

NONLINEAR DYNAMICS IN HIGH FREQUENCY INTRA-DAY FINANCIAL DATA: EVIDENCE FOR THE UK LONG GILT FUTURES MARKET

David G McMillan¹ and Alan E H Speight^{2,*}

July 1999

Abstract

Tests against the null of linearity indicate smooth transition autoregressive nonlinearities in the conditional mean of intra-day UK long gilt futures returns at the five and fifteen minute frequencies. The higher frequency model entails a first-order autoregressive process with switching intercept. The lower frequency model is first-order autoregressive for returns near zero, but a near random-walk for large returns, consistent with the rapid extraction of profitable opportunities in excess of friction transaction cost boundaries. These nonlinearities are robust to the presence of asymmetric and component structures in conditional variance, but suggest that the potential for predictable regularities are confined to small price movements over fine time intervals.

Keywords: Futures Contract, High Frequency, Smooth
Transition Threshold, Conditional Volatility.
JEL Classification: G12, G13, G14, C22.

¹ Department of Economics, University of St Andrews, Fife, KY16 9AL, UK

² Department of Economics, University of Wales, Swansea, SA2 8PP, UK.

* Corresponding Author: tel: (+44) 1792-205678; fax: (+44) 1792-295872;
e-mail: a.speight@swan.ac.uk

CRACKED TRADING SOFTWARE

70+ DVD's FOR SALE & EXCHANGE

www.traders-software.com

www.forex-warez.com

www.trading-software-collection.com

www.tradestation-download-free.com

Contacts

andreybbrv@gmail.com

andreybbrv@yandex.ru

Skype: andreybbrv

1. Introduction

Over the past decade and a half, the genre of models of generalised autoregressive conditional heteroscedasticity (GARCH: Engle, 1982; Bollerslev, 1986) have provided the dominant means for modelling nonlinear dependence in financial data, largely due to their empirical success in capturing the time-varying conditional volatility characteristic of the returns distributions of many financial assets.¹ A popular and theoretically appealing explanation for the presence of ARCH effects in asset returns, embodied in the mixture of distributions hypothesis, is that returns evolve as a subordinate stochastic process such that the distribution of returns follows a mixture of normals with changing variance, the rate of new information arrival providing the stochastic mixing variable. Thereby, asset prices evolve at different rates during identical intervals of time according to the flow of new information, and the distribution of returns, when measured over fixed time intervals, appears kurtotic. As suggested by Diebold (1986), the empirical success of ARCH-type models may then lie in their ability to capture serially correlation in the time-series properties of the mixing variable, the flow of information.² In extension of this approach, the recent examination of high-frequency intra-day data has prompted several researchers to suggest that volatility may more accurately be characterised by heterogeneous components reflecting heterogeneous information flows (Andersen and Bollerslev, 1997a), or perhaps the actions of heterogeneous market traders (Müller et. al., 1997).

The analysis of high frequency intra-day data also raises a further consideration. Namely, the potential for the conditional mean process for high-frequency returns data to be more accurately described by a non-linear process.³ Whilst there has been extensive investigation of non-linearity in conditional mean in many macroeconomic time series, mostly associated with increasing recognition of the potentially asymmetric nature of the business cycle, relatively little research has been conducted

seeking to identify, model or explain stochastic non-linear conditional mean structure in financial market data.⁴ One reason for this is the lack of substantive linear structure in daily or lower frequency financial data, market returns at such frequencies typically approximating random walk processes, since linear structure is generally a prerequisite for the conduct of formal statistical tests against the null hypothesis of linearity.⁵ Moreover, a well defined non-linear conditional mean structure for security returns over a period of a day, for example, would potentially allow informed market participants to secure systematic profits.⁶ In contrast with such lower frequency data, intra-day data affords the linear structure which must precede consideration of non-linearity whilst not necessarily being inconsistent with market efficiency given the short time intervals over which such processes are found to extend. Particularly since there must exist some time interval at sufficiently high frequency over which market prices are brought to equilibrium following disturbance due to new information, especially in the context of the gradual dissemination of information, noise trading, or transaction costs. These rationales for the presence of linear structure, and the latter in particular, also provide rationales for the presence of non-linear structure. Especially that of threshold form, where the parameters of a linear model are permitted to change through time due to a switching rule defined over past price movements relative to some threshold value.

In the investigation of intra-day long gilt futures returns data reported here, we therefore consider both linear and nonlinear conditional mean structures. For the latter, we adopt the smooth transition autoregressive (STAR) model (Chan and Tong, 1986; Teräsvirta and Anderson, 1992; Granger and Teräsvirta, 1993; Teräsvirta, 1994) which allows for differing market dynamics according to the magnitude of returns, motivated by considerations of market frictions, such as noise trading and transactions costs, which create a band of price movements around the equilibrium price with

arbitrageurs only actively trading when deviations from equilibrium become sufficiently large. Following confirmatory preliminary tests for the presence of threshold non-linearities, STAR conditional mean estimates are reported. The robustness of that nonlinear mean structure to the presence of ARCH effects is examined through joint estimation under maximum likelihood using one of two extensions of the basic GARCH framework which permit conditional variance asymmetry or heterogeneity respectively. The former is provided by the exponential-GARCH (EGARCH) model of Nelson (1991), which has a correspondence with the informational flow hypothesis discussed above, whilst the latter is provided by the Engle and Lee (1993) component-GARCH (CGARCH) model, which permits the decomposition of conditional volatility into long-run and short-run elements, in keeping with recently advanced notions of volatility heterogeneity in intra-day financial data.

The remainder of the paper is organised as follows. In the following section we outline the empirical models to be estimated and further discuss their properties and relationship to issues of market dynamics. Section 3 describes the data and institutional setting from which it is drawn, provides nonparametric kernel density estimates of the data distributions and reports the results of preliminary tests for nonlinearity in conditional mean. Section 4 discusses issues of model specification and evaluation, and reports conditional mean and variance estimates. Section 5 provides a summary of our findings and their interpretation, and concludes by noting their implications for considerations of market efficiency and the activities of market agents.

2. Models

2.1. Market Frictions, Threshold Nonlinearities and the ESTAR Model

An issue which has received much attention in the empirical finance literature of late, and which offers

an appealing explanation for asymmetries in market returns, is related to the phenomenon of ‘noise-trading’. The rationale generally offered for the existence of noise trading is that it allows privately informed traders to profitably exploit their informational advantage, without which market efficiency would not be assured (eg. Kyle, 1985). That rationale does not, however, explain the reasons for noise trading, on which there are differing views. Thus, noise trading may be regarded as resulting either from rational agents trading for liquidity and hedging purposes, consistent with a fully-rational efficient-markets perspective (Diamond and Verrechia, 1981; Ausubel, 1990a,b; Biasis and Hillion, 1994; Dow, 1995; Dow and Gorton, 1994, 1996), or as the actions of irrational (or not-fully rational) agents trading on beliefs and sentiments that are not justified by news concerning underlying fundamentals (Black, 1986; Schleifer and Summers, 1990; De Long et. al., 1990). An interesting alternative interpretation recently offered by Dow and Gorton (1997) suggests that delegated portfolio managers may engage in noise trading in order to appease clients or managers who are unable to distinguish purposeful inaction from non-purposeful inaction, as a result of which the amount of noise trading can be large compared to the amount of hedging volume and Pareto improving.

Whatever the underlying reasons for noise trading, its existence means that profitable opportunities will arise for privately informed and arbitrage traders. In early recognition of the potential nonlinear consequences of such trading activities, Cootner (1962) notes that the activities of noise traders will cause prices to hit upper or lower ‘reflecting barriers’ around equilibrium, and thus trigger arbitrage activities by informed traders which push prices back to equilibrium. The existence and position of such barriers will likely depend on the existence and size of market frictions such as transactions costs, giving rise to a band of price movements around the equilibrium price with fully rational traders only actively trading when deviations from equilibrium are sufficiently large to make

arbitrage trade profitable (He and Modest, 1995). Such opportunities are unlikely to be long-lived, existing only for as long as reassessment of underlying fundamentals in the light of news may warrant. However, while the actions of individual traders may be represented by a simple threshold model which imposes an abrupt switch in behaviour, only if all traders act simultaneously will this also be the observed market outcome. For a market of many traders acting at slightly different times a smooth transition model is therefore more appropriate than a ‘heaviside’ threshold model.

In previous examinations of intra-day asset price volatility, the differenced logarithm of the asset price has typically been modelled as a linear autoregressive (AR) process of order p , such that the asset return, $r_t = \log(P_t/P_{t-1})$, is described by:

$$(1) \quad r_t = \alpha_0 + \sum_{i=1}^p \alpha_i r_{t-i} + \varepsilon_t.$$

In order to investigate the possibility of threshold nonlinearities due to noise trading of the form described above, we consider the nonlinear STAR(p) generalisation of (1), expressed in general form (Teräsvirta and Anderson, 1992; Granger and Teräsvirta, 1993) as:

$$(2) \quad r_t = \alpha_0 + \sum_{i=1}^p \alpha_i r_{t-i} + F(r_{t-d})(\theta_0 + \sum_{i=1}^p \theta_i r_{t-i}) + \varepsilon_t$$

where $F(\cdot)$ denotes a transition function defined over a transition variable, provided here by the lagged return value, r_{t-d} , where d is the delay parameter. One interpretation of (2) is that r_t is described by the linear model in the second term on some occasions, and by that process with the addition of the potentially non-linear component in the compound third term on other occasions. Alternatively, the

components $F(r_{t-d})(\theta_0)$ and $F(r_{t-d})(\sum_{i=1}^p \theta_i r_{t-i})$ may be interpreted as rendering the intercepts

and autoregressive parameters of the model time-varying, and (2) therefore as belonging to the class of state-dependent models (Priestley, 1988). The transition function utilized here is of the exponential form:

$$(3) \quad F(r_{t-d}) = 1 - \exp(-\gamma (r_{t-d} - c)^2); \quad \gamma > 0,$$

where γ is a smoothing or transition parameter and c a threshold parameter, the combination of (2) and (3) yielding the exponential-STAR (ESTAR) model, whereby the parameters in (4) change symmetrically about c with r_{t-d} , such that as $(r_{t-d} - c) \rightarrow 0$, $F(r_{t-d}) \rightarrow 0$, and as $(r_{t-d} - c) \rightarrow \infty$, $F(r_{t-d}) \rightarrow 1$, whilst as either $\gamma \rightarrow 0$ or $\gamma \rightarrow \infty$ the model reduces to the linear AR form.⁷ Thus, the ESTAR model implies that the dynamic process for moderate returns will differ from that for larger returns, irrespective of sign.⁸

A practical problem frequently encountered in the estimation of STAR models concerns convergence and precision in estimates of the smoothing or transition parameter, γ . In particular, a large γ value results in a steep slope for the transition function at c , and a large number of observations in the neighbourhood of c are in principle required in order to estimate γ accurately. Consequently, with changes in γ having only a minor effect upon the transition function, the convergence of γ can prove problematic. A solution to this problem, suggested by Teräsvirta (1994) and adopted in estimation here, is to scale the smoothing parameter by the variance of the transition variable, $\sigma^2(r_{t-d})$, yielding the revised transition function:

$$(3') \quad F(r_{t,h}) = 1 - \exp[-\gamma (r_{t,h} - c)^2 / \sigma^2(r_{t,h})]$$

with appropriate adjustment required in interpretation of the resulting estimate of γ .

2.2. *The Exponential-GARCH (EGARCH) Model*

The initial model of conditional volatility examined is the exponential GARCH (EGARCH) model of Nelson (1991). The selection of the EGARCH model is motivated by its close relationship with the mixture of distributions hypothesis, originally due to Clark (1973), which views the variability of security prices as arising from differences in information arrival rates. The standard model assumes a fixed number of traders possessing different expectations and risk profiles, resulting in different reservation prices. Market clearing requires that the equilibrium price be the average of these reservation prices. Information arrival then causes traders to adjust their reservation prices, which in turn causes trade, which then changes the market price. Under the assumption that these price changes are normally distributed, it has been demonstrated that the aggregate of price changes and traded volume are jointly stochastic independent normals (Tauchen and Pitts, 1983; Gallant Hsieh and Tauchen, 1991). Where information events vary over time, price changes at the daily frequency, for example, are the sum over intra-day price changes. By appeal to the Central Limit Theorem, aggregated price changes are then described by mixtures of independent normals, where mixing depends on the rate of information arrival. In keeping with this framework, following Nelson (1990, 1991), the EGARCH model has lognormal conditional variance in continuous time, with the implication that as the sampling interval becomes finer in discrete time, the distribution of innovations approaches a conditionally normal mixture of distributions, thereby formally linking the EGARCH and mixture of distributions approaches.⁹

Notationally, let the asset return r_t have an expected return m_t (given by the conditional expectation of either the AR or ESTAR model defined above), and conditional variance given by $h_t^2 = \text{var}(r_t | \Psi_{t-1}) = E((r_t - m_t)^2 | \Psi_{t-1})$, where Ψ_{t-1} defines the set of all information available at time $t-1$. The first-order EGARCH model, which is also the appropriate empirical model order further below, is then given by:

$$(4) \quad \log(h_t^2) = \omega + \alpha (|\varepsilon_{t-1}| / |h_{t-1}|) + \zeta (\varepsilon_{t-1} / h_{t-1}) + \beta \log(h_{t-1}^2)$$

where the logarithmic form ensures conditional variance non-negativity without the necessity of constraining the coefficients of the model. Regarding the coefficients of (4), the parameter α captures the volatility clustering effect that is characteristic of ARCH processes, a positive value indicating that large (small) shocks tend to follow large (small) shocks of random sign, while the parameter β captures the degree of persistence in shocks to volatility, with half-life decay given by $\log(0.5) / \log(\beta)$. The potentially asymmetric effect of positive and negative shocks on conditional variance is captured by a non-zero value for the parameter ζ . For $\zeta \neq 0$, $\log(h_t^2)$ responds asymmetrically to $(\varepsilon_{t-1} / h_{t-1})$ in a piecewise linear manner: where that ratio is positive, $\log(h_t^2)$ is linear in $(\varepsilon_{t-1} / h_{t-1})$ with slope $(\zeta + \alpha)$, whilst for $(\varepsilon_{t-1} / h_{t-1}) < 0$, $\log(h_t^2)$ is linear in $(\varepsilon_{t-1} / h_{t-1})$ with slope $(\zeta - \alpha)$.

2.3. The Component-GARCH (CGARCH) Model

While the preceding EGARCH representation of volatility is based on assumed homogeneity of the price discovery process, it has recently been suggested that intra-day returns volatility may more

realistically comprise heterogeneous components (eg. Andersen and Bollerslev, 1997a). Such components may reflect differing market reactions to differing sources and types of news, or the differing reactions of market agents with heterogeneous positions and time horizons to the same items of news (Müller et. al., 1997). On either view, returns volatility will consequently be dominated by transient or short-run volatility over higher data frequencies and by more persistent or long-run volatility over lower data frequencies.

In order to examine the data for the possible presence of such components we implement the component-GARCH model of Engle and Lee (1993) which facilitates the decomposition of volatility into a long-run or (inter-day) component, and a short-run (intra-day) component.¹⁰ This (necessarily first-order) CGARCH model is given by the joint process:

$$(5a) \quad h_t^2 = \alpha_t + \alpha (\varepsilon_{t-1}^2 - \alpha_{t-1}) + \beta (h_{t-1}^2 - \alpha_{t-1})$$

$$(5b) \quad \alpha_t = \omega + \rho \alpha_{t-1} + \phi (\varepsilon_{t-1}^2 - h_{t-1}^2)$$

where the forecasting error $(\varepsilon_t^2 - h_t^2)$ serves as the driving force for the time-dependent movement of the long-run component, α_t , and the difference between the conditional variance and long-run volatility, $h_t^2 - \alpha_t$, defines the short-run component. The initial impact of a shock to the transitory component is quantified by ϕ , while ρ indicates the degree of memory in the transitory component, the sum of these parameters providing a measure of transitory shock persistence. The initial effect of a shock to the permanent component is given by α , with persistence measured by the autoregressive root, β , and where $1 > \rho > (\alpha + \beta)$ 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. The conditional variance is covariance stationary provided that the permanent component and the transitory component are both covariance stationary, as satisfied by $\rho < 1$ and $(\alpha + \beta) < 1$ respectively, while the additional restriction of non-negativity on the model parameters ensures that h_t^2 is non-negative as long as q_t is non-negative.¹¹

3. Data and Preliminary Diagnostics

3.1. Data and Market Background

The data analysed here consists of the prices of UK government bond (Long Gilt) futures contracts traded on the London International Financial Futures and Options Exchange (LIFFE), which is also the data source.¹² The Long Gilt futures contract is of interest as a heavily traded investment and hedging instrument, the main users of which LIFFE identifies as market makers, institutional investors and issuers of long-term debt; for purposes of hedging, investment, asset allocation, portfolio insurance and duration adjustment, such activities being primarily driven by consideration of long-run factors and underlying fundamentals. A further feature of the Long Gilt futures market is its low margin requirement, which encourages a degree of short-term speculation and provides circumstances conducive to noise-trading of the manner described in the previous section.

The sample covers the period 24th January 1992 to 30th June 1995. The contract price data, p_t , is sampled at five and fifteen minute intervals and transformed to yield the returns series, $r_t = \log(p_t/p_{t-1})$, with the overnight return excluded so as to ensure consistent time-series.¹³ With 846 trading days in the sample period, this yields 80,163 observations at the five minute frequency, and

26,721 observations at the fifteen minute frequency.¹⁴

As has been noted elsewhere, high frequency intra-day data is strongly characterised by high-frequency periodicity corresponding to proximity in time to market opening and closing, macroeconomic and other systematic news releases and other factors, and where the strength of these intra-day effects is such that failing to adjust for them can result in misleading analysis of the dynamic dependencies in the data (Goodhart et. al., 1993; Andersen and Bollerslev, 1997b; Guillaume et. al., 1997; Goodhart and O'Hara, 1997). Prior to estimation, we therefore follow Andersen and Bollerslev (1997b) in standardising returns by the mean absolute value for each intra-day time interval, at both the both five and fifteen minute frequencies.^{15, 16} Summary statistics for the data, both before and after adjustment by standardisation, including measures of central tendency, skewness, kurtosis, tests of normality, and selective correlogram values for the levels and squares of the series, are reported in Table 1. Self-evidently, adjustment increases the range and standard deviation of the underlying series, which has the indirect benefit of aiding parameter convergence in estimation. Otherwise the basic properties of the data are little affected. The distributional properties of the adjusted data are further illustrated in Figure 1, which depicts the results of nonparametric Epanechnikov kernel density estimation for both data frequencies, where bandwidth selection is determined according to the data-based criteria of Silverman (1986). The 'peakedness' relative to the normal indicated by the kurtosis statistics in Table 1 is clearly obvious in both distributions, and further motivates the consideration of GARCH processes below. Additionally evident are the 'peaked shoulders' in the distributions, also present in the comparable distributions of the unadjusted data, and most pronounced in the fifteen minute frequency data, which suggests a concentration of data points a margin either side of the zero mean, and more so on the upper side of the distribution. This property further suggests to us the influence of significant market frictions,

such that beyond small return values a range of price changes become more pronounced and numerous, and reinforces our consideration of threshold models able to accommodate this feature below. Before proceeding to the estimation of such models, however, we first consider formal statistical tests for the presence of such nonlinearities.

3.2. Preliminary Diagnostics

The specification of preliminary linear AR(p) models is determined by reference to the autocorrelation and partial autocorrelation functions, the Schwarz criterion, the estimated log-likelihood, and residual tests for serial correlation.¹⁷ This identification procedure indicates that an AR(2) process is appropriate at the five minute frequency, whilst an AR(1) model is appropriate at the fifteen minute frequency. Model estimates for these specifications are reported in the first column of results in Tables 2 and 3 respectively. At both frequencies, autoregressive parameters are negative and significant, parameter values confirming the absence of long-lived persistence or drift in returns.¹⁸

Given appropriately specified AR models, we test for the presence of conditional mean nonlinearity following the procedure detailed in Teräsvirta and Anderson (1992), Granger and Teräsvirta (1993) and Teräsvirta (1994). This entails testing for threshold nonlinearities against the null of linearity over a range of suitable possible values for the delay parameter d . The corresponding LM-type test of AR(p) linearity assuming known d is equivalent to the test of the null hypothesis of linearity $H_0: \beta_{2j} = \beta_{3j} = \beta_{4j} = 0$ ($j = 1, \dots, p$), against the alternative in the following artificial regression:

$$(6) \quad y_t = \beta_0 + \sum_{j=1}^p \beta_{1j} y_{t-j} + \sum_{j=1}^p \beta_{2j} y_{t-j} y_{t-d} + \sum_{j=1}^p \beta_{3j} y_{t-j} y_{t-d}^2 + \sum_{j=1}^p \beta_{4j} y_{t-j} y_{t-d}^3$$

The test statistic, computed as $LM = T(SSR_0 - SSR_1)/SSR_0$ where T denotes the sample size, SSR_0 the sum of squared residuals from the linear AR(p) model and SSR_1 the sum of squared residuals obtained from (6), is asymptotically distributed as $\chi^2_{(p/2)(p+1)+2p^2}$ where d is unknown. Where linearity is rejected for more than one value of the delay parameter, then d is determined such that $p(d) = \arg \min_{1 \leq d \leq D} p(d)$, where $p(d)$ refers to the probability value at which the null of linearity is marginally rejected.¹⁹ Application of these tests for all possible delay values $1 \leq d \leq 4$ for both data frequencies confirm rejection of the null hypothesis of linearity in favour of STAR nonlinearity with application of the minimum $p(d)$ rule indicating $d=2$ at the five minute frequency and $d=1$ at the fifteen minute frequency.^{20, 21} Given this diagnostic support for non-linear STAR models over linear AR alternatives as descriptions of conditional mean structure in long gilt futures returns at frequencies of both five and fifteen minutes, we proceed to full estimation of those models in the following section.

4. Results

4.1. Model Identification and Evaluation

Estimation of all models reported below is by iterative non-linear least squares. The validity of the estimated models is appraised on the basis of the significance of autoregressive terms and examination of coefficient estimates, in particular ensuring that the transition value, c , is within the range of $\{r_t\}$. The Akaike and Schwarz information criteria are also used to guide selection amongst competing models (Teräsvirta, 1994). The properties of the model residuals are also examined, both for departures from

normality and for remaining ARCH effects. We also examine the dynamic properties of the regimes corresponding to $F(r_{t-2}) = 0$ and $F(r_{t-2}) = 1$ by inspecting the roots of the relevant characteristic polynomials, as well as the dynamic properties of the full models. In the absence of a general analytical solution, the latter procedure is performed numerically, using data generated from the estimated model after setting the error term to zero, with a sequence of observed values of the series acting as starting values, several of the latter being considered. For the models under investigation, this may result in a unique stable equilibrium, a limit cycle such that a set of values repeat themselves perpetually, chaotic realisations, whereby a small change in initial values results in divergent but stable limit points, or explosive values (in which case the model is rejected).

4.2. Nonlinear Dependence in Conditional Mean and Conditional Variance

Preliminary estimates of ESTAR models of nonlinear dependence in conditional mean alone are reported in the fourth column of results for each frequency in Tables 2 and 3. The properties of these models are broadly similar in terms of specification, parameter sign and magnitude to those which obtain under joint conditional mean and conditional variance estimation, with the exception that the estimated transition parameters are strictly statistically insignificant suggesting a degree of misspecification due to the conditional variance structure not being modelled (though see the discussion in 2.1 above), and the remainder of our discussion therefore focuses on jointly estimated models of nonlinear dependence in both conditional mean and variance. ESTAR-EGARCH and ESTAR-CGARCH estimation results are reported in the fourth and fifth columns of Tables 2 and 3, with corresponding AR-EGARCH and AR-CGARCH estimation results reported in columns two and three of those Tables for purposes of comparison.

4.2.1. Five Minute Frequency Results

At the five minute frequency, of immediate note is the reduction in model order relative to the linear case. For the resulting ESTAR(1)-EGARCH model, the central regime corresponding to $F(r_{t-d}) \sim 0$, which arises as $(r_{t-d} - c) \rightarrow 0$, is described by an AR(1) process, whilst the outer regimes corresponding to $F(r_{t-d}) \neq 0$ as $(r_{t-d} - c) \rightarrow \pm\infty$ invokes an additional AR(0) process.²² Thus, there is significant negative autocorrelation in returns irrespective of size, with the nonlinearity present being described by a shifting intercept dependent on the magnitude of returns relative to the interval norm. The latter specifically implies significant negative drift in returns in the neighbourhood of the threshold value, but, on net, a tendency towards positive drift for larger returns of either sign as $F(r_{t-d}) \rightarrow 1$. Both regimes of the model are trivially stationary, and the full model is characterised by stable roots and a near-zero unique limit point (0.0044), with rapid adjustment to equilibrium within approximately eleven periods, or fifty-five minutes. The estimated value of the threshold parameter, c , suggests that the central regime characterises returns that are around one and three-quarter times higher than the interval norm (standardised in relation to the intra-day interval average).

Transition between regimes is dictated by the estimated transition function, which is portrayed in Figure 2(a). The estimated transition parameter value of 0.14 (or 0.32 after reversing the scale transformation), significant at the ten per cent level, suggests a moderate speed of transition between regimes, and therefore a tendency for returns to sojourn in the centre regime. The minimum of the function corresponds with the threshold parameter, its width in the neighbourhood of c determines the range of the central regime, whilst its steepness (symmetric about c) determines the speed of transition between the centre and outer regimes. The mid-points between regimes occur for the data values (-

1.5, 5.0), expressed as multiples of interval means. That is, the mid-way transition point between the centre and outer regimes is passed when returns are falling by one-and-a-half times their average for that intra-day five minute interval, or rising by five times their normal interval value. An approximate measure of the width of the centre regime, corresponding to $F(r_{t-d}) < 0.2$, yields the pair of data points (0.0, 3.5); that is returns ranging from zero to three-and-a-half times the relevant five minute interval norm. Finally, and consistent with AR-EGARCH model estimates, conditional variance parameters for the five minute frequency ESTAR-EGARCH model indicate a high and significant measure of persistence in shocks to volatility of 0.97, implying half-life decay in just under two hours, and significant volatility clustering, but no evidence of significant asymmetry in volatility with respect to shocks of differing sign.

ESTAR-CGARCH model estimates at the five minute frequency confirm the magnitude and significance of the ESTAR parameters discussed above, but with some increase in the threshold parameter, decrease in the transition parameter, and the additional significance of the centre regime intercept. The preceding discussion therefore mostly continues to hold, other than that the mid-points between regimes now occurs for data values (0.5, 7.5), or returns of one-half and seven-and-a-half times their interval average, while the central regime has a width corresponding to $F(r_{t-d}) < 0.2$ of (1.9, 5.8), or returns of approximately two to six times their average interval value. This estimated transition function is portrayed in panel (b) of Figure 2. Concerning the CGARCH parameter estimates, the initial effect of a shock to the permanent component of volatility, as quantified by the parameter N , is fairly modest at under 0.2, while the autoregressive root, D , is strongly significant at over 0.99, suggesting very strong persistence in the effect of such shocks, with a half-life decay of approximately four days. Both parameters of the transitory component are significant and provide a joint persistence

measure of over 0.9, implying a half-life decay in shocks to transitory volatility of approximately 35 minutes.

In a comparison across estimated models at the five minute frequency, the log-likelihood is clearly maximized in the ESTAR-CGARCH case. Testing between linear and nonlinear mean specifications at the five minute frequency cannot be conducted using likelihood ratio tests due to the non-nested nature of the models arising from the difference in the AR and ESTAR autoregressive orders. Likelihood ratio testing between EGARCH and CGARCH specifications is also not possible. We therefore discriminate between these non-nested models on the basis of information criteria minimization. Employing both the Akaike information criterion (AIC) and Schwarz (Bayesian) information criterion (BIC), CGARCH variance specifications are preferred amongst both AR and ESTAR models when considered separately. However, between those AR-CGARCH and ESTAR-CGARCH models, while the BIC marginally favours the former, the AIC marginally favours the latter. Residual diagnostics indicate residual non-normality, primarily due to excess kurtosis, reinforcing the use of Bollerslev-Wooldridge robust standard errors in the appraisal of parameter significance conducted above. However, LM tests indicate the presence of remaining ARCH effects for all models, though only of first order form for both CGARCH models.²³

4.2.2. Fifteen Minute Frequency Results

At the fifteen-minute frequency, a STAR(1)-EGARCH model again holds, but now with significant first order autoregressive parameters and insignificant intercepts for both $\beta(r_{t-d}) = 0$ and $\beta(r_{t-d}) \neq 0$.

Moreover, estimated autoregressive parameter values are approximately equal but of opposing sign such that, for $\beta(r_{t-d}) = 1$, returns in the outer regimes are described by driftless random walks, whilst

for $F(r_{t-d}) = 0$, returns described by the central regime are characterised by significant negative autocorrelation. Transition between these regimes is again governed by the estimated transition function parameters, ζ and c , which yield the transition function depicted in Figure 3(a). The estimated threshold value is again positive in value, but now statistically insignificantly different from zero. The estimated transition parameter of around 0.25 (or 1.51 after reversing the scale transformation), again significant at the ten per cent level, suggests a far greater speed of transition between regimes than at the five minute frequency. The mid-points between regimes occur for the values (-2.5, 5.5), such that mid-way transition between the centre and outer regimes occurs when returns are falling by two-and-a-half times their interval average, or rising by five-and-a-half times their interval average. The approximate width of the centre regime, corresponding to $F(r_{t-d}) < 0.2$, is delimited by return values relative to interval norms of (-0.7, 3.8). The model is again characterised by stable roots in each regime, with a near-zero unique limit point (0.0153) achieved within eight periods, though the majority of adjustment to equilibrium occurs in only four periods, or one hour. Conditional variance parameters for the fifteen minute ESTAR-EGARCH (and AR-EGARCH) model continue to indicate a high and significant measure of persistence in shocks to volatility at over 0.98, implying a shock persistence half-life of over nine hours, or more than a full trading day, and significant volatility clustering, but again no evidence of significant asymmetry in volatility with respect to shocks of differing sign.

ESTAR-CGARCH estimates confirm the preceding mean model interpretation for fifteen minute returns, though the threshold parameter is much reduced and continues to be statistically indistinguishable from zero, whilst the transition parameter is increased and significant at the conventional 5% probability level. The transition function, depicted in Figure 3(b), is more closely centred on zero, with faster transition between regimes dictated by the ζ -estimate of 0.44 (2.66 after

scale transformation reversal). The regime transition mid-points now correspond to the data values (-2.5, 3.5), and the central regime range measure identified by $F(r_t, \hat{\mu}) < 0.2$ to (-1.1, 2.2). Thus, returns rapidly move to a random walk process once they have fallen by more than their average absolute value, or risen by more than twice their average value. Concerning return volatility, CGARCH parameters are again significant throughout, with very strong permanent component persistence now implying a shock half-life of approximately ten days at the fifteen-minute frequency, whilst transitory component shock persistence exhibits a half-life of almost exactly one-hour.

Finally, across fifteen minute frequency models, evaluation of linear mean versus nonlinear mean models is possible on the basis of likelihood ratio tests, and the nonlinear ESTAR alternative is consistently favoured.²⁴ In discriminating across all models, both the AIC and BIC criteria clearly favour the ESTAR-CGARCH specification. Moreover, residual ARCH effects are insignificant for the ESTAR-CGARCH model at all lag lengths, suggesting that all volatility structure is adequately captured.²⁵ The broader interpretation and implications of these findings are discussed in the following concluding section.

5. Summary and Implications

Motivated by considerations of market frictions and heterogeneities in information flows and market agents, the empirical evidence reported here has sought to identify the source of nonlinear dependence in futures returns, with particular regard to the potential for such dependence to arise either in conditional mean or conditional variance, and separately or jointly. Preliminary tests against the null of linearity indicate the presence of smooth transition autoregressive nonlinearity in the conditional mean of UK long gilt futures returns at both the five and fifteen minute frequency. At the five minute

frequency, the estimated linear model is second-order autoregressive, whilst the nonlinear STAR model consists of a first-order autoregressive process with switching intercept. That structure is robust to the joint estimation of conditional variance processes of either exponential-GARCH or component-GARCH type. The former confirms the presence of significant clustering and persistence in conditional volatility, whilst the latter entails the successful decomposition of volatility into a long-lived permanent component and a more ephemeral transitory component. At the fifteen minute frequency, both AR and STAR processes are of first-order, the nonlinear process exhibiting negative autocorrelation for small returns near zero, but with cancelling coefficients consistent with near random-walk behaviour for larger returns of either sign. This structure is also robust to the joint presence of EGARCH or CGARCH conditional variance processes, with the STAR-CGARCH specification being unambiguously favoured on the basis of model selection criteria and residual diagnostics.

The persistence of return movements at the five minute frequency, and for larger returns especially, strongly suggests that the market does not adjust to equilibrium within that fine high frequency time interval. The persistence of smaller returns but not larger returns at the fifteen minute frequency suggests that the market is slow to respond to small price movements, but that the profitable opportunities implied by larger movements are mostly eliminated within the quarter-hour, with the greater part of convergence to full equilibrium in the absence of further shocks being achieved in approximately one-hour. The significance of the component structure to volatility is particularly pertinent in the light of recent arguments suggesting its existence is due to heterogeneity in information flows or heterogeneity in trader types. In conjunction with the empirical findings reported here, these considerations lead us to conclude that long gilt futures market returns are driven by the response of heterogeneous traders to heterogeneous information flows, possibly with a degree of noise trading in

response to smaller return values, but with the fairly rapid extraction of profitable opportunities consistent with weak-form market efficiency following larger price movements when measured relative to the relevant intra-day time interval average.

Our findings also have broader implications for considerations of market efficiency and the use of technical analysis. The existence of nonlinearities in market returns might generally be expected to allow the potential for predictable regularities. What is demonstrated here is that such regularities are confined to only a very high frequency of time interval and only small movements in prices. Nevertheless, for those frequencies and range of price movements, the potential for tapping those regularities, possibly through the use of technical analysis or trading rules, remains. The efficient markets hypothesis may therefore not be expected to hold at the higher intra-day frequencies as the mechanisms by which markets adjust to equilibrium are at work, and the apparent widespread use of technical analysis in financial markets that has been documented may receive some empirical support.²⁶

On that level, the results reported here may also be interpreted as illustrating the rate at which market prices impound new information over the higher intra-day frequencies, particularly if some information may initially be private prior its market dissemination.

**Table 1. Summary Statistics: Unadjusted and Adjusted Data -
Five and Fifteen Minute Frequencies**

Frequency:	Five Minutes		Fifteen Minutes	
Data:	Unadjusted	Adjusted	Unadjusted	Adjusted
Mean	2.55x10 ⁻⁷	0.000298	8.13x10 ⁻⁷	0.000971
Median	0.000000	0.000000	0.000000	0.000000
Standard Deviation	0.000562	1.5116	0.000923	2.4490
Minimum	0.0062	-22.0515	-0.0097	-22.5075
Maximum	0.0090	26.1938	0.0103	24.6209
Skewness	0.06	0.07	0.01	-0.01
Kurtosis	11.58	11.11	11.40	9.72
Normality	246,163.7	219,835.0	78,540.5	50,216.04
Q ₁	441.77	579.30	45.4	53.36
Q ₁₀	468.53	618.51	74.13	75.96
Q ₂₀	499.89	640.74	87.58	89.75
Q ₁ ²	2,028.6	1,962.1	749.19	809.05
Q ₁₀ ²	10,569.0	11,139.0	3,055.9	3,395.1
Q ₂₀ ²	14,947.0	15,026.0	3,775.6	5,148.7

Notes: For data description and details of the adjustment procedure used to accommodate intra-daily 'seasonal' patterns through scaling by time interval mean values, see Section 3.1. Summary statistics are mostly self-explanatory. Additionally, 'Normality' is the Jarque-Bera test of the null hypothesis of normality, distributed as χ^2_2 which is clearly rejected throughout. $Q_i(t-1, 10, 20)$ are selected values from the correlogram of the data, and $Q_i^2(t-1, 10, 20)$ are selected values from the correlogram of the squares of the data.

Table 2.
Model Estimation Results - Five Minute Frequency.

Model/ Parameter	AR	AR- EGARCH	AR- CGARCH	ESTAR	ESTAR- EGARCH	ESTAR- CGARCH
B0	0.0003 (0.0053)	0.0297 (0.0472)	0.0073* (0.0044)	-0.1068* (0.0361)	-0.0957 (0.0120)	-0.1337* (0.0365)
B1	-0.0869* (0.0056)	-0.1068* (0.0062)	-0.1091* (0.0040)	-0.0872* (0.0056)	-0.1052* (0.0051)	-0.1091* (0.0040)
B2	-0.0228* (0.0053)	-0.0414* (0.0048)	-0.0392* (0.0041)			
20				0.1885* (0.0546)	0.2814* (0.1053)	0.2529* (0.0564)
C				0.1629 (0.1348)	0.1400** (0.0763)	0.1286** (0.0765)
c				3.5510* (1.2136)	1.7663* (0.8961)	3.8721* (1.0110)
T		-0.0942* (0.0086)	1.9949* (0.1428)		-0.0927* (0.0066)	1.9950* (0.1425)
D			0.9973* (0.0005)			0.9973* (0.0005)
N			0.0160* (0.0021)			0.0158* (0.0021)
"		0.1569* (0.0076)	0.0743* (0.0048)		0.1591* (0.0076)	0.0741* (0.0048)
\$		0.9712* (0.0095)	0.8297* (0.0116)		0.9668* (0.0066)	0.8307* (0.0115)
-		-0.0037 (0.0051)			-0.0084 (0.0052)	
Log L	-146554.0	-137455.3	-136588.9	-146542.1	-137417.1	-136584.3
AIC	3.6566	3.4300	3.4081	3.6563	3.4288	3.4080
BIC	3.6569	3.4308	3.4090	3.6569	3.4298	3.4092
Skew	0.07	-0.06	0.14	0.07	0.03	0.14
Kurt	11.51	16.33	10.61	11.50	12.99	10.61
JB	242158.5*	593868.9*	193590.5*		333435.4*	193467.6*
A₁	1671.11*	31.32*	8.96*	1663.16*	51.00*	9.18*

A_{10}	5254.39*	40.31*	10.37		67.21*	10.78
A_{20}	5504.87*	47.43*	17.06		73.84*	17.44

Notes: For model mnemonics and specifications, see Section 2. Additionally, LogL denotes the maximized log likelihood value, AIC and BIC denote the Akaike and Schwarz (Bayesian) information criteria, Skew and Kurt are regular measures of skewness and kurtosis respectively, JB is the Jarque-Bera test of the null of normality, distributed as χ^2_3 , and A_i is the regular ARCH LM test for lags $i=1,10,20$ distributed as χ^2_i . Asterisk(s) denote significance at the 5%(10%) level.

Table 3.
Model Estimation Results - Fifteen Minute Frequency

Model/ Parameter	AR	AR- EGARCH	AR- CGARCH	ESTAR	ESTAR- EGARCH	ESTAR- CGARCH
B0	0.0009 (0.0150)	0.0122 (0.0138)	0.0206* (0.0123)	0.0175 (0.0180)	0.0427 (0.0323)	0.0486* (0.0201)
B1	-0.0448* (0.0097)	-0.0651* (0.0077)	-0.0713* (0.0072)	-0.1073* (0.0312)	-0.1252* (0.0264)	-0.1393* (0.0198)
20				-0.1040 (0.1618)	0.0201 (0.2901)	-0.0823 (0.1302)
21				0.0919* (0.0394)	0.1338* (0.0280)	0.1229* (0.0256)
C				0.4152 (0.3494)	0.2525** (0.1452)	0.4438* (0.2240)
c				-0.2741 (1.6757)	1.5153 (1.5076)	0.5498 (0.9730)
T		-0.0784* (0.0075)	5.4146* (0.6261)		-0.0787* (0.0071)	5.2752* (0.6256)
D			0.9969* (0.0009)			0.9975* (0.0007)
N			0.0169* (0.0024)			0.0139* (0.0023)
"		0.1505* (0.0121)	0.0914* (0.0117)		0.1478* (0.0113)	0.0896* (0.0110)
\$		0.9808* (0.0025)	0.7316* (0.0319)		0.9820* (0.0024)	0.7559* (0.0285)
-		-0.0111 (0.0083)			-0.0114 (0.0083)	
Log L	-61811.76	-58971.20	-58744.66	-61799.07	-58951.49	-58728.26
AIC	4.6268	4.4154	4.3990	4.6261	4.4133	4.3967
BIC	4.6274	4.4172	4.4011	4.6280	4.4164	4.4000
Skew	-0.01	0.08	0.03	0.002	0.09	0.03

Kurt	9.90	8.91	8.85	9.82	8.87	8.81
JB	52946.09*	38866.79*	38076.52*		38332.03*	37534.74*
A₁	817.02*	142.57*	4.29*	822.10*	125.74*	3.51
A₁₀	1727.35*	175.47*	9.97		161.93*	9.23
A₂₀	1962.56*	184.97*	11.45		171.53*	10.52
<p>Notes: For model mnemonics and specifications, see Section 2. Additionally, LogL denotes the maximized log likelihood value, AIC and BIC denote the Akaike and Schwarz (Bayesian) information criteria, Skew and Kurt are regular measures of skewness and kurtosis respectively, JB is the Jarque-Bera test of the null of normality, distributed as χ^2_3, and A_i is the regular ARCH LM test for lags i=1,10,20, distributed as χ^2_i. Asterisk(s) denote significance at the 5%(10%) level.</p>						

Notes:

1. For reviews, see Bollerslev, Chou and Kroner (1992), Bera and Higgins (1993), and Bollerslev, Engle and Nelson (1994).
2. Other explanations for time-varying volatility based on considerations of market microstructure include that of Kyle (1985) whereby information held by an informed trader is transmitted into prices gradually through information diffusion. An alternative rationale is provided the theoretical model of Timmermann (1995), where the source of volatility clustering is incomplete learning and limited knowledge of the process generating fundamentals. Other explanations which have led to extensions of the ARCH model, relate to the influence of macroeconomic volatility as revealed through such variables as the interest rate (Glosten, Jagannathan and Runkle, 1993), the money supply and oil prices (Engel and Rodriguez, 1989) and various measures of the state of the business cycle (Schwert, 1989), while the models of Sentana (1995) and Bera, Higgins and Lee (1992) have been afforded a random coefficient interpretation.
3. There has been interest in testing high frequency data for the presence of deterministic nonlinear dynamics of chaotic form, for which there would appear to be little evidence (Vassilicos, 1990; Vassilicos, Demos and Tata, 1992; Vassilicos and Demos, 1994; Abhayankar, Copeland and Wong, 1995, 1997). Such tests have, however, suggested the presence of strong stochastic nonlinearities, though this has typically been ascribed to the presence of ARCH effects without explicit consideration of the nonlinear mean alternative. For an excellent review of issues and applications associated with high frequency financial data, see Goodhart and O'Hara (1997).
4. The few exceptions to this claim have typically been concerned with exchange rate data. For example, Kräger and Kugler (1993) examine the performance of threshold models using weekly exchange rate data from the 1980's. Peel and Speight (1994) model inter-war exchange rates using threshold models and the bilinear model (Granger and Andersen, 1978), while Peel and Speight (1996) model East European black-market exchange rates using the bilinear model. Bera and Higgins (1997) examine bilinear models for US stock prices, and the pound-dollar exchange rate. Coakley and Fuertes (1997), Obstfeld and Taylor (1997), and O'Connell and Wei (1997) examine real exchange rates during the post-war float using threshold models, while Coakley and Fuertes (1998) do likewise for nominal exchange rates.
5. This observation is compounded by the fact that if the true data generating process is indeed non-linear, then fitting a linear model will result in a longer lag length that required by the correct non-linear

specification (Granger and Teräsvirta, 1993).

6. Additionally, the consideration of a non-linear conditional mean model for asset prices has often been regarded as providing a competing potential explanation for the non-linear dependence implied by GARCH models of volatility (eg. Kräger and Kugler, 1993; Bera and Higgins, 1997) rather than as providing a complementary explanation (eg. Weiss, 1984; Peel, Lane and Raeburn, 1997; Peel and Speight, 1998).

7. The ESTAR model may also be interpreted as a generalisation of the earlier exponential autoregressive (EAR) model of Haggan and Ozaki (1981), the more restrictive EAR case being obtained under $\theta_0 = c = 1$, that restriction making the E(ST)AR model location invariant.

8. Alternatives specifications for the transition function include the logistic and heaviside functions, yielding LSTAR and TAR models respectively. The former has the capacity to accommodate the latter, but more generally permits smooth transition between differing dynamics associated with positive and negative signs for $(r_{t-d} - c)$. The ESTAR specification is preferred here for the reasons given later in the text. However, the LSTAR model is also considered as an empirical alternative, the results of which are noted further below.

9. For further details of the mixture of distributions interpretation of ARCH models see, for example, Bera and Higgins (1993, pp.324-7). For further discussion of the interpretation of the EGARCH model as a discrete time approximation to an underlying diffusion model expressed in continuous time see, for example, Bollerslev et. al. (1994, pp. 2994-6).

10. The CGARCH model may also be regarded as a variant of the threshold GARCH model proposed by Rabemananjara and Zakoian (1993). On the interpretation of the CGARCH model as a diffusion approximation process, see Engle and Lee (1996).

11. By substitution using (6) it is readily shown that the component model may be alternatively expressed as a GARCH(2,2) model, reducing to the GARCH(1,1) case if $\alpha = \beta = 0$ or $\rho = \phi = 0$. The GARCH model is thus only capable of describing at most one element of the more general condition variance component specification and represents the long-run component only under the specific conditions $\alpha = \beta = 0, \rho = 1$. It is due to this limitation of the basic GARCH form, and its representation as a special case of the CGARCH model, that the basic GARCH form is not given explicit consideration here. It should also be noted that whilst the CGARCH model be extended to asymmetric form, given the lack of empirical support for conditional variance asymmetry in the EGARCH model estimates reported further below, we do not pursue the ACGARCH form here.

12. A financial futures contract is formally defined as an agreement to exchange a specified quantity and quality of an underlying asset at a specified date in the future for a price agreed at the time the contract is traded. The contract can either require physical delivery of the underlying asset or can be cash settled. A cash settled contract requires a cash amount to be paid on the delivery date, that sum reflecting the difference between the initial futures price and the price of the underlying asset at settlement. The price agreed when the futures contract is traded is not paid or received in full nor does the underlying product change hands at this point. Instead, margin is lodged by both the buyer and seller of the contract. This margin acts as financial surety that they can, should they need to, fulfil their side of the contract. On the last trading day of a futures contract, the price of the contract will converge to the price of the underlying asset. Prior to expiry these two prices may be different. This is primarily due to the different financial circumstances caused by having a position in a futures contract rather than in the

underlying asset, such as the absence of interest or coupon payments, and the fact that when a future is purchased there is a much smaller cash outlay, and the difference in outlay can be invested to earn interest. However, arbitrage ensures that futures price movements are closely correlated with movements in the price of the underlying asset.

13. LIFFE futures contracts have four, quarterly, delivery months, in March, June September and December. Since several contracts may be traded simultaneously, and given the continuous series requirement, a decision must be made as to which contract price to take at any given point in time. In the case of Long Gilt futures the contract can be delivered at the seller's discretion on any business day in the delivery month up to two days prior to the last business day of that month. In their examinations of bond futures, Becker, Finnerty and Kopecky (1993, 1995, 1996) use the nearest delivery date contract until two days from the start of the delivery month at which point they switch to using the next nearest delivery contract. We follow Abhyankar et. al. (1995) and Buckle et. al. (1998) in basing the choice of contract on traded volume. Thus, the switch of contract here occurs on the day on which volume in the second nearest contract exceeds traded volume in the nearest contract. This is approximately one month before expiry of the nearest contract but is not a fixed distance from expiry. Having already excluded the overnight return there is no further requirement to adjust the series as a result of the change in contracts. Nevertheless, the effect of this splicing of contracts on the estimates reported in Section 4 is tested for as a matter of empirical robustness, the results of which are noted where appropriate.

14. Our attention is restricted to the five and fifteen minute frequencies due to the lack of linear autoregressive structure at lower frequencies, for the reasons set out in the Introduction. A possible objection to the use of high frequency fixed interval intra-day transactions data, is that no transactions may occur during some intervals such that the very measurement of returns becomes problematic. This issue is not peculiar to intra-day data, since the problem of sporadic trading also arises to some degree in lower frequency data, but it is potentially more acute at the intra-day frequency. However, for the heavily traded contract analysed here the problem does not arise. Over the entire data set, zero return and volume incidences account for only 3% and 0.4% of data points at the five and fifteen minute frequencies respectively.

15. Various alternative adjustments for systematic intra-day effects have been proposed in the literature, including the use of interval dummies (Baillie and Bollerslev, 1990, 1991), time-scaling (Dacorogna et. al., 1993), Fourier transforms (Anderson and Bollerslev, 1994) and artificial neural networks (Lo, 1994). In order not to compound the potential nonlinearities we are testing for, we forego the latter approaches in favour of the methodology described in the text. For a more detailed discussion of the intra-day deterministic patterns in LIFFE futures returns and volume data, including that analysed here, see ap Gwilym, McMillan and Speight (1999), and for an examination of alternative intra-day seasonal adjustment methods in LIFFE FTSE-100 futures in particular, see McMillan and Speight (1999).

16. Given that the floor trading times for Long Gilt futures changed on August 1st 1994 from 8:30am-4:15pm to 8:00am-4:15pm, this standardisation exercise is conducted separately for both sub-periods so as not to confound the intra-day deterministic pattern. The effect of this change in trading times on the estimates reported in Section 4 is tested for as a matter of empirical robustness, the results of which are noted where appropriate.

17. To check for the absence of remaining serial correlation the heteroscedasticity-robust LM test of Wooldridge (1990) was employed. Results available upon request.

18. The presence of negative autocorrelation in high frequency returns is not uncommon, various explanations for which have been proposed, including inventory concerns of market makers and bid-ask bounce. For further discussion see Goodhart and O'Hara (1997, pp.95-6).

19. The rationale for such a selection method is as follows. When d is correctly specified, (6) is the appropriate auxiliary regression against the non-linear alternative. If another d is selected, then (6) is mis-specified. The power of the corresponding test against the mis-specified non-linear model will thus be weaker than the power of the test based upon the correctly specified auxiliary regression.

20. The exceptions to linearity rejection at the 5% significance level arise for $d=1$ at the higher frequency and for $d=2$ and $d=3$ at the lower frequency. Test statistic values for increasing values of d (probability values in parentheses) are, at the five minute frequency, 14.82 (0.19), 55.74 (5.67×10^{-8}), 37.45 (9.68×10^{-5}), 21.83 (0.03), and at the fifteen minute frequency, 42.16 (3.71×10^{-9}), 3.86 (0.28), 7.43 (0.06), 10.38 (0.02).

21. On the basis of a series of nested tests using (6) and given the limitation of assuming d is known, in which case the LM test statistic is asymptotically distributed as χ^2_{3p} , it is also possible to discriminate between the ESTAR model and the LSTAR variant noted previously. The sequence of hypotheses tested are: (i) $H_{01}: \beta_{4j} = 0$; (ii) $H_{02}: \beta_{3j} = 0 \mid \beta_{4j} = 0$; (iii) $H_{03}: \beta_{2j} = 0 \mid \beta_{3j} = \beta_{4j} = 0; (j = 1, \dots, p)$. The rationale behind this sequence is based on interpreting the coefficients β_{4j} as functions of the parameters of the STAR model. Thus, if (i) is accepted and (ii) is rejected the ESTAR variant may be preferred. If either (i) is rejected, or (i) and (ii) are accepted and (iii) is rejected, the LSTAR variant may be preferred. Adopting the d values indicated in the text, test statistics values at the five-minute frequency are (i) 1.94, (ii) 43.93, and (iii) 9.88, relative to χ^2_6 critical values of 12.59 and 10.64 at the 5% and 10% significance levels respectively, clearly confirming preference for an ESTAR specification. At the fifteen minute frequency, test statistic values are (i) 29.06, (ii) 9.61, and (iii) 3.51, relative to χ^2_3 critical values of 7.81 and 6.25 at the 5% and 10% significance levels. The latter are somewhat inconclusive, and whilst this may imply a possible LSTAR specification, in estimation that model was found to behave poorly in terms of parameter identification, and is in any event not well motivated in the context addressed here, having typically been applied previously in the analysis of macroeconomic business cycle asymmetry. The LSTAR alternative is therefore not considered further here, though further details are available upon request.

22. Autoregressive parameters associated with non-zero values of the transition function terms proved insignificant in estimation.

23. In an attempt to model these effects longer lag lengths were examined. Whilst a GARCH(2,2) specification was found to adequately capture remaining ARCH effects, and provides further support for our analysis of the CGARCH model (see n.11 above), the estimated model failed to satisfy necessary GARCH parameter restrictions, and is therefore not discussed further.

24. Likelihood ratio test values are, for the following comparisons, AR versus ESTAR, 25.38, AR-EGARCH versus ESTAR-EGARCH, 39.41, AR-CGARCH versus ESTAR-CGARCH, 32.80, relative to 5% and 1% critical values of 9.49 and 13.28 respectively.

25. A range of further tests confirm that the results reported in the text are robust to consideration of a number of issues raised in preceding notes. First, Wald tests for the significance of a dummy associated with the extended opening of the long gilt futures market in August 1994 (see n.16) are insignificant in all conditional mean equations for both frequencies. Second, conditional mean dummies associated with

contract rollover dates in the construction of continuous returns series (see n.13), whilst significant but with no conformity of sign, have no or negligible impact on other parameter estimates. Third, conditional mean and variance dummies reflecting the time-to-maturity of specific futures contracts, whilst similarly generally significant but of no consistent sign, also have little or no effect on model parameter estimates. Full details of these robustness exercises, for both data frequencies and all models reported in the text, are available on request.

26. See, for example, Taylor and Allen (1992).