2001) using a generalized bivariate mixture model and by Gillemot, Farmer, and Lillo (2005) using a nonparametric specification.

The relative forecasting performance of the various models is also investigated. Following Andersen and Bollerslev (1998), a regression of the realized variances constructed from intraday returns on the fitted variances produced by each model confirms that the SC-EGARCH(1, 1) model does a poor job of capturing the relation between volume and ARCH effects. When the realized variances are regressed on the fitted variances from the SC-GARCH(1, 1) model, much lower R^2 values are obtained than when the fitted variances from the basic EGARCH(1, 1) model are used. In contrast, the fitted variances from the SC-EGARCH(2, 2) model outperform those from both of these models as well as those from the basic EGARCH(2, 2) model. The R^2 values are in line with those generated by the bivariate stochastic autoregressive volatility specification examined in Fleming et al. (2006). As the superior performance of the SC-EGARCH(2, 2) model is primarily attributable to the undiminished role of ARCH effects, the results suggest that one must look beyond volume to identify the features of the trading process that give rise to ARCH effects in daily stock returns.

The rest of the study is organized as follows. In the second section the specification of GARCH models with stochastic covariates is discussed, the potential problems with the SC-GARCH(1, 1) specification are highlighted, and the two-component SC-EGARCH(2, 2) model used in the empirical analysis is

introduced. In the third section the data set is described, the construction of the realized variances used for the model comparisons is explained, and the empirical results are presented. The last section offers a few concluding remarks.

BACKGROUND AND METHODOLOGY

Low-order GARCH and EGARCH models are among the most widely applied models in economics and finance. Indeed, they have become benchmarks in the volatility forecasting literature.¹ It is not surprising, therefore, that a number of studies investigate whether various stochastic variables explain changes in return volatility by including them as covariates in a GARCH(1, 1) or EGARCH(1, 1) model. Despite the appeal of the methodology, it is subject to a specification issue that has gone unrecognized in the literature. The issue arises from an implicit constraint that makes it difficult to draw reliable inferences from the model-fitting results. To illustrate, suppose a model of the following form is specified:

$$R_t = \mu + \sqrt{h_t} z_t \quad (1)$$

$$h_t = \omega + \beta h_{t-1} + \alpha r_{t-1}^2 + \gamma X_t \quad (2)$$

where R_t is the daily stock return, X_t is the covariate of interest, z_t is an i.i.d. $N(0, 1)$ standardized innovation, and $r_t = R_t - \mu$ is the demeaned return. Expanding the expression for h_t via recursive substitution yields

$$h_t = \omega \sum_{i=1}^{\infty} \beta^{i-1} + \sum_{i=1}^{\infty} \beta^{i-1} (\alpha r_{t-i}^2 + \gamma X_{t-i}) + \gamma X_t. \quad (3)$$

Thus, the coefficient on X_{t-i} is constrained to decline with the lag length i at the same rate as the coefficient on r_{t-i}^2 .

To see why this constraint is an important issue, suppose there is a positive correlation between the volatility of daily stock returns and the contemporaneous realization of the covariate. This is consistent with $\gamma > 0$. If one fits the model and finds that the estimate of γ is positive and statistically significant, but the estimates of α and β are statistically indistinguishable from zero, should this be interpreted as support for the hypothesis that the covariate explains ARCH effects? The answer is unclear given the constraint. Suppose

¹Hansen and Lunde (2005a), for example, compared the forecasting performance of 330 ARCH-type models to that of the GARCH(1, 1) model. They found no evidence that more sophisticated models outperform the GARCH(1, 1) model in forecasting the volatility of daily currency returns.

that the lags of X_t provide no information about return volatility beyond that contained in lagged squared returns. In this case, putting any weight on the lags of X_t adds noise to the conditional variances. However, if one wants to allow for a relation between h_t and X_t while giving zero weight to the lagged X_t s, the constraint forces β to equal zero. Therefore, if the relation between h_t and X_t is strong enough, the covariate can drive ARCH effects out of the model even if it does not capture volatility persistence. This specification issue is a concern for any investigation based on an SC-GARCH(1, 1) or SC-EGARCH(1, 1) model.

Using a Two-Component Model to Relax the Constraint

The problem with the model in Equations (1) and (2) is that it constrains the joint dynamics of return volatility and the covariate in a way that may not be empirically plausible. A straightforward remedy is proposed: base the analysis on a higher-order GARCH model that allows for more complex dynamics. As the attention is restricted to the case of a single covariate, a model with two sources of volatility shocks that could potentially generate two different levels of volatility persistence is developed. Specifically, a two-component SC-GARCH model that is both parsimonious and capable of producing the necessary complex dynamics is proposed.

The SC-GARCH model is specified in terms of logarithms to avoid problems with enforcing nonnegativity. In particular, an SC-EGARCH model of the following form is considered:

$$r_t = \sqrt{h_t} z_{rt} \quad (4)$$

$$\Delta \log h_t = \Delta m_t + \kappa_h(m_{t-1} - \log h_{t-1}) + \sigma_h u_{t-1} + \gamma_h w_t \quad (5)$$

$$\Delta m_t = \kappa_m(s - m_{t-1}) + \sigma_m u_{t-1} + \gamma_m w_t \quad (6)$$

where Δ denotes the first-difference operator, $u_t = (|z_{rt}| - E[|z_{rt}|]) / \sqrt{\text{var}(|z_{rt}|)}$, and $w_t = (\log X_t - E[\log X_t]) / \sqrt{\text{var}(\log X_t)}$. To see the origins of the model, suppose $\kappa_m = \kappa_h$ and $\gamma_h = \sigma_m = \gamma_m = 0$. In this case, Equations (4)–(6) collapse to

$$r_t = \sqrt{h_t} z_{rt} \quad (7)$$

$$\Delta \log h_t = \kappa_h(s - \log h_{t-1}) + \sigma_h u_{t-1} \quad (8)$$

which is simply an EGARCH(1, 1) model expressed in a form that lends a convenient interpretation to each parameter.² Specifically, s is the unconditional mean of $\log h_t$, κ_h determines the speed at which $\log h_t$ reverts toward s , and σ_h is the volatility of the innovations to $\log h_t$.³

Now consider the full model in Equations (4)–(6). Its underlying structure is still that of an EGARCH specification, but instead of reverting toward a fixed mean s , the log variance is pulled toward a stochastic mean m_t whose dynamics are described by an autoregressive process. The idea behind this generalization, which follows Engle and Lee (1999), is that m_t captures low-frequency variations in volatility, whereas high-frequency variations are captured by $\log h_t - m_t$. This gives the model the flexibility to incorporate volatility shocks that decay at a different rate than ARCH effects. Suppose, for example, that the covariate has a very transitory impact on volatility, with most of the volatility persistence owing to ARCH effects. One would expect to find that κ_h is large relative to κ_m , σ_h is small relative to γ_h , and σ_m is large relative to γ_m .

Additional insights into the dynamic properties of the model are gained by expressing the conditional variance function in a way that eliminates m_t from explicit consideration. This is accomplished by substituting Equation (6) into Equation (5), and then substituting for m_{t-1} in the resulting expression using the original Equation (5). After consolidating terms

$$\Delta \log h_t = \kappa_1(s - \log h_{t-1}) + \kappa_2(s - \log h_{t-2}) + \sigma_1 u_{t-1} + \sigma_2 u_{t-2} + \gamma_1 w_t + \gamma_2 w_{t-1} \quad (9)$$

is obtained, where $\kappa_1 = \kappa_h + \kappa_m - 1$, $\kappa_2 = (1 - \kappa_h)(1 - \kappa_m)$, $\sigma_1 = \sigma_h + \sigma_m$, $\sigma_2 = -(\kappa_h \sigma_m + \kappa_m \sigma_h)$, $\gamma_1 = \gamma_h + \gamma_m$, and $\gamma_2 = -(\kappa_h \gamma_m + \kappa_m \gamma_h)$. Hence, the model has an SC-EGARCH(2, 2) representation. This representation highlights both the similarities and the differences between the methodology used here and the approach used in earlier research. Both add stochastic covariates to a standard GARCH model, but the approach here specifies a model with sufficient flexibility to simultaneously capture both transitory and persistent volatility shocks. This flexibility should produce more robust results.

Model Comparisons Using Realized Variances

Ultimately the objective is to compare how well the different models capture the dynamics of volatility. In the empirical investigation these comparisons are conducted using realized variances. The concept of realized variance was introduced

²Unlike the EGARCH specification of Nelson (1991), the model above does not allow for leverage effects. This is simply for ease of exposition; allowing them does not have much impact on the findings.

³ $|z_{it}|$ and $\log X_t$ are standardized to make it easier to compare the coefficient estimates across firms. This has no effect on the dynamic implications of the model.

by Merton (1980). Let $R_{t,i,m}$, $i = 1, \dots, m$, denote the intraday returns on day t over m equally spaced intervals. The realized variance on day t is the sum of the squared returns:

$$RV_t = \sum_{i=1}^m R_{t,i,m}^2. \quad (10)$$

The realized variance should be close to the true variance provided that certain conditions are satisfied. For example, if returns are generated by a continuous-time process with instantaneous volatility σ_t , then it is natural to use the integrated variance $IV_t = \int_0^1 \sigma_{t+\tau}^2 d\tau$ as a measure of the daily variance. Under weak regularity conditions, $RV_t - IV_t \rightarrow 0$ almost surely as $m \rightarrow \infty$ (for details, see Andersen, Bollerslev, Diebold, & Labys, 2001; Barndorff-Nielsen & Shephard, 2002). This suggests that by increasing the frequency at which the returns are sampled, consistent nonparametric estimates of the integrated variance that in principle are arbitrarily efficient can be constructed.⁴

Andersen and Bollerslev (1998) used realized variances to assess whether standard volatility models generate accurate forecasts. Their approach consisted of regressing the realized variances on the fitted variances produced by a volatility model estimated using daily returns.⁵ To apply their approach, regressions of the following form are fitted:

$$\log RV_t = a + b \log \hat{h}_t + e_t \quad (11)$$

where $\log \hat{h}_t$ denotes the fitted log variance for day t produced by one of the EGARCH models. Although the regression R^2 will be biased toward zero because the variance of $\log RV_t$ is greater than the variance of the true log volatility (for details, see Andersen, Bollerslev, & Meddahi, 2005), this does not affect the model comparisons because the ratio of the R^2 values produced by different models is bias free.

EMPIRICAL ILLUSTRATION USING TRADING VOLUME

The importance of the constraint in practice is illustrated by reexamining the role of trading volume in explaining ARCH effects in daily stock returns. This

⁴Obviously the true price process is unobservable in practice and realized variances constructed according to Equation (10) can be biased by the influence of microstructure effects on observed prices and the absence of high-frequency returns during the nontrading periods overnight and on weekends. The approach for dealing with these issues is discussed in the next section.

⁵An alternative approach would be to use measures based on absolute values (realized absolute value or realized power) for this purpose. Some researchers argue that these measures outperform realized variances in predicting changes in quadratic variation. See, for example, Forsberg and Ghysels (2007) and Ghysels, Santa-Clara, and Valkanov (2006).

application is especially interesting because conflicting results reported in the literature can potentially be explained by the presence of the constraint. In an early and influential study, Lamoureux and Lastrapes (1990) added contemporaneous trading volume to a GARCH(1, 1) model and found that the coefficients on lagged squared returns become statistically insignificant for 16 of the 20 firms in their sample. They concluded that lagged squared returns contain little information about return volatility beyond that contained in trading volume. Other studies, such as Liesenfeld (1998) and Fleming et al. (2006), used different methodologies and found starkly different results. Although these researchers take issue with the Lamoureux and Lastrapes (1990) analysis by arguing that trading volume is endogenous, the question of whether the SC-GARCH(1, 1) model used by Lamoureux and Lastrapes (1990) delivers reliable inferences apart from the bias induced by endogeneity has not been considered. This question is addressed by investigating the impact of relaxing the constraint.⁶

The Data Set

The 20 stocks in the MMI are used for the empirical analysis.⁷ These stocks are widely held by both individual and institutional investors and generally exhibit a high level of trading activity. The data set consists of daily returns, trading volumes, and realized variances. Intraday observations are obtained on transaction prices and trading volume from the Trade and Quote (TAQ) database of the New York Stock Exchange and information on daily returns, stock splits, and dividends from the Center for Research in Security Prices (CRSP) daily stock price file. The sample period is January 4, 1993–December 31, 2003 (2,770 observations).⁸

In the TAQ database, records with an out-of-sequence time stamp, a zero price, a correction code greater than two (indicating errors and corrections), or a condition code (indicating nonstandard settlement) are deleted. In addition, two screens intended to identify and eliminate price reporting errors are

⁶As the consequences of ignoring endogeneity are well known, the attention is confined to the impact of relaxing the constraint. Like any GARCH model with contemporaneous volume as a covariate, the specification may produce biased parameter estimates. However, the dynamics implied by the SC-EGARCH(2, 2) estimates are similar to those documented by Fleming et al. (2006) using a state-space approach that accounts for the endogeneity of volume, and any bias should have little impact on the model comparisons that are the main focus of the analysis.

⁷These firms are American Express (AXP), AT&T (T), ChevronTexaco (CVX), Coca-Cola (KO), Disney (DIS), Dow Chemical (DOW), DuPont (DD), Eastman Kodak (EK), Exxon-Mobil (XOM), General Electric (GE), General Motors (GM), International Business Machines (IBM), International Paper (IP), Johnson & Johnson (JNJ), McDonald's (MCD), Merck (MRK), 3M (MMM), Philip Morris (MO), Procter and Gamble (PG), and Sears (S).

⁸Philip Morris did not open on May 25, 1994, in advance of a board meeting regarding a proposal to split the firm's food and tobacco businesses. This date is excluded from the sample.

applied. First, prices that are more than 20% higher or lower than the previous transaction price are excluded. Second, prices that imply a price change greater than two percent in magnitude that are immediately followed by a price reversal greater than two percent in magnitude are flagged. The flagged price is excluded if the implied price change is more than two times the next largest price change for the day, or if the price falls outside the day's high–low range (ignoring the flagged price) by more than the next largest price change for the day. The remaining TAQ records are used to construct the data set.

To calculate daily trading volume, the volume for all transactions in the day is aggregated. This figure is adjusted for stock splits and stock dividends using information from the CRSP daily stock price file and then detrended to obtain the volume series used to fit the models.⁹ The daily realized variances are estimated by combining the trading-day realized variance with the overnight squared return using the approach of Hansen and Lunde (2005b). This is accomplished in two steps. First, the trading-day realized variance is obtained using the Newey and West (1987) correction proposed by Hansen and Lunde (2004). This yields an unbiased estimator of the integrated variance even for returns sampled at very high frequencies. A 30-second sampling frequency is used for returns and a 30-minute window length is used for the Newey–West correction. Second, a weighted sum of the trading-day realized variance and the squared nontrading-period return is taken to obtain the full-day realized variance.¹⁰ The weights placed on the trading- and nontrading-period variance estimators follow Hansen and Lunde (2005b). Details are in the Appendix.

Estimation and Inference for the EGARCH(1, 1) Models

The analysis begins by fitting the basic EGARCH(1, 1) model in Equations (7) and (8). Specifically, the parameters are estimated via maximum likelihood and the Bollerslev and Wooldridge (1992) approach is used to compute robust standard errors. Table I reports the parameter estimates and *t*-ratios along with several specification diagnostics. As expected, the model only partially accounts for the fat tails that characterize the distribution of daily returns. The excess

⁹The volume series is detrended by extracting a quadratic time trend via ordinary least squares regression. Although other methods would provide more flexibility in fitting the trend, they would be more prone to overfitting as well, which could inadvertently remove components of volume that are important to the volume–volatility relation. Replacing the adjusted volume series with the original series has little impact on the empirical results.

¹⁰The nontrading-period return, which is computed using the last transaction price for the previous day and the first price on the current day, is adjusted for cash dividends and stock distributions reported in the CRSP database.

TABLE I
EGARCH(1, 1) Model

Firm	Estimates			t-Ratios			Diagnostics			
	s	κ_h	σ_h	s	κ_h	σ_h	\mathcal{L}	R^2	ρ	C_K
AXP	1.56	0.03	0.11	12.2	2.7	5.5	-5,893.3	0.12	0.98	1.35
CVX	0.89	0.02	0.07	6.8	1.8	3.9	-4,950.8	0.06	0.98	0.87
DD	1.45	0.01	0.07	7.1	1.2	2.4	-5,512.3	0.09	0.99	2.01
DIS	1.87	0.01	0.07	7.6	1.3	2.1	-5,873.6	0.09	0.99	5.82
DOW	1.61	0.01	0.08	5.0	1.8	3.9	-5,352.6	0.15	0.99	1.78
EK	1.48	0.34	0.22	12.3	2.3	3.9	-5,820.4	0.03	0.51	15.69
GE	1.05	0.01	0.07	2.9	1.1	3.1	-5,280.5	0.13	0.99	1.26
GM	1.48	0.03	0.08	13.6	2.1	4.3	-5,798.8	0.07	0.98	1.44
IBM	1.92	0.02	0.08	8.7	2.0	3.8	-5,984.5	0.05	0.98	4.02
IP	1.47	0.01	0.05	6.2	1.0	2.1	-5,630.9	0.10	0.99	1.44
JNJ	1.08	0.02	0.09	7.2	2.3	4.3	-5,181.9	0.08	0.98	1.61
KO	1.21	0.01	0.06	4.9	1.6	3.7	-5,173.9	0.10	0.99	1.81
MCD	1.48	0.01	0.06	6.8	2.3	4.0	-5,417.6	0.06	0.99	2.25
MMM	1.20	0.01	0.04	7.4	1.7	2.3	-5,107.4	0.07	0.99	3.98
MO	2.95	0.00	0.04	1.5	0.7	4.0	-5,870.4	0.05	0.99	10.16
MRK	1.38	0.01	0.04	7.7	1.3	2.7	-5,531.5	0.05	0.99	2.18
PG	1.56	0.00	0.06	2.7	1.9	5.6	-5,175.5	0.09	0.99	10.72
S	1.90	0.03	0.11	11.2	1.9	3.4	-6,150.0	0.07	0.96	4.02
T	2.17	0.01	0.06	6.9	1.1	1.7	-5,909.6	0.13	0.99	7.59
XOM	0.87	0.01	0.07	3.8	2.0	4.4	-4,798.7	0.11	0.99	1.09

Note. The table reports the results of fitting an EGARCH(1, 1) model to daily percentage returns on the MMI stocks. The model is of the form

$$r_t = \sqrt{h_t} z_t$$
$$\Delta \log h_t = \kappa_h (s - \log h_{t-1}) + \sigma_h u_{t-1}$$

where r_t is the demeaned return for day t , $u_t = (|z_t| - E[|z_t|]) / \sqrt{\text{var}(|z_t|)}$, and $z_t \sim NID(0,1)$. The model is fitted via maximum likelihood. The table reports the parameter estimates, the associated t -ratios, the maximized value of log-likelihood (\mathcal{L}), the sample R^2 for a regression of $\{r_t\}_{t=1}^T$ on the fitted conditional volatilities (R^2), the first-order sample autocorrelation of the fitted conditional volatilities (ρ), and the coefficient of excess kurtosis for the standardized returns (C_K). The t -ratios are based on robust standard errors. The sample period is January 5, 1993–December 31, 2003. MMI, major market index.

kurtosis of the standardized returns is positive for all firms, with especially large values for Eastman Kodak, Phillip Morris, and Procter and Gamble. Finding a few large values is not unusual, however, given the extreme returns that occasionally occur for individual stocks.

The results clearly indicate high levels of volatility persistence. The estimate of κ_h is close to zero for all firms except Eastman Kodak and only seven of the estimates have t -ratios of 2.0 or greater. In addition, the first-order sample autocorrelation of the fitted conditional volatilities is 0.96 or higher for all firms except Eastman Kodak. The diagnostics suggest that the low estimate of persistence for Eastman Kodak is probably due to a small number of influential observations. Eastman Kodak has the largest excess kurtosis of any firm, and an

examination of the data reveals several instances of daily returns between 10 and 20% in magnitude. Overall, the model-fitting results are consistent with those of previous studies in the volatility modeling literature (see, e.g., Kim & Kon, 1994).

The model itself does not appear to have much explanatory power. The R^2 value for a regression of the absolute demeaned returns on the fitted conditional volatilities ranges from 3% for Eastman Kodak to 15% for Dow Chemical.¹¹ Of course, as Andersen and Bollerslev (1998) pointed out, such regressions are expected to produce relatively low R^2 values because absolute returns are a noisy proxy for volatility. The realized variance regressions considered later provide a better benchmark for assessing how well the EGARCH(1, 1) specification captures volatility dynamics.

Next the EGARCH(1, 1) model in which contemporaneous volume is specified as a covariate is considered. Table II reports the model-fitting results. The most striking change from Table I is a sharp increase in the estimates of κ_h together with a sharp decline in the estimates of ρ , the first-order sample autocorrelation of the fitted conditional volatilities. The estimate of κ_h exceeds one for 19 of the 20 stocks and most of the estimates are highly statistically significant. The largest ρ estimate now is just 0.63 (Dow Chemical), whereas the smallest estimate is 0.25 (ChevronTexaco). Clearly, the addition of volume as an explanatory variable produces a marked drop in the degree of volatility persistence implied by the model.

More generally, the results confirm that volume is a significant factor in explaining contemporaneous volatility. In most cases, volume enters the model with a t -ratio greater than 15, the log-likelihood values are substantially higher than those in Table I, and the excess kurtosis is substantially lower as well. Incorporating volume also produces a jump in the R^2 values for a regression of the absolute demeaned returns on the fitted conditional volatilities. The majority of the R^2 values in Table II are greater than 20% (the largest is 39% for Procter and Gamble), whereas the majority of those in Table I are less than 10%. This increase in explanatory power suggests that the contemporaneous relation between volume and volatility is quite strong.

Despite these results, it would be premature to conclude that volume accounts for or subsumes ARCH effects in daily returns. The most obvious indication of this is that the absolute standardized return still enters the volume-augmented model with a positive and statistically significant coefficient for almost all of the stocks. In many cases, the t -ratio on the σ_h estimate is five or greater. The question is how to interpret this evidence given the implicit constraint imposed by EGARCH(1, 1) model. One possibility is that, even if the

¹¹Absolute, rather than squared, demeaned returns are used in these regressions so that the results are less sensitive to outliers. See Davidian and Carroll (1987).

TABLE II
SC-EGARCH(1, 1) Model

Firm	Estimates				t-Ratios				Diagnostics			
	s	κ_h	σ_h	γ_h	s	κ_h	σ_h	γ_h	\mathcal{L}	R^2	ρ	C_K
AXP	1.32	1.19	0.09	0.74	44.7	24.1	3.6	21.3	-5,757.2	0.19	0.41	0.21
CVX	0.66	1.16	0.04	0.58	21.9	20.0	1.6	15.7	-4,836.1	0.12	0.25	0.37
DD	1.07	1.17	0.13	0.63	32.5	14.3	4.6	17.3	-5,410.6	0.15	0.35	0.62
DIS	1.21	1.23	0.13	0.78	39.8	30.5	5.7	22.6	-5,600.4	0.25	0.31	0.22
DOW	1.52	0.01	0.08	0.01	5.0	1.4	3.3	0.7	-5,351.5	0.15	0.99	1.76
EK	0.90	1.22	0.09	0.89	27.9	31.0	3.3	27.8	-5,171.8	0.37	0.33	0.59
GE	0.91	1.20	0.07	0.72	29.9	19.1	2.8	21.2	-5,190.0	0.21	0.37	0.38
GM	1.21	1.31	0.07	0.72	39.8	30.3	2.9	19.0	-5,598.9	0.19	0.30	0.40
IBM	1.23	1.26	0.08	0.81	38.7	32.1	2.9	22.3	-5,624.0	0.28	0.35	0.57
IP	1.15	1.10	0.16	0.63	35.3	17.5	6.2	17.9	-5,522.5	0.18	0.45	0.46
JNJ	0.75	1.17	0.07	0.68	25.7	21.2	2.8	19.8	-4,965.6	0.24	0.39	0.17
KO	0.74	1.14	0.05	0.73	26.4	21.7	2.1	23.0	-4,900.3	0.25	0.41	0.04
MCD	0.90	1.24	0.14	0.69	28.4	26.2	5.7	20.7	-5,177.4	0.23	0.28	0.45
MMM	0.67	1.27	0.06	0.74	21.4	26.0	2.3	19.9	-4,855.5	0.21	0.28	0.56
MO	1.00	1.26	0.05	0.96	17.3	28.6	1.2	22.1	-5,311.6	0.37	0.34	6.96
MRK	0.95	1.29	0.05	0.73	34.8	29.5	2.2	25.0	-5,245.8	0.25	0.27	0.04
PG	0.71	1.15	0.12	0.68	22.8	23.1	4.6	20.1	-4,910.2	0.39	0.39	0.35
S	1.42	1.34	0.05	0.77	44.7	26.4	2.2	22.5	-5,894.6	0.26	0.29	0.67
T	1.39	1.17	0.15	0.68	33.0	7.4	2.8	10.0	-5,831.4	0.24	0.41	2.45
XOM	0.56	1.02	0.15	0.62	18.1	15.0	5.5	17.1	-4,705.7	0.16	0.44	0.21

Note. The table reports the results of fitting an EGARCH(1, 1) model with volume specified as a covariate to daily percentage returns on the MMI stocks. The model is of the form

$$r_t = \sqrt{h_t} z_t$$
$$\Delta \log h_t = \kappa_h (s - \log h_{t-1}) + \sigma_h u_{t-1} + \gamma_h w_t$$

where r_t is the demeaned return for day t , $u_t = (|z_t| - E[|z_t|]) / \sqrt{\text{var}(|z_t|)}$, $w_t = (\log X_t - E[\log X_t]) / \sqrt{\text{var}(\log X_t)}$, X_t denotes the daily volume, and $z_t \sim NID(0,1)$. The model is fitted via maximum likelihood. The table reports the parameter estimates, the associated t -ratios, the maximized value of log-likelihood (\mathcal{L}), the sample R^2 for a regression of $\{|r_t|\}_{t=1}^T$ on the fitted conditional volatilities (R^2), the first-order sample autocorrelation of the fitted conditional volatilities (ρ), and the coefficient of excess kurtosis for the standardized returns (C_K). The t -ratios are based on robust standard errors. The sample period is January 5, 1993–December 31, 2003. MMI, major market index.

constraint were relaxed, ARCH effects would make only a small contribution to the explanatory power of the volume-augmented model. This would support the conclusions drawn by Lamoureux and Lastrapes (1990). Alternatively, the constraint could be masking the true contribution of ARCH effects to volatility dynamics. A more detailed analysis of this issue is considered next.

Estimation and Inference for the EGARCH(2, 2) Models

To investigate the effects of relaxing the constraint imposed by the SC-EGARCH(1, 1) model, a more flexible econometric specification that nests

the SC-EGARCH(1, 1) model as a special case is estimated. In particular, an EGARCH(2, 2) model that allows for both short- and long-term volatility components is considered. The model is initially estimated without incorporating volume to assess how its empirical implications differ from those of the basic EGARCH(1, 1) model. Table III reports the results.

In general, the EGARCH(2, 2) model fits better than the EGARCH(1, 1) model reported in Table I. Most of the t -ratios for the κ_h and σ_h estimates are greater than two and the increase in the log-likelihood is statistically significant at the five percent level for a majority of the stocks. Nonetheless, allowing for two volatility components does not have a major impact on the volatility dynamics

TABLE III
EGARCH(2, 2) Model

Firm	Estimates					t -Ratios					Diagnostics			
	κ_h	σ_h	s	κ_m	σ_m	κ_h	σ_h	s	κ_m	σ_m	\mathcal{L}	R^2	ρ	C_K
AXP	0.10	0.10	1.59	0.00	0.03	3.2	4.7	8.0	1.4	2.4	-5,884.4	0.13	0.97	1.27
CVX	1.73	0.03	0.90	0.01	0.06	9.9	1.2	6.5	1.8	4.0	-4,949.8	0.06	0.96	0.87
DD	0.22	0.11	1.42	0.00	0.03	2.5	3.9	4.7	1.1	2.9	-5,497.1	0.09	0.95	1.63
DIS	1.37	0.08	1.88	0.01	0.05	7.1	2.3	7.3	1.8	2.4	-5,866.6	0.10	0.91	5.75
DOW	0.14	0.07	1.62	0.00	0.04	0.5	1.6	5.2	1.3	1.9	-5,345.3	0.15	0.98	1.63
EK	0.44	0.21	1.73	0.00	0.01	1.8	3.4	9.7	2.7	2.4	-5,810.9	0.03	0.48	16.32
GE	0.13	0.08	-0.13	0.00	0.03	1.4	3.3	-0.3	-0.8	2.9	-5,268.1	0.13	0.98	1.00
GM	0.06	0.07	1.54	0.00	0.02	1.9	3.5	10.0	0.9	1.1	-5,795.1	0.07	0.97	1.41
IBM	0.09	0.05	1.99	0.01	0.05	1.1	1.8	8.7	2.5	3.4	-5,982.4	0.06	0.97	4.08
IP	0.17	0.08	1.36	0.00	0.03	2.5	2.8	4.3	1.2	4.3	-5,618.0	0.11	0.97	1.20
JNJ	0.86	0.08	1.10	0.02	0.07	3.8	2.3	6.6	2.2	3.8	-5,178.0	0.08	0.92	1.60
KO	0.41	0.09	1.18	0.00	0.04	2.7	3.0	4.1	1.6	5.1	-5,164.3	0.10	0.96	1.70
MCD	1.07	0.11	1.46	0.01	0.05	3.4	3.1	6.7	2.2	3.5	-5,404.6	0.07	0.87	2.47
MMM	1.04	0.10	1.18	0.01	0.04	1.1	1.9	7.9	2.0	2.3	-5,098.5	0.07	0.89	3.62
MO	0.51	0.11	3.11	0.00	0.03	1.7	2.9	1.9	1.0	4.7	-5,848.7	0.06	0.91	10.68
MRK	0.63	0.06	1.41	0.01	0.03	0.2	1.5	5.3	0.6	1.1	-5,528.5	0.05	0.95	2.19
PG	1.46	0.04	1.54	0.00	0.06	3.1	1.0	2.6	1.9	5.0	-5,173.5	0.09	0.96	11.03
S	1.90	-0.02	1.89	0.04	0.11	28.8	-1.1	11.9	1.8	3.2	-6,148.3	0.07	0.95	3.82
T	0.21	0.11	2.03	0.00	0.03	3.8	4.7	7.2	1.9	3.4	-5,886.6	0.14	0.96	7.18
XOM	0.30	0.09	0.88	0.01	0.05	3.5	3.7	3.2	2.0	5.4	-4,788.4	0.11	0.96	1.01

Note. The table reports the results of fitting an EGARCH(2, 2) model to daily percentage returns on the MMI stocks. The model has a two-component representation of the form

$$\begin{aligned}
 r_t &= \sqrt{h_t} z_t \\
 \Delta \log h_t &= \Delta m_t + \kappa_h(m_{t-1} - \log h_{t-1}) + \sigma_h u_{t-1} \\
 \Delta m_t &= \kappa_m(s - m_{t-1}) + \sigma_m u_{t-1}
 \end{aligned}$$

where r_t is the demeaned return for day t , $u_t = (|z_t| - E[|z_t|]) / \sqrt{\text{var}(|z_t|)}$, and $z_t \sim NID(0,1)$. The model is fitted via maximum likelihood. The table reports the parameter estimates, the associated t -ratios, the maximized value of log-likelihood (\mathcal{L}), the sample R^2 for a regression of $\{|r_t|\}_{t=1}^T$ on the fitted conditional volatilities (R^2), the first-order sample autocorrelation of the fitted conditional volatilities (ρ), and the coefficient of excess kurtosis for the standardized returns (C_K). The t -ratios are based on robust standard errors. The sample period is January 5, 1993–December 31, 2003. MMI, major market index.

implied by the model. Although some decline is observed in the first-order sample autocorrelation of the fitted conditional volatilities, the autocorrelation still exceeds 0.90 for 17 of the 20 stocks. Similarly, the R^2 for a regression of the absolute demeaned returns on the fitted conditional volatilities suggests little increase in the explanatory power of the model. Overall these results point to a relatively modest improvement in the goodness of fit.¹²

Now the primary issue is considered, which is how the EGARCH(2, 2) model performs once volume is specified as a covariate. Table IV reports the results. Two aspects of the results stand out immediately. First, all of the log-likelihood values are significantly higher than those for the SC-EGARCH(1, 1) model reported in Table II. The average increase in log-likelihood across stocks is 103. Second, all of the R^2 values are substantially higher as well. Most of the increases are in the range of 6–10 percentage points, with increases of 13 percentage points for three of the firms. These findings point to a clear increase in explanatory power relative to the SC-EGARCH(1, 1) model.

An interesting pattern is also seen in the coefficient estimates. All of the κ_h and γ_h estimates are positive, highly statistically significant, and comparable in magnitude to the corresponding estimates in Table II. More importantly, only four of the σ_h estimates are significantly different from zero at the five percent level. This indicates that the short-term dynamics of log volatility, which are captured by $\log h_t - m_t$, are explained almost exclusively by volume. In contrast, only one of the γ_m estimates is statistically significant at the five percent level, whereas the ς , κ_m , and σ_m estimates are similar to the corresponding estimates in Table I. This indicates that the long-term dynamics of log volatility, which are captured by m_t , are explained almost exclusively by the absolute standardized returns.

These findings suggest a much different role for volume than the results obtained using the SC-EGARCH(1, 1) model. Specifically, nothing in the results for the SC-EGARCH(2, 2) model indicates that volume accounts for or subsumes ARCH effects in daily returns. On the contrary, it is found that ARCH effects are a key determinant of long-term volatility dynamics and that the long-term component of volatility displays the high level of persistence typically reported in the ARCH literature. Moreover, the evidence suggests that the long-term component of volatility in the SC-EGARCH(2, 2) model behaves similar to the conditional volatility implied by the basic EGARCH(1, 1) model.

Figure 1 illustrates this point more clearly. The figure compares the fitted values produced by the two models for American Express, the first stock alphabetically in the MMI. Panels A and B plot the fitted values of $\log h_t$ and m_t ,

¹²Christoffersen, Jacobs, and Wang (2005) used a two-component GARCH specification in an option pricing context and found that it substantially outperforms their benchmark single-component model.

TABLE IV
SC-EGARCH(2, 2) Model

Firm	Estimates						t-Ratios						Diagnostics					
	κ_h	σ_h	γ_h	s	κ_m	σ_m	γ_m	κ_h	σ_h	γ_h	s	κ_m	σ_m	γ_m	\mathcal{L}	R^2	ρ	C_K
AXP	1.21	-0.05	0.72	1.01	0.01	0.06	0.00	32.4	-1.8	20.7	6.0	1.9	4.1	0.7	-5,618.9	0.29	0.62	-0.03
CVX	1.15	-0.07	0.60	0.48	0.01	0.05	0.00	23.9	-2.4	18.2	3.2	1.5	3.5	0.1	-4,730.4	0.18	0.45	0.02
DD	1.25	0.01	0.62	0.78	0.01	0.06	0.00	24.8	0.3	20.3	3.6	1.3	3.0	-0.3	-5,246.3	0.25	0.57	0.06
DIS	1.24	0.01	0.79	0.78	0.00	0.05	0.00	36.2	0.2	23.5	4.0	2.4	6.6	-0.5	-5,437.6	0.35	0.49	-0.18
DOW	1.14	-0.01	0.61	0.92	0.00	0.06	0.00	21.5	-0.5	18.3	3.9	1.9	5.5	-1.4	-5,112.2	0.27	0.69	0.39
EK	1.18	-0.03	0.96	0.83	0.06	0.10	-0.02	28.5	-0.7	26.3	12.6	1.9	3.5	-1.6	-5,093.0	0.43	0.35	0.50
GE	1.24	-0.05	0.67	-0.32	0.00	0.03	0.00	28.4	-1.9	23.3	-1.0	1.8	5.2	1.6	-4,992.9	0.34	0.65	-0.29
GM	1.30	-0.03	0.76	1.05	0.02	0.06	-0.01	33.5	-1.0	23.9	11.2	2.9	5.3	-1.6	-5,492.0	0.26	0.41	-0.11
IBM	1.21	-0.09	0.89	0.83	0.01	0.06	-0.01	36.7	-3.2	24.6	5.2	2.3	5.0	-2.4	-5,429.8	0.41	0.44	-0.18
IP	1.20	0.02	0.64	0.56	0.00	0.04	0.00	26.8	0.5	19.4	0.9	0.6	1.7	-0.6	-5,377.7	0.27	0.65	0.05
JNJ	1.21	-0.03	0.70	0.61	0.02	0.06	0.00	29.3	-1.1	19.8	6.5	3.2	6.2	-1.4	-4,883.8	0.30	0.46	-0.10
KO	1.21	-0.04	0.69	0.37	0.01	0.04	0.00	30.7	-1.5	24.2	2.3	2.5	5.8	0.9	-4,835.8	0.34	0.56	-0.26
MCD	1.26	0.04	0.72	0.69	0.01	0.05	0.00	29.3	1.2	24.3	3.9	1.4	3.1	-0.6	-5,059.6	0.32	0.41	0.11
MMM	1.27	-0.02	0.73	0.49	0.01	0.04	0.00	27.4	-0.7	22.8	2.1	1.3	3.4	-0.9	-4,733.2	0.29	0.47	0.11
MO	1.26	-0.12	1.02	0.93	0.03	0.10	-0.01	43.6	-2.4	26.7	5.3	1.7	3.6	-0.9	-5,218.1	0.41	0.41	8.23
MRK	1.28	-0.03	0.73	0.69	0.01	0.06	0.00	32.2	-0.9	25.5	5.8	2.5	5.5	0.6	-5,154.3	0.30	0.41	-0.30
PG	1.18	0.00	0.72	0.17	0.00	0.05	0.00	28.8	-0.1	22.0	0.6	1.3	4.1	-0.5	-4,762.9	0.45	0.46	-0.07
S	1.26	-0.06	0.81	1.24	0.01	0.08	-0.01	29.5	-2.1	24.4	9.8	2.4	5.1	-1.6	-5,743.7	0.36	0.45	0.00
T	1.26	-0.02	0.83	0.82	0.01	0.07	-0.01	34.8	-0.8	24.4	2.7	1.7	5.3	-1.5	-5,420.7	0.37	0.60	0.18
XOM	1.13	-0.01	0.60	0.08	0.01	0.05	0.00	24.6	-0.4	21.9	0.3	1.5	5.5	1.0	-4,552.5	0.26	0.66	-0.25

Note. The table reports the results of fitting an EGARCH(2, 2) model with volume specified as a covariate to daily percentage returns on the MMI stocks. The model has a two-component representation of the form

$$r_t = \sqrt{h_t} z_{rt}$$

$$\Delta \log h_t = \Delta m_t + \kappa_h(m_{t-1} - \log h_{t-1}) + \sigma_h u_{t-1} + \gamma_h w_t$$

$$\Delta m_t = \kappa_m(s - m_{t-1}) + \sigma_m u_{t-1} + \gamma_m w_t$$

where r_t is the demeaned return for day t , $u_t = (|z_{rt}| - E[|z_{rt}|] / \sqrt{\text{var}(\log X_t)})$, $w_t = (\log X_t - E[\log X_t]) / \sqrt{\text{var}(\log X_t)}$, X_t denotes the daily volume, and $z_{rt} \sim N(0, 1)$. The model is fitted via maximum likelihood. The table reports the parameter estimates, the associated t-ratios, the maximized value of log-likelihood (\mathcal{L}), the sample R^2 or a regression of $\{r_t\}_{t=1}^T$ on the fitted conditional volatilities (R^2), the first-order sample autocorrelation of the fitted conditional volatilities (ρ), and the coefficient of excess kurtosis for the standardized returns (C_K). The t-ratios are based on robust standard errors. The sample period is January 5, 1993–December 31, 2003. MMI, major market index.

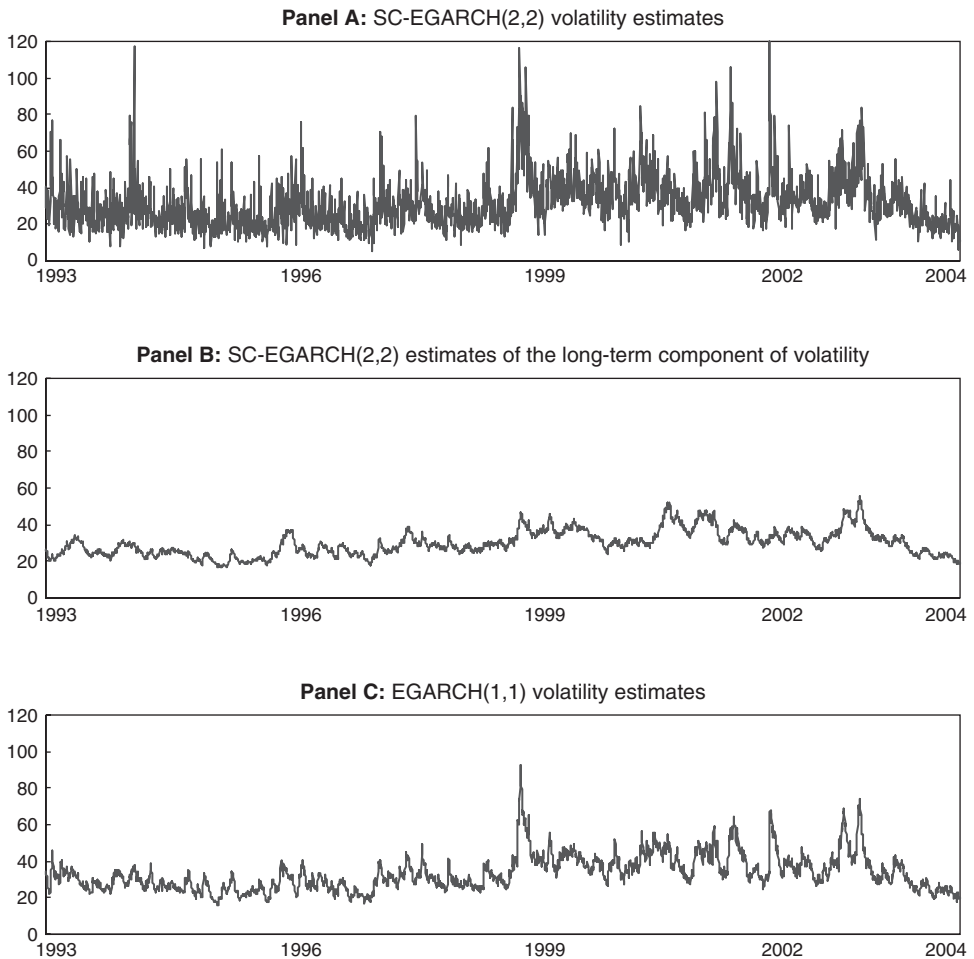


FIGURE 1

Comparison of volatility estimates for American Express. The figure plots the daily volatility estimates for American Express under the EGARCH(1, 1) model and the SC-EGARCH(2, 2) model with volume specified as a covariate. Panel A shows the fitted volatility estimates for the SC-EGARCH(2, 2) model, Panel B shows the fitted long-term component of volatility under the SC-EGARCH(2, 2) model, and Panel C shows the fitted volatility estimates for the EGARCH(1, 1) model. Each series is expressed as an annualized percentage volatility. The sample period is January 5, 1993–December 31, 2003.

respectively, for the SC-EGARCH(2, 2) model. Panel C plots the fitted values of $\log h_t$ for the basic EGARCH(1, 1) model. The fitted log volatility in Panel A is highly variable, reflecting the strong short-term impact of trading volume. However, there are also indications of an underlying autoregressive structure that seems characteristic of a slowly mean-reverting process. Once the associated component of log volatility (Panel B) is isolated, it is found that it tracks closely with the fitted values from the basic EGARCH(1, 1) model (Panel C).

Note that regardless of the model, incorporating volume does produce a large drop in the persistence of the fitted volatility series as found by Lamoureux and Lastrapes (1990) and others. Consider the ρ estimates in Table IV. Although they are higher than the corresponding estimates in Table II, they are still well below the values reported in Tables I and III. As ARCH effects appear to be undiminished for the model analyzed in Table IV, this has two implications. First, the short-term component of volatility is much less persistent than is typical of the fitted volatility from ARCH models. Second, short-term dynamics account for a substantial fraction of the total variation in the volatility of daily returns.

Table V provides additional evidence on the short- and long-term volatility dynamics. The first four columns report the sample variance of the fitted $\log h_t$ for each of the models in Tables I–IV. Not surprisingly, both of the volume-augmented models imply substantially more variation in $\log h_t$. The more

TABLE V
Volume Versus ARCH Effects

Firm	<i>Estimated var(log h_t) for the Different Models</i>				<i>Components of var(log h_t) for the SC-EGARCH(2, 2) Model</i>		
	EGARCH (1, 1)	SC-EGARCH (1, 1)	EGARCH (2, 2)	SC-EGARCH (2, 2)	Short term	Long term	Interaction
AXP	0.301	0.484	0.315	0.663	0.425	0.214	0.024
CVX	0.164	0.304	0.168	0.444	0.320	0.153	−0.028
DD	0.271	0.371	0.289	0.581	0.318	0.268	−0.004
DIS	0.281	0.527	0.287	0.736	0.511	0.257	−0.031
DOW	0.516	0.445	0.509	0.790	0.327	0.478	−0.016
EK	0.129	0.654	0.133	0.724	0.777	0.141	−0.194
GE	0.439	0.448	0.509	0.741	0.371	0.358	0.012
GM	0.148	0.415	0.151	0.534	0.450	0.148	−0.063
IBM	0.217	0.526	0.217	0.729	0.634	0.299	−0.204
IP	0.293	0.392	0.304	0.598	0.344	0.260	−0.006
JNJ	0.190	0.411	0.195	0.504	0.406	0.124	−0.026
KO	0.280	0.475	0.290	0.606	0.393	0.173	0.040
MCD	0.214	0.416	0.218	0.543	0.425	0.190	−0.072
MMM	0.169	0.448	0.180	0.597	0.424	0.255	−0.018
MO	0.274	0.731	0.291	0.890	0.800	0.160	−0.070
MRK	0.115	0.432	0.119	0.537	0.421	0.137	−0.020
PG	0.369	0.431	0.367	0.691	0.435	0.251	0.005
S	0.230	0.453	0.228	0.613	0.509	0.227	−0.123
T	0.516	0.418	0.511	1.054	0.535	0.625	−0.105
XOM	0.328	0.415	0.343	0.627	0.329	0.268	0.030

Note. The table examines the extent to which trading volume captures ARCH effects in daily returns on the MMI stocks. The sample variance of the fitted log volatility series from the models in Tables I–IV is reported. In addition, the variance for the model in Table IV is decomposed into three components—short term, long term, and interaction—using the relation $\text{var}(\log h_t) = \text{var}(\log h_t - m_t) + \text{var}(m_t) + 2\text{cov}(\log h_t - m_t, m_t)$. The sample period is January 5, 1993–December 31, 2003. MMI, major market index.

interesting comparison is between these two models. The SC-EGARCH(2, 2) model yields the higher value for every stock, and the difference is often 20% or more. This is indicative of the impact of relaxing the constraint on the decay rates imposed by the SC-EGARCH(1, 1) model. In the absence of the constraint, it becomes apparent that ARCH effects make an important contribution to the dynamics of the log variance series.

The three remaining columns of the table decompose the sample variance of the fitted $\log h_t$ for the SC-EGARCH(2, 2) model into three components—short term, long term, and interaction—using the relation $\text{var}(\log h_t) = \text{var}(\log h_t - m_t) + \text{var}(m_t) + 2\text{cov}(\log h_t - m_t, m_t)$. The results show that most of the variation in $\log h_t$ is short term in nature. But this is not due to the absence of strong ARCH effects. To see this, compare the variance of m_t in column six with the variance of the fitted $\log h_t$ from the EGARCH(1, 1) model in column one. The two sets of figures are similar, which is consistent with the evidence from Figure 1. In general, the long-term component of volatility tends to closely mimic the conditional volatility implied by the basic EGARCH(1, 1) model.

The relation between the short- and long-term volatility components is also of interest. The interaction term in the variance decomposition is negative for most of the firms. However, with the exception of AT&T, Eastman Kodak, and General Motors, the correlation between the components is such that a regression of one on the other would yield an R^2 of less than five percent. Therefore, it seems that the short- and long-term components of volatility are largely unrelated. As the former is driven primarily by contemporaneous volume and the latter by lagged absolute returns, this lack of correlation is broadly consistent with volatility following a stochastic autoregressive process in which the unpredictable volatility shocks are strongly associated with the contemporaneous level of trading activity. This is consistent with Fleming et al. (2006).

Regression-Based Model Comparisons

Table VI provides direct evidence on how well the various models capture the dynamics of volatility. The table reports the R^2 values for a regression of the log realized variances on the fitted log variances from each of the four EGARCH specifications. The R^2 values for the basic EGARCH(1, 1) model range from 16.8% for Eastman Kodak to 50.1% for AT&T. This range is roughly consistent with the evidence reported by Andersen and Bollerslev (1998) for a GARCH(1, 1) model. As the basic model captures up to 50% of the variation in the log realized variances, it provides a reasonable benchmark for assessing the performance of the other three models.

TABLE VI
Realized Variance Regressions

Firm	Regression R^2 for the Different Models			
	EGARCH(1, 1)	SC-EGARCH(1, 1)	EGARCH(2, 2)	SC-EGARCH(2, 2)
AXP	0.376	0.348	0.384	0.555
CVX	0.283	0.170	0.284	0.377
DD	0.405	0.203	0.417	0.520
DIS	0.403	0.238	0.411	0.507
DOW	0.486	0.279	0.496	0.562
EK	0.168	0.318	0.223	0.439
GE	0.482	0.282	0.493	0.636
GM	0.240	0.261	0.246	0.424
IBM	0.309	0.283	0.322	0.578
IP	0.447	0.209	0.463	0.493
JNJ	0.232	0.293	0.243	0.404
KO	0.369	0.291	0.383	0.521
MCD	0.241	0.216	0.257	0.426
MMM	0.340	0.238	0.354	0.491
MO	0.267	0.416	0.303	0.529
MRK	0.246	0.289	0.259	0.488
PG	0.358	0.281	0.362	0.521
S	0.320	0.203	0.311	0.440
T	0.501	0.221	0.522	0.613
XOM	0.433	0.226	0.445	0.520

Note. The table reports the R^2 for the regression

$$\log RV_t = a + b \log \hat{h}_t + e_t$$

where RV_t is the realized variance for day t and $\log \hat{h}_t$ is the fitted log variance for day t for each of the models in Tables I–IV. The realized variance is constructed using the full-day Newey–West estimator described in the Appendix, with a 30-second sampling frequency and a window length of 30 minutes. The sample period is January 5, 1993–December 31, 2003.

If contemporaneous volume largely subsumes ARCH effects, then one should find that the SC-EGARCH(1, 1) model performs at least as well as the basic model. It is found that this is not the case. The SC-GARCH(1, 1) model produces a lower R^2 value for 15 of the 20 stocks and, in some cases, the reduction exceeds 20 percentage points. Apparently, the addition of trading volume forces the model to place too little weight on the lagged absolute returns, leading to variance estimates that have a lower correlation with the realized variances than the estimates from the basic model. Thus, the results support the earlier conclusions about the shortcomings of the SC EGARCH(1, 1) specification.

The R^2 values for the EGARCH(2, 2) model are similar to those for the EGARCH(1, 1) model. However, adding contemporaneous volume to the EGARCH(2, 2) specification leads to a substantial increase in the R^2 value for

most of the firms, typically on the order of 10–20 percentage points. In most cases, the R^2 for the SC-EGARCH(2, 2) model exceeds 50%. This finding confirms the need to properly account for the short-term impact of the information contained in daily volume to uncover the true nature of the relation between ARCH effects and trading volume.

CONCLUSIONS

The specification of GARCH models with trading volume as a covariate is more complex than it initially appears. Even if the most commonly cited specification issue—bias arising from the endogeneity of volume—can be reasonably ignored, the models typically used in the literature impose a constraint that makes it difficult to draw reliable inferences from the model-fitting results. In particular, they restrict the half-life of a volatility shock to be the same regardless of its source. A careful investigation reveals that this restriction is strongly rejected by the data and that once the constraint on decay rates is relaxed, specifying contemporaneous volume as a covariate does little to diminish the importance of lagged squared returns in capturing the dynamics of volatility.

More generally, the analysis suggests that any GARCH(1, 1) or EGARCH(1, 1) model with stochastic covariates has the potential to produce unreliable inferences if the covariates capture a component of volatility distinct from that captured by lagged squared returns. Researchers should be cautious about using these models in the absence of suitable robustness checks. Robustness could be established, for example, by fitting a higher-order GARCH model, such as a two-component specification, and conducting model comparisons using standard diagnostic measures. It should be readily apparent from the model comparisons whether the constraint is a concern.

APPENDIX

This Appendix describes the approach for constructing realized variances, including the choice of sampling frequency and the method of dealing with trading- versus nontrading-period returns. In theory the realized variances should be constructed by sampling returns as frequently as possible. As the sampling frequency increases, however, returns become more negatively serially correlated due to market microstructure effects, which leads to biased variance estimates. Moreover, high-frequency returns are not available on weekends or overnight. These issues are dealt with separately: first the realized variance for the trading day is constructed using an estimator that is robust to serial correlation in returns, and then the full-day realized variance is constructed by

combining the trading-day realized variance with the squared nontrading-period return.

The realized variance for the trading day is constructed using the Newey–West (1987) estimator proposed by Hansen and Lunde (2004):

$$RV_{t[o]} = \sum_{i=1}^m R_{t_{i,m}}^2 + 2 \sum_{j=1}^q \left(1 - \frac{j}{q+1}\right) \sum_{i=1}^{m-j} R_{t_{i,m}} R_{t_{j,m}} \quad (A1)$$

where q denotes the window length for the autocovariance terms. As this estimator is consistent in the presence of serial correlation, it allows one to sample returns at a higher frequency and thereby incorporate information that might otherwise be lost. The full-day realized variance is obtained by combining $RV_{t[o]}$ with the squared nontrading-period return, $R_{t[c]}^2$, using the weighting scheme proposed by Hansen and Lunde (2005b). They considered the class of conditionally unbiased estimators that is linear in $RV_{t[o]}$ and $R_{t[c]}^2$ and showed that the following weights deliver the lowest mean-squared error:

$$RV_t = \varphi \frac{\psi}{\psi_o} RV_{t[o]} + (1 - \varphi) \frac{\psi}{\psi_c} R_{t[c]}^2 \quad (A2)$$

where

$$\varphi = \frac{\psi_o^2 \eta_c^2 - \psi_o \psi_c \eta_{oc}}{\psi_c^2 \eta_o^2 + \psi_o^2 \eta_c^2 - 2\psi_o \psi_c \eta_{oc}} \quad (A3)$$

and $\psi = E(R_t^2)$, $\psi_o = E(RV_{t[o]})$, $\psi_c = E(R_{t[c]}^2)$, $\eta_o^2 = \text{var}(RV_{t[o]})$, $\eta_c^2 = \text{var}(R_{t[c]}^2)$, and $\eta_{oc} = \text{cov}(RV_{t[o]}, R_{t[c]}^2)$. Note that the ratios ψ/ψ_o and ψ/ψ_c ensure that the full-day realized variance has the same unconditional mean as the squared close-to-close return, whereas φ determines the weights placed on the trading- and nontrading-period variance estimators. In general, φ should be close to one because variance is typically lower during the nontrading period than the trading period, and $R_{t[c]}^2$ is an imprecise estimator of the nontrading-period variance. This can most easily be seen by assuming $\eta_{oc} = 0$.

To implement Equations (A1) and (A2), intraday transaction prices from the TAQ database are used. The price filters described in The Data Set section are applied to eliminate obvious reporting errors and then the remaining prices are used to construct returns. The trading day for stocks is usually 390 minutes in length (9:30A.M. to 4:00P.M. Eastern Standard Time). Sampling frequencies as high as $m = 780$ (i.e., 30-second returns) are considered. For a given choice of m , one needs to find the price at the beginning and end of each $m/390$ -minute interval. The intervals start with the first price in the TAQ database for that day, which is treated as the beginning price for

the interval in which it occurs. The price at the end of this and each successive interval is then estimated by linear interpolation of the prices nearest (on either side) to the end of the interval (see Andersen & Bollerslev, 1997). If one or more prices occur exactly at the end of the interval, the average of these prices is used. The last transaction price of the day is used as the price at the end of the last interval. The returns are constructed by differencing these log prices. As expected, the returns have a negative first-order serial correlation that increases with the sampling frequency. The average correlation coefficient across the 20 MMI stocks is -0.07 for five-minute returns and -0.15 for 30-second returns.¹³

The intraday returns are used to construct $RV_{t[o]}$ using values of q that correspond to four different window lengths: 0, 15, 30, and 60 minutes. Using a window length of 0, the bias caused by microstructure effects is readily apparent. Realized variances constructed using five-minute returns, which is common practice in the literature, are on average 13% greater than the average squared open-to-close return. The bias is much worse at higher sampling frequencies. However, increasing the window length counteracts the bias. Using a 15-minute window, the realized variances are still noticeably biased; but, using a 30-minute window, the average realized variances at every sampling frequency are within two percent of the average squared open-to-close return. Increasing the window length further (e.g., 60 minutes) substantially increases the standard deviation of the realized variances, as including unnecessary covariance terms in Equation (A1) reduces efficiency. Based on these results, the realized variances constructed using 30-second returns and a 30-minute window length are used in the construction of the full-day realized variances.

To obtain the full-day realized variances, the sample analogs of ψ , ψ_o , ψ_c , η_o^2 , η_c^2 , and η_{oc} are substituted into Equations (A2) and (A3). Hansen and Lunde (2005b) suggested removing outliers from the estimation to avoid obtaining a negative weight on $R_{t[c]}^2$. Accordingly, days in which either $RV_{t[o]}$ or $R_{t[c]}^2$ is among the largest 0.5% of the observations for each stock are excluded. The average φ estimate for the 20 stocks is 0.92. By comparison, the ratio of the average squared close-to-close return to the average squared close-to-open return indicates that approximately 20% of the daily variance occurs during the nontrading period. However, the φ estimate gives less weight than this to the nontrading-period variance estimate because the trading-period variance estimate is much more precise.

¹³These serial correlation coefficients (based on interpolated prices) are substantially smaller than those obtained using the last transaction price in each intraday time interval. This is true even if an MA(1) model is used to filter returns as in Andersen, Bollerslev, Diebold, and Ebens (2001).

BIBLIOGRAPHY

- Andersen, T. G., & Bollerslev, T. (1997). Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance*, 4, 115–158.
- Andersen, T. G., & Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39, 885–905.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., & Ebens, H. (2001). The distribution of stock return volatility. *Journal of Financial Economics*, 61, 43–76.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2001). The distribution of realized exchange rate volatility. *Journal of the American Statistical Association*, 96, 42–55.
- Andersen, T. G., Bollerslev, T., & Meddahi, N. (2005). Correcting the errors: Volatility forecast evaluation using high-frequency data and realized volatilities. *Econometrica*, 73, 279–296.
- Barndorff-Nielsen, O. E., & Shephard, N. (2002). Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society, Series B*, 64, 253–280.
- Blair, B. J., Poon, S., & Taylor, S. J. (2001). Forecasting S&P 100 volatility: The incremental information content of implied volatilities and high frequency returns. *Journal of Econometrics*, 105, 5–26.
- Bollerslev, T., & Wooldridge, J. M. (1992). Quasi-maximum likelihood estimation and inference in dynamic models with time varying covariances. *Econometric Reviews*, 11, 143–172.
- Christoffersen, P., Jacobs, K., & Wang, Y. (2005). Option valuation with long-run and short-run volatility components (working paper). McGill University.
- Davidian, M., & Carroll, R. J. (1987). Variance function estimation. *Journal of the American Statistical Association*, 82, 1079–1091.
- Day, T. E., & Lewis, C. M. (1992). Stock market volatility and the information content of stock index options. *Journal of Econometrics*, 52, 267–287.
- Dominguez, K. (1998). Central bank intervention and exchange rate volatility. *Journal of International Money and Finance*, 17, 161–190.
- Engle, R. F., & Lee, G. J. (1999). A permanent and transitory component model of stock return volatility. In R. F. Engle, & H. White (Eds.), *Cointegration, causality, and forecasting: A festschrift in honour of Clive W.J. Granger*. Oxford: Oxford University Press.
- Engle, R. F., & Patton, A. (2001). What good is a volatility model? *Quantitative Finance*, 1, 237–245.
- Fleming, J., Kirby, C., & Ostdiek, B. (2006). Stochastic volatility, trading volume, and the daily flow of information. *Journal of Business*, 79, 1551–1590.
- Forsberg, L., & Ghysels, E. (2007). Why do absolute returns predict volatility so well? *Journal of Financial Econometrics*, 5, 31–67.
- Fujihara, R., & Mougoue, M. (1997). Linear dependence, nonlinear dependence and petroleum futures market efficiency. *Journal of Futures Markets*, 17, 75–99.
- Gallo, G., & Pacini, B. (2000). The effects of trading activity on market volatility. *European Journal of Finance*, 6, 163–175.
- Ghysels, E., Santa-Clara, P., & Valkanov, R. (2006). Predicting volatility: Getting the most out of return data sampled at different frequencies. *Journal of Econometrics*, 131, 59–95.

- Gillemot, L., Farmer, J. D., & Lillo, F. (2005). There's more to volatility than volume (working paper). Santa Fe Institute.
- Girma, P., & Mougoue, M. (2002). An empirical examination of the relation between futures spreads volatility, volume, and open interest. *Journal of Futures Markets*, 22, 1083–1102.
- Glosten, L., Jagannathan, R., & Runkle, D. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance*, 48, 1779–1801.
- Hagiwara, M., & Herce, M. (1999). Endogenous exchange rate volatility, trading volume and interest rate differentials in a model of portfolio selection. *Review of International Economics*, 7, 202–218.
- Hansen, P. R., & Lunde, A. (2004). An unbiased measure of realized variance (working paper). Stanford University.
- Hansen, P. R., & Lunde, A. (2005a). A forecast comparison of volatility models: Does anything beat a GARCH(1, 1)? *Journal of Applied Econometrics*, 20, 873–889.
- Hansen, P. R., & Lunde, A. (2005b). A realized variance for the whole day based on intermittent high-frequency data. *Journal of Financial Econometrics*, 3, 525–554.
- Hodrick, R. (1989). Risk, uncertainty, and exchange rates. *Journal of Monetary Economics*, 23, 433–459.
- Kim, D., & Kon, S. J. (1994). Alternative models for the conditional heteroscedasticity of stock returns. *Journal of Business*, 67, 563–598.
- Lamoureux, C. G., & Lastrapes, W. D. (1990). Heteroskedasticity in stock return data: Volume versus GARCH effects. *Journal of Finance*, 45, 221–229.
- Lamoureux, C. G., & Lastrapes, W. D. (1993). Forecasting stock-return variance: Toward an understanding of stochastic implied volatility. *Review of Financial Studies*, 6, 293–326.
- Liesenfeld, R. (1998). Dynamic bivariate mixture models: Modeling the behavior of prices and trading volume. *Journal of Business and Economic Statistics*, 16, 101–109.
- Liesenfeld, R. (2001). A generalized bivariate mixture model for stock price volatility and trading volume. *Journal of Econometrics*, 104, 141–178.
- Marsh, T. A., & Wagner, N. (2005). Surprise volume and heteroskedasticity in equity market returns. *Quantitative Finance*, 5, 153–168.
- Merton, R. C. (1980). On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics*, 8, 323–361.
- Nelson, D. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 59, 347–370.
- Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55, 703–708.