TIME SERIES MODELS FOR EXCHANGE RATE AND AGRICULTURAL PRICE FORECASTS

Girard W. Bradshaw and David Orden*

Virginia Polytechnic Institute and State University

Introduction

The Time Series Approach to Forecasting

There has been an historical dichotomy in the econometrics of forecasting literature that has admitted two approaches to the building of forecasting models. These two approaches are (1) Structural econometric/causal forecasting models, and (2) Time series models (Kennedy 1985; Granger and Newbold 1986). Recently, some authors have highlighted the compatibility and complimentarity of these two approaches, but to a large degree the two literatures remain separate.[1]

Structural econometric models are specified, as is well known, by appeal to prevailing economic theory. They consist of a set of dependent variables (the variables to be forecasted) and a set of independent variables which are used to "explain" or account for the variation in the dependent variables. These models aim to capture the structural relationships, identified from theoretical investigations, among the variables in the economy, often employing numerous overidentifying restrictions in the process. The popularity of large-scale simultaneous equation models of this type reached a peak in the 1960s and early 1970s. They continue to be widely used in commercial forecasting and, to an extent, in research. However, in the late 1970s, forecasters using these models, in particular macroeconometric models, were confounded by their failure to accurately predict simultaneous high inflation and unemployment levels (Lucas and Sargent 1979). This break-down in forecast accuracy opened the door for simpler, less costly, and more accurate alternative forecasting models. Time series models offer one such alternative.

Time series models are built on the premise that a time series has a particular recurring statistical history which can be modelled and then exploited for the purpose of forecasting. The unique statistical history is used to project forward the likely path of the time series, thus generating an extrapolative forecast. Behind the idea of time series forecasting is the eclectic view that we may not know enough about the true structure of the economy to construct a detailed structural econometric model that will forecast well (see, for example, Sims 1980).

For illustrative purposes we shall delineate two classes of time series models, those that do not allow for dynamic interactions among variables (univariate) and those that do (multivariate).

---

* Graduate Research Assistant and Assistant Professor, Department of Agricultural Economics, Virginia Polytechnic Institute and State University (Virginia Tech), Blacksburg, Virginia 24061.
The authors wish to thank Rick Ashley, Doug McTaggart, and John Robertson for helpful comments, and Ed Allen, Crops Section, ERS/USDA and David Stallings, Leader, Demand and Trade Section, ERS/USDA for providing all of the price and exchange rate data. Funding for this research was provided by the Agriculture and Rural Economics Division of the Economic Research Service, U.S. Department of Agriculture, under Cooperative Agreement 58-319S-5-00402.

[1] The Structural Econometric Modelling Time Series Analysis (SEMTSA) approach of Zellner (1979), and papers by Harvey (1981), Davidson et al. (1978), and Hendry (1978) are examples of work which emphasize the need for more synthesis. Blanchard and Watson (1986), Bernanke (1986), Sims(1986) and Fackler(1988) have addressed the identification of dynamic simultaneous equation structural models and their relationship to multivariate time series models.

Univariate time series models express the variation in a time series as a function of autoregressive terms (past own values) and moving average terms (contemporaneous and past errors):[2]

$$X_t = \phi_1 X_{t-1} + \dots + \phi_p X_{t-p} + \varepsilon_t - \theta_1 \varepsilon_{t-1} - \dots - \theta_q \varepsilon_{t-q} \tag{1.1}$$

Multivariate time series models reflect the importance of "the influence of other observable variables known or suspected to be related to the series of interest" (Kling and Bessler 1985). The multivariate time series models to be used herein will be vector autoregressions (VARs) augmented, where appropriate, by error correction terms to form vector error-correction models (ECMs), as discussed below. A VAR model does not impose *a priori* restrictions such as exogeneity or functional form as used for the identification of structural simultaneous equation models. Instead, a VAR is a reduced-form model in which interactions that are present in the data emerge on their own. If $X_t' = (X_{1t},\dots, X_{mt})$ is a vector of variables that we wish to model with a VAR, under the conditions of joint-stationarity and ergodicity (see Granger and Newbold 1986) X_t' has a vector autoregressive representation:

$$\Phi(B)X_t = \mathrm{E}_t \tag{1.2}$$

where $\Phi(B)$ is an $m \times m$ infinite matrix function in the backshift operator, and E_t is a vector of well-behaved error terms.[3] Each element of Φ(B) follows the structure:

$$\phi_{ij}(B) = \sum_{k=0}^{\infty} \phi_{ij,k} B^k$$

This infinite AR structure is then *approximated* by a finite autoregression for empirical estimation. Choosing the lag length in a VAR is an important issue in empirical work, and will also be addressed in our paper.

In practice, the question of whether a traditional structural econometric model or a time series model is better for a particular forecasting project turns on the validity of the prior information that we have.[4] If a particular economic theory is "true," it would be unwise not to use that available information. [5] A univariate ARIMA model incorporates no prior *economic* information thus would be a poor choice as a forecasting model in the face of a structural econometric model imposing valid identification restrictions (Prothero and Wallis 1976). In fact, however, ARIMA models frequently out-perform structural econometric models in forecasting.[6]

The multivariate time series approach asserts that the truth lies somewhere between the traditional simultaneous equation approach and atheoretical univariate time series models. Through economic theory, variables can be identified that have a high prior probability of having an im-

---

2 Hence the name ARIMA models. The "I" in ARIMA designates that the time series is modelled in differences of the original series. This is done to induce stationarity, an important topic to be covered in detail below.

3 The backshift operator, also called the lag operator, performs the following operation, $B^k X_t = X_{t-k}$, on either a single random variable or a vector of variables.

4 One way to test the general credibility of identification restrictions would be to conduct a forecasting experiment where the structural model competes with various specifications of multivariate time series models. Zellner (1982) argues that structural models that have excellent in-sample fit must still prove their ability to forecast well in order for them to be useful contributions to economic science.

5 Rick Ashley points out that if the forecasts of such explanatory variables are poor enough, it may be unwise to include these variables in the forecasting model even if the theory supporting their relevence is valid.

6 An example was presented in Nelson (1972) where it was demonstrated that a univariate ARIMA model could out-forecast the Federal Reserve-MIT-Penn quarterly model of the U.S. economy. See also Cooper(1972).

portant effect on the variables to be forecasted, though we are not quite certain how these interrelationships are manifest in, say, particular functional forms or exclusion and exogeneity restrictions. VAR models can be viewed as quasi-time series models in that they are specified using some *a priori* information from economic theory to guide the selection of the variables to be included in a specification in order to produce a model with superior forecasting accuracy.

Multivariate Time Series Analysis in Agricultural Economics

Much of the recent agricultural economics literature that is concerned with use of time series models reflects an interest in relationships among various economic time series. Bessler (1984a, 1984b) discusses the methodological aspects of fitting VARs and examines dynamic economic relationships in the hog market and the relationship of relative prices and money in Brazil, respectively.

In this study, we focus on the role of the exchange rate in explaining variation in agricultural commodity prices. Orden (1986) investigated the dynamic effects of macroeconomic shocks on U.S. agriculture using VAR models. He found evidence that movements in the real exchange rate have substantial impacts on agricultural exports and real prices received by farmers. Orden's study used policy analysis techniques developed by Sims (1980, 1986), and employed Sims' lag selection criterion. Thornton and Batten (1985) found that, for money-income relationships, different lag structures can change the outcome of exogeneity tests. Therefore, we might place value on the consideration of alternative test procedures for lag length selection and such procedures will be used herein.

From a purely forecasting perspective, Bessler and Babula (1987) found the real exchange rate to have little impact on improving the accuracy of forecasts of wheat exports. The real exchange rate did have a notable effect on increasing accuracy of forecasts of real wheat prices.[7] The impact of real exchange rates on forecasts of real wheat, corn, and soybean prices will be a primary focus of this paper.[8]

As Bessler and Babula's results suggest, multivariate time series models are not always able to out-perform univariate models in out-of-sample forecasting ability (Litterman 1984, Brandt and Bessler 1984). This is puzzling as one might expect that a multivariate model should forecast at least as well as a univariate model because it theoretically encompasses that univariate model.

The failure of multivariate models to out-forecast univariate models is likely the result of ignoring very important characteristics of time series *data*. For example, Litterman has argued that aggregate economic data suffers from a low *signal-to-noise ratio* meaning that the useful (systematic) variation in the time series is obscured by purely random fluctuations. This random noise overpowers the useful signal, i.e. the variation that can be used to explain variations in another variable in a multivariate model. Litterman(1986) explains that the parameters in a VAR will likely fit both the useful systematic variation as well as the random variation, resulting in an overparameterized model. The random variation, however, is not useful for forecasting. The task of the forecaster, then, is to devise a way to filter out the random noise to reveal the variation that is systematic. The multivariate forecasting problem is thus a trade-off between oversimplification and overparameterization.

---

[7] Bessler and Babula use a decomposition of forecast error variance to isolate the effect of the exchange rate on the wheat price.

[8] More generally, our study is pursuing the investigation of the effects of real exchange rates (and other macroeconomic variables) on price and export-quantity forecasts. Bessler and Babula's results are anomalous as we would expect that a change in a real commodity price would be associated with a change in exports of that commodity.

Model Specification Issues

Litterman (1976) proposes a way of filtering the noise from the signal through the use of Bayesian priors.[9] The point of using Bayesian priors is to attack the overparameterization problem inherent in unrestricted VARs where the modeller can quickly run out of degrees of freedom even in a moderately large sample. Litterman suggests that it is reasonable to expect that coefficients on long lags are more likely to be zero than coefficients on shorter lags. By specifying Normal prior distributions about zero with smaller standard deviations for the coefficients as lag length increases, long lags are allowed to enter the equations at the margin only if there is strong evidence for such relationships in the data (see Litterman 1984, and Doan, Litterman, and Sims 1983 for discussions of Bayesian VAR models for forecasting).

Bessler (1986) addresses the use in forecasting of a nonsymmetric (i.e., the prior on cross variable effects is different from the prior on own lags) random walk Bayesian prior on the coefficients in a VAR model for the U.S. hog market.[10] He finds the VAR with this prior out-performs a univariate autoregression, an unrestricted VAR, and a Bayesian VAR with a symmetric prior. Kling and Bessler (1985) also found that Litterman's Bayesian VAR forecasted very well for macroeconomic data.

An alternative to the Bayesian procedure for obtaining a more parsimonious VAR representation is to use a multivariate statistical decision criterion for the choice of lag length. Lutkepohl (1985) has investigated the use of 12 such statistical decision rules in a monte carlo simulation. He found several of them to be quite accurate in choosing the correct lag length. Lutkepohl's results indicate that the Bayesian Information Criterion of Schwartz (1978) and the criterion of Hannan and Quinn (1979) are the most accurate given a moderately large sample size. These decision criteria are applicable in both a univariate and a multivariate lag selection problem, the multivariate being the general case.

Hsiao (1979) has developed a procedure to help overcome the overparameterization problem, *as well as allow for more realistic differing lag structures in each equation of the system.* Hsaio's procedure uses the Final Prediction Error criterion of Akaike (1971), though any of a number of available statistical rules could be used as the underlying decision criterion in his procedure (see Judge et al. 1985 p. 675). Though Hsiao's procedure is not without fault, it is a useful procedure for modelling restricted VARs.[11] Fewer parameters in a VAR allows the remaining parameters to be estimated with more degrees of freedom, hence more accurately. Hsiao's method for reducing the number of parameters to be estimated in a VAR is closer to the time series philosophy of allowing the data to determine the model specification than the Bayesian procedure which forces the modeller to choose a prior to impose.[12]

Another reason given for the poor performance of VAR representations of economic time series is the lack of attention given to data issues inherent in time series econometrics. For example, if the time series to be modelled contains a deterministic trend or a unit root, or if the variance of the series is not constant throughout, the series is said to be *non-stationary*. Many economic time series, especially macroeconomic time series, are in fact non-stationary in their levels (Wasserfallen 1986; Nelson and Plosser 1982). Because multiple time series theory relys on stationarity for its validity, the modelling of nonstationary series *as if they were stationary* will produce undesireable results, one of which can be poor forecasting performance. Dickey et al. (1986) show that, for univariate models, forecasts from a nonstationary model do not decay to the sample mean as the forecast horizon increases, and the forecast standard errors will diverge to $+\infty$ instead of converging

---

[9] A prior is an informed belief that the modeller brings to the modelling exercise. The priors appear in the form of probability distributions on the coefficients.

[10] The random walk prior is justified for many macroeconomic and financial variables (e.g., Nelson and Plosser 1982).

[11] Webb (1985) applies the Akaike Information Criterion (AIC) within his own procedure to the choice of lag length in a VAR. He notes a consistent improvement in the forecasting accuracy of his specification using the AIC over an unrestricted VAR.

[12] Doan, Litterman, and Sims (1983) propose the quasi-Bayesian approach of using the data to select an optimal prior.

to the series' standard deviation. This divergence of the forecast-error variance should be reflected in the forecast variance component of the RMSE of the forecasts from nonstationary models. At longer horizons, then, increasingly poorer forecasts from nonstationary models should be observed relative to stationary models.

Another undesireable result is the appearance of spurious relationships when nonstationary variables are regressed on each other (Granger and Newbold 1974, 1986). This is a problem that is now gaining long overdue recognition. Phillips (1986) has demonstrated, using large-sample theory, that when non-stationary series are regressed upon one another the t-tests of significance are biased toward rejecting the null hypothesis of no relationship. In related work, Phillips and Durlauf (1986) demonstrate that the asymptotic theory for non-stationary multiple time series departs significantly from classical theory. In the case of non-stationary series, the asymptotic distribution of the usual test criteria is nuisance parameter dependent, meaning that classical test statistic distributions are no longer applicable.

A third consideration is that when sets of random variables are being modelled, as in a VAR, attention must be paid not only to the stationarity of individual variables but to possible equilibrium relationships among the variables. These equilibrium relationships are manifest when two or more non-stationary variables have a linear combination that is stationary. Such economic variables are then said to be *co-integrated.* The consequences of ignoring co-integration include a loss in forecasting efficiency as important prior information will not be included in the model specification. In fact, when co-integration is present, the usual VAR representation is inappropriate (Engle and Granger 1987).

The importance of these data issues makes it necessary to have a strategy for investigating the time series properties of the variables to be modelled *before a particular model is chosen.* On the basis of these time series properties, appropriate forecasting models, in our case possibly bivariate commodity price-exchange rate models, can be specified.

Stationarity and Cointegration

The Univariate Case

It is well known in the time series literature that the time series being modelled must be stationary for there to be a linear model representation (see Wold 1954, Judge et al. 1985). Stationarity requires that the mean and variance of the series be finite and time invariant and that the covariance between any two values of the process depend solely on the distance between these values in time and not on time itself.

Unfortunately, the levels of many economic time series appear to be nonstationary (Nelson and Plosser 1982, Nelson and Kang 1984). Hence, in order to apply linear models such as ARMA models, a time series must be tested for and possibly transformed to stationarity. A useful starting point is to examine a time plot of the raw data series. If the series exhibits fluctuations that are more violent for a particular segment of the series than for others, the series very likely is variance nonstationary, i.e., there is not a constant variance throughout the series. The most common method for inducing variance stationarity in a series is to take the natural logarithm of each observation. This transformation will reduce the swings of the levels which constitute the variance nonstationarity and often yields a series that is a good approximation to one having constant variance.

A more insidious form of nonstationarity, however, is nonstationarity in the mean of a series. In this case, the series shows no propensity to return to, or move around, a particular fixed level. When a series has this lack of affinity for a mean, and the movement seems to be in a particular direction, the series is often said to exhibit a "trend". In this paper, "trend" will be reserved for a *deterministic* functional dependence on time. For example, consider a series that has two parts, a deterministic linear trend and a residual representing the stationary component which includes all of the interesting variation that we wish to model:

$$X_t = \alpha + \beta t + \varepsilon_t \qquad (2.1)$$

Often, (2.1) is estimated as a linear regression model, and the residuals are then treated as a stationary series that has well-defined variance, covariances, and autocorrelations (Nelson 1984). This would mean that ε_t could be modelled as an ARMA process after the trend was removed from X_t. The function need not be linear, however, which leads to the more general representation:

$$X_t = f(t) + \varepsilon_t \tag{2.2}$$

The relation in (2.2) is called a trend stationary process (TSP); X_t is stationary *around the trend function $f(t)$*. It is important to note that this is just one hypothesis concerning the manifestation of nonstationarity, and indeed there are problems with this particular hypothesis. Even if we could know *a priori* that the variable X_t is a TSP, there is little chance that the actual functional form could ever be accurately specified. If the nonstationarity is not correctly modelled, the residuals in (2.2) will not be stationary. In addition, over the course of a time series, we may observe local upward trends followed by local downward trends. Thus, a global OLS trend would not be an accurate representation of the nonstationarity.

An alternative hypothesis about the way nonstationarity in mean arises was introduced by Box and Jenkins (1970). They view nonstationarity not as a manifestation of deterministic functions of time, but as the *accumulation of random shocks*. In this case, the first differences of the series are stationary. This kind of process takes the form:

$$X_t = X_{t-1} + D + \varepsilon_t \tag{2.3}$$

Where ε_t is a stationary series with zero mean and constant finite variance and D is the fixed mean of the first differences, often called the drift parameter.[13] The level of the series at any given time t is equal to the previous level of the series, plus the drift, plus the random shock. The series is cumulative, or additive, in its level. This additivity exhibits itself as an apparent "trend". Equation (2.3) is said to belong to the difference stationary class of processes (DSP).

DSPs are also called "integrated" processes, the word "integrated" reflecting the additive nature of the series. The following three definitions will be useful in following sections (from Granger 1986, p.216):

*Definition 2.1a.*If a time series Z_t needs no differencing to become stationary, it is called *integrated of order zero* which is denoted $Z_t \sim I(0)$.

*Definition 2.1b.*If a time series Z_t must be differenced d times to become I(0), it is called *integrated of order d* which is denoted $Z_t \sim I(d)$.

*Definition 2.1c.*Let Δ^b represent b applications of the difference operator. If $Z_t \sim I(d)$ then the b^{th} differenced series is $\Delta^b Z_t \sim I(d-b)$.

Dickey(1975), Fuller (1976), and Dickey and Fuller(1979, 1981) have developed a series of tests (henceforth DF tests) for discriminating between the hypothisis that a series is a TSP and the hypothesis that it is a DSP. Their tests only entertain a DSP that is integrated of order one. The procedure is to perform OLS on the model:

$$X_t = \alpha + \beta t + \rho X_{t-1} + \varepsilon_t \tag{2.4}$$

The null hypothesis for the first test, τ_μ, is that $\rho = 1$, or that x_t contains a unit root and is nonstationary, with an alternative model that the series is generated by a stationary autoregression ($\rho < 1$) with drift. The null hypothesis for the second test, τ_τ, is again that $\rho = 1$, with drift in the null model and an alternative model that the series is generated by stationary autoregression around a linear time trend (drift plus a time parameter). Fuller (1976 p.373) has tabulated critical values

[13] The simplest member of this class of processes is the random walk where ε_t would be a white noise process (zero mean, finite variance and zero covariance between any two values separated in time), and the drift would be zero.

for τ_μ, and τ_τ, both of which are "t-ratios", $(\rho - 1)/\sigma_\rho$, that follow *nonstandard distributions.* The rejection regions are given by small values of τ_μ or τ_τ.

Dickey and Fuller also describe two likelihood ratio tests for the joint null hypothesis of a simple random walk. In the first of these tests, Φ_1, the null hypothesis is $(\alpha, \rho) = (0,1)$ in a model that is assumed not to include a time parameter. In the second test, Φ_2, the null hypothesis is $(\alpha, \beta, \rho) = (0, 0, 1)$ in a model that may have a linear time trend. Finally, Dickey and Fuller describe a likelihood ratio test for the joint null hypothesis of a random walk with drift $(\alpha, \beta, \rho) = (\alpha, 0, 1)$ in a model that again includes a time parameter. The rejection regions are for large values of the test statistics, and critical values are found in Dickey and Fuller (1981, p. 1069).

A question that arises with the DF test is whether it is appropriate to model X_t as AR(1) or a random walk, as the error, ε_t, in (2.4) may not be empirical white noise. For example, if there is evidence of moving average behavior, a higher order autoregression may be needed to approximate the dynamics of the X_t process. Consequently, a more general model is often fit. This results in an augmented Dickey-Fuller test (ADF) based on the model:

$$X_t = \alpha + \beta t + \rho X_{t-1} + \sum_{i=1}^{p} \phi_i \Delta X_{t-i} + \varepsilon_t \tag{2.5}$$

where lags of ΔX_t are added until ε_t is white noise. The hypotheses to be tested about the properties of the series are the same for this specification as for model (2.4). Fuller (1976) and Dickey and Fuller (1981) show that their tests also apply to these higher order autoregressions.

The Multivariate Case

Linear vector time series models, such as VARs, are only applicable to stationary vector time series (Judge et al. 1985). A vector time series $X_t' = (X_{1t}, X_{2t}, ..., X_{mt})$, is stationary when each series is individually stationary in mean and variance. In addition, all covariances, whether intraseries (an autocovariance) or interseries (across every pair of the m variables in the vector process) must be independent of t and depend only on the *time displacement* between observations.

As discussed above, there is reason to believe that many economic time series are nonstationary. Hence, the practitioner is faced with the problem of how to apply the theory of vector linear models to nonstationary time series. Extrapolating from the univariate case, a practical solution would seem to be to examine the univariate time series properties of each series in the vector to be modelled and use appropriate transformations to reduce each individual series to stationarity. Such an approach is advocated in some articles on VAR modelling (e.g., Hsaio 1979). Hsaio logged and differenced each series in the bivariate money-income relationship before proceeding with specifying a VAR. However, differencing of the individual series has been criticized by others (e.g. Taio and Box 1981; Lutkepohl 1982). The difficulty noted is that while each individual series may be nonstationary, "for vector time series, linear combinations of the components of $[X_t']$ may often be stationary, and simultaneous differencing of all series can lead to unnecessary complications in model fitting" (Taio and Box 1981 p.804). This phenomenon of linear combinations of nonstationary series being stationary has been termed co-integration (Granger 1980; Granger and Weiss 1983; Engle and Granger 1986; Engle and Yoo 1987). Essentially, if there exist linear combinations of the individual nonstationary series that are stationary, differencing each series individually will result in a system that is overdifferenced. If this is the case, the system will no longer have a multivariate linear time series representation with an invertible moving average. Intuitively, if a system is co-integrated, estimating a model in differences ignores the equilibrium relationships among the nonstationary variables that contain important information. Modelling the co-integration restrictions, then, should help a model produce forecasts that are more accurate than a model in which the restrictions are ignored (Engle and Yoo 1987).

For integrated processes of order greater than zero (i.e. nonstationary series) the use of statistical techniques which assume stationarity can give incorrect results in the multivariate case as well as in the univariate case. To illustrate the issues involved, consider the *static* regression:

$$Y_t = \alpha + \beta' X_t + \varepsilon_t \tag{2.6}$$

where β' is a vector of coefficients and X_t is a vector of regressors. Suppose that Y_t and the X_t's are each $\sim I(1)$. Rearranging (2.6):

$$\varepsilon_t = Y_t - \alpha - \beta' X_t \tag{2.7}$$

In general the linear combination in (2.7) will yield $\varepsilon_t \sim I(1)$ because ε_t is a linear combination of I(1) series. Hence the residuals will be nonstationary.[14] Phillips (1986) and Phillips and Durlauf (1986) investigate the effect of using integrated processes in static multivariate regressions such as (2.6), and in multiple time series regressions such as VARs. They conclude, using large sample asymptotics, that the distributions associated with the usual inferential statistics *do not follow the same distributions that they would under stationarity*. For the case of static multiple regression, Phillips proves that the coefficients of the regression do not converge in probability to constants as the sample size goes to infinity, as is the case when the variables are stationary; that is, the variables have no limiting distribution (Phillips 1986a; Banerjee et al. 1986). Phillips also shows that the distributions of the t-ratios diverge as the sample size goes to infinity. This means that no asymptotically correct critical values exist for conventional significance tests. For critical values from conventional asymptotics, the rejection rate for the null hypothesis will increase with sample size (Phillips 1986, p.318). These results confirm the monte carlo evidence in Granger and Newbold (1974, 1986). The bias toward wrongly rejecting is their concept of spurious regressions.[15] Granger and Newbold illustrate the problem by regressing independent random walks on one another. They note that using the usual t-test (designed under the maintained hypothesis that the variables involved are stationary) at five-percent level of significance will, on average, lead to wrongly rejecting the null hypothesis *three-fourths* of the time. Where the number of independent variables is greater than one, Granger and Newbold (1974) report a bias in F-tests toward wrongly rejecting the joint null that all coefficients are zero from 76 percent to 96 percent of the time, with the rejection rate increasing with the number of included variables.[16]

Dynamic multivariate time series regressions with integrated processes, as opposed to static regressions, are investigated by Phillips and Durlauf (1986). They find that OLS does provide *consistent* estimates of the regression coefficients in this case. However, these estimates are not asymptotically normally distributed. An important result is that the limiting covariance matrices for the estimated coefficients have distributions that depend on the number of variables in the system. These nuisance parameter dependencies invalidate the usual classical significance tests. New statistical tests must therefore be devised which are free of nuisance parameter dependencies.

Thus, nonstationarity in a multivariate regression, such as a VAR, can cause serious problems for statistical inference. In order to avoid being fooled by spurious relationships, making invalid conclusions based on the application of the wrong asymptotic theory, or generating poor forecasts, one must insure that the series involved are stationary.

Co-integration and its implications

In a static multivariate regression, finding an ARIMA representation for the residuals and/or differencing the included variables should eliminate the occurance of invalid conclusions on the basis of classical inferential techniques. For time series regressions with integrated processes, consideration of co-integration plays a vital role in deciding what to do about nonstationarity.

---

[14] For the case where $Y_t \sim I(1)$ and $X_t \sim I(0)$ the residual in (2.7) will be I(1) as well.

[15] Yule (1926) was the first to formally investigate this phenomenon, often called "spurious" or "nonsense" correlations. Yule examined the correlations between unrelated series. When the series were stationary, no correlation was observed, as expected. For I(1) series, the correlation distribution indicated a high degree of linear association, and for I(2) series the most often encountered correlations were ± 1.

[16] To underscore their results, Granger and Newbold work with statistically independent variables. However, the distributional results proved by Phillips(1986) also apply to correlated time series. The crucial results are that the coefficients do not converge to constants and that the distributions of the test statistics diverge as the sample size increases to infinity.

Consider again a regression relationship such as (2.6) in which each variable is I(1) and we assume, for illustration, that there is only one regressor. We would expect that the residuals in this series would be I(1) as they are a linear combination of I(1) variables. However, in the special case where there exists a unique constant, say, γ, such that the two I(1) series have a unique linear combination:

$$z_t = Y_t - \alpha - \gamma X_t \tag{2.8}$$

that is stationary (more precisely $z_t \sim I(0)$). In this case, X_t and Y_t are co-integrated, with co-integrating constant γ.[17] In the two variable case, γ will be unique. For vectors of more than two time series there may be multiple vectors of co-integrating constants and co-integrating relationships among the variables may not be unique (see Engle and Granger 1987).

The intuition behind co-integration is that some economic variables move together through time, hence the co-integration relationship can be thought of as an equilibrium relationship. It says that two (or more) variables which have unbounded variance and no constant mean, have a linear combination that has finite variance and a constant mean. Consequently, the variable z_t in (2.8) can be said to measure departures from long-run equilibrium between the two series. Granger (1986) cites prices and wages, the money supply and prices, government income and expenditure (perhaps only at the state or local level), and the imports and exports of a country as pairs of variables that may be cointegrated.

It makes sense that we should test economic variables for cointegration relationships and then make use of the resulting information in model specification. Engle and Granger (1987) have proposed tests of the null hypothesis of no co-integration against the alternative of co-integration. The tests are based on the residual in (2.8) being I(0) if the series are co-integrated.[18] That we are interested in whether the errors are I(0) or I(1) suggests DF and ADF tests be applied to the residuals obtained by estimating this "co-integrating regression".

However, the co-integrating regression yields both an estimate of the co-integrating parameter, $\hat{\gamma}$, and the residual series, $\hat{z}_t$.[19] Since $\hat{z}_t$ can be obtained only after first obtaining $\hat{\gamma}$, it has a dependency on the estimate of the co-integrating parameter. In the unit root test on the $\hat{z}_t$ series from the co-integrating regression, the large sample behavior of the "t-statistic" has nuisance parameter dependencies which stem from this dependency. These are manifest as a dependency of the "t-statistic" on the number of variables in the co-integrating regression (Engle and Yoo 1987). The critical values in Fuller (1976) and in Dickey and Fuller (1981) used for the usual DF and ADF unit root tests do not apply for the co-integration test because they do not account for these nuisance parameter dependencies. New critical values, dependent on the number of variables in the vector time series, are provided in tables in Engle and Yoo (1987).[20]

If the test for co-integration is unable to reject the null hypothesis of no cointegration the appropriate model is one in first differences, as each of the variables are I(1) and they have no linear combination(s) that are I(0). If co-integration is found, a model is needed that includes this infor-

---

[17] In general, for any pair of series X_{1t} and X_{2t} both $\sim I(d)$, if there exists a linear combination (2.7) such that $z_t \sim I(d-b)$ with $b > 0$, the pair are co-integrated of order $d-b$, denoted $(X_{1t}, X_{2t}) \sim CI(d,b)$.

[18] Note that each individual series must have an order of integration equal to the other's for co-integration to make sense. To insure this, one should use tests such as the DF and ADF as well as ACF and partial autocorrelation function plots of the raw and differenced series to determine the order of integration of each individual series. If such a preliminary examination strongly indicates that the series have differing orders of integration, a formal co-integration test is unnecessary.

[19] Stock(1987) has shown that if Y_t and X_t are cointegrated, OLS estimates of γ are highly efficient and super consistent; that is, as $T \rightarrow \infty$ $\hat{\gamma}$ will converge to its true value twice as rapidly as would be the case for a usual OLS parameter estimate in a similar, stationary, regression.

[20] For the cointegration test, we can ignore the trend functions and the hypothesis becomes I(1) vs. I(0) with the relevent statistic being the t-ratio on the parameter ρ.

mation. Granger (1982) and Engle and Granger (1987) prove that a bivariate co-integrated system has an *error-correction model* representation:[21]

$$\Delta X_t = -\xi_1 z_{t-1} + \alpha(B)\Delta X_t + \beta(B)\Delta Y_t + \varepsilon_{1t} \tag{2.9a}$$

$$\Delta Y_t = -\xi_2 z_{t-1} + \gamma(B)\Delta X_t + \phi(B)\Delta Y_t + \varepsilon_{2t} \tag{2.9b}$$

with $|\xi_1| + |\xi_2| \neq 0$. The co-integration is captured in (2.9) uniquely through the $\hat{z}_{t-1}$ term which is obtained from the cointegrating regression. This representation captures the co-integration in terms of the levels of the co-integrated variables (Engle and Granger 1987). The levels enter the equation as last period's departure from long-run equilibrium. Specifying a VAR in differences, if the variables are co-integrated, can be thought of as a specification error because the error correcting terms ($-\xi_1 z_{t-1}$ and $-\xi_2 z_{t-1}$) are incorrectly excluded from the equations.

The model selection process, then, is a process that, to a large degree, depends on the information we can extract from the data concerning its time series properties. In the following section, we consider the data analysis techniques discussed above to help in specifying univariate models for forecasting agricultural commodity prices and bivariate models with prices and exchange rates. The models suggested by the data analysis will then be compared to various popular time series specifications to see whether these techniques make a difference by providing better forecasts.

Forecasting Models for Agricultural Prices

Data Description

The price data used herein are average monthly cash prices of No.1 Hard Red Winter Wheat at Kansas City, No.2 Yellow Corn at Chicago, and No.1 Yellow Soybeans at the Illinois Processor, deflated by the U.S. CPI. The CPI was taken from various issues of the *Survey of Current Business*. The price data were obtained from the Crops Section of the USDA, and run from January, 1974 through August, 1987. A post-1973 sample period was chosen so that exchange rate effects would be observed only over the period of flexible market-determined rates.

The exchange rate data includes the crop-specific real trade-weighted exchange rates for wheat, corn and soybean exports calculated by the Demand and Trade Section of the USDA. The overall index is calculated as follows: The weights for the indices are average value shares of U.S. commercial exports from 1976-78. The current real exchange rate for each country is computed by taking the ratio of the same period CPI in the U.S. to that of the country in question and multiplying by the period average spot rate. The percent change from the base value is then multiplied by the weight. These weighted changes are then summed into a total which is the real index.

Analysis of the Data and Tests for Unit Roots

After an examination of the time plots we conclude that each series should be expressed in natural logarithms to compensate for an apparently nonstationary variance.[22] The time plots of the individual series also indicated possible nonstationarity in mean. Estimates of the autocorrelations and partial autocorrelations of each series provide a useful starting point for evaluating this nonstationarity and are reported in table 1.[23] In each case, the autocorrelations are large at low lags

[21] Granger(1986) and Engle and Granger(1987) also discuss the error correction models for vectors of more than two time series.

[22] Henceforth when a price variable is referred to as "wheat price" it should be understood that it is the logarithm of the real wheat price, and similarly for other prices and the exchange rates.

[23] The autocorrelations and partial autocorrelations found in tables 1 and 2 are calculated over the entire sample period 1974:1 to 1987:8. The forecasting models will be calculated over a smaller sample period, hence it might be argued that we should inspect these values instead. We use the full sample so that we can assimilate the most information possible. All the results concerning stationarity should be (and are) robust to this slight change in sample period.

Table 1. Estimated autocorrelations and partial autocorrelations on logged data, lags 1-24, 1974:1-1987:8

	Wheat Price		Corn Price		Soybean Price		Wheat Ex. Rate		Corn Ex. Rate		Soybean Ex. Rate	
Lag	ACF	PACF	ACF	PACF	ACF	PACF	ACF	PACF	ACF	PACF	ACF	PACF
1	.95	.95	.96	.96	.97	.97	.99	.99	.99	.99	.99	.99
2	.90	-.14	.91	-.13	.92	-.22	.98	-.20	.97	-.24	.97	-.24
3	.85	.10	.87	.04	.87	.00	.96	-.13	.96	-.07	.96	-.06
4	.82	.14	.83	.07	.83	.12	.95	-.02	.94	-.09	.94	-.10
5	.80	.03	.80	-.04	.79	.01	.93	-.02	.92	-.00	.91	-.01
6	.77	-.04	.76	-.05	.76	.01	.92	-.06	.90	-.05	.89	-.07
7	.74	.01	.71	-.07	.73	-.06	.90	.04	.87	-.02	.87	-.03
8	.71	.03	.67	.01	.70	.03	.88	-.01	.85	-.04	.84	-.03
9	.69	-.02	.64	.06	.67	.00	.86	-.06	.82	-.04	.82	-.05
10	.66	-.10	.60	-.03	.63	-.11	.85	.09	.80	.01	.79	.01
11	.62	-.03	.57	-.07	.59	.01	.83	-.01	.77	-.09	.76	-.08
12	.58	-.05	.52	-.07	.56	-.01	.81	.10	.75	-.06	.73	-.02
13	.54	-.04	.48	.00	.53	.02	.80	-.07	.72	-.07	.70	-.05
14	.51	-.03	.45	.03	.50	.10	.78	-.11	.69	-.05	.67	-.04
15	.48	.05	.42	.11	.49	.02	.77	-.07	.66	-.08	.64	-.10
16	.46	.07	.40	.02	.48	.07	.75	.04	.62	-.11	.61	-.11
17	.44	.05	.38	.03	.47	.06	.73	-.06	.58	-.02	.57	-.04
18	.43	.03	.36	.00	.47	.06	.71	.02	.55	.12	.53	.11
19	.41	-.04	.35	-.02	.47	-.02	.69	-.03	.51	-.06	.50	-.05
20	.38	-.08	.33	-.05	.46	-.00	.67	-.13	.48	.15	.47	.15
21	.35	-.02	.31	-.00	.45	-.00	.65	.02	.45	.03	.44	.03
22	.32	-.03	.29	.01	.44	-.05	.63	-.09	.42	-.07	.41	-.09
23	.30	.00	.27	-.05	.43	-.04	.60	-.08	.39	-.02	.37	-.04
24	.28	-.01	.25	-.02	.42	.04	.58	-.10	.37	-.01	.34	-.01

Estimates are based on full sample from 1974:1 to 1987:8. Standard errors are approximately 1.96/SQRT (167) = $\pm$ 0.15; where 167 is sample size.

and decay quite slowly. Autocorrelations at lag 24 are all significant. By comparison, the autocorrelations and partial autocorrelations calculated for the first differences of the series (reported in table 2), decay quickly to insignificance by at most the fifth lag. We therefore can be reasonably confident that none of the series contains more than one unit root.

The autocorrelations are also useful for detecting seasonality. Large, significant, autocorrelations will show up at the seasonal lags (called seasonal spikes) if seasonal patterns are indicated. There is no such behavior indicated in table 1, however, nonstationarity can often mask seasonal spikes in the autocorrelations. An inspection of the autocorrelations of the first differences of each series indicates very slight evidence for seasonality in the case of the corn price only. The seasonal spike at lag 12, however, is barely significant. In addition, there is no evidence of a seasonal spike at the second seasonal lag, 24, which leads us to conclude that there is not sufficient evidence to warrant a seasonal transformation.

As the autocorrelations of the series to be modelled indicate possible nonstationarity, DF and ADF tests for unit roots are conducted as outlined above. In every case, the simple DF regression, which contains no lags of the dependent variable, shows signs of serial correlation in the residuals according to the Ljung-Box Q-Statistic (see Ljung and Box 1978).[24] To remedy this lack of fit, higher order autoregressions were estimated. For the wheat price and the wheat exchange rate a two-lag model proves sufficient to eliminate serial correlation of the errors. In the other cases, a one-lag model appears sufficient.

The results of unit root tests are summarized in table 3. There is convincing evidence of the presence of unit roots in the levels of each of the variables we consider. The evidence for the presence of unit roots is most conclusive for the wheat price and each of the exchange rates. In these four cases, none of the test statistics is rejected, and we therefore conclude that each contains a unit root.

For the corn price, the value of the test statistic, τ_μ, suggests that we do not reject the null hypothesis of a unit root. The statistic, τ_τ, provides further evidence of a unit root, though at a smaller level of confidence. The statistic Φ_1 for the joint-null hypothesis that $(\alpha, \rho) = (0,1)$ is not rejected and the statistic Φ_2 for the joint-null $(\alpha, \beta, \rho) = (0,0,1)$ is also not rejected. Lastly, the joint test for the null $(\alpha, \beta, \rho) = (\alpha,0,1)$ is rejected at the 10-percent level, though this does not provide strong counter-evidence. We therefore conclude that the corn price is nonstationary as a result of the presence of a unit root.

The unit root tests on the soybean price are less conclusive. The first "t-ratio" test for the null of a unit root, τ_μ, is rejected at the 10-percent level. The test statistic, τ_τ , is rejected at the 1-percent level. The joint tests, however, are inconclusive. The statistic, Φ_1, with a null of $(\alpha, \rho) = (0,1)$ is not rejected. The test of $(\alpha, \beta, \rho) = (0,0,1)$, Φ_2, is rejected at the 5-percent level. Finally, the test of $(\alpha, \beta, \rho) = (\alpha, 0,1)$, Φ_3, is rejected at the 1-percent level. It is critical to remember that a rejection of the joint-null which includes a restriction that $\rho = 1$ does not necessarily imply that we are rejecting that particular restriction. To sort out the conflicting evidence concerning soybean price we consider comparisons of the empirical power of the unit root tests. Dickey and Fuller(1981) rank the tests, denoted by their corresponding statistics, on the basis of their power as follows: $\Phi_1 > \Phi_3 > \Phi_2$ and $\Phi_3 > \tau_\tau$. For soybean price, the more powerful test provides an indication that we should not reject the null of a unit root, and consequently that will be our conclusion.[25]

The unit root tests imply that for each of the series we considered the nonstationary behavior is a result of the presence of a unit root. However, there is ambiguity as to the presence of a drift parameters in the specifications. We therefore test the significance of the drift term in an autoregression of each differenced series.[26] For the wheat price there is marginal evidence for inclu-

---

[24] The null hypothesis is no serial correlation, hence model adequacy is rejected for large values of this statistic.

[25] Schwert (1987) argues that specifying the correct ARMA structure, and not just an autoregressive approximation, is necessary to avoid possibly wrongly rejecting the null hypothesis of a unit root.

[26] Stock and Watson (1987) also suggest checking the difference specifications for quadratic trends as well,

Table 2. Estimated autocorrelations and partial autocorrelations on first differences of logged data, lags 1-24, 1974:1 - 1987:8

	Wheat Price		Corn Price		Soybean Price		Wheat Ex. Rate		Corn Ex. Rate		Soybean Ex. Rate	
Lag	ACF	PACF	ACF	PACF	ACF	PACF	ACF	PACF	ACF	PACF	ACF	PACF
2	-.09	-.27	.07	-.13	.05	-.10	.20	.14	.12	.05	.11	.04
3	-.17	-.04	-.20	-.22	-.18	-.19	.13	.06	.12	.09	.13	.10
4	-.07	.00	-.18	.01	-.18	-.05	.06	-.01	.06	.00	.07	.01
5	-.05	-.08	-.04	.06	-.12	-.04	-.03	-.07	.08	.05	.09	.06
6	-.11	-.11	-.06	-.15	-.11	-.11	-.01	-.01	.04	-.00	.06	.01
7	-.19	-.14	-.11	-.11	-.09	-.06	-.01	.01	.05	.03	.04	.01
8	-.06	.04	-.21	-.13	-.02	-.01	-.01	.00	.09	.07	.09	.07
9	.05	-.02	-.12	.01	.08	.05	.01	.02	.02	-.04	.03	-.03
10	.15	.10	.07	.09	.06	-.05	-.07	-.08	.07	.06	.07	.06
11	.13	.03	.20	.07	.03	-.02	.07	.10	.09	.04	.06	.01
12	.13	.11	.22	.06	-.01	-.01	.09	.08	.10	.06	.07	.05
13	.04	-.03	.07	-.03	-.11	-.12	.10	.06	.11	.06	.09	.04
14	.01	.04	-.09	-.09	-.14	-.09	.06	.00	.14	.08	.13	.10
15	-.07	-.06	-.16	-.06	-.14	-.07	.04	-.03	.13	.05	.12	.05
16	-.07	.03	-.06	.05	-.09	-.06	.03	.01	.08	.00	.10	.03
17	-.04	.01	-.04	-.08	-.10	-.14	-.06	-.07	-.10	-.17	-.08	-.17
18	.03	.06	-.04	-.02	-.03	-.05	.02	.05	.03	.06	.02	.04
19	.12	.15	-.04	.05	-.01	-.08	-.01	.00	-.12	-.18	-.13	-.19
20	.07	-.04	-.03	.01	.10	.03	-.02	-.03	-.09	-.03	-.10	-.03
21	-.06	-.06	.02	-.02	.13	.01	.13	.17	.04	.06	.05	.08
22	-.12	-.10	.07	-.01	.07	-.05	.08	.03	.00	-.00	.02	.01
23	-.05	.01	.06	-.04	-.04	-.08	.04	-.02	.04	.03	.03	.04
24	.01	-.04	-.03	-.09	-.11	-.09	.09	.04	-.01	-.04	-.00	-.04

Estimates are based on full sample from 1974:1 to 1987:8. Standard errors of the estimates are 1.96/SQRT (166) = ± 0.15; where 167 is sample size.

Table 3. Values of test statistics in Dickey-Fuller unit root tests, 1974:1-1985:4

Test Statistic	Wheat Price(2)	Corn Price(1)	Soybean Price(1)	Wheat Ex. Rate(2)	Corn Ex. Rate(1)	Soybean Ex. Rat(1)
τ_u	-2.28	-2.30	-2.38 [c]	1.13	-.002	-.09
τ_τ	-2.81	-3.25 [c]	-4.33 [a]	-.57	1.41	-1.45
Φ_1	3.39	3.21	3.15	1.65	.48	.40
Φ_2	3.77	4.08	6.14 [b]	3.73	2.08	2.08
Φ_3	4.30	5.58 [c]	9.36 [a]	4.50	2.62	2.71

Note: The number in parentheses beside each variable name indicates the number of lags of Δx_t in the Dickey-Fuller regression. Critical values of the test statistics are from Fuller (1976, p.373), and Dickey and Fuller (1981, p.1063).

[a] reject at 1-percent significance level.

[b] reject at 5-percent significance level.

[c] reject at 10-percent significance level.

sion of a drift term. The t-statistic is 1.69, so the null hypothesis that the constant is zero can be rejected at the 10-percent level. For the other agricultural prices and the exchange rates, the constant is not significant in any of the equations, even at the 10-percent level.

The next step in our data analysis was to test for co-integration between the series in the bivariate relationships we wish to model. Each of the exchange rates are regressed on their commodity price counterparts. The residuals from these regressions are then tested for unit roots against the tables in Engle and Yoo (1987).

The results for the co-integration tests are found in table 4. The first row of table 4 displays the t-ratios for the simple DF test for co-integration. Rows 2 through 6 display t-ratios for the ADF tests for co-integration with the indicated number of lags. According to the Q-statistic serial correlation is removed with one additional lag of the residuals. We entertain higher lags because, for the wheat case, the test results change for higher order autoregressions unlike the stationarity tests where higher order autoregressions provided results that were consistent with those in table 3. The null hypothesis of no co-integration is rejected for the relationship of wheat price and the wheat exchange rate at the 5-percent level by the DF test and the ADF test *at lag 1 only*. (in this case adding a second lag to the ADF regression reduces the value of the test statistic by 35-percent). In all other cases, we do not reject the null of no co-integration.

If the wheat price is co-integrated with the wheat exchange rate, a reverse of the co-integrating regression, with the exchange rate as the dependent variable, should also yield residuals that are I(0) providing a check on the robustness of the initial results. For the wheat price-wheat exchange rate relationship, however, this does not hold true. A co-integration ADF test on the residuals from the reverse regression is unable to reject the null of no co-integration for any lag length. This casts considerable doubt on the initial ADF test results.

Based on the DF tests and the cointegration test results, we maintain that the correct time series specifications are in differences. Nevertheless, for the sake of comparison, we will estimate an error-correction model for the wheat price-exchange rate system.

Estimation and Evaluation of Forecasting Models

The experiments we conduct are to compare forecasts from univariate and multivariate models chosen on the basis of unit root, drift, and cointegration tests to results from other models. In order to abstract from the issue of how lag length selection can determine the accuracy of forecasts (see Lutkepohl 1985) and concentrate on the issue of how ignoring the time series properties of the data can impinge on the forecasting accuracy of the estimated models, we apply two lag length selection criteria to each of the models in the experiment. We have chosen the Bayesian Information Criterion of Schwartz(1978), denoted by "BIC," and the criterion of Hannan and Quinn(1979), denoted "HQ." We do not claim that either criteria will select the "best" forecasting model. We are only hoping to select univariate and multivariate models in a consistent manner, so that the resulting specifications can be compared to each other on criteria other than lag length. We compare the relative gain or loss in forecasting accuracy as measured by root mean square error (RMSE) at different forecast horizons. For each of the commodity prices and each exchange rate we estimate a univariate model in differences (UD) with and without a constant, and a univariate model in levels with a trend (UL). In addition, for each bivariate commodity price-exchange system we estimate a VAR in differences (VARD), and a VAR in levels with a trend in each equation (VARL).[27] For the wheat case, we also estimate an error correction model (ECM) and a VAR that includes a constant in the price equation but not in the exchange rate equation (SUR); this model is estimated by seemingly unrelated regressions.

Table 5 shows the lag lengths selected by the criteria we have chosen for each of the models indicated. The BIC and HQ criteria agree on lag length in 10 of the 12 univariate models, and 2 of the 3 bivariate systems. Where the two criteria disagree on lag length we report RMSEs for these

---

however, we follow Nelson and Plosser(1982) who take the position that for a log-differenced series to have a deterministic trend would imply that rates of change are ever increasing ($\beta > 0$) or ever decreasing ($\beta < 0$), a curious hypothesis for most economic variables, except, perhaps, a controlled variable such as the money supply.

[27] The motivation for including linear trends in levels models is found in Sims(1980, p.18).

Table 4. Test statistics from DF and ADF tests for co-integration, 1974:1-1987:8

Test	What Price/Wheat Ex. Rate System	Corn Price/Corn Ex. Rate System	Soybean Price/Soybean Ex. Rate System
Dickey-Fuller			
$t_{\hat{p}}$	-3.53 [a]	-.209	-1.41
Augmented DF			
1-lag	-3.54 [a]	-1.16	-2.16
2-lag	-2.61	-.741	-1.88
3-lag	-1.87	-.059	1.03
4-lag	-1.61	-.317	-1.31
5-lag	-2.16	-.760	1.19

Note: Critical values are interpolated from Engle and Yoo (1987) for sample size of 167.

[a] reject null of no cointegration at 5-percent significance level.

Table 5. Lag lengths chosen by BIC and HQ criterion for various forecasting model specifications, 1974-1985:4

Unvariate Autoregression

Model	Wheat Price	Corn Price	Soybean Price	Wheat Ex. Rate	Corn Ex. Rate	Soybean Ex. Rate
UL	6	2	2	2	2	2
UD	5	1(3) [a]	1(6) [a]	1	1	1

Vector Autoregressions

Model	Wheat Price/ Wheat Ex. Rate	Corn Price/ Corn Ex. Rate	Soybean Price / Soybean Ex. Rate
VARL	2(3) [a]	2	2
VARD	2(3) [a]	1	1
ECM	1(2) [a]	NA	NA

Note: See text for descriptions of the alternative models.

[a] For cases in which criterion do not select the same specification the lag length chosen by BIC is given first with that chosen by HQ in parentheses.

models based on the best out-of-sample forecasts in a competition between the BIC and HQ determined models.

The models are estimated based on the sample period 1974-1985:4, the starting month in 1974 depending on the lag length for a particular model. A post-sample period of 28 observations from 1985:5-1987:8 was held out to calculate RMSEs. All reported RMSEs are calculated based on forecasts of the *log-levels.*

Table 6 shows RMSEs calculated over the horizons 1, 2, 3, 6, 12, 15, and 18 for each of the models forecasting the commodity prices. The first set of comparisons is between price forecasts from UL models and the UD models consistent with our specification tests (i.e., with drift in the wheat price autoregression, and without drift in the case of corn and soybeans). On the basis of the gain/loss in accuracy from using the UD models instead of the UL models, we observe that in all but 3 of 21 cases, the difference specification dominates. At short horizons, (1,2,3), the gain in accuracy ranges from 4.2-percent for forecasts of the corn price at the 1-step horizon, to a gain in accuracy of 46-percent for soybean price at the 3-step horizon. At longer horizons most of the gains are even larger. For forecasts of the wheat price at horizons 12, 15, and 18, the gains in forecast accuracy from using the UD specification are 117-percent, 93-percent, and 142-percent respectively. Gains in forecast accuracy of 96-percent, 54-percent, and 29-percent at horizons 12, 15, and 18 are realized for the soybean price as well. For the corn price the UD model without constant is more accurate than the UL model at horizons through 12 months, but less accurate at longer horizons.

Accuracy gains are also realized for forecasts of each exchange rate when the UD models consistent with our specification tests (i.e., without drift) compete with the UL models (see table 7). At low horizons, the percentage gains in forecast accuracy from using the UD model instead of the UL model range from a low of 6-percent at the 1-step-ahead forecast of the wheat exchange rate, to 61-percent at the 3-step forecast of the corn exchange rate. At longer horizons, the gains are consistently above 50-percent ranging to a high of 289-percent at the 18-step forecast of the wheat exchange rate.

One curious result from the univariate regressions concerns the impact of including a drift term in the difference model. Our specification tests provided only marginal evidence for inclusion of a constant in the wheat price equation and rejected inclusion of a constant in the other price equations and the three exchange rate equations. As expected, forecast accuracy of the UD model for wheat price is improved substantially by inclusion of a drift term, especially at long forecast horizons. Forecast accuracy for corn and soybean prices is also improved by inclusion of a constant in the UD model. Inclusion of a constant worsens forecast accuracy for the exchange rates (which is again consistent with our specification tests).

Turning to the bivariate models, a comparison of the forecasts of the commodity prices from the VARL and the VARD models corroborates the results from the univariate model comparisons. The corn and soybean price VARD models without constant have lower forecast RMSEs than the corresponding VARL models in all but 1 of 14 cases. Again, inclusion of a constant seems to improve the forecasts of the commodity prices from the VARD models. For wheat price, including a constant in both equations of the VAR substantially improves forecast accuracy, consistent with our specification tests. Very slight additional gains at the longer horizons are obtained by including the constant only in the price equation (see SUR in table 6). The error correction model performs very poorly, confirming our suspicion of a type I error in the co-integration DF and 1-lag ADF tests.

A comparison of forecasts of the exchange rate from the VARL and VARD models also suggests a gain in forecast accuracy from appropriately accounting for unit roots. At every horizon, a gain in forecast accuracy is realized by using the difference specified model instead of the model in levels with trend. The difference model without constant forecasts better than the difference model with a constant for all three exchange rates.

The Role of the Exchange Rate

To investigate the role of the exchange rate in forecasting commodity prices, we first compare the forecasts of the wheat price from the VARL model with those from the UL model. Interestingly, we observe lower RMSEs at all horizons for the VARL model. The exchange rate, when included in the model, improves the accuracy of the forecasts of the wheat price by 6-percent

Table 6. RMSEs of forecasts of wheat price, corn price, and soybean price, alternative forecasting models, 1985:5-1987:8

Horizon	obs	VARL	VARD			ECM	UL	UD	
			SUR	w/ Const	w/o Const			w/ Const	w/o Const
		Wheat Price							
1	28	.050	.050	.050	.051	.207	.053	.051	.052
2	27	.090	.088	.088	.092	.474	.096	.089	.093
3	26	.118	.114	.114	.122	.251	.127	.113	.121
6	23	.146	.141	.141	.155	1.55	.164	.133	.149
12	17	.163	.104	.107	.186	2.77	.233	.107	.182
15	14	.168	.137	.142	.255	3.27	.281	.145	.253
18	11	.125	.102	.107	.235	3.81	.264	.109	.233
		Corn Price							
1	28	.075		.070	.070		.074	.067	.071
2	27	.149		.138	.140		.147	.130	.139
3	26	.203		.185	.190		.200	.177	.188
6	23	.286		.231	.245		.280	.220	.238
12	17	.421		.307	.365		.407	.306	.353
15	14	.484		.355	.480		.462	.400	.468
18	11	.493		.401	.458		.467	.403	.484
		Soybean Price							
1	28	.039		.029	.029		.036	.028	.029
2	27	.078		.049	.050		.020	.046	.050
3	26	.112		.069	.070		.102	.064	.070
6	23	.160		.084	.086		.149	.073	.086
12	17	.145		.054	.081		.159	.048	.081
15	14	.135		.062	.107		.165	.067	.107
18	11	.119		.087	.122		.159	.088	.123

Table 7. RMSEs of forecasts of wheat exchange rate, corn exchange rate, and soybean exchange rate alternative forecasting models, 1985:5-1987:8

Horizon	obs	VARL	VARD		UL	UD	
			Wheat Exchange Rate				
			w/ Const	w/o Const		w/ Const	w/o Const
1	28	.019	.017	.017	.019	.018	.018
2	27	.029	.025	.025	.031	.026	.025
3	26	.041	.033	.031	.043	.033	.031
6	23	.073	.059	.054	.082	.058	.053
12	17	.117	.082	.071	.135	.083	.072
15	14	.138	.083	.065	.159	.084	.065
18	11	.150	.077	.045	.179	.078	.046
			Corn Exchange Rate				
			w/ Const	w/o Const		w/ Const	w/o Const
1	28	.022	.020	.019	.025	.020	.019
2	27	.042	.035	.032	.048	.035	.033
3	26	.061	.049	.044	.071	.049	.044
6	23	.126	.097	.085	.144	.098	.086
12	17	.255	.188	.163	.268	.189	.164
15	14	.311	.233	.201	.323	.233	.202
18	11	.358	.280	.241	.376	.280	.242
			Soybean Exchange Rate				
			w/ Const	w/o Const		w/ Const	w/o Const
1	28	.027	.022	.022	.027	.022	.022
2	27	.052	.038	.036	.051	.039	.036
3	26	.076	.053	.048	.075	.054	.049
6	23	.154	.104	.074	.152	.106	.085
12	17	.285	.202	.129	.284	.250	.180
15	14	.342	.249	.220	.341	.253	.221
18	11	.395	.298	.262	.394	.303	.263

at the 1-step horizon, 43-percent at the 12-step horizon, and 67-percent and 111-percent, respectively, at the 15 and 18-step horizons. This result is consistent with Bessler and Babula (1987); their models were estimated in levels as well. However, when we compare the UD model with the VARD model for wheat price, a contrary result emerges. The VARD model for the wheat price shows *no forecasting superiority* to the UD model with a constant. The UD model preferred by our specification tests (with constant) dominates the VARL model as well, providing further evidence against the position that the exchange rate can help us forecast the wheat price. The exchange rate does not help to forecast the wheat price when account is taken of nonstationarity in the data.

Forecasts of soybean prices from the VARL model appear to corroborate the results for forecasts of wheat prices from the levels specifications, again giving evidence that the exchange rate matters in a forecasting context. At short to medium horizons, the UL model dominates VARL, at longer horizons, however, the VARL forecasts *improve* relative to forecasts from the UL model, showing a 27-percent and 34-percent gain in accuracy at horizons 15 and 18. This outcome is again not observed in the difference-specified models. The VARD models for the soybean price-exchange rate system do not outperform the UD models for the soybean price. Once again, the UD model bests the VARL model in all but one case, while the UD model with a constant provides better forecasts in all cases.

Finally, for the corn price, we get no evidence of the exchange rate improving forecasts. The VARL provides somewhat less accurate forecasts than the UL, especially at long horizons. The VARD models with and without constant perform about the same as the corresponding UD models.

Summary and Conclusions

The objectives of our paper have been to examine the appropriate specification of forecasting models with respect to possible nonstationarity in time series data, and to investigate the effects of the exchange rate in forecasting agricultural prices. In particular, we have been interested in the role of nonstationarity in evaluating whether incorporating exchange rates in bivariate models with agricultural commodity prices improves price forecasts compared to univariate models.

We find that when careful attention is paid to the unit root properties of the data, better forecasting models can be constructed than when these properties are ignored. Our specification tests suggested that each price and exchange rate series be modelled in differences, with a constant only in the wheat price model. The UD models perform better than the UL models for all three exchange rates and for the corn and the soybean prices, whether a constant is included in the difference specification or not. For wheat prices, forecasts from the UD model with a constant are much better than forecasts from either the UL model or a difference model without a constant.

The results from the comparisons among univariate models are reinforced by comparing the VARL models to the VARD models. Among forecasts of the exchange rates and the commodity prices, the VARD models produced lower RMSEs than the VARL models in 41 of 42 cases. Thus, both the comparisons among univariate and bivariate models confirm the theoretical result that forecasts from nonstationary models are sub-optimal. Our results argue for testing for nonstationarity and cointegration and specifying models appropriately before estimating their parameters and making forecasts.

A further consequence of ignoring nonstationarity arises when we examine the role of the exchange rate in forecasting agricultural prices. If we had only examined the forecasting proficiency of the VARL and UL models, without recognizing the possibility of unit roots, we likely would have concluded that inclusion of the exchange rate in a bivariate model *improves* price forecasts. This conclusion, though it has been reported in the literature, is suspect. In our analysis, forecasts from the VARL models for wheat and soybean prices outperform UL models for these prices. But these VARL models are beaten in out-of-sample forecasting performance by UD and VARD models. Further, the VARD models do not improve on forecasts from the UD models. This suggests that when the information in the data is used efficiently, in this case by removing unit roots, the exchange rate does not help to forecast prices.

The broad issue raised by our analysis concerns the implications of the result that incorporating an exchange rate in bivariate models with wheat, corn or soybean prices does not improve forecasts from univariate models. This result is perhaps not all that surprising since arbitrage in competitive asset markets may lead prices themselves to reflect all that is known, at a given moment in time, about their own future. The failure of the exchange rate to improve upon univariate price forecasts does not necessarily imply that macroeconomic factors are unimportant to agriculture. We have touched only a small piece of a complex problem. Our results suggest that macroeconomic shocks reflected in exchange rates may affect agricultural prices--other asset prices--simultaneously. To uncover the macroeconomic impacts may require a more articulate identification of the macroeconomic shocks than is conveyed simply by associating exchange rate shocks with macroeconomics, and price shocks with agriculture in a reduced form model. This is a more subtle result than concluding that macroeconomic phenomena matter to agriculture because exchange rates improve agricultural price forecasts, when the latter result may arise only from the failure to address nonstationarity in the data. Capturing macroeconomic impacts on agriculture in dynamic models with time series data remains a challenging area of research.

References

Akaike, H. (1971), Autoregressive Model Fitting for Control, Annals of the Institute of Statistical Mathematics, 23, 163-180.

Ashley, R.A., On the Relative Worth of Recent Macroeconomic Forecasts, forthcoming in the International Journal of Forecasting.

Banerjee, A. et al. (1986), Exploring Equilibrium Relationships in Econometrics through Static Models: Some Monte Carlo Evidence, Oxford Bulletin of Econometrics and Statistics, 48, 253-277.

Bernanke, B.S. (1986), Alternatiave Explanations of the Money-Income Correlation, NBER Working Paper # 1842, 49-100.

Bessler, (1984a), Relative Prices and Money: A Vector Autoregression on Brazilian Data, American Journal of Agricultural Economics, 66, 25-30.

Bessler, D.A. (1984b), Analysis of Dynamic Economic Relationships: An Application to the U.S. Hog Market, Canadian Journal of Agricultural Economics, 32, 109-24.

Bessler, D.A. (1986), Forecasting Vector Autoregressions with Bayesian Priors, American Journal of Agricultural Economics, 3, 144-151.

Bessler, D.A. and R. Babula (1987), Forecasting Wheat Exports: Do Exchange Rates Matter? Journal of Business and Economic Statistics, 5, 397-406.

Blanchard, 0.J. and M.W. Watson (1986), Are Business Cycles all Alike? The American Business Cycle, Continuity and Change, Ed. Robert J. Gordon, University of Chicago Press, 123-179.

Box, G.E.P. and G.M. Jenkins (1970), Time Series Analysis, Forecasting and Control, San Francisco: Holden Day, First Edition.

Box, G.E.P. and G.M. Jenkins (1976), Time Series Analysis, Forecasting and Control, San Francisco: Holden Day, Second Edition.

Brandt, J.A. and D.A. Bessler (1984), Vector Autoregressions on U.S. Hog Prices, North Central Journal of Agricultural Economics, 6, 29-36.

Cooper, R.L. (1972), The Predictive Performance of Quarterly Econometric models of the United States, In Econometric Models of Cyclical Behavior, Ed. B.G. Hickman, N.Y.: Columbia University Press.

Davidson, J.E.H. et al. (1978), Econometric Modelling of the Aggregate Time Series Relationship between Consumers' Expenditure and Income in the United Kingdom, Economic Journal, 88, 661-692.

Dickey, D.A. (1975), Estimation and Hypothesis Testing in Nonstationary Time Series, Ph.D. dissertation, Iowa State University.

Dickey, D.A. et. al. (1986), Unit Roots in Time Series Models: Tests and Implications, American Statistician, 40,1, 12-26.

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

______ (1981). Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root, Econometrica, 49, 1057-1072.

Doan, T., R. Litterman, and C. Sims (1983), Forecasting and Conditional Projections Using Realistic Prior Distributions, Econometric Reviews, 3, 1-100.

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

Engle, R.F. and B.S. Yoo (1987), Forecasting and Testing in Co-Integrated Systems, Journal of Econometrics, 35. 143-159.

Fackler, P. (1988), Vector Autoregressive Techniques for Structural Analysis, WP # 113 North Carolina State University Dept. of Economics and Business.

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

Granger, C.W.J. (1981), Some Properties of Time Series Data and Their Use in Econometric Model Specification, Journal of Econometrics, 16, 121-130.

Granger, C.W.J. (1983), Co-Integrated Variables and Error Correcting Models, Discussion Paper (UCSD).

Granger, C.W.J. (1986), Developments in the Study of Co-Integrated Variables, Oxford Bulletin of Economics and Statistics, 8, 213-228.

Granger, C.W.J. and A.A. Weiss (1983), Time Series Analysis of Error Correcting Models in S. Karlin et al. eds., Studies in Econometrics, Time Series, and Multivariate Analysis, New York: Academic Press.

Granger, C.W.J. and P. Newbold (1974), Spurious Regressions in Econometrics, Journal of Econometrics, 2, 111-120.

Granger, C.W.J. and P. Newbold (1986), Forecasting Economic Time Series, New York: Academic Press, Second Edition.

Hall, G. (1986), An Application of the Granger and Engle Two-Step Estimation Procedure to United Kingdom Aggregate Wage Data, Oxford Bulletin of Economics and Statistics, 48, 229-239.

Hannan, E.J. and B.G. Quinn (1979), The Determination of the Order of an Autogregression, Journal of the Royal Statistical Sociey, B, 41, 190-195.

Harvey, A.C. (1981), The Econometric Analysis of Time Series, Deddington, Great Britain: Phillip Allen.

Hendry, D.F. and G.E. Mizon, (1978), Serial Correlation as a Convenient Simplification, Not a Nusiance: A Comment on a Study of the Demand for Money by the Bank of England, Economic Journal, 88, 549-63.

Hsaio, C. (1979), Autoregressive Modeling of Canadian Money and Income Data, Journal of the American Statistical Association, 74, 553-560.

Judge, G.G. et al. (1985), The Theory and Practice of Econometrics, New York: Wiley.

Kennedy, P. (1985), A Guide to Econometrics, Cambridge, MA: MIT Press.

Kling, J.L. and D.A. Bessler (1985), A Comparison of Multivariate Forecasting Procedures for Economic Time Series, International Journal of Forecasting, 1, 5-24.

Litterman, R.B. (1976), Techniques of Forecasting Using Vector Autoregressions, Unpublished Ph.D. dissertation, University of Minnesota, Department of Economics.

Litterman, R.B. (1984), Specifying Vector Autogressions for Macroeconomic Forecasting, Staff Report 92, Federal Reserve Bank of Minneapolis, Research Dept.

Litterman, R.B. (1986), A Statistical Approach to Economic Forecasting, Journal of Business and Economic Statistics, 4, 1-4.

Ljung, G.M. and G.E.P. Box (1978), on a Measure of Lack of Fit in Time Series Models, Biometrika, 65, 297-303.

Lucas, R.E. and T.J. Sargent (1979), Beyond Keynesian Macroeconomics, in Rational Expectations and Econometric Practice, Minneapolis: University of Minnesota Press.

Lutkepohl, H. (1982), Differencing Multiple Time Series: Another Look of Canadian Money and Income Data, Journal of Time Series Analysis, 3, 235-243.

Lutkepohl, H. (1985), Comparison of Criterion for Estimating the Order of a Vector Autoregressive Process, Journal of Time Series Analysis, 6, 35-52.

Meese, R. A., and K. Singleton (1982), On Unit Roots and the Empirical Modeling of Exchange Rates, Journal of Finance, 37, 1029-1035.

Nelson, C. R. (1972), The Predictive Performance of the FRB-MIT-PENN Model of the U.S. Economy, American Economic Review, 62, 902-917.

Nelson, C. R. and H. Kang (1984), Pitfalls in the Use of Time as an Explanatory Variable in a Regression, Journal of Business and Economic Statistics, 2, 73-82.

Nelson, C. R., and C. I. Plosser (1982), Trends and Random Walks in Macroeconomic Time Series, Journal of Monetary Economics, 10, 139-162.

Orden, D. (1986), Agriculture, Trade, and Macroeconomics: The U.S. Case, Journal of Policy Modeling, 8, 27-51.

Phillips, P.C.B. (1986), Understanding Spurious Regressions in Econometrics, Journal of Econometrics, 33, 311-340.

Phillips, P.C.B., and R.N. Durlauf (1986), Multiple Time Series Regression with Integrated Processes, Review of Economic Studies, 23, 473-495.

Prothero, D. L., and K. F. Wallis (1976), Modelling Macroeconomic Time Series (with discussion), Journal of the Royal Statistical Society, Series A, 139, 468-500.

Schwartz, G. (1978), Estimating the Dimensions of a Model, The Annals of Statistics, 6, 461-464.

Schwert, G.W. (1987), Effects of Model Specification on Tests for Unit Roots in Macroeconomic Data, Journal of Monetary Economics, 20, 73-103.

Sims, C. A. (1980), Macroeconomics and Reality, Econometrica, 48, 1-48.

Sims, C.A. (1986), Are Forecasting Models Useable for Policy Analysis? Quarterly Review, Federal Reserve Bank of Minneapolis, Winter, 2-16.

Stock, J.H. (1987), Asymptotic Properties of Least Squares Estimators of Co-Integrating Vectors, Econometrica, 55, 1035-1056.

Stock, J.H., and M.W. Watson (1987), Interpreting the Evidence on Money-Income Causality, NBER Working Paper # 2228.

Thornton, D., and D. Batten (1985), Lag Length Selection and Tests of Granger Causality Between Money and Income, Journal of Money Credit and Banking, 17, 164-178.

Taio, G.C., and G.E.P. Box (1981), Modeling Multiple Time Series with Applications, Journal of the American Statistical Association, 76, 802-816.

U.S. Department of Commerce, Survey of Current Business, various issues 1974-1987.

Wasserfallen, W. W. (1986), Non-stationarities in Macroeconomic Time Series--Further Evidence and Implications, Canadian Journal of Economics, 19, 498-510.

Webb, R.H. (1985), Toward More Accurate Macroeconomic Forecasts from Vector Autoregressions, Economic Review, Federal Reserve Bank of Richmond, July/August, 3-17.

Wold, H. (1954), A Study in the Analysis of Stationary Time Series, Uppsala: Almquist and Wicksell, 2nd Ed.

Yule, G.U. (1926), Why Do We Sometimes Get Nonsense Correlations Between Time Series? - A Study in Sampling and the Nature of Time Series, Journal of the Royal Statistical Society, 9, 1-64.

Zellner A. (1982), Basic Issues in Econometrics, Past and Present, in Basic Issues in Econometrics, Chicago: University of Chicago Press, 26-34.

Zellner A. (1929), Statistical Analysis of Econometric Models, in Basic Issues in Econometrics, Chicago: University of Chicago Press, 83-119.