Time Series Econometrics and Commodity Price Analysis: A Review

Robert J. Myers*

1. Introduction

The econometric analysis of commodity prices has a long and distinguished history highlighted by structural investigations of supply and demand systems; forecasting supplies and prices; evaluating market efficiency and the effects of price policies, etc., (see reviews by Tomek and Robinson 1977, and Tomek and Myers 1993). Much of this research relies on a standard set of econometric methods, as outlined in books such as Theil (1971) and Johnston (1984). However, evolving developments in time series econometrics have cast doubt over the use of standard econometric methods for estimating commodity market models. The goals in this paper are to review some developments taking place in the time series literature and to discuss their implications for modelling commodity prices and markets.

The time series developments that will be discussed include stochastic trends (unit roots); common stochastic trends driving multiple time series (cointegration); and time-varying volatility in the innovations of time series (conditional heteroscedasticity). None of these developments are new and all have been discussed in econometric textbooks such as Harvey (1990), Lutkepohl (1992), Cuthbertson, Hall, and Taylor (1992), and Griffiths, Hill and Judge (1993). While the developments themselves are well known, however, the resulting implications for commodity price analysis appear not to be widely appreciated. The present paper is intended to help correct this situation and provide a modern perspective on econometric modelling of commodity markets using time series data.

The paper is divided into three parts. First, the characteristic time series properties of commodity prices are examined within the context of the time series literature. Second, some of the implications which these properties have for the econometric analysis of commodity prices and markets are discussed. Third, the paper reports a simulation study which highlights some of the econometric problems arising from the time series properties of commodity prices.

2. Characteristic Time Series Properties of Commodity Prices

Prices of different commodities are influenced by distinct forces and therefore will behave somewhat differently. Nevertheless, there are several characteristic time series properties which many commodity price series seem to share in common. In this section of the paper, some of these characteristic properties of commodity price data are outlined and discussed.

2.1 High Volatility

Prices of primary commodities are often highly volatile, particularly compared to prices of manufactured consumer goods (Newbery and Stiglitz

* Department of Agricultural Economics, Michigan State University, East Lansing MI, 48824, USA. This review is a revision of my invited paper "Time-Series Econometrics and Commodity Price Analysis" presented at the 36th Annual Conference of the Australian Agricultural Economics Society, Australian National University, Canberra, February 10-11, 1992. That paper was written while I was at the Department of Economics, University of Queensland. Comments and suggestions from Steve Hanson and the reviewers are gratefully acknowledged.

Review coordinated by Garry Griffith.

1981). This volatility can pose major problems to industry participants and policymakers, particularly in countries whose export earnings and GDP depend heavily on sales of primary commodities.

One interesting question surrounding commodity price movements involves the extent to which high volatility is indicative of market inefficiency. It has been argued that commodity prices are flexible while wages and prices of manufactured goods are fixed or rigid, at least in the short run (Okun 1975; Rausser 1985). Under this interpretation, commodity prices "overshoot" their long-run equilibrium levels and high commodity price volatility indicates economic inefficiency (Dornbusch 1976; Frankel 1986).

Another issue surrounding commodity price volatility is management of the resulting price risks. Market instruments, such as futures and options, are available to some market participants but government regulation, in the form of price stabilisation schemes, has also been common. It is now widely understood that the efficiency of market mechanisms for dealing with risk depends on the completeness of the market structure (Hart 1975; 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 then some risks are uninsurable and there may be a role for government in designing risk sharing mechanisms which improve the distribution of risk throughout the economy (Newbery and Stiglitz 1981; Innes and Rausser 1989).

2.2 Stochastic Trends

Another characteristic property of many commodity price series, at least when sampled at high frequencies (daily, weekly, or even monthly intervals) is that they appear to contain stochastic trends (Ardeni 1989; Baillie and Myers 1991; Goodwin 1992; Goodwin and Schroeder 1991). A stochastic trend increases by some fixed amount on average but in any given period the 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:

$$(1) \qquad r_t - r_{t-1} = \mu + \varepsilon_t$$

where the drift parameter μ is the average change in r_t each period; and ε_t is a serially uncorrelated random shock. If a commodity price followed a pure stochastic trend then r_t would represent the price level and $r_t - r_{t-1}$ would represent the price change from period to period.

A commodity price p_t may contain a stochastic trend but also be subject to stationary deviations around the stochastic trend. In this case, the price can be written as the sum of a stochastic trend component r_t and a stationary component z_t:

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

where z_t is a stationary stochastic process. Given a set of data on commodity prices only, it may be difficult to explicitly separate out the stochastic trend component from the stationary component (see Stock and Watson 1988). Nevertheless, representing prices as the sum of a pure stochastic trend component, and a component representing stationary deviations around the trend, is a useful way of conceptualizing the time-series properties of many commodity prices. For example, 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 stochastic trend component and a stationary component.

Over the past decade, important advances have been made in developing statistical tests for stochastic trends. The original unit root tests developed by Dickey and Fuller (1979, 1981) are based on regressions of the form:

(3a) $p_t = \hat{\alpha} p_{t-1} + \hat{\varepsilon}_t$

(3b) $p_t = \mu^* + \alpha^* p_{t-1} + \varepsilon^* t$

(3c) $p_t = \tilde{\mu} + \tilde{\alpha} p_{t-1} + \tilde{\beta} t + \tilde{\varepsilon}_t$

for t=1,2,.., T. The Dickey-Fuller tests use standard t and F statistics computed from these regressions but the statistics follow a non-standard distribution. Full details of the tests and appropriate critical values are available in Fuller (1976, pp.366-382) and Dickey and Fuller (1981).

A frequent problem in applications of Dickey-Fuller tests is that the residuals from the regressions (3) are autocorrelated, thus violating an important assumption underlying construction of the tests. The augmented Dickey-Fuller tests are designed to correct this problem. In augmented Dickey-Fuller tests the regression equations (3) are expanded by including lagged differences of the dependent variables as additional explanatory variables. The testing procedure then proceeds exactly as before, with all of the relevant statistics having the same limiting distributions.

Unfortunately, the augmented Dickey-Fuller tests are not able to deal effectively with all of the distributional problems typically found in applications. Suppose, for example, that the error term has moving average as well as autoregressive terms. Then inclusion of any finite number of lagged differences will not eliminate autocorrelation in the residuals. Furthermore, suppose that the variance of the residuals changes over time. Then this heterogeneity in the distribution of the residuals will lead to biased and inconsistent test results. To help overcome these problems, Phillips (1987) and Perron (1988) have developed an alternative set of tests which are more robust to autocorrelation and heterogeneity in the distribution of the residuals. The tests are based on regressions like (3) and focus on the same null and alternative hypotheses. However, the relevant test statistics are not just standard *t* and *F* statistics. Full details of the formulae for constructing these Phillips-Perron test statistics, and information on the appropriate distributions to use to test the hypotheses, are available in Perron (1988).

A problem with both the Dickey-Fuller and Phillips-Perron 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 stochastic trend (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. The KPSS test can be a useful consistency check for the more conventional Dickey-Fuller and Phillips-Perron tests.

Many of these unit root tests have been applied to commodity price data and results indicate that commodity price data sampled at high frequencies show consistent evidence of stochastic trends (e.g. Ardeni 1989; Baillie and Myers 1992; Goodwin 1992; Goodwin and Schroeder 1991). However, the evidence is far less clear for low frequency (annual) data. The explanation for this discrepancy may lie in the smaller number of annual observations typically available and/or in the low power of unit root tests. Deaton and Laroque (1992) outline a theoretical model for storable commodities in which the equilibrium price process shifts between two regimes depending on whether speculative inventories are positive or zero. The positive storage regime features a stochastic trend but even infrequent shifts to the zero storage regime results in a price process that is stationary in the long run. They also argue that establishing the long-run properties of commodity price series using the number of observations typically available in practice is going to be extremely difficult.

2.3 Comovements in Commodity Price Series

Many commodity prices share a tendency to move together over time, even when the commodities themselves are largely unrelated in both production and consumption (Pindyck and Rotemberg 1990). There are three main reasons for such comovements. First, it could be that supply and demand shocks to any one commodity 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 similarly. Common macroeconomic shocks undoubtedly explain some of the comovement among commodity prices but 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.

One way to formalise the idea of comovements among commodity prices is to use the theory of cointegrated stochastic processes. We have already seen that many commodity prices can be represented as the sum of a stochastic trend and stationary deviations around the trend. In these circumstances two commodity prices are said to be cointegrated if they share the same stochastic trend:

(4a) $p_{1t} = r_t + z_{1t}$

(4b) $p_{2t} = \delta r_t + z_{2t}$

where p_{it} is commodity price i; r_t is the common stochastic trend; z_{it} is stationary component *i*; and δ is a scaling parameter. Substituting (4a) into (4b) gives

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

where $z_t = z_{2t} - \delta z_{1t}$.

The first component $p_{2t} = \delta p_{1t}$ of (5) represents a long-run equilibrium relationship between the two prices resulting from their common stochastic trend[1]. The two prices return to this linear relationship in the long run, even though prices may deviate from the relationship in the short run as z_t varies. It is important to realize that the long-run equilibrium relationship has no causal interpretation in the usual sense. Thus, when there is a shock to the long-run equilibrium it makes no sense to think of holding one variable, say p_{1t}, fixed and computing the path of the other, p_{2t}, back to the long-run equilibrium. In practice, both p_{1t} and p_{2t} adjust to the shock and all one can say is that eventually the long-run equilibrium relationship will be re-established. Thus, there is no unique "dependent variable" in (5) and the choice of a "dependent variable" is really just an arbitrary normalisation.

The parameter δ which characterises the nature of the long-run equilibrium relationship can be estimated by applying ordinary least squares (OLS) to (5). Perhaps surprisingly, OLS estimates the parameters of this *cointegrating regression* consistently, even when z_t is autocorrelated, heteroscedastic, 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 because every other linear combination of p_{1t} and p_{2t}, besides that represented by the long-run equilibrium, has asymptotically infinite variance. Nevertheless, OLS estimates of a cointegrating regression generally follow a nonstandard distribution theory, even asymptotically, and so one must

[1] Because z_t is not necessarily assumed to have zero mean then a constant term can be added to the long-run equilibrium relationship without any loss of generality.

be very careful undertaking hypothesis tests. Moreover, OLS is a biased estimator of the cointegrating regression in small samples, and this small sample bias can be large. Details of estimation and testing of cointegrating regressions can be found in Engle and Granger (1987), Engle and Yoo (1987), Harvey (1990), and Cuthbertson, Hall and Taylor (1992).

A limitation of univariate OLS based tests for cointegration is that there is no systematic way to investigate all possible cointegrating vectors in a multivariate system. To overcome this problem, Johansen (1988, 1991) and Johansen and Juselius (1990) have developed maximum likelihood methods for testing and estimating cointegrating vectors in a multivariate framework. Johansen uses moment and cross-moment matrices from auxiliary regressions to test the null hypothesis that an (nx1) vector of variables, say a vector of commodity prices p_t, has at most k cointegrating vectors, and to estimate the resulting "error correction model" with cointegration restrictions imposed. Full details of Johansen's approach are available in the references already cited.

Empirical tests for cointegration among commodity prices have provided mixed results. Goodwin and Schroeder (1991) find evidence of cointegration among regional US cattle prices and Goodwin (1992), using Johansen's multivariate testing framework, supports the hypothesis of one cointegrating vector among five international wheat prices. However, Ardeni (1989) investigates several commodity prices, including wheat, at different locations and concludes that the evidence for cointegration is weak, even when considering prices for the same commodity at different locations.

2.4 Time-Varying Volatility

Some of the earliest research on the distributions of commodity prices assumed that price changes are independent draws from an identical normal distribution. It soon became apparent, however, that the volatility of price changes varies over time as the series moves between volatile periods, where large price changes tend to be followed by other large changes, and tranquil periods, where small price changes tend to be followed by other small changes. This temporal instability in the variance of commodity price data has become a well-known feature in empirical studies using high frequency data (e.g. Baillie and Myers 1992; Yang and Brorsen 1992).

It should be clear that time-varying volatility can be 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 has been removed by first differencing. Time-varying volatility in commodity prices 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.

Bollerslev (1986) generalised ARCH to the 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) $h_t = \omega + \alpha e^2_{t-1} + \beta h_{t-1}$

In this example, price changes equal 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.

ARCH and GARCH models are highly nonlinear and maximum likelihood is the usual estimation approach (Harvey 1990). However, Engle (1982) has developed a simple Lagrange multiplier test for conditional heteroscedasticity which can be undertaken using OLS procedures. Applications of this testing and estimation framework to high frequency commodity price data have led to consistent rejections of the no ARCH effects hypothesis, suggesting that conditional heteroscedasticity is a common characteristic of commodity price data sampled at high frequencies (e.g. Baillie and Myers 1991; Yang and Brorsen 1992).

2.5 Excess Kurtosis

Early research on the distribution of commodity prices also assumed that price changes are normally (or lognormally) 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; Deaton and Laroque 1992).

ARCH models lead to a partial solution of 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 (Baillie and Myers 1992; Yang and Brorsen 1992). One solution 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 capture the excess kurtosis inherent in commodity price data.

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 identically and independently distributed errors (often normal as well) is also invoked. As just discussed, however, these assumptions are inappropriate for many commodity price data. Commodity prices may have 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.

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 relatively 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 a simple textbook model of commodity supply and demand:

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

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

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

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. For now, the supply and demand disturbances are defined as generally as possible: they may be autocorrelated so as to generate market dynamics; they may be heteroscedastic reflecting time-varying volatility; and they may be stationary or have stochastic trends depending on how the variables in the system interact with one another. All variables are in logarithms, so that the γ_{ij} coefficients can be interpreted as supply and demand elasticities. Furthermore, to simplify the presentation it is assumed that any deterministic components (mean and/or deterministic trends) have been removed prior to model specification. The variables in this model could easily be interpreted as vectors, and/or lagged prices and quantities could be included in the model, without changing the substance of the discussion which follows.

The supply and demand shift variables are assumed to follow 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} may be correlated with each other but not with the structural disturbances u_{1t} and u_{2t}. Again, these processes could be allowed to be more complicated without changing the substance of the arguments which follow. Note that the exogenous variables may have a stochastic trend or be stationary, depending on the value of ρ_i. If $\rho_1=1$ and k_t and x_t are cointegrated then Equation (8b) is replaced by

(9) $x_t = \delta k_t + \varepsilon_{2t}$

where δ is the scaling parameter defining the long-run equilibrium relationship between the supply and demand shift variables.

The model is exactly identified and has reduced form:

(10a) $p_t = \pi_{11}x_t + \pi_{12}k_t + v_{1t}$

(10b) $y_t = \pi_{21}x_t + \pi_{22}k_t + v_{2t}$

where the π_{ij} are functions of the structural parameters; and v_{1t} and v_{2t} are correlated functions of the structural errors. Four different cases are now examined, each characterised by different assumptions about which variables are stationary and which have stochastic trends.

Case 1: Stationary Regressors

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, although the IV estimator is consistent and asymptotically normal, its small sample properties are generally unknown.

Case 2: Two Distinct Stochastic Trends

Now suppose that the supply and demand shift variables follow distinct random walks, $\rho_1 = \rho_2 = 1$, and therefore are not cointegrated. Then by the reduced form, price and quantity generally have stochastic trends as well because they are linear combinations 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 supply and demand equations represent spurious regressions in the sense of Granger and Newbold (1974). Results from IV (or OLS) 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 nonstandard and application of standard techniques will lead to major problems.

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 cointegrating relationships and the relevant supply and demand elasticities represent the long-run relationship between the series (i.e. the cointegrating vectors). Because the structural disturbances may be autocorrelated then the system may be subject to short-run dynamics and so the elasticities represented by the cointegrating vectors are best interpreted as long-run elasticities.

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 1990). However, the IV estimator is biased in small samples and its asymptotic distribution 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), other than that given by the supply equation itself, have asymptotically infinite variance. Because OLS minimises the residual variance it moves quickly to the finite residual variance defined by the supply equation parameters. Like the IV estimator, the OLS estimator of cointegrated supply and demand equations is biased in small samples and its asymptotic distribution is generally not normal. Thus, conventional inference cannot proceed as usual for this estimator either.

Two main approaches to overcoming problems with standard OLS and IV estimation of cointegrating regressions have been developed. The first is full information maximum likelihood (FIML). Phillips (1991) and Johansen (1991) have shown that FIML provides optimal inference, provided the restrictions implied by cointegration are imposed during estimation. Small sample bias is reduced and the usual asymptotic hypothesis tests are generally applicable. Furthermore, Gonzalo (1989) provides monte carlo evidence suggesting that FIML estimation of cointegrated systems provides superior estimates (lower small sample bias and more accurate standard errors), even in small samples with errors that deviate from normality. The second approach is to "fully modify" OLS and/or IV estimators so as to reduce small sample bias and allow conventional asymptotic inference (Phillips and Hansen 1990). The "fully modified" estimators have asymptotic mixed normal distributions which allow general hypothesis tests using conventional techniques. Furthermore, the "fully modified" OLS and IV estimators are asymptotically equivalent to FIML (Phillips 1991).

Because OLS and IV are both consistent estimators of cointegrating regressions, it is of interest to determine which gives better results for cointegrated supply and demand systems in small samples. Monte carlo simulations have shown that the performance of the two estimators in small samples depends on the signal to noise ratio (Phillips and Hanson 1990). If the variance of the innovations in the stochastic trends which drive the long-run behavior of the variables is high relative to the variance of the short-term dynamics (the structural disturbances in the supply and demand model) then the signal to noise ratio is high, small sample bias is low, and OLS and IV both provide good estimates of the cointegrating vector (although hypothesis testing is nonstandard). If the variance of the innovations in the stochastic trends is low relative to the variance of the short-term dynamics, then the signal to noise ratio is low and IV estimation generally provides a better estimate of the cointegrating vector than OLS. The reason is that the simultaneity bias in OLS, which goes to zero asymptotically, can be quite high in small samples when the signal to noise ratio is low. IV estimation helps reduce this small sample simultaneity bias. Nevertheless, standard errors computed in the conventional way are again subject to bias and so hypothesis testing is nonstandard, even asymptotically. Some monte carlo evidence supporting these conclusions for the simple commodity market model is presented below.

Case 3: One Stochastic Trend With Stationary Supply Shifter

Now 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. One approach is to simply 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 quickly to the true parameter value, though it remains biased in small samples and 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 variable. The supply elasticity could also be estimated by a multiple regression of quantity on price and the supply shifter k_t, or by using conventional IV techniques. The relative performance of these estimators in small samples should depend on the signal to noise ratio as discussed above, although in all three cases conventional standard errors are not to be trusted. The preferred alternative, of course, is to use maximum likelihood and impose cointegration restrictions during estimation.

Case 4: One Stochastic Trend With Cointegrated Supply and Demand Shifters

Finally, consider the possibility that the 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.[2] In this case the reduced form quantity equation (10b) has a stationary dependent variable and a linear combination of the stochastically trending regressors is stationary. In the structural form, the dependent variable quantity in the supply and demand equations is stationary but the explana-

[2] It is, of course, possible that the supply and demand shift variables are cointegrated but the cointegrating relationship is not represented by either reduced form equation, so that equilibrium price and quantity both have stochastic trends. In this case, however, the supply and demand equations are either spurious regressions or cointegrating regressions and so estimation and inference reverts to Case 2 which has already been discussed.

tory variables have a stochastic trend. Thus $\gamma_{11}p_t + \gamma_{12}k_t$ and $\gamma_{21}p_t + \gamma_{22}x_t$ represent stationary linear combinations (cointegrating relationships) and the model can be transformed so that all of the coefficients of interest can be written as coefficients on stationary variables. Sims, Stock and Watson (1990) have shown that when a regression contains explanatory variables with stochastic trends, but the equation can be rearranged so that all of the coefficients of interest can be written as the coefficient on a stationary variable, then conventional estimation and distribution theory is applicable.

The Sims, Stock and Watson (1990) result implies that when the supply and demand shifters are cointegrated in a way that leaves quantity stationary then conventional inference is applicable. Thus, the IV estimator should perform better than OLS in this case because IV accounts for the simultaneity between price and quantity in the conventional estimation framework, while OLS does not. The monte carlo evidence presented below supports these conclusions.

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. 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.

It is important to note, however, that time-varying volatility and excess kurtosis may not be as crucial in 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 these distributional problems. On the other hand, it is crucial that hypothesis testing procedures be robust to heteroscedasticity and excess kurtosis. Furthermore, heteroscedasticity and excess kurtosis violate the assumptions underlying Johansen's multivariate cointegration testing and estimation framework.

4. A Monte Carlo Simulation

A monte carlo experiment was undertaken to highlight some of the effects from using OLS and IV estimators of commodity market models under various forms of cointegration. The data generating process used in the simulation is a simple version of the commodity market model:

$$\text{(11a)} \quad \begin{aligned} y_t &= p_t + k_t + e_t\,; \\ e_t &= 0.5e_{t-1} + u_{1t} \end{aligned} \qquad \text{(Supply)}$$

$$\text{(11b)} \quad y_t = -p_t + x_t + u_{2t} \qquad \text{(Demand)}$$

$$\text{(11c)} \quad k_t = \rho_1 k_{t-1} + \lambda\varepsilon_{1t}$$

$$\text{(11d)} \quad x_t = \rho_2 x_{t-1} + \lambda\varepsilon_{2t}$$

Each of the disturbance terms u_{it} and ε_{it} are identically and independently distributed N(0,1) variables. The supply disturbance is autocorrelated to introduce the effects of short-run dynamics into the results. The parameters ρ_1 and ρ_2 are set equal to different values depending on which case is being investigated. The same four cases discussed above were simulated. For Case 1, where all variables are stationary, then $\rho_1 = \rho_2 = 0$. For Case 2, where there are two distinct stochastic trends, then $\rho_1 = \rho_2 = 1$. For Case 3, where the supply shifter is stationary but the demand shifter has a stochastic trend, then $\rho_1 = 0$ but $\rho_2 = 1$. Finally for Case 4, where the supply and demand shift variables are cointegrated then $\rho_1 = 1$ and (11d) is replaced by $x_t = -k_t + \lambda\varepsilon_{2t}$. Notice that this particular cointegration relationship ensures that quantity is stationary while price has a stochastic trend.

The parameter λ determines the signal to noise ratio. If λ is large then the variance of the random walks are large relative to the variance of supply and demand disturbances (high signal to noise ratio) while if λ is small the reverse occurs. The experiment was repeated for the two values, $\lambda = 1$ and $\lambda = 10$.

Each time the model was simulated, a series of observations on four independent N(0,1) variables was drawn using the random number generator RNDN in GAUSS. These random numbers were then applied to Equation (11) to generate a series of observations on (y_t, p_t, k_t, e_t), using zero as a start up value for x_t, k_t and e_t. The number of observations generated exceeds the desired sample size by 100 and then the first 100 observations of each sample are discarded to ensure that startup values do not have a major impact on the results. IV and OLS estimators were then applied to the data set and their performance compared. The whole process was then repeated a total of 10,000 times and summary results tabulated.

It is assumed that the aim of the exercise is to estimate the price elasticity of supply. In most cases two alternative estimators were applied. First is the conventional IV (2SLS) estimator of the supply equation, using k_t and x_t as instruments for p_t. Second is an OLS regression of y_t on (p_t, k_t). For Case 3, where k_t is stationary, a simple OLS regression of y_t on p_t is also applied. Because the supply equation is a cointegrating regression, and the supply shift variable k_t is stationary, the simple OLS regression should generate a consistent estimate of the supply elasticity in this case. The estimation was repeated for two sample sizes, 500 and 50. The large sample size is used to gain some insight into the asymptotic properties of the estimators while the small sample size is used to investigate their small sample properties.

The large sample results are reported in Table 1. In the case of low signal to noise ratio ($\lambda = 1$) 2SLS clearly outperforms OLS when all variables are stationary. 2SLS has a mean estimate very close to the true supply elasticity of one, has lower root mean square error (RMSE) than the OLS estimator, and rejects the null hypothesis that the supply elasticity equals one 4.91 per cent of the time, which is quite close to the theoretical value of 5 per cent. Because price and quantity are simultaneously determined the OLS estimator is very biased and so the null hypothesis that the supply elasticity is one is rejected 100 per cent of the time under OLS.

The situation is quite different in Case 2 where there are two distinct stochastic trends and the supply equation is a cointegrating regression. 2SLS still performs best but OLS does much better. The bias in the OLS estimator is reduced significantly and would go to zero asymptotically. Notice, however, that both the OLS and 2SLS estimators require nonstandard inference procedures, even when sample sizes are large. The 2SLS estimator rejects the null that the supply elasticity is one 26 per cent of the time using a standard t-test, when the actual number of rejections should be 5 per cent. For OLS this figure is 83 per cent. Similar results occur in Case 3 when there is only one stochastic trend. In this case, however, the supply shift variable is stationary and so excluding it from the regression should not affect the consistency of OLS. Indeed, a simple regression of quantity on price gives a mean supply elasticity estimate of 0.91 which is quite close to the true value of one. Nevertheless, the simple OLS regression has more bias and higher mean square error than the multiple regression including the supply shifter. Thus, it appears that there is an advantage to including additional relevant regressors, even when they are stationary.

In case 4, where the supply and demand shifters are cointegrated and quantity is stationary then standard inference procedures are applicable. This can be seen by noting the return of the bias in the OLS estimator and the percent of rejections in the 2SLS estimator returning to its theoretical value of 5 per cent.

Table 1: Large Sample Simulation Results for Supply Elasticity Estimation[a]

Case and Estimator	Mean Supply Elasticity Estimate	Root Mean Square Error	Percent of Rejections at the 5% Level[b]
Low Signal to Noise Ratio ($\lambda = 1$)			
Case 1: All Variables Stationary			
2SLS	1.01	0.105	4.91
OLS	0.20	0.799	100.00
Case 2: Two Distinct Stochastic Trends			
2SLS	1.00	0.030	26.32
OLS	0.93	0.086	83.30
Case 3: One Stochastic Trend With Stationary Supply Shifter			
2SLS	1.00	0.026	26.46
OLS	0.95	0.068	75.43
Simple OLS	0.91	0.111	88.93
Case 4: One Stochastic Trend With Cointegrated Supply and Demand Shifters			
2SLS	1.01	0.105	4.99
OLS	0.20	0.797	100.00
High Signal to Noise Ratio ($\lambda = 10$)			
Case 1: All Variables Stationary			
2SLS	1.00	0.010	5.24
OLS	0.97	0.028	72.64
Case 2: Two Distinct Stochastic Trends			
2SLS	1.00	0.003	25.41
OLS	1.00	0.003	27.25
Case 3: One Stochastic Trend With Stationary Supply Shifter			
2SLS	1.00	0.003	26.28
OLS	1.00	0.003	27.17
Simple OLS	0.96	0.054	76.29
Case 4: One Stochastic Trend With Cointegrated Supply and Demand Shifters			
2SLS	1.00	0.010	5.38
OLS	0.97	0.028	72.61

a Results from estimating the supply elasticity with 500 observations and 10,000 repetitions of the experiment. Case 1 is $\rho_1 = \rho_2 = 0$;; Case 2 is $\rho_1 = \rho_2 = 1$; Case 3 is $\rho_1 = 0$ and $\rho_2 = 1$; Case 4 is $\rho_1 = 1$ and $x_t = -k_t + \lambda\varepsilon_{2t}$.

b Percent rejections of H_0: $\lambda_{11} = 1.0$ using the standard regression t-test and the conventional 5% significance level in a two-sided test.

Large sample results when there is a high signal to noise ratio are reported in the second part of Table 1. Results are very similar to the case of low signal to noise ratio except that there is almost no difference between 2SLS and OLS when the supply elasticity is a cointegrating regression (Cases 2 and 3) and there is a smaller difference between simple OLS (without the supply shifter) and multiple OLS (with the supply shifter) when the supply shift variable is stationary (Case 3). Notice, however, that standard inference procedures can still lead to errors when the supply equation is a cointegrating regression, and that this is true irrespective of whether OLS or 2SLS is used (see the number of rejections far exceeding the theoretical value of 5 per cent in Cases 2 and 3 of Table 1).

Table 2: Small Sample Simulation Results for Supply Elasticity Estimation[a]

Case and Estimator	Mean Supply Elasticity Estimate	Root Mean Square Error	Percent of Rejections at the 5% Level[b]
Low Signal to Noise Ratio ($\lambda = 1$)			
Case 1: All Variables Stationary			
2SLS	1.08	0.494	5.66
OLS	0.22	0.795	99.82
Case 2: Two Distinct Stochastic Trends			
2SLS	1.04	0.295	20.24
OLS	0.63	0.425	76.60
Case 3: One Stochastic Trend With Stationary Supply Shifter			
2SLS	1.03	0.263	21.44
OLS	0.69	0.375	70.52
Simple OLS	0.51	0.553	85.57
Case 4: One Stochastic Trend With Cointegrated Supply and Demand Shifters			
2SLS	1.07	0.486	5.58
OLS	0.24	0.780	99.75
High Signal to Noise Ratio ($\lambda = 10$)			
Case 1: All Variables Stationary			
2SLS	1.00	0.033	5.76
OLS	0.97	0.041	11.91
Case 2: Two Distinct Stochastic Trends			
2SLS	1.00	0.026	23.13
OLS	0.99	0.026	22.51
Case 3: One Stochastic Trend With Stationary Supply Shifter			
2SLS	1.00	0.023	24.37
OLS	1.00	0.024	23.64
Simple OLS	0.68	0.379	73.03
Case 4: One Stochastic Trend With Cointegrated Supply and Demand Shifters			
2SLS	1.00	0.033	5.76
OLS	0.98	0.040	11.84

[a] Results from estimating the supply elasticity with 50 observations and 10,000 repetitions of the experiment. Case 1 is $\rho_1 = \rho_2 = 0$;; Case 2 is $\rho_1 = \rho_2 = 1$; Case 3 is $\rho_1 = 0$ and $\rho_2 = 1$; Case 4 is $\rho_1 = 1$ and $x_t = -k_t + \lambda\varepsilon_{2t}$.

[b] Percent rejections of H_0: $\lambda_{11} = 1.0$ using the standard regression t-test and the conventional 5% significance level in a two-sided test.

Results from the small sample simulations are reported in Table 2. The pattern of these results is very similar to the large sample results. The main difference is that, in small samples, OLS estimation of cointegrating regressions seems to perform even worse relative to 2SLS than in the case of large samples, particularly under low signal to noise ratio. The reason for this is that the small sample size exacerbates the simultaneity bias in the OLS estimator. 2SLS helps reduce this bias by accounting for the simultaneous determination of price and quantity. Once again, however, standard inference procedures are not applicable with either OLS or 2SLS when estimating a cointegrating regression (see the high rate of rejections in Cases 2 and 3).

The main conclusion from these results is that 2SLS (or some other IV procedure) is generally preferable to OLS estimation of cointegrating regressions in commodity market models featuring simultaneity. OLS may perform as well as 2SLS when sample sizes are large and there is a high signal to noise ratio. In general however, the small sample simultaneity bias in the OLS estimator causes its performance to deteriorate relative to 2SLS, which takes this simultaneity bias into account. Standard inference procedures are not applicable with *either* estimator, irrespective of sample size, and a FIML or "fully modified" IV approach is required for proper hypothesis testing.

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 prices. These time series characteristics have 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. Time-varying volatility and excess kurtosis in commodity prices also need to be properly accounted for if estimation is to proceed with full 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.

References

ARDENI, P.G. (1989), "Does the Law of One Price Really Hold?", *American Journal of Agricultural Economics* 71(3), 661-669.

BAILLIE, R.T. and MYERS, R.J. (1991), "Bivariate GARCH Estimation of the Optimal Commodity Futures Hedge", *Journal of Applied Econometrics* 6(2), 109-124.

BEVERIDGE, S. and NELSON, C.R. (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(1), 151-174.

BOLLERSLEV, T. (1986), "Generalized Autoregressive Conditional Heteroscedasticity", *Journal of Econometrics* 31(3), 307-327.

CUTHBERTSON, K., HALL, S.G. and TAYLOR, M.P. (1992), *Applied Econometric Techniques*, University of Michigan Press, Ann Arbor.

DEATON, A. and LAROQUE, G. (1992), "On the Behavior of Commodity Prices", *Review of Economic Studies* 59(1), 1-24.

DICKEY, D.A. and FULLER, W.A. (1979), "Distribution of the Estimators for Autoregressive Time Series with a Unit Root", *Journal of the American Statistical Association* 74(366), 427-431.

DICKEY, D.A. and FULLER, W.A. (1981), "Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root", *Econometrica* 49(5), 1057-1072.

DORNBUSCH, R. (1976), "Expectations and Exchange Rate Dynamics", *Journal of Political Economy* 84(6), 1161-1176.

ENGLE, R.F. (1982), "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of U.K. Inflation", *Econometrica* 50(4), 987-1008.

ENGLE, R.F. and GRANGER, C.W.J. (1987), "Co-Integration and Error Correction: Representation, Estimation and Testing", *Econometrica* 55(2), 251-276.

ENGLE, R.F. and YOO, B.S. (1987), "Forecasting and Testing in Cointegrated Systems", *Journal of Econometrics* 35(1), 143-159.

FRANKEL, J.A. (1986), "Expectations and Commodity Price Dynamics: The Overshooting Model", *American Journal of Agricultural Economics* 68(2), 344-350.

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

GONZALO, J. (1989), *Comparison of Five Alternative Methods of Estimating Long-Run Equilibrium Relationships*, Discussion Paper 89-55, University of California, San Diego.

GOODWIN, B.K. (1992), "Multivariate Cointegration Tests and the Law of One Price in International Wheat Markets", *Review of Agricultural Economics* 14(1), 117-124.

GOODWIN, B.K. and SCHROEDER, T.C. (1991), "Cointegration Tests and Spatial Price Linkages in Regional Cattle Markets", *American Journal of Agricultural Economics* 73(2), 452-464.

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

GRANGER, C.W.J. and NEWBOLD, P. (1974), "Spurious Regressions in Econometrics", *Journal of Econometrics* 2(2), 111-120.

GRIFFITHS, W.E., HILL, R.C. and JUDGE, G. (1992), *Learning and Practicing Econometrics*, Wiley, New York.

HARVEY, A. (1990), *The Econometric Analysis of Time Series*, 2nd Edition, MIT Press, Cambridge.

HART, O.D. (1975), "On the Optimality of Equilibrium When the Market Structure is Incomplete", *Journal of Economic Theory* 11(2), 418-443.

INNES, R.D. and RAUSSER, G. (1989), "Incomplete Markets and Government Agricultural Policy", *American Journal of Agricultural Economics* 71(4), 915-931.

JOHANSEN, S. (1988), "Statistical Analysis of Cointegration Vectors", *Journal of Economic Dynamics and Control* 12(2), 231-254.

JOHANSEN, S. (1991), "Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models", *Econometrica* 59(6), 1551-1580.

JOHANSEN, S. and JUSELIUS, K. (1990), "Maximum Likelihood and Inference on Cointegration--With Applications to the Demand for Money", *Oxford Bulletin of Economics and Statistics* 52(2), 169-210.

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

KWIATKOWSKI, D., PHILLIPS, P.C.B., SCHMIDT, P. and SHIN, Y. (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* 54(1), 159-178.

LUTKEPOHL, H. (1992), *Introduction to Multiple Time Series Analysis*, Springer, New York.

MYERS, R.J. (1988), "The Value of Ideal Contingency Markets in Agriculture", *American Journal of Agricultural Economics* 70(2), 255-267.

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

OKUN, A. (1975), "Inflation: Its Mechanics and Welfare Costs", *Brookings Papers on Economic Activity* 2, 351-390.

PERRON, P. (1988), "Trends and Random Walks in Macroeconomic Time Series: Further Evidence from a New Approach", *Journal of Economic Dynamics and Control* 12(2), 277-301.

PHILLIPS, P.C.B. (1987), "Time Series Regression with a Unit Root", *Econometrica* 55(2), 277-301.

PHILLIPS, P.C.B. (1991), "Optimal Inference in Cointegrated Systems", *Econometrica* 59(2), 283-306.

PHILLIPS, P.C.B. and HANSEN, B.E. (1990), "Statistical Inference in Instrumental Variables Regression with I(1) Processes", *Review of Economic Studies* 57(1), 99-125.

PINDYCK, R.S. and ROTEMBERG, J. (1990), "The Excess Comovement of Commodity Prices", *Economic Journal* 100(403), 1173-1189.

RAUSSER, G.C. (1985), "Macroeconomics and U.S. Agricultural Policy", in B.L. Gardner (ed.), *US Agricultural Policy: The 1985 Farm Legislation*, American Enterprise Institute for Public Policy Research, Washington, DC.

SIMS, C.A., STOCK, J.H. and WATSON, M.W. (1990), "Inference in Linear Time-Series Models with Some Unit Roots", *Econometrica* 58(1), 113-144.

STOCK, J.H. and WATSON, M.W. (1988), "Variable Trends in Economic Time Series", *Journal of Economic Perspectives* 2(3), 147-174.

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

TOMEK, W.G. and MYERS, R.J. (1993), "Empirical Analysis of Agricultural Commodity Prices: A Viewpoint", *Review of Agricultural Economics* 15(1), 181-202.

TOMEK, W.G. and ROBINSON, K.L. (1977), "Agricultural Price Analysis and Outlook", in L.R. Martin (ed.), *A Survey of Agricultural Economics Literature*, Volume 1, University of Minnesota Press, Minneapolis, 327-409.

YANG, S.R. and BRORSEN, B.W. (1992), "Nonlinear Dynamics of Daily Cash Prices", *American Journal of Agricultural Economics* 74(3), 706-715.Table 1: Large Sample Simulation Results for Supply Elasticity Estimation.