A TEST OF THE COST-OF-CARRY RELATIONSHIP USING THE LONDON METAL EXCHANGE LEAD CONTRACT

RICHARD HEANEY

INTRODUCTION

This article tests the relationship between the major elements of the cost-of-carry relationship, spot price, futures price, interest rate to maturity, and stock level effects. The article draws on commodity prices obtained from the London Metals Exchange (LME) which exhibit useful characteristics for a test of the cost-of-carry model. Finally, the multivariate error-correction method of Johansen (1988) is relied upon for statistical testing.

Testing the cost-of-carry model poses four major problems. First, in most futures markets there are a limited number of maturity dates. This is apparent in Edwards and Canter (1995), Cho and McDougall (1990), and Fama and French (1987, 1988). As a result, futures price time series generally exhibit reducing time to maturity. Brenner and Kroner (1995) argue that this affects observed futures price variance because the futures price variance must tend toward zero as the time to maturity approaches zero. Although not a key problem in the cost-of-carry model tests reported

Helpful comments and suggestions were provided by Stephen Taylor, Peter Pope, John O'Hanlon, T. S. Ho, and Michael Mumford, participants at workshops held at the University of Lancaster, Australian National University, Monash University, and Macquarie University, and participants at the Northern Finance Association Conference held in Quebec, 1996.

■ *Richard Heaney is a Reader in the Department of Commerce at the Australian National University, Canberra, Australia.*

The Journal of Futures Markets, Vol. 18, No. 2, 177–200 (1998)

CCC 0270-7314/98/020177-24

in this article, this characteristic is important for unit root tests conducted on weekly observations. The futures prices used in this study avoid this problem because the futures prices are written with a fixed time to maturity of three months.

Second, futures prices are generally observed on a futures exchange and spot prices and stock level information must be obtained from other sources and other markets. The LME is of interest because both spot and futures contracts can be traded at the same time and in the same market. Further, stock levels are reported to traders on a regular basis and these stock level observations, in aggregate form, are used to capture stock level effects on the cost-of-carry relationship. The stock level represents the stocks accumulated and held in LME-approved warehouses as a result of trading on the LME and is not a measure of total world lead stock levels. Rather, it provides a measure of the stock level for the specific grades of lead deliverable under LME spot or futures contracts. Thus the futures price, the spot price, and the stock levels are all obtained from the one market for the one commodity.

Third, interest rate information is often difficult to obtain for tests of futures contract pricing. Generally, the required interest rate time to maturity reduces as the futures contract maturity date approaches and, in theory, this requires a complete term structure estimate for each trading day to allow matching of the maturity of the interest rate security with the futures contract maturity. The quotation of fixed time to maturity futures prices on the LME simplifies the observation of interest rates. In this paper, the interest rate carry cost is estimated using freely available three-month interest rates.

Fourth, the major components of the cost of carrying a commodity need to be identified. Although storage costs and insurance may be modeled using a constant term, the effect of convenience yield and inventory risk premium cannot be dealt with so easily. A model of the stock level effects, based on the total stock available in LME-approved warehouses, is included in the cost-of-carry model for this analysis.

Invariably, a single equation model is assumed for tests of the cost-of-carry model whether the tests are direct or indirect in nature. The variables underlying the cost-of-carry relationship are generally found to be nonstationary. Thus, assumptions implicit in single equation tests may not be valid. The Johansen method provides a test of the adequacy of the single equation cost-of-carry model while also providing tests of the cost-of-carry model restrictions. The method does this within a multivariate framework which specifically caters to the statistical problems associated with nonstationary time series.

The lead contract and the data used in the analysis are described in the following section. The third section deals with the cost-of-carry model and identifies the possibility of stock level effects. The statistical tests are described in the fourth section and the results of unit root and cointegration tests are discussed in the fifth. Conclusions and some suggestions for future research are presented finally.

THE LME LEAD CONTRACT AND DATA USED IN ANALYSIS

The commodities copper, tin, lead, and zinc were all traded over the study period, 1976–1995, but all except the lead contract were subject to changes in the definition of the deliverable spot asset. For example, the copper contract changed to High Grade in November 1981 and then to Grade A in June 1986. The tin contract changed to 99.85% in June 1989. This occurred after a complete cessation of trading in the tin contract with the collapse of the International Tin Council in the 1980s, (Sephton and Cochrane, 1991). Finally, a number of changes occurred with the zinc contract, including the introduction of a new zinc contract in June 1986 (Sephton and Cochrane, 1991). Further, contracts instituted after 1975, such as the aluminum contracts and the nickel contract, were not included because of the limited time series available for these contracts. As a result, this paper focuses on the lead contract which was officially introduced to the LME in 1920, has traded continuously except for closure during the Second World War, and exhibits no fundamental change in definition of the deliverable asset over the study period.

Quarterly data from March 1976 to June 1995 are used in analysis. Because LME-approved lead stocks held at LME-approved warehouses are observed at weekly intervals, spot price, futures price, and interest rates are matched against the stock level to give a weekly time series. Quarterly observations are generated using the observation for the last week in March, June, September, or December. The lead spot and futures prices are unofficial LME market closing prices expressed in United Kingdom Pounds (GBP) and obtained from Datastream.[1] Weekly observations of LME-approved stocks held at LME-approved warehouses are obtained from the LME. These are similar in nature to the certified stocks described in Wright and Williams (1989) though the LME stock levels are driven both by spot (prompt) and futures deliveries. The U.K. Treasury bill three-month mid-rate and the Eurocurrency Sterling three-month

[1]Price changes more than four standard deviations from the mean are checked against the Metal Bulletin prices and, where inconsistent, the Metal Bulletin price is used.

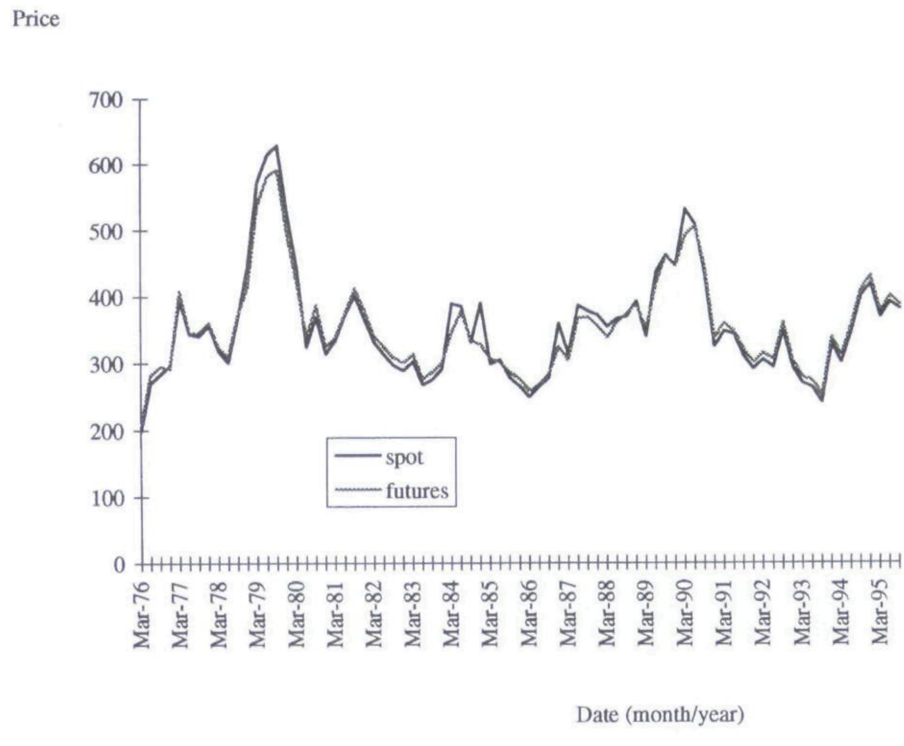

FIGURE 1
Lead spot prices and futures prices (GBPs per tonne).

mid-rate are used as estimates of the risk-free rate.

The time series characteristics of the data can be noted from Figures 1–3. Figure 1 plots the lead spot and futures prices over the study period. Spot and futures prices tend to track each other fairly closely over the study period, although the futures price generally exceeds the spot price where stock levels are high, and the reverse occurs when stock levels are unusually low. This is particularly evident in Figure 2, where two graphs are provided. The first plots the relative basis and the second is a plot of the level of LME lead stocks. In this figure a negative relative basis is only observed during periods of relatively low stock levels. The U.K. Treasury bill three-month mid-rate and the Eurocurrency Sterling three-month mid-rate are also plotted in Figure 3. Descriptive statistics for these variables are reported in Tables I and II.

One characteristic of markets, commodity markets included, is the possibility of manipulation.[2] One such attempt led to the closure of the

[2]I am indebted to Steven Taylor, who pointed out this possibility. This is arguably a characteristic of commodity markets though it is virtually impossible to identify ex ante and difficult to identify ex

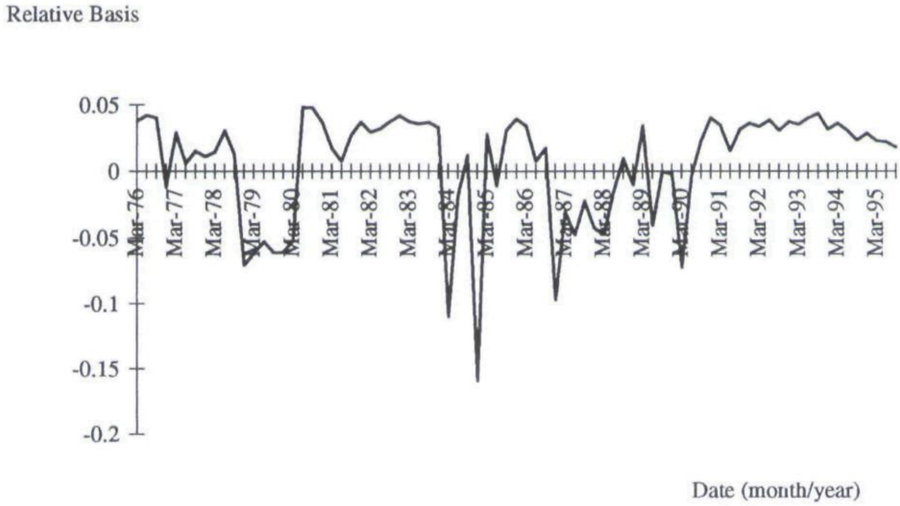

A

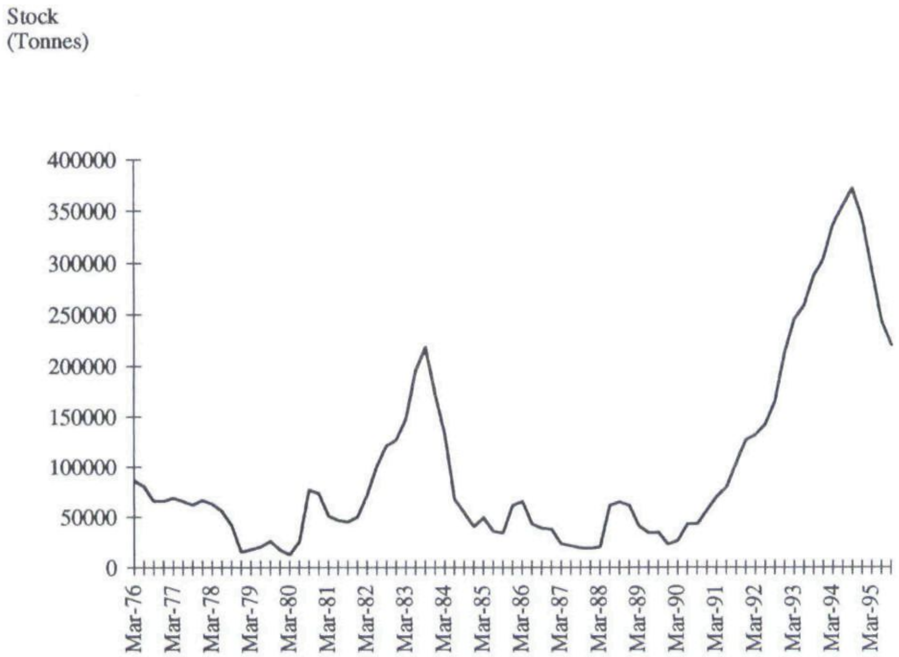

B

FIGURE 2
A: Lead relative basis [(futures − spot)/spot]. B: LME-approved lead stocks.

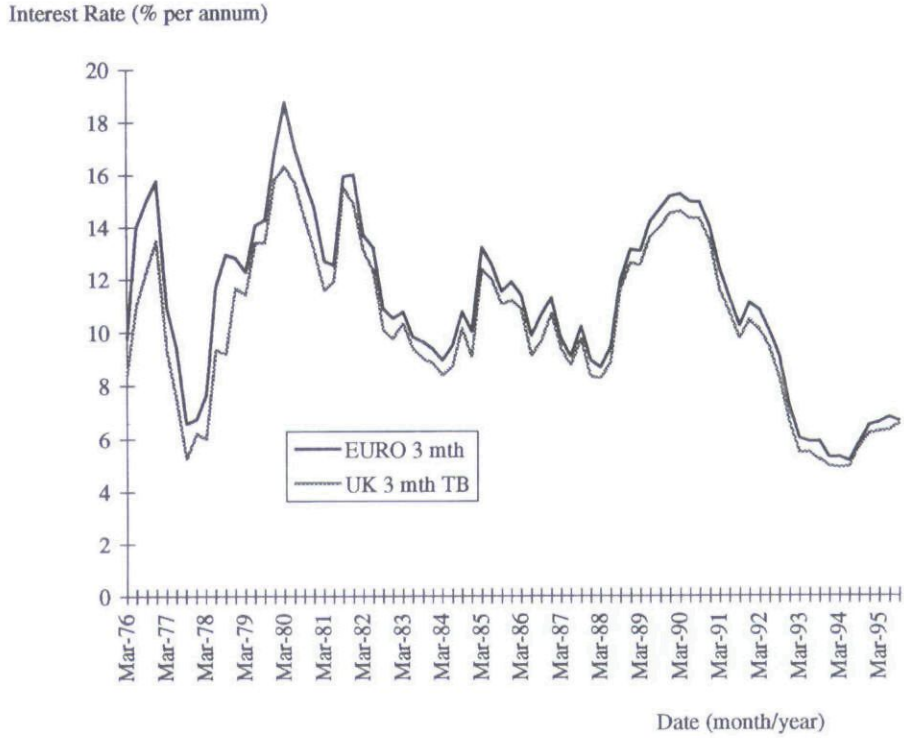

FIGURE 3

U.K. Treasury bill three-month mid-rate and Euro Sterling three-month mid-rate. Note: Interest rate is the three-month mid-rate expressed as a nominal rate per annum.

LME tin market for some years in the 1980s. Bailey and Ng (1991) argue that manipulation could result in the existence of a default premium in traded prices and it is argued that the default premium effect may be reduced with the introduction of a futures margining and deposits system. Though a margining system was introduced at the LME in 1986, it is not expected to have a material effect in this analysis (Benninga and Protopapadakis, 1994; Cornell and Reinganum, 1981; French, 1983; Pindyk, 1993). Regardless, the boundaries between manipulation and ordinary

post other than in the most blatant examples. Manipulation has occurred from time to time since the LME's inception, although the most obvious form of manipulation, based on acquiring available stocks and thereby restricting supply, is limited by the traders' ability to fund commodity acquisitions. This form of manipulation suggests traders may require a price discount to compensate for the possibility of future artificially inflated prices, but this is not the only form of manipulation. An alternative includes taking short positions in futures contracts and selling stocks, during periods when the spot is thinly traded. The sale of stocks forces down the spot price and maximizes profits from the futures contracts. This form of manipulation might be consistent with a premium impounded in the spot price.

TABLE I

Descriptive Statistics for Levels
(Quarterly Observations, N = 79; Weekly Observations, N = 1029)

	Mean	SD	Excess Skewness	Kurtosis	Maximum	Minimum
U.K. Treasury bill three-month interest rates						
Quarterly	0.0253	0.0074	−0.0368	−0.7982	0.0399	0.0122
Weekly	0.0254	0.0072	−0.1155	−0.7543	0.0401	0.0110
Euro Sterling three-month interest rates						
Quarterly	0.0274	0.0079	−0.0294	−0.6674	0.0458	0.0128
Weekly	0.0274	0.0078	−0.0419	−0.5446	0.0509	0.0126
Natural log of the stock level						
Quarterly	11.1390	0.8655	0.2391	−0.7041	12.8260	9.4375
Weekly	11.1180	0.8747	0.1880	−0.6647	12.8260	8.9522
Natural log of the spot price						
Quarterly	5.8436	0.2135	0.5745	0.7785	6.4425	5.2920
Weekly	5.8343	0.2197	0.4616	1.0584	6.6908	5.1052
Natural log of the three-month futures price						
Quarterly	5.8509	0.1914	0.4882	0.7615	6.3793	5.3291
Weekly	5.8424	0.1971	0.3126	1.0999	6.4599	5.1446

Notes: The U.K. Treasury bill and Euro Sterling three-month mid-rates are continuously compounding rates of return for the three-month period. The stock level is the number of tonnes of LME-approved stock held in LME-approved warehouses. The interest rates, futures prices, and spot prices are obtained from Datastream, while the stock levels are obtained from the LME.

trading become blurred, especially where large transactions are involved. With greater internationalization of the metal markets, greater efficiency in production and transport technologies, and fairly widespread mining,[3] the likelihood of manipulation in the lead market is fairly low (Radetzki, 1990).

THE PRICING OF FUTURES CONTRACTS

Futures pricing can be viewed in terms of the cost of storage. Working (1949) introduces this concept by analyzing the prices of two futures contracts with differing maturity dates. It is argued that the futures price difference is the cost of carrying the underlying asset from one futures contract maturity date to the next. This futures price difference, or price of storage, may be either positive or negative though Working (1949)

[3]LME-accepted brands of lead are derived from a wide range of countries including Australia, Africa, Canada, Europe, Russia, South America, and the United States of America.

TABLE II

Autocorrelations over Levels for Lags 1–6 (Quarterly Observations, N = 79; Weekly Observations, N = 1029)

	$\rho1$	$\rho2$	$\rho3$	$\rho4$	$\rho5$	$\rho6$	χ^2 (6)
U.K. Treasury bill three-month interest rates							
Quarterly	0.889[a]	0.754[a]	0.593[a]	0.469[a]	0.347[a]	0.264[a]	177.02[a]
Weekly	0.993[a]	0.985[a]	0.977[a]	0.967[a]	0.957[a]	0.947[a]	5853.46[a]
Euro Sterling three-month interest rates							
Quarterly	0.884[a]	0.728[a]	0.568[a]	0.442[a]	0.346[a]	0.294[a]	169.96[a]
Weekly	0.990[a]	0.981[a]	0.970[a]	0.960[a]	0.949[a]	0.939[a]	5779.20[a]
Natural log of the stock level							
Quarterly	0.922[a]	0.805[a]	0.701[a]	0.607[a]	0.506[a]	0.409[a]	233.20[a]
Weekly	0.997[a]	0.993[a]	0.988[a]	0.982[a]	0.975[a]	0.969[a]	6009.41[a]
Natural log of the spot price							
Quarterly	0.751[a]	0.574[a]	0.337[a]	0.112	−0.003	−0.039	84.51[a]
Weekly	0.977[a]	0.955[a]	0.933[a]	0.909[a]	0.886[a]	0.863[a]	5269.66[a]
Natural log of the three-month futures price							
Quarterly	0.757[a]	0.571[a]	0.364[a]	0.150	0.048	0.008	87.43[a]
Weekly	0.979[a]	0.959[a]	0.938[a]	0.917[a]	0.894[a]	0.871[a]	5333.94[a]

Notes: Approximate critical values for the correlation coefficients are 0.225 at the 5% level for the quarterly times series and 0.062 for the weekly time series. The chi-square critical value is χ^2 (6, 5%) = 12.592.
[a]Significant at the 5% level.

argues that the magnitude of the price of storage is dependent on the level of stocks. The same argument can be applied to the difference between the spot price and a futures price observed at one time. In this case, the storage period is now the time to maturity of the futures contract.

The early literature focuses on the relationship between the spot price and futures price, assuming the current futures price is equal to the expected spot price (Brennan, 1958; Telser, 1958). Brennan (1958) argues that the supply of storage schedule is essentially fixed while the demand schedule varies over time. Thus, demand changes map out the supply of storage schedule determining the relationship between spot price and futures price. Working (1949) and Brennan (1958) describe an upward sloping supply of storage schedule consisting of three components: physical storage costs, convenience yield, and risk premium for holding stocks.[4]

The convenience yield is the yield to suppliers arising from physical possession of the stocks. Where stocks are held, specific client needs can

[4]Working (1949) alludes to the effect of costs of storage exceeding the physical storage cost as the level of total stocks increases. Brennan (1958) defines the effect as marginal risk aversion cost.

be met and sudden increases in demand can be accommodated without disrupting production. This characteristic of stocks becomes more valuable as total stocks diminish. As the convenience yield increases with falls in the level of stocks, the total costs of storage may become negative. At high levels of stock it is argued that the convenience yield approaches zero, while the risk premium for holding stocks tends to increase. The higher the level of stocks, the more risky the investment in stocks and the greater the compensation required for holding the stocks. The result is a nonlinear cost-of-storage function which explains the relationship between price spread and stock level.

Extensions to this model include the effect of changing stocks levels (Weymar, 1966), government intervention (Kofi, 1973; Telser, 1958), and the effect of changes in consumption on convenience yields (Telser, 1958). These factors are not modeled explicitly in this study. Stocks levels tend to be fairly stable over the relatively short period of three months to futures contract maturity and so the Weymar effect is not expected to be substantial in this study. Further, the pricing problems arising from government intervention observed by Kofi (1973) and Telser (1958) in the United States should not affect this study given the U.K. government's aversion to market regulation over the study period. Finally, consumption of lead, deliverable under LME contracts, is not included in the analysis due to the lack of appropriate consumption data and the rather mixed results reported for the cotton and wheat contracts as reported in Telser (1958).

Although not specifically addressed in this article, a large body of literature focuses on the relationship between the current futures price and the expected futures price, or expected spot price, at maturity (Fort and Quirk, 1988; Lien, 1987; Hirshleifer, 1988, 1989). The more recent theoretical papers suggest that the assumption that current futures price is the best estimate of the maturity futures price, or maturity spot price, is not always appropriate. Empirical tests of the predictive ability of LME commodity forwards also exhibit considerable variation. These include Goss (1981, 1983), Hsieh and Kulatilaka (1982), MacDonald and Taylor (1988a, 1988b, 1989), Sephton and Cochrane (1990, 1991), Chowdhury (1991), and Moore and Cullen (1995). Similar tests have been applied to a range of commodity and financial futures traded on markets other than the LME (Crowder and Hamed, 1993; Cooper, 1993; Krehbiel and Adkins, 1993; Moore, 1994). The variation in results may be due to differences in statistical techniques and time periods chosen for analysis though it appears the matter is still unsettled. Although this paper does not specifically deal with the question of predictive ability, it is possible that stock level effects could explain some of the variation in predictive ability observed in the literature.

This article deals with the relationship between the currently observed spot price and the currently observed futures price. The purchaser of the metal can either take out a futures contract or purchase the metal now and store it. This suggests a pricing mechanism called the cost-of-carry model. The cost-of-carry model can be expressed in terms of discrete rates of return or continuously compounding returns. Continuously compounding returns are chosen for analysis here to simplify statistical analysis as this ensures the model is log-linear and therefore amenable to widely used cointegration tests. This choice also simplifies the modeling of storage costs as it assumes costs are proportional to the spot price, thus mitigating the problem of modeling fixed costs which change over the study period. The cost-of-carry model is defined:

$$F(t, T) = P(t) \exp[rf(t,T) + s(t, T) + sle(t)] \tag{1}$$

where $F(t, T)$ = futures price quoted at time t, maturing at time T; $P(t)$ = underlying metal price quoted at time t; $rf(t, T)$ = continuously compounding risk-free rate of return for the period t to T; $s(t, T)$ = costs of storage from time t to T; $sle(t)$ = stock level effect at time t; and exp[.] = exponential function.

Taking logs gives a linear model:

$$\ln[F(t, T)] = \ln[P(t)] + rf(t, T) + s(t, T) + sle(t) \tag{2}$$

The risk-free rate of return represents the opportunity cost of investing cash in the metal now rather than using a futures contract. Storage costs include e.g., insurance, warehousing, and transport, and are assumed to be a fixed proportion of the spot price. Stock level effects including convenience yield and risk premium, described in Working (1949), Brennan (1958), Telser (1958), French (1986), Fama and French (1987, 1988), and Cho and McDougal (1990), accrue to holding stocks but do not flow through to the futures contract.

Tests for stock level effects include indirect tests (Fama and French, 1988) and direct tests (Brennan, 1958). Indirect tests provide some insight into the existence of stock level effects, but these tests do not provide specific models for use in pricing or hedging decisions. Direct tests explicitly model the effect, generally assuming stock level effects are some linear or log-linear function of stock levels. The use of LME time series in this article tends to minimize the major difficulties with the direct method: identification of an appropriate stock measure and identification of the price of the commodity underlying the futures contract.

The stock level effects are modeled as a linear function of the log of stocks giving a stock effect:

$$sle(t) = \alpha \ln(S(t)) - \gamma \tag{3}$$

where γ, α = the parameters of the model and ln(.) = the natural logarithm.

The restriction $\alpha > 0$ is required to ensure the model is consistent with the Working (1949) convenience yield and risk premium effect.[5] The term γ must be negatively signed if the futures price is to fall below the spot price.

Combining eqs. (1)–(3) gives the final cost-of-carry model:

$$\ln(F(t, T)) = \ln(P(t)) + rf(t, T) + \gamma^* + \alpha \ln(S(t)) \tag{4}$$

where γ^* = the convenience yield parameter adjusted for storage costs $(s(t, T) - \gamma)$.

When stocks are sufficiently low, the convenience yield tends to overwhelm the other storage costs resulting in a futures price less than the spot price. If the stock levels are sufficiently high, this effect is small and the difference between the futures price and the spot price is essentially the cost of storage plus any stock risk premium effect.

STATISTICAL TESTS: UNIT ROOT TESTS AND TESTS FOR COINTEGRATION

The autocorrelations reported in Table II suggest the possibility of unit root or near unit root processes. This can have implications for later analysis. Therefore, the augmented Dickey-Fuller and the Phillips-Perron tests are run to test for the existence of one unit root in the time series. The augmented Dickey-Fuller and the Phillips-Perron test statistics are obtained from the regression:

$$y_t = \mu + \alpha y_{t-1} + \tau(t - N/2) + \varepsilon_t \tag{5}$$

where y_t = time series observation at time t; μ = intercept term; α = lagged level term parameter; t = time trend variable ranging from 1 to N; N = total number of observations; and ε_t = regression residuals.

[5]If $\alpha > 0$, then the convenience yield (cy) component of stock level effects behaves as expected: $dcy(t)/dS(t) = \alpha/S(t) > 0$ and $d^2cy(t)/dS(t)^2 = -\alpha/(S(t)^2) < 0$.

The augmented Dickey-Fuller t-test, $(Z(t_\alpha))$, adjusted for serial correlation, and the Phillips-Perron t-test, $(Z(t_\alpha))$, and "slope" parameter test, $(Z(\alpha))$, adjusted for serial correlation and heteroskedasticity, are reported in the next section.[6]

It is likely that the cost-of-carry model will hold in a long-term sense, given the size of the market and the number of active traders. If one or more of the variables making up the cost-of-carry relationship changes, due to some exogenous shock, traders will react quickly to the possibility of profitable trading, leading to a return to this long-run relationship. The cost-of-carry model says nothing about which variable (or variables) adjusts when a shock leads to a divergence from the cost-of-carry relationship. Although it is unlikely that interest rates would change, it is possible that some combination of the spot price, futures price, and inventory levels may change as arbitrage occurs.

A single equation approach to testing the cost-of-carry relationship could be implemented, but the assumption of exogeneity implicit in this approach (Ericsson, 1992) appears to place inappropriate a priori restrictions on the relationship between the cost-of-carry variables. Is it possible to say a priori that the spot price and stock level are unaffected by the cost-of-carry divergence? If the variables are not exogenous, a single equation model is invalid (Bannerjee and Hendry, 1992; Kremers, Ericsson, and Dolado, 1992; Urbain, 1992).

The method described in Johansen (1988) and Johansen and Juselius (1990) provides an alternative to single equation models which does not assume exogeneity of the underlying variables and can deal with unit root or near unit root processes (Bannerjee and Hendry, 1992). The Johansen approach also provides statistical tests for the number of cointegrating vectors and cointegrating vector restrictions.

Following Johansen (1988), the vector autoregression (VAR) is defined:

$$X_t = A_1 X_{t-1} + \ldots + A_{k-1} X_{t-k+1} + A_{t-k} X_{t-k} + \mu + \varepsilon_t \qquad (6)$$

This can be rearranged to give:

$$\Delta X_t = \Gamma_1 \Delta X_{t-1} + \ldots + \Gamma_{k-1} \Delta X_{t-k+1} + \Pi X_{t-k} + \mu + \varepsilon_t \qquad (7)$$

where

[6]Due to the nature of the Phillips-Perron test, it is able to deal with a wide range of serial correlation and heteroskedasticity in the residuals. The two tests are reported to allow for the possibility of size distortions with the Phillips-Perron test when the residuals follow a negative moving average process (Banerjee, Dolado, Galbraith, and Hendry, 1993).

X_t = the vector of variables $(F(t, T)\ P(t)\ rf(t, T)\ S(t))'$

$\Gamma_i = (A_1 + \ldots + A_i - I)$

I = identity matrix

$\Pi = (A_1 + \ldots + A_k - I)$ and $\alpha\beta' = \Pi$

Δ = change in the vector of variables $(\Delta X_t = X_t - X_{t-1})$

The term Π is a measure of the long-run relationship between the variables remaining after removal of short-term effects. In effect the rank, r, of this matrix, Π, determines the number of cointegrating vectors. Given the four variables used in this study, there are five possible values of r. If $r = 0$, the variables are nonstationary and there are no cointegrating vectors. If $r = 4$, all four variables are stationary. If $0 < r < 4$, there are r cointegrating vectors. Once the number of cointegrating vectors is identified, the matrix, Π, can then be defined in terms of two matrices, α and β, where $\alpha\beta' = \Pi$. The $(4 \times r)$ matrix, α, provides estimates of the speed of adjustment or loading attached to the variables and thus generates an estimate of the rate at which the variables react to departures from the long-run relationship. If the variables are exogenous, the speed of adjustment is zero. The $(4 \times r)$ matrix, β, is the set of cointegrating vectors, which may range from one cointegrating vector through to a maximum of three cointegrating vectors, depending on the number of cointegrating vectors identified in the VAR.

Before tests may be undertaken, the number of lags, k, must be selected. This choice focuses on the VAR to ensure that it captures the time series nature of the data. Enders (1995) suggests a set of three criteria for use in this choice: the method suggested in Sims (1980), the Akaike Information Criterion (AIC) and the Schwartz Bayesian Criterion (SBC). A further criterion used in this paper relies on chi-square tests for serial correlation in the VAR residuals. Cheung and Lai (1993) observe that cointegration tests are more sensitive to underparameterization and so a conservative approach is taken to avoid the inclusion of insufficient lagged terms.

Once the number of lag terms in the VAR is identified it is necessary to statistically test whether the intercept term lies wholly within the cointegrating vector. This is an important question because both storage costs and the convenience yield model indicate that the cointegrating vector should include an intercept term. Further, from a statistical testing point of view, Johansen and Juselius (1990), Lloyd (1994), and Dickey and Rossana (1994) show that this choice determines the critical values used in tests for the number of cointegrating vectors. The adjusted VAR takes the form:

$$\Delta X_t = \Gamma_1 \Delta X_{t-1} + \ldots + \Gamma_{k-1} \Delta X_{t-k+1} + \Pi^* X^*_{t-k} + \varepsilon_t \qquad (8)$$

where $X^*_{t-k} = (X_t\ 1)$. The X^*_{t-k} matrix includes a vector of ones to capture the intercept term. Johansen and Juselius (1990) describe a likelihood ratio test to test whether the intercept term in the VAR, μ, is validly restricted to the cointegrating vector. The first equation is the unrestricted eq. (7) with the intercept appearing in the VAR and the second equation is the restricted eq. (8) with the intercept constrained to the cointegrating vector. The test is a likelihood ratio test defined in terms of eigenvalues based on the matrix Π from both the unrestricted and restricted equations. The eigenvalues from the restricted and unrestricted models are ranked from largest to smallest, with $\lambda_1 > \lambda_2 > \ldots > \lambda_p$ for the unrestricted equation and $\lambda_1^* > \lambda_2^* > \ldots > \lambda_p^*$ for the restricted equation. The likelihood ratio test[7] takes the form:

$$-2 \ln(Q) = -N \sum_{i=r+1}^{p} \ln[(1 - \lambda_i^*)/(1 - \lambda_i)] \qquad (9)$$

Given the VAR form which best fits the data, it is possible to test for the number of cointegrating vectors and also test the cost-of-carry restrictions which apply to cointegrating vectors. Johansen (1988) and Johansen and Juselius (1990) describe the test for the number of cointegrating vectors, which is a likelihood ratio test based on eigenvalues from the Π matrix. There are two tests generally reported. The first is the trace test and the second is the maximal eigenvalue test. These are defined as follows:

$$\begin{aligned} &\text{Trace} = -N^*\{\ln(1 - \lambda_{r+1}) + \ldots + \ln(1 - \lambda_p)\} \\ &\text{Maximal eigenvalue} = -N^* \ln(1 - \lambda_{r+1}) \end{aligned} \qquad (10)$$

where p = the number of equations in the VAR, one each for the futures price, spot price, interest rate, and stock level; r = the number of cointegrating vectors; λ_i = the ith eigenvalue where the eigenvalues are ranked $\lambda_1 > \lambda_2 > \ldots > \lambda_n$; and N = the number of observations adjusted for observations lost due to inclusion of lags.

Statistical tests of the cost-of-carry model restrictions may be conducted also. These tests concentrate on the cointegrating vector, β. The statistical test again concentrates on the difference between the restricted and the unrestricted model. The chi-square test is defined:

$$-2 \ln(Q) = -N \ln[(1 - \lambda_i^*)/(1 - \lambda_i)] \qquad (11)$$

[7]This test is χ^2 distributed with $(p - r)$ degrees of freedom.

where the maximum eigenvalue, λ_1^*, is obtained from the restricted model, and the maximum eigenvalue, λ_1, is obtained from the unrestricted model. The number of degrees of freedom for the test is $(p + 1 - s)$, where p is the number of equations in the VAR. $(p + 1)$ is used to include the effect of the constant term in the cointegrating vector. The parameter s refers to the number of columns in the restriction matrix. For a test of equality of the futures, spot, and interest rate parameters, the restriction matrix has three columns. The rows of the matrix are, respectively, (1 0 0), (−1 0 0), (−1 0 0), (0 1 0), and (0 0 1). This ensures that there are only three parameters to estimate: one for the three variables futures price, spot price, and interest rate, and one each for the stock level and the constant.

Tests can be conducted also on the speed of adjustment parameters, α. This provides the facility to test for exogeneity. If the speed of adjustment parameter for a variable is zero, the cointegrating vector does not explain changes in the variable and the variable is said to be exogenous with respect to the cointegrating vector (Ericsson, 1992; Kremers et al., 1992; Urbain, 1992). Given eq. (4), the cost-of-carry model divergence is defined:

$$\ln(F(t, T)) - \ln(P(t)) - rf(t, T) - \gamma^* - \alpha \ln(S(t)) \qquad (12)$$

For example, the speed of adjustment for the interest rate variable is zero if the interest rate is unaffected by cost-of-carry model divergence. Alternatively, if the futures price always adjusts to reduce cost-of-carry model divergence, the futures price speed of adjustment is negative. Further, the speed of adjustment for the spot price is positive if the spot price adjusts to remove divergence. The speed of adjustment test is based also on eigenvalues drawn from restricted and unrestricted models, though, in this case, the speed of adjustment elements is set to zero in the restricted model. The degrees of freedom for this test is $(p - m)$, where p is the number of equations and m is the number of columns in the restrictions matrix. For example, $(p - m)$ equals three for the restriction, $\alpha 1 = \alpha 2 = \alpha 3 = 0$, and one for the restriction, $\alpha 4 = 0$.

EMPIRICAL RESULTS

Unit Root Tests

Results of the unit root tests[8] are reported in Table III. The null hypothesis of one unit root cannot be rejected at the 5% level for the interest

[8]These tests are conducted using the SHAZAM COINT instruction, which calculates the Phillips-Perron adjustment using the Newey and West (1987) variance estimation technique.

TABLE III

Unit Root Tests: Augmented Dickey-Fuller Test Statistic, $Z(t_\alpha)$, and Phillips-Perron Test Statistic, $Z(\alpha)$ (Quarterly and Weekly Observations)

	Quarterly Time Series			*Weekly Time Series*		
	1 lag	*5 lags*	*10 lags*	*12 lags*	*26 lags*	*52 lags*
U.K. Treasury bill three-month interest rates						
ADF (t_α) statistics	−2.45	−2.40	−2.60	−2.58	−2.68	−2.65
PP(α) statistics	−11.11	−12.33	−10.84	−10.72	−11.46	−12.17
PP(t_α) statistics	−2.46	−2.58	−2.43	−2.31	−2.39	−2.46
Euro Sterling three-month interest rates						
ADF (t_α) statistics	−2.87	−2.02	−2.06	−3.10	−2.95	−2.57
PP(α) statistics	−13.77	−13.86	−12.47	−13.72	−15.36	−15.76
PP(t_α) statistics	−2.74	−2.75	−2.62	−2.64	−2.79	−2.83
Natural log of the stock level						
ADF(t_α) statistics	−2.63	−2.18	−2.33	−2.34	−2.46	−2.13
PP(α) statistics	−9.08	−9.23	−8.17	−7.15	−8.66	−8.63
PP(t_α) statistics	−2.21	−2.23	−2.11	−1.95	−2.14	−2.13
Natural log of the spot price						
ADF(t_α) statistics	−3.05	−2.58	−3.00	−3.73[a]	−3.12	−3.75[a]
PP(α) statistics	−19.97	−22.61[a]	−19.53	−23.94[a]	−24.13[a]	−27.42[a]
PP(t_α) statistics	−3.69[a]	−3.85[a]	−3.67[a]	−3.93[a]	−3.94[a]	−4.13[a]
Natural log of the three-month futures price						
ADF(t_α) statistics	−3.05	−2.55	−2.93	−3.83[a]	−3.08	−3.79[a]
PP(α) statistics	−19.97	−22.38[a]	−20.34[a]	−23.40[a]	−23.91[a]	−26.40[a]
PP(t_α) statistics	−3.71[a]	−3.85[a]	−3.73[a]	−3.95[a]	−3.98[a]	−4.12[a]

Notes: The model includes intercept and time trend. These tests are repeated using a model which includes intercept, but no time trend, with little change in results. The t-statistic is reported for the augmented Dickey-Fuller (ADF) test, and the t-statistic and parameter value are reported for the Phillips-Perron (PP) test. Critical values for the 5% level of significance for the t-statistic used in the augmented Dickey-Fuller test are −3.41 for weekly and −3.47 for quarterly data and for the Phillips-Perron parameter estimate-based test the critical values are −21.7 for weekly and −20.3 for quarterly data (Fuller, 1976). Number of lags or window size used in the test is indicated in the column heading.
[a]Significant at the 5% level.

rate and stock level time series and this result does not appear to be affected by the chosen lag structure. The finding of one unit root in the interest rate time series is also consistent with the literature (Hall, Anderson, and Granger, 1992; MacDonald and Murphy, 1989; Shea, 1992).

The results of unit root tests for the lead spot price and the lead forward price vary with the statistical test used and the time interval chosen. With quarterly observations the augmented Dickey-Fuller test cannot reject the null of one unit root in the time series, although the Phillips-Perron tests reject the unit root hypothesis where lags of 5 quarters or 10 quarters are chosen. For weekly observations, the Phillips-Perron tests reject the unit root hypothesis given lags of 12, 26, and 52 weeks, though rejection varies with the lag choice for the Dickey-Fuller test. The Dickey-Fuller results for weekly data are unexpected given that Pierse and Snell (1995) suggest the power of unit root tests is not expected to change with increases in sampling frequency.[9] It is also possible that the Phillips-Perron test results may be due to size distortions which are greatest when residuals follow a negative first-order moving average process (Banerjee et al., 1993), though there is little evidence of this process in the residuals.

The evidence of rejection of the null hypothesis of a unit root is fairly rare in the commodity price literature. For example, Chowdhury (1991), Franses and Kofman (1991), Krehbiel and Adkins (1993), MacDonald and Taylor (1988b), and Moore and Cullen (1995) do not reject the null hypothesis of one unit root in commodity price time series, although Moore and Cullen (1995) suggest the evidence for lead is "marginal" when using a model which includes intercept and time trend.

Although the augmented Dickey-Fuller results for quarterly data are consistent with one unit root, this does not follow for weekly data nor for the Phillips-Perron tests.[10] The finding of one unit root in the spot and futures prices may be a function of the study period chosen, the choice of the statistical test, and lag length selection, though further analysis of this question is left to future research. Regardless of the existence of unit roots in the process, the correlations reported in Table II suggest that the time series are persistent and at least indicative of near unit root processes if not unit root processes.

Cointegration Tests

There is little evidence of use of the Johansen test in tests of the cost-of-carry model, though this technique has been used in tests of market efficiency where the test focuses on cointegration between the futures price and the spot price. In both Krehbiel and Adkins (1993) and Crowder and Hamed (1993), cointegration between the spot and futures price is rejected while rejection is observed also with analyses based on the simpler Engle and Granger (1987) approach (see Chowdhury, 1991; MacDonald and Taylor, 1988a). The level of rejection may not be surprising where costs of carry are nonstationary, given the discussion in Brenner and Kroner (1995). If one or more of the carrying cost components are themselves

[9]This may be due to the fairly small sample size of 79 observations for the quarterly time series.
[10]This discrepancy may be due to differences in variance calculation method or choice of time series, though the test type does not seem to matter as both the slope parameter-based test and adjusted t-test provide similar results.

TABLE IV

Eigenvalues, Trace, and Maximal Eigenvalue Tests
U.K. Treasury Note Three-Month Mid-Rate

Null	*Eigenvalue*	*Maximal*	*Critical Value*	*Trace*	*Critical Value*	*N*
$r = 0$	0.33022	30.86[a]	28.14	61.87[a]	53.12	77
$r = 1$	0.17992	15.27	22.00	31.01	34.91	
$r = 2$	0.13920	11.54	15.67	15.74	19.96	
$r = 3$	0.05303	4.20	9.24	4.20	9.24	

Notes: The tests are performed sequentially at $r = 0$, then $r = 1$, and so on, until the test is no longer rejected. The trace statistic tests for the number of cointegrating vectors. At each value of r the null hypothesis is r or less than r cointegrating vectors. For example, if $r = 0$ is rejected but $r = 1$ cannot be rejected, then there is one cointegrating vector. If $r = 0$ and $r = 1$ are rejected but $r = 2$ is not rejected, then there are two cointegrating vectors. The maximal eigenvalue test has a null of r cointegrating vectors against an alternative of $r + 1$ cointegrating vectors. Thus, if the maximal eigenvalue test is rejected at $t = 0$ but not elsewhere, only one cointegrating vector exists in the process. The two statistics are defined:

$$\text{Trace} = -N^* \{\ln(1 - \lambda_{r+1}) + \ldots + \ln(1 - \lambda_p)\}$$
$$\text{Maximal eigenvalue} = -N^* \ln[(1 - \lambda_{r+1})]$$

where $r =$ the number of cointegrating vectors between the four variables; $p = 4$, the number of equations in the VAR; λ_i = the ith eigenvalue, and $N = 77$, the number of observations adjusted for observations lost due to the use of lagged terms (two observations). The eigenvalues are obtained from a VAR model with intercept constrained to the cointegrating vector and the set of equations for the one lag model takes the form:

$$\Delta X_t = \Gamma_1 \Delta X_{t-1} + \Pi^* X^*_{t-2} + \varepsilon_t$$

where X_t is the matrix of variables $[F(tT)\ p(t)\ rf(tT)\ S(t)]'$ and $X^*_t = [X_t\ 1]'$ includes a vector of ones to capture the intercept term. Γ_i is a parameter matrix and $\alpha\beta'^* = \Pi^*$. The term Π^* is a measure of the long-run relationship between the levels of the variables once the autoregressive term has removed the short-term effects. The $(4 \times r)$ matrix, α, provides estimates of the speed of adjustment toward the long-run relationship and the $(5 \times r)$ matrix, β^*, is the set of cointegrating vectors. The critical values at the 5% level are obtained from Osterwald-Lenum (1992).
[a]Statistically significant at the 5% level.

nonstationary, then a cointegrating relationship between the spot price and the futures price is not expected to exist if the cost-of-carry relationship holds. Rather, the cointegrating relationship is predicted to exist between the components which make up the cost-of-carry relationship, futures price, spot price, interest rate to maturity, and stock level effects.

The Johansen tests are conducted on quarterly data. The interval choice is an important one because too large an observation interval leads to the systematic sampling problems described in Weiss (1984). Given the futures prices are for three-month delivery, the choice of quarterly data seems reasonable. The Johansen tests are run using both the U.K. Treasury bill three-month mid-rate and the Euro Sterling three-month mid-rate as proxies for the risk-free rate. As there is little change in the results, only the U.K. Treasury bill three-month mid-rate results are reported in Tables IV and V. The tests are repeated with two VAR models. The first includes one lagged term and the second includes two lagged

TABLE V

Estimates of Cointegrating Vectors and Speed of Adjustment vectors
U.K. Treasury Note Three-Month Mid-Rate
(α and β Matrices where $\alpha\beta' = \Pi$)

Futures β_0	*Spot* β_1	*Interest rate* β_2	*Stock* β_3	*Constant* β_4
Normalized cointegrating vectors (normalized eigenvectors, β)				
1	−0.962	−1.596	−0.027	0.115
Normalized average speed of adjustment parameters (α)				
0.924	1.944	−0.033	−0.615	—

		Restricted cointegrating vectors				
Tests of restrictions on parameters		β_0	β_1	β_2	β_3	β_4
$\beta_0 = \beta_1 = \beta_2$	$\chi^2(2)$ 1.46	{1	−1	−1	−0.026	0.306}
Test statistic critical values $\chi^2(2, 10\%) = 4.605$, $\chi^2(2, 5\%) = 5.991$						
$\beta_0 = \beta_1$	$\chi^2(1) = 1.41$	{1	−1	−1.232	−0.027	0.328}
$\beta_0 = \beta_2$	$\chi^2(1) = 0.43$	{1	−0.968	−1	−0.024	0.096}
$\beta_3 = 0$	$\chi^2(1) = 9.94$[a]	{1	−0.997	1.665	0	−0.063}
Test statistic critical values $\chi^2(1, 10\%) = 2.706$, $\chi^2(1, 5\%) = 3.841$						
Tests of restrictions on parameters—restrictions on speed of adjustment or loading vector						
$\alpha_1 = \alpha_2 = \alpha_3 = 0$	$\chi^2(3) = 6.1$[b]					
Test statistic critical values $\chi^2(3, 10\%) = 6.251$, $\chi^2(3, 5\%) = 7.815$						
$\alpha_1 = \alpha_3 = 0$	$\chi^2(2) = 4.63$[b]					
Test statistic critical values $\chi^2(2, 10\%) = 4.605$, $\chi^2(2, 5\%) = 5.991$						
$\alpha_0 = 0$	$\chi^2(1) = 1.34$					
$\alpha_1 = 0$	$\chi^2(1) = 4.63$[a]					
$\alpha_2 = 0$	$\chi^2(1) = 2.00$					
$\alpha_3 = 0$	$\chi^2(1) = 0.11$					
Test statistic critical values $\chi^2(1, 10\%) = 2.706$, $\chi^2(1, 5\%) = 3.841$						

Notes: The var described in Table IV underlies the results reported in this table. The elements in cointegrating vectors (associated with the largest eigenvalue) are normalized by dividing all the parameters in the cointegrating vector by the parameter matching the futures price. The degrees of freedom for tests of cointegrating vector restrictions is $r(p + 1 - s)$, where r = number of cointegrating vectors (one), $p = 4$, number of equations in the VAR with the addition of one to allow for the constant term included in the cointegrating vector, and s = number of columns in the restriction matrix. Elements in the average speed of adjustment vector measure the sensitivity of each of the cost-of-carry variables to cost-of-carry pricing discrepancies. These values are normalized in line with the cointegrating vector normalization. The number of degrees of freedom for the test of the speed of adjustment, or loading, is $r(p - m)$, where m is the number of restrictions imposed on the speed of adjustment vector.

[a]Statistically significant at the 5% level.

[b]Statistically significant at the 10% level.

terms. At least one lag term is required given the serial correlation evident in the residuals and the use of more than two lags has a substantial impact on the degrees of freedom given a sample size of 79 quarterly observations. Results are reported in Tables IV and V for the one lag VAR though the following discussion refers to both models where required.

Tests for the number of cointegrating vectors are reported in Table IV. Only one cointegrating vector is identified at the 5% level of significance. Similar results are obtained with the use of Euro Sterling three-month mid-rates. Further, the lag choice has little impact on these statistical tests. This is consistent with the cost-of-carry model, which predicts only one cointegrating vector between the cost-of-carry variables. A test of the null hypothesis that the constant term is located within the cointegrating vector cannot be rejected.[11] This result is also consistent with cost-of-carry model prediction where it is assumed there are non-interest rate-based storage costs and a fixed component arising from stock level effects.

The cost-of-carry model is rearranged into vector form for cointegration testing, giving the predicted cointegrating vector $[+1\ -1\ -1\ -\alpha\ -\gamma^*]'$. In Table V the spot price parameter of -0.962 is close to -1.000 with failure to reject equality of the futures price parameter and the spot price parameter at the 5% level of significance. This result varies little with lag choice or interest rate proxy choice. The interest rate parameter of -1.596 is somewhat greater than expected though the hypothesis that the futures price parameter is equal to the interest rate parameters is not rejected at the 5% level of significance, suggesting the interest rate parameter is not measured with a great deal of precision. The magnitude of this variable is also rather sensitive to lag choice with the inclusion of two lags resulting in a parameter value of -0.700.

The stock level parameter of -0.027 is statistically significantly different from zero, exhibiting little variation with lag choice or interest rate proxy. This is an important result because it identifies stock level as a statistically significant component of the cost-of-carry model for the lead contract. The importance of the stock level effect is further highlighted by the reversal of the parameter sign for interest rate and constant term when the stock level parameter is constrained to zero. The intercept value of 0.11, though little affected by interest rate choice, does vary somewhat with lag choice, increasing to 0.40 with the two lag term model. Both the intercept and stock level parameter estimates are consistent with the Working (1949) and Brennan (1958) cost-of-carry model.

The speed of adjustment parameters are also reported in Table V. Chi-square tests, with the null of zero speed of adjustment, cannot be rejected at the 5% level of significance except for the spot price. The spot price speed of adjustment parameter of 1.944 suggests quite rapid adjustment in the spot price toward correction of cost-of-carry model di-

[11]The test statistic value is 0.51, well below the critical value, $\chi^2(3, 5\%) = 7.815$.

vergence. The remaining speed of adjustment parameters are insignificant though the parameter values suggest these variables tend to increase the level of cost-of-carry divergence rather than decrease it with values of 0.92 for the futures price, -0.03 for the interest rate, and -0.61 for the stock effect. Statistical significance of the speed of adjustment parameters is sensitive to the number of lags included in the model. The two lag model identifies statistically significant speed of adjustment parameters for both the spot price and the futures price. In this case, the futures price speed of adjustment parameter is 1.817 and the spot price parameter is 2.586. Again, the rather perverse result is noted where the futures price reaction leads to further cost-of-carry model divergence, while the spot price reaction tends to reduce the divergence with the magnitude of the spot price reaction almost twice that of the futures price reaction.

If the one lag analysis is viewed in isolation, a researcher would assume that a single equation error-correction model (ECM) is suitable for analysis, but the result for the two lag model makes the issue less clear. Uncertainty about the speed of adjustment test result suggests that the use of the Johansen technique is justified. Although the method is more complex, it allows for interactions between the error-correction vector and the underlying cost-of-carry variables without the need for assumptions about the exogeneity of chosen explanatory variables.

CONCLUSIONS

This article provides a direct test of the cost-of-carry model using time series of matched quarterly stock level, futures prices, and spot prices for the commodity lead, which is traded on the LME. This is a unique data set because spot price, futures price, and stock levels are obtained from the same market and reported futures prices have a fixed time to maturity of three months.

Unit root test results for spot and futures prices are somewhat tentative with evidence both for and against unit root processes in these variables, while interest rates and stock levels appear to follow unit root processes. Nevertheless, the assumption of unit root or near unit root processes does not appear to be too extreme and so the Johansen cointegration test is selected for analysis of the cost-of-carry model. There is evidence of one cointegrating vector which consists of futures price, spot price, interest rate to maturity, stock level, and a constant term. Consistent with the cost-of-carry model, statistical tests cannot reject equality of the parameters associated with the futures price, spot price, and the interest rate to maturity. Tests also indicate the existence of a stock level effect. Thus, in line with Working (1949) and Brennan (1958), the difference between futures price and spot price varies with stock levels and the level of interest rates.

BIBLIOGRAPHY

Bailey, W., and Ng, E. (1991): "Default Premiums in Commodity Markets: Theory and Evidence," *Journal of Finance,* 46:1071–1093.

Banerjee, A., Dolado, J. J., Galbraith, J. W., and Hendry, D. F. (1993): *Cointegration, Error-Correction, and the Econometric Analysis of Non-Stationary Data.* Oxford: Oxford University Press.

Bannerjee, A., and Hendry, D. F. (1992): "Testing Integration and Cointegration: An Overview," *Oxford Bulletin of Economics and Statistics,* 54:225–255.

Benninga, S., and Protopapadakis, A. (1994): "Forward and Futures Prices with Markovian Interest-Rate Processes," *Journal of Business,* 67:401–421.

Brennan, M. J. (1958): "The Supply of Storage," *The American Economic Review,* 48:50–72.

Brenner, R. J., and Kroner, F. K. (1995): "Arbitrage, Cointegration, and Testing the Unbiasedness Hypothesis in Financial Markets," *Journal of Financial and Quantitative Analysis,* 30:23–42.

Cheung, Y.-W., and Lai, K. S. (1993): "Finite-Sample Sizes of Johansen's Likelihood Ratio Tests for Cointegration," *Oxford Bulletin of Economics and Statistics,* 55:313–328.

Cho, D. W., and McDougall, G. S. (1990): "The Supply of Storage in Energy Futures Markets," *The Journal of Futures Markets,* 10:611–621.

Chowdhury, A. R. (1991): "Futures Market Efficiency: Evidence from Cointegration Tests," *The Journal of Futures Markets,* 11:577–589.

Cooper, R. (1993): "Risk Premia in the Futures and Forward Markets," *The Journal of Futures Markets,* 13:357–371.

Cornell, B., and Reinganum, M. R. (1981): "Forward and Futures Prices: Evidence from the Foreign Exchange Markets," *Journal of Finance,* 36:1035–1045.

Crowder, W. J., and Hamed, A. (1993): "A Cointegration Test for Oil Futures Market Efficiency," *The Journal of Futures Markets,* 13:933–941.

Dickey, D. A., and Rossana, R. J. (1994): "Cointegrated Time Series: A Guide to Estimation and Hypothesis Testing," *Oxford Bulletin of Economics and Statistics,* 56:325–353.

Edwards, F. R., and Canter, M. S. (1995): "The Collapse of Metallgesellschaft: Unhedgeable Risks, Poor Hedging Strategy, or Just Bad Luck?" *The Journal of Futures Markets,* 15:211–264.

Enders, W. (1995): *Applied Econometric Time Series.* New York: John Wiley & Sons.

Engle, R. F., and Granger, C. W. J. (1987): "Cointegration, and Error Correction: Representation, Estimation and Testing," *Econometrica,* 55:251–276.

Ericsson, N. R. (1992): "Cointegration, Exogeneity, and Policy Analysis: An Overview," *Journal of Policy Modelling,* 14:251–280.

Fama, E. F., and French, K. R. (1987): "Detecting Spot Price Forecasts in Futures Prices," *Journal of Business,* 60:55–73.
Fama, E. F., and French, K. R. (1988): "Business Cycles and the Behaviour of Metals Prices," *Journal of Finance,* 43:1075–1093.
Fort, R., and Quirk, J. (1988): "Normal Backwardation and the Stocks Effect," *Journal of Political Economy,* 96:81–99.
Franses, P. H., and Kofman, P. (1991): "An Empirical Test for Parities between Metal Prices at the LME," *The Journal of Futures Markets,* 11:729–736.
French, K. R. (1983): "A Comparison of Futures and Forward Prices," *Journal of Financial Economics,* 12:311–342.
French, K. R. (1986): "Detecting Spot Price Forecasts in Futures Prices," *Journal of Business,* 59:39–54.
Fuller, W. (1976): *Introduction to Statistical Time Series.* New York: John Wiley & Sons.
Goss, B. A. (1981): "The Forward Pricing Function of the London Metal Exchange," *Applied Economics,* 13:133–150.
Goss, B. A. (1983): "The Semi-Strong Form Efficiency of the London Metal Exchange," *Applied Economics,* 15:681–698.
Hall, A. D., Anderson, H. M., and Granger, C. W. J. (1992): "A Cointegration Analysis of Treasury Bill Yields," *The Review of Economics and Statistics,* 74:116–126.
Hirshleifer, D. (1988): "Risk, Futures Pricing, and the Organisation of Production in Commodity Markets," *Journal of Political Economy,* 96:1206–1220.
Hirshleifer, D. (1989): "Futures Trading, Storage, and the Division of Risk: A Multiperiod Analysis," *The Economic Journal,* 99:700–719.
Hsieh, D. A., and Kulatilaka, N. (1982): "Rational Expectations and Risk Premia in Forward Markets: Primary Metals at the London Metals Exchange," *Journal of Finance,* 37:1199–1207.
Johansen, S. (1988): "Statistical Analysis of Cointegration Vectors," *Journal of Economic Dynamics and Control,* 12:231–254.
Johansen, S., and Juselius, K. (1990): "Maximum Likelihood Estimation and Inference on Cointegration with Application to the Demand for Money," *Oxford Bulletin of Economics and Statistics,* 52:169–209.
Kofi, T. A. (1973): "A Framework for Comparing the Efficiency of Futures Markets," *American Journal of Agricultural Econometrics,* 55:584–594.
Krehbiel, T., and Adkins, L. C. (1993): "Cointegration Tests of the Unbiased Expectations Hypothesis in Metals Markets," *The Journal of Futures Markets,* 13:753–763.
Kremers, J. J. M., Ericsson, N. R., and Dolado, J. J. (1992): "The Power of Cointegration Tests," *Oxford Bulletin of Economics and Statistics,* 54:325–348.
Lien, D.-H. D. (1987): "The Stocks Effect in Commodity Futures Markets: An Empirical Study," *The Journal of Futures Markets,* 7:637–652.
Lloyd, T. (1994): "Testing a Present Value Model of Agricultural Land Values," *Oxford Bulletin of Economics and Statistics,* 56:209–223.
Mac Donald, R., and Murphy, P. D. (1989): "Testing for the Long Run Relationship between Nominal Interest Rates and Inflation Using Cointegration Techniques," *Applied Economics,* 21:439–447.

Mac Donald, R., and Taylor, M. P. (1988a): "Testing Rational Expectations and Efficiency in the London Metal Exchange," *Oxford Bulletin of Economics and Statistics,* 50:41–52.

Mac Donald, R., and Taylor, M. P. (1988b): "Metals Prices, Efficiency and Cointegration: Some Evidence from the London Metals Exchange," *Bulletin of Economic Research,* 40:235–239.

Mac Donald, R., and Taylor, M. P. (1989): "Rational Expectations, Risk and Efficiency in the London Metal Exchange: An Empirical Analysis," *Applied Economics,* 21:143–153.

Moore, M. J. (1994): "Testing for Unbiasedness in Forward Markets," *The Manchester School,* 62:67–78.

Moore, M. J., and Cullen, U. (1995): "Speculative Efficiency on the London Metal Exchange," *The Manchester School,* 63:236–236.

Newey, W., and West, K. (1987): "A Simple, Positive Semi-Definite, Heteroscedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica,* 55:703–708.

Osterwald-Lenum, M. (1992): "A Note with Quantiles of the Asymptotic Distribution of the Maximum Likelihood Cointegration Rank Test Statistics," *Oxford Bulletin of Economics and Statistics,* 54:461–472.

Pierce, R. G., and Snell, A. J. (1995): "Temporal Aggregation and the Power of Tests for a Unit Root," *Journal of Econometrics,* 65:333–345.

Pindyck, R. S. (1993): "The Present Value Model of Rational Commodity Pricing," *The Economic Journal,* 103:511–530.

Radetzki, M. (1990): *A Guide to Primary Commodities in the World Economy,* Cambridge, MA: Basil Blackwell.

Sephton, P. S., and Cochrane, D. K. (1990): "A Note on the Efficiency of the London Metal Exchange," *Economic Letters,* 33:341–345.

Sephton, P. S., and Cochrane, D. K. (1991): "The Efficiency of the London Metal Exchange: Another Look at the Evidence," *Applied Economics,* 23:669–674.

Shea, G. S. (1992): "Benchmarking the Expectations Hypothesis of the Interest Rate Term Structure: An Analysis of Cointegration Vectors," *Journal of Business and Economics Statistics,* 10:347–366.

Sims, C. A. (1980): "Macroeconomics and Reality," *Econometrica,* 48:1–48.

Telser, L. G. (1958): "Futures Trading and the Storage of Cotton and Wheat," *Journal of Political Economy,* 66:233–255.

Urbain, J.-P. (1992): "On Weak Exogeneity in Error Correction Models," *Oxford Bulletin of Economics and Statistics,* 54:187–207.

Weiss, A. A. (1984): "Systematic Sampling and Temporal Aggregation in Time Series Models," *Journal of Econometrics,* 26:271–281.

Weymar, F. H. (1966): "The Supply of Storage Revisited," *The American Economic Review,* 56:1226–1234.

Working, H. (1949): "The Theory of the Price of Storage," *The American Economic Review,* 34:1254–1262.

Wright, B. D., and Williams, J. C. (1989): "A Theory of Negative Prices for Storage," *The Journal of Futures Markets,* 9:1–13.