Spurious Long Memory in Commodity Futures:
Implications for Agribusiness Option Pricing
Selected Paper #174577 prepared for presentation
at the American Agricultural Economics Association Annual Meeting, Portland, OR, July 29-August 1, 2007

Gabriel J. Power[*][†]and Calum G. Turvey

Texas A&M University and Cornell University

Presented: July 31st, 2007

Abstract

Long memory, and more precisely fractionally integration, has been put forward as an explanation for the persistence of shocks in a number of economic time series data as well as to reconcile misleading findings of unit roots in data that should be stationary. Recent evidence suggests that long memory characterizes not commodity futures prices but rather price volatility (generally defined as L_p norms of price logreturns). One implication of long memory in volatility is the mispricing of options written on commodity futures, the consequence of which is that fractional Brownian motion should replace geometric Brownian motion as the building block for option pricing solutions. This paper asks whether findings of long memory in volatility might be spurious and caused either by fragile and inaccurate estimation methods and standard errors, by correlated short memory dynamics, or by alternative data generating processes proven to generate the illusion of long memory. We find that for nine out of eleven agricultural commodities for which futures contracts are traded, long memory is spurious but is not caused by the effect of short memory. Alternative explanations are addressed and implications for option pricing are highlighted.

JEL Classification Codes: C52, C53, G12, G13, Q13, Q14.

1

1 Introduction

The price of a financial option is frequently quoted in terms of its implied volatility. This is an unobserved parameter that solves, for a going price and set of observable characteristics, the famous and widely-used Black-Scholes-Merton formula (Black and Scholes 1973; Merton 1973). The measurement of volatility remains an active and diverse area of research in both academia and industry. A central concern is whether volatility rapidly or slowly recovers from shocks that affect its magnitude. In commodity markets, options are written on futures contracts. A number of papers found that futures prices appeared to be more persistent than they ought to be, a finding that challenged existing theories on the mechanics of commodity futures. More recently, the literature has concluded that persistence, or long memory, in commodity prices was spurious and explained by a combination of measurement error and potential long memory in the volatility of futures. As a result, it remains unclear whether the price volatility of commodity futures is characterized by long memory.

This paper addresses one set of causes and consequences of option pricing bias in commodity markets, namely long memory (persistence) in futures price volatility. The principal aim of this work is to determine whether empirical findings of long memory in commodity futures prices and volatility are spurious. The leading alternative explanations that are considered include the effect of correlated short memory dynamics generally measured as ARMA parameters and the presence of structural breaks or level shifts in the data (Smith, 2005; Banerjee and Urga, 2005; Perron, 2006).

The main finding of this paper is that apparent long memory in commodity futures price volatility is almost invariably spurious but is not caused by the effect of short memory dynamics. Rather, a better candidate would be a Markov-switching or stochastic break model, either of which could generate the appearance of long memory.

The paper takes the following steps to answer the question. A measure of volatility is constructed using the daily price range and based on Alizadeh, Brandt and Diebold (2002). While less accurate than the realized volatility computed from intra-day high-frequency tick data, this measure appears superior to the traditionally used log-return based volatility estimates. It seems the best choice given that intra-day data are seldom available for agricultural commodity futures. An estimator of the fractional difference long memory parameter d (for a process that is $I(d)$) based on a wavelet transform is used. Wavelets are ideally suited to distinguish short from long memory and also to detect the fractal signature of long memory. As a result, this estimator avoids the bias that affects certain other long memory estimators such as the frequently-used Geweke-Porter-Hudak (GPH) approach. It is also consistent and efficient, unlike the GPH

estimator and has better small-sample properties than traditional approximate Whittle-type estimators. Standard errors are computed from an analytical formula that incorporates the sample covariance of the short memory parameters to avoid being understated. Robustness checks include a separate estimation using only Wednesday observations (sampled weekly) to consider the presence of "day of the week effects" as well as a simultaneous, partially overlapping time series approach developed by Smith (2005) to avoid the effect of nonlinear dynamics introduced by splicing together futures data when a contract expires and the second-nearby maturity becomes the nearby. Simple Likelihood Ratio tests are computed to evaluate whether the long memory parameter is significant and the results are contrasted with the evidence from Wald and modified KPSS and Phillips-Perron tests designed to consider the presence of spurious long memory. Semi-parametric wavelet-based estimators of long memory in the tradition of the Hurst-Mandelbrot R/S analysis are considered, but recent work suggests that neither bootstrap nor Monte Carlo standard errors and confidence intervals are reliable. Although we find evidence of long memory using these estimators, there is no way to confidently test hypotheses. Finally, out-of-sample forecasting is conducted using true and spurious long memory models as candidates and results are interpreted in the decision theoretic loss function framework of Elliott, Granger and Timmermann (2006).

2   Long Memory in Commodity Futures Data

The main contribution of this paper is to present new, robust estimates, along with small, efficient standard errors, of the long memory parameter identified in commodity futures price volatility. These results contribute to an active and growing literature in agricultural economics to better understand the relationship between commodity futures and options through improved models of price volatility and measures of serial dependence. A large estimate of long memory in futures price volatility implies a potentially large bias in the classic Black-Scholes option pricing method and its American option counterpart. Option pricing based on the Black-Scholes model assumes that the underlying asset (here, commodity futures) is reasonably well described by Geometric Brownian Motion (GBM), which means the natural logarithm of the asset price behaves in the continuous-time limit as an IID random walk with drift.

One typical violation of the Black-Scholes model in futures price sample data is volatility clustering (Myers and Hanson 1993), generally addressed by using GARCH models (Bollerslev 1986). However, this short-range dependence does not appear to affect much option pricing solutions (Roberts 2002).

Long-range dependence, or long memory, implies the Black-Scholes option pricing solution is fundamentally biased (Rogers 1997; Sottinen 2001), as the underlying asset is better described by fractional Brownian

motion, a more general stochastic process that nests geometric Brownian motion as a special case (Cox and Miller 1965).

2.1 The Option Pricing Bias from Long Memory

Understanding the behavior of futures prices is central to commodity risk management (Tomek 1997; Tomek and Peterson 2001). A better understanding of the process driving futures prices influences hedging and inventory decisions, spot price discovery, and the use of commodity options written on futures.

A large and still active literature suggests that long memory or persistence in commodity futures price data is significant and therefore of practical consequence Jin and Frechette 2004; Corazza, Malliaris and Nardelli 1998; Crato and Ray; Cromwell, Labys and Kouassi; Helms and Rosenman 1984; Peterson, Ma and Ritchey; Wei and Leuthold, 2000). Long memory in time series is characterized by a hyperbolic (slow) rate of decay (i.e. persistence) of the impulse response coefficients, instead of the usual geometric (faster) rate of decay. A long memory process is I(d), or fractionally integrated of order $d \in (-1, 1)$. The case $d = 1$ is well-known: such a process has a unit root and "permanent memory."

This concern has replaced the empirical search for a unit root in commodity price data. It is now well-established that agricultural commodity price time series are unlikely to contain a unit root (Wang and Tomek, 2007). Support for this conclusion comes both from theory (Tomek, 1994) and the econometric literature on the low power of unit root tests in the presence of either structural breaks or long memory.

How important is the bias caused by long memory on option pricing? Ohanissian, Russell and Tsay (2004) find that it can cause options to be mispriced by as much as 67%. In contrast, Roberts (2002) finds that commodity option pricing is unaffected by short-range dependence such as ARCH effects.

3 Commodity Futures Price Data

The data consist of business daily observations of agricultural commodity futures prices for contracts on coffee, cotton, cocoa, sugar #11, frozen concentrated orange juice, hard red winter wheat, soybeans, corn, canola, live cattle, and lean hogs (formerly live hogs). The observations cover the years 1987-2007 but vary slightly between commodities. The contracts chosen include both storable and non-storable commodities. Commodities may be categorized as non-storable, storable with large inventories ("overhangs") or storable with small inventories. These categories also lead to testable predictions of futures forecasting accuracy. Futures provide an unbiased forecasting measure for non-storable commodities (as well as other instruments

such as Federal Funds). For storable commodities with large inventories, futures prices incorporate a cost of carry (storage plus interest), and perhaps a convenience yield (this need not be the case, see e.g., Brennan, Williams and Wright, 1997). Storable commodities with small inventories can be described by two cases. If futures prices are higher than spot prices ("contango") then the analysis follows the large-inventory case. But if futures prices are lower than spot prices ("backwardation"), we can apply the analysis as if it were non-storable.

Descriptive statistics for the volatility proxy used, the log range, are presented in Table 1. Figures 1 through 11 provide time series plots of the daily log range for each commodity futures contract over the time period 1988-2004. Although the log range for commodity futures prices is not perfectly distributed Gaussian , it is nonetheless closer to normality than are logreturns and powers of logreturns (Alizadeh, Brandt and Diebold 2002). International commodities traded at the New York Board of Trade, such as cocoa, coffee and cotton, are more volatile, skewed and leptokurtic (heavy-tailed) than are more domestic commodities such as Chicago Board of Trade grains and Chicago Mercantile Exchange lean meats.

Data for the years 2005, 2006 and 2007 are reserved for out-of-sample forecasting (at least 500 observations for each commodity). This leaves more than 4000 unused observations for each commodity for forecasting exercises. For estimation purposes, $T = 4096$. Although data are available going back several decades, the possibility of slow, difficult to identify structural breaks suggests it is better to use only relatively recent data (Alston and Chalfant 1988, 1991).

Commodity futures contracts are traded until the 15th of the contract month (or the last business day before the 15th). To avoid near-maturity effects and delivery risk bias, observations for contracts in their own expiry month are discarded. Contracts are therefore rolled-over (spliced) approximately 15 days before they expire.

In providing diagnostic statistics, Kim and White's (2001) approach is followed to compute robust measures of skewness and kurtosis. Their work shows that once outliers are accounted for, estimated sample higher order moments in financial time series are generally Gaussian Normal. These results help explain why observed power laws and Pareto-stable sample properties do not resist aggregation over longer time periods.

A number of robustness checks are performed. First, to control for calendar effects such as the "weekend" anomaly (French 1980; Thaler 1987; Gibbons and Hess 1981; Kamara 1997), we repeat estimation for one commodity using only the Wednesday observation (i.e. weekly sampling). Two reasons suggest however calendar effects need not be a problem. Empirical work has found these anomalies seem to have dissapeared since 1975 (Connolly 1989) or since 1987 (Fortune 1998), and the earliest data used begins in 1987. Also,

once unintentional data snooping is accounted for, calendar effects are not statistically significant (Sullivan, Timmermann and White, 2001).

Second, splicing contracts introduces spurious nonlinear dynamics into the pooled time series dataset (Smith, 2005). For a single commodity, the Partially Overlapping Time Series method, using Kalman filtering, is followed to evaluate the impact of splicing bias on the long memory parameter estimate and standard error.

Standard time series diagnostic tests are performed on the data (Augmented Dickey-Fuller, Phillips-Perron, KPSS, Variance Ratio) to evaluate its sample properties and ensure comparability of our data with data used by other researchers. Test results show that in levels we cannot reject the null of a unit root (ADF) while in differences we cannot reject the null of no unit root (KPSS). Nonstationarity in levels and stationarity in differences is a typical finding but Wang and Tomek (2007) show that commodity prices in levels should not be characterized by a unit root and such a test result is misleading and a consequence of failing to control for level shifts in the data. In log-return or log-range form stationarity is confirmed and ARCH effects (volatility clustering) are present. Details of these tests are provided in the Appendix.

3.1 Wavelets Distinguish Short from Long Memory

A substantial difficulty associated with estimating the long memory parameter (H or d) is that it is, even asymptotically, correlated with short memory dynamics, such as AR and MA parameters (Tanaka 1999). As a result, both the point estimate of the long memory parameter d and its standard errors are biased. Several papers in the literature focus only on the estimation of long memory and do not consider the impact of short memory dynamics.

One solution to this problem is to use an estimation method based on wavelet functions (Gencay, Selcuk, and Whitcher 2001). This is because wavelets by construction are able to separate long memory from short memory dependence, or more generally, variation in a signal or time series that occurs at different timescales (Percival and Walden 2001).

To briefly explain the concept of wavelet functions, we follow the exposition by Tanaka (2004) and by Crowley (2007). A rigorous treatment is provided by Vidakovic (1998).

Consider a continuous-valued time series signal $x \in L^2$ that is square-integrable, such that:

$$||x^2|| = \int_{-\infty}^{\infty} x^2(t)dt \tag{1}$$

6

and a wavelet function denoted $\psi(t)$ such that

$$\int_{-\infty}^{\infty} \psi(t)dt = 0 \qquad (2)$$

and

$$\int_{-\infty}^{\infty} \psi^2(t)dt = 0 \qquad (3)$$

The two conditions mean the wavelet function oscillates around zero and has unit energy (variance). A large number of functions satisfy these two and other regularity conditions and therefore qualify as wavelets. For a given wavelet function, a Continuous Wavelet Transform may be defined from translations and dilations of the base wavelet function. For example, the Continuous Wavelet Transform using the Haar wavelet is:

$$C_{a,b}(\psi, x) = \frac{1}{\sqrt{a}} \int_{-\infty}^{\infty} \psi\left(\frac{t-b}{a}\right) x(t)dt \qquad (4)$$

where a is the dilation parameter and b is the translation parameter. The two parameters imply that the wavelet transform is localized in both time and scale, where scale may be interpreted in economic time series as a time horizon or timescale. In contrast, the Fourier transform has a single localization parameter, in frequency.

For applications to discrete-valued time series, a Discrete Wavelet Transform may be obtained. The relationship between Continuous and Discrete Wavelet Transforms is not unlike the relationship between asymptotic theory and small-sample properties. That is, it is convenient to work with the continuous transform to prove theorems and establish properties, but not to use with actual data.

A powerful result is that the application of a wavelet transform to time series data, followed by an application of a corresponding inverse wavelet transform to the wavelet coefficients produced in the first stage, produces the following relationship between the original data and a number of nearly orthogonal time series of the same length and valued in the same unit as the original data:

$$x(t) = s(t) + \sum_{j=1}^{J} d_j(t) \qquad (5)$$

where each $d_j(t)$ is a time series corresponding to a proportion of the original data associated with a specific timescale j, and where $s(t)$ is the long-run trend of the data. The number J is limited by the number of observations in the original data and is typically, though not necessarily, bounded by: $J^{max} = \log_2(N)$ for a number N of observations. The different timescale vectors of data are essentially decorrelated and any loss of statistical information is limited by machine precision (e.g. double precision in Matlab).

The implication is, wavelet-based estimates of long memory are unbiased even when the time series data short memory parameters are either ignored or inaccurately estimated. Moreover, long memory and short memory parameters can be independently and accurately estimated. To the best of our knowledge, the only work to adopt wavelet-based estimators to examine long memory in agricultural commodity price data is Elder and Jin (2007).

However, where they focus on contrasting long memory estimation results between different wavelet and non-wavelet based methods, this work uses only wavelet-based estimators but also estimates jointly short memory parameters and tests for spurious long memory. Lastly, the variable used for volatility is the daily logreturn in their work and the daily log-range in this work.

4 Identifying Spurious Long Memory

One persistent problem found in the literature is the identification of long memory in the futures contracts data. Early evidence of long memory in cash and futures prices has been overturned and recent advances have focused on long memory in the volatility of futures prices (e.g. Jin and Frechette 2004).

In time series econometrics, long memory is increasingly defined as fractional integration, or $I(d)$, which is only one type of long memory process (Granger, 2000). It is well understood that the aggregation of short memory (e.g. ARMA) time series data results in long memory (Granger 1980, 1990). Chambers (1998) proves that true long memory processes have a long memory parameter that is invariant under time aggregation, a fact that is helpful for hypothesis testing. Apparent long memory can also be the consequence of structural breaks and level shifts, and these two phenomena may provide better candidate models in economics (Granger, 2005; Diebold and Inoue, 2001). Therefore, disentangling structural breaks and other stochastic, singular shocks from long memory in time series data is a difficult problem (see e.g. Banerjee and Urga, ed., 2005, *Journal of Econometrics* symposium; Perron, 2006).

Spurious findings of long memory may be explained by the method of estimation, by level shifts, structural breaks and regime switches, or by inefficient standard errors and confidence intervals (Chambers, 1998; Diebold and Inoue 2001; Shimotsu, 2006; Zivot and Andrews, 1992). In the case of agricultural commodities, Lordkipandize (2004, p. 82) finds that soybean and corn futures price volatility is primarily caused by seasonality and maturity effects rather than long memory. This finding is in contrast with Cromwell, Labys and Kouassi's (2000) results that at least six commodity price series are fractionally integrated.

The estimators generally used in the literature are not necessarily robust. The popular Geweke-Porter-

8

Hudak semiparametric estimator is both inconsistent and inefficient (Robinson 1995a,b). Moreover, Smith (2005) shows that the GPH estimator is heavily biased in the presence of level shifts in the data and suggests a new, nearly-unbiased version of this estimator. He shows that this bias explains an apparently large ($d = 0.79$) estimate of long memory in relative soybean prices. As a result, once level shifts have been accounted for, estimates of long memory are not statistically different from zero.

Based on the widely used GARCH model of time-varying volatility (Bollerslev 1986), the long memory FIGARCH parametric estimator (Bollerslev and Mikkelsen 1996) is both fragile to misspecified short memory parameters and is also unreliable as a measure of long memory (Davidson 2004). Finally, the accuracy of Quasi-MLE stochastic volatility long memory models (Breidt, Crato and de Lima 1998) has been questioned because their estimation is fragile to the presence of non-Gaussian innovations (Alizadeh, Brandt and Diebold, 2002; Andersen and Sorensen, 1997).

Riedi (2003) shows that confidence intervals on Hurst long memory estimates are only reliable under quite restrictive conditions, and Franco and Reisen (2007) find through simulation that bootstrapped standard errors on estimated parameters of fractional integration are not satisfactory. Turvey (2007) shows that for all but two agricultural commodities, the data generating process is consistent with white noise innovations rather than fractional Gaussian noise (as would be the case under long memory).

These results suggest the need for a more robust investigation of long memory.

5   Semiparametric Wavelet Estimation of Long Memory

The literature on estimation of long memory in time series data finds that semiparametric estimators, frequently used in the natural sciences, are superior to parametric (maximum likelihood) estimators under model misspecification (Boes et al. 1989) But MLE is preferable when the model is correctly specified (Cheung 1993).

Wavelets properties make them ideal for detecting self-similarity in stochastic processes, including long memory such as fractional Brownian motion. An alternative to parametric estimation is to follow a similar approach to Hurst's original rescaled range (R/S) estimation of long memory.

The classic work of Mandelbrot and Van Ness (1968) showed how to estimate long memory based on Hurst's H Rescaled-Range analysis. Lo (1991) improves upon Hurst's and Mandelbrot's Rescaled-Range approach by making it robust to heteroskedasticity in the data, but it remains less practical to use in a parametric time series setting and both confidence intervals and hypothesis testing are unreliable. As

wavelets are ideally suited to detect the self-similar fractal signature of several types of long memory, Abry, Veitch and Flandrin (1998), Abry and Veitch (1999), and Teyssiere and Abry (2006) develop improved, unbiased, and more robust semiparametric estimators of long range dependence based on the Hurst measure H.

Taqqu (2003) shows how the Hurst coefficient can be easily obtained from an application of a wavelet transform to time series data. The method consists of first computing the Discrete Wavelet Transform of the relevant time series data, which produces a vector of wavelet coefficients. Then the wavelet coefficients, each of which is associated with a timescale, are squared and regressed over the base-2 (dyadic) logarithm of the timescales. The slope coefficient is in theory proportional to H. First, using the wavelet transform log-2 timescale coefficients are extracted from the data. Then, the wavelet coefficients are regressed on the log-2 scale in a simple affine model. Finally, the Hurst parameter is computed directly from the slope parameter estimated in the regression of wavelet coefficients over log-2 scale.

Although straightforward to compute and frequently used in the natural sciences, this approach is of limited usefulness in economics because it has been shown that both analytical and bootstrap standard errors and confidence intervals are unreliable when applied to this estimation method (Riedi 2003; Franco and Reisen 2007).

Jensen (1998) shows how to use a similar method to obtain an OLS estimator of the fractional difference parameter d, which is directly related to the Hurst coefficient H. His method is computationally straightforward, but unlike maximum likelihood estimators is not full information.

In summary, results from the application of wavelet-based semi-parametric estimators provide evidence in favor of long memory, but these are not pursued further because of the inability to conduct reliable hypothesis testing.

6   Parametric Wavelet Estimation of Long Memory

Following Granger's (1980) and Hosking's (1981) formal definitions of long memory in the ARMA time series framework, Sowell (1992) obtained an exact MLE for fractionally integrated processes. Computation, however, requires inverting a dense covariance matrix, which is seldom possible when using large datasets. For this reason, the Geweke-Porter-Hudak (1983) approach remains popular, despite evidence evidence that the GPH estimator is inconsistent and inefficient (Agiagoglou, Newbold and Wohar, 1992; Robinson, 1995). Robinson suggests a semiparametric local Whittle estimator based on Kunsch (1987). Feasible exact ML

estimators suffer from a large bias as the sample size grows (Cheung and Diebold, 1994).

Hosking (1984) suggests a Cumulative Sum of Squares estimator that is asymptotically equivalent to Sowell's (1992) exact MLE. However, Chung and Baillie (1993) show that the CSS estimator is severely biased in small to moderate-sized samples. Chung (1996a,b) derives asymptotic results for the CSS estimator of a generalized ARFIMA(p,d,q) process including an analytical formula for standard errors, which is used in this paper.

The ability of wavelet functions to orthogonalize stochastic processes helps distinguish long memory from short memory (ARMA) components as well as from change-points or structural breaks (Percival and Walden, 2000). Moreover, wavelet-based estimators of long memory are not affected by the presence of an unknown mean μ (Jensen 2000), unlike time domain exact ML estimators. This is helpful because the sample mean is an inaccurate estimator of the population mean in the presence of long memory (Beran 1994; Cheung and Diebold 1994). Wavelets create a sparse representation of the covariance matrix, which greatly simplifies computational solution of the log-likelihood. A large number of wavelet-based estimators exist. McCoy and Walden's (1996) wavelet-based MLE of long memory parameters has been improved upon by Percival and Bruce (1998) to include polynomial trends, by Jensen (2000) for robustness to contaiminated data and by Craigmile, Guttorp and Percival (2005) to make it robust against trend contamination. Whitcher (2004) includes a seasonal component, and Jensen (1998, 1999) suggests a method to jointly estimate long memory and short memory parameters.

The first estimator used here has the smallest Mean Squared Error in the class of semiparametric estimators, is computationally efficient, only slightly affected by the boundary effects due to the finiteness of the data sample, and is robust to misspecification of trend and short memory (ARMA) parameters.

The general long memory process to be estimated is:

$$\Phi(L)(1 - 2\eta L + L^2)^d (Y_t - \mu) = \Theta(L)\epsilon_t \tag{6}$$

which includes both autoregressive $\Phi(L)$ and moving average $\Theta(L)$ polynomials, a fractional order of integration d (Hosking 1981) as well as a seasonal persistence process (Gray, Zhang and Woodward 1989) which is equivalently a power series known as Gegenbauer polynomials (Rainville, 1960).

Gegenbauer polynomials enable the long memory parameter to be associated with seasonality, an advantageous option to study economic time series. In this paper, however, time series are sampled daily and attempts to include seasonal structure are left for an extension of this work.

First is considered a simple $I(d)$ process, or ARFIMA(0,d,0). In effect this model can be seen as a

restricted case of the more general ARFIMA(p,d,q) which is estimated in the second round and can be compared using Likelihood Ratio tests.

Then is applied an estimator that also obtains short memory parameter values and coefficients for seasonality, which are useful for grain futures contracts.

6.1 Log-Range Measure of Volatility

A daily range-based measure of volatility is used instead of the traditional daily logreturns (computed as deviations in absolute value or squares from the long-run mean). There are several reasons for this choice.

The ideal measure of daily integrated volatility is realized volatility computed from ultra high-frequency tick data, with evidence from stock and foreign exchange markets (Anderson, Bollerslev, Diebold et al. 2001a, 2001b, 2003; Barndorff-Nielsen and Sheppard, 2002). However, the lower volume of trade in commodity markets implies realized volatility is not always a feasible option. Anderson and Bollerslev (1998) show that range-based volatility is nearly as accurate as computing the realized volatility from tick-by-tick data (see e.g. Anderson, Bollerslev, Diebold et al. 2001a,b for evidence that realized volatility is the ideal measure of daily integrated volatility).

What is more, the log-range is very well approximated by the Gaussian Normal distribution (Alizadeh, Brandt and Diebold, 2002; Brandt and Jones, 2006), which improves forecasting efficiency and accuracy using QMLE and Kalman filtering. QMLE estimation using a logreturn-based volatility is highly inefficient (Kim, Shephard and Chib, 1998; Andersen and Sorensen, 1997).

Logreturns are a very noisy proxy for price variation and are contaminated by measurement error (Parkinson, 1980; Garman and Klass 1980; Rogers and Satchell, 1991). As a result, the log-range based volatility measure is more efficient (has a smaller variance) than logreturn-based volatility measures.

Lastly, the typical measure of volatility $E[|r|^\theta]$, for r the logreturn deviation from the long-run mean and $\theta \in \mathbb{Z}$, is not well supported by choice theory as a proxy for risk (Machina, 1987; Levy, 1992). Granger (2000) suggests following Nyquist's (1983) L_p norm argument if one must use the logreturn volatility measure. Then, if the data are leptokurtic, the L_1 norm and $\theta = 1$ are preferable, while the L_2 norm and $\theta = 2$ are appropriate when the data are approximately Gaussian.

The log-range, in the general case of intra-daily observations, is defined as:

$$\ln |f(s_{iH,(i+1)H})| = \ln(\sup F_t - \inf F_t) \text{for} t \in [iH, (i+1)H] \tag{7}$$

12

Assuming the price of an asset is approximately log-Normal, the variance of daily returns equals the diffusion term in the Geometric Brownian Motion representation of the asset price, which is constant in the classic Black-Scholes-Merton model but may be considered stochastic.

Parkinson (1977, 1980) suggests the following volatility measure approximating daily realized variance of returns. This measure uses the range of daily log-prices.

$$0.361 * (\sup \ln(F_{t,\tau}) - \inf \ln(F_{t,\tau}))^2 \tag{8}$$

This measure of daily volatility is the closest approximate to actually computing realized volatility, according to Alizadeh, Brandt and Diebold (2002) and it is the variable used to estimate long memory. To simplify computation and interpretation, the actual variable used is one thousand times Parkinson's range-based daily volatility measure.

Open and close prices are not incorporated as they do not improve accuracy of results and they introduce undesirable market microstructure effects (Brown, 1990; Alizadeh, 1998).

The wavelet ML estimator used relies on the Haar(4)-based wavelet transform (Daubechies 1992). Jensen (2000) reports simulation results from using a range of different wavelet transforms in the likelihood estimator and concludes it has the smallest small-sample bias and (Root) Mean Squared Error. For a general ARFIMA(p,d,q) process with white noise innovations

$$\tilde{x}(t) = x(t) + \eta(t) \tag{9}$$

the concentrated log-likelihood (Jensen, 2000) is:

$$\log \mathcal{L}(\theta) = -\frac{1}{2} \sum_{m \in \mathcal{M}} \sum_{\mathcal{N}(m)} \left(\frac{\langle \tilde{x}, \psi_{m,n} \rangle^2}{\sigma_{\tilde{m},n}} + \ln 2\pi \sigma_{\tilde{m},n} \right) \tag{10}$$

where $\theta = (d, \sigma^2, \sigma_\eta^2)'$ is the vector of parameters (long memory, ARFIMA variance, and white noise variance) and $\langle \tilde{x}, \psi_{m,n} \rangle$ is the vector of wavelet coefficients resulting from the convolution of wavelet functions with the original data.

The EM algorithm can be used to improve convergence of the estimator (Wornell and Oppenheim 1992). This algorithm will improve the likelihood with every step but is not guaranteed to converge to the best estimates.

Standard errors for the long memory parameter are computed following analytical solutions from Chung (1996a,b) and Tanaka (1999). It is necessary to first estimate d, then estimate the ARMA parameters

$\phi(L), \theta(L)$ and finally obtain the information matrix for the ARMA parameters (see e.g. Hamilton 1994, pp. 142-144). Then, accurate standard errors for d can be computed. For the general case of an ARFIMA(p,d,q) process, standard errors are computed as follows:

$$\text{se}(\hat{d}) = T^{-1/2}\sqrt{\frac{6}{\pi^2}} \tag{11}$$

For the case of seasonal persistence, which is not addressed here, standard errors are computed as follows:

$$\text{se}(\hat{d}) = \frac{[2(\frac{\pi^2}{3} - \pi\arccos(0.5) + \arccos^2)]^{-1/2}}{T^{1/2}} \tag{12}$$

7 Long Memory Estimates

The results of the restricted I(d) wavelet MLE specification are presented in Table 2. The estimated long memory parameter d is 0.309 for live cattle, 0.320 for lean hogs, 0.321 for soybeans, 0.304 for corn, 0.431 for wheat, 0.436 for canola, 0.258 for coffee, 0.271 for cocoa, 0.290 for cotton, 0.194 for orange juice, and 0.279 for sugar #11.

The naive standard errors in Table 2 are identical for all commodities because they do not depend on the Information Matrix but rather are a function of the number of observations ($T = 4096$ in all cases) and the ARMA parameters, here restricted to be zero.

After estimating the ARMA parameters and computing their Information Matrix, correct standard errors are obtained and presented in Table 2. Surprisingly, the correct standard errors for the fractional difference parameter d are only slightly affected by the presence of ARMA terms. For example, the correct standard error for the long memory parameter in soybean futures price volatility is 0.0122, hardly different from the naive standard error 0.0105.

A Wald test can be applied to evaluate the null that the fractional difference parameter d is zero, equivalently that there is no long memory (Tanaka 1999). Even though the wavelet-based estimator is robust to the presence of misspecified short memory parameters, the test is only accurate if these ARMA (or GARCH) terms are included in the Information Matrix:

$$\frac{\hat{d} - d_0}{\sqrt{T^{-1}\omega^{-2}}} \tag{13}$$

where the complete Information matrix for both long and short memory parameter estimators is:

14

$$\omega^2 = \left(\frac{\pi^2}{6} - \kappa' \mathcal{I}^{-1}(\phi, \theta) \kappa \right) \tag{14}$$

where $\mathcal{I}(\phi, \theta)$ is the Information matrix for only the ARMA terms and where the term κ is computed from the expansion of the ARMA(p,q) lag polynomials (applying Wold's Theorem after verifying that invertibility holds by looking at the roots of the MA polynomial):

$$\kappa' = \left(\sum_{j=1}^{\infty} j^{-1} \phi_{j-1}, \sum_{j=1}^{\infty} j^{-1} \phi_{j-2}, ..., \sum_{j=1}^{\infty} j^{-1} \phi_{j-p}, -\sum_{j=1}^{\infty} j^{-1} \theta_{j-1}, -\sum_{j=1}^{\infty} j^{-1} \theta_{j-2}, ..., -\sum_{j=1}^{\infty} j^{-1} \theta_{j-q} \right) \tag{15}$$

therefore κ is (p+q)x1.

The ARMA information matrix is computed using the BHHH Hessian estimator (Berndt, Hall, Hall and Hausman 1974):

$$\hat{\mathcal{I}}(\hat{\phi}, \hat{\theta}) = \hat{\mathbf{G}}'\hat{\mathbf{G}} \tag{16}$$

where $\hat{\mathbf{G}}$ is the T x (p+q+1) matrix of estimated scores.

However, the Wald test has a size problem and tends to over-reject the null of $d = 0$. Similarly, Likelihood Ratio tests do not seem very effective against spurious long memory. They may be computed to compare the restricted I(d) to the unrestricted ARFIMA (p, d, q) model to evaluate the significance of the short memory parameters. In addition, a LR test can be used to evaluate the restricted ARMA (p, q) against the unrestricted ARFIMA (p, d, q) to evaluate the significance of the fractional difference, long memory parameter d.

8 Short Memory Estimates

The estimates of the long memory fractional difference parameter when estimated using a wavelet transformation of the covariance matrix are unaffected by short memory dynamics. This has the advantage of enabling a two-step estimation procedure, which should improve the convergence of the likelihood by reducing the number of parameters to be estimated together, As a consequence, rather than estimate simultaneously all parameters $(d, \phi, \theta)'$ as does Jensen (1998), first the long memory parameter d is estimated, and then the short memory ARMA parameters $(\phi(L), \theta(L)$ are estimated using the correctly fractionally differenced data. The differencing procedure is similar to taking differences of a dataset that is originally in levels, but here we need to take a fractional difference, which must be solved numerically. Then a standard full-information ARMA likelihood procedure can be used to efficiently estimate the short memory ARMA parameters using the fractionally differenced data.

Once the short memory parameters are estimated, their Information matrix can be used to obtain the correct standard errors for the long memory parameter d and these are presented in Table 2.

When data in levels are integrated of order one, $I(1)$, the underlying process is nonstationary and traditional statistical inference is invalid, hence the need to difference the data. For data that are integrated of order $d < 0.5$, $I(d)$, the underlying process is stationary while it is nonstationary when the data are $I(d)$ for $d \geq 0.5$. The stationarity of the differenced data is verified using an appropriate ADF-GLS test (Elliott, Rothemberg and Stock 1996).

To fractionally difference the time series data, the following binomial formula (Hosking 1981) is used:

$$\Delta^d Y_t = \sum_{j=1}^{k} \frac{\Gamma(j - d)}{\Gamma(j + 1)\Gamma(-d)} Y_{t-j} \tag{17}$$

Since working with Gamma functions is unwieldy, Stirling's approximation is used to simplify computations (Abramowitz and Stegun 1972, p. 257):

$$\frac{\Gamma(k + \alpha)}{\Gamma(k + \beta)} = \lim_{k \to \infty} k^{\alpha - \beta}(1 + O(k^{-1})) \tag{18}$$

A fast numerical solution is to use the Gauss hypergeometric function:

$$_2F_1(d, 1; 1; L) = \left(\frac{1}{1 - L}\right)^d \tag{19}$$

This can be implemented using the statistical analysis language R (Reisen 1999; Fraley et al. 2006).

Once the properly fractionally differenced data are obtained, a standard ARMA model is fitted by exact maximum likelihood (e.g. Hamilton 1994, pp. 132-133) in a state-space framework using the Kalman filter and assuming Gaussian innovations. The estimates and standard errors are compared with results obtained from using Haslett and Raftery's (1989) older method of joint ARFIMA parameter estimation, which is faster but less robust.

Model selection of ARMA parameters is based on pairwise Likelihood Ratio tests between a larger unrestricted model and a smaller restricted model, always using the 1% level of significance. Akaike and Schwartz Information Criteria are computed but these may be less reliable because they have been found to overparameterize the model.

The results, parameter estimates and standard errors, are presented in Table 2.

16

9 Testing for Spurious Long Memory

The literature on testing between unit roots (or long memory) and structural breaks or level shifts is vast (Banerjee and Urga 2005; Perron 2006).

Two recently developed approaches are discussed in the following section and one is selected to be applied to all commodity datasets.

9.1 Shimotsu's Approach

Shimotsu (2006) suggests two simple but effective tests for spurious long memory. The first is a Wald test that uses the invariance to time aggregation property of true long memory processes shown by Chambers (1998) and that is valid for both stationary and nonstationary long memory processes. The second consists of fractionally differencing the data (using an estimate of the memory parameter d) and then computing tests on these data for both a null of $I(0)$ (KPSS test) and a null of $I(1)$ (Phillips-Perron Z_t test).

Shimotsu considers as alternative data generating processes three cases: (i) a mean plus noise process, (ii) Engle and Smith's (1999) stochastic permanent break model, and (iii) a Markov-switching model. A powerful testing procedure is then to consider Likelihood Ratio (or Score) tests to evaluate the null hypothesis of a long memory model (with or without ARMA parameters) against an alternative of a precisely specified structurally changing or regime-switching model. In the absence of a clearly supported specification, however, Shimotsu's second test is used. It should be mentioned that important class of models, jump-diffusions, is ignored in this analysis. Yet on theoretical ground it appears a more plausible class of models than does long memory (e.g. Granger 2003, 2005). Jump-diffusion models are increasingly used and particularly useful to link futures with options-on-futures (see e.g. Koekebakker and Lien 2004; Saphores, Khalaf and Pelletier 2002). However, these models are substantially more difficult to work with and hypothesis testing is complicated by the presence of nuisance parameters that must be dealt with through simulation methods (e.g. Khalaf, Saphores and Bilodeau 2003).

To evaluate the hypothesis of true long memory against a unknown forms of spurious long memory, adjusted KPSS and Phillips-Perron tests are applied to the fractionally differenced and appropriately demeaned data. Critical values are provided by Shimotsu (2006, Table 2).

Suppose the data generating process appears to be long memory $I(d)$ but is in fact $I(1)$. This could be a mean plus noise or a stochastic break model. Then taking the d^{th} difference will result in a new process that is $I(\alpha)$ where $\alpha = 1 - d$. The KPSS test will correctly reject the null that this new, fractionally differenced

process is $I(0)$ but the Phillips-Perron test will fail to reject the null that the new process is $I(1)$, so we learn that the process is most likely spurious long memory. Suppose on the other hand the data generating process is characterized by true long memory d. Then, taking the d^{th} difference will result in a $I(0)$ process and the KPSS test will correctly fail to reject the $I(0)$ null, while the Phillips-Perron test will correctly reject the $I(1)$ null.

The results, presented in Table 2, are summarized as follows.

There is strong evidence of true long memory for KCBOT wheat and WCE canola futures volatility, and the evidence for CBOT corn futures is ambiguous as neither test null hypothesis can be rejected. For the other commodities, however, the results imply long memory is spurious and better explained by a nearly nonstationary model such as Engle and Smith's (1999) stochastic break process.

In conclusion, since the wavelet-based estimator is robust to the presence of short memory dynamics, findings of spurious memory for most of the commodities suggests other dynamics must be responsible for the illusion of persistence in volatility. One leading candidate model that is addressed in the next section, is a class of Markov-switching models that generates spurious long memory.

9.2   Ohanissian, Russell and Tsay's Approach

Ohanissian, Russell and Tsay (2005) suggest a simple test of spurious long memory that is based on Hausman's (1978) result that an efficient estimator must have zero asymptotic covariance with any other consistent, asymptotically normal estimator. Under the null of long memory, the covariance of two estimates of long memory for the same data but aggregated two different ways will asymptotically equal the variance of the long memory estimator for the less-aggregated data. The test's limitations are that it relies on the GPH long memory estimator which is, even with its improvements by Robinson and by Smith, markedly inferior to wavelet and Whittle-type estimators. Also, it is designed for applications to ultra-high-frequency financial data ("tick" observations), and only has good power when a large dataset is available such that the number of different aggregation levels M is high. A large number of aggregation levels can be used to obtain better test power This test value is compared with a critical value distributed as Chi-Squared(M). Using this test, the authors find that long memory in foreign exchange rate time series data is true, not spurious.

If the less aggregated data consists of the original daily observations, and the aggregation level m results in a number of ordinates $l^{(m)}$, generally chosen to be $\sqrt{\frac{n}{m}}$ then:

$$\lim_{n \to \infty} 4l^{(m_i)} \left(Cov(\hat{d}^{(m_i)}, \hat{d}^{(m_j)}) - Var(\hat{d}^{(m_i)}) \right) = o(1) \tag{20}$$

18

For example, in our data $n = 4096$, $m = 8$, so $l^{(m)} = 22.627$. The test statistic is $\widehat{W} = (T\hat{d})'(T\Lambda T)^{-1}(T\hat{d})$.

The results in Table 2 suggest we can reject the null of true long memory for all commodities at the 5% level of significance except for wheat, canola and cocoa, for which we can only reject the null at the 10% level. These test results provide additional support for the claim that Kansas City wheat and Winnipeg canola futures volatility is characterized by true long memory, while for all other commodities the long memory is spurious. The critical test values are Chi-Square$(M-1)$ for M aggregation levels used in the covariance, including the original dataset. In particular here they are 3.841 (5%) and 6.635 (1%).

10 Conclusion

This paper investigates one important cause and consequence of bias in commodity futures option pricing and contributes to the active literature on the robust estimation of long memory in commodity futures price volatility using a novel empirical strategy that also enables the computation of efficient standard errors for the long memory parameter jointly with the unbiased estimation of short memory parameters.

There is evidence of long memory for all commodity futures contracts in the log-range volatility of prices. However, it is smaller in magnitude than previous research has found and, most importantly, the results appear to be spurious for all commodities except for Kansas City Board of Trade wheat futures and Winnipeg Commodity Exchange canola futures. The results are inconclusive for Chicago Board of Trade corn futures, so a non-nested hypothesis test is constructed using as an alternative a Markov-switching process that generates spurious long memory.

Although the time series model used is relatively simple, we refer to Lordkipandize (2004) who estimates a much larger, stochastic volatility model of commodity derivative prices. She concludes that once breaks and seasonality are properly accounted for, the effect of long memory is inconsequential. Her analysis was only applied to corn and soybeans futures, however.

What this paper's findings imply is that true long memory is unlikely to be a good description of the data generating process underlying agricultural commodity futures prices and volatility. But since spurious long memory is frequently found, models of commodity futures should be used that are known to mimic the serial dependence of long memory processes.

Option pricing in agribusiness is therefore unlikely to gain much by using fractional Brownian motion and fractional noise as building blocks instead of the classic Black-Scholes-Merton model, but the results in this paper provide support for the expanding volume of research on jump-diffusion models of option pricing.

A References

Abramowitz, M. and I.A. Stegun, eds. (1972). Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables. New York: Dover.

Abry, P. , P. Flandrin, M.S. Taqqu and D. Veitch, (2000). Wavelets for the analysis, estimation, and synthesis of scaling data. In: K. Park and W. Willinger, Editors, Self Similar Network Traffic and Performance Evaluation, Wiley, New York, pp. 39-88.

Abry, P., P. Flandrin, M.S. Taqqu and D. Veitch (2002). Self-similarity and long-range dependence through the wavelet lens. In: P. Doukhan, G. Oppenheim and M.S. Taqqu, Editors, Long-Range Dependence Theory and Applications, Birkhauser, Boston (2002), pp. 527-556.

Abry, P.; P. Flandrin, M.S. Taqqu, D. Veitch (2003), Self-similarity and long-range dependence through the wavelet lens, Theory and applications of long-range dependence, Birkhsuser.

Agiakloglou, C., Newbold, P., Wohar, M., (1992). Bias in an estimator of the fractional difference parameter. Journal of Time Series Analysis 14, 235-246.

Alizadeh, S., M.W. Brandt, F.X. Diebold (2002). Range-Based Estimation of Stochastic Volatility Models. Journal of Finance 57(3), 1047-1091.

Alston, J.M., and J.A. Chalfant (1988). Accounting for Changes in Tastes, J. Pol. Econ., 96(2): 391-410.

Baillie, R.T., T. Bollerslev, and H.O. Mikkelsen (1996). Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity. Journal of Econometrics 73(1): 5-59.

Banerjee, A. and Giovanni Urga (2005). Modelling structural breaks, long memory and stock market volatility: an overview, Journal of Econometrics 129: 1-34

Barkoulas, J., W.C. Labys, and J. Onochie. Fractional Dynamics in International Commodity Prices. Journal of Futures Markets 17(April 1997):161-89.

Beran, J. (1994). J. Beran, Statistics for Long-Memory Processes, Chapman & Hall, New York (1994).

Bollerslev, T. Generalized Autoregressive Conditional Heteroskedasticity. Journal of Econometrics 31(April 1986):307-27.

Bollerslev, T., R.F. Engle, and D.B. Nelson. (1994) ARCH Models. In R.F. Engle and D. McFadden, eds. Handbook of Econometrics, vol. 4. Amsterdam: Elsevier Science B. V., 1994, pp. 2959-3038.

Bollerslev, T., R.Y. Chou, and K.F. Kroner. ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence. Journal of Econometrics 52(1): 5-59.

Brandt, M.W. and C. Jones (2006). Forecasting Volatility with Range-Based EGARCH Models, Journal of Business and Economic Statistics, 24: 470-486.

Breidt, J. F., Crato, N., and de Lima, P. J. F. (1998). On the detection and estimation of long-memory in stochastic volatility, Journal of Econometrics 83, 325-348, 1998.

Chambers, M.J. (1998). Long memory and aggregation in macroeconomic time series. International Economic Review 39, 1053-1072.

Cheung, Y.W. and F.X. Diebold (1994). On Maximum likelihood estimation of the differencing parameter of fractionally-integrated noise with unknown mean, Journal of Econometrics, 62(2): 301-316.

Chung, C., (1996). A generalized fractionally integrated autoregressive moving-average process. Journal of Time Series Analysis 17, 111-140.

Chung, C., (1996). Estimating a generalized long memory process. Journal of Econometrics 73, 237-259.

Cohen, A., Daubechies, I., and Feauveau, J.C. (1992) Biorthogonal basis of compactly supported wavelets, Communications in Pure and Applied Mathematics, vol. 45, 485-560.

Corazza, M., A.G. Malliaris, and C. Nardelli. (1997) Searching for Fractal Structure in Agricultural Futures Markets. Journal of Futures Markets 17: 433-73.

Craigmile, P.F. (2005), Approximate wavelet based simulation of long memory processes, Journal of Statistical Computation and Simulation, 75 (5), 363-80.

Craigmile, P.F., P. Guttorp, and D. B. Percival (2005). Wavelet based estimation for polynomial contaminated fractionally differenced processes, IEEE Transactions on Signal Processing, 53 (8), 3151-3161.

Crato, N., and B.K. Ray. (2000). Memory in Returns and Volatilities of Futures Contracts. Journal of Futures Markets 20: 525-43.

Crowley, P. (2007). An Intuitive guide to wavelets, Journal of Economic Surveys, forthcoming.

Daubechies, I. (1992), Ten lectures on wavelets, SIAM.

Davidson, J.E. (2004). Moment and memory properties of linear conditional heteroskedasticity models, and a new model, Journal of Business and Economics Statistics 22 (1): 16-29.

Decoster, G.P., W.C. Labys, and D.W. Mitchell (1992). Evidence of Chaos in Commodity Futures Prices. Journal of Futures Markets 12:291-305.

Diebold, F.X., Inoue, A., (2001). Long memory and regime switching. Journal of Econometrics 105, 131-159.

Elliott, G., T.J. Rothemberg, J.H. Stock (1996). Efficient Tests for an Autoregressive Unit Root, Econometrica 64(4): 813-836.

Engle, R.F. (1982). Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation. Econometrica 50:987-1008.

Engle, R.F., and A.D. Smith (1999). Stochastic permanent breaks. Review of Economics and Statistics 81, 553-574.

Garman and Klass (1980). On the Estimation of Security Price Volatilities from Historical Data, Journal of Business, Vol. 53, No. 1, pp. 67-78.

Gencay, R., F. Selcuk and B. Whitcher (2001). Scaling Properties of Foreign Exchange Volatility. Physica A: Statistical Mechanics 289 (1-2), 249-266.

Gencay, R., F. Selcuk and B. Whitcher (2003). Systematic Risk and Timescales. Quantitative Finance, 3 (2), 108-116.

Geweke, J., and S. Porter-Hudak (1983). The Estimation and Application of Long Memory Time Series Models. Journal of Time Series Analysis 4(): 221-38.

Granger, C.W. (1980). Long Memory Relationships and the Aggregation of Dynamic Models. Journal of Econometrics 14(): 227-38.

Granger, C.W., and R. Joyeaux (1981). An Introduction to Long-Memory Time Series Models and Fractional Differencing. Journal of Time Series Analysis 1(): 15-29.

Granger, C.W.J., (2002). Long memory, volatility, risk and distribution. UCSD Department of Economics working paper.

Hamilton, J.D. (1994). Time Series Analysis. Princeton, NJ: Princeton University Press:.

Helms, B.P., F.R. Kaen, and R.E. Rosenman (1984). Memory in Commodity Futures Contracts. Journal of Futures Markets 4():559-67.

Hosking, J.R. (1981). Fractional Differencing. Biometrika 68: 165-76.

Hurst, H.E. (1951). The Long-Term Dependence in Stock Returns. Transactions of the American Society of Civil Engineers 116 (): 770-99.

Jensen, M.J., (1999). An Approximate Wavelet MLE of Short and Long Memory Parameters, Studies in Nonlinear Dynamics and Econometrics, 3, 239-253.

Jensen, M.J., (1999). Using Wavelets to obtain a consistent ordinary least squares estimator of the long memory parameter. Journal of Forecasting 18, 17-32.

Jensen, M.J., (2000). An Alternative Maximum Likelihood Estimator of Long-Memory Processes Using Compactly Supported Wavelets, Journal of Economic Dynamics and Control, 24, 361-387.

Jensen, M.J., and B. Whitcher (2000). Wavelet Estimation of a Local Long-Memory Parameter, Exploration Geophysics, 31, 89-98.

Kao, G.W., and C.K. Ma (1992). Memories, Heteroscedasticity, and Price Limit in Currency Futures Markets. Journal of Futures Markets 12(1992):679-92.

Khalaf, L., J.-D. Saphores, and J.-F. Bilodeau, (2003). Simulation-Based Exact Jump Tests in Models with Conditional Heteroskedasticity, Journal of Economic Dynamics and Control 28(3): 531-553.

Koekebakker, S. and G. Lien (2004). Volatility and Price Jumps in Agricultural Futures Prices: Evidence from Wheat Options, American Journal of Agricultural Economics, Vol. 86, No. 4, pp. 1018-1031.

Kwiatkowski, D., P. Phillips, P Schmidt, and Y. Shin (1992). Testing the Null of Stationarity Against the Alternative of a Unit Root. Journal of Econometrics 54(): 159-78.

Lo, A.W. (1991). Long-Term Memory in Stock Market Prices. Econometrica 59(): 1279-313.

Lordkipandize, N. (2004). Modeling and Estimation of Long Memory in Stochastic Volatility: Application to Options on Futures Contracts. PhD Dissertation, Cornell University, January 2004.

Mallat, S. (1998), A wavelet tour of signal processing, Academic Press.

Mandelbrot, B.B. (1963). The Variation of Certain Speculative Prices. Journal of Business 36(October 1963a): 394-419.

McCoy, E. J., and A. T. Walden (1996). Wavelet analysis and synthesis of stationary long-memory processes. Journal of Computational and Graphical Statistics, 5: 1-31.

Meyer, Y. (1990), Ondelettes et oprateurs, Tome 1, Hermann Ed. (English translation: Wavelets and operators, Cambridge Univ. Press, 1993.)

Newey, W. and K.D. West, (1987). A simple positive definite, heteroskedasticity and autocorrelation consistent covariance matrix, Econometrica 55 (1987), pp. 703-708.

Ng, S. and P. Perron (1998). Lag Length Selection and the Construction of Unit Root Tests with Good Size and Power. Econometrica, 69(6), 1519-1554.

Ogden, R.T. (1997), Essential wavelets for statistical applications and data analysis, Birkhauser.

Ohanissian, A., J.R. Russell, and R. Tsay (2004). True or Spurious Long Memory? A New Test. University of Chicago Graduate School of Business working paper.

Parkinson, M. (1980). The Extreme Value Method for Estimating the Variance of the Rate of Return. Journal of Business, Vol. 53, No. 1, pp. 61-65

Percival, D. and A.G. Bruce (1998). Wavelet-Based Approximate Maximum Likelihood Estimation for Trend-Contaminated Fractional Difference. University of Washington mimeo.

Percival, D.B. and A.T. Walden, (2000), Wavelet Methods for Time Series Analysis, Cambridge University Press, Cambridge.

Percival, D.B. and H.O. Mofjeld (1997). Analysis of Subtidal Coastal Sea Level Fluctuations Using Wavelets, Journal of the American Statistical Association, 92, no. 439, pp. 868-80.

Perron, P. (1989). The Great crash, the oil price shock, and the unit root hypothesis, Econometrica 57, pp. 1361-1401.

Perron, P. (2006). Dealing with Structural Breaks, in: Palgrave Handbook of Econometrics, Vol. 1: Econometric Theory, K. Patterson and T.C. Mills (eds.), Palgrave Macmillan, 2006, 278-352.

Phillips P.C.B. and P. Perron, (1988) Testing for a unit root in time series regression, Biometrika 75 (1988), pp. 335-346.

Roberts, Matthew C. (2002). Homogenizing Heterogeneous Data: Inducing ARCH effects in futures data. Selected Paper presented at the annual meeting of the American Agricultural Economics Association. Tampa, Florida.

Robinson, P.M. (1994). Time series with strong dependence. In: C.A. Sims, Editor, Advances in Econometrics, Sixth World Congress, Cambridge University Press, Cambridge (1994), pp. 47-95.

Robinson, P.M. (1995). Gaussian semiparametric estimation of long range dependence, The Annals of Statistics 23, pp. 1630-1661.

Robinson, P.M. (1995). Log-Periodogram Regression of Time Series with Long Range Dependence. Annals of Statistics 23(): 1040-72.

Robinson, P.M. (2001). The Memory of stochastic volatility models, Journal of Econometrics 101 (2001), pp. 195-218.

Rogers, L.C.G. (1997). Arbitrage with Fractional Brownian Motion, Mathematical Finance 7 (1), 95-105.

Rogers, L.C.G. and Satchell (1991). Estimating Variance From High, Low and Closing Prices. Annals of Applied Probability,

Vol. 1, No. 4 , pp. 504-512.

Saphores J.-D., L. Khalaf, and D. Pelletier, 2002. On ARCH and Jumps in Natural Resource Prices, An Application To Pacific Northwest Stumpage Prices, American Journal of Agricultural Economics, 84(2): 387-400.

Shimotsu, K. (2006). Simple But Effective Tests of Spurious Long Memory, QED working paper #1101, Queen's University Department of Economics.

Smith, A. (2005). Level Shifts and the Illusion of Long Memory in Economic Time Series, Journal of Business and Economic Statistics, 2005, 23(3): 321-335.

Smith, A. (2005). Partially Overlapping Time Series: A New Model for Volatility Dynamics in Commodity Futures, Journal of Applied Econometrics, 20(3): 405-422.

Sorensen, Carsten. (2002) Modeling Seasonality in Agricultural Commodity Futures, Journal of Futures Markets, vol. 22: 393-426.

Sottinen, T. (2001) Fractional Brownian motion, random walks and binary market models, Finance and Stochastics 5(3).

Tanaka, K. (1999). The Nonstationary Fractional Unit Root. Econometric Theory, 15: 549-582.

Tanaka, K. (2004). Frequency domain and wavelet-based estimation for long memory plus noise models, in: Harvey, A.C., S.J. Koopman and N. Shephard, ed., State Space and Unobserved Component Models: Theory and Applications, Cambridge University Press: Cambridge, U.K.

Teyssiere, G. (2003). Interaction Models for common long-range dependence in asset price volatility. In: G. Rangarajan and M. Ding, Editors, Processes with Long Range Correlations: Theory and Applications Lecture Notes in Physics Vol. 621, Springer, Berlin (2003), pp. 251-269.

Turvey, C.G. (2007). A Note on Scaled Variance Ratio Estimation of the Hurst Exponent with Application to Agricultural Commodity Prices, Physica A, Vol.377(1): 155-165.

Veitch, D. and P. Abry (1999). A Wavelet based joint estimator of the parameters of long-range dependence, IEEE Transactions on Information Theory 45 (1999), pp. 878-897.

Wang, D. and W. Tomek (2007). Commodity Prices and Unit Root Tests, American Journal of Agricultural Economics, forthcoming.

Wei, A., and R.M. Leuthold (2000). Agricultural Futures Prices and Long Memory Processes. Working paper, University of Illinois Office for Futures and Options Research, Paper Number 00-04, April 2000.

Whitcher, B., P. Guttorp, and D. B. Percival (2000). Multiscale Detection and Location of Multiple Variance Changes in the Presence of Long Memory. Journal of Statistical Computation and Simulation, 68 (1), 65-88.

Zivot, E., and D.W.K. Andrews (1992). Further Evidence on the great crash, the oil-price shock, and the unit root hypothesis. Journal of Business and Economic Statistics 10, 251-270.

Table 1: Descriptive statistics of log-range price volatility in commodity futures contract time series data, T=4266, daily observations from 2/1988 to 1/2005

Futures contract	Mean	Std dev.	Skewness (Normal=0)	Kurtosis (Normal=3)
CBOT corn	0.015	0.008	2.070	8.417
CBOT soybeans	0.015	0.008	1.932	6.719
CME lean hogs	0.017	0.009	2.315	14.013
CME live cattle	0.011	0.005	1.494	3.077
KCBOT wheat	0.015	0.009	1.690	4.776
WCE canola	0.012	0.007	1.544	4.368
NYBOT cocoa	0.022	0.013	1.723	5.149
NYBOT coffee	0.028	0.018	2.175	9.204
NYBOT FCOJ	0.020	0.014	3.071	19.649
NYBOT cotton	0.018	0.011	2.295	12.069
NYBOT sugar#11	0.026	0.016	2.562	16.308

Table 2(a): Log-range volatility ARFIMA (p,d,q) model estimates, standard errors and hypothesis test results, for CME, CBOT, KCBOT and WCE commodities

Commodity futures contract	CBOT corn	CBOT soybeans	CME lean hogs	CME live cattle	KCBOT wheat	WCE canola
Long memory d	0.304	0.321	0.32	0.309	0.431	0.436
Correct standard error for d	0.0122	0.0122	0.0122	0.0122	0.0121	0.0123
Naïve standard error for d	0.0155	0.0155	0.0155	0.0155	0.0155	0.0155
Intercept	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001
AR1	-0.002	0.409	0.148	0.767	0.564	-0.630
	(0.202)	(0.111)	(0.176)	(0.073)	(0.036)	(0.062)
AR2	0.634	-0.014	-0.455	0.918	0.083	0.336
	(0.153)	(0.126)	(0.219)	(0.087)	(0.021)	(0.056)
AR3	0.190	0.649	-0.236	-0.65		
	(0.151)	(0.125)	(0.104)	(0.059)		
AR4	0.093	-0.125	0.536	-0.063		
	(0.028)	(0.065)	(0.13)	(0.024)		
AR5			0.042			
			(0.228)			
AR6			0.615			
			(0.126)			
AR7			0.140			
			(0.018)			
MA1	-0.146	-0.631	-0.265	-0.931	-0.834	0.392
	(0.203)	(0.109)	(0.176)	(0.07)	(0.032)	(0.056)
MA2	-0.661	0.09	0.463	-0.805		-0.539
	(0.169)	(0.136)	(0.239)	(0.094)		(0.048)
MA3	-0.073	-0.682	0.204	0.774		
	(0.015)	(0.130)	(0.120)	(0.067)		
MA4		0.341	-0.531			
		(0.07)	(0.134)			
MA5			0.044			
			(0.244)			
MA6			-0.598			
			(0.148)			

Seasonal (sinusoidal) coefficients are very small and not significantly different from zero, they are therefore omitted.

	CBOT corn	CBOT soybeans	CME lean hogs	CME live cattle	KCBOT wheat	WCE canola
Log-likelihood	14906.4	15144.8	14979.1	16885.3	8112.46	15035.5
Wald test Ho:d=0, model I(d)	24.95	26.35	26.27	25.36	23.15	35.79
Wald test Ho: d=0, model ARFIMA (p,d,q)	24.86	26.25	26.17	25.27	23.06	35.66
Shimotsu's adjusted KPSS test, Ho: d=0	0.14*	0.61***	0.46***	1.04***	0.04	0.02
Shimotsu's Phillips-Perron Z test, Ho: d=1	-1.46	-0.30	-0.60	0.52	-3.63**	-4.20***
Ohanissian-Russell-Tsay test, Ho: true long memory	4.42**	4.55**	4.06**	3.92**	2.82*	3.10*
Long memory true?	**Unclear**	**No**	**No**	**No**	**Yes**	**Yes**

Critical test values (exact values were computed and used in the analysis but approximate values are included here for convenience, source: Shimotsu 2006 Table 2): adjusted KPSS test 0.135 (10%), 0.17 (5%), 0.26 (1%) and adjusted Phillips-Perron Z test -3.09 (10%), -3.36 (5%), -3.9 (1%); and Chi-Square(1) for Ohanissian-Russell-Tsay Wald-type test 2.706 (10%), 3.841 (5%) and 6.635 (1%).

Table 2(b): Log-range volatility ARFIMA (p,d,q) model estimates, standard errors and hypothesis test results, for NYBOT commodities

Commodity futures contract	NYBOT cocoa	NYBOT coffee	NYBOT FCOJ	NYBOT cotton	NYBOT sugar#11
Long memory d	0.271	0.258	0.194	0.29	0.279
Correct standard error for d	0.0120	0.0119	0.0118	0.120	0.121
Naïve standard error for d	0.0155	0.0155	0.0155	0.0155	0.0155
Intercept	<0.001	<0.001	<0.001	<0.001	<0.001
AR1	0.624	0.001	0.308	0.337	0.529
	(0.145)	(0.110)	(0.541)	(0.033)	(0.073)
AR2	0.649	0.067	0.677	0.429	-0.109
	(0.383)	(0.059)	(0.539)	(0.037)	(0.031)
AR3	-0.058	0.141	-0.848	0.967	
	(0.236)	(0.031)	(0.045)	(0.032)	
AR4	-0.367	0.365	-0.105	-0.425	
	(0.172)	(0.032)	(0.017)	(0.076)	
AR5		0.634			
		(0.062)			
AR6		-0.526			
		(0.113)			
AR7					
MA1	-0.717	-0.065	-0.309	-0.442	-0.649
	(0.158)	(0.104)	(0.541)	(0.031)	(0.067)
MA2	-0.611	-0.065	-0.672	-0.412	0.104
	(0.400)	(0.057)	(0.526)	(0.037)	(0.031)
MA3	0.136	-0.122	0.013	0.887	-0.970
	(0.256)	(0.036)	(0.018)	(0.050)	(0.032)
MA4	0.352	-0.344		0.564	
	(0.184)	(0.033)		(0.069)	
MA5		-0.572			
		(0.058)			
MA6		0.592			
		(0.104)			

Seasonal (sinusoidal) coefficients are very small and not significantly different from zero, they are therefore omitted.

Log-likelihood	12902.4	11677.9	12427.04	13733.3	11791.5
Wald test Ho:d=0, model I(d)	22.24	21.18	15.92	23.80	22.90
Wald test Ho: d=0, model ARFIMA (p,d,q)	22.16	21.10	15.87	23.72	22.82
Shimotsu's adjusted KPSS test, Ho: d=0	0.48***	0.75***	0.76***	0.40***	0.28***
Shimotsu's Phillips-Perron Z test, Ho: d=1	-0.61	-0.58	-0.61	-0.80	-0.95
Ohanissian-Russell-Tsay test, Ho: true long memory	3.80*	5.21**	3.91**	4.50**	4.27**
Long memory true?	**No**	**No**	**No**	**No**	**No**

Critical test values (approximate, source: Shimotsu 2006 Table 2): adjusted KPSS test 0.135 (10%), 0.17 (5%), 0.26 (1%) and adjusted Phillips-Perron Z test -3.09 (10%), -3.36 (5%), -3.9 (1%); and Chi-Square(1) for Ohanissian-Russell-Tsay Wald-type test 2.706 (10%), 3.841 (5%) and 6.635 (1%).

Figure 1: Time series plot of daily log-range price volatility, CME lean hogs futures
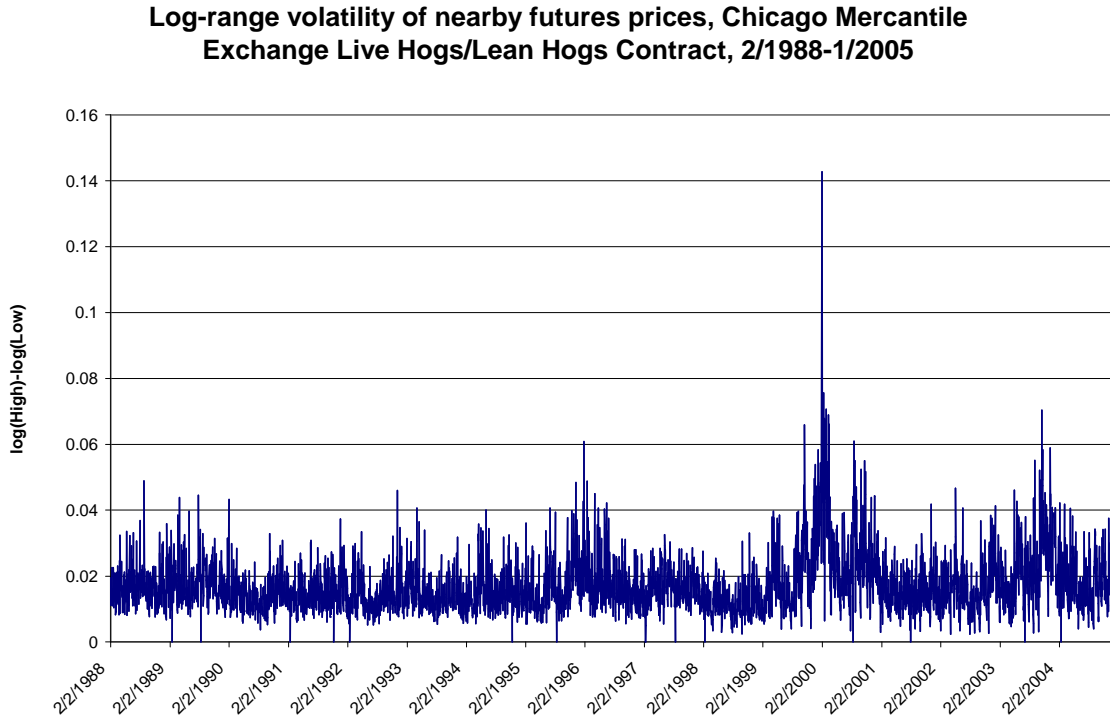
Log-range volatility of nearby futures prices, Chicago Mercantile Exchange Live Hogs/Lean Hogs Contract, 2/1988-1/2005

Figure 2: Time series plot of daily log-range price volatility, CME live cattle futures

Log-range price volatility of nearby futures prices, Chicago Mercantile Exchange Live Cattle Contract, 2/1988-1/2005

Figure 3: Time series plot of daily log-range price volatility, CBOT soybeans futures

Log-range price volatility of nearby futures prices,
Chicago Board of Trade Soybeans Contract, 2/1988-1/2005

Figure 4: Time series plot of daily log-range price volatility, CBOT corn futures

Log-range volatility of nearby futures prices,
Chicago Board of Trade corn contract, 2/1988-1/2005

Figure 5: Time series plot of daily log-range price volatility, KCBOT wheat futures


Log-range volatility of nearby futures prices,
Kansas City Board of Trade wheat futures, 2/1988-1/2005

Figure 6: Time series plot of daily log-range price volatility, WCE canola futures

Log-range volatility of nearby futures prices,
Winnipeg Commodity Exchange canola contract, 2/1988-1/2005

Figure 7: Time series plot of daily log-range price volatility, NYBoT cocoa futures

Log-range volatility of nearby futures prices,
New York Board of Trade cocoa contract, 2/1988-1/2005

Figure 8: Time series plot of daily log-range price volatility, NYBoT coffee futures

Log-range volatility of nearby futures prices,
New York Board of Trade coffee contract, 2/1988-1/2005

Figure 9: Time series plot of daily log-range price volatility, NYBoT cotton futures

Log-range volatility of nearby futures prices,
New York Board of Trade cotton contract, 2/1988-1/2005

Figure 10: Time series plot of daily log-range price volatility, NYBoT sugar#11 futures

Log-range volatility of nearby futures prices,
New York Board of Trade sugar#11 contract, 2/1988-1/2005
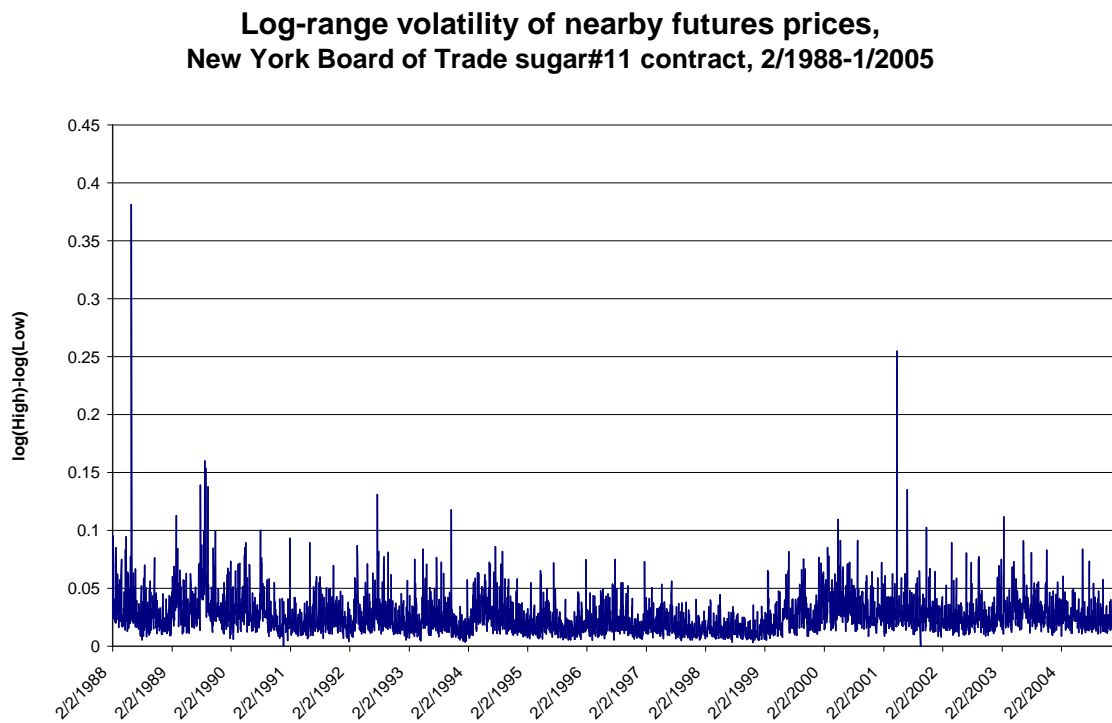
Figure 11: Time series plot of daily log-range price volatility, NYBoT frozen concentrated orange juice futures

Log-range volatility of nearby futures prices, New York Board of
Trade frozen concentrated orange juice contract, 2/1988-1/2005