Journal of Business & Economic Statistics, April 1994, Vol. 12, No. 2

Endogenous Trading Volume and Momentum in Stock-Return Volatility

Christopher G. Lamoureux
John M. Olin School of Business, Washington University in St. Louis, St. Louis, MO 63130-4899

William D. Lastrapes
Department of Economics, Terry College of Business, The University of Georgia, Athens, GA 30602

This article examines the ability of volume data to shed light on the source of persistence in stock-return volatility. A mixture model, in which a latent common factor restricts the joint density of volume and returns, is used to relax the assumption of exogenous volume used in previous studies. We use a point-in-time signal-extraction procedure to identify this latent process and a calibrated simulation to conduct analysis of the viability of the model to explain important properties of the data. Using daily returns and volume on individual stocks, our procedure cannot accommodate serial dependence in squared returns.

KEY WORDS: Autoregressive conditional heteroscedasticity model; Mixture model; Random-effects model; Small-sample analysis.

Consider the following model for daily stock returns and trading volume:

$$r_t = \sigma_1 Z_{1t}\sqrt{F_t}, \tag{1}$$

$$V_t = \mu_2 F_t + \sigma_2 Z_{2t}\sqrt{F_t}, \tag{2}$$

and

$$F_t = \alpha_0 + \alpha F_{t-1} + \phi_t, \quad F_t \geq 0, \tag{3}$$

where r_t is the stock return on day t, V_t is daily trading volume, F_t is a latent mixing variable, Z_1 and Z_2 are mutually and serially independent stochastic processes with zero mean and unit variance, and ϕ_t is a serially independent random variable with zero mean that is restricted to ensure that F is always nonnegative. Equations (1) and (2) correspond to the mixture model of Tauchen and Pitts (1983), which interprets the latent process as the number of daily information arrivals to the market. Equation (3) generalizes their assumption that the information-arrival process is serially uncorrelated. The primary restriction of the model is that information arrival is a common factor affecting both daily returns and volume.

This article focuses on a key implication of the model: The dynamic properties of observed stock returns depend only on the dynamic properties of the latent mixing variable, F. For example, the model implies that the first-order autocovariance of squared returns is $\text{cov}(r_t^2, r_{t-1}^2) = \sigma_1^4\text{var}(F)\alpha$. For $\alpha = 0$, returns are serially independent. For $\alpha > 0$, however, squared returns are positively serially correlated. Such volatility persistence is an actual feature of stock-return data and was first documented by Mandelbrot (1963). It has more recently been modeled as a generalized autoregressive conditional heteroscedasticity (GARCH) process (Bollerslev, Chou, and Kroner 1992). In this article, we use daily trading volume to ask whether the documented persistence in daily individual stock-return variance can be attributed to the particular data-generating process in Equations (1), (2), and (3).

The use of volume data to learn about the dynamics of returns makes this an interesting exercise. As noted by Ross (1987), arbitrage theories of asset price behavior, typical of theoretical finance, say nothing about the role of trading volume. Although recent theoretical models have been developed to guide intuition about the relationship between volume and returns, "they have not evolved sufficiently to guide the specification of an empirical model of daily stock market data" (Gallant, Rossi, and Tauchen 1992, p. 202). Therefore, the task of current research is to develop a suitable empirical framework for using volume to understand market behavior. This article considers the ability of the simple mixture model to provide such a framework.

The model places many restrictions on the joint density of returns and volume. We focus on persistence in return variance because GARCH models have captivated empirical work in finance, despite their lack of theoretical rationale. This observation motivates the work of Lamoureux and Lastrapes (1990), who demonstrated that contemporaneous volume is a sufficient statistic for the entire history of squared returns in a GARCH specification. Although that result is of interest, it does not shed light on the ultimate source of variance persistence because inference relies on the assumption that volume is weakly exogenous. The empirical strategy here allows us to examine the ability of the mixture model to motivate GARCH models without conditioning on contemporaneous volume data.

Because of the difficulties of applying maximum likelihood techniques to estimate the joint density of r, V, and F when the latent variable is serially dependent, we employ a tractable signal-extraction strategy to isolate the model's implications regarding serial dependence in returns. The strategy entails extracting, for each observation in the sample, the value of F_t that sets the observed values r_t^2 and V_t as close as possible to the respective conditional means implied by the mixture model. Although ignoring the time series properties of the data in this step is inefficient, full-information extraction procedures are likely to be computationally burdensome and intractable. Given the estimated values for the latent process, we construct, according to Equation (1), an adjusted return series that is proportional to the ratio of returns to the square root of the extracted latent variable. Under the null hypothesis that the mixture model is valid, F_t filters the time dependencies from the return series so that adjusted returns do not exhibit persistence in variance. We construct statistics to examine this hypothesis. Because the asymptotic sampling distributions of the statistics derived from our estimation procedure are unknown, we use nonparametric simulations, calibrated to the data, to analyze the model's implications on the data-generating process. A similar strategy for conducting inference in small samples is found, for example, in the work of Cecchetti, Lam, and Mark (1990).

Section 1 details the empirical procedure. Section 2 describes the data, the small-sample properties of the extracted series and test statistics, and the results regarding the ability of the model to account for persistence in return variance. Section 3 concludes by discussing the implications of the results.

1. EMPIRICAL METHOD

The analysis involves extracting an estimated F process, adjusting returns according to the mixture model, and then testing for serial dependence in adjusted returns. The first step relies only on the common factor restriction implied by Equations (1) and (2); the level of volume and squared returns are influenced by a single factor and, furthermore, their serial dependence is due exclusively to the time series properties of the common factor. Thus let

$$e_t = \begin{pmatrix} e_{1t} \\ e_{2t} \end{pmatrix} = \begin{pmatrix} r_t^2 - E(r_t^2|F_t, \theta) \\ V_t - E(V_t|F_t, \theta) \end{pmatrix} \tag{4}$$

and

$$E(e_t e_t'|F_t, \theta) = \Sigma_t, \tag{5}$$

where $\theta = [\sigma_1\, \mu_2\, \sigma_2]$ and, from (1) and (2), $E(r_t^2|F_t, \theta) = \sigma_1^2 F_t$ and $E(V_t|F_t, \theta) = \mu_2 F_t$. Separately, for $t = 1 \ldots T$, we choose F_t to minimize, conditional on a consistent estimator for $\theta(\hat{\theta})$, the following conditional moment criterion:

$$L_t = \hat{e}_t' \hat{\Sigma}_t^{-1} \hat{e}_t = \frac{(r_t^2 - \hat{\sigma}_1^2 F_t)^2}{2(\hat{\sigma}_1^2 F_t)^2} + \frac{(V_t - \hat{\mu}_2 F_t)^2}{\hat{\sigma}_2^2 F_t}. \tag{6}$$

The second equality uses the fact that Σ_t is diagonal, due to the assumption that Z_1 and Z_2 are independent. The minimization is accomplished using a Newton–Raphson procedure with analytic derivatives, which always converged in two to four iterations.

Because the location of F cannot be identified, we normalize F to have unit unconditional mean (i.e., we set $\alpha_0 = 1 - \alpha$). Then the following unconditional moments are implied by the model:

$$E(V) = \mu_2, \tag{7}$$

$$\text{var}(r) = \sigma_1^2, \tag{8}$$

$$\frac{\text{cov}(r^2, V)}{\sigma_1^2 \mu_2} = \sigma_F^2, \tag{9}$$

and

$$\text{var}(V) - \mu_2^2 \sigma_F^2 = \sigma_2^2. \tag{10}$$

Consistent estimators of the parameters in θ are just-identified from these equations using method of moments, assuming that F is a stationary and ergodic process. Note that the *unconditional* covariance between squared returns and volume is nonzero and is exploited to obtain estimates of σ_2, despite the diagonality of Σ_t. This unconditional covariation is due solely to the joint dependence of returns and volume on the common factor F.

Given the extracted F_t series, $\hat{F}_t$, the adjusted return series can be constructed according to Equation (1):

$$\hat{Z}_{1t} = \frac{r_t}{\hat{\sigma}_1 \sqrt{\hat{F}_t}}, \tag{11}$$

where $\hat{\sigma}_1$ is the moment estimator of σ_1. We then ask whether the overidentifying restriction of the model that Z_{1t} has no persistence in variance seems plausible in light of the simulations. Thus we test whether the portion of squared returns not attributable to F_t is serially independent by examining the magnitude of the Ljung–Box Q statistic (see Granger and Newbold 1986, p. 100) for $\hat{Z}_{1t}^2$.

Even though we use only (r_t, V_t) to extract F_t, there is a sense in which our analysis depends on the entire sample and hence may benefit from increasing the sample size. First, the moment estimators from the first step use all of the sample. Second, our analysis focuses on the *serial correlation* of $\hat{Z}_1^2$, not on the value of $\hat{Z}_1$ itself. If the error in the extraction is large but serially independent, we can still hope to improve our understanding of the time series properties of Z_1 by adding additional data. Nevertheless, the asymptotic properties of the test statistic are unknown. We therefore conduct our analysis by simulating the sampling distribution of the test statistic under the null hypothesis that the full model is valid. The simulations are calibrated to the data and are described in Section 2.

We summarize the serial dependence in returns, in addition to the Q statistics, by estimating a GARCH (1, 1) model:

$$y_t = a_0 + a_1 y_{t-1} + \epsilon_t, \tag{12}$$

$$\epsilon_t \sim N(0, h_t), \tag{13}$$

and

$$h_t = \gamma_0 + \gamma_1 \epsilon_{t-1}^2 + \gamma_2 h_{t-1}, \tag{14}$$

where y_t is an observable random variable. To establish that our sample of returns exhibits GARCH persistence in variance, we report estimates of the model for $y = r$. We then examine whether this persistence remains after conditioning on the mixture model by repeating the estimation for $y = \hat{Z}_1$.

2. RESULTS

2.1 Simulations and Small-Sample Properties

We apply the procedure to a sample of daily returns and volume data for 10 individual companies covering the period January 3, 1967, through December 30, 1987 (4,858 trading days). All analysis is conducted on seven nonoverlapping subperiods consisting of 727 to 759 observations. The returns are cash dividend adjusted, close-to-close (last transaction of the day), taken from the Center for Research in Security Prices tapes. Volume is the number of shares traded in the day, taken from the Standard & Poor's *Daily Stock Price Guide*. In the event of a stock split, volume data on the ex-split day and subsequent days in the sample are divided by the split factor. There are no stock dividends in the sample.

Because the distributional forms of the random shocks Z_1, Z_2, and ϕ are not specified by the model, we use the data to determine the nature of the densities used in the simulations. The calibrated simulation proceeds in four steps. In the first, we choose a representative company/period (Bristol-Myers, period 1—1/2/70 to 12/29/72), perform the extraction procedure, and filter the implied variables $\hat{Z}_1$, $\hat{Z}_2$, and $\hat{F}$. The filters used are an (autoregressive) AR(1) in mean for $\hat{Z}_1$ and $\hat{Z}_2$, AR(3) in mean for $\hat{F}_t$, and GARCH(1, 1) for conditional variance in $\hat{Z}_1$ and $\hat{Z}_2$. The time series filters are specified based on the autocorrelation functions of the extracted variables.

In the second step, we randomly sample with replacement from the standardized residuals of the three time series models. These random draws are then used to simulate the three stochastic variables in the model over the sample period of 755 observations. Importantly, Z_{1t} is generated as a serially independent process because this is the maintained null hypothesis. However, Z_{2t} is recolored to allow for the serial dependence found in the data because such dependence is not inconsistent with serial independence of Z_1. Zero unconditional mean and unit unconditional variance of Z_{2t} is maintained in the simulation. F_t is simulated according to Equation (3), where α is set, alternately, to .5, .9, and .99 and the unconditional variance of the simulated F process is set equal to the sample variance of the extracted series of the representative company/period. Negative draws on F occur infrequently throughout the simulation. In such cases, F is set to its absolute value, which amounts to truncating the distribution of ϕ to ensure positive F.

Third, given the simulated values of Z_1, Z_2, and F, data for r and V are simulated over the sample period according to Equations (1) and (2). Values for σ_1 and σ_2 are set to .014 and 223, respectively (see Table 2, Sec. 2.2, Company/Period 3/1), but μ_2 equals 335 so that μ_2/σ_2 equals its grand mean over companies and subperiods of 1.5. To ensure that simulated volume is always positive, we truncate $\hat{Z}_2$, forcing its minimum value to be -1.13. This truncation affects the lower 7% tail of the empirical density of $\hat{Z}_2$.

Finally, we perform on the simulated data exactly the extraction and estimation procedures used on the actual data. In many cases, the moment estimator of σ_2 is negative. In such cases, we use a shrinkage estimator for σ_2. Specifically, from the raw data, we have moment estimators of the θ parameters for each of 10 companies in each of seven distinct periods. In 55 of these 70 cases, the moment estimate of σ_2 is nonnegative. The shrinkage estimator for σ_2 is $.667 \times \hat{\mu}_2$, where .667 represents the average ratio of $\hat{\sigma}_2$ to $\hat{\mu}_2$ across the 55 nonnegative cases (its grand mean).

The four steps are replicated 5,000 times to generate sampling densities, from which the properties of this empirical procedure can be determined. Table 1 reports these sampling densities for statistics when $\alpha = .5$ and .99. The nature of the results is similar for $\alpha = .9$. Despite the apparent bias in $\hat{\sigma}_2$ (especially for $\alpha = .99$), the median correlation coefficient between the simulated F_t process and the extracted $\hat{F}_t$ is close to 1, which is one measure of the effectiveness of the extraction procedure for the calibrated parameter values. This result is not sensitive to the degree of serial correlation in F; thus it is not likely that this correlation is spurious. Because $\hat{F}$ is an accurate estimator of F in the simulations, the properties of $\hat{Z}_1$ will locally mimic those of the true Z_1, even if the extraction procedure is not consistent. Note also from the Q_{25} statistics, which reflect the tests for serial correlation up to 25 lags, that the serial dependence in $\hat{F}$ closely mimics that of the actual F process.

From the distribution of $Q_{25}(r^2)$, it is clear that the model generates persistence in variance: The median value exceeds 37.65, the 95th percentile of the χ^2 distribution with 25 df, even when the persistence in F is relatively small (i.e., for $\alpha =$ .50). Because the extraction procedure is not asymptotically motivated, there is no reason to expect the Q statistics for $\hat{Z}_1$ and $\hat{Z}_1^2$ to be distributed as χ^2. Indeed, for $\alpha = .99$, the Q statistic for $\hat{Z}_1^2$ has values that exceed those of the χ^2_{25} (although for $\alpha = .5$ the empirical density is a much closer match). This density is used to examine the viability of the mixture model as the source of the persistence in squared returns in the following section.

The table also shows the actual moment estimator of $\sigma_2^2(\hat{\sigma}_2^{2*})$, as well as the estimator that uses the shrinkage technique instead of negative values ($\hat{\sigma}_2$). We obtain negative moment estimates of σ_2^2 in approximately 35% of the simulated samples. Thus the moment estimator may contain little information about the true value. This result implies that estimating a negative $\hat{\sigma}_2^2$ from moment conditions is not direct evidence against the model, and may cast doubt on standard generalized method of moments (GMM) inference for this model (Richardson and Smith 1992). Furthermore, it suggests that we should consider the sensitivity of our analysis to this estimator. In the next section, we therefore report results for the moment estimator and the shrinkage estimator for all companies and all subperiods.

2.2 Serial Dependence in Adjusted Returns

Table 2 contains the tests for variance persistence and the conditioning values of the model parameters. Only those cases in which GARCH persistence is evident in the return data are reported. As an illustration, consider the case of Bristol-Myers, subperiod 1. The Q_{25} statistic of 145.22 for

Table 1. Percentiles From Simulated Densities (over 5,000 replications) Under the Null Hypothesis

α	Statistic	5%ile	10%ile	25%ile	50%ile	75%ile	90%ile	95%ile
.99	$\text{cor}(F, \hat{F})$	.9332	.9417	.9552	.9661	.9731	.9784	.9811
.50	$\text{cor}(F, \hat{F})$	.9717	.9742	.9784	.9826	.9862	.9885	.9897
.99	cv_F^2	1.401	1.500	1.717	2.019	2.392	2.767	3.030
.50	cv_F^2	1.663	1.793	2.076	2.572	3.017	3.429	3.645
.99	$\hat{\sigma}_F^2$	1.254	1.374	1.581	1.895	2.268	2.637	2.923
.50	$\hat{\sigma}_F^2$	1.433	1.574	1.840	2.313	2.989	3.689	4.118
.99	$\hat{\sigma}_{\hat{F}}^2$	1.419	1.526	1.729	2.038	2.388	2.739	2.983
.50	$\hat{\sigma}_{\hat{F}}^2$	1.702	1.843	2.130	2.620	3.068	3.481	3.713
.99	$\hat{\sigma}_1$	.013	.013	.014	.014	.014	.014	.015
.50	$\hat{\sigma}_1$	.013	.013	.014	.014	.014	.015	.015
.99	$\hat{\mu}_2$	348.1	350.7	355.0	360.5	366.7	372.9	377.1
.50	$\hat{\sigma}_2$	341.6	343.0	345.5	348.1	350.8	353.0	354.4
.99	$\hat{\sigma}_2$	30.5	42.9	65.9	102.4	134.4	153.9	165.6
.50	$\hat{\sigma}_2$	70.8	99.5	164.3	197.0	211.0	261.1	300.6
.99	$\hat{\sigma}_2^{2*}$	−12,211	−8,337	−2,725	2,283	6,898	11,855	15,504
.50	$\hat{\sigma}_2^{2*}$	−98,777	−67,265	−24,724	5,847	35,818	68,166	90,378
.99	$Q_{25}(F)$	1,811.8	2,207.6	3,100.8	4,275.0	5,845.7	7,670.4	8,851.7
.50	$Q_{25}(F)$	130.1	140.5	161.7	186.2	215.1	250.5	279.3
.99	$Q_{25}(\hat{F})$	1,401.2	1,755.5	2,588.4	3,717.6	5,138.2	6,940.0	8,087.1
.50	$Q_{25}(\hat{F})$	120.7	131.3	153.6	178.3	208.0	243.5	270.2
.99	$Q_{25}(r^2)$	234.9	305.6	464.3	707.0	1,042.4	1,456.4	1,765.5
.50	$Q_{25}(r^2)$	19.3	23.7	33.7	48.6	71.3	100.6	122.0
.99	$Q_{25}(\hat{Z}_1)$	14.5	16.4	19.9	24.3	29.5	34.8	38.5
.50	$Q_{25}(\hat{Z}_1)$	14.7	16.6	20.0	24.4	29.2	34.9	37.9
.99	$Q_{25}(\hat{Z}_1^2)$	15.9	18.0	21.9	27.4	33.9	40.9	46.4
.50	$Q_{25}(\hat{Z}_1^2)$	14.7	16.7	20.0	24.4	29.5	34.9	38.2
.99	$Q_{25}(\hat{Z}_2)$	14.8	16.7	20.5	25.4	31.6	38.3	43.4
.50	$Q_{25}(\hat{Z}_2)$	14.3	16.2	19.8	24.3	29.4	35.0	30.0

NOTE: Sample size is 755. The actual parameter values are $\sigma_1 = .014$, $\mu_2 = 355$, $\sigma_2 = 222.931$, and the unconditional variance of F is 1.129. cv_F^2 is the sample squared coefficient of variation of the simulated F process, $\hat{\sigma}_F^2$ is the moment estimator of F from the simulated sample, and $\hat{\sigma}_{\hat{F}}^2$ is the sample variance of $\hat{F}$, the extracted latent common factor. $\hat{Z}_{1t} = r_t/(\hat{\sigma}_1\sqrt{\hat{F}_t})$. $\hat{\sigma}_1$, $\hat{\mu}_2$, and $\hat{\sigma}_2$ are the method-of-moment estimators. If the moment estimator for σ_2^2 is negative, then $\hat{\sigma}_2$ is redefined to be $\hat{\mu}_2/1.5$. The unadjusted moment estimator is reported as $\hat{\sigma}_2^{2*}$. $Q_{25}(X)$ represents the Ljung–Box test for serial dependence in X up to 25 lags.

the squared residuals of actual returns suggests the presence of persistence in variance. If this persistence is due to the mixture model under consideration, then this statistic for $y = \hat{Z}_1$ (using the moment estimator for σ_2) and $y = \hat{Z}_1^s$ (using the shrinkage estimator) would be drawn in repeated sampling from a distribution such as those characterized in Table 1. The comparison of the Q_{25} values of 420 and 463 with the 95% value of the density of Q_{25} under the null (i.e., 46.4 and 38.2 from Table 1) casts doubt on the plausibility of the mixture model—with parameter configurations as in the simulations—as the data-generating process. The result is the same when we examine the estimated GARCH coefficients. The sum of the GARCH coefficients is virtually the same for the estimated Z_1 variables as the value of .96 for the raw-return series. This finding is general throughout the table. In 10 of the reported cases, the Q statistic for the adjusted squared return falls below the 5% critical value when $\alpha = .99$; however, GARCH persistence remains in each of these cases.

It is possible that the results stem from a single day being insufficiently long to effectively purge noise from returns (French and Roll 1986). To examine this possibility, we repeat our extraction on 10-day returns, using the first 2,600 days only, in which there is little evidence of a trend in volume. Even with the 10-day returns and volume, the results (not reported) parallel those obtained using daily data.

We also perform, but do not report, out-of-sample comparisons and encompassing tests, which exploit a primary advantage of the extraction procedure over other strategies. The exercise entails using the extracted F and the model to forecast return variance and volume out-of-sample. We compare these forecasts to those of basic multivariate and univariate time series models of returns and volume. The forecasting ability of the model is generally worse than the time series models. In addition, encompassing regressions for both squared returns and volume indicate that the extraction-based forecast contains no incremental explanatory power over the other forecast models. These results are consistent with the preceding findings.

Table 2. Serial Dependence in (Adjusted) Returns

C/N		$y = r$	$y = \hat{Z}_1$	$y = \hat{Z}_1^s$	$\hat{\sigma}_1$	$\hat{\mu}_2$	$\hat{\sigma}_2$	$\hat{\sigma}_2^s$	$\hat{\sigma}_F^2$	$\hat{\sigma}_F^2$
1/1	γ_1	.05	.02	.02	.015	151.93	157.58	98.02	0.013	0.969
	γ_2	.90	.96	.96						
	Q_{25}	106.65	45.51	47.79						
1/3	γ_1	.11	.08	.08	.016	326.20	186.91	210.45	.109	.711
	γ_2	.83	.84	.84						
	Q_{25}	148.13	113.99	110.35						
1/4	γ_1	.04	.04	.04	.017	398.33	198.94	256.99	.265	.690
	γ_2	.96	.95	.95						
	Q_{25}	81.76	100.95	99.19						
1/6	γ_1	.10	—	.05	.021	1,170.89	—	650.97	1.021	.824
	γ_2	.89	—	.92						
	Q_{25}	282.45	—	209.55						
2/0	γ_1	.26	.17	.17	.014	183.78	110.93	118.57	.344	.821
	γ_2	.46	.54	.54						
	Q_{25}	48.41	66.37	66.04						
2/1	γ_1	.09	.09	.09	.011	202.87	106.78	130.89	.015	.650
	γ_2	.90	.89	.89						
	Q_{25}	304.78	596.36	569.79						
2/2	γ_1	.11	.07	.07	.016	285.14	40.81	183.96	.333	.589
	γ_2	.83	.89	.89						
	Q_{25}	49.50	83.90	83.51						
2/3	γ_1	.06	.05	.04	.011	343.25	128.44	221.45	.219	.605
	γ_2	.88	.88	.88						
	Q_{25}	57.20	38.34	40.43						
2/4	γ_1	.08	—	.05	.020	654.93	—	369.03	.343	.593
	γ_2	.92	—	.95						
	Q_{25}	214.45	—	193.05						
2/5	γ_1	.03	.03	.03	.019	926.54	327.64	597.78	.268	.629
	γ_2	.97	.97	.97						
	Q_{25}	161.44	274.37	266.93						
3/1	γ_1	.08	.07	.07	.014	259.15	222.93	167.19	.777	1.129
	γ_2	.87	.90	.90						
	Q_{25}	145.22	419.68	463.16						
3/2	γ_1	.06	.06	.06	.019	228.97	103.03	147.73	.309	.696
	γ_2	.89	.92	.92						
	Q_{25}	218.38	337.67	324.81						
3/3	γ_1	.23	.10	.10	.015	325.27	213.45	209.85	.301	.843
	γ_2	.37	.41	.41						
	Q_{25}	77.50	37.00	37.20						
3/6	γ_1	.19	—	.08	.021	1,041.05	—	574.84	1.097	.923
	γ_2	.82	—	.90						
	Q_{25}	341.86	—	724.91						

(continued)

3. CONCLUSION

In this article, we have attempted to uncover the source of persistence in shocks to stock-return variance, using daily volume data but without assuming that volume is exogenous. The maintained null hypothesis is that the dynamics of daily return variance are due solely to daily persistence in the latent speed of arrival of information to the market, which leads to similar dynamics in the level of trading volume. The small-sample approach of the article, which relies on nonparametric simulations calibrated to the data, indicates that accounting for serial dependence in the information-arrival process does not eliminate GARCH persistence in variance. Accordingly, we find that there is a distinction between the actual data and that generated from the mixture model of Equations (1)–(3). This finding is consistent with the results of Richardson and Smith (1992), who used GMM to test the mixture model but did not account for time dependencies in the data. Yet their rejection of the mixture model may be due to the lack of information contained in the moment estimator of σ_2, as suggested by our simulation results. Our evidence against the model is robust to this potential problem and isolates the inability of the model to jointly accommodate the dynamic properties of squared returns and volume.

Our analysis has focused on a particular version of the mixture model and does not apply to generalizations. One such extension to the mixture model of (1), (2), and (3) involves a more complex dynamic structure of the noise portion of volume (Z_2). For example, liquidity traders, facing nontrivial transactions costs, may trade only after large price moves (in

Table 2. (continued)

C/N		$y = r$	$y = \hat{Z}_1$	$y = \hat{Z}_1^s$	$\hat{\sigma}_1$	$\hat{\mu}_2$	$\hat{\sigma}_2$	$\hat{\sigma}_2^s$	$\hat{\sigma}_F^2$	$\hat{\sigma}_F^2$
4/0	γ_1	.09	.10	.11	.017	104.76	136.89	67.58	.208	1.197
	γ_2	.84	.80	.76						
	Q_{25}	110.37	79.41	93.08						
4/2	γ_1	.13	.08	.11	.025	112.49	320.77	72.57	.469	2.350
	γ_2	.74	.90	.86						
	Q_{25}	137.18	275.71	317.57						
4/3	γ_1	.14	.09	.09	.017	166.47	124.91	107.40	.470	.963
	γ_2	.56	.65	.68						
	Q_{25}	56.53	31.58	32.17						
4/4	γ_1	.14	.07	.08	.018	253.72	186.56	163.69	.353	1.884
	γ_2	.78	.86	.85						
	Q_{25}	80.51	78.94	81.62						
4/5	γ_1	.04	.05	.05	.022	619.45	585.58	399.65	.599	1.104
	γ_2	.91	.95	.95						
	Q_{25}	74.26	185.00	173.47						
5/1	γ_1	.06	.06	.06	.032	26.96	17.46	17.40	.817	1.040
	γ_2	.87	.92	.92						
	Q_{25}	88.91	88.50	88.60						
5/2	γ_1	.03	.05	.05	.048	25.80	35.75	16.64	1.584	1.565
	γ_2	.97	.95	.94						
	Q_{25}y	48.86	237.70	391.97						
5/6	γ_1	.20	.09	.09	.026	44.11	48.12	28.46	1.072	1.455
	γ_2	.56	.79	.90						
	Q_{25}	192.23	92.49	119.99						
6/1	γ_1	.08	.11	.12	.013	66.65	108.73	43.00	.238	1.459
	γ_2	.85	.81	.82						
	Q_{25}	82.28	136.07	187.89						
6/2	γ_1	.08	.12	.12	.015	66.07	45.84	42.62	.304	.893
	γ_2	.90	.71	.70						
	Q_{25}	120.58	110.39	112.78						
6/3	γ_1	.08	.03	.03	.010	114.75	115.44	74.03	.092	.945
	γ_2	.70	.88	.87						
	Q_{25}	54.09	29.95	29.54						
6/4	γ_1	.13	.08	.09	.015	192.79	270.93	124.38	.180	1.269
	γ_2	.72	.86	.86						
	Q_{25}	425.82	159.14	216.33						
6/5	γ_1	.10	—	.07	.017	934.13	—	336.13	1.966	1.372
	γ_2	.87	—	.87						
	Q_{25}	103.89	—	77.48						
6/6	γ_1	.20	.06	.08	.014	1,317.94	2,769.96	850.28	.356	1.931
	γ_2	.66	.76	.72						
	Q_{25}	92.98	37.08	44.69						

(continued)

either direction) because the price must cross through a no-trade zone (Davis and Norman 1990). This behavior implies that Z_{2t} depends on F_{t-j} and/or Z_{1t-j}, $(j = 1, \ldots)$. Furthermore, to the extent that this behavior is public knowledge, returns are independent of these dynamics; that is, Z_1 remains serially independent. In this case, the null hypothesis of interest is valid, but our extraction procedure will introduce time dependence in Z_1. We verified this conjecture by imposing dependence of Z_{2t} on the absolute value of r_{t-1} and simulating as previously. The Q statistics for $\hat{Z}_1^2$ are much larger than those generated under the null and reported in Table 1. These results indicate a need for future research along these lines.

A skeptical reader may feel that the analysis is not valid because the asymptotic properties of the tests are unknown and that the simulations that we have analyzed and reported are unrepresentative. Even from this perspective, our analysis sheds new light on important characteristics of the data: *Contemporaneous* squared returns and volume are noisy predictors of future return volatility. That is, r_t^2 and V_t are not very useful instruments in predicting r_{t+1}^2 in the context of the mixture model. This result is consistent with the finding in GARCH specification searches (see Bollerslev et al. 1992) that long lags of r^2 are required for forecasting r_{t+1}^2. Contrast this result, however, with the finding of Lamoureux and Lastrapes (1990) that, for purposes of predicting r_t^2, V_t is a sufficient statistic for r_{t-j}, $(j = 1, \ldots)$. If indeed our analysis is distorted because the data are much noisier than the simulated environments, then this would motivate an attempt to reduce the noise in the extraction by using the entire sample

Table 2. (continued)

C/N		$y = r$	$y = \hat{Z}_1$	$y = \hat{Z}_1^s$	$\hat{\sigma}_1$	$\hat{\mu}_2$	$\hat{\sigma}_2$	$\hat{\sigma}_2^s$	$\hat{\sigma}_F^2$	$\hat{\sigma}_F^2$
7/0	γ_1	.06	—	.06	.021	203.49	—	86.45	1.566	1.102
	γ_2	.92	—	.91						
	Q_{25}	126.08	—	100.59						
7/2	γ_1	.13	.07	.07	.021	90.61	92.42	58.46	.419	1.81
	γ_2	.62	.84	.85						
	Q_{25}	84.11	96.00	98.44						
7/3	γ_1	.06	.04	.04	.012	176.42	151.02	113.82	.431	1.051
	γ_2	.78	.87	.89						
	Q_{25}	46.16	33.61	34.91						
7/4	γ_1	.05	.04	.04	.020	338.83	289.25	218.61	.251	.934
	γ_2	.85	.92	.92						
	Q_{25}	44.83	43.04	47.55						
7/6	γ_1	.16	.07	.07	.024	862.41	474.53	556.39	.493	.950
	γ_2	.03	.89	.90						
	Q_{25}	130.22	279.61	264.01						
8/2	γ_1	.15	.09	.09	.040	94.85	78.05	61.19	1.206	1.281
	γ_2	.72	.90	.90						
	Q_{25}	68.79	294.44	299.19						
8/5	γ_1	.06	.04	.04	.035	308.13	199.63	198.79	.867	1.063
	γ_2	.88	.95	.95						
	Q_{25}	64.37	57.40	57.45						
8/6	γ_1	.21	—	.07	.047	348.01	—	143.87	1.611	1.118
	γ_2	.53	—	.89						
	Q_{25}	226.89	—	107.28						
9/1	γ_1	.05	.05	.05	.024	153.72	129.82	99.17	.326	1.004
	γ_2	.94	.94	.94						
	Q_{25}	297.26	360.73	384.96						
9/2	γ_1	.05	.02	.02	.022	238.73	78.39	154.02	.366	.671
	γ_2	.92	.98	.98						
	Q_{25}	94.77	47.63	43.48						
9/6	γ_1	.18	—	.04	.024	880.66	—	465.16	.782	.735
	γ_2	.66	—	.94						
	Q_{25}	145.25	—	73.62						
10/1	γ_1	.02	.12	.14	.016	14.20	21.30	9.16	.546	1.553
	γ_2	.98	.83	.82						
	Q_{25}	115.97	225.06	286.35						
10/2	γ_1	.06	.03	.03	.028	12.95	35.28	8.35	.835	2.452
	γ_2	.94	.94	.96						
	Q_{25}	109.17	29.65	91.17						
10/3	γ_1	.10	.05	.05	.021	18.99	12.73	12.25	.699	1.014
	γ_2	.85	.93	.93						
	Q_{25}	254.88	281.13	287.58						
10/4	γ_1	.06	.04	.04	.023	18.93	15.45	12.21	.665	1.107
	γ_2	.89	.91	.91						
	Q_{25}	52.75	30.85	32.33						
10/5	γ_1	.12	.09	.09	.021	33.12	23.47	21.37	.884	1.077
	γ_2	.83	.86	.86						
	Q_{25}	102.28	151.10	153.45						
10/6	γ_1	.14	—	.08	.021	66.56	—	25.16	3.925	1.456
	γ_2	.75	—	.88						
	Q_{25}	540.69	—	140.73						

NOTE: The model estimated is Equations (12), (13), and (14) in the text. C denotes company: 1 Alcoa, 2 Amoco, 3 Bristol-Meyers, 4 Champion, 5 Helene Curtis, 6 DPL (Dayton Power and Light), 7 Enserch, 8 Genesco, 9 Hewlett-Packard, 10 Kansas City Southern. N denotes subperiod: 0—1-3-67 to 12-31-69, 1—1-2-70 to 12-29-72, 2—1-2-73 to 12-31-75, 3—1-2-76 to 12-29-75, 4—1-2-79 to 12-31-81, 5—1-4-82 to 12-31-84, 6—1-2-85 to 12-30-87. $\hat{Z}_{1t} = r_t/(\hat{\sigma}_1\sqrt{\hat{F}_t})$, and $\hat{Z}_{1t}^s = r_t/(\hat{\sigma}_1\sqrt{\hat{F}_t^s})$, where $\hat{F}_t$ is extracted using the moment estimator of σ_2 and $\hat{F}_t^s$ is extracted using the shrinkage estimator of σ_2. $\hat{\sigma}_2^s$ is the shrinkage estimate of σ_2, $\hat{\sigma}_F^2$ is the sample variance from extracted F series, adjusted for sample mean (moment estimator is used when available). — denotes a case in which the moment estimator of σ_2 is negative. See note to Table 1 for other definitions.

jointly to extract the F process. One interpretation of the point-in-time method developed in this article is that it is an initial step in an adaptive estimation strategy, providing a first look at the properties of the latent F process.

ACKNOWLEDGMENTS

We are grateful to Yacine Aït-Sahalia, Torben Andersen, Kerry Back, Tim Bollerslev, Sid Chib, John Cochrane, Phil Dybvig, Bruce Lehmann, Mahesh Maheswaran, and George Tauchen for useful discussions.

[Received April 1993. Revised October 1993.]

REFERENCES

Bollerslev, T., Chou, R.Y., and Kroner, K.F. (1992), "ARCH Modeling in Finance," *Journal of Econometrics*, 52, 5–59.

Cecchetti, S.G., Lam, P.-S., and Mark, N.C. (1990), "Mean Reversion in Equilibrium Asset Prices," *American Economic Review*, 80, 398–418.

Davis, M.H.A., and Norman, A.R. (1990), "Portfolio Selection With Transactions Costs," *Mathematics of Operations Research*, 15, 676–713.

French, K.R., and Roll, R. (1986), "Stock Return Variances: The Arrival of Information and the Reaction of Traders," *Journal of Financial Economics*, 17, 3–26.

Gallant, A.R., Rossi, P.E., and Tauchen, G. (1992), "Stock Prices and Volume," *Review of Financial Studies*, 5, 199–242.

Granger, C.W.J., and Newbold, P. (1986), *Forecasting Economic Time Series* (2nd ed.), San Diego: Academic Press.

Lamoureux, C.G., and Lastrapes, W.D. (1990), "Heteroscedasticity in Stock Return Data: Volume vs. GARCH Effects," *Journal of Finance*, 45, 487–498.

Mandelbrot, B. (1963), "The Variation of Certain Speculative Prices," *The Journal of Business*, 36, 394–419.

Richardson, M., and Smith, T. (1992), "A Direct Test of the Mixture of Distributions Hypothesis: Measuring the Information Flow Throughout the Day," mimeo, Wharton School, Finance Dept.

Ross, S.A. (1987), "The Interrelations of Finance and Economics: Theoretical Perspectives," *American Economic Review*, 77, 29–34.

Tauchen, G.E., and Pitts, M. (1983), "The Price Variability-Volume Relationship on Speculative Markets," *Econometrica*, 51, 485–505.