INTRA-DAY VOLATILITY COMPONENTS IN FTSE-100 STOCK INDEX FUTURES

ALAN E. H. SPEIGHT*
DAVID G. MCMILLAN
OWAIN AP GWILYM

Recent research has suggested that intra-day volatility may contain both short-run and long-run components due to the existence of heterogeneous information flows or heterogeneous market agents (Andersen & Bollerslev, 1997a, 1997b; Müller et al., 1997). We report direct evidence for the existence of such a volatility decomposition in intra-day UK FTSE-100 futures returns data at frequencies of one hour and higher using the permanent–transitory component variance model of Engle and Lee (1993). Moreover, the transitory component identified exhibits rapid decay, volatility at the half-day frequency being completely dominated by the highly persistent permanent component. The model also is able to capture all dependency within the data at frequencies of one hour and lower. © 2000 John Wiley & Sons, Inc. Jrl Fut Mark 20:425–444, 2000.

INTRODUCTION

In the analysis of high-frequency financial data, several researchers recently have suggested that heterogeneous market volatility components

*Correspondence author, Senior Lecturer, Department of Economics, University of Wales, Swansea, SA2 8PP, UK; e-mail: a.speight@swan.ac.uk

■ *Alan E. H. Speight is a Senior Lecturer in the Department of Economics at the University of Wales in Swansea, United Kingdom.*

■ *David G. McGillan is a Lecturer in the Department of Economics at the University of St. Andrews in the United Kingdom.*

■ *Owain ap Gwilym is a Lecturer in the Department of Management at the University of Southampton in the United Kingdom.*

The Journal of Futures Markets, Vol. 20, No. 5, 425–444 (2000)

may exist at the intra-day level (e.g., Andersen & Bollerslev, 1997a, 1997b; Müller et al., 1997).[1] This suggestion has been motivated in large part by the apparently paradoxical observation that while the degree of intertemporal volatility dependence is very high in inter-day and lower frequency returns, the persistence of intra-day volatility shocks is extremely low in comparison, the latter often exhibiting half-lives of only several hours.[2] In particular, Andersen and Bollerslev (1997b) have suggested that market volatility may reflect the aggregation of numerous independent volatility components, each of which is endowed with a particular dependence structure due to the arrival of heterogeneous information. This 'heterogeneous information' extension of the 'information-flow' (or 'mixture-of-distributions hypothesis': MDH; Clark, 1973) approach to market volatility imparts both short-run and long-run volatility effects. If the decay of the short-run volatility component dominates over intra-day frequencies, and the long-run volatility component dominates over inter-day and lower frequencies, the aggregation of such component processes then gives rise to the (near-) integrated and long-memory lower-frequency dependencies that have been shown to characterize many returns volatilities. Using an alternative approach, Müller et al. (1997) argue that such volatility structure may arise due to heterogeneous traders rather than heterogeneous information flows. That is, due to market participants possessing different time horizons, such that short-term traders evaluate the market at a higher frequency and have shorter memory than long-term traders, again resulting in a component structure to volatility.[3]

However, as Andersen and Bollerslev note, while strong intra-day volatility periodicity and clustering have been documented, and they themselves report qualitative supporting evidence from high-frequency intra-day foreign exchange returns, there is little direct evidence to date

[1] Also see, for example, Guillaume, Pictet, and Dacorogna (1995) and Ghose and Kroner (1996). For a survey of the intra-day literature, see Goodhart and O'Hara (1997).

[2] See the references given in the text, n.1, and Dacorogna, Müller, Olsen, and Pictet (1997). Relatedly, it also has been reported that in the estimation of GARCH models for intra-day data at different frequencies, coefficient estimates obtained are in conflict with the theoretical aggregation results for GARCH models established by Drost and Nijman (1993) and the consistency theorems established by Nelson (1990, 1992). See, for example, Andersen and Bollerslev (1997a).

[3] As a consequence of this distinction in the activities of traders, Müller et al. (1997) argue that volatility measured at different time resolutions will reflect the perceptions and actions of different market components (i.e., different types of traders). While the similarity of this decomposition with the component model of Engle and Lee (1993) investigated here is noted by Müller et al., that empirical model is not analyzed further by those authors, a heterogenous-ARCH (HARCH) model being implemented empirically (that model, in fact, being a special case of the QARCH model introduced by Sentana, 1995). For related alternative component-type model estimates for inter-day stock return volatility, see den Hertog (1994), Engle and Lee (1993), and Palm and Urbain (1995), and for interest rate volatility, Jones, Lamont, and Lumsdaine (1998).

for the existence of heterogeneous volatility components of the kind implied by the extended MDH in intra-day volatility processes.[4] Perhaps the most obvious approach to evaluating such hypotheses would be to extend an existing agenda in the literature (e.g., Goodhart, Hall, Henry, & Pesaran, 1993) and attempt to associate specific volatility components with the economic factors generating explicit information flows. Unfortunately, as Andersen and Bollerslev also note, such an approach would not seem feasible in high-frequency intra-day data given the very many and diverse factors generating such information flows. However, a tractable alternative is readily offered by the statistical unobserved component model for conditional variance introduced by Engle and Lee (1993), which is comprised of permanent and transitory components, and is similar in spirit to the Beveridge–Nelson (1981) decomposition for economic time series.[5] This paper applies such a component model to intra-day U.K. stock index futures market returns data at various frequencies in an effort to determine whether permanent and transitory components can be explicitly identified in such data and, where present, whether the persistence of short-run volatility diminishes as the intra-day frequency decreases. In anticipation of the results reported below, we find that such a decomposition indeed does hold for U.K. futures at very high intra-day frequencies, whilst the permanent volatility component dominates transitory volatility at the half-day frequency. Further consideration of the role of information flows, as proxied by trading volume, offers empirical support to the heterogeneous information MDH interpretation advocated by Andersen and Bollerslev (1997a, 1997b) at all but the higher intra-day frequencies considered, the heterogeneous volatility structure at those higher frequencies being more likely attributable to a particular

[4]Specifically, Andersen and Bollerslev (1997b) provide qualitative supporting evidence from various sampling frequencies of five-minute Deutschmark–U.S. Dollar exchange rate returns, October 1992–September 1993. In particular, they demonstrate that the dynamic response of such returns varies systematically across identifiable macroeconomic announcements, which is consistent with the notion of distinct underlying components. They further show, by examining the autocorrelation function of those absolute returns, that high persistence is recorded for time intervals of around five to ten minutes. However, for returns over three and four hours, they note low persistence and, indeed, negative values in the autocorrelation function (which they term antipersistence) as the short-run effect decays. When the series is aggregated to a daily frequency, strong persistence (in long-run volatility) again is recorded. Andersen and Bollerslev attempt to formally separate these effects through the use of a band pass filter in the frequency domain and demonstrate that the volatility process exhibits long-memory dependency in terms of a degree of fractional integration in absolute returns which is invariant to sampling frequency, as implied by the theoretical aggregation of such components under suitable conditions. See Andersen and Bollerslev (1997b) for further details. Also see n.3 above.

[5]The component model also provides a natural extension of the GARCH model in terms of generalizing the latter's ARIMA representation for the squared residuals. For further details, see Bollerslev (1986) and Engle and Lee (1993).

class of traders with very short intra-day horizons as suggested by Müller et al. (1997). Our results also allow us to assess the extent to which homogenous GARCH processes understate volatility persistence at the higher intra-day frequencies, and to determine at what frequencies GARCH and component models are able to adequately capture all non-linear dependence in the data.

The remainder of the paper is organized as follows. The Institutional Considerations section considers the institutional background to the futures market investigated here, with particular reference to issues of liquidity services and information trading that further motivate consideration of heterogeneous volatility components. Market Structure and Data follows, which describes the specific market structure from which the data is drawn and some basic properties of the data. The section titled The Component Model outlines the structure and properties of the component model and its relationship to the basic GARCH model. Conditional Volatility Results reports GARCH and component model estimates and residual diagnostics. The results of introducing volume into these specifications in an effort to determine whether the component structure to volatility is driven by heterogeneous information flows or heterogeneous market agents are presented under Volatility–Volume Interaction. And finally, under Summary and Conclusions, we summarize our findings and offer some concluding remarks.

INSTITUTIONAL CONSIDERATIONS

While it is generally recognized that information flows are the predominant cause of price changes in asset markets, public information is likely to have a more important role than private information in stock index futures markets. As Tse (1999) notes, under open outcry trading, informed traders are easily identified by other traders in the pit, thus the adverse selection problem induced by any information asymmetry may not be severe. Moreover, since the underlying asset is a basket of stocks, a diversification effect reduces the impact of both private and stock-specific information, and whereas it may be optimal to trade in the underlying market if stock-specific information is held, market-wide information may be exploited more easily and cheaply using index futures. Indeed, numerous studies have identified the impact of public information in the form of macroeconomic announcements on prices and volatility in futures markets.[6] It is not unreasonable to expect that different news

[6]For a recent survey, see ap Gwilym and Sutcliffe (1999).

events, in the sense of news concerning different macroeconomic fundamentals, will impart differing volatility dynamics, in the sense of differing longevities in their impact on the volatility of market prices. Such heterogeneous information flows potentially can give rise to volatility components in accord with the hypothesis of Andersen and Bollerslev (1997a, 1997b) described in the Introduction.

Another important aspect of most futures markets that differentiate them from other securities markets is the absence of designated market makers. Liquidity, in the sense of a willingness to acquire a risky position by trading, generally is provided by local traders (agents trading for their own account), especially when order imbalances develop. Though not obliged to do so, these agents act as market makers through their presence in the market and their very short trading horizon. For instance, local traders can stimulate trade in a static market by narrowing the traded bid-ask spread. Floor traders in a futures market can be classified as scalpers, day traders, or position traders. Scalpers and day traders both have intra-day trading horizons, with scalpers often holding positions for only a few minutes. Both types of trader typically are anxious to close positions before the end of the trading day. In contrast, position traders have longer inter-day horizons. These heterogeneous trading horizons potentially can give rise to volatility components in accord with the hypothesis of Müller et al. (1997) outlined in the Introduction.

Heterogeneities in futures market volatility therefore might be expected to arise from both heterogeneous information flows and heterogeneous trader horizons. Moreover, information trading and liquidity services in futures markets can be expected to be interwoven, particularly at short intervals, and be inherently difficult to separate empirically.[7] One means of resolving their separate effects, given a suitable measure or proxy for information flows, is to examine the ability of such a proxy to explain observed volatility. Traded volume is regarded widely as providing such a proxy, and provides the basis for a discriminatory test procedure described later. Prior to the description of the empirical model and results, and augmentation to accommodate consideration of volume, we first describe the properties of the data analyzed here and its market context in the following section.

MARKET STRUCTURE AND DATA

The underlying dataset includes all trades on FTSE-100 stock index futures contracts from January 1992 to June 1995. During the period stud-

[7]We are grateful to an anonymous referee for drawing our attention to this issue.

ied, futures contracts on the U.K. FTSE-100 stock market index were traded at the London International Financial Futures and Options Exchange (LIFFE), under an open outcry trading system, from 08:35–16:10GMT.[8] This trading system is very similar to those at U.S. futures markets (e.g. the Chicago Mercantile Exchange, CME). The data source is the LIFFE Time and Sales CD-ROM, which contains information on the time to the nearest second, contract type, commodity code, delivery month, price, transaction code (bid, ask or trade), and traded volume. Time and sales data from U.S. futures exchanges, such as the CME and the Chicago Board of Trade, only includes data on a particular trade if it involved a change in price from the last trade. The data available from LIFFE is potentially far more informative since it contains a complete record of trade data, regardless of whether a price change was involved or not. Hence, transactions at the same price also are recorded and the complete trading history is available.

The delivery months for this contract are March, June, September, and December, with the nearest three available for trading. All contract maturities are traded in the same pit, and there could be up to fifty price makers active at any time. Since multiple maturities are traded at any given time, criteria must be set in order to create a continuous returns series. We follow the volume crossover procedure in using price data on the front month contract until traded volume in that contract is exceeded by traded volume in the second nearest contract. For stock index futures, this date tends to be very close to the expiry date of the front month contract. In our sample, rollover occurred on the day immediately prior to the expiry date.

The contract has a minimum price movement of 0.5 index points, which had a value of £12.50 during the period studied.[9] At the average price for FTSE-100 futures in our sample of 2966.4, the tick value is 1.7 basis points $[12.5 \div (25 \times 2966.4)]$. Table I highlights the discreteness of price changes for the underlying transaction-to-transaction price data over this sample period. The first point to note is a particular feature of the FTSE-100 futures contract, which is that 'odd' tick prices (e.g., 3000.5) are rarely used. Around 98% of traded prices are at full index points, since the market does not appear to require the additional price resolution of the half index points. The result of this is that there are very

[8]Since May 1999, this contract has been traded electronically on LIFFE Connect.

[9]The value of the minimum price movement was subsequently reduced to £5. The first contract with the new unit of trading was the June 1998 contract, with notice to the market issued in December 1997. This reflects the gains made by the underlying index over time that resulted in an increase in the nominal value of the contract.

TABLE I

Values Taken by the Price Change in FTSE100 Stock Index Futures: January 1992 to June 1995

	Transaction-by-Transaction		*At 5-Minute Intervals*	
Price Change (in ticks)	*Number of Observations*	*% of Sample*	*Number of Observations*	*% of Sample*
<−6	651	0.07	5266	6.84
−6	2324	0.27	4246	5.52
−5	49	0.01	129	0.17
−4	20305	2.33	7600	9.87
−3	607	0.07	274	0.36
−2	196523	22.54	11983	15.57
−1	7447	0.85	476	0.62
0	416805	47.80	16772	21.79
+1	7285	0.84	497	0.65
+2	196289	22.51	12310	15.99
+3	544	0.06	256	0.33
+4	20128	2.31	7610	9.88
+5	45	0.01	109	0.14
+6	2335	0.27	4344	5.64
>+6	698	0.08	5114	6.64
Total	872035	100.00	76986	100.00

few price changes occurring at ±1, 3, 5, etc. ticks.[10] Table I also considers the discreteness of price changes over the five-minute intervals used in constructing the basic returns series analyzed below. The problem of discreteness is less severe in this case, and a larger range of values (beyond ±6 ticks) is observed, as one would intuitively expect over a longer interval. While it remains the case that division by the price to generate returns does not result in a truly continuous variable at high frequencies, additional lengthening of the differencing interval (see below) further reduces the problem of discreteness.[11]

The tick data is sampled at five-minute intervals to generate the basic returns series $r_t = \log(P_t/P_{t-1})$, where P is the price of the security. For the first interval of each day, P_{t-1} is the opening trade price of the day. Thus, the overnight return is excluded to ensure a consistent time series. Volume, V_t, is simply the total traded volume for each five-minute interval.

[10]For a detailed investigation of this clustering phenomenon, see ap Gwilym, Clare, and Thomas (1998).

[11]A caveat on our analysis is that a comparison of different models of the variance of a continuous process may be revealing which model better approximates this fundamental discreteness, a problem that is generic to high-frequency studies. We are grateful to an anonymous referee for raising this point.

The 846 trading days in the sample yield 76,896 observations at the five-minute frequency. This data also is aggregated over an additional three frequencies to generate fifteen-minute, one-hour, and half-day returns and volume series, with respective sample sizes of 25,662, 6415, and 1710. In order to avoid confounding the dynamic dependencies in the data by the strong intra-day volatility patterns that have been documented widely, we follow Andersen and Bollerslev (1997b) in standardizing these series by the average absolute return or traded volume, respectively, for each particular intra-day frequency interval.[12]

THE COMPONENT MODEL

Following Engle and Lee (1993), let r_t denote the return on an asset, the expected return being m_t, and define the conditional variance of that return as $h_t \equiv Var(r_t|\Omega_{t-1}) = E[(r_t - m_t)^2|\Omega_{t-1}]$ where $\Omega t-1$denotes the set of all information available at time $t-1$. The simple GARCH(1,1) process (Bollerslev, 1986) is then defined by:

$$r_t = m_t + \varepsilon_t \quad (1)$$

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

where $(\omega, \alpha, \beta)$are fixed parameters, ε_t serially is uncorrelated with zero mean and conditional variance h_t as already stated, and the standardized error, $z_t = \varepsilon_t/\sqrt{h_t}$, is identically and independently distributed (*iid*) with zero mean and unit variance. To illustrate the extension of the component model over the GARCH model, consider the multi-step forecast of the conditional variance in the GARCH(1,1) model in eq. (1). Defining the multi-step variance forecast conditional on Ω_{t-1} as $h_{t+k} \equiv Var(r_{t+k}|\Omega_{t-1})$, and given the assumption that the returns process r_t is covariance stationary (i.e., $\alpha + \beta < 1$), the GARCH(1,1) multi-step conditional variance forecast is given by $h_{t+k} = \omega[1 - (\alpha+\beta)^k]/(1-\alpha-\beta)$, which, as $k \to \infty$, converges on the unconditional variance $[\omega/(1-\alpha-\beta)] = Var\ (r_t) \equiv \sigma^2$, allowing the GARCH(1,1) model to be re-expressed as:

$$h_t = \sigma^2 + \alpha(\varepsilon_{t-1}^2 - \sigma^2) + \beta(h_{t-1} - \sigma^2) \quad (3)$$

where the terms in parentheses have expected values of zero, reflecting the constancy of volatility in the long run. In contrast, the component model extends the expression in eq. (3) to allow the possibility that long-

[12]For a closer examination of the intra-day periodicity in FTSE-100 index and other LIFFE futures, see ap Gwilym, McMillan, and Speight (1999).

run volatility is not constant. That is, by allowing a time-varying permanent component, q_t, and its lagged value, to replace the constant long-run volatility, σ^2, above, where the lagged forecasting error $(\varepsilon_{t-1}^2 - h_{t-1})$ serves as the driving force for the time-dependent movement of that permanent component:

$$h_t = q_t + \alpha(\varepsilon_{t-1}^2 - q_{t-1}) + \beta(h_{t-1} - q_{t-1}) \tag{4a}$$

$$q_t = \omega + \rho q_{t-1} + \phi(\varepsilon_{t-1}^2 - h_{t-1}) \tag{4b}$$

and the autoregressive root, $0 < (\alpha+\beta) < \rho \leq 1$, accommodates the often empirically relevant case of (near-) integration in volatility for ρ values of (close to) unity. Thus, conditional variance is decomposed into a permanent or long-run component, and a transitory or short-run component defined simply as $(h_t - q_t)$. The conditional variance is covariance stationary in this model if the permanent component and the transitory component are both covariance stationary, as satisfied by $\rho < 1$ and $(\alpha+\beta) < 1$, respectively, those values also quantifying the persistence of shocks to these component processes. For $1 > \rho > (\alpha+\beta)$, the transitory component decays more quickly than the permanent component such that the latter dominates forecasts of the conditional variance as the forecasting horizon is extended, and eventually converges on a constant so long as the permanent component is stationary: $h_{t+k} = q_{t+k} = \omega/(1-\rho)$ as $k \rightarrow \infty$, for $0 < \rho < 1$.[13] Note that the component model reduces to the GARCH(1,1) model if either $\alpha=\beta=0$, or $\rho=\phi=0$. Thus, the GARCH model only is capable of describing, at most, one element of the more general conditional variance component specification, and only represents the permanent component under the specific conditions $\alpha=\beta=0$, $\rho=1$.[14]

[13]Further, by substitution using eq. (4) following Engle and Lee, note that the component model may be expressed alternatively as either a GARCH(2,2) model, or a GARCH(1,1) model with time-varying intercept, the latter being:

$$h_t = [\omega + (\rho - \alpha - \beta)q_{t-1}] + (\alpha + \phi)\varepsilon_{t-1}^2 + (\beta - \phi)h_{t-1}$$

such that for $\omega > 0$, $\alpha > 0$, $\beta > \phi > 0$, $1 > \rho > (\alpha + \beta) > 0$, the conditional variance h_t is ensured to be non-negative so long as q_t is non-negative. Since substitution also allows the permanent component to be expressed as a GARCH(2,2) process, the results of Nelson and Cao (1992) may be used to verify constraints for the non-negativity of q_t, which in turn can be shown also to be satisfied under the restrictions already given. For further details of stationarity and non-negativity conditions, see Engle and Lee (1993).

[14]These conditions, and the latter especially, are self-evident in the comparison of the model in eq. (4), and its GARCH(2,2) representation obtained by substitution using eq. (4) under the conditions given in the text and the property that $h_{t+k} - q_{t+k}$ as $k \rightarrow 0$ for $\rho = 1$. For further details, see Engle and Lee (1993).

TABLE II

Volatility Model Estimates

		Coefficient Estimates				
Frequency	*Model*	ω	ρ	ϕ	α	β
Five-Minutes	GARCH	0.0271*	–	–	0.0509*	0.9366*
		(0.0026)			(0.0022)	(0.0030)
	Component	1.9395*	0.9987*	0.0057*	0.0958*	0.8376*
		(0.1682)	(0.0004)	(0.0013)	(0.0064)	(0.0093)
Fifteen-Minutes	GARCH	0.1129*	–	–	0.0564*	0.9247*
		(0.0138)			(0.0041)	(0.0052)
	Component	5.3336*	0.9982*	0.0075*	0.0967*	0.7227*
		(0.5218)	(0.0006)	(0.0014)	(0.0074)	(0.0267)
One-Hour	GARCH	0.1812*	–	–	0.0327*	0.9575*
		(0.0470)			(0.0050)	(0.0065)
	Component	17.4195*	0.9955*	0.0151*	0.0504*	0.5281*
		(2.0736)	(0.0017)	(0.0031)	0.0191	(0.2103)
Half-Day	GARCH	0.6307	–	–	0.0236*	0.9692*
		(0.4188)			(0.0065)	(0.0096)
	Component	83.9615*	0.9928*	0.0258*	−0.0255	0.1907
		(13.8626)	(0.0048)	(0.0068)	(0.0226)	(1.0452)

For model specifications see text, eqs. (2), (4), and (5); all standard errors, in parentheses, adjusted by the method of Bollerslev and Woolridge (1992).
* indicates asymptotic coefficient significance at the 5% level.

CONDITIONAL VOLATILITY RESULTS

Coefficient estimates for both GARCH and component models obtained by maximum likelihood, together with Bollerslev and Wooldridge (1992) non-normality robust standard errors, are reported in Table II for each of the four data frequencies noted above.[15] Residual diagnostics for both models are reported in Table III, and include moment measures, Jarque–Bera tests for departures from normality, Engle ARCH-LM tests, and BDS tests of the null that the series in question are *iid* against an unspecified alternative.[16]

[15]Estimation is by maximum likelihood, using the Marquardt (1963) iterative nonlinear algorithm. Conditional means, m_t, and conditional variances are estimated jointly, the former having autoregressive form. Details of conditional mean estimates are suppressed here in the interests of brevity, but are available on request.

[16]More formally, the Brock, Dechert, and Scheinkman (BDS; 1987) test statistic is based upon a measure of spatial correlation in m-dimensional space known as the 'correlation integral' (Grassberger and Procaccia, 1983) and defined as:

$$W_{m,T}(d) = T^{0.5}[C_{m,T}(d) - C_{1,T}(d)^{m}]/\sigma_{m,T}(d)$$

where σ is the sample standard deviation of the data, and $C_{m,T}(k)$ is the sample correlation integral given 'embedding dimension', m, and distance, d. In applications to *iid* series, the BDS statistic is

TABLE III

Residual Diagnostics

											BDS(2,d)		BDS(3,d)		BDS(4,d)	
Frequency	*Model*	*Mean*	*S.D.*	*Sk.*	*Ku.*	*JB*	A_1	A_4	A_8	A_{12}	$d = \sigma$	$d = \sigma/2$	$d = \sigma$	$d = \sigma/2$	$d = \sigma$	$d = \sigma/2$
Five	GARCH	0.001	1.00	0.24	12.08	264,893.40*	57.55*	57.82*	60.15*	64.73*	11.90*	15.78*	13.60*	18.81*	13.97*	21.03*
Minutes	Component	−0.001	1.00	0.27	13.56	358,539.70*	5.73*	10.60*	13.05	15.34	6.79*	11.52*	7.33*	13.42*	7.07*	14.82*
Fifteen	GARCH	−0.004	1.00	0.10	6.58	13,601.22*	26.54*	27.87*	33.01*	40.24*	4.19*	4.71*	4.31*	5.37*	5.83*	6.57*
Minutes	Component	−0.004	1.00	0.11	6.96	16,609.02*	0.61	1.08	4.44	17.42	1.88	2.51*	1.53	2.81*	1.86	3.72*
One	GARCH	−0.010	1.00	−0.02	5.21	1,542.87*	11.54*	14.43*	15.11	18.17	1.54	1.73	1.35	1.44	1.86	1.79
Hour	Component	−0.010	1.00	0.01	5.01	1,287.14*	0.34	0.64	1.04	4.04	−0.97	−0.58	−1.51	−1.21	−1.09	−0.96
Half-	GARCH	−0.010	1.00	0.13	3.82	52.13*	0.10	2.10	5.39	8.31	−0.40	−0.56	−0.70	−0.70	−0.61	−1.14
Day	Component	−0.010	1.00	0.12	3.82	51.73*	0.01	2.09	5.40	8.40	−0.08	−0.23	−0.34	−0.32	−0.26	−0.76

'Sk.' and 'Ku.' denote measures of the second and third moments of skewness and kurtosis, on the basis of which the Jarque-Bera test for normality is calculated; 'JB' $\sim \chi^2_2$; 'A_i' denotes the i-th order Engle (1982) ARCH test, distributed as χ^2_i; 'BDS' denotes the Brock et al. (1987) test for departures from iid defined over (m, d) where m denotes embedding dimension and d distance (determined with reference to the sample residual standard deviation, σ), asymptotically distributed as $N(0,1)$.

* indicates asymptotic test significance at the 5% level (with the exception of moment measures).

The preliminary GARCH(1,1) estimates confirm the presence of persistence in volatility at each of the four intra-day frequencies considered, at 0.988, 0.981, 0.990, and 0.993 at the five-minute, fifteen-minute, one-hour, and half-day frequencies, respectively, with corresponding half-lives of 4.6 hours, 9 hours, 70 hours, and 360 hours, again in order of declining frequency. These results are in keeping with previous research reporting that volatility shock half-lives calculated from GARCH model estimates are extremely low in high-frequency intra-day financial data, and particularly so here at the five-minute and fifteen-minute frequencies.[17] Residual diagnostics for these models, and for the estimated component models discussed below, indicate that the degree of non-normality decreases as the frequency of the data lowers, but remains statistically significant at all frequencies for both models, validating our use of robust standard errors throughout. For the GARCH models specifically, remaining diagnostics indicate the presence of residual ARCH structure at all frequencies other than the half-day, these results being broadly confirmed by BDS statistics.[18] As noted above, the component model implies the presence of a higher order GARCH structure, consistent with these residual diagnostics for GARCH(1,1) models, and therefore we proceed to component model estimation.[19]

asymptotically distributed as a standard normal, $W \sim N(0,1)$, and usefully for our purposes it has been demonstrated that the test remains valid in application to residuals from regression models [BDS, 1987; Brock, Hsieh, & LeBaron (BHL), 1991; Hsieh, 1989]. As a test for departures from *iid* in the residuals of linear and non-linear systems, the BDS statistic therefore provides a portmanteau test for adequacy of fit against unspecified residual structure. We follow the recommendations of BHL in reporting BDS diagnostics for embedding dimensions of $m = 2,3,4$ and, for each m, $d = (\sigma, \sigma/2)$. However, the asymptotic distribution does not approximate well the finite distribution following GARCH estimation. While a desirable course of action would be to obtain the empirical distribution of the BDS statistic under the null hypothesis of *iid*, and thus to compute critical values by bootstrapping the data, the large number of data points involved rendered this process overburdensome. Therefore, in determining significance, we also are guided by Monte Carlo simulation results for the GARCH(1,1) model reported by Brock, Hsieh, and LeBaron (BHL; 1991). For further details of the BDS statistic and its properties, see Brock et al. (1991) and Brock and Potter (1993). Also see n.18 and n.24 below.

[17]See the Introduction, and the references cited in n.1 and n.2.

[18]Note that while the significance of the BDS statistics reported in Table III is appraised relative to the asymptotic $N(0,1)$ value of ± 1.96, as mentioned in n.16, the asymptotic distribution does not approximate well the BDS statistic applied to the standardized residuals of GARCH models. Therefore, we also use the simulation quantiles reported by BHL as guidance to the actual sizes of the BDS statistics applied to GARCH model residuals. The significance of test statistics indicated in Table 3 and the inference made in the text continue to hold on these alternative critical values, with the addition of significance in the BDS(2, σ) statistic at the one-hour frequency.

[19]The GARCH(2,2) models implied by the component model, as noted under The Component Model, were also estimated and found to possesss significant coefficient estimates throughout, with the exception of the half-day frequency, consistent with the component model estimates reported below. These results are not reported here in the interest of brevity, and given some uncertainty as to the stationarity of those models in the absence of parameter conditions for such general GARCH models in the presence of negative coefficients (in contrast to known conditions for lower-order GARCH

Component model results are reported in the second row of results for each frequency in Table II. At the higher frequencies of five minutes, fifteen minutes, and one hour, these results are supportive of the permanent–transitory component decomposition, with all parameters statistically significant and satisfying the inequality constraints noted in The Component Model section. The persistence of shocks to the permanent component is very high, in excess of 0.99 at all frequencies, while the persistence of shocks to the transitory components declines with decreasing frequency, from 0.933 in five-minute returns to 0.819 in fifteen-minute returns, 0.579 in one-hour returns, and 0.165 in half-day returns. For the three higher frequencies, these results imply transitory component half-lives of 0.84 hours, 0.87 hours, and 1.27 hours at the five-minute, fifteen-minute, and one-hour frequencies, respectively, indicating full decay of a shock to the transitory component within a few hours. Indeed, transitory component parameter estimates are statistically insignificant at the half-day frequency. Corresponding permanent component half-lives in order of decreasing frequency are 44 hours, 96 hours, and 154 hours, with a permanent component half-life of 360 hours (or approximately 48 days) at the half-day frequency.[20] Thus, while the effect of a shock to the permanent component conditions volatility over several days or weeks, a shock to transitory volatility is dissipated within the half-day. Moreover, in comparison with these half-lives calculated using permanent component persistence measures, the GARCH models reported above understate volatility shock half-lives by a factor of over ten in five-minute and fifteen-minute returns, and by more than two at the hourly frequency. On the basis of *t*-statistics employing robust standard errors, the hypothesis of integration in variance ($\rho = 1$) can be rejected for all frequencies other than the half-day at the 5% significance level, with test values in order of decreasing frequency of -3.25, -3.00, -2.65, and -1.50. This result is confirmed by Wald tests of the null hypothesis that persistence is integrated for both the GARCH and component models; only at the half-day frequency can that null not be rejected for both models, test statistics being 1.94 and 2.12, respectively.[21] Indeed, persistence measures for the component and GARCH models are identical to four decimal places in this case. Therefore, these results imply complete dissipation of the transitory component by the half-day frequency and, from

models, and for strictly positive higher-order models: see for example, Bougerol & Picard, 1992; Drost & Nijman, 1993; Nelson & Cao, 1992).

[20]All half-lives calculated using $[\ln(0.5)/\ln(\eta)]$, where η is the measure of persistence, given by $(\alpha + \beta)$ in the GARCH model and the transitory component, and ρ in the permanent component.

[21]The χ^2_1 critical value is 3.84. The null of variance integration is rejected strongly at all higher frequencies: test statistics available on request.

inspection of eq. (6) under ($\alpha = \beta = 0$, $\rho = 1$), reversion to an integrated-GARCH model at that frequency.[22,23]

Turning to residual diagnostics for the component models in Table III, while BDS statistics indicate some remaining residual structure at both the five- and fifteen-minute returns intervals, residual ARCH effects are present at the 5% significance level only in five-minute returns. These test statistics are reduced substantially at the lower frequencies, and suggest that the component model is able to capture all structure within FTSE-100 futures returns observed at a frequency of one hour or lower.[24] Thus, while alternative variance specifications to those employed here may prove fruitful in further modeling of five-minute returns, and there may be further nonlinearity in fifteen-minute returns, the component model provides a satisfactory model of volatility in hourly returns, the GARCH model only adequately capturing returns dependency at the half-day frequency following the dissipation of transitory volatility.

VOLATILITY–VOLUME INTERACTION

The results of the preceding section lend broad support to the existence of a component structure in intra-day FTSE-100 futures volatility. As noted in the Introduction, the rationalization for the existence of components in intra-day conditional volatility offered by Andersen and Bollerslev (1997b) lies in a heterogeneous information flow generalization of the mixture-of-distributions hypothesis (MDH) advanced by Clark (1973).[25] In order to appraise the validity of this interpretation, we follow

[22]Reduction of the component model to the GARCH(1,1) form is confirmed by a likelihood ratio test of the null hypothesis parameter restrictions, which cannot be rejected on the basis of a test statistic value of 0.23 relative to the critical value of 5.99. Preference for the GARCH model at the half-day frequency also is indicated by BIC values. At all other frequencies, BIC values and likelihood ratio tests strongly support the component model over the GARCH model. Further details available on request.

[23]Given the capacity of the component model to capture change in the unconditional variance that is manifest as time variation in the conditional variance intercept, as expressed in eq. (6), the failure of the permanent–transitory decomposition to hold and reversion to an integrated-GARCH(1,1) model at the half-day frequency also may be taken as supportive of the view that long-memory characteristics are an inherent property of the returns-generating process at that frequency, rather than the result of externally induced structural change in the volatility process as suggested by, for example, Lamoureux and Lastrapes (1990).

[24]Further to n.16, it should be noted that the finite distribution of the BDS statistic following component model estimation is unknown, and asymptotic critical values are employed here in the absence of bootstrap values. However, we do not believe that the use of asymptotic values unduly influences our results, particularly in view of the very large number of observations in our samples and the extremely low BDS values obtained.

[25]Under the MDH proposed by Clark (1973), asset prices obey a subordinated stochastic process, with prices evolving at different rates according to the amount of information becoming available to the market, and at faster rates given the arrival of unexpected information. Thus, returns are antic-

and extend the methodology of previous researchers (e.g., Lamoureux & Lastrapes, 1990) in testing for the significance of information arrival, as proxied by (unexpected) trading volume, V_t^u, in causing volatility clustering and associated effects.[26] Following the argument of Lamoureux and Lastrapes, under the information flow hypothesis volume should be significant and such as to substantially reduce volatility persistence. Thus, we consider the inclusion of unexpected volume in the component model to yield, after the simplification of eq. (4):

$$h_t = \omega + \rho q_{t-1} + \phi(\varepsilon_{t-1}^2 - h_{t-1}) + \alpha(\varepsilon_{t-1}^2 - q_{t-1}) + \beta(h_{t-1} - q_{t-1}) + \delta V_t^u \tag{5}$$

with attention to the significance of, and the resulting impact on, the magnitude and significance of the permanent and transitory component persistence parameters, ρ and $(\alpha + \beta)$, respectively.

The results of augmented model estimation for those frequencies at which a component structure already has been indicated, namely hourly and higher, are presented in Table IV. Unexpected volume, as proxied by the residuals from appropriately specified autoregressive conditional mean models, is statistically significant throughout, with a consistently positive effect on volatility that is increasing in magnitude with decreasing frequency.[27] While permanent component volatility persistence remains statistically significant at all frequencies, its magnitude is reduced. This reduction is most pronounced at the hourly frequency, where persistence is approximately halved, to 0.50, with more muted reductions to 0.89 and 0.96 at the fifteen- and five-minute frequencies, respectively. A further effect of the introduction of unexpected volume is that transitory com-

ipated to follow a mixture of distributions in which the rate of (serially correlated) information arrival to the market is the mixing variable, the stochastic properties of which the GARCH model reflects (e.g., Diebold, 1986). For a particularly lucid account of the MDH and related issues, see Omran and McKenzie (1996). Generally supportive evidence for the MDH has previously been reported at the daily level for individual stocks (see for example Lamoureux & Lastrapes for U.S. data, and Omran & McKenzie, 1996, for U.K. data). However, several recent studies examining futures data have reported less convincing evidence. Najand and Yung (1991) examining Treasury-bond futures, Foster (1995) examining crude oil futures, and Jacobs and Onochie (1998) examining several international futures markets all report evidence that, while volume is positive and significant in a GARCH equation, its inclusion has little effect on the magnitude and significance of the GARCH parameters themselves.

[26]The use of unexpected volume follows the observation by Bollerslev, Chou, and Kroner (1992), amongst others, that a problem with the use of absolute contemporaneous volume lies in potential simultaneity bias arising from the possible endogeneity of volume due to potential contemporaneous correlation between volume and volatility. Results based on absolute volume therefore are not reported here, but are available on request.

[27]Autoregressive conditional mean models and residual diagnostics for volume are suppressed here in the interests of brevity, but further details are available from the authors on request.

TABLE IV

Augmented Component Model Estimates

	Coefficient Estimates					
Frequency	ω	ρ	ϕ	α	β	δ
Five Minutes	2.1338*	0.9644*	0.0602*	0.0610*	0.1089	0.1708*
	(0.0909)	(0.0025)	(0.0038)	(0.0091)	(0.1322)	(0.0107)
Fifteen minutes	5.5214*	0.8914*	0.2302	−0.1645	1.0250*	0.3463*
	(0.1838)	(0.0128)	(0.4014)	(0.4024)	(0.4702)	(0.0329)
One Hour	16.7570*	0.5039*	1.3910	−1.3343	1.7290	0.9730*
	(0.4849)	(0.0305)	(4.4643)	(4.4617)	(4.6081)	(0.0646)

For model specification see text, equation (5); all standard errors, in parentheses, adjusted by the method of Bollerslev and Wooldridge (1992).
* indicates asymptotic coefficient significance at the 5% level.

ponent parameters are less-well identified and behaved. Nevertheless, the conditions for covariance stationarity noted under The Component Model continue to be satisfied throughout by virtue of $0 < (\alpha + \beta) < \rho < 1$.[28] Moreover, transitory component parameters both are insignificant statistically at the hourly frequency, while one or the other of the transitory component parameters are insignificant at the higher frequencies.

These results lend support to a heterogeneous information flow interpretation of the component structure in volatility at the hourly level, but are less confirmatory at the higher frequencies of five and fifteen minutes. One explanation for this lack of support at the higher frequencies may be that unexpected volume over those fine time intervals carries little relevant information content, and that possibility should not be overlooked. However, an alternative and appealing explanation in the context of the futures market investigated here, as described under Institutional Considerations, is that the component volatility structure at the higher frequencies reflects the activities of local traders with short intra-day trading horizons, such as scalpers, acting as market makers by providing liquidity services in return for small gains from the bid-ask spread, particularly in circumstances where order imbalances arise. The volatility induced by such activities therefore is not based on information flows, in keeping with the higher frequency volatility–volume results reported above.[29]

[28]Further, while the sufficient conditions for conditional variance non-negativity given in n. 13 above are violated by instances of transitory component parameter negativity at the fifteen-minute and hourly frequencies, $\alpha < 0$, inspection of the estimated models confirms that the conditional variance is strictly positive in estimation throughout.

[29]Another possible explanation is that the increasing explanatory power of volume with increasing

SUMMARY AND CONCLUSIONS

Andersen and Bollerslev (1997b) have recently suggested that the arrival of heterogeneous information to a financial market imparts both short-run and long-run volatility, such that the short-run effects dominate over high-frequency intra-day intervals while the impact of highly persistent processes dominate over longer horizons. Müller et al. (1997) have alternatively suggested that such structure results from heterogeneous market agents characterized by different trading horizons. However, the standard GARCH model implicitly assumes homogeneity of the price discovery process and is unable to capture these effects. In order to appraise these contentions, we empirically test for explicit volatility decomposition using the variance component model of Engle and Lee (1993) in application to intra-day UK FTSE-100 futures data. The existence of a component structure to volatility is supported, in the sense of the coexistence of a transitory component to volatility at intra-day frequencies of one hour and higher that decays within the half-day, and permanent volatility that decays over a much longer horizon. The component model is also able to capture all structure within the data on the basis of residual tests at frequencies of one hour and lower.

Inspection of an augmented version of the component model that includes unexpected volume as a proxy for news arrival suggests that the heterogeneous information flow hypothesis of Andersen and Bollerslev provides a compelling explanation of the component structure in hourly volatility, and the transitory component in particular. However, at frequencies higher than hourly, intra-day heterogeneous information flows appear unable to provide an adequate explanation for the existence of components in volatility, the likely alternative explanation for which may be the existence of heterogeneous market agents, as suggested by Müller et al., and local liquidity traders with very short trading horizons in particular. At the half-day frequency, the permanent–transitory decomposition no longer is supported, transitory component parameters being insignificant while the permanent component is sufficiently persistent that the null of integration in variance cannot be rejected. Only at the half-day frequency does the component model therefore reduce to a GARCH model; GARCH models at higher frequencies imply substantial under-estimates of volatility shock half-lives, with obvious consequences for vol-

temporal aggregation reflects the gradual or sequential, albeit quick, dissemination of private information and associated trading across the market. However, this explanation is perhaps tenuous in the context of the open outcry system as described under Institutional Considerations and Market Structure and Data.

atility forecasts and derivative pricing methods based on such estimates. The extension of the analysis conducted here to high-frequency returns data for alternative assets, and to the issue of low-frequency volatility forecasting based on the extraction of the latent permanent or long-memory component in high-frequency data, therefore provide interesting avenues for further investigation.

BIBLIOGRAPHY

Andersen T. G., & Bollerslev, T. (1997a). Intra-day periodicity and volatility persistence in financial markets. Journal of Empirical Finance, 4, 115–158.

Andersen T. G., & Bollerslev, T. (1997b). Heterogeneous information arrivals and return volatility dynamics: Uncovering the long-run in high frequency returns. Journal of Finance, 52, 975–1005.

ap Gwilym, O., Clare, A. D., & Thomas, S. H. (1998). Extreme price clustering in the London equity index futures and options markets, Journal of Banking and Finance, 22, 1193–1206.

ap Gwilym, O., McMillan, D. G., & Speight, A. E. H. (1999). The intraday relationship between volume and volatility in LIFFE futures markets. Applied Financial Economics, 9, 593–604

ap Gwilym, O., & C. M. S. Sutcliffe (1999). High frequency financial market data, Risk Books: London.

Beveridge, S., & Nelson, C. R. (1981). A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the business cycle., Journal of Monetary Economics, 7, 151–174.

Bollerslev, T. (1986). Generalised autoregressive heteroscedasticity. Journal of Econometrics, 31, 307–327.

Bollerslev, T., Chou, R. Y., & Kroner, K. F. (1992). ARCH modelling in finance. Journal of Econometrics, 52, 5–59.

Bollerslev, T., & Wooldridge, J. M. (1992). Quasi maximum likelihood estimation and inference in dynamic models with time varying covariances. Econometric Reviews, 11, 143–172.

Bougerol, P., & Picard, N. (1992). Stationarity of GARCH processes and of some nonnegative time series. Journal of Econometrics, 52, 115–128.

Brock, W., Dechert, W. D., & Scheinkman, J. (1987). A test for independence based on the correlation dimension. University of Wisconsin, Economics Working Paper No. SSRI-8702.

Brock, W. A., Hsieh, D. A., & LeBaron, B. (1991). Nonlinear dynamics, chaos, and instability: Statistical theory and economic evidence. Massachusetts: MIT Press.

Brock, W. A., & Potter, S. A. (1993). Non-linear time series and macroeconomics. Handbook of Statistics, 11, 195–229.

Clark, P. K. (1973) A subordinate stochastic process model with finite variance for speculative prices. Econometrica, 41, 135–155.

Dacorogna, M. M., Müller, U. A., Olsen, R. B., & Pictet, O. V. (1997). Modelling short-term volatility with GARCH and HARCH models. Olsen and Associates Research Group, Zurich, mimeo.

den Hertog, R. G. J. (1994). Pricing of permanent and transitory volatility for US stock returns: A composite GARCH model. Economics Letters, 44, 421–426.

Diebold, F. X. (1986). Modelling the persistence of conditional variances: A comment. Econometric Reviews, 5, 51–56.

Drost, F. C., & Nijman, T. (1993). Temporal aggregation of GARCH processes. Econometrica, 61, 909–927.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. Econometrica, 50, 987–1007.

Engle, R. F., & Lee, G. G. J. (1993). A permanent and transitory component model of stock return volatility. Department of Economics, UCSD, Discussion Paper No: 92-44R.

Foster, A. J. (1995). Volume-volatility relationships for crude oil futures markets. Journal of Futures Markets, 15, 929–951.

Ghose, D., & Kroner, K. F. (1996), Components of volatility in foreign exchange markets: An empirical analysis of high frequency data, Department of economics. University of Arizona, mimeo.

Goodhart, C. A. E., Hall, S. G., Henry, S. G. B., & Pesaran, B. (1993) News effects in a high-frequency model of the sterling-dollar exchange rate. Journal of Applied Econometrics, 8, 1–13.

Goodhart, C. A. E., & O'Hara, M. (1997), High frequency data in financial markets: Issues and applications. Journal of Empirical Finance, 4, 73–114.

Grassberger, P., & Procaccia, I. (1983). Measuring the strangeness of strange attractors. Physica D, 9, 189–208.

Guillaume, D. M., Pictet, O. V., & Dacorogna, M. M. (1995). On the intra-daily performance of GARCH processes. Olsen and Associates Research Group, Zurich, mimeo.

Hsieh, D. A. (1989) Modelling heteroscedasticity in daily foreign exchange rates. Journal of Business and Economic Statistics, 7, 307–317.

Jacobs, M. and Onochie, J. (1998). A bivariate generalised autoregressive conditional heteroscedsticity-in-mean study and the relationship between return variability and trading volume in international futures markets. Journal of Futures Markets, 18, 379–397.

Jones, C., Lamont, O., & Lumsdaine, R. (1998). Macroeconomic news and bond market volatility. Journal of Financial Economics, 47, 315–337.

Lamoureux, C. G., & Lastrapes, W. D. (1990). Persistence in variance, structural change and the GARCH model. Journal of Business and Economic Statistics, 8, 225–234.

Marquardt, D. W. (1963). An algorithm for least squares estimation of nonlinear parameters. Journal of Society of Industrial Applied Mathematics, 11, 431–441.

Müller, U. A., Dacorogna, M. M., Davé, R. D., Olsen, R. B., Pictet, O. V., & von Weizsäcker, J. E. (1997). Volatilities of different time resolutions - analyzing the dynamics of market components, Journal of Empirical Finance, 4, 213–239.

Najand, M. and Yung, K. (1991). A GARCH examination of the relationship between volume and price variability in futures markets. Journal of Futures Markets, 11, 613–621.

Nelson, D. B. (1990). ARCH models as diffusion approximations. Journal of Econometrics, 45, 7–38.

Nelson, D. B. (1992). Fitting and forecasting with misspecified ARCH models I: Getting the right variance with the wrong model. Journal of Econometrics, 52, 61–90.

Nelson, D. B., & Cao, C. Q. (1992). Inequality constraints in the univariate GARCH model. Journal of Business and Economic Statistics, 10, 229–235.

Omran, M. F. and McKenzie, E. (1996). Heteroscedasticity in stock market prices revisited: Unexpected volume versus GARCH effects. University of Sterling, Department of Accounting and Finance, unpublished manuscript.

OPalm, F. C., & Urbain, J-P. (1995) Common trends and transitory components of stock price volatility. University of Limburg, Department of Economics, mimeo.

Sentana, E. (1995). Quadratic ARCH models. Review of Economic Studies, 62, 639–61.

Tse, Y. (1999). Market microstructure of FTSE-100 index futures: An intraday empirical analysis. Journal of Futures Markets, 19, 31–58.