

**LONG MEMORY, THE “TAYLOR EFFECT” AND
INTRADAY VOLATILITY IN
COMMODITY FUTURES MARKETS**

Celso Brunetti*
Department of Business Studies
The University of Edinburgh

September 1999
First Draft: November 1995

ABSTRACT

This paper investigates long-term dependence in commodity futures markets. Using daily futures returns on cocoa, coffee and sugar, we show that FIGARCH models are able to adequately describe both the long and the short run characteristics of commodity market volatility. The paper also considers three measures of risk – squared returns, absolute returns and ranges (intraday volatility) - and finds that they exhibit the long memory property. The range shows the strongest autocorrelation structure. Moreover, there is evidence of the so-called “Taylor effect”.

Key words: long memory; fractional differencing; ARFIMA; FIGARCH; squared returns; absolute returns; Taylor effect; range.

* Acknowledgements: I am grateful to Christopher Gilbert for helpful comments and to Andrew Harvey, Marten Sola and Stephen Taylor for useful suggestions and to the London International Financial Futures Exchange for provision of the data. All errors remain my responsibility.

1. Introduction

The phenomenon of long memory has been known since the time that ancient Egyptian hydrologists studied the flows and inflows of the river Nile.¹ The idea is very simple and states that the effects of an event (shock) persist over a long period of time. It is well documented that some economic and financial time series exhibit long range dependence and many authors have considered the issue of a long memory component in economic data.² The so-called long memory property is usually defined in relation to the autocorrelations of the process by requiring that the dependence between distant observations, although small, be significantly different from zero. Technically, a long memory process is characterised by a fractional degree of integration - that is to say the degree of integration is less than one but greater than zero - and exhibits hyperbolically decaying autocorrelations.³ In other words, the autocorrelation function shows persistence that is consistent with neither an I(0) process nor an I(1) process.

It is well known that asset prices returns contain little serial correlation, in accord with the efficient markets hypothesis; however, their volatilities exhibit a much richer structure. There is a lot of evidence showing that conditional volatility of returns on asset prices displays long memory or long range dependence.⁴ On this regard, particularly important are the works of Andersen and Bollerslev (1997; 1998). Using the mixture of distribution hypothesis, they interpreted volatility as a combination of heterogeneous information arrivals. Although, each of the information flow process exhibits short memory, the volatility process is a long memory process. Therefore, they provided evidence that the long memory characteristic of the volatility process "... constitute an intrinsic feature of the returns generating process, rather than a manifestation of occasional structural shifts" (Andersen and Bollerslev (1997), page 975).⁵

The first goal of this paper is to verify that commodity price volatility follows a long memory process. For this purpose, we use both Generalised AutoRegressive Conditional Heteroscedasticity (GARCH) models (Bollerslev (1986), Taylor (1986)) and

¹ It is even mentioned in the Bible and is referred to as the Joseph Effect: "Seven years of great abundance are coming throughout the land of Egypt, but seven years of famine will follow them" (Genesis, 41, 28-29).

² For an extensive survey of the major econometric works on long memory processes see Baillie (1996).

³ For an extensive description of long memory processes see Beran (1994).

⁴ Among the others, see Ding *et al.* (1993), de Lima and Crato (1993), Psaradakis and Sola (1995) and Bollerslev and Mikkelsen (1996) for evidence of long memory in stock market volatility.

⁵ See Granger and Ding (1996).

Fractionally Integrated GARCH (FIGARCH) models (Baillie *et al.* (1996)). We argue that if a time series is characterised by a long memory behaviour, a good model must take account of such behaviour. In order to illustrate this proposition we use two types of models, the first model (GARCH) does not take account of the long memory effect, while the second (FIGARCH) is characterised by the presence of a long memory component.

Taylor (1986) found that absolute returns of speculative assets have significant positive serial correlation over long lags. He also found that autocorrelations of absolute returns are greater than that of squared returns. Ding *et al.* (1993), Granger and Ding (1995) and Ding and Granger (1996) also found evidence that the power transformation of the absolute returns, $|s_t|^a$, has high autocorrelation for long lags and that this property is strongest when $a = 1$.⁶ Granger and Ding (1995), using the results of Luce (1960), showed that the expected absolute return, and any power transformation of this return, may be interpreted as a measure of risk. There is therefore evidence that *the absolute return measure of risk* contains more serial correlation than *the squared return measure of risk*. They refer to this phenomenon as the “Taylor effect”. Following Parkinson (1980) and Garman and Klass (1980), we consider a third measure of risk which uses the highest and the lowest prices in a day: The daily range. We will often refer to this measure as intraday volatility. The second goal of this paper is to see which of these three measures of risk exhibits the strongest autocorrelation structure. The paper also models absolute returns, squared returns and intraday volatility using Fractionally Integrated AutoRegressive Moving Average (ARFIMA) models (Granger and Joyeux (1980) and Hosking (1981)). Using the ARFIMA models we propose to examine whether the three measures of risk considered in this paper exhibit long memory. There is large literature on both absolute returns and squared returns, but little attention has been paid to extreme values volatility measures. Parkinson (1980) and Garman and Klass (1980) showed that high/low prices contain more information than close prices. If intraday volatility will also exhibit a richer autocorrelation structure than squared and absolute returns, this may motivate consideration of this measure of risk in modelling financial market volatility.

The paper is organised as follows. Section 2 briefly defines long memory. Section 3 presents the ARFIMA and the FIGARCH models. Section 4 describes the data and

examines the empirical autocorrelations of the series. Section 5 contains and discusses the results from estimation of GARCH and FIGARCH models for commodity futures returns. Section 6 contains and discusses the ARFIMA estimates for three measures of risk - absolute returns, squared returns and ranges. Section 7 concludes.

2. Defining Long Memory

From an empirical point of view, the presence of long memory may be defined in terms of the observed autocorrelations, which show high dependence between very distant observations. The autocorrelations of a long memory process is consistent with an essentially stationary process but takes far longer to decay than the exponential rate associated with the stationary ARMA class of processes. Defining the autocorrelation between observation at time t and observation at time $t - j$ as ρ_j , long memory processes are characterised by the following property:⁷

$$\lim_{t \rightarrow \infty} \sum_{j=-t}^t |\rho_j| = \infty. \quad (1)$$

Fractionally integrated processes are long memory processes given the property in (1).

It is possible to define the fractionally integrated process y_t as follows,

$$(1-L)^d y_t = u_t \quad (2)$$

where d is not integer and represents the fractional order of integration and L is the lag operator. For values of d less than $1/2$ and positive, the process y_t is long memory, and its autocorrelations are all positive and exhibit a hyperbolic rate of decay. For $-0.5 < d < 0$ the process has short memory according to (1). It is possible to note from (2) that fractionally integrated processes are intermediate between $I(0)$ and $I(1)$ processes. Following Granger and Joyeux (1980) and Hosking (1981), we may rewrite (2) as follows,

⁶ Ding and Granger (1996) used various speculative returns and found that for foreign exchange rate returns, this property is strongest when $a = 1/4$.

⁷ Note that the long memory property might be defined also in terms of the spectral density (see Beran (1994)). An alternative definition of long memory is in terms of Wold decomposition. For a survey see Baillie (1996).

$$(1-L)^d(y_t - \mu) = \varepsilon_t \quad (3)$$

where μ is the mean of the process y_t , $E(\varepsilon_t) = 0$, $E(\varepsilon_t^2) = \sigma^2$ and $E(\varepsilon_t \varepsilon_s) = 0$ for $s \neq t$. Equation (3) defines a fractional white noise process. The fractional difference operator $(1-L)^d$ is defined as follows,

$$(1-L)^d = \sum_{j=0}^{\infty} \frac{\Gamma(j-d)L^j}{\Gamma(j+1)\Gamma(-d)} \quad (4)$$

where $\Gamma(\cdot)$ is the standard gamma function.

The autocorrelation function of (3) at lag j is equal to

$$\rho_j = \frac{\Gamma(j+d)\Gamma(1-d)}{\Gamma(j-d)\Gamma(d)} . \quad (5)$$

The asymptotic approximation of Equation (5) is given by

$$\begin{aligned} \rho_j &\approx c j^{2d-1} \\ \text{where } c &= \frac{\Gamma(1-d)}{\Gamma(d)} . \end{aligned} \quad (5a)$$

Hence the autocorrelation coefficients exhibit slow hyperbolic decay for large j .

3. ARFIMA and FIGARCH Models

In order to introduce the ARFIMA models we start from the stable ARMA class of processes. Following Box and Jenkins (1976) we define an AutoRegressive Moving Average, ARMA (p,q), model as follows:

$$\phi(L)(y_t - \mu) = \theta(L)\varepsilon_t \quad (6)$$

where $\phi(L) = 1 - \sum_{j=1}^p \phi_j L^j$, $\theta(L) = 1 + \sum_{j=1}^q \theta_j L^j$, μ is the mean of y_t , y_t is an I(0)

process and ε_t is white noise. If y_t is I(1), Equation (6) holds after differentiation,

$$\phi(L)(1-L)(y_t - \mu) = \theta(L)\varepsilon_t . \quad (7)$$

Equation (7) represents an ARIMA (p,1,q) model.

Some economic data exhibit an autocorrelation structure which is between those of I(0) and I(1) processes. In other words, the autocorrelations of the original series have the

appearance of being non stationary, while the differenced series appears over-differenced. This consideration led Granger and Joyeux (1980) and Hosking (1981) separately to the formulation of the fractionally integrated ARMA processes, or ARFIMA (p,d,q),

$$\phi(L)(1-L)^d(y_t - \mu) = \theta(L)\varepsilon_t \quad (8)$$

where d is the fractionally differenced parameter. Granger and Joyeux (1980) and Hosking (1981) showed that the autocorrelation coefficients of an ARFIMA model exhibit a slow hyperbolic rate of decay. For $d < 1$ the process is mean reverting and is covariance stationary for $-0.5 < d < 0.5$. The process exhibits long memory for $0 < d < 0.5$. From Equation (8), it is evident that ARFIMA specification reduces to a stable ARMA process when $d = 0$ and to an ARIMA model for $d = 1$. ARFIMA class of models is very flexible and captures both long and short memory components of a process. In fact, the parameter d accounts for the long memory component, while the $\phi(L)$ and the $\theta(L)$ polynomials capture the short run dynamic.

Sowell (1992) derives the exact MLE of the ARFIMA process assuming a normal distribution for the error term. However, Sowell's (1992) log-likelihood is computationally demanding. Chung and Baillie (1993) proposed an approximate MLE based on the minimisation of the Conditional Sum of Square (CSS) estimator given by,

$$S = \frac{1}{2} \log(\sigma^2) + \frac{1}{2} \sigma^2 \sum_{t=1}^T [\phi(L)\theta(L)^{-1}(1-L)^d(y_t - \mu)] \quad (9)$$

where T is the data sample size (number of observations). In the estimates reported in Section 6 we use the CSS estimator.

The same issues that arise in modelling long run dependencies in the first moment of a process, become relevant also for the second moment. Volatility has been modelled by the GARCH class of models (Bollerslev (1986)). GARCH models are widely used and are able to capture one of the main characteristics in observed data, volatility clustering.

We may define the GARCH (p,q) process as follows,

$$\begin{aligned} s_t &= \mu_t + h_t^{1/2} \varepsilon_t \\ \varepsilon_t &\sim N(0,1) \\ h_t &= \omega + \alpha(L)\varepsilon_t^2 + \beta(L)h_t \end{aligned} \quad (10)$$

where s_t is the asset return at time t , μ_t could be any regression function for the conditional mean of the process, $\alpha(L)$ and $\beta(L)$ are lag polynomials of order q and p and ε_t is i.i.d. The process is covariance stationary if $\sum_{j=1}^q \alpha_j + \sum_{j=1}^p \beta_j < 1$.⁸ Restrictions need to

be imposed on the parameters of the conditional variance equation to ensure that the conditional variance is strictly positive.⁹ Ding and Granger (1996) provided an analytical expression for the autocorrelation function of GARCH models and showed that it decreases exponentially. Defining the innovations in the conditional variance process as

$$v_t = \varepsilon_t^2 - h_t \quad (11)$$

the GARCH(p,q) model in (10) might be expressed as an ARMA(m,p) process in ε_t^2 ,

$$\phi(L)\varepsilon_t^2 = \omega + [1 - \beta(L)]v_t \quad (12)$$

where m is equal to $\max(p,q)$ and $\phi(L) = [1 - \alpha(L) - \beta(L)]$. When $\phi(L)$ in (12) contains a unit root, the GARCH process becomes an Integrated GARCH, or IGARCH, (Engle and Bollerslev (1986)),

$$\phi(L)(1-L)\varepsilon_t^2 = \omega + [1 - \beta(L)]v_t. \quad (13)$$

We might interpret Equation (13) as the ARIMA representation of the IGARCH process.

Baillie, Bollerslev and Mikkelsen (1996) introduced long memory in the conditional variance of a GARCH model and proposed the Fractionally Integrated GARCH, or FIGARCH (p,d,q) model, where the conditional variance is specified as follows,¹⁰

$$\phi(L)(1-L)^d \varepsilon_t^2 = \omega + [1 - \beta(L)]v_t. \quad (14)$$

The FIGARCH model in (14) reduces to a GARCH model when $d = 0$ (Equation (12)) and to an IGARCH model when $d = 1$ (Equation (13)). The FIGARCH model imposes an ARFIMA structure on ε_t^2 . The similarities between GARCH, IGARCH and

⁸ The stationary conditions reported here rely on the underlying assumption of absence of autocorrelation. Bera, Higgins and Lee (1990) demonstrated that the presence of autocorrelation leads to a different stationary condition.

⁹ Nelson and Cao (1992) showed that not all the parameters of the conditional variance need to be positive to guarantee a strictly positiveness of the variance.

¹⁰ Independent research by Ding and Granger (1996) leads to a closely related model. Analysing the sample autocorrelation of various asset returns, Ding and Granger (1996) noted that there are different volatility components. Assuming an infinite number of volatility components, they proposed a model which is similar to the FIGARCH specification. Bollerslev and Mikkelsen (1996) extended the FIGARCH process to a log transformation of the conditional variance and proposed the Fractionally Integrated Exponential GARCH (see Nelson (1991)).

FIGARCH models for the conditional variance and, respectively, ARMA, ARIMA and ARFIMA processes for the conditional mean are obvious.

We may rewrite the FIGARCH (p,d,q) model in (14) as follows,

$$[1 - \beta(L)]h_t = \omega + [1 - \beta(L) - \phi(L)(1 - L)^d] \varepsilon_t^2 \quad (15)$$

or equivalently,

$$h_t = \frac{\omega}{1 - \beta(L)} + \lambda(L) \varepsilon_t^2 \quad (16)$$

where $\lambda(L) = 1 - \frac{(1 - \phi(L))(1 - L)^d}{1 - \beta(L)}$.

The FIGARCH process in (16) implies a hyperbolic rate of decay for the lagged squared innovations, which is a characteristic of long memory processes. General conditions for the conditional variance to be positive are difficult to establish. However, Bollerslev and Mikkelsen (1996) derive sufficient conditions for the case of a FIGARCH(1,d,1) as $\beta - d \leq \phi \leq \frac{1}{3}(2 - d)$ and $d \left(\phi - \frac{1}{2}(1 - d) \right) \leq \beta(\phi - \beta + d)$. Under assumption of normally distributed error terms, the parameters of the FIGARCH (p,d,q) process can be obtained by maximising the (quasi) Maximum Likelihood given by,

$$\text{Loglik}(\theta, \varepsilon_t) = -\frac{T \log(2\pi)}{2} - \frac{1}{2} \sum_{t=1}^T \left(\log(h_t) + \frac{\varepsilon_t^2}{h_t} \right) \quad (17)$$

where T is the number of observations used for the estimation. As for the ARFIMA class of models for the conditional mean of a process, the FIGARCH(p,d,q) model captures both long memory effect (the d parameter) and short run effects ($\phi(L)$ and $\beta(L)$ polynomials).

4. The Data

We analyse futures contracts on cocoa, coffee and sugar from the London International Futures Financial Exchange (LIFFE). We choose to analyse commodity markets because these markets have seen an increasing investor interest over recent years. The data set analysed consists of daily observations from 3 January 1989 for cocoa and sugar, and 2 January 1988 for coffee, to 31 December 1997, a total of 2274

observations for cocoa and sugar, and 2527 observations for coffee. This period includes 1994, a year which saw substantial portfolio investment in commodities. For each commodity we consider the second delivery contract.¹¹ Each contract starts trading well before the delivery month. A futures contract is in the second delivery when is preceded by the contract closest to maturity. On 30 January 1996 the cocoa May 1996 contract is in the second position and the March 1996 contract is in the first position. The latter is the contract nearest to maturity. We did not consider the first nearby contract (March 1996, in the previous example) because it might contain delivery distortion. Moreover, in the markets analysed, liquidity concentrate on the second delivery contracts. Figures 1 and 2 graph, respectively, the monthly average volume and open interest for the first, the second and the third delivery contracts for the cocoa market. The second delivery contract shows both higher volume and higher open interest when compared with the first and the third delivery contracts.¹² For each commodity we compute returns, absolute returns, squared returns and intraday volatility.

Following standard practice, we define returns as $s_t = \ln(F_t / F_{t-1})$, where F_t is the futures price at day t , absolute returns as $|s_t|$ and squared returns as s_t^2 . Daily returns over rolling dates are computed by using the same future contract. For example, on 1 August 1997 the cocoa December 1997 contract enters the second delivery; the return for that day is calculated by using the 31 July 1997 price for the December 1997 contract. For intraday volatility we use the following formula,

$$\sigma_t^{Range} = \frac{\ln(H_t) - \ln(L_t)}{2 \ln(2)} \quad (18)$$

where H_t and L_t are, respectively, the highest and the lowest prices at day t . This measure expresses the variability of price changes in the course of one day. Equation (18) estimates the innovation standard deviation under the assumption that the logarithm of the asset price follows a Brownian motion with zero drift. Parkinson (1980) and Garman and Klass (1980) argue that high/low prices contain more information regarding volatility than close prices. They show that extreme-value estimators are much more efficient than close-close estimator (squared returns). However, Garman and Klass (1980) recognise that discrete-time trading gives rise to a downward bias to ranges

¹¹ The LIFFE futures contracts on sugar and cocoa are denominated in U.S. Dollars and U.K. Sterling respectively. The futures contracts on coffee are denominated in U.S. Dollars since March 1991, before that date they were quoted in U.K. Sterling.

volatility estimates.¹³ The bias is inversely related to trading volumes. The commodities we are analysing have high volumes. Therefore, the bias should be negligible. Beckers (1983) showed that high and low prices contain information unavailable in closing prices and this information is particularly useful in forecasting future variance. Taylor (1987) demonstrated that forecast computed from daily ranges are empirically better than forecast calculated from daily closing prices. Wiggins (1991) proved that Parkinson's estimator of variance is significantly more efficient than the close-close estimator; however, he underlines the fact that recording errors might reverse this result.¹⁴

We first check for seasonality in volatility, which may be due to the harvest cycle. Figure 3 graphs the monthly average ranges for the three commodities analysed in this paper. There is no clear evidence of a seasonal pattern in the volatility figures. The London commodity markets have traditionally been primarily oriented towards industry (producers and consumers) and trade (dealers); but recently new operators (banks, funds, etc.) have entered in these markets. Therefore portfolio demand for these assets has increased.¹⁵ At the beginning of our sample, volatility in the cocoa and the sugar market moderately increases during the months June/July (harvest period), but this feature disappears in the later years when portfolio demand in these markets increased. The coffee market appears the most volatile all over the period and does not show any evidence of seasonality.

Table 1 reports summary statistics for the returns, squared returns, absolute returns and intraday volatility for cocoa, coffee and sugar. The numbers in parenthesis refer to the summary statistics of the log intraday volatility. In the first part of Table 1 we report mean, standard deviation, skewness, kurtosis and the standard χ^2 normality test.

Only the returns on cocoa exhibit a negative mean. Coffee has the highest standard deviation for all the series. The coffee market, in fact, was hit by the unprecedented coincidence of two severe frosts in a single crop year, 1994, which

¹² This is true also when the first contract is not in its delivery month. We obtain similar results for coffee and sugar.

¹³ The downward bias is due to the fact that true high and low prices sometimes remain unobserved. Marsh and Rosenfeld (1986) also show that non-trading can reduce the efficiency of extreme-value estimators, particularly for low price stocks.

¹⁴ In this regard the sugar data may be problematic: many observations on volume and open interest are missing, and there appear to be some mistakes in the recorded prices.

¹⁵ See Gilbert and Brunetti (1996).

resulted in high price movements.¹⁶ As expected, returns, squared returns and absolute returns are not normally distributed. The log range (numbers in parenthesis) shows evidence of normality for cocoa and coffee, suggesting that the rescaled range seems to be well fitted by a log-normal distribution. This is a very useful property for empirical application of this measure of risk.

In the second part of Table 1 we report the Ljung-Box Q statistics up to 10, 20 and 50 lags - denominated by Q(10), Q(20) and Q(50) respectively – the ADF and KPSS tests. In the Q test the null hypothesis implies independently and identically distributed data.¹⁷ Table 1 shows that the null of i.i.d. is only rejected for the returns of sugar and then only in the case of ten lags. For all the other series there is evidence of serial correlation. It is interesting to note that the values of the Q test for the range is the highest for all commodities.

The last two rows of Table 1 report the ADF and the KPSS tests. In the Augmented Dickey-Fuller (ADF) test (Dickey and Fuller (1981)) the null hypothesis is that of a unit root, $I(1)$, against an alternative hypothesis of $I(0)$. The KPSS test of Kwiatkowski, Phillips, Schmidt and Shin (1992), is designed to test an $I(0)$ null hypothesis versus an $I(1)$ alternative.

Table 1 shows that, for the return series for all commodities, the ADF test rejects the null of $I(1)$ and the KPSS fails to reject the null of $I(0)$. Hence there is evidence that the return series are $I(0)$. More interestingly, for squared returns, absolute returns and ranges Table 1 gives rejection of the $I(1)$ null using the ADF test and also rejection of the $I(0)$ null from the KPSS test. We might interpret these results as a possible evidence for long memory. In other words, if the squared returns, absolute returns and ranges are neither $I(1)$ nor $I(0)$, they may be $I(d)$. This issue will be further investigated in what follows.

¹⁶ See Gilbert and Brunetti (1996).

¹⁷ The Q test statistic assumes that the variance of the process is constant (homoskedasticity), so that the reported critical values are merely indicative. Diebold (1988), among the others, noted that the presence of ARCH may give rise to spurious significance of the portmanteau test. Nevertheless, Bollerslev and Mikkelsen (1996) showed that the Q test is still valid in detecting serial correlation.

4.1. The Sample Autocorrelations

Table 2 gives the sample autocorrelations for lags 1 to 5 and 10, 20 and 50. The numbers in the second row are twice the standard errors of the sample autocorrelation under the assumption that each series is independently and identically distributed and has finite variance. The final four rows of Table 2 show the averages of the sample autocorrelations from lags 1 to 10, 11 to 20, 21 to 50 and 51 to 100. For all the commodities, the returns series exhibit little autocorrelation; the first two lags autocorrelations, although small, are significantly different from zero, the only exception being lag 1 for sugar. The second lag autocorrelations for the return series are negative implying mean reversion.

Table 2 shows that the different commodities share some commonalities as well as they exhibit different behaviours. For example, squared returns and absolute returns for cocoa have much lower autocorrelations than those for the corresponding series for coffee and sugar.

A common result is that the sample autocorrelations of absolute returns are greater than the sample autocorrelations of squared returns.¹⁸ Taylor (1986) obtained the same results for the same commodities but for a different data sample.¹⁹ Granger and Ding (1995) refer to this property as the Taylor effect.²⁰ They also showed that absolute returns may be modelled using the conditional exponential distribution. This distribution has the property that the mean is equal to the standard deviation, a feature which is evident in our data (see Table 1). Moreover, the skewness and the kurtosis of the absolute returns of cocoa and sugar are very close to those of the exponential distribution – which are equal to two and nine respectively. We therefore tested whether absolute returns of cocoa, coffee and sugar have an exponential distribution.

¹⁸ The only exceptions are for coffee at lags 2, 3 and 5.

¹⁹ Taylor (1986) analysed many spot and futures prices of different financial assets. The data sample he used covers mainly the decade 1971 – 1980.

²⁰ Based on these findings, for each commodity, we further examined the sample autocorrelation of the transformed absolute returns $|s_t|^a$ for $a = 0.25, 0.50, 0.75, 1.25, 1.50, 1.75, 2$ for lags 1 to 100 and still found that $|s_t|^a$ have the largest autocorrelation for $a = 1$ or close to 1.

Comparing the empirical frequency distribution with the theoretical²¹ distribution, we easily accepted the null hypothesis of exponential distribution for absolute returns for cocoa and sugar, but not for coffee.²² From Table 2 it also appears that the range is the most autocorrelated of all the series considered in this paper. This result is not surprising. Taylor (1987) provides the theoretical autocorrelation of daily ranges²³ and proves that ranges are more autocorrelated than absolute returns. He also shows that both absolute returns and ranges exhibit the same persistence, in a forecast sense, when an ARMA(1,1) model is used for forecasting.

Figures 4, 5 and 6 show the correlograms for up to fifty lags for cocoa, coffee²⁴ and sugar respectively. The parallel lines show the 95% confidence interval for the estimated sample autocorrelation under the assumption that each series is an i.i.d. process. Focusing only on squared returns, absolute returns and ranges, it is seen that the latter exhibits higher autocorrelation structure than the other two. Moreover, the sample autocorrelation function for the range decreases very fast at the beginning and then very slowly. Interestingly, absolute returns and ranges for sugar show similar patterns and, for some lags, absolute returns have higher autocorrelations than ranges.

The overall picture is that commodity data exhibit the so-called Taylor effect - findings that are in line with those of Taylor (1986), Ding *et al.* (1993), Granger and Ding (1995) and Ding and Granger (1996). Moreover, the intraday volatility measure of risk shows characteristics of long memory more than the squared returns and absolute returns measures of risk. The ranking in terms of autocorrelation – from the lowest to the highest – is: returns, squared returns, absolute returns, ranges.

²¹ We divided the observations in ten bins of equal width. As mean of the theoretical distribution we used the sample mean.

²² The χ^2 test for cocoa, sugar and coffee is equal to 13.86, 13.97 and 226.04 respectively - the 95% $\chi^2(9)$ critical value is 16.9.

²³ See also Parkinson (1980).

²⁴ For coffee we replace the returns for 27 June 1994 and 11 July 1994, the days of the two frosts, with the averages of the returns for the preceding and the following days, on the basis that the large price movements those days were a discrete jump and had no implications for future prices.

5. GARCH and FIGARCH Model Estimations

In applied work, it has been frequently demonstrated that the GARCH (1,1) process is able to represent the majority of financial time series (Bera and Higgins (1993)). Following this practise we focused on GARCH (1,1) and FIGARCH (1,d,1) models. The mean specification varies across the different commodities. For cocoa the mean equation has three autoregressive components at lags one, six and ten, while for coffee the autoregressive components are at lags one and three. The conditional mean equation for sugar involves an AR(2) and an AR(4) parameters. For the estimation of the FIGARCH models we followed the approach proposed by Ballie *et al.* (1996). They truncated the infinite ARCH polynomial in Equation (16) at lag one thousand and fixed the pre-sample values of ε_t^2 , $t = 0, -1, -2, -1000$, at the unconditional sample variance. This approach is not innocuous. In fact, for $d > 0$, the population variance does not exist. However, Ballie *et al.* (1996) show that this approach does not affect the asymptotic distributions of the estimators. Robust standard errors have been computed following the procedures proposed by Bollerslev and Wooldridge (1992).

For each commodity we estimated two models: GARCH(1,1) and FIGARCH(1,1). In order to choose among these models – i.e. which model best describes the data – we used the Akaike Information Criterion (AIC) and the Schwarz Information Criterion (SIC).²⁵ Monte Carlo simulations show that these criteria may be effectively used in discriminating between GARCH and FIGARCH alternatives (Bollerslev and Mikkelsen (1996)).

Table 3 shows the estimation results.

One may note that there is evidence of low order autocorrelation in the daily returns. Estimation of the long memory parameter in the FIGARCH model may be biased if the model for the mean returns is misspecified. This motivates our effort to find the proper specification for the conditional mean equation.

²⁵ The AIC and SIC to be minimised are defined as follows,

$$AIC = -2\log L(\hat{\theta}) + 2\gamma$$

$$SIC = -2\log L(\hat{\theta}) + \gamma \log(T)$$

where γ denotes the number of estimated parameters, $L(\hat{\theta})$ is the value of the Log-likelihood and T is the number of observations used for the estimates.

Table 3
Quasi Maximum Likelihood Estimates for GARCH(1,1) and
FIGARCH(1,d,1) for Cocoa, Coffee and Sugar

	Cocoa δ_1 : AR(1) δ_2 : AR(6) δ_3 : AR(10)		Coffee δ_1 : AR(1) δ_2 : AR(3)		Sugar δ_1 : AR(2) δ_{12} : AR(4)	
	GARCH	FIGARCH	GARCH	FIGARCH	GARCH	FIGARCH
i (%)	-0.054 (0.029)	-0.047 (0.029)	-0.022 (0.035)	-0.019 (0.034)	0.023 (0.019)	0.027 (0.019)
δ_1	0.059 (0.022)	0.056 (0.024)	0.064 (0.022)	0.064 (0.023)	-0.053 (0.022)	-0.053 (0.022)
δ_2	-0.061 (0.022)	-0.064 (0.022)	0.050 (0.021)	0.046 (0.021)	0.034 (0.022)	0.039 (0.022)
δ_3	0.062 (0.022)	0.063 (0.022)				
ω	0.040 (0.013)	0.097 (0.033)	0.118 (0.026)	0.157 (0.053)	0.008 (0.003)	0.022 (0.009)
α_1	0.044 (0.008)		0.074 (0.009)		0.042 (0.006)	
β_1	0.938 (0.012)	0.690 (0.062)	0.898 (0.013)	0.646 (0.075)	0.952 (0.073)	0.723 (0.052)
ϕ_1		0.452 (0.063)		0.400 (0.071)		0.489 (0.061)
d		0.332 (0.062)		0.390 (0.062)		0.355 (0.046)
AIC	8010	7996	10299	10290	6421	6400
SIC	8057	8043	10347	10345	6468	6454

Robust standard errors in parentheses.
Returns series were rescaled before estimation, so that $s_t = s_t * 100$.

The sum of the estimates of α_1 and β_1 in the GARCH model is very close to one for all three commodities, indicating that the volatility process is highly persistent. In particular, the estimates of β_1 in the GARCH model is very high, indicating a strong autoregressive component in the conditional variance process. Focusing now on the FIGARCH results, the fractional differencing parameter, d , is estimated as significantly different from zero, implying fractional integration.

Moreover, the estimated values of d are similar across the different commodities. The conditions for the conditional variance to remain positive are always satisfied (see Section 3). The β_1 estimate falls considerably as one moves from GARCH to FIGARCH, in line with the findings of Baillie, Bollerslev and Mikkelsen (1996). They claim that, in the presence of long memory, there is an upward bias in the GARCH estimates due to the fact that the GARCH model does not take into account the long

memory component of the volatility process. Moreover, according to the AIC and SIC criteria, the FIGARCH models fit the data series better than the GARCH models.²⁶

Table 4 provides diagnostics to assess the adequacy of the estimated models. In particular, Table 4 records the skewness and the kurtosis of the standardised residuals of each model and the Q test up to 10 and 20 lags for both standardised residuals and squared standardised residuals – denoted by Q(10), Q(20), Q²(10) and Q²(20).

Table 4 Models Diagnostics						
	Cocoa		Coffee		Sugar	
	GARCH	FIGARCH	GARCH	FIGARCH	GARCH	FIGARCH
Skewness	0.191	0.182	0.080	0.185	0.000	-0.020
Kurtosis	5.501	5.458	7.385	7.592	6.482	6.231
Q(10)	10.898	11.541	12.079	12.135	10.843	9.849
Q(20)	22.951	24.317	16.750	16.270	22.253	22.117
Q ² (10)	7.733	6.863	6.224	5.050	9.080	5.741
Q ² (20)	19.448	17.247	47.508 (*)	68.684 (*)	13.252	10.415

Skewness, Kurtosis and Q(10) refer to standardise residuals.
Q²(10) refers to squared standardised residuals.
(*) indicates rejection of the null of i.i.d.

The standardised residuals exhibit less skewness and kurtosis than the returns (see Table 1).²⁷ Perhaps of greater importance, the Q statistics²⁸ fails to reject the null hypothesis of independently and identically distributed standardised residuals and squared standardise residuals, the only exception being the Q²(20) for coffee.

In the FIGARCH model, the long memory parameter is applied to the squared error term. Taking out the mean and the autoregressive component of the return series, the squared error term coincides with the squared return. Hence, the FIGARCH model estimates provided evidence that the squared returns exhibit long memory. In the

²⁶ For $\omega > 0$ and $0 < d \leq 1$, the second moment of the unconditional distribution of the error terms does not exist (it is infinite) (see. Baillie *et al* (1996)). Therefore, caution should be used in interpreting Tables 1 and 2.

²⁷ The only exception is skewness for cocoa.

²⁸ Following Diebold (1988) and Bollerslev and Mikkelsen (1996), the critical values for the Q-test are drawn from a χ^2 distribution with degrees of freedom adjusted by subtracting the number of estimated parameters from the number of autocorrelations being tested equal to zero.

following section this issue will be further analysed together with the issue of which of the risk measures considered in this paper shows evidence of higher temporal dependence.

6. ARFIMA Estimates

The paper considers, as already said, three measures of risk: squared returns, absolute returns and ranges. In Section 4 we analysed the autocorrelations structure of these series and found evidence of long memory. To further check this result we modelled the three measures of risk using the ARFIMA class of models. There are many ways of testing for the presence of long memory. In particular, it is possible to distinguish between parametric and semiparametric tests. We disregard the latter and consider only parametric test because “Despite the amount of theoretical work in attempting to derive robust semiparametric estimators of long memory parameter, there is substantial evidence documenting their poor performance in terms of bias and mean squared error.” Baillie (1996, page 35). This choice is supported by the consideration that ARFIMA models describe the long-run dynamic of the conditional mean in the same way in which FIGARCH class of models does that with the conditional variance.

For each commodity and for each measure of risk, we estimated the following ARFIMA(1,d,1) model using the Conditional Sum of Square (CSS) estimator described in Equation (9),

$$(1 - \phi_1 L)(1 - L)^d (y_t - \mu) = (1 + \theta_1 L)\varepsilon_t. \quad (19)$$

Table 5 shows the estimation results.

All coefficients in Table 5 are significantly different from zero. For all commodities and for all series analysed, there is a positive autoregressive component (ϕ_1) and a negative moving average component (θ_1). More importantly, the fractional differencing parameter is always significant and positive and lies in the range 0.273 - 0.417, implying both stationarity and long memory.

Table 5 Approximate Maximum Likelihood Estimates – ARFIMA for Squared Returns, Absolute Returns and Ranges Cocoa, Coffee and Sugar									
	Cocoa			Coffee			Sugar		
	Sq. Ret.	Abs. Ret.	Range	Sq. Ret.	Abs. Ret.	Range	Sq. Ret.	Abs. Ret.	Range
ϕ_1	0.500 (0.061)	0.415 (0.056)	0.366 (0.067)	0.373 (0.065)	0.180 (0.015)	0.182 (0.035)	0.388 (0.066)	0.386 (0.063)	0.680 (0.169)
θ_1	-0.704 (0.063)	-0.674 (0.068)	-0.602 (0.081)	-0.613 (0.089)	-0.351 (0.177)	-0.259 (0.037)	-0.604 (0.075)	-0.628 (0.073)	-0.114 (0.020)
d	0.273 (0.063)	0.323 (0.062)	0.417 (0.052)	0.285 (0.062)	0.281 (0.045)	0.348 (0.042)	0.298 (0.050)	0.348 (0.051)	0.304 (0.037)

Standard errors in parentheses.

These results are in line with those of the FIGARCH estimates reported in Table 3.²⁹ With the exception of sugar, the range exhibits the highest value for the long memory parameter. Absolute returns show a greater value for d than the squared returns except for the case of coffee where, however, the two parameters are very close (0.281 and 0.285).

Having found that the three measures of risk considered in this paper all exhibit long memory, this raises the question of which of these measures shows higher temporal dependence. For these purpose we consider the theoretical autocorrelations for a fractionally white noise process given by Equation (5a). We substitute in that expression the estimated values of d from the ARFIMA models. We assume $j = 50, 100, 1000, 2000$. Table 6 reports the results of this simple exercise. The numbers in the second row are twice the standard errors of the sample autocorrelation.

The value of the theoretical autocorrelation increases with the value of d , which is an obvious conclusion. Comparing squared returns and absolute returns for cocoa and sugar, Table 6 clearly shows the Taylor effect. The higher estimated values of the long memory component in absolute returns than in square returns, impart a stronger structure in the theoretical autocorrelations of the former than of the latter. Not surprisingly, in the case of coffee, the theoretical autocorrelations of absolute returns and squared returns are very similar due to the fact that the estimated values of d are almost the same. Focusing

²⁹ Given the similarities between squared returns and squared error terms, we were expecting a value of d , from the ARFIMA estimates for the squared returns to be similar to that of FIGARCH estimates. However, the differences between these two values might be due to the specification of the mean equation of the FIGARCH model.

on ranges and absolute returns, it is evident from Table 6 that the former shows much more dependence than the latter for cocoa and coffee, the reverse is true for sugar.

Table 6									
Theoretical Autocorrelations for Squared Returns, Absolute Returns and Range for Estimated Values of d Cocoa, Coffee and Sugar									
	Cocoa			Coffee			Sugar		
	0.042			0.040			0.042		
	Sq. Ret.	Abs. Ret.	Range	Sq. Ret.	Abs. Ret.	Range	Sq. Ret.	Abs. Ret.	Range
D	0.273	0.323	0.417	0.285	0.281	0.348	0.298	0.348	0.304
ρ_{50}	0.064	0.121	0.376	0.075	0.071	0.164	0.088	0.164	0.095
ρ_{100}	0.047	0.095	0.335	0.056	0.053	0.133	0.067	0.133	0.073
ρ_{1000}	0.017	0.042	0.229	0.021	0.019	0.066	0.026	0.066	0.029
ρ_{2000}	0.012	0.033	0.204	0.015	0.014	0.053	0.020	0.053	0.022

$\rho_j = c_j^{2d-1}$; $c = [\Gamma(1-d)]/[\Gamma(d)]$; $j = 50, 100, 1000, 2000$.

7. Conclusions

In this paper we have analysed commodity futures prices. In particular we focused on cocoa, coffee and sugar futures contracts quoted at the London International Futures Financial Exchange. The returns were modelled through GARCH and FIGARCH models. We found evidence supporting the FIGARCH models, in the sense that FIGARCH models fit the data series better than the GARCH models. FIGARCH specification, using the tools of long memory processes, is able to capture both long and short run characteristics of the volatility process. Estimates of the fractional degree of integration parameter were found to be significantly different from zero, indicating that commodity futures price volatility is a long memory process, thus rejecting the GARCH specification.

For each commodity, the paper also considers three measures of risk, squared returns, absolute returns and scaled ranges. The sample autocorrelations and the ARFIMA estimates confirm that the three series exhibit long memory. Absolute returns are found to be much more autocorrelated than squared returns. This provides evidence of the so-

called Taylor effect. The range has very high dependence between distant observations and shows the highest autocorrelation.

Many models have been proposed in the literature to accommodate the findings that asset returns show little serial correlation but absolute returns and their power transformations all have a strong autocorrelation structure. All these models consider either squared returns - Bollerslev (1986), Baillie, Bollerslev and Mikkelsen (1996) – or absolute returns and their power transformations – Taylor (1986), Ding, Granger and Engle (1993), Ding and Granger (1996). This paper suggests that, at least for commodity data, the range exhibits a much richer structure than squared returns and absolute returns. Moreover, there is evidence that the range is well approximated by a log-normal distribution. These findings motivate future consideration of the range as volatility measure.

REFERENCES

- Andersen, T. G., and T., Bollerslev (1997), “Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns”, *Journal of Finance*, 3, 975 – 1005.
- Andersen, T. G., and T., Bollerslev (1998), “Deutsche Mark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies”, *Journal of Finance*, 219 – 265.
- Baillie, R. T. (1996), “Long Memory Processes and Fractional Integration in Econometrics”, *Journal of Econometrics*, 73, 5 - 59.
- Baillie, R. T., Chung, C. F., and M. A., Tieslau (1996), “Analysing Inflation by the Fractionally Integrated ARFIMA-GARCH Model”, *Journal of Applied Econometrics*, 11, 23 - 40.
- Baillie, R. T., Bollerslev, T., and H. O. A., Mikkelsen (1996), “Fractionally Integrated Generalised Autoregressive Conditional Heteroskedasticity”, *Journal of Econometrics*, 74, 3 - 30.
- Beckers, S. (1983), “Variances of Security Price Returns Based on High, Low, and Closing Prices”, *Journal of Business*, 56, 97 - 112.
- Bera, A. K., Higgins, M. L., and S., Lee (1990), “On the Formulation of a General Structure for Conditional Heteroskedasticity”, Mimeo.
- Bera, A. K., and M. L., Higgins (1993), “ARCH Models: Properties, Estimation and Testing”, *Journal of Economic Survey*, 7, 305 - 366.
- Beran, J. (1994), *Statistics for Long Memory Processes*, Chapman & Hall, London.
- Bollerslev, T. (1986), “Generalised Autoregressive Conditional Heteroskedasticity”, *Journal of Econometrics*, 31, 307 - 327.
- Bollerslev, T., Chou, R. Y., and K. F., Kroner (1992), “ARCH Modelling in Finance: A Review of the Theory and Empirical Evidence”, *Journal of Econometrics*, 52, 5 - 59.

- Bollerslev, T., and J. M., Wooldridge (1992), "Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time-Varying Covariances", *Econometric Reviews*, 11, 143 - 172.
- Bollerslev, T., Engle, R. F., and D. B., Nelson (1994), "ARCH Models", in Engle, R. F., and D. L., McFadden, *Handbook of Econometrics*, Elsevier, vol. IV, 2959 - 3038.
- Bollerslev, T., and H. O. A., Mikkelsen (1996), "Modelling and Pricing Long Memory in Stock Market Volatility", *Journal of Econometrics*, 73, 151 - 184.
- Box, G. E. P., and G. M., Jenkins (1976), *Time Series Analysis: Forecasting and Control*, rev. Ed. San Francisco: Holden-Day.
- Chung, C. F. (1994), "A Note on Calculating the Autocovariances of Fractionally Integrated ARMA Models", *Economics Letters*, 45, 293 - 297.
- Chung, C. F., and R. T., Baillie (1993), "Small Sample Bias in Conditional Sum of Squares Estimators of Fractionally Integrated ARMA Models", *Empirical Economics*, 18, 791 - 806.
- de Lima, P. J. F., and N., Crato (1993), "Long-Range Dependence in the Conditional Variance of Stock Returns", August 1993 Joint Statistical Meeting, San Francisco. Proceedings of the Business and Economic Statistics Section.
- Dickey, D. A., and W. A., Fuller (1981), "Likelihood Ratio Statistics for Autoregressive Time series with a Unit Root", *Econometrica*, 49, 1057 - 1072.
- Diebold, F. X. (1988), *Empirical Modelling of Exchange Rate Dynamics*, Springer-Verlag, New York.
- Ding, Z., Granger, C. W. J., and R. E., Engle (1993), "A Long Memory Property of Stock Market Returns and a New Model", *Journal of Empirical Finance*, 1, 83 - 106.
- Ding, Z., and C. W. J., Granger (1996), "Modelling Volatility Persistence of Speculative Returns: A New Approach", *Journal of Econometrics*, 73, 185 - 215.
- Engle, R. F., and T., Bollerslev (1986), "Modelling the Persistence of Conditional Variances", *Econometric Reviews*, 5, 1 - 50.
- Engle, R. F., Ito, T., and W. L., Lin (1990), "Meteor Showers or Heat Waves? Heteroskedastic Intra-Daily Volatility in the Foreign Exchange Market", *Econometrica*, 58, 525 - 542.
- Gallant, A. R., Rossi, P. E., and G., Tauchen (1993), "Nonlinear Dynamic Structures", *Econometrica*, 61, 871 - 907.
- Garman, M. B., and M. J., Klass (1980), "On the Estimation of Security Price Volatility from Historical Data", *Journal of Business*, 53, 67 - 78.
- Gilbert, C. L., and C., Brunetti (1996), "Commodity Price Volatility in the Nineties", in Smith, K.T., and P., Kennison eds, *Commodity Derivatives and Finance*, Euromoney Books, London, 1 - 22.
- Granger, C. W. J. (1980), "Long-Memory Relationships and the Aggregation of Dynamic Models", *Journal of Econometrics*, 14, 227 - 238.
- Granger, C. W. J., and R., Joyeux (1980), "An Introduction to Long-Memory Time Series Models and Fractional Differencing", *Journal of Time Series Analysis*, 1, 15 - 39.
- Granger, C. W. J., and Z., Ding (1995), "Some Properties of Absolute Return; An Alternative Measure of Risk", *Annales D'Economie et de Statistique*, 40, 67 - 91.
- Granger, C. W. J., and Z., Ding (1996), "Varieties of Long Memory Models", *Journal of Econometrics*, 73, 61 - 77.

- Hosking, J. R. M. (1981), "Fractional Differencing", *Biometrika*, 68, 165 - 176.
- Kwiatkowsky, D., Phillips, P. C. B., Schmidt, P., and Y., Shin (1992), "Testing the Null Hypotheses of Stationary against the Alternative of a Unit Root: How Sure Are We That Economic Time Series Have a Unit Root?", *Journal of Econometrics*, 54, 159 - 178.
- Lo, A. W. (1991) "Long Term Memory in Stock Market Prices" *Econometrica*, 59, 1279 - 1313.
- Luce, R. D. (1960), "Several Possible Measures of Risk", *Theory and Decision*, 12, 217 - 228.
- Marsh, T. A., and E. R., Rosenfeld (1986), "Non-Trading, Market Making, and Estimates of Stock Price Volatility", *Journal of Financial Economics*, 15, 359 - 372.
- Nelson, D. B. (1991), "Conditional Heteroskedasticity in Asset Returns: A New Approach", *Econometrica*, 59, 347 - 370.
- Nelson, D. B., and C. Q., Cao (1992), "Inequality Constraints in the Univariate GARCH Model", *Journal of Business and Economic Statistics*, 10, 229 - 235.
- Pagan, A. R. (1996), "The Econometrics of Financial Markets", *Journal of Empirical Finance*, 3, 15 - 102.
- Pagan, A. R., and H., Sabau (1992), "Consistency Tests for Heteroskedasticity and Risk Models", *Estudios Economicos*, 7, 3 - 30.
- Parkinson, M. (1980), "The Extreme Value Method for Estimating the Variance of the Rate of Return", *Journal of Business*, 53, 61 - 65.
- Psaradakis, Z., and M., Sola (1995), "Modelling Long Memory in Stock Market Volatility: A Fractionally Integrated Generalised ARCH Approach", Discussion Paper in Financial Economics No. FE 2/95, Department of Economics, Birkbeck College, University of London.
- Robinson, P. M. (1981), "Testing for Strong Serial Correlation and Dynamic Conditional Heteroskedasticity in Multiple Regression", *Journal of Econometrics*, 47, 67 - 84.
- Sowell, F. (1992), "Maximum Likelihood Estimation of Stationary Univariate Fractionally Integrated Time Series Models", *Journal of Econometrics*, 53, 165 - 188.
- Taylor, S. (1986), *Modelling Financial Time Series*, John Wiley & Sons, New York.
- Taylor, S. J. (1987), "Forecasting the Volatility of Currency Exchange Rates", *International Journal of Forecasting*, 3, 159 - 170.
- Wiggins, J. B. (1991), "Empirical tests of the Bias and Efficiency of the Extreme-Value Variance Estimator for Common Stocks", *Journal of Business*, 64, 417 - 432.