Volatility Forecasts, Trading Volume and the ARCH vs Option-Implied Volatility Tradeoff

R. Glen Donaldson
University of British Columbia

Mark J. Kamstra
Simon Fraser University

Version: February 12, 2001

Please send correspondence to: R. Glen Donaldson, Faculty of Commerce and Business Administration, University of British Columbia, Vancouver, BC, Canada, V6T 1Z2.
Voice: 604-822-8344. Fax: 604-822-8521. Email: glen.donaldson@commerce.ubc.ca.

We thank Ron Giammarino, Lisa Kramer, Alan Kraus and participants in various seminars and conferences for helpful discussions and the Social Sciences and Humanities and Research Council of Canada for financial support. The usual disclaimer applies.

Volatility Forecasts, Trading Volume and the ARCH vs Option-Implied Volatility Tradeoff

Abstract

Market expectations of future return volatility play a crucial role in finance; so too does our understanding of the process by which information is incorporated in security prices through the trading process. This paper seeks to learn something about both of these issues by investigating empirically the role of trading volume (a) in predicting the relative informativeness of volatility forecasts produced by ARCH models versus the volatility forecasts derived from option prices, and (b) in improving volatility forecasts produced by ARCH and option models and combinations of models. We find that if trading volume was low during period $t-1$ then ARCH is much more important than options for forecasting future stock market volatility. Conversely, if volume was high during period $t-1$, then option-implied volatility is much more important than ARCH for forecasting future volatility. Our findings reveal an important regime-switching role for trading volume and suggest that option markets may be more efficient in high volume states. Results from various tests also uncover possible sources of volume-related nonlinearity in the relationship between past and future return innovations as captured by asymmetric ARCH models.

1 Introduction and Overview

Market expectations of future return volatility play a crucial role in finance; so too does our understanding of the process by which information is incorporated in security prices through the trading process. This paper seeks to learn something about both of these issues by investigating empirically the role of trading volume (a) in predicting the relative informativeness of volatility forecasts produced by ARCH models versus the volatility forecasts derived from option prices, and (b) in improving volatility forecasts produced by ARCH and option models and combinations of models.

Previous studies have reported that trading volume cannot forecast volatility directly. In this paper we uncover a new result: that volume does indeed have predictive power for forecasting volatility, with volume playing the role of a regime-switching variable between states in which option-implied volatility is more or less informative than ARCH for volatility forecasting. Indeed, we find that the accuracy of volatility forecasts can be significantly improved by accounting for the volume regime effect and by combining information from ARCH models and option prices accordingly. This finding is made possible because of the novel way we incorporate trading volume into our functional forms and because, while previous papers have added either trading volume or option-implied volatility (but not both) to ARCH models, our study is the first we know of to consider all three factors together.

Traditional ARCH models – including GARCH, EGARCH, and so forth – forecast future return volatility given only information on lagged return innovations.[1] Previous papers that have studied the relationship between return-based volatility forecasts and trading volume have found that, while volume and returns are correlated contemporaneously, *lagged* volume has no power to forecast *future* volatility once the effects of lagged return innovations have been fully accounted for.[2] In other words, results from previous research suggests that in an

[1] Bollerslev, Cho and Kroner (1992) review the traditional ARCH literature. Examples and references to more recent work in areas such as seminonparametric ARCH modeling can be found in Donaldson and Kamstra (1997) and Engle and Ng (1993).

[2] See Brooks (1998), Heimstra and Jones (1994), Lamoureux and Lastrapes (1994) and Richardson and Smith (1994) for an analysis of *lagged* volume effects and returns. The relationship between *contemporaneous* volume and returns is outside the scope of our paper and thus the contemporaneous volume literature is not cited here. However, references can be found the in the aforementioned articles and in Karpoff (1987).

ARCH model that already accounts for the impact of lagged return innovations on future volatility, lagged volume will have no marginal power to forecast future volatility. We find a different result on our paper. In particular, when we interact lagged volume with option-implied volatility in an augmented ARCH model, we uncover a crucial and significant role for lagged trading volume in forecasting future return volatility. We find this important volume effect, where others have not, because we include implied volatility in our ARCH investigation whereas previous volume/return studies have not.

Previous papers that have added option-implied volatility to ARCH have done so primarily to investigate the efficiency of option markets, not to improve ARCH forecasts *per se*. If option markets are efficient then option prices will contain all available information concerning the expected future volatility of underlying prices – including any information used by ARCH models – and thus volatility forecasts implied by option prices should encompass volatility forecasts from ARCH models. However, most studies to date have found that option-implied volatility cannot encompass ARCH in one-day-ahead volatility forecasting and have thus concluded that either the option market is not efficient or that the option-pricing models employed are misspecified, or are at least problematic for short-term forecasting.[3] We find a different result. In particular, we show that option-implied volatility can indeed encompass ARCH in some cases even at short horizons, but only if the effects of lagged trading volume are also accounted for.

The data and basic model specifications we employ in our investigation are presented in Section 2 below. In Section 3 we combine ARCH volatility forecasts with both option-implied volatility forecasts and lagged trading volume and observe the importance of each input in various environments. Here we obtain an interesting result. If trading volume was "lower than normal" during period $t-1$, then the best forecast of time t volatility is found by combining the ARCH forecast from period $t-1$ with the option-implied volatility forecast

[3]See, for example, Canina and Figlewski (1993), Day and Lewis (1992), and Lamoureux and Lastrapes (1993). Note that Christensen and Prabhala (1998) and Fleming (1998) have found that option-implied volatility can outperform ARCH at longer horizons (e.g., one-month-ahead forecasts in Christensen and Prabhala (1998)) once certain biases are accounted for (as in Fleming (1998)). We discuss this more fully below.

from period $t-1$, with approximately twice as much weight being given to ARCH as to options. Conversely, if trading volume was "higher than normal" during period $t-1$, then the best forecast of time t volatility is obtained from lagged option prices alone, with almost no weight given to ARCH. Results from a simple combining exercise therefore suggest that option-implied volatility forecasts dominate ARCH forecasts when volume is high, but in low volume states ARCH provides a more reliable forecast than option-implied volatility.

Our results have several interesting implications. First, unlike most previous research on option-volatility versus ARCH, we find that volatility forecasts derived from option prices can indeed encompass ARCH volatility forecasts even at short horizons, although this result only obtains following high volume days. This suggests that the option market may be more efficient on high volume days (and/or that option pricing models are less misspecified on high volume days). It also suggests that option prices may be more informative on high volume days since in high volume periods option-implied volatility provides a better forecast of the future variance of underlying returns. This empirical finding is consistent with the spirit of theoretical work by authors such as Admati and Pfleiderer (1988) and He and Wang (1995), which suggests that market prices may be more "informative" during high volume periods. Indeed, our forecast combining approach suggests a new avenue for examining the issue of price informativeness in financial markets.

Second, as noted above, previous research suggests that lagged trading volume has no power to forecast future volatility once the effects of lagged return innovations are accounted for. Conversely, we are able to uncover a key role for lagged volume as a switching variable between informativeness regimes, whereby ARCH is a more accurate volatility forecaster on low volume days and option-implied volatility is a more accurate forecaster on high volume days.

Third, we find that the ability of ARCH models to remove autoregressive conditional heteroskedasticity from the data can be improved significantly by adding the lagged volume switching variable and option-implied volatility to standard ARCH models. This finding is investigated further in Section 4 of our paper, in which we build and test various volume-option-augmented ARCH models. In Section 4 we also search for nonlinear effects and other time-series features that may help to explain our findings. Among other interesting results, we find that when volume and options are added to ARCH models various nonlinear representations of lagged return innovations are no longer needed to fit the data within an ARCH modeling context. This suggests that ARCH modelers may profit from expanding the traditional ARCH information set to include volume, options, and other types of information in addition to the history of lagged return innovations. Section 5 concludes.

2 Basic Models and Data

Define stock returns, R_t, as the first difference of the log daily closing value of the S&P 500 Stock Price Index. $E(R_t|I_{t-1})$ is then the conditional forecast of this return such that $R_t = E(R_t|I_{t-1}) + \epsilon_t$, where I_{t-1} is the date $t-1$ conditioning information set on which date t forecasts are based and the additive forecast error, ϵ_t, has zero mean and conditional variance $E(\epsilon_t^2|I_{t-1}) = \sigma_t^2$. A variety of specification tests on the daily S&P 500 data post-1987 (which is the time period we study due to the fact that index option data prior to 1988 is too thin to be used for extracting reliable implied volatilities) reveals that expected returns are appropriately modeled with a constant, such that $R_t = \mu + \epsilon_t$. As a baseline in our own investigations below we therefore employ the constant expected returns specification.[4]

A well-documented feature of stock market data is that the return innovations, ϵ_t, appear to be drawn from a time-dependent heteroskedastic distribution. An important goal of the conditional volatility literature is to capture this feature of the data with the appropriate model for the conditional variance process so as to produce a forecasted variance, $\hat{\sigma}_t^2$, along with a return forecast error, $\hat{\epsilon}_t$, such that the standardized residuals, $\hat{\epsilon}_t/\hat{\sigma}_t$, are homoskedastic and independent.

[4] In the interest of robustness we investigated a wide variety of alternative specifications for the expected return, including specifications in which expected returns are modeled as a simple AR(1) process, or as a more complicated seasonal process with dummy variables for months of the year, days of the week, and each gap of i days between trading days, plus MA and AR terms as necessary to completely whiten the data (Data employed are from 1988-1995). Our volatility results did not change appreciably. We therefore report the simple constant specification in the analysis below.

2.1 ARCH

In the ARCH family of models the conditioning information set traditionally used to make volatility forecasts contains only the history of ϵ. Lagged ϵ^2 are included to capture the feature of volatility clustering; i.e., future volatility is related to lagged squared return innovations. Levels of lagged ϵ are also sometimes employed to capture the perception that volatility may be related in an asymmetric way to lagged return innovations, with sharp drops in stock prices causing more future volatility than upturns cause.[5] One specification that has gained particular popularity is the asymmetric Sign-GARCH model of Glosten, Jagannathan and Runkle (1993), shown in equations (1)-(3) below. This is the basic model we employ in our study and for simplicity is referred to as simply "ARCH" throughout the remainder of the paper.[6]

$$R_t = \mu + \epsilon_t \quad ; \quad \epsilon_t \sim (0, h_t^2) \tag{1}$$

$$h_t^2 = \alpha + \beta h_{t-1}^2 + \gamma \epsilon_{t-1}^2 + \delta D_{t-1} \epsilon_{t-1}^2 \tag{2}$$

$$D_{t-1} = \begin{cases} 1 & if \ \epsilon_{t-1} < 0 \\ 0 & otherwise \end{cases} \tag{3}$$

Parameter estimates for this model, estimated under Maximum Likelihood, are reported in Section 4 below, along with comparisons to other models.[7] Results from various in-sample and out-of-sample diagnostic tests and performance evaluations are also presented below.

2.2 Option-Implied Volatility

We obtain our option-implied volatility forecast for day t from the date $t-1$ closing price of call options on the S&P 500 stock price index. This variable, noted S_{t-1}^2, therefore provides us

[5]See Engle and Ng (1993) for further discussion of asymmetric ARCH effects.

[6]For robustness we investigated a variety of other ARCH-type specifications, as well as a variety of different lag lengths (e.g., higher order models). Our results were qualitatively unchanged.

[7]The time period studied is dictated by the availability of option data, as discussed below.

with an indication of the option market's forecast of future return volatility given information available at time $t-1$. There is of course some debate over which of the many possible option pricing models is most appropriate for extracting an estimate of implied volatility from the option price data. To help answer to this question, Dumas, Fleming and Whaley (1998) compare a variety of simple and complex option pricing models on the S&P 500 option data and find that simple Black-Scholes provides the most accurate volatility forecasts. From this finding Dumas, Fleming and Whaley (1998) conclude that, "...the simple Black-Scholes constant volatility model performs best of all the volatility function specifications! In sum, simpler is better [for volatility forecasting]." In our study we therefore use Black-Scholes option-implied volatility.[8]

---

[8] Following tradition in this literature (e.g., Canina and Figlewski (1993), Day and Lewis (1992), Dumas, Fleming and Whaley (1998), etc.), we filtered the option price data to exclude all contracts: traded before 1988, with trading volume less than 50, trading greater than \$15 in- or out-of-the-money, priced less than \$0.25, and with less than four days to maturity, in an effort to eliminate potential biases in the implied volatility estimates associated with thin-trading problems. From the option contracts that remained we then constructed the set of contracts with the shortest time to maturity at each date and, from this set, selected the closest to-the-money contract to provide the option-implied volatility estimate. Note that this procedure requires using options with more than one day to maturity to calculate the Black-Scholes forecast of volatility one day ahead. Fortunately, this slight mismatch between the forecast horizon and time to maturity of the option is not a serious problem in our particular application given that we are using Black-Scholes. In particular, the Black-Scholes model does *not* assume that volatility is different at each moment in time and that the implied volatility one estimates from an option's price is the average of these time-varying volatilities over a contract's life. Instead, Black-Scholes assumes that volatility is identical at every moment in time during an option's life. Thus, given the theory behind Black-Scholes, multi-day options can be safely employed to extract a *theoretically consistent* one-day-ahead Black-Scholes volatility forecast, as we have done in this paper. Note that the same cannot be said for implied volatilities extracted from more complex models that explicitly permit time varying volatility, which another good reason to use simple Black-Scholes implied volatilities in our particular application. (Even if Black Scholes is not perfect, Hull and White (1987) found that the potential biases from the Black-Scholes model relative to a stochastic volatility model become insignificant when options are close to the money and time to maturity is short, as is the case in our data set.

Note that, although we have done our best to obtain accurate implied volatilities, no one can ever claim perfection in such an exercise. Fortunately, perfection is not essential in our particular study. In most ARCH-option volatility papers, where the focus is on testing market efficiency, perfection of the option-pricing model is crucial to success of the study since part of the null hypothesis is that any apparent deviations from efficiency are *not* caused by model imperfections. In our paper the bar is much lower in that our objective is simply to see if some – even potentially flawed – volatility indicator derived from option prices might contain information that might help ARCH models make better forecasts once volume effects are also accounted for. Thus, although we have attempted to eliminate biases and misspecifications in our implied volatilities, we note that the validity of our investigation is not dependent on our implied volatility estimates being perfect.

2.3 Volume

The final variable we consider is trading volume at the NYSE. Since our objective is to forecast volatility, we are interested in *lagged* volume; i.e., $Volume_{t-1}$, or some function of $Volume_{t-1}$, as an element of the I_{t-1} forecasting information set. Since one purpose of our investigation is to determine whether ARCH and options behave differently on high versus low volume days, we first consider the high/low volume indicator variable V_{t-1}, where:

$$V_{t-1} = \begin{cases} 1 & if\ Volume_{t-1} \geq \frac{1}{(n-1)} \sum_{i=2}^{n} Volume_{t-i} \\ 0 & otherwise \end{cases} \tag{4}$$

We set $n = 5$ so that the variable V_{t-1} is "one" if lagged volume is above its one-week lagged moving average, and is "zero" otherwise.[9] As shall be discussed below, this simple discrete indicator variable works as well as, and in many cases better than, continuous alternatives.

3 Combining Forecasts

3.1 The Combining Model and Results

Perhaps the most obvious way to isolate and compare ARCH, option and volume effects is to estimate a simple linear combination of the ARCH forecasts and option-implied forecasts, along with our high/low volume switching variable. We therefore first use equations (1)-(3) and returns information from day $t-1$ to calculate an ARCH volatility forecast for day t, and denote this conditional volatility forecast as $\hat{h}_t^2$, where the "hat" denotes "forecast conditional on time $t-1$ information". We next use option prices from day $t-1$ to obtain an implied volatility estimate for day $t-1$, which we take as the option-implied volatility forecast for day t, noted $\hat{S}_t^2$ (i.e., $\hat{S}_t^2 = S_{t-1}^2$).[10] We then calculate our volume variable, V_{t-1} from

[9] We also investigated a variety of other lag lengths, including a one-month lagged moving average, instead of one week. Results from the estimations and tests discussed below were not affected qualitatively.

[10] Black-Scholes assumes *constant* instantaneous volatility over the life of the option and thus Black-Scholes volatility is at any given point in time forecasted to be the same for every one of the option's N remaining periods. Of course, the forecast of this constant Black-Scholes volatility is updated/changed each day as the forecaster's information set is updated, but at any fixed point in time Black-Scholes volatility assumes that volatility over the option's remaining life is constant, which supports our use of $\hat{S}_t^2 = S_{t-1}^2$ in our paper.

equation (4). Finally, we combine these three variables in a joint mean-variance Maximum Likelihood regression to obtain our combined volatility forecast, $\hat{\sigma}_t^2$. Parameter estimates from the combining regression are reported in Table 1 below (with t-tests, asymptotic standard errors in parentheses and Bollerslev-Wooldridge robust standard errors in brackets). Note that all parameters, except the intercepts and intercept dummy, are significant at conventional levels.[11,12] We use the obvious one-sided alternative for the t-tests as the signs of the coefficients are indicated a priori, and we report t-tests based on the asymptotic standard errors.[13]

Table 1
Parameter Estimates from the Combining Regression
{t-tests based on asymptotic standard errors in braces}
(Asymptotic standard errors in parentheses)
[Bollerslev-Wooldridge robust standard errors in brackets]

$\hat{\sigma}_t^2 =$	0.5E-5	$-$	0.5E-5 V_{t-1}	$+$	0.57 $\hat{h}_t^2$	$-$	$0.59V_{t-1}\hat{h}_t^2$	$+$	$0.33\hat{S}_t^2$	$+$	$0.73V_{t-1}\hat{S}_t^2$
	{1.00}		{0.71}		{4.12***}		{2.57***}		{2.06**}		{3.17***}
	(0.5E-5)		(0.7E-5)		(0.17)		(0.23)		(0.16)		(0.23)
	[0.9E-5]		[0.1E-4]		[0.27]		[0.42]		[0.24]		[0.47]

$\hat{\sigma}_t^2$ = stock market return volatility on day t
V_{t-1} = indicator variable - trading volume higher than average on day $t-1$
$\hat{h}_t^2$ = forecasted volatility from an ARCH model, conditional on $t-1$ information
$\hat{S}_t^2$ = forecasted volatility from option prices, conditional on $t-1$ information

* Significant at the 10% level, one-sided test.
** Significant at the 5% level, one-sided test.
*** Significant at the 1% level, one-sided test.

---

[11] Log Likelihood = 6860. The jointly estimated mean parameter from equation (1) is $\mu = 0.000382$, with standard error of 0.000155. Other diagnostics and results are reported and discussed in the following section.

[12] Forecast combining has a long tradition in the forecasting literature. For a general literature review see Clemens (1989). For a more recent example see Donaldson and Kamstra (1996).

[13] We use asymptotic rather than Bollerslev-Wooldridge robust standard errors for our t-tests because t-tests based on the latter are well-known to be overly conservative in many cases. Use of the Bollerslev-Wooldridge robust standard errors does not change the qualitative results or the conclusions drawn below.

From the theoretical studies on trading volume cited in the introduction to this paper and discussed more fully below (e.g., Admati and Pfleiderer (1988), He and Wang (1995), etc.), we would expect market prices may be more "informative" during high volume periods and thus would expect option-implied volatility, which is based on market prices, to forecast volatility more accurately, while ARCH models based on a long history of lagged prices would not necessarily prove helpful in forecasting. Thus, it is more appropriate to have our tests in this paper be one-tailed tests (with the null hypothesis of no difference between high- and low-volume states, and the alternative hypothesis that option-implied volatility forecasts improve more than ARCH forecasts improve in high-volume states) than it is to employ two-tailed tests. Comments concerning "significance" in the discussions that follow are therefore based on one-tailed tests. Readers who prefer a two-tailed alternative may compute significance levels based on the reported standard errors.

From the parameter estimates in Table 1 we see that, in low volume states when $V_{t-1} = 0$, ARCH is almost twice as important as option-implied for forecasting future volatility; i.e., the parameter on $\hat{h}_t^2$ (0.57) is almost twice as large as the parameter on $\hat{S}_t^2$ (0.33). Conversely, in high volume states when $V_{t-1} = 1$, the weight on ARCH falls to approximately zero – i.e., the parameter on $\hat{h}_t^2$ plus the parameter on $V_{t-1}\hat{h}_t^2$ is roughly zero $(0.57 - 0.59 \approx 0)$ – while the weight on option-implied rises to approximately one; i.e., the parameters on $\hat{S}_t^2$ and $V_{t-1}\hat{S}_t^2$ sum to roughly unity $(0.33 + 0.73 \approx 1)$. In other words, *in high volume states option-implied volatility dominates ARCH, while in low volume states ARCH appears more important than option-implied volatility.*[14]

We conducted several tests to check the robustness of this result. First, we investigated different definitions of volume (e.g., equation (4) based on a one-month lag instead of a one-week lag), different definitions of option-implied volatility (e.g., weighted average volatility across moneyness instead of volatility implied from an at-the-money option), and different definitions of ARCH (e.g., different models, different lag lengths, etc.). The re-

[14]Note that there is no restriction that the coefficients on ARCH and options sum to unity in our combining regression and thus the regression intercepts can be nonzero, which is the case above. See Clemens (1989) for a discussion of the pros and cons of restricting the intercept in a combining regression.

sulting regression coefficients changed somewhat from specification to specification, but the basic finding remained: option-implied volatility dominates ARCH in high volume states and ARCH dominates option-implied volatility in low volume states.

Second, we separated our data into various subsamples and re-ran the regression. Again, coefficients changed somewhat from subperiod to subperiod, but the key result was qualitatively robust.

Third, instead of using the variance of the residual from equation (1) as the dependent variable in the combining regression, we tried defining the dependent variable as option-implied volatility from day t. (Thus, we used ARCH, volume and option-implied volatility from day $t-1$ to forecast option-implied volatility from day t.) Here again we found our familiar result: option-implied volatility is a better forecaster of future volatility relative to ARCH when volume is high, no matter how the term "volatility" is defined. Regardless of the specification for volatility, this core result was robust.

Finally, we considered conducting various option-based tests on our volatility forecasts, including tests for pricing/hedging effectiveness, but did not do so because pricing/hedging effectiveness testing cannot be undertaken legitimately given the menu of option pricing models currently in existence. To test the pricing/hedging effectiveness of combining ARCH and volume effects in an option pricing framework, one would first need to build an option-pricing model that explicitly allows for both volume and ARCH effects, and then test the pricing/hedging effectiveness of that volume-ARCH-option pricing model relative to option pricing models that do not contain volume and/or ARCH effects (for a similar approach to testing pricing/hedging effectiveness of various competing models see Bakshi, Cao and Chen (1997)). In particular, it would *not* be legitimate to simply plug volatility forecasts from various ARCH-volume models into the Black-Scholes Model (or indeed into any other currently available option pricing model) and observe its resulting pricing/hedging performance. This is because no option-pricing model we know of permits volume effects in the model and thus the option-pricing model, volatility forecasting procedure and test procedure would be inconsistent with each other. While the mechanical exercise is of course feasible, the resulting performance, good or bad, would vary unpredictably with new data since there is no way to reliably establish the properties of such an internally inconsistent methodology.

3.2 Discussion

As stated in the introduction to our paper, we are not attempting to test market efficiency in our study. It is, however, interesting to note that on high volume days option-implied volatility dominates ARCH and thus that, using the yardstick suggested by previous authors, the market may indeed be efficient when enough information is flowing into the market (assuming volume is a good proxy for information flow, as it is commonly assumed – see for instance Admati and Pfleiderer (1988)). The failure of option-implied volatility to dominate ARCH on low volume days might suggest that, if the market is indeed somewhat inefficient, it may only be so when there is comparatively little information flowing. Alternatively, our results could be interpreted to reveal that the Black-Scholes model is misspecified in some way that is most clearly seen on low volume days and that the market is always efficient.

There appears to be at least two possible (not necessarily mutually exclusive) explanations for our finding that option-implied volatility provides a better volatility forecast relative to ARCH following high volume days: (a) the informativeness of the ARCH volatility forecast declines in high volume states, and/or (b) the informativeness of option-implied volatility increases in high volume states.

We begin by investigating whether there is any change in average volatility following high versus low volume days since this could potentially lead ARCH to under-forecast future volatility following heavy volume days and thus help to explain why ARCH does worse relative to options following high volume days. To examine this possibility we compared average squared errors from equation (1) and average volatility forecasts from our ARCH model on day t when volume was high versus low on day $t-1$. We found that when volume was high on day $t-1$ the average day t squared error is 0.0000630 and is 0.0000635 when volume was low on day $t-1$, with no statistically significant difference between these two numbers, and that when volume was high on day $t-1$ the average day t ARCH forecast is 0.0000651 and is 0.0000681 when volume was low on day $t-1$. In other words, squared pricing errors are no larger following high volume days than following low volume days and the ARCH model does not tend to under-forecast volatility on days following either low or high volume. We see essentially the same result from the insignificant intercepts in the combining regression in Table 1 (and in the regressions in Table 2 below). This suggests that the ARCH versus options effect is not coming from average volatility levels and thus that any explanation of our results is more likely to rest on intertemporal correlations between forecasted and realized volatility.

To investigate correlation effects we computed the simple correlation between realized volatility at time t (as measured by time t's squared return innovation) and the ARCH volatility forecast at time $t-1$. This correlation equals 29% when volume is low on day $t-1$ and is 17% when volume on day $t-1$ is high. In other words, ARCH volatility provides a "better" (i.e., more highly correlated) forecast for future volatility following a low volume market. It therefore seems likely that the ARCH versus options effect is at least partially driven by ARCH doing worse on high volume days than on low volume days. Conversely, the simple correlation between realized volatility at time t (again, as measured by time t's squared return innovation) and option-implied volatility at time $t-1$ equals 21% when volume is low on day $t-1$ and is 24% when volume on day $t-1$ is high. In other words, option-implied volatility provides a "better" (i.e., more highly correlated) forecast for future volatility when the options are traded in a high volume market. Thus, it appears that our combining regression findings are driven by both options doing better on high volume days and by ARCH doing worse. This finding is consistent with research by authors such as Admati and Pfleiderer (1988) and He and Wang (1995) which suggests that prices should be more informative on high volume days. Indeed, our test procedure suggests a possible new approach for examining empirically the relative informativeness of market prices.

4 Augmented ARCH

Forecast combining was appropriate above because our purpose was to reveal starkly the basic ARCH-volume-option relationship. Forecast combining is also a useful tool for situations in which the econometrician possesses the forecasts produced by various models but not the information sets used to produce the forecasts.[15] However, we do possess the information set on which at least the ARCH forecasts are based and thus, to produce optimal volatility forecasts, we should ideally add option and volume information to the ARCH model directly and estimate an augmented ARCH mega-model. We therefore investigate augmented ARCH models in this section.

4.1 The Model

The augmented-ARCH model we employ is given below, in which R_t is the daily stock return and S_t^2 is option-implied return variance.

$$R_t = \mu + \epsilon_t \quad ; \quad \epsilon_t \sim (0, \sigma_t^2) \tag{5}$$

$$\begin{aligned}\sigma_t^2 = \alpha_0 + \alpha_1 V_{t-1} + \beta_0 \sigma_{t-1}^2 + \beta_1 V_{t-1} \sigma_{t-1}^2 + \gamma_0 \epsilon_{t-1}^2 + \gamma_1 V_{t-1} \epsilon_{t-1}^2 \\ + \delta_0 D_{t-1} \epsilon_{t-1}^2 + \delta_1 V_{t-1} D_{t-1} \epsilon_{t-1}^2 + \phi_0 S_{t-1}^2 + \phi_1 V_{t-1} S_{t-1}^2\end{aligned} \tag{6}$$

$$D_{t-1} = \begin{cases} 1 & if\ \epsilon_{t-1} < 0 \\ 0 & otherwise \end{cases} \tag{7}$$

$$V_{t-1} = \begin{cases} 1 & if\ Volume_{t-1} \geq \frac{1}{(n-1)} \sum_{i=2}^{n} Volume_{t-i} \\ 0 & otherwise \end{cases} \tag{8}$$

Equations (5)-(8) are estimated jointly under Maximum Likelihood with $n = 5$, as are a number of interesting restricted and extended versions of (5)-(8). Results are reported below.

[15] See Clemens (1989).

To better understand the effects of adding lagged volume and implied volatilities to ARCH, we investigated all possible combinations and permutations within our model, including: each variable alone, each possible combination, variables interacting with each other, and so forth. We also expanded our model to investigate a variety of different functional forms for the conditional volatility, including variables – and groups of variables – added and interacted nonlinearly. Table 2 presents estimation results for a small collection of models which revealed the most interesting information concerning the effects of volume and implied volatility. Summary statistics for these most interesting models follow in Table 3. In all cases our core result – that options provide better forecasts relative to ARCH on high volume days than on low volume days – remains qualitatively robust, although different models reveal different types of information about the nature of our core result.

4.2 Parameter Estimates

Panel A of Table 2 below reports parameter estimates (with asymptotic standard errors in parentheses and Bollerslev-Wooldridge robust standard errors in square brackets) for the most interesting specifications contained within equations (5)-(8). Panel B reports common diagnostics for each model in question. These diagnostics include: the model log likelihood, the p-value from a traditional Ljung-Box (1978) test for symmetric ARCH at 24 lags, and the p-values from an Engle-Ng (1993) Sign Bias Test, Negative Sign Bias Test, Positive Sign Bias Test and Joint Sign Bias Test – all at 5 lags – for the presence of asymmetric ARCH effects. P-values below $Z\%$ therefore reject, with $Z\%$ confidence, the null hypothesis that there is no ARCH effect in favor of the alternative hypothesis that the model in question has uncaptured ARCH.

Table 2
Estimation Results and Diagnostics for
Volatility Models based on Equations (5)-(8)

Panel A: Parameter Estimates
(Asymptotic Standard Errors)
[Bollerslev-Wooldridge Robust Standard Errors]

Column: Parameter Name	1 Options Only	2 ARCH Model	3 ARCH plus Options	4 ARCH with Options and Volume	5 Full Model
μ	0.00033** (0.00014) [0.00016]	0.00043*** (0.00016) [0.00016]	0.00036*** (0.00015) [0.00015]	0.00037*** (0.00015) [0.00015]	0.00037*** (0.00015) [0.00015]
α_0		5.68E-7*** (1.15E-7) [1.62E-7]	2.77E-6** (1.33E-6) [2.47E-6]	3.04E-6** (1.38E-6) [2.52E-6]	6.11E-6** (2.71E-6) [5.18E-6]
α_1					-628E-8 (5E-6) [7.91E-6]
β_0		0.96748*** (0.0047) [0.00694]	0.59819*** (0.05965) [0.09629]	0.79072*** (0.04692) [0.07098]	0.68266*** (0.07616) [0.14434]
β_1				-0.425*** (0.16028) [0.23354]	-0.2128 (0.2043) [0.31094]
γ_0		0.01514*** (0.00631) [0.00756]	-0.0565*** (0.01353) [0.01858]	-0.0358** (0.01671) [0.02958]	-0.0655* (0.04769) [0.07434]
γ_1					0.01807 (0.05158) [0.07412]
δ_0		0.01096 (0.00872) [0.01402]	0.06927*** (0.02294) [0.03066]	0.04026* (0.02544) [0.03845]	0.04051 (0.06251) [0.111]
δ_1					0.01432 (0.06708) [0.11191]
ϕ_0			0.40286*** (0.14372) [0.26401]	0.13342 (0.1889) [0.27736]	0.20431 (0.20108) [0.35277]
ϕ_1				0.55495*** (0.22146) [0.34119]	0.45054** (0.2304) [0.38546]

* Significant at the 10% level, one-sided test.
** Significant at the 5% level, one-sided test.
*** Significant at the 1% level, one-sided test.

Table 2 Continued
Estimation Results and Diagnostics for
Volatility Models based on Equations (5)-(8)

Panel B: Diagnostics (Test P-values)					
Column: Test or Statistic	1 Options Only	2 ARCH Model	3 ARCH plus Options	4 ARCH with Options and Volume	5 Full Model
Log Like	6832.256	6820.543	6864.994	6868.812	6869.46
ARCH Test	0.723	0.969	0.881	0.771	0.812
Sign Bias	0.142	0.242	0.051	0.078	0.035
Neg. Bias	0.826	0.529	0.557	0.651	0.367
Pos. Bias	0.025	0.211	0.040	0.043	0.016
Joint Test	0.155	0.296	0.255	0.219	0.136

We begin our analysis of Table 2 by considering the results from the basic Glosten, Jagannathan and Runkle (1993) specification, as reported in Column 2 (labeled "ARCH"). Note from Panel A Column 2 that all the parameter estimates from the ARCH model are of the expected sign and magnitude and, from Panel B, that the model passes all standard specification tests at conventional significance levels (e.g., there are no p-values below 0.050 in Panel B, Column 2). Note in particular from Panel A that the parameter on lagged conditional volatility, β_0, is close to unity, which reveals the highly persistent nature of stock return volatility. Also note that the asymmetry parameter δ, which is intended to capture asymmetric volatility effects, is positive as expected, although not significant in our data. The insignificance of δ in the ARCH column (Column 2) of Table 2 might lead one to conclude that negative return innovations – i.e., negative ϵ – do not affect volatility in a significantly differently way than positive innovations. However, as shall soon be demonstrated, the standard GJR Sign-GARCH model is sufficiently misspecified that such conclusions from this model may be misleading.

Column 1 ("Options Only") of Table 2 reports results from option-implied volatility alone; i.e., results from the model in (5)-(8) estimated with all parameters in equation (7) set to zero except for ϕ_0, which is constrained to equal unity. The interesting results from Column 1 come from Panel B. Note in particular that the log likelihood from option-implied volatility alone exceeds the log likelihood from ARCH in Column 2. Also note from Panel B that option-implied volatility passes all of the various ARCH tests at the 1% significance level, and fails only the Positive Sign Bias ARCH test at 5% (p-value = 0.025, which is less than 0.050).

Column 3 of Table 2 reports results obtained by adding option-implied volatility to ARCH to produce an ARCH plus Option model. Recall that the parameter on lagged conditional volatility in the Column 2's basic ARCH model was $\beta_0 = 0.9675$. Note that, in Column 3's ARCH+Option Model, $\beta_0 = 0.5982$ and that this decline in β_0 is offset by the implied volatility parameter in Column 3 of $\phi_0 = 0.4029$ (i.e., $\beta_0 + \phi_0$ from Column 3 approximately equals β_0 from Column 2). In other words, as one might expect, Column 3 reveals that S^2_{t-1} and σ^2_{t-1} appear to trade off against each other somewhat. As S^2_{t-1} is added to ARCH, the parameter on σ^2_{t-1} declines by an offsetting amount. Note, however, that this tradeoff does *not* imply that S^2_{t-1} is devoid of additional information over what is contained in σ^2_{t-1}; i.e., S^2_{t-1} does not simply "replace" some part of σ^2_{t-1}. On the contrary, option-implied volatility does indeed add new information to ARCH, as evidenced by the sizable increase in the log likelihood value from Column 3 over either Column 1 or Column 2; an increase that leads to the ARCH+Option Model being preferred to either ARCH or Options alone, by the Akaike Information Criterion and other standard selection criteria. This finding suggests that the standard ARCH model, which omits option information, may be misspecified and thus potentially misleading. For example, note that the parameter on innovation asymmetry, δ_0, is now significant in Column 3 where it was insignificant in Column 2. This suggests that positive and negative return shocks may indeed have different effects on future volatility and that these effects can be more easily seen in a more carefully specified ARCH model that includes option information.[16]

---

[16]Note from Column 3 in Table 2 that the sign on the asymmetry dummy variable, δ_0, is positive as one might expect, but that γ_0, the coefficient on ϵ^2_{t-1} is negative and of approximately opposite magnitude. In other words, a negative shock does not move volatility (i.e., $\gamma_0+\delta_0 \approx 0$) while a positive shock lowers volatility. This may seem somewhat perplexing given that the basic ARCH specification in Column 2 suggests that *any* shock increases volatility. However, note that the model in Column 3 includes option-implied volatility and that this variable will already contain at least part of the market's response to the return innovation. Thus, γ_0 and δ_0 are more accurately interpreted as capturing the shock response not already accounted for in S^2_{t-1}, in which case γ_0 and δ_0 could reasonably have the values we observe in Column 3 of Table 2.

The next series of models we investigated are based on lagged trading volume and on various functions of lagged trading volume (e.g., log volume, change in volume, our volume dummy variable from equation (4), etc.) added to the standard ARCH model. The (unreported) results of these volume investigations revealed that, either alone or when added to a standard ARCH model, lagged volume has no power to *predict* volatility. From this one might be tempted to conclude, as previous researchers have concluded, that lagged volume has no power to forecast future volatility once the effects of lagged return innovations have been accounted for. However, such a conclusion would be premature. It would be more accurate to argue that, while volume cannot by itself forecast volatility, it does play an important regime-switching role, interacting with other variables in the model as we have already seen.

The far right column in Table 2 reports parameter estimates and diagnostics for the "Full Model" from Equations (5)-(8), in which volume is interacted with every possible variable in our information set. Here we see the result from the previous section that option-implied volatility receives increased weight on high volume days (i.e., $\phi_1 > 0$). As one might expect, the Full Model shows typical signs of over-fitting the data, including an almost complete absence of statistically significant regression coefficients. We therefore estimate in Column 4 of Table 2 a restricted version, called "ARCH with Options and Volume", in which volume is interacted with only option prices and the lagged conditional variance. Here we see that, on high volume days, the weight on option-implied volatility increases and the weight on the lagged conditional variance decreases (i.e., $\beta_1 < 0$ and $\phi_1 > 0$). In other words, the 'ARCH with Options and Volume" model from Table 2 confirms our findings from the simple combining exercise we conducted in Table 1 above.

Finally, we investigated a number of nonlinear extensions to equations (5)-(8), including models based on the Artificial Neural Network specifications. Interestingly, results from various (unreported) tests revealed that adding lagged volume and lagged option volatility to ARCH models removes the need to explicitly model nonlinear effects in lagged return innovations. In other words, volume, and especially options, seem to account for the nonlinear effects otherwise omitted from standard ARCH models. This suggests that ARCH modelers may benefit at least as much from expanding the traditional ARCH information set to include variables such as option-implied volatility and volume as from building ever more complex nonlinear models based only on lagged return innovations.

4.3 Summary Statistics

Table 3: Summary Statistics

Method	Variance Forecasts					Standardized Returns			
	Mean 10^{-4}	Std 10^{-4}	Skew	Kurtosis	RMSE 10^{-4}	Std	Skew	Kurtosis	Lnl
Const Var	0.632	1.902	15.207	338.12	1.901	1.000	-0.635	10.060	6690
ARCH	0.665	0.631	5.161	35.472	1.863	1.000	-0.670	9.281	6820
Option	0.573	0.414	2.479	12.904	1.851	1.071	-0.529	6.817	6832
ARCH+Opt	0.613	0.384	2.212	10.065	1.838	1.000	-0.488	6.606	6865
Combined	0.625	0.431	3.071	17.257	1.846	1.000	-0.508	6.771	6853
Part Model	0.613	0.401	2.648	13.876	1.834	1.000	-0.442	6.185	6868
Full Model	0.613	0.401	2.664	14.210	1.836	1.000	-0.430	6.102	6869

Table 3 presents some useful summary statistics on all of Table 2's models. The first row in Panel A of Table 3, marked "Const Var", reports results from a model assuming constant variance; i.e., equations (5)-(8) with all variance parameters zero except α_0 (this is essentially the raw data). The row marked "ARCH" is for the basic Glosten, Jagannathan and Runkle (1993) Sign-ARCH model in equations (1)-(3). "Option" signifies variance defined as lagged option-implied volatility; i.e., equations (5)-(8) with all variance parameters zero except ϕ_0. "ARCH+Opt" is options added to ARCH, but without volume; i.e., equations (5)-(8) with all parameters subscript 1 set to zero. The "Combined" model noted in Table 3 is the model from Section 3 of this paper, in which ARCH forecasts are first constructed separately and then combined with both volume and options. "Part Model" is the restricted version of the Full Model in which $\alpha_1 = \gamma_1 = \delta_1 = 0$, as in Column 4 of Table 2 above. The last row of Table 3 reports results from the "Full Model", which is the maximum likelihood combination of ARCH, volume and options obtained by estimating equations (5)-(8) will all variables in play, as in Column 5 of Table 2 above.

Of particular interest in Table 3 are results from the columns on Standardized Returns; i.e., $\hat{\epsilon}_t/\hat{\sigma}_t$. A desirable model is one that produces standardized returns with less (excess) kurtosis than the raw data. By this criterion the models with ARCH, volume and options (i.e., the Part and Full models) perform best (i.e., deliver the lowest kurtosis), while basic ARCH does worst. Table 3 also presents the root mean squared errors from the various models. Models with smaller squared/absolute errors are preferred. Again models with ARCH, volume and options appear to perform best; i.e., have the lowest mean squared error. Basic ARCH again does worst. Table 3 therefore confirms our conclusion from Table 2 that, by adding volume and implied volatility to a basic ARCH model, we are able to capture important effects that a basic ARCH model alone cannot capture. Furthermore, adding both volume and options to ARCH does better than adding either options or volume alone to ARCH.

5 Summary and Conclusions

Previous studies have reported that trading volume cannot forecast volatility directly. In this paper we uncover a new result: that volume does indeed have predictive power for forecasting volatility, with volume playing the role of a regime-switching variable between states in which option-implied volatility is more or less informative than ARCH for volatility forecasting. Indeed, we find that the accuracy of volatility forecasts can be significantly improved by accounting for the volume regime effect and by combining information from ARCH models and option prices accordingly. This finding is made possible because of the novel way we incorporate trading volume into our functional forms and because, while previous papers have added either trading volume or option-implied volatility (but not both) to ARCH models, our study is the first we know of to consider all three factors together.

Results produced by our investigation reveal that if trading volume was "lower than normal" during period $t-1$ then the best forecast of time t volatility is found by combining the ARCH forecast with the option-implied volatility forecast, with more weight being given to ARCH than to options. Conversely, if trading volume was "higher than normal" during period $t-1$, then the best forecast of time t volatility is obtained by placing more weight on options and less on ARCH. This result is robust to a variety of perturbations in sample period and model specification and seems to be largely driven by improvements in the quality of option-implied volatility forecasts during high volume periods.

Results from the combining exercise in Section 3 reveal that option prices contain better information about future volatility on high volume days than on low volume days. This suggests that market prices contain more information in high volume periods than in low volume periods (indeed, our work suggests a new way to test the relative informativeness of market prices in various volume regimes). Our results also suggest either that option markets are more efficient in high volume periods – in that not only do prices contain more information, they contain it more accurately – or else that option pricing models are less misspecified in high volume periods.

Results from the various ARCH tests in Section 4 further reveal that adding option and volume information to ARCH models greatly improves the ARCH models' forecasting performance and that, when volume and options are both added to ARCH models, various nonlinear representations of lagged return innovations are no longer needed to fit the data within an ARCH modeling context. This suggests that ARCH modelers may profit from expanding the traditional ARCH information set to include volume, options, and other types of information in addition to the history of lagged return innovations.

References

Admati, A. and P. Pfleiderer, 1988, "A Theory of Intraday Patterns: Volume and Price Variability," *Review of Financial Studies* 1, 3-40.

Bakshi, G.; C. Cao and Z. Chen, 1997, "Empirical Performance of Alternative Option Pricing Models," *Journal of Finance* 52, 105-134.

Bessembinder, H.; K. Chan and P.J. Seguin, 1996, "An Empirical Examination of Information, Differences of Opinion and Trading Activity," *Journal of Financial Economics* 40, 105-134.

Bollerslev, T., R.Y. Chou, and K.F. Kroner, 1992, "ARCH Modelling in Finance: A Review of the Theory and Empirical Evidence," *Journal of Econometrics* 52, 5-59.

Bollerslev, T. and J.M. Wooldridge, 1992, "Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time-Varying Covariances," *Econometric Reviews* 11, 143-172.

Brooks, 1998, "Predicting Stock Market Volatility: Can Market Volume Help?" *Journal of Forecasting* 17, 59-80.

Canina, L. and S. Figlewski, 1993, "The Information Content of Implied Volatility," *Review of Financial Studies* 6, 659-681.

Christensen, B.J. and N.R. Prabhala, 1998, "The Relation Between Implied and Realized Volatility," *Journal of Financial Economics* 50, 125-150.

Clemen, R.T., 1989, "Combining Forecasts: A Review and Annotated Bibliography," *International Journal of Forecasting* 5, 559-583.

Day, T. and C. Lewis, 1992, "Stock Market Volatility and the Information Content of Index Options," *Journal of Econometrics* 52, 267-287.

Donaldson, R.G. and M. Kamstra, 1997, "An Artificial Neural Network-GARCH Model for International Stock Return Volatility," *Journal of Empirical Finance* 4, 17-46.

Donaldson, R.G. and M. Kamstra, 1996, "Forecast Combining with Neural Networks," *Journal of Forecasting* 15, 49-61.

Dumas, B., J. Fleming, and R. Whaley, 1998, "Implied Volatility Functions: Empirical Tests," *Journal of Finance* 53, 2059-2106.

Engle, R.F. and V.K. Ng, 1993, "Measuring and Testing the Impact of News on Volatility," *Journal of Finance* 48, 1749-1778.

Fleming, J., 1998, "The Quality of Market Volatility Forecasts Implied by S&P 100 Index Option Prices," *Journal of Empirical Finance* 5, 317-345.

Glosten, L.R., R. Jagannathan and D.E. Runkle, 1993, "The Relationship Between Expected Value and the Volatility of the Nominal Excess Return On Stocks," *Journal of Finance* 48, 1779-1801.

Harris, M. and A. Raviv, 1993, "Differences of Opinion Make a Horse Race," *Review of Financial Studies* 6, 473-506.

He, H. and J. Wang, 1995, "Differential Information and the Dynamic Behavior of Stock Trading Volume," *Review of Financial Studies* 8, 919-972.

Heimstra, C. and J. Jones, 1994, "Testing for Linear and Nonlinear Granger Causality in the Stock Price-Volume Relation," *Journal of Finance* 49, 1639-1664.

Hull, J. and A. White, 1987, "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance* 62, 281-300.

Karpoff, J., 1987, "The Relation Between Price Changes and Trading Volume: A Survey," *Journal of Financial and Quantitative Analysis* 22, 109-126.

Lamoureux, C. and W. Lastrapes, 1993, "Forecasting Stock-Return Variance: Toward and Understanding of Stochastic Implied Volatility," *Review of Financial Studies* 6, 293-326.

Lamoureux, C. and W. Lastrapes, 1994, "Endogenous Trading Volume and Momentum in Stock Return Volatility," *Journal of Business and Economic Statistics* 12, 253-260.

Ljung, G.M. and G.E.P. Box, 1978, "On a Measure of Lack of Fit in Time Series Models," *Biometrika* 65, 297-303.

Richardson, M. and T. Smith, 1994, "A Direct Test of the Mixture of Distributions Hypothesis: Measuring the Daily Flow of Information," *Journal of Financial and Quantitative Analysis* 29, 101-116.