TIME SERIES ECONOMETRICS AND COMMODITY PRICE ANALYSIS*

by

Robert J. Myers

Department of Economics
The University of Queensland

February 1992

* Invited paper at the 36th Annual Conference of the Australian Agricultural Economics Society, Australian National University, Canberra, February 10-12, 1992.

1. Introduction

Econometric analysis of commodity prices has a long and distinguished history stretching back to the birth of econometrics itself as an emerging science in the 1920s and 1930s (Working, 1922; Schultz, 1938). Since then, a very large literature has developed focusing on estimating commodity supply and demand systems; forecasting commodity supplies and prices; and evaluating the effects of commodity pricing policies. Much of this research relies on a standard set of econometric methods, as outlined in books such as Theil (1971) and Johnston (1984).

The goal in this paper is not to provide a detailed survey of the literature on econometric analysis of commodity prices. This has been done elsewhere (e.g. Tomek and Robinson, 1977) and, in any case, is well beyond the scope of what can be achieved in the limited time and space available here. Rather, the aims are to comment on some recent developments taking place in the time-series econometrics literature and discuss their implications for modelling the behaviour of commodity prices. The thesis of the paper is that developments in the time-series literature have important implications for modelling commodity prices, and that these implications often have not been fully appreciated by those undertaking commodity price analysis.

The time-series developments that will be discussed include stochastic trends (unit roots) in economic time series; common stochastic trends driving multiple time series (cointegration); and time-varying volatility in the innovations of economic time series (conditional heteroscedasticity). None of these developments are all that new. The pioneering work on stochastic trends was undertaken by David Dickey and Wayne Fuller in the 1970s, while Engle and Granger's seminal paper on cointegration was widely circulated well prior to its publication date of 1987. Similarly, the original papers on conditional heteroscedasticity by Robert Engle and Tim Bollerslev were published in the early 1980s and have since spawned an impressive literature, some of which relates to commodity prices (Bollerslev, Chou and Kroner, 1990; Baillie and Myers, 1991; Myers, 1991). But while the developments themselves are well known, the resulting implications for the econometric analysis of commodity prices appear not to be fully appreciated in may cases. The present paper is designed to correct this situation and provide a modern perspective on econometric modelling of commodity prices using time-series data.

The paper is divided into three parts. First, in Section 2, the characteristic time-series properties of commodity price series are outlined. It turns out that many commodity price series exhibit similar time-series properties, even when they are quite separate and would appear not to be closely related to one another. The second part of the paper, in Section 3, discusses some of the implications which these characteristic properties of commodity price series have for the econometric analysis of commodity markets. It is shown that conventional statistical inference can be very misleading in econometric models containing commodity prices. The third and final part of the paper in Section 4 reports a simulation study which highlights some of the econometric problems arising from the time-series properties of commodity prices.

2. Characteristic Properties of Commodity Price Series

Prices of different commodities obviously are influenced by distinct forces and therefore behave somewhat differently. Nevertheless, there are several characteristic properties which most commodity price series seem to share in common. Perhaps surprisingly, both storable commodities and those considered difficult to store appear to share the same broad features. In this section of the paper, these characteristic properties of commodity price series are outlined and discussed. The properties are illustrated using cattle and wheat prices as an example, two quite distinct commodities which are of major importance to Australian agriculture.

2.1 High Volatility

Prices of primary commodities generally are much more volatile than prices of manufactured consumer goods. Figure 1 shows monthly cattle and wheat prices over a twenty-year period beginning in 1970.[1] As well as the considerable month to month volatility, major secular upswings (1973 and 1978-79) and downswings (1974 and 1976) have occurred in both price series. These swings suggest the possibility of regime changes from one set of economic fundamentals driving the market to another.

This volatility in commodity prices poses major problems to industry participants and policymakers, particularly in countries whose export earnings and GDP depend heavily on sales of primary commodities. An understanding of commodity price movements is therefore essential to the sound management of macroeconomic fluctuations in general, and to the microeconomic management of agricultural price policies and risk sharing mechanisms.

One especially interesting question stemming from recent research on commodity price volatility involves the extent to which commodity price volatility can be accounted for by changes in the economic fundamentals underlying the market. If most of the commodity price volatility can be accounted for by changes in economic fundamentals then markets are said to be efficient in guiding resource allocation decisions. On the other hand, if there is considerable 'excess volatility' beyond that which can be accounted for by changes in economic fundamentals, then commodity markets are being driven by fads or speculative bubbles and commodity prices become inefficient resource allocation signals.

Research on this issue is still at an early stage but work by Pindyck and Rotemberg (1990) suggests that there is 'excess comovement' in commodity prices, in the sense that different commodity prices move together much more closely than would be expected based on changes in economic fundamentals alone.

Even if the high volatility characterising commodity prices does accurately reflect changes in underlying economic fundamentals, there remains the question of how the resulting price risks can best be managed. It is now widely understood that the efficiency of market mechanisms in these circumstances relies on the completeness of the market structure (Newbery and Stiglitz, 1981; Myers, 1988). If markets are complete, and all relevant risks are therefore insurable on competitive markets, then the market mechanism remains economically efficient. However, if markets are incomplete and some risks are therefore uninsurable, then there may be a role for government in designing risk sharing mechanisms which improve the distribution of risk throughout the economy.

2.2 Stochastic Trends

Another characteristic property of high frequency commodity price series (series sampled at daily, weekly or monthly intervals) is that they appear to contain trends which change randomly over time. Looking again at Figure 1, for example, it is obvious that forecasts from a deterministic linear time trend estimated using data from the beginning of 1970 to the end of 1973 would have failed miserably in predicting the plunge in commodity prices which occurred in 1974 and the subsequent reversal to higher price levels. If commodity prices were the result of random deviations about an unchanging linear trend then this kind of price behaviour would never be observed.

An alternative to a deterministic linear time trend, which increases by a fixed amount every period, is to assume a stochastic trend, which increases by some fixed amount on average but in any given period the change in trend deviates from the average by some unpredictable random amount (Stock and Watson, 1988). Formally, this notion of a stochastic trend can be modelled as a random walk with drift:

$$(1) \qquad w_t = \mu + w_{t-1} + \varepsilon_t$$

where the drift parameter μ is the average predictable increase in w_t each period, and ε_t is a serially uncorrelated random shock to the trend.

When a commodity price p_t contains a stochastic trend then the price can be written as the sum of a random walk component and a stationary component:

$$(2) \qquad p_t = w_t + z_t$$

where z_t has finite variances and autocovariances and a distribution which does not depend on time. In this case, w_t represents the stochastic trend and z_t represents deviations or cyclical swings away from trend.

It turns out that (2) is a simple yet powerful model for explaining high frequency commodity price data. Indeed, Beveridge and Nelson (1981) have shown that any variable which can be modelled as an autoregressive integrated moving average (ARIMA) process with order of integration one (i.e. requires first differencing to induce stationarity), has a representation as the sum of a random walk and a stationary component, as in (2). Thus, all of the research which has found once integrated ARIMA models do a good job of modelling high frequency commodity price series (e.g. Baillie and Myers 1991) is consistent with the idea that commodity prices are made up of a stochastic trend and stationary deviations around trend.

Over the past decade, important advances have been made in the development of formal statistical tests for the presence of stochastic trends in time-series data. The original tests developed by Dickey and Fuller (1979, 1981) have been improved by Phillips (1987) and Perron (1988) so that they are now more robust to autocorrelation and heteroscedasticity in the errors, two problems which plague high frequency economic time-series data. Both the Dickey-Fuller and Phillips-Perron approaches test the null hypothesis of a stochastic trend against the alternative that the series is stationary, or stationary around a deterministic linear time trend. The difference between them is that the Phillips-Perron framework allows for a very general error structure that may be autocorrelated and heterogeneously distributed. Full details of the tests and test statistics can be found in Perron (1988).

Results from applying the Dickey-Fuller and Phillips-Perron unit root tests to monthly cattle and wheat prices are shown in Tables 1 and 2 respectively. Both tests are based on a regression of the commodity price on a constant, time trend, and the lagged commodity price. Under the null hypothesis of a stochastic trend then the coefficient on lagged price is one but the usual t-statistic has a nonstandard distribution and so special tables tabulated in Fuller (1976) must be used. The lag length ℓ of the test refers to the number of lagged first differences of price included in the Dickey-Fuller test or, in the case of the Phillips-Perron test, ℓ is the lag truncation parameter in the Newey and West (1987) formula for generating consistent covariance matrix estimates. The null hypothesis of a stochastic trend cannot be rejected for either cattle or wheat, even at the 10% level of significance.

One of the problems with the Dickey–Fuller and Phillips–Perron unit root tests is that the stochastic trend is the null hypothesis. This ensures that a stochastic trend is accepted unless there is strong evidence against it. It could be, however, that a stochastic trend cannot be rejected simply because the data are not very informative about whether or not there is a unit root (i.e. standard unit root tests have low power against the alternative that the series is stationary but with a root that is close to unity). In response, a new test has been developed by Kwiatkowski, Phillips, Schmidt and Shin (1992) which tests the null hypothesis of stationarity against the alternative that the series has a stochastic trend.

Results from applying the KPSS test to cattle and wheat prices are shown at the bottom of Tables 1 and 2 respectively. The estimated model encompasses a random walk and the test is a Lagrange multiplier test of the null hypothesis that the innovations in the random walk have zero variance. Full details of the test are available in Kwiatkowski, Phillips, Schmidt and Shin (1992). The null hypothesis of stationarity is soundly rejected in the cattle and wheat price data irrespective of the lag truncation parameter ℓ used to generate consistent covariance matrix estimates for the test.

Because a stochastic trend in cattle and wheat prices cannot be rejected, but stationarity can, then the preponderance of evidence points toward a stochastic trend. Similar results have been found for other commodity price data sampled at high frequencies. It should be remembered, however, that a stochastic trend is less evident in commodity price data sampled at annual intervals. Indeed, Deaton and Laroque (1990) argue convincingly that we should expect commodity prices sampled at annual frequencies to be stationary for theoretical reasons, although they admit that this will be difficult to discover empirically because of the small number of annual observations typically available.

Furthermore, stochastic trends and stationarity are not the only time–series models that commodity prices might follow. For example, Hamilton (1989) has developed a model of a time series as a sequence of stochastic segmented time trends which can generate long swings in the data and is an alternative to the random walk notion of stochastic trends presented above. Because this model is not encompassed in the standard unit root hypothesis testing framework then the standard tests really have nothing to say regarding its applicability.

The question of whether or not commodity prices have a stochastic trend has, as we shall see below, crucial implications for how we go about econometric analysis of commodity market data. The existence of stochastic trends distorts many of the usual distributional assumptions and implications typical in commodity price analysis. Thus, great care must be taken in testing for stochastic trends and in the econometric analysis of data that contain stochastic trends.

2.3 Comovements in Commodity Price Series

A careful comparison of the commodity prices in Figure 1 reveals that, while there are major differences between the cattle and wheat price series, there is also a persistent tendency for them to move together over time. For example, the two commodity prices followed each other very closely over the 1970–73 period, and again throughout most of the 1980s. And while cattle prices were much more volatile than wheat prices through the ups and downs of the 1974–78 period, there remains a tendency for the two prices to move in the same direction. This leaning towards comovements in prices is a property shared by many different apparently unrelated commodities (Pindyck and Rotemberg, 1990).

There are three possible reasons for persistent comovements in commodity prices. First, it could be that supply and demand shocks to any one commodity market spill over into other related commodities causing a group of commodity prices to move together. While this is a logical explanation for commodities which are strongly related to one another, either in production or consumption (e.g. wheat and rice), it cannot explain comovements between largely unrelated commodities (e.g. cattle and copper). Second, common macroeconomic shocks to, say, the money supply or interest rates could be affecting all commodity prices in a similar way. Common macroeconomic shocks undoubtedly explain some of the comovement among commodity prices. However, research by Pindyck and Rotemberg (1990) suggests that macroeconomic shocks can only explain a small fraction of the actual comovement in commodity prices. A third possibility is that market speculation and overreaction causes spillovers between commodity markets that cannot be accounted for by changing microeconomic fundamentals or common macroeconomic shocks. In this interpretation there is 'excess comovement' among commodity prices leading to commodity price volatility that is greater than it 'ought to be'. As discussed earlier, research on this issue is still at an early stage.

One way to formalise the idea of comovements among commodity prices is to invoke the theory of cointegrated stochastic processes. We have already seen that most commodity price series can be represented as the sum of a stochastic trend and stationary deviations about trend, as in (1) and (2). Under these circumstances two commodity prices are said to be cointegrated if they share the same stochastic trend:

(3a) $p_{1t} = w_t + z_{1t}$

(3b) $p_{2t} = \delta w_t + z_{2t}$

where p_{it} is commodity price i; z_{it} is stationary component i; and δ is a parameter which identifies a long-run relationship between the two prices. Substituting (3a) into (3b) gives

(4) $p_{2t} = \delta p_{1t} + z_t$

where $z_t = z_{2t} - \delta z_{1t}$. In (4), the first component $p_{2t} = \delta p_{1t}$ represents a long-run equilibrium relationship between the two prices, resulting from their common stochastic trend, while z_t represents temporary stationary deviations from this long-run equilibrium relationship.

The parameter δ which characterises the nature of the long-run equilibrium relationship between the series can be estimated by applying ordinary least squares (OLS) to (4). Perhaps surprisingly, OLS estimates the parameters of this 'cointegrating regression' consistently under very general assumptions about the statistical properties of z_t, even when z_t is autocorrelated and correlated with p_{1t}. In fact, the OLS estimate converges to the true parameter value at a faster rate than in the usual case of stationary regressors for reasons that will be explained later. Nevertheless, OLS estimates of a cointegrating regression generally follow a nonstandard distribution theory and so one must be very careful in undertaking hypothesis tests using results from cointegrating regressions. More on this later.

It should be emphasised that cointegration is not the only reason for comovement in commodity prices. For example, the stochastic trends driving two prices could be distinct but highly correlated. In the very long run these prices would eventually diverge and become unrelated but they could show considerable comovement in small finite data series. Or alternatively, the stochastic trends driving two prices could be completely unrelated but there might be strong correlations between their stationary components. Such series might exhibit considerable comovement in the short-run even though they are unrelated in the long run.

OLS based tests for cointegration have been developed by Engle and Granger (1987). The first step is to estimate the cointegrating regression (4) via OLS. Then the residuals from this regression are tested for the existence of a stochastic trend. If the hypothesis of a stochastic trend in the residuals is rejected then the series are cointegrated because they each have stochastic trends but a linear combination of the series is stationary. This indicates the series have common stochastic trends, with the cointegrating regression providing an estimate of the long-run equilibrium relationship between the series. On the other hand, if the hypothesis of a stochastic trend in the residuals cannot be rejected then the series are not cointegrated and are therefore driven by distinct, though possibly highly correlated, stochastic trends. Testing for cointegrating vectors in multivariate systems is more difficult but appropriate techniques have been developed by Johansen (1988).

Unfortunately, testing for a stochastic trend in the residuals from a cointegrating regression is not just a straightforward application of the Dickey-Fuller and Phillips-Perron unit root tests. The problem is that the cointegrating parameter δ must be estimated, but will only be identified by the OLS estimator when the null hypothesis of no cointegration is false. Because OLS seeks the δ which minimises the residual variance, and therefore is most likely to generate stationary residuals, then the standard Dickey-Fuller and Phillips-Perron tests are biased towards rejecting the null too often (Engle and Granger, 1987). This problem can be corrected by using the same test statistic but an alternative distribution to conduct the test. An appropriate asymptotic distribution has been tabulated by Phillips and Ouliaris (1990).

Results from applying the Dickey-Fuller and Phillips-Perron unit root tests to residuals from an OLS regression of wheat prices on cattle prices are shown in Table 3. The critical values are from the distribution tabulated by Phillips and Ouliaris (1990). The results do not support cointegration between cattle and wheat prices, even though both series have stochastic trends and appear to move together regularly. These results are consistent with some unpublished research from the World Bank which indicates that cointegration between commodity prices is the exception rather than the rule (for an exception see Vogelvang, 1990). Thus, comovement among commodity prices appears to be due mainly to short-run correlations in the data, rather than the existence of a common stochastic trend which keeps the series in a long-run equilibrium relationship.

Despite the finding that commodity prices often are not cointegrated, stochastic trends and cointegration play a crucial role in interpreting econometric studies of supply and demand for commodities. We have already seen that commodity prices generally have stochastic trends. Thus, if structural supply and demand equations include price variables and have stationary errors then the commodity price must be cointegrated with at least one other variable in the system. This issue is addressed in detail further below.

2.4 Time–Varying Volatility

Some of the earliest research on the distribution of commodity prices assumed that price changes are independent draws from an identical normal distribution. It soon became apparent, however, that this simple model missed some important characteristics of commodity price data. In particular, it was found that the volatility of price changes varies over time as the series moves between highly volatile periods, where large price changes tend to be followed by other large changes (of either sign) and less volatile periods, where small price changes tend to be followed by other small changes (again of either sign). This temporal instability in the variance of commodity price innovations has become a well–known feature in empirical studies.

It should be clear that time–varying volatility can be quite consistent with the existence of a stochastic trend in commodity prices. In fact, most of the research on time–varying volatility has focused on price changes, which implies that the stochastic trend in the series has been removed by first differencing. Once the stochastic trend has been removed, many commodity price series exhibit little residual autocorrelation. This does not necessarily mean, however, that the variance of the price innovations (e.g. price changes) is constant. If the series experiences time–varying volatility then the variance of price innovations will change over time in response to different shocks to the commodity market.

Time–varying volatility in commodity price series leads to autocorrelation patterns in the conditional variance of price innovations, where the variance is conditional on an information set available at the time forecasts are being formed. Engle (1982) has termed this 'conditional heteroscedasticity' and developed the autoregressive conditional heteroscedasticity (ARCH) model to capture such effects. A sequence of innovations e_t follow an ARCH(q) model if they can be represented:

(5a) $e_t|\Omega_{t-1} \sim D(0, h_t)$

(5b) $h_t = \omega + \sum_{i=1}^{q} \alpha_i e_{t-i}^2$

where D is some arbitrary distribution with mean zero and variance h_t, conditional on a set of information Ω_{t-1} available at t−1. The conditional variance h_t is a weighted sum of past squared innovations, leading to autocorrelation in the squared residuals and time-varying conditional variances.

Bollersiev (1986) has generalised the ARCH model to the so-called generalised autoregressive conditional heteroscedasticity (GARCH) model by including lagged conditional variances as well as lagged squared innovations in the equation explaining conditional variance movements. A simple yet useful example of a GARCH model is the GARCH (1,1) model of price changes:

(6a) $p_t - p_{t-1} = \mu + e_t$

(6b) $e_t|\Omega_{t-1} \sim D(0, h_t)$

(6c) $h_t = \omega + \alpha e_{t-1}^2 + \beta h_{t-1}$

In this example, price changes are equal to a constant μ plus a serially uncorrelated error e_t. The error e_t is sampled from some arbitrary distribution D with mean zero and variance h_t, conditional on a set of information Ω_{t-1}; and the conditional variance h_t

evolves based on last period's conditional variance and the realised value of last period's squared innovation e_{t-1}^2. Note that (6c) allows for a wide range of temporal patterns in the conditional variance of price innovations. This model has been applied successfully to a number of primary commodity prices by Baillie and Myers (1991).

A simple test for conditional heteroscedasticity has been developed by Engle (1982) and is applied to the cattle and wheat price series analysed here. The test involves first estimating an ARIMA model to account for any autocorrelation in the series. Then the residuals from the ARIMA model are squared and regressed on q lagged squared residuals. Under the null hypothesis of no ARCH effects against an alternative of an ARCH(q) model then the statistic TR^2, where T is the sample size and R^2 is the coefficient of determination from the lagged squared residuals regression, is distributed chi-square with q degrees of freedom. Results from applying this test to the cattle and wheat data used here are shown in Table 4. The null hypothesis of no ARCH effects is soundly rejected in the case of cattle prices. The evidence is less decisive in the case of wheat prices but there does appear to be significant second-order autocorrelation in the squared residual series. Thus, both commodity prices seem to exhibit time-varying volatility and, as we shall see, this has implications for econometric analysis of commodity markets.

2.5 Excess Kurtosis

Another weakness of early research on the distribution of commodity prices was the assumption that price changes are normally distributed. Subsequent work has shown that the tails of empirical price distributions appear to be much fatter than the normal, indicating excess kurtosis in commodity price changes (Gordon, 1985). This excess kurtosis cannot be handled within the confines of the usual normal distributional assumption.

Perhaps surprisingly, the ARCH model leads to a partial solution to the excess kurtosis problem. Even if the conditional distribution of price changes is assumed normal in the ARCH model, then the unconditional distribution is not normal and, in fact, has fatter tails than the normal (Engle, 1982). Thus, ARCH and GARCH models go part of the way towards accounting for apparent non-normalities in the empirical distribution of commodity prices. Nevertheless, empirical research has shown that even the ARCH and GARCH models fail to capture all of the excess kurtosis in commodity prices, if the assumption of a normal conditional distribution for price innovations is maintained.

One solution which has been quite effective is to assume the conditional distribution of price innovations in the GARCH model follows a t-distribution with degrees of freedom treated as a parameter to be estimated. This provides the model with the necessary flexibility to fully capture the excess kurtosis inherent in commodity price data. The GARCH model with conditional t-distribution has done a good job of representing commodity price data for a range of different commodities (Baillie and Myers, 1991).

To investigate excess kurtosis in commodity prices the residuals from the ARIMA models of cattle and wheat prices used to conduct the ARCH test were standardised by dividing by their respective standard errors. If the distribution of the errors was normal then the distribution of the standardised errors would be standard normal with kurtosis equal to three. The sample kurtosis of the actual standardised errors was found to be 4.66 for cattle and 9.12 for wheat, thus casting considerable doubt on the normality of the residuals, especially for wheat.

3. Implications for Econometric Analysis of Commodity Markets

Statistical inference in econometric studies of commodity markets is usually conducted under the assumption that all variables are stationary (no stochastic trends), and often the additional assumption of normally distributed errors is also invoked. We have just seen, however, that these assumptions are inappropriate for commodity prices. Commodity prices typically feature stochastic trends and may share trends with other commodity market variables (i.e. be cointegrated). Commodity prices also tend to experience time-varying volatility and follow distributions that have excess kurtosis compared to the normal.

In this section, the implications which some of these empirical attributes of commodity prices have for conventional econometric analysis of commodity markets are outlined and discussed. It turns out that in some cases these attributes create significant difficulties in estimation and inference.

3.1 High Volatility

High volatility of commodity prices is perhaps their least problematic attribute from the perspective of undertaking appropriate statistical inference. Certainly, highly volatile prices may be difficult to explain using standard econometric models and techniques, so that the R^2 in equations trying to explain price movements may be quite small. This in itself, however, poses no particular statistical problems. The challenge presented by highly volatile commodity prices lies in explaining why the volatility occurs and deciding what, if anything, needs to be done to alleviate any undesirable consequences. Since this is not the focus of the current paper a discussion of these issues is left for another time.

3.2 Stochastic Trends and Cointegration

Stochastic trends and cointegration can be much more important for statistical inference. The issues are addressed within the context of the most simple textbook model of commodity supply and demand. The model is:

(7a) $y_t^s = \gamma_{11} p_t + \gamma_{12} k_t + u_{1t}$ (Supply)

(7b) $y_t^d = \gamma_{21} p_t + \gamma_{22} x_t + u_{2t}$ (Demand)

(7c) $y_t^s = y_t^d$ (Market equilibrium)

where y_t^s is quantity supplied; y_t^d is quantity demanded; p_t is the commodity price; k_t is a supply shifter (e.g. technical change); x_t is a demand shifter (e.g. income); and u_{1t} and u_{2t} are random supply and demand disturbances. All variables are in logarithms, so that the γ_{ij} coefficients represent supply and demand elasticities. Variables are also expressed as deviations from their respective means to eliminate constant terms. The variables in this model could easily be interpreted as vectors, and lagged prices and quantities included in the model, without changing the substance of the discussion which follows.

The supply and demand shift variables are assumed to follow strictly exogenous autoregressive processes:

(8a) $k_t = \rho_1 k_{t-1} + \varepsilon_{1t}$

(8b) $x_t = \rho_2 x_{t-1} + \varepsilon_{2t}$

where the random disturbances ε_{1t} and ε_{2t} are uncorrelated with the u_{it}. It is important to note that $|\rho_i| = 1$ as well as $|\rho_i| < 1$ is allowed so that the exogenous variables may have a stochastic trend or be stationary, depending on the value of ρ_i.

The model is exactly identified and has reduced form:

(9a) $p_t = \pi_{11}x_t + \pi_{12}k_t + v_{1t}$

(9b) $y_t = \pi_{21}x_t + \pi_{22}k_t + v_{2t}$

where $\pi_{11} = \gamma_{22}/(\gamma_{11} - \gamma_{21})$; $\pi_{12} = -\gamma_{12}/(\gamma_{11} - \gamma_{21})$; $\pi_{21} = \gamma_{11}\gamma_{22}/(\gamma_{11} - \gamma_{21})$; $\pi_{22} = -\gamma_{12}\gamma_{21}/(\gamma_{11} - \gamma_{21})$; and v_{1t} and v_{2t} are correlated functions of the structural errors u_{1t} and u_{2t}. The reduced form plays an important role in conventional estimation and inference in the model.

The conventional approach to estimation when $|\rho_i| < 1$ and all variables are stationary is to apply OLS to the reduced form and an instrumental variables (IV) estimator, such as two-stage least squares (2SLS), to the structural equations. The structural equations can also be estimated via systems methods such as three stage least squares or full information maximum likelihood. However, these systems methods have the disadvantage that any misspecification in one equation can spill over and cause problems in estimation of every equation. If all of the variables are stationary then IV estimation and inference usually takes place on the basis of asymptotic results because the small sample properties of the IV estimator are generally unknown.

Now suppose that the supply and demand shift variables in the commodity market model (7) follow independent random walks, $\rho_1 = \rho_2 = 1$, and are therefore not cointegrated. Then by the reduced form (9) price and quantity both have stochastic trends as well because they are a linear combination of k_t and x_t. In this case, the supply and demand equations (7a) and (7b) each represent linear combinations of stochastically trending variables. There are two possibilities. The first is that the linear combinations represented by the supply and demand equations (i.e. the structural errors u_{it}) themselves have a stochastic trend. In this case there is no long-run relationship between the variables and the result is a 'spurious regression' in the sense of Granger and Newbold (1974). Results from IV estimation of such equations are notoriously unreliable because estimated coefficients are not consistent and the R^2 converges to a random number. Clearly, estimation and inference in this case is frought with difficulties and the application of standard techniques will lead to major errors.

The second possibility is that the structural disturbances u_{it} are stationary and the supply and demand equations therefore represent stationary linear combinations of stochastically trending variables. In other words, (y_t, p_t, k_t) and (y_t, p_t, x_t) are two cointegrated vectors with the relevant supply and demand elasticities representing the long-run relationship between the series.

What are the implications of this scenario for estimation and statistical inference? Fortunately, the conventional IV estimator applied to the supply and demand equations remains consistent (Phillips and Hansen, 1988). In fact, the IV estimator applied to a cointegrating regression converges more rapidly to the true parameter values than it would if all of the variables were stationary. On the other hand, the asymptotic distribution of the IV estimator is generally not normal, as it would be with stationary variables. Thus, normal distribution theory cannot be used in hypothesis testing, even when relying on large sample results. This is an obvious problem for conventional inference in these types of models.

Furthermore, it turns out that the IV estimator of the cointegrated supply and demand equations is not the only consistent estimator. In particular, simple OLS is also consistent and converges rapidly to the true parameter values, despite the obvious simultaneity problem. The reason for this 'super consistency' is that all linear combinations of, for example, the supply equation variables (y_t, p_t, k_t) <u>other than that given by the supply equation itself</u> have asymptotically infinite variance. Because OLS minimises the residual variance it moves very quickly to the finite residual variance defined by the supply equation parameters.

Like the IV estimator, the asymptotic distribution of the OLS estimator of cointegrated relationships is not normal either. Thus conventional inference cannot proceed as usual even when relying on large sample results. However, Hansen and Phillips (1988) have developed methods for 'modifying' OLS (and IV) estimators and standard errors so as to reduce small sample bias and allow conventional asymptotic inference.

The choice between IV and OLS estimators of cointegrated supply and demand equations comes down to two issues. First, which is easier to apply? OLS is obviously easier than IV, although with modern computer technology the difference must be described as marginal. Second, which estimator performs better in small samples? Monte Carlo simulations have shown that this depends on the signal to noise ratio. If the variance of the innovations in the random walks which drive the long-run behaviour of the variables is high relative to the variance of the short-term dynamics (the structural disturbances in the commodity market model) then the signal to noise ratio is high, small sample bias is negligible, and OLS works well. Modified OLS estimates will then allow conventional inference. When the variance of the random walks is relatively low, however, then the signal to noise ratio is low and the IV estimator will generally work better, although again modified estimators are required before conventional inference can be applied. Some Monte Carlo evidence supporting these conclusions for our commodity market model is presented later.

To see the potential power and utility of accounting for stochastic trends and cointegration in commodity market studies, consider the case when the supply shift variable k_t is stationary but the demand shift variable x_t has a stochastic trend. By the reduced form, both price and quantity depend on x_t so these variables both have stochastic trends as well. Nevertheless the supply equation (7a) defines a linear combination of price and quantity $y_t - \gamma_{11}p_t$ which is stationary, and the cointegrating vector $(1, -\gamma_{11})$ defines the price elasticity of supply.

Now suppose one wanted to estimate the supply elasticity. Instead of worrying about conventional identification and simultaneous equations bias, all that is required is to run an OLS regression of quantity on price. Because all other linear combinations of quantity and price (besides the supply equation) have infinite variance, this simple bivariate linear regression gives a consistent estimate of the supply elasticity. The estimate also converges very quickly to the true parameter value, though one would not trust standard errors computed with the usual formula when undertaking statistical inference. All of this occurs despite the fact that price and quantity are determined simultaneously, and that the regression is 'misspecified' by exclusion of the supply shift variables in the regression.

Finally, consider the possibility that commodity price has a stochastic trend but that the quantity variable is stationary. This implies that the supply and demand shift variables both have stochastic trends but that the linear combination $\pi_{21}x_t + \pi_{22}k_t$ is stationary.

In this case, the reduced form price equation (9a) is a cointegrating regression, and therefore inference must proceed with caution But the reduced form quantity equation (9b) has a stationary dependent variable and a linear combination of the stochastically trending regressors is stationary. In these circumstances we can proceed as if all variables are stationary and conventional inference is applicable (Stock and Watson, 1988; Sims, Stock and Watson, 1990).

In the structural form, the dependent variable quantity in the supply and demand equations is stationary but the explanatory variables have a stochastic trend. Thus $\gamma_{11}p_t + \gamma_{12}k_t$ and $\gamma_{21}p_t + \gamma_{22}x_t$ represent cointegrating relationships and all of the coefficients of interest can be written as coefficients on stationary variables. This implies that conventional inference is applicable and that the IV estimator may perform better than OLS because price and quantity are simultaneously determined.

Most commodity prices appear to be characterised by a stochastic trend and this has crucial implications for estimation and inference. If quantity is stationary then price must be cointegrated with supply and demand shifters in the commodity market model. Thus, supply and demand equations have stationary dependent variables and a linear combination of the stochastically trending explanatory variables is stationary. In this case, IV estimation of the structural form leads to conventional asymptotic inference. On the other hand, if price and quantity both have stochastic trends then the supply and demand equations represent cointegrating regressions and the conventional asymptotic distribution theory for the IV estimator breaks down. The IV estimator remains consistent but so is OLS. Great care must be taken in using estimated standard errors from either estimator in these circumstances. The performance of the estimators in small samples depends on the variance of the random walks driving the system compared to the variance of the stationary disturbances.

3.3 Time-Varying Volatility and Excess Kurtosis

Time-varying volatility in commodity prices has the same general effect on statistical inference as any other form of heteroscedasticity. In particular, the standard OLS and IV estimators remain unbiased and consistent. However, there is a loss of efficiency and estimated standard errors may be biased. Indeed, Engle (1982) points out that the main problem with applying OLS to a model with ARCH disturbances is the resulting loss in efficiency.

It is important to note, however, that time-varying volatility has little effect on OLS and IV estimation of cointegrating regressions. The reason is that the 'super consistency' of such estimators causes rapid convergence to the true parameter values irrespective of the presence of heteroscedasticity. On the other hand, it is important that any 'modified' estimates used to undertake statistical inference in cointegrating regressions be robust to the presence of heteroscedasticity, because heteroscedasticity will affect the standard errors.

Excess kurtosis causes problems whenever inference requires a particular distributional assumption on the disturbance terms. Although the normal is typically chosen, the actual distribution of commodity prices appears to have fatter tails than the normal. This can be a particular problem in maximum likelihood estimation of commodity market models.

Time-varying volatility and excess kurtosis can be modelled parametrically using GARCH models and a conditional t-distribution for price innovations. If properly specified, such models should lead to increased efficiency and more accurate standard errors in commodity market estimation. Some nonparametric methods which are more robust to specification error have also been developed (e.g. Pagan and Ullah, 1988).

4. A Monte Carlo Simulation

A Monte Carlo experiment was undertaken to highlight some of the issues raised in the discussion of stochastic trends and cointegration above. The data generating process used in the simulation is a simple version of the commodity market model:

(10a) $y_t = 0.5p_t + k_t + u_{1t}$ (Supply)

(10b) $y_t = -0.5p_t + x_t + u_{2t}$ (Demand)

(10c) $k_t = \varepsilon_{1t}$

(9d) $x_t = x_{t-1} + \lambda\varepsilon_{2t}$

where all of the disturbance terms u_{it} and ε_{it} are identically and independently distributed N(0,1) variables. Notice that the parameter λ determines the signal to noise ratio. If λ is large then the variance of the random walk is large relative to the variance of stationary disturbances (high signal to noise ratio) while if λ is small the reverse occurs.

Each time the model was simulated, a series of 50 observations on 4 independent N(0,1) variables was drawn using the random number generator RNDN in GAUSS. These random numbers were then applied to Equation (10) to generate a series of 50 observations on (y_t, p_t, k_t, x_t), using $x_0 = 0$ as a start up value for x_t. Alternative estimators were then applied to the data set and their performance compared. The whole process was then repeated a total of 1000 times and summary results tabulated. The aim is to get an idea of how alternative estimators perform in small samples under stochastic trends and cointegration.

It is assumed that the aim of the exercise is to estimate the price elasticity of supply. Three alternative estimators are compared. First there is the conventional IV (2SLS) estimator of the supply equation, using k_t and x_t as instruments for p_t. Second is a simple bivariate OLS regression of y_t on p_t. Because the supply shift variable k_t is stationary, this should generate a consistent estimate of the supply elasticity. Third is a multivariate OLS regression of y_t on $(p_t\ k_t)$. Inclusion of the stationary regressor k_t should not affect the consistency of the supply elasticity estimate from the cointegrating regression. In small samples, however, inclusion of the k_t variable in the regression may lead to a different performance in inference. The experiment was repeated for two values of λ, $\lambda = 1$ and $\lambda = 10$.

Results from the experiment are reported in Table 5. In the case of low signal to noise ratio $(\lambda = 1)$ conventional 2SLS clearly performs best. The mean supply elasticity estimate is very close to the true value of 0.5 and it has a significantly lower root mean square error (RMSE) than the OLS estimators. Nevertheless, the standard errors from the 2SLS estimators cannot be trusted because the mean standard error estimate differs considerably from the actual RMSE of the estimator and the null hypothesis that the supply elasticity equals 5.0 is rejected 19% of the time when conventional distribution theory says it should be only 5%. Bivariate OLS performs the worst under low signal to noise with considerable bias in both the estimated supply elasticity and its standard error.

In the case of high signal to noise ratio, however, a simple bivariate regression of quantity on price gives a remarkably good estimate of the supply elasticity. The mean estimate is very close to the true value of 0.5 and the mean estimated standard error from a standard regression package is very close to the actual RMSE of 0.011. Using conventional inference the null hypothesis $\gamma_{11} = 0.5$ is rejected 9% of the time when the figure should be 5%. Thus, gre are still must be taken in undertaking inference in this model. The 2SLS and multivariate OLS estimators also perform well although the distribution theory for the 2SLS estimator is more distorted than the OLS estimators with the null hypothesis being rejected 18% of the time when it should be 5%.

These simulation results highlight two important facts about estimating supply and demand equations when prices have stochastic trends and are cointegrated with other variables. First, standard errors from the usual IV techniques, such as 2SLS, are not to be trusted and significant errors can be made in hypothesis testing if one naively uses standard tests and standard distribution theory. Second, in the case of high signal to noise ratios then simple OLS estimation of equations in a simultaneous equation system can provide remarkably good results. Here too, however, standard errors are not to be trusted and great care must be taken in hypothesis testing and statistical inference.

5. Concluding Comments

Developments in time-series econometrics have proceeded at a remarkable pace over the past decade and the full implications for econometric analysis of commodity markets are only just beginning to be widely understood. Many commodity prices appear to contain stochastic trends and be cointegrated with other commodity market variables. Time-varying volatility and excess kurtosis also characterise many commodity price series. These time-series characteristics have several significant implications for econometric analysis of commodity markets.

Most importantly, the presence of stochastic trends raises a number of econometric pitfalls for the unwary, particularly when it comes to inference regarding supply and demand elasticity estimates. New approaches are needed to account for the special distributional characteristics of stochastically trending variables. Time-varying volatility and excess kurtosis in commodity prices also need to be properly accounted for if estimation is to proceed with maximum efficiency. Indeed, to avoid major errors in estimation, inference and interpretation, commodity market analysts must become increasingly aware of the time-series characteristics of their data, and of the resulting implications for the use of various econometric methods and techniques.

Notes

1. Cattle prices in cents per kilogram are from January 1970 through June 1990 and are average monthly prices at Brisbane markets from various issues of the Australian Meat and Livestock Corporation's *Annual Statistical Review*. Wheat prices in U.S. dollars per ton are from January 1970 through February 1987 and are average Australian export prices from various issues of the International Wheat Council's *World Wheat Statistics*.

References

Baillie, R.T. and R.J. Myers (1991), 'Bivariate GARCH Estimation of the Optimal Commodity Futures Hedge', *Journal of Applied Econometrics*, 6, p.109–124.

Beveridge, S. and C.R. Nelson (1981), 'A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the "Business Cycle"', *Journal of Monetary Economics*, 7, p.151–174.

Bollerslev, T. (1986), 'Generalized Autoregressive Conditional Heteroscedasticity', *Journal of Econometrics*, 31, p.307–327.

Bollerslev, T., R.Y. Chou and K.F. Kroner (1990), *ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence*, Working Paper No. 97, Department of Finance, Kellogg Graduate School of Management, Northwestern University.

Deaton, A. and G. Laroque (1990), *On the Behavior of Commodity Prices*, Working Paper, June 1990.

Dickey, D.A. and W.A. Fuller (1979), 'Distribution of the Estimators for Autoregressive Time Series with a Unit Root', *Journal of the American Statistical Association*, 74, p.427–431.

Dickey, D.A. and W.A. Fuller (1981), 'Likelihood Ratio Sttistics for Autoregressive Time Series with a Unit Root', *Econometrica*, 49, p.1057–1072.

Engle, R.F. (1982), 'Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of U.K. Inflation', *Econometrica*, 50, p.987–1008.

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

Fuller, W.A. (1976), *Introduction to Statistical Time Series*, New York: John Wiley and Sons.

Gordon, J.D. (1985), *The Distribution of Daily Changes in Commodity Futures Prices*, Technical Bulletin No. 1536, Economic Research Service, USDA.

Granger, C.W.J. and P. Newbold (1974), 'Spurious Regressions in Econometrics', *Journal of Economics*, 2, p.111–120.

Hamilton, J.D. (1989), 'A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle', *Econometrica*, 57, p.357–384.

Hansen, B.E. and P.C.B. Phillips (1988), *Estimation and Inference in Models of Cointegration: A Simulation Study*, Cowles Foundation Discussion Paper No. 881.

Johansen, S. (1988), 'Statistical Analysis of Cointegration Vectors', *Journal of Economic Dynamics and Control*, 12, p.231–254.

Johnston, J. (1984), *Econometric Methods*, 3rd edition, New York: McGraw Hill.

Kwiatkowski, D., P.C.B. Phillips, P. Schmidt and Y. Chin (1992), 'Testing the Null Hypothesis of Stationarity Against the Alternative of a Unit Root: How Sure Are We that Economic Time Series have a Unit Root?', *Journal of Econometrics*, forthcoming.

Myers, R.J. (1991), 'Estimating Time-Varying Optimal Hedge Ratios on Futures Markets', *Journal of Futures Markets*, 11, p.39–53.

Myers, R.J. (1988), 'The Value of Ideal Contingency Markets in Agriculture', *American Journal of Agricultural Economics*, 70, p.255–267.

Newbery, D.M.G. and J.E. Stiglitz (1981), *The Theory of Commodity Price Stabilization*, Oxford: Oxford University Press.

Newey, W.K. and K.D. West (1987), 'A Simple, Positive Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix', *Econometrica*, 55, p.703–708.

Pagan, A. and A. Ullah (1988), 'The Econometric Analysis of Models with Risk Terms', *Journal of Applied Econometrics*, 3, p.87–105.

Perron, P. (1988), 'Trends and Random Walks in Macroeconomic Time Series: Further Evidence from a New Approach', *Journal of Economic Dynamics and Control*, 12, p.277–301.

Phillips, P.C.B. (1987), 'Time Series Regression with a Unit Root', *Econometrica*, 55, p.277–301.

Phillips, P.C.B. and B.E. Hansen (1988), *Statistical Inference in Instrumental Variables Regression with I(1) Processes*, Cowles Foundation Discussion Paper No. 869.

Phillips, P.C.B. and S. Ouliaris (1990), 'Asymptotic Properties of Residual Based Tests for Cointegration', *Econometrica*, 58, p.165–193.

Pindyck, R.S. and J. Rotemberg (1990), 'The Excess Comovement of Commodity Prices', *Economic Journal*, 100, p.1173–1189.

Schultz, H. (1938), *The Theory and Measurement of Demand*, Chicago: University of Chicago Press.

Sims, C.A., J.H. Stock and M.W. Watson (1990), 'Inference in Linear Time–Series Models with Some Unit Roots', *Econometrica*, 58, p.113–144.

Stock, J.H. and M.W. Watson (1988), 'Variable Trends in Economic Time Series', *Journal of Economic Perspectives*, 2, p.147–174.

Theil, H. (1971), *Principles of Econometrics*, New York: John Wiley and Sons.

Tomek, W.G. and K.L. Robinson (1977), 'Agricultural Price Analysis and Outlook', in L.R. Martin (ed.), *A Survey of Agricultural Economics Literature*, Volume 1, Minneapolis: University of Minnesota Press, p.327–409.

Vogelvang, B. (1990), *Testing for Cointegration with Spot Prices of Some Related Agricultural Commodities*, Research Memorandum 1990-1, Vrije Universiteit, Amsterdam.

Working, H. (1922), *Factors Determining the Price of Potatoes in St. Paul and Minneapolis*, University of Minnesota Agricultural Experiment Station Bulletin No. 10.

Table 1

Testing for Stochastic Trends in Cattle Prices[a]

Test	Statistic Under Lag Value			Critical Values	
	ℓ=0	ℓ=4	ℓ=12	5%	10%
Dickey–Fuller					
$p_t = \mu + \alpha p_{t-1} + \beta t + \varepsilon_t$					
$H_0 : \alpha=1$	−1.76	−2.27	−2.94	−3.43	−3.13
$H_0 : \alpha=1;\ \beta=0$	1.66	2.61	4.40	6.34	5.39
Phillips–Perron					
$p_t = \mu + \alpha p_{t-1} + \beta t + \varepsilon_t$					
$H_0 : \alpha=1$	−1.77	−2.15	−2.41	−3.43	−3.13
$H_0 : \alpha=1;\ \beta=0$	1.68	2.38	2.96	6.43	5.39
KPSS					
$p_t = \mu + w_t + \beta t + \varepsilon_t$					
$w_t = w_{t-1} + e_t$					
$H_0 : \sigma_e^2 = 0$	1.62	0.34	0.15	0.15	0.12

a The null hypothesis is rejected if the statistic is greater in absolute value than the critical value. Critical values are from Fuller (1976), Dickey and Fuller (1981), and Kwiatkowski, Phillips, Schmidt and Shin (1992).

Table 2

Testing for Stochastic Trends in Wheat Prices[a]

Test	Statistic Under Lag Value			Critical Values	
	ℓ=0	ℓ=4	ℓ=12	5%	10%
Dickey–Fuller					
$p_t = \mu + \alpha p_{t-1} + \beta t + \varepsilon_t$					
$H_0 : \alpha=1$	–1.13	–1.67	–2.40	–3.43	–3.13
$H_0 : \alpha=1;\ \beta=0$	3.01	3.07	3.60	6.34	5.39
Phillips–Perron					
$p_t = \mu + \alpha p_{t-1} + \beta t + \varepsilon_t$					
$H_0 : \alpha=1$	–1.14	–1.53	–1.58	–3.43	–3.13
$H_0 : \alpha=1;\ \beta=0$	3.06	2.56	2.59	6.34	5.39
KPSS					
$p_t = \mu + w_t + \beta t + \varepsilon_t$ $w_t = w_{t-1} + e_t$					
$H_0 : \sigma_e^2 = 0$	2.55	0.54	0.23	0.15	0.12

a The null hypothesis is rejected if the statistic is greater in absolute value than the critical value. Critical values are from Fuller (1976), Dickey and Fuller (1981), and Kwiatkowski, Phillips, Schmidt and Shin (1992).

Table 3

Testing for Cointegration Between Cattle and Wheat Prices[a]

$$p_{1t} = \mu + \delta p_{2t} + z_t$$

Test	Statistic Under Lag Value			Critical Values	
	$\ell=0$	$\ell=4$	$\ell=12$	5%	10%
Dickey–Fuller					
$\hat{z}_t = \alpha\hat{z}_{t-1} + \varepsilon_t$					
$H_0 : \alpha=1$	−0.29	−0.72	−1.09	−3.37	−3.07
Phillips–Perron					
$\hat{z}_t = \alpha\hat{z}_{t-1} + \varepsilon_t$					
$H_0 : \alpha=1$	−0.30	−0.68	−0.94	−3.37	−3.07

a The null hypothesis of no cointegration is rejected if the statistic is greater in absolute value than the critical value. Critical values are from Phillips and Ouliaris (1990).

Table 4

Testing for Conditional Heteroscedasticity in Cattle and Wheat Prices

Model	Statistic TR^2	Critical Values 5%	10%
Cattle			
$(1-a_1L-a_2L^2)\Delta p_t = (1+b_1L+b_2L^2+b_5L^5)e_t$			
$e_t \mid \Omega_{t-1} \sim D(0,h_t)$			
$h_t = \omega + \sum_{i=1}^{q} \alpha_i e_{t-i}^2$			
q=1	11.14	3.84	2.71
q=2	15.08	5.99	4.61
q=3	20.14	7.81	6.25
q=4	22.80	9.49	7.78
q=5	23.56	11.07	9.24
Wheat			
$(1-a_1L-a_2L^2)\Delta p_t = (1+b_1L+b_2L^2+b_8L^8)e_t$			
$e_t \mid \Omega_{t-1} \sim D(0,h_t)$			
$h_t = \omega + \sum_{i=1}^{q} \alpha_i e_{t-i}^2$			
q=1	0.00	3.84	2.71
q=2	5.84	5.99	4.61
q=3	6.74	7.81	6.25
q=4	7.39	9.49	7.78
q=5	7.48	11.07	9.24

Table 5

Simulation Results[a]

Estimator	Mean Estimate	Root Mean Square Error	Mean Estimate of Standard Error	Percent of Rejections at the 5% Level[b]
Case 1: Low Signal to Noise Ratio ($\delta = 1$)				
2SLS	0.507	0.0792	0.045	18.9
Bivariate OLS	0.275	0.2549	0.058	85.1
Multivariate OLS	0.367	0.1581	0.047	67.8
Case 2: High Signal to Noise Ratio ($\delta = 10$)				
2SLS	0.499	0.0067	0.004	17.9
Bivariate OLS	0.495	0.0110	0.009	9.1
Multivariate OLS	0.497	0.0072	0.006	7.3

a. Results are from estimating the supply elasticity in model (9) with 50 observations and 1000 repetitions of the experiment.

b. Percent of rejections of $H_0 : \hat{\gamma}_{11} = 0.5$ using the standard regression t-statistic and the conventional 5% significance level in a two-sided test.

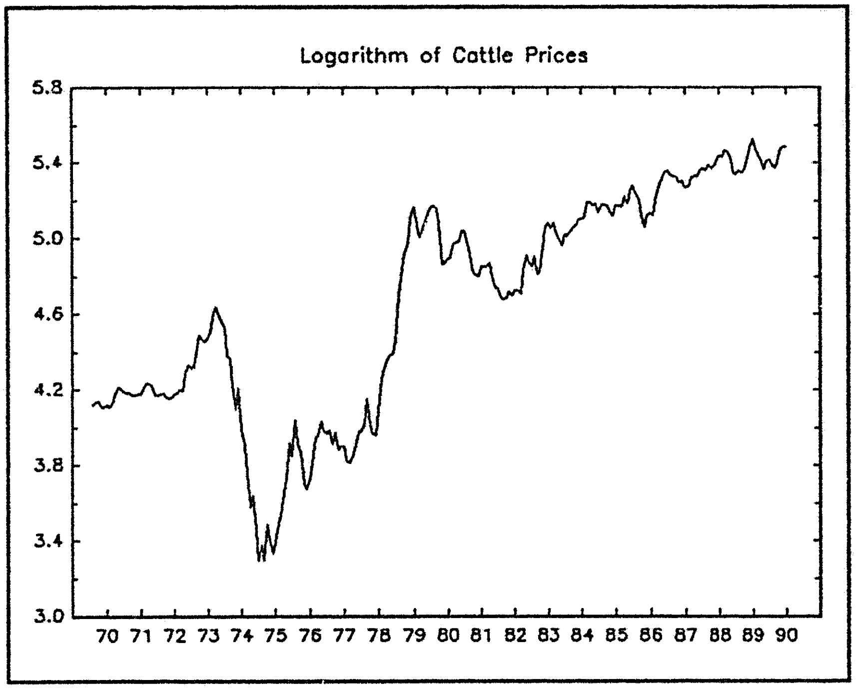

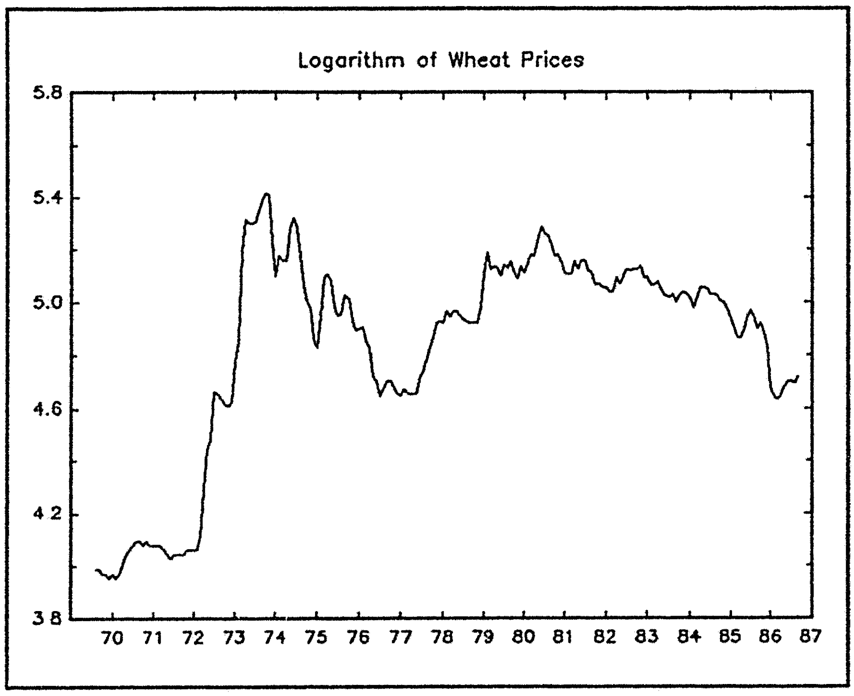

Figure 1. Logarithm of Monthly Commodity Prices