Searching for the Missing Link: High Frequency Price Dynamics and Autocorrelations For Seasonally Produced Commodities

Craig Pirrong
John M. Olin School of Business
Washington University
St. Louis, MO 63119

July 21, 1999

Abstract. Recent research on the behavior of commodity prices demonstrates that storage alone is insufficient to induce price autocorrelations of the magnitude observed in empirical data. Researchers have posited a high autocorrelation in commodity demand to remedy this deficiency. However, a traditional storage model with a high demand autocorrelation cannot explain salient features of commodity futures prices for seasonally produced goods, specifically the high correlation between old crop and new crop futures prices and the responsiveness of old crop prices to news about the expected harvest. Incorporating intertemporal substitution in consumption can explain these features of commodity futures prices. Intertemporal substitution provides an additional linkage between old crop and new crop prices that overcomes the deficiencies in the standard storage model which assumes no such substitution effect.

JEL Classification: G13, D90, Q11

1 Introduction

What links commodity prices over time? The optimal storage of commodities and the consequent behavior of commodity prices represents one of the most basic problems in the economics of intertemporal resource allocation. Unfortunately, this "Ur problem" of finance is as yet unresolved. Commodity prices are highly autocorrelated, but no one has been able to explain why.

Recent contributions to the theory of commodity price behavior examine the sources of the high autocorrelation in commodity prices using non-linear rational expectations models.[1] Early models posit i.i.d. net demand shocks. These models imply that storage induces some autocorrelation in the data, but less than observed in actual prices.[2] Later models attempt to address the failure of models with i.i.d. demand shocks to fit the data by assuming autocorrelated net demand disturbances. High autocorrelations in net demand are required to explain the high autocorrelations in prices.[3] Since there is little evidence of autocorrelation in output for agricultural commodities, one implication of these results is that high autocorrelation in demand shocks is required to explain high autocorrelation in prices.

This article examines whether demand autocorrelation is indeed the crucial intertemporal link between commodity prices. To do so, I depart from the received approach that (1) generally examines price behavior at low frequencies (e.g., annually) and (2) does not (with a few exceptions) distinguish between commodities that are produced continuously (e.g., copper) from those that are produced periodically (e.g., corn).[4] In contrast, I analyze the dynamics of seasonally produced commodities at a higher frequency. This approach exploits new sources of information and provides new insights on the forces linking prices over time.

Specifically, the focus on seasonally produced commodities permits the untangling of the effects of supply shocks and demand shocks on price movements. Moreover, modeling price dynamics at a frequency that exceeds the frequency of production generates implications about how spot and futures

[1] Deaton and Laroque (1992, 1995, 1996), Williams and Wright (1991), Chambers and Bailey (1996).

[2] Deaton and Laroque (1992), Williams and Wright (1991).

[3] Deaton and Laroque (1995, 1996), Chambers and Bailey (1996).

[4] Chambers and Bailey (1996) demonstrate the importance of explicit modeling of periodicity in net demand shocks. Routledge, Seppi and Spatt (1997) model forward curves for seasonal commodities using a dynamic programming model.

prices respond to supply and demand shocks at different times of the year. These include implications regarding the correlation between futures prices with different expiration dates, and the responses of futures prices with different expiration dates to supply shocks. These implications can be tested using futures price data that are available at high (e.g., daily) frequencies. For example, the daily correlation (during the months of March and April) between the price of corn for delivery in May (prior to the harvest) and the price for delivery in the following December (after the harvest) provides information on the correlations of prices across crop years. The May future is referred to as an "old crop" contract because it matures prior to the new harvest, whereas the December future is called a "new crop" contract because it expires soon after the harvest is completed.

To achieve these results, this article analyzes an infinite horizon model in which production occurs during a single period during the year (the "Fall") and consumption occurs in all four periods ("Winter," "Spring," and "Summer" as well as "Fall.") Agents decide how much to consume after the harvest and in each subsequent period by solving a stochastic dynamic program.[5] Unlike received models in the storage-price dynamics literature, the specification herein posits two distinct shocks, a demand shock and a supply shock. Demand shocks occur each period/season and may be autocorrelated. Agents receive information about the distribution of output at the next harvest in each period; all uncertainty about this year's output is resolved in the "Fall" when the harvest occurs. This separation of supply and demand shocks permits a more precise determination of the source of intertemporal price linkages for storable, seasonally produced commodities. Moreover, unlike received models in the literature, the theory permits intertemporal substitution in consumption of the seasonally produced good.

The model is solved numerically for a wide variety of parameter values. These solutions generate several testable hypotheses regarding the behavior of commodity prices:

- If there is no intertemporal substitution, the correlation between old crop and new crop futures prices prior to the harvest are low if supply shocks are important.

- If there is no intertemporal substitution, old crop futures prices typ-

[5] A consumption decision implies a storage decision since carry-in is given at each point in time.

ically do not respond appreciably to shocks to the expected havest. That is, old crop futures prices will change far less in response to the arrival of new information about the next harvest than will new crop futures prices.

- If there is significant intertemporal substitution in consumption then old crop-new crop futures price correlations will be high, and the arrival of information about the size of the impending harvest will cause old crop and new crop futures prices to move in the same direction by similar amounts.

These hypotheses are tested using data on futures prices for seasonal commodities. The behavior of these prices is inconsistent with the implications of the high demand autocorrelation-no substitution model. Old crop-new crop futures correlations are high, old crop and new crop futures prices exhibit approximately equal sensitivity to news about the size of the impending harvest, and old crop prices are more variable on days when the most reliable forecasts regarding crop size are released. Put differently, if there is no intertemporal substitution, even highly autocorrelated demand and supply disturbances explain neither the high correlation between old crop and new crop futures prices observed prior to the harvest, nor the sensitivity of old crop futures prices to information about future harvests. However, extensive intertemporal substitutibility can generate old crop-new crop futures price correlations and old crop price sensitivity to expected supply shocks that are far closer in magnitude to those observed empirically for major commodities. Thus, demand considerations may be important in explaining high frequency price correlations, but the persistence of demand shocks cannot be the crucial feature of demand that influences commodity price dynamics. Instead, intertemporal substitutability is a plausible candidate for the "missing link" between commodity prices over time. In any event, the results show that some intertemporal linkage other than storage and autocorrelated demand is required to explain the high frequency dynamics of seasonally produced commodity prices.

Explicit consideration of the implications of the seasonality of production explains these findings. Solution of the dynamic program describing the storage/consumption decision quite sensibly implies that it is seldom optimal to carry inventory from the period immediately preceding the harvest to the harvest period. As a result, storage and arbitrage do not link old crop and new crop prices. Shocks to the expected harvest therefore have little impact on old crop prices. If shocks to the harvest are large in magnitude, this lack of connection of prices via storage implies that new crop and old crop prices will exhibit little correlation even if demand is highly autocorrelated. In contrast, intertemporal substitution provides an alternative channel by which shocks to the anticipated harvest influence prices prior to the harvest. Thus, anticipation during the spring of a large harvest in the fall causes both spot prices and new crop futures prices to fall in the spring. This can create correlations of the magnitude observed in the data.

The remainder of this paper is organized as follows. Section 2 presents a stochastic dynamic programming model describing optimal storage of a seasonally produced commodity in a competitive market with multiple decision periods per year and describes the numerical methods used to solve this model. Section 3 analyzes the model and generates testable implications regarding the correlation between new crop and old crop futures prices and the responsiveness of old crop prices to news about the size of the coming harvest. Section 4 presents events pertaining to these predictions. This evidence shows that there is a relatively high degree of correlation between these prices and that harvest-related news impacts old crop prices. These implications are inconsistent with the basic storage model but can be explained by a storage model with intertemporal substitution. Section 5 summarizes the work.

2 Optimal Competitive Storage for a Seasonally Produced Commodity

2.1 Introduction

The literature on commodity price dynamics has relied upon a standard storage model that assumes that consumption and production occur every period. Moreover, the empirical tests of this storage model typically employ spot price data only.

This received approach has several limitations. First, the existing approach posits net demand shocks that consist of a demand shock minus a supply shock and cannot disentangle the effects of supply shocks from those of demand disturbances. These models can describe the behavior of prices only at relatively low frequency corresponding to the frequency of production. Second, by utilizing only spot price data, the received empirical efforts have ignored a potentially valuable source of information–futures prices. Third, testing of the model is extremely difficult unless one assumes a very simple structure for net disturbances. With i.i.d. net shocks the basic storage model implies that prices follow a particular autoregression that can serve as the basis for empirical testing. Such a straightforward implication does not hold with any time dependence in the disturbances. Estimating a model with time dependent disturbances entails acute econometric difficulties.

In this section I analyze a model that addresses these difficulties. This expanded model explicitly incorporates seasonality in production. Due to seasonality, production does not occur every period, even though the commodity is consumed every period. This seasonal structure generates new implications about the relation between prices, supply shocks and demand shocks. In particular, it permits me to disentangle the effects of supply and demand shocks. This facilitates a test of the model. Moreover, I derive testable implications regarding the behavior of futures prices. This permits the use of a new source of information to test the implications of storage models. These implications can be tested using simple statistical techniques.

The objective of this analysis is to develop a framework that generates predictions regarding the correlations between old-crop and new-crop futures prices, and the responsiveness of old-crop prices to information about the size of the impending harvest. The methodology is in some respects similar to that of Deaton-Laroque (1992 and 1996), Williams-Wright (1991), and Miranda-Rui (1999), all of whom use simulations to evaluate the ability of a numerical solution of the model to generate time series dynamics that mimic the behavior exhibited by actual commodity prices.[6] I also evaluate the ability of a numerically solved storage model to mimic salient features of spot and futures prices under assumed parameter values, but instead of using simulation, I approximate the correlations and price sensitivities of interest numerically. This section shows that these approximations can be derived using estimates of the partial derivatives of the pricing functions. These approximations generate implications (derived in section 3) that are tested in section 4.

[6]Numerous papers in the economic growth and equity premium puzzle literatures also use numerical simulations under assumed parameter values to determine the ability of the models to explain stylized facts about economic dynamics.

2.2 A Framework for Determining Spot-Forward Correlations and Sensitivities to Harvest Shocks

Consider an infinite horizon storage model. Consumption occurs every period, $t, t+\Delta t, t+2\Delta t, \ldots$. A commodity is produced at times $\tau, 2\tau, 3\tau, \ldots$, where $\tau \geq 2\Delta t$. The fact that production does not occur every period formalizes the seasonality in the economy. The commodity is storable. Since the commodity is produced only periodically, to consume at any $t \in [j\tau + \Delta t, (j+1)\tau - \Delta t]$ agents must have stored positive amounts of the commodity at all dates $t' \in [j\tau, t - \Delta t]$. Call x_t the amount of inventory carried-in to time t. At t, agents decide how much of this commodity to consume and how much to store for future consumption.

Output is random. For each $t \in [j\tau + \Delta t, (j+1)\tau - \Delta t]$, competitive, risk neutral agents have information about the next harvest. At such t, the expected value of the next harvest is $\bar{H} + \sigma_z Z_t$ where $\bar{H}$ is the unconditional expectation of the harvest, σ_z is a parameter that describes the rate of information flow regarding the size of the incoming harvest, and Z_t is a standardized shock to the expected harvest.

The dynamics of the standardized shock to the expected harvest at $(j+1)\tau$ for $t \in [j\tau + \Delta t, (j+1)\tau]$ are:

$$\Delta Z_t = \alpha(x_t, Z_t, \eta_t)\Delta t + z_t$$

where z_t is a bounded i.i.d. variate with variance Δt. For $t \leq (j+1)\tau$ the expectation of the the harvest at $(j+2)\tau$ is the unconditional expectation $\bar{H}$. The $\alpha(.)$ function describes the drift of the standardized harvest forecast shock. This generalized form includes standard processes (e.g., random walk, AR1) as special cases.

Demand is also random; η_t denotes a standardized demand shock at t. Inverse demand at any time t is given by $D(q_t, \eta_t, \bar{P}_{t+\Delta t})$, where $\bar{P}_{t+\Delta t}$ is the time t expectation of the price at $t + \Delta t$. The inclusion of next period's expected price means that current and future consumption may be substitutes. The dynamics of the standardized demand shock η are:

$$\Delta \eta_t = \mu(x_t, Z_t, \eta_t)\Delta t + e_t$$

where e_t is another bounded i.i.d. variate with variance Δt that is uncorrelated with z_t. The $\mu(.)$ function gives the drift of demand shock processes.

Determination of a competitive equilibrium requires solution of a stochastic dynamic programming problem. Consider the solution of the optimal storage problem for this commodity. Call $P(x_t, Z_t, \eta_t, t)$ the time t spot price as a function of the state variables (carry-in, expected harvest, and demand) implied by the solution to this program; this price function depends on the time index because of the seasonality in the problem. Due to seasonality, knowledge of the state variables alone is insufficient to determine price; the price function should change as time passes and the harvest nears. Similarly, call $F(x, Z, \eta, t, T)$ the forward price for the commodity delivered at $T > t$ as of t. Because agents are risk neutral, this forward price equals the expected spot price at T conditional on time t information.

A first-order Taylor's expansion of the spot and forward pricing functions implies:[7]

$$\Delta P_t = \theta_P(x_t, Z_t, \eta_t, t)\Delta t + P_Z z_t + P_\eta e_t \tag{1}$$

and

$$\Delta F_t = \theta_F(x_t, Z_t, \eta_t, t, T)\Delta t + F_Z z_t + F_\eta e_t \tag{2}$$

where subscripts indicate partial derivatives with respect to the subscripted variable.[8] These imply that the variance of the spot price is:

$$\sigma_P^2(x_t, Z_t, \eta_t, t) \equiv E(\Delta P_t^2)/\Delta t = P_Z^2 + P_\eta^2 \tag{3}$$

and the variance of the forward price is:

$$\sigma_F^2(x_t, Z_t, \eta_t, t) \equiv E(\Delta F_t^2)/\Delta t = F_Z^2 + F_\eta^2 \tag{4}$$

Moreover, the correlation between ΔP_t and ΔF_t is:[9]

$$corr(\Delta P_t, \Delta F_t) = \frac{F_Z P_Z + F_\eta P_\eta}{\sqrt{(P_Z^2 + P_\eta^2)(F_Z^2 + F_\eta^2)}} \tag{5}$$

It is also possible to estimate the slope coefficient of a regression of the spot price against the forward price. This coefficient is:

$$\frac{cov(\Delta P_t, \Delta F_t)}{var(\Delta F_t)} = \frac{F_Z P_Z + F_\eta P_\eta}{F_Z^2 + F_\eta^2}$$

[7]A first-order expansion is sufficient because this study focuses on variances and correlations; second order terms are $o(\Delta t)$ in the variance and correlation expressions.

[8]θ_P and θ_F are drift functions. They are irrelevant to the determination of correlations and are therefore ignored hereafter.

[9]Note that multiple shocks are required in order to obtain correlations other than -1, 0, or +1.

The foregoing implies that the spot and forward correlations and spot-forward regression coefficients depend on the sensitivity of prices to the random variables Z_t and η_t. Determination of the relevant partial derivatives that measure these sensitivities requires a solution of the pricing functions for each period. The dynamic programming problem that generates the pricing functions cannot be solved in closed form, so the $P(.)$ and $F(.)$ functions cannot be determined analytically. Instead, it is necessary to solve the problem numerically and approximate the price functions. The next section describes the method for solving this seasonal storage problem numerically.

2.3 Numerical Implementation

In theory, there is a separate price function for each time period between harvests. Computational considerations limit the number of discrete periods between harvests that can be analyzed. Herein I assume that there is a harvest period followed by three non-harvest periods, each with its own price function, making a total of four price functions. It is intuitive and convenient refer to a set of four periods including one harvest period as a year, and to refer to the periods as "seasons."[10] A year is normalized to be of length 1. The seasons are numbered one through four. Therefore, each season is of length $\Delta t = 1/4$. Production occurs in the fourth season of each year. Consumption occurs in each season of each year. At season i in year t, $i = 1, \ldots, 4$, inverse demand for the commodity is:

$$p_{4(t-1)+i} = D_i(q_{4(t-1)+i}, \bar{p}_{4(t-1)+i+1}, \eta_{4(t-1)+i}) \tag{6}$$

In (6), $q_{4(t-1)+i}$ is the amount consumed in season i in year t $D_i(.)$ is the demand function for season i, $p_{4(t-1)+i}$ is the spot price of the commodity in year t and season i, $\bar{p}_{4(t-1)+i+1}$ is next period's expected spot price (where the expectation is conditional on information available in season i of year t), and $\eta_{4(t-1)+i}$ is a demand shock. This specification permits intertemporal substitutibility in consumption if $\partial D_i / \partial \bar{p}_{4(t-1)+i+1} \neq 0$.

It is assumed that the demand shock follows an AR1 process:

$$\eta_{4(t-1)+i} = \rho \eta_{4(t-1)+i-1} + e_{4(t-1)+i} \tag{7}$$

[10]This framework is similar to that of Chambers-Bailey (1997) and Routledge, Seppi, and Spatt (1997). Williams and Wright (1991) implement a similar methodology for a commodity with seasonal demand but no seasonality in supply. None of these models considers separate supply and demand shocks.

where $e_{4(t-1)+i}$ is a bounded i.i.d. disturbance with variance Δt.

In each season $j \leq 3$ agents receive a signal about the output expected in the next harvest period (i.e., the next season 4). As before, the unconditional expectation of the harvest is $\bar{H}$. In period 1, the expected harvest is $E_1 H = \bar{H} + \sigma_z z_1$, where z_1 is the signal observed in period 1. In period 2 the expected harvest is $E_2 H = \bar{H} + \sigma_z(z_1 + z_2)$. In period 3 the expected harvest is $E_3 H = \bar{H} + \sigma_z(z_1 + z_2 + z_3)$. The actual harvest that occurs in period 4 is $H = \bar{H} + \sigma_z(z_1 + z_2 + z_3 + z_4)$. The z's are bounded i.i.d. variates with variance Δt.

In this framework, the harvest forecast follows a random walk. This set-up formalizes the notion that agents receive information about the harvest throughout the crop year. As an example, winter precipitation, spring planting conditions, summer weather during plant pollination, and fall field conditions all influence the size of the harvest. Moreover, the precision of agents' expectations increases as the crop year proceeds; agents have better information about the next harvest of corn in the summer than in the winter. Finally, forecasts are unbiased.[11]

Immediately after the harvest, the expected harvest in the following year is the unconditional expectation $\bar{H}$. That is, successive harvests are uncorrelated. Moreover, the η and the z are uncorrelated.

In period 4, competitive risk neutral agents must decide how to allocate the available harvest between current consumption and storage. Since there is no production in the next three succeeding seasons, all consumption prior to the next harvest must be supplied from stocks held from this harvest. In period 4 of year t, when the demand shock is η_4 agents endowed with a harvest H and carry-in from previous harvests of x choose carry-out $s_4 \geq 0$ so that

$$P_4(H, x, \eta) = \max[\beta E P_1(z_1, s_4, \eta_1), D(H + x, E P_1(z_1, 0, \eta_1), \eta_4)]. \qquad (8)$$

In this expression, $P_4(.)$ is the equilibrium price in period 4 as a function of the harvest, carry-in and the demand shock, and $P_1(.)$ is the equilibrium price in period 1 as a function of the harvest signal received in period 1, the demand shock in period 1, and carry-in available in period 1. $\beta < 1$ is a discount factor. This expression states that agents choose storage in period 4 so that

[11] The specification of the evolution of the expected harvest assumes that the flow of information is constant in each period. A more general set-up would allow for variations in the flow of information throughout the year, e.g., $H = \bar{H} + \sigma_{z1} z_1 + \sigma_{z2} z_2 + \sigma_{z3} z_3 + \sigma_{z4} z_4$.

the spot price in period 4 equals the discounted expected future price in period 1 if storage is positive, and equals the marginal value of consumption measured at $H + x$ if nothing is carried over. If this expression did not hold, an arbitrage opportunity would exist. For example, if $P_4 < \beta E P_1$, risk neutral agents could buy the spot commodity in period 4, store it until period 1, and expect to sell it at a higher price next period.[12]

Similarly, in period 1 of year t, given a standardized expected harvest shock Z_1, initial carry-in x and demand shock η_1, carry-out $s_1 \geq 0$ is determined by:

$$P_1(Z_1, x, \eta_1) = \max[\beta E P_2(Z_2, s_1, \eta_2), D(x, E P_2(Z_2, 0, \eta_2), \eta_1)]. \qquad (9)$$

where $Z_2 = Z_1 + z_2$ and expectations are over z_2 and η_2. In this expression, the price function $P_1(.)$ gives the period 1 price as a function of the standardized shock to the expected harvest Z_1, carry-in x, and demand shock η_1. Given harvest shock Z_2, carry-in x and demand shock η_2, carry-out in period 2 of year t, $s_2 \geq 0$, solves:

$$P_2(Z_2, x, \eta_2) = \max[\beta E P_3(Z_3, s_2, \eta_3), D_2(x, E P_3(Z_3, 0, \eta_3), \eta_1)] \qquad (10)$$

where $Z_3 = Z_2 + z_3$ and expectations are over z_3 and η_3. Finally, given harvest shock Z_3, carry-in x and demand shock η_3, carry-out in period 3 of year t, $s_3 \geq 0$, solves:

$$P_3(Z_3, s_3, \eta_3) = \max[\beta E P_4(Z_4, s_3, \eta_4), D_3(x, E P_4(Z_4, 0, \eta_4), \eta_3)] \qquad (11)$$

where $Z_4 = Z_3 + z_4$, and expectations are over z_4 and η_4. Again, these expressions mean that competitive storers choose carry-out to equate the current spot price and next period's expected price.

The basic contours of the numerical solution of a recursive dynamic economic model of this sort are fairly well understood, so only a brief description of the methodology is required. There are three state variables: the standardized shock to the expected harvest Z, the demand shock η, and initial

[12] Deaton-Laroque (1992, 1996) and Chambers-Bailey (1996) show that if net demand shocks are uncorrelated, storage is zero if the spot price exceeds some constant "cutoff price" equal to the unconditional expectation of the following period's price assuming zero storage. This cutoff price feature implies that prices follow an autoregression that can be estimated. This implication does not follow if either harvest shocks or demand shocks are autocorrelated.

inventory x. The first step is to discretize the state variables to create a grid in Z, η, and x for each of the four seasons.

The values of the grid in the standardized harvest shock Z are bounded by -2.4 and +2.4 and are equally spaced with increments ΔZ. There are N_Z points along the grid,

$$\mathcal{Z} = \{-2.4, -2.4 + \Delta Z, -2.4 + 2\Delta Z, \ldots, 2.4 - \Delta Z, 2.4\}.$$

Define $Z_j = -2.4 + (j-1)\Delta Z$. In period j the expected size of the next harvest is $\bar{H} + \sigma_z Z_j$.

The transition probability matrix for Z during seasons 1, 2, and 3 is constructed as follows:

$$\pi_i(Z_j, Z_k) \equiv \Pr[Z_{4(t-1)+i+1} = Z_k | Z_{4(t-1)+i} = Z_j] = N(\epsilon_1) - N(\epsilon_2) \qquad (12)$$

for $i = 1, 2, 3$, and where

$$\epsilon_1 = .5[Z_{k+1} + Z_k] - Z_j$$

$$\epsilon_2 = .5[Z_k + Z_{k-1}] - Z_j$$

and $N(.)$ denotes the standard normal distribution function. That is, if $Z = Z_j$, (i.e., the initial value of Z is at the j'th point on the Z grid) the probability that the value of Z in the next period is $Z = Z_k$ (i.e., is at the k'th point on the grid) equals the probability that a standard normal variate falls in an interval of length ΔZ centered on $Z_k - Z_j$.[13]

In season 4, Z is an i.i.d. variate (because successive harvests are not correlated). Therefore, the probability of transiting from $Z = Z_j$ in season 4 to $Z = Z_k$ in season 1 is

$$\pi_4(Z_j, Z_k) = N[.5(Z_k + Z_{k+1})] - N[.5(Z_k + Z_{k-1})].$$

The values of η are also bounded by -2.4 and +2.4 and are equally spaced with increments $\Delta\eta$. There are N_η points in the η dimension of the grid. Call $\eta_j = -2.4 + (j-1)\Delta\eta$. In the discretization,

$$p(\eta_j, \eta_k) = \Pr[\eta_{4(t-1)+i+1} = \eta_k | \eta_{4(t-1)+i} = \eta_j] = N(e_1) - N(e_2) \qquad (13)$$

where

$$e_1 = .5(\eta_{k+1} + \eta_k) - \rho\eta_j$$

[13] See Deaton and Laroque (1995, 1996) for a similar discretization of an AR1 process.

and

$$e_2 = .5(\eta_k + \eta_{k-1}) - \rho\eta_j$$

That is, the probability that the demand shock will equal η_k given that its previous value was η_j is equal to the probability that an i.i.d. standard normal variate will fall in an interval of length $\Delta\eta$ centered on $\eta_k - \rho\eta_j$.

For each value of Z and η in each of the four grids, values of carry-in x are spread equally along an interval $[0, x^i_{max}]$. There is a different value of x^i_{max} for each season i to reflect the fact that carry-in should be larger in early seasons (e.g., season 1) than later seasons (e.g., season 4). The x^i_{max} are determined in a trial and error process that ensures that equilibrium storage in season i never exceeds x^i_{max} in long simulation runs.

In the numerical analysis, demand functions are linear in consumption, demand shock, and expected future price. For $i = 1, 2, 3$,

$$D_i(x - s, Z, \eta) = a_i - b_i(x - s) + \sigma_\eta\eta + c_i\bar{P}_{i+1} \tag{14}$$

where s is the amount stored, x is carry-in, $\bar{P}_{i+1}$ is next period's expected price, and $a_i > 0$, $b_i > 0$, c_i and $\sigma_\eta > 0$ are parameters; $c_i > 0$ indicates that current and future consumption are substitutes. Note that $\bar{P}_{i+1}$ is a function of s and the state variables Z and η:

$$\bar{P}_{i+1} = \sum_{k=1}^{N_z}\sum_{j=1}^{N_\eta} \pi_i(Z, Z_k)p(\eta, \eta_j)P_{i+1}(s, Z_k, \eta_j).$$

For $i = 4$:

$$D_4(x - s, Z, \eta) = a_4 - b_4[\bar{H} + \sigma_z Z + x - s] + \sigma_\eta\eta + c_4\bar{P}_1 \tag{15}$$

and

$$\bar{P}_1 = \sum_{k=1}^{N_z}\sum_{j=1}^{N_\eta} \pi_4(Z, Z_k)p(\eta, \eta_j)P_1(s, Z_k, \eta_j).$$

Given the grids and demand functions, an initial guess for the functions P_i, $i = 1, \ldots, 4$ is formed in three steps. (There is one such function for each pair $\{Z, \eta\}$ on each season's grid.) First, for each value of x it is assumed that $s_1(x, Z, \eta) = 2x/3$; that $s_2(x, Z, \eta) = .5x$; that $s_3(x, Z, \eta) = x$; and $s_4(x, Z, \eta) = .25[x + \bar{H} + \sigma_z Z]$. That is, it is initially assumed that agents consume one-third of their carry-in in season 1, half their carry-in in season 2, all their carry-in in season 3, and one-quarter of carry-in plus production in season 4. Second, given this guess, the relevant equation from (13)-(14) is solved for s_i. This determines price for each of the points along the x_i axis of the grid. A fourth order polynomial in x_i is fit to these prices using OLS, and the resulting polynomial function $P_i(x, Z, \eta)$ is used as the initial guess for the pricing function for this $\{Z, \eta\}$ pair.

Given initial guesses for the price functions, for each Z, η, and x in each grid, in season $i = 1, 2, 3$ the following equation is solved (using Newton-Raphson) for s:

$$D_i(x - s, Z, \eta) = \beta \sum_{k=1}^{N_z} \sum_{j=1}^{N_\eta} \pi_i(Z, Z_k) p(\eta, \eta_j) P_{i+1}(s, Z_k, \eta_j) = \beta \bar{P}_{i+1}. \quad (16)$$

For $i = 4$

$$D_4(x - s, Z, \eta) = \beta \sum_{k=1}^{N_z} \sum_{j=1}^{N_\eta} \pi_4(Z, Z_k) p(\eta, \eta_j) P_1(s, Z_k, \eta_j) = \beta \bar{P}_1. \quad (17)$$

These equations equate the spot price to the discounted expected future price. If the s that solves the relevant equation is positive, the spot price at this grid point is set equal to $D_i(x - s, Z, \eta)$. If the s that solves the relevant equation is negative, then price at this grid point is set equal to $D_i(x, Z, \eta)$ because storage must be non-negative and thus consumption equals carry-in in this case. After determining prices on each node of the grid, for each $\{Z, \eta\}$ for each season a fourth order polynomial in x is fit to these prices using OLS. These new polynomials are used as the $P_i(x, Z, \eta)$ in the next iteration and the process is repeated. The process stops when the average absolute percentage price change between iterations is small (e.g., .001 percent).

Upon convergence to a fixed point, this process defines four spot price functions $P_i^*(x, Z, \eta)$. Using these functions it is possible to study commodity price dynamics and correlations between spot and forward prices by using them to approximate the partial derivatives that determine the correlation in (5). For example, in season 1, $\partial P_1^* / \partial Z$ is used to estimate P_Z during the first season. Similarly, in season 2 one can use:

$$F_{2,3} = \sum_{k=1}^{N_z} \sum_{j=1}^{N_\eta} \pi_3(Z, Z_k) p(\eta, \eta_j) P_3(s, Z_k, \eta_j) \quad (18)$$

to approximate the forward price for delivery in season 3 as of season 2. Again, $\partial F_{2,3} / \partial Z$ can be used as the estimate of F_Z in (4) or (5).

Consistent with the discretization, I use finite difference approximations to calculate the necessary partial derivatives. For $N_Z > j > 1$,

$$\frac{\partial P_i^*(x, Z_j, \eta_k)}{\partial Z} \approx P_Z \approx \frac{P_i^*(x, Z_{j+1}, \eta_k) - P_i^*(x, Z_{j-1}, \eta_k)}{2\Delta Z}$$

For $j = 1$,

$$\frac{\partial P_i^*(x, Z_j, \eta_k)}{\partial Z} \approx P_Z \approx \frac{P_i^*(x, Z_2, \eta_k) - P_i^*(x, Z_1, \eta_k)}{\Delta Z}$$

and for $j = N_Z$

$$\frac{\partial P_i^*(x, Z_j, \eta_k)}{\partial Z} \approx P_Z \approx \frac{P_i^*(x, Z_{N_z}, \eta_k) - P_i^*(x, Z_{N_z-1}, \eta_k)}{\Delta Z}$$

The partial derivatives with respect to η and for the F_Z and F_η are similarly calculated.

3 Results and Testable Implications

The foregoing model is solved for a variety of different parameters. The solutions tell a consistent story regardless of the parameters, so hereafter I present results for some representative cases. In the first case, $a_i = 60$, $b_i = .5$, for $i = 1, \ldots, 4$, $\sigma_z = 40$, $\sigma_\eta = 2.25$, and $\rho = .9$. The high value of ρ indicates that demand is highly autocorrelated; Deaton and Laroque (1996) require values of net demand autocorrelation of approximately this magnitude to fit the observed autocorrelations in annual commodity prices. In the first case, $c_i = 0$. That is, there is no intertemporal substitution. For this case, only storage and demand autocorrelation can induce relations between old crop and new crop prices. In the second case, $c_i = .9$ and $\rho = 0$; all other parameters are identical to those used in the first case. Thus, in the second case demand autocorrelation cannot explain any correlation between new crop and old crop prices; only storage and intertemporal substitution can connect these prices.

For $c_i = 0$, correlation between the spot price in season three and the forward price for delivery in season four (the harvest period) implied by the solution to the programming problem and (5) ranges from between .3 and .5 for values of carry-in observed with some frequency in long simulations; the correlation depends on η and Z, with the highest correlation observed when these shocks are both equal to their unconditional expected values of zero. Correlations are near 1 when the amount carried into season 3 is immense. Simulation runs of 5000 periods indicate that the carry-in required to generate these large correlations occur less than .3 percent of the time. Therefore, the carry-in required to generate spot-forward correlations that exceed .5 by a sizable margin is so large that it would almost never be observed in practice in this model economy. This implies that although high positive correlations are possible in a model with high demand autocorrelation but no intertemporal substitution, they are extremely rare.

This result is readily explained. It is almost always uneconomic to carry inventory from season 3, the summer, to season 4, the harvest period; only when period 3 carry-in is huge (which occurs only if the previous harvest was huge), the new harvest is expected to be very small, and demand is low will it prove optimal to carry inventory into the harvest period. In simulations, inventory is carried from season 3 to season 4 only about .3 percent of the time. Since inventory is not typically carried across crop years, arbitrage and storage do not link old crop and new crop prices. That is, $P_3 > \beta EP_4$ almost always. This implies that old crop prices should be little affected by news about the harvest, i.e., Z shocks, but that new crop prices will respond to these shocks. This tends to reduce the correlation between old crop and new crop prices. Indeed, an increase in σ_z increases the sensitivity of new crop prices to harvest shocks. Thus, if shocks to output are an important source of uncertainty in the market (as is plausible for agricultural commodities) correlations between new crop and old crop prices will be low. Only if demand shocks predominate and supply shocks are unimportant does high demand autocorrelation suffice to create high correlations between old crop and new crop prices. This is implausible for agricultural commodities and inconsistent with the fact that announcements about the expected future harvest have significant effects on both new crop and old crop prices, as will be documented in section 4.

In terms of expression (5), $P_Z \approx 0$ implies:

$$corr(\Delta P_t, \Delta F_t) \approx \frac{F_\eta}{\sqrt{(F_Z^2 + F_\eta^2)}} \tag{19}$$

where P is an old crop futures price and F is a new crop futures price.[14] If

[14] Hereafter, P will refer exclusively to an old crop spot price and F will refer to a new crop futures price.

supply shocks have larger effects on season 4 prices than demand shocks then $F_Z > F_\eta$, implying that this correlation is less than .5.

It is possible to generate correlations above .5 by increasing the importance of demand shocks relative to the importance of supply shocks, i.e., by increasing σ_η and decreasing σ_Z.[15] Even though this adjustment does affect correlations, it does not change a crucial implication of the model. Specifically, without interemporal substitution, old crop futures prices do not respond appreciably to changes in the size of the expected harvest regardless of the relative values of σ_z and σ_η. That is, P_Z, the partial derivative of the old crop spot price with respect to a shock to expected harvest, is nearly zero unless carry-in is far above its normal observed range regardless of the relative magnitudes of σ_Z and σ_η.[16] This is true because the absence of storage over the harvest implies that storage cannot transmit output shocks to old crop prices. Thus, the storage model with autocorrelated demand but no intertemporal substitution implies that old crop futures prices will not respond to shocks to the size of the expected harvest.

The values of the slope coefficient in a regression of the spot price change against the futures price change generated by this model lie in the .3 to .5 range for this version of the model. The coefficients for values of carry-in observed with some frequency in long simulations are closer to the lower end of this range.

Allowing for intertemporal substitution leads to substantially different price behavior. In this case, correlations between spot and forward prices are far above zero even for small values of carry-in. Given the assumed parameters, the observed correlation between the spot price and new crop futures price during season 3 ranges between .70 to .95 for a range of values of carry-in within 2 standard deviations of mean carry-in calculated from a simulation of the behavior of this economy over 5000 periods. As will be seen in section 4, these correlations are of the magnitude typically observed in the data.

The intertemporal substitution model also implies that spot prices re-

[15] The appendix shows that a large variance of the demand shock relative to the variance of the supply shock is also required to rationalize the findings of Deaton and Laroque. This large ratio is implausible for seasonally produced agricultural commodities.

[16] When $c_i = 0$ estimated partial derivatives P_Z for different values of carry-in, Z, and η vary tightly around zero regardless of the values for σ_z and σ_η. It is likely that any deviation from zero is due to the fact that price surfaces and partial derivatives are approximated numerically.

spond to shocks to expected harvest (i.e., Z shocks), although the magnitude implied by the model with $c_i = .9$ is somewhat smaller than that found in empirical data reported in section 4. The values of P_Z/F_Z range between .75 and .85. At carry-in equal to its median value in long simulations, a harvest shock that causes a one unit change in the forward price leads to a change of about .8 units (in the same direction) in the spot price.

These values for P_Z/F_Z far larger than those implied by the model without intertemporal substitution, but are still less than 1.00. It is possible to generate $P_Z/F_Z = 1$ only if $c_i = 1$, that is, if this quarter's consumption and next quarter's consumption are perfect substitutes. In this case, consumption in any quarter does not change when expected futures prices change.

The values of the slope coefficient in a regression of the spot price change against the futures price change generated by this model lie in the .65 to .85 range for the model with intertemporal substitution. The coefficient at the median carry-in is .75.

The values for correlations and slope coefficients for the intertemporal substitution model are similar to those found in empirical data that will be examined in section 4. This is in contrast to the correlation and P_Z/F_Z values generated by the high demand autocorrelation-no intertemporal substitution model.

Repeated experimentation indicates that three conditions are required to produce (1) correlations in the .7-.9 range, (2) large relative price sensitivities to harvest shocks and (3) slope coefficients that are close to 1.00. First, c_i must be close 1.00. Second, ρ must be relatively small but positive. Third, the variability of demand shocks must be sufficiently high. As an example, $c_i = .9$, $\rho = .3$, $\sigma_\eta = 9$ and $\sigma_z = 40$ generate spot-futures price relations that are similar to those documented in section 4 below. With these parameter values, the spot price-new crop futures price correlation during season 3 ranges between .75 and .85 for values of carry-in observed with some frequency in long simulations. The P_Z/F_Z ratio is approximately .85, and the slope coefficient ranges between .95 and 1.05.

The reasons these conditions must hold are straightforward. Large values of c_i are required to generate values of P_Z/F_Z of the magnitude implied by regressions of old crop price changes on new crop price changes when information about the harvest is released.

If the c_i are close to 1.00 and demand autocorrelation ρ is also large, however, new crop and old crop futures prices are almost perfectly correlated. In this case, since a demand shock that occurs in period 3 is expected to persist to period 4, the magnitude of the response of both old crop and new crop prices to this demand shock is similar. The high degree of intertemporal substitution causes both new crop and old crop prices to respond similarly to supply shocks. In terms of expression (5), $F_\eta \approx P_\eta$ when ρ is close to 1.00. Since $F_Z \approx P_Z$ when there is a high degree of intertemporal substitution,

$$corr(\Delta P_t, \Delta F_t) \approx \frac{P_Z^2 + P_\eta^2}{\sqrt{(P_Z^2 + P_\eta^2)^2}} = 1$$

This implies that to produce correlations clustering falling in the range .7-.9 there must exist some shock that influences old crop prices but not new crop prices. A relatively transient demand disturbance can have this effect; since there is typically no carryover across crop years, a demand shock that exhibits relatively little persistence which occurs in the old crop year has a smaller effect on the futures price than the spot price. With such a demand shock, $F_\eta < P_\eta$, which by (5) allows a correlation less than 1.00 when $F_Z \approx P_Z$.

Moreover, this demand shock must have a variance that is sufficiently high to cause divergences between movements in new crop prices and old crop prices that are large enough and frequent enough to overcome the fact that harvest shocks cause new crop and old prices to covary closely when current and future consumption are almost perfect substitutes. For a given value of σ_Z, if σ_η is too small, P_η and F_η are very small relative to P_Z and F_Z. Again, by (5) this implies that the spot-futures correlation is close to 1.00 in this case. The variance of the demand shock also affects the slope coefficient. Given a spot-forward correlation of around .8, to generate a slope coefficient of approximately 1.00 the variance of the spot price must exceed the variance of the new crop futures price during season 3. A high variance demand shock that exhibits little persistence can induce this disparity in variances.

In brief, examination of an augmented storage model that allows for seasonality in production and intertemporal substitution generates testable implications about co-movements in old crop and new crop futures prices for seasonally produced commodities such as grains. Specifically, in the absence of intertemporal substitution, (1) the correlation between old crop and new crop futures prices should be low (.5 or lower) if supply shocks are an important source of uncertainty, and (2) old crop futures prices should exhibit virtually no reaction to changes in the size of the expected harvest regardless of the relative importance of supply and demand shocks. That is, new crop futures prices should respond far more to shocks to the size of the expected harvest than old crop prices. In contrast, with strong intertemporal substitution the model predicts high correlations between old crop and new crop futures prices and nearly equal responses of old crop and new crop futures prices to supply shocks. Moreover, it is possible to find parameter values such that on average new crop and old crop futures prices change by the same amount.

The following section examines futures price data for seasonally produced commodities to test these predictions.

4 Tests of the Storage Model Using Old Crop and New Crop Futures Prices

The foregoing theoretical analysis demonstrates that explicit analysis of the implications of seasonality on the co-movements of new crop and old crop futures prices generates new insights and empirical implications from the standard storage model. Moreover, this approach makes it possible to draw upon a new source of data–futures prices–to test storage models. This section exploits this new data to test the implications of storage theory.

A few words about this data are in order. At any given time of year, futures contracts for multiple delivery dates are traded on a wide variety of seasonally produced commodities. These include wheat (CBOT, MGE, KCBOT), corn (CBOT), oats (CBOT), soybeans (CBOT), canola (WCE), and cotton (NYCE).[17] Of particular importance is the fact that there is simultaneous trading of futures contracts that expire prior to and after the harvest of each commodity. For example, during April of each year the CBOT trades corn futures contracts for both May delivery and December delivery.

Viewing futures prices as expected spot prices, the new crop futures prices represent the market's expectation of what the spot price for the commodity will be at harvest time.[18] Therefore, correlations between old crop futures

[17]CBOT is the Chicago Board of Trade; KCBOT is the Kansas City Board of Trade; MGE is the Minneapolis Grain Exchange; NYCE is the New York Cotton Exchange; WCE is the Winnipeg Commodity Exchange.

[18]Futures prices may embed a risk premium in addition to an expected spot price. Heretofore, much effort has been expended to identify risk premia in commodity futures prices with little success. I will therefore treat futures prices as expected spot prices. Even if futures prices do embed a risk premium, as long as this risk premium is relatively stable

prices and new crop futures prices measure the intertemporal linkage between prices across crop years. For example, the correlation measured during April between the daily change in the May corn futures price and the December corn futures price should be high (low) when prices are highly (not highly) correlated across crop years.

Table 1 presents correlations between old crop and new crop futures price changes from the 1981-1996 period for eight seasonally produced commodities: canola, corn, cotton, oats, soybeans, hard red winter (HRW) wheat, hard spring (HS) wheat, and soft red winter (SRW) wheat. The correlations are calculated as follows. For each commodity, a two-month period ending approximately 5 months prior to the harvest is selected. As an example, the period for corn is April and May. For each day t in the sample period, the daily (market close-to-market close) change in the old crop (May) futures price is calculated. This is $\Delta F_{o,t} = F_{o,t} - F_{o,t-1}$. The daily change in the new crop (December) futures price, $\Delta F_{n,t} = F_{n,t} - F_{n,t-1}$ is similarly calculated. The correlation between these price changes is then estimated. For each commodity, a correlation is estimated from a sample that pools daily price changes from the two month period within that single year.[19]

For each commodity, panel 1 of the table reports the portion from each year in which the correlation is measured, the old crop expiration month used, and the new crop expiration month used. Panel 2 reports the correlations by commodity and by year.

The correlations for individual years reported in Table 1 are generally around .7 or .8, and for every commodity studied there is a year in which the correlation exceeds .9. Correlations almost always exceed .3. Only for cotton does one observe a correlation below .25 (in a single year).[20] Similar results obtain for different choices of the old crop contract month and the time period in which the correlation is estimated. For example, similar correlations are obtained when using the changes in the September futures price during June

the price changes used to calculate correlations below will primarily depend on changes in expected spot prices. In this case all implications of the analysis in the text follow.

[19] These results are not sensitive to the choice of period in which the correlation is estimated or the old crop contract month. For example, the correlation between the September and December corn futures contracts during July-August is similar to the May-December correlation reported in Table 1.

[20] The low correlation for canola in 1994 is likely due to a squeeze of the June contract. A squeeze reduces the correlation between nearby and deferred futures prices. In essence, a squeeze is an extreme, highly transient, demand shock.

through August as the old crop contract for corn or soybeans.

The new crop-old crop futures correlation results parallel those reported for annual data in Deaton and Laroque (1992, 1995, 1996). They find what they consider to be high price autocorrelations in annual data; the magnitude of their correlation estimates are similar to the magnitude of the correlations reported here for daily futures prices. Deaton and Laroque suggest that such high correlations can be explained only if there is considerable autocorrelation in the underlying demand and supply processes. Since the past harvest is fixed in the futures data, however, autocorrelations in output cannot explain high correlations between new crop and old crop futures prices.

These futures price correlations are inconsistent with the implications of the storage model without intertemporal substitution analyzed earlier. That model implies low correlations for seasonally produced goods if supply shocks are important sources of uncertainty, as they clearly are for the commodities studied. These results imply that high autocorrelation in demand is not sufficient to explain high old crop-new crop futures price correlations even in the presence of storage. Instead, some other intertemporal linkage–such as intertemporal substitution–is required to explain the persistence in price shocks for seasonally produced commodities.

Evidence on the relative responses of old crop and new crop futures prices to shocks to the expected harvest provides even stronger evidence that is inconsistent with the basic storage model with autocorrelated demand. To quantify this relation, I examine futures price movements on days when the government discloses expected and actual acreage planted for each crop or issues Crop Production Reports that forecast expected harvests for corn, soybeans, and spring wheat. The acreage and production reports for these crops are released in June, July, August, and September. USDA forecasts are widely considered to be the most reliable available predictors of the coming harvest. Examining new crop-old crop price relations for these days is intended to isolate the effect of supply information from the effect of demand information. If supply information influences old crop prices, one would expect a positive, significant relation between new crop and old crop price changes on these dates. Moreover, one would expect that old crop futures prices are more volatile on crop announcement days than on days lacking any release if crop report information affects old crop futures prices. However, if information related to harvest size does not affect old crop prices (as the no-substitution model implies), the old crop futurs prices should be no more volatile on these days than other days.

Table 2 presents slope coefficients from regressions of the change in old crop futures prices on the change in the new crop price for corn, soybeans, and hard spring wheat for days on which the United States Department of Agriculture releases reports that disclose the forecasts for the next harvest. For each commodity, the September price is used as the old crop price; choice of the September contract allows use of announcements from all four reporting months. For corn and spring wheat, the December price is the new crop price. For soybeans, the November price is the new crop price. For purposes of comparison, the table also reports slope coefficients from regressions estimated from samples including all days in June-September that the USDA issued neither an acreage nor a crop production report.[21]

Since prior to 1995 the USDA released reports immediately following the close of futures trading on the release date, for 1981-1995 the difference between the futures price at the opening of trading on the day following the release date and the futures price at the close of trading on the release date measures the post-report price change. In 1996, the USDA began to release reports prior to the opening of trading. As a consequence I use the close-to-close price change to measure the post-release price responses for this year.

Note that the coefficients in all regressions are positive and significant, with p-values very close to 0. Indeed, for all three commodities, these coefficients are almost exactly one, indicating that old crop prices and new crop prices respond almost identically to news about the harvest contained in the USDA reports. One cannot reject the hypothesis that the slope coefficients for the announcement and non-announcement days are the same at the one-percent confidence level. Moreover, the R^2's imply new crop and old crop prices for corn and soybeans are very highly correlated on these days. It is interesting to note that for all three commodities, old crop and new crop prices are more highly correlated on USDA announcement days than on other days, as indicated by the higher R^2's on announcement days.

If one interprets the slope coefficient in the announcement day regression as a measure of P_Z/F_Z, these relative price responses are actually somewhat larger than those simulated in section 3 even with a very high coefficient of intertemporal substitution (i.e., a high c_i). Thus, although the storage

[21]Since information relating to expected harvest is revealed on these other days during the growing period (through weather reports, for instance) the regression using the no-announcement sample should also reveal some information about the relative responsiveness of old and new crop prices to output-related shocks.

model with intertemporal substitution comes far closer to mimicking the behavior of actual old crop and new crop futures prices than the model without substitution, one could interpret the results in Table 2 to mean that actual old crop prices respond more to shocks about the size of the expected harvest than the model prices.[22]

The size of the slope coefficients on the non-announcement days in Table 2 are two-to-three times larger than the slope coefficients generated by the no-intertemporal substitution model with large demand autocorrelation. The slope coefficients in the regressions are approximately equal to those presented in section 3 when $c_i = .9$, $\rho = .3$, $\sigma_\eta = 9$ and $\sigma_Z = 40$ were assumed. Thus, the intertemporal substitution model can mimic the real data for these three commodities whereas I have been unable to find parameters such that no-substitution model does so.

The variances of old crop and new crop price changes on USDA report release dates also indicate that information about the size of the harvest has a significant effect on old crop prices. Table 3 reports the variances of corn, soybean, and spring wheat futures price changes for both old crop and new crop futures for days on which the USDA announces acreage or crop forecasts and on all other days in June through September in the 1981-1996 period. Note that the variances are uniformly higher on crop forecast release dates. F-tests indicate that the variances are different at the .995 confidence level. This implies that the flow of information relevant to price determination is higher on these dates than on other dates for both new crop and old crop prices. Indeed, the variance ratios are very similar for new crop and old crop futures prices. This suggests that harvest-related information is of similar importance in determining both old crop and new crop prices. This provides further evidence that information about the size of the new crop is an important determinant of old crop futures prices. This is inconsistent with the implications of the high demand autocorrelation-no intertemporal substitution version of the storage model.

The foregoing evidence demonstrates clearly that (1) information about the harvest tends to cause old crop and new crop prices to move in the same

[22]This finding could result from the fact that demand shocks also occur on harvest announcement days. The slope coefficient in the announcement day regression provides an uncontaminated measure of P_Z/F_Z only if there is no demand related information released on the announcement day. Therefore, the near unity of the slope coefficient is not necessarily inconsistent with the finding of section 3 that $P_Z/F_Z < 1$ in the numerically solved storage model.

direction and by about the same amount, and (2) information about the coming harvest influences old crop futures prices. These results are inconsistent with the implications of the basic storage model even when there is substantial demand autocorrelation. In contrast, they are consistent with the patterns simulated in the model with intertemporal substitution. Thus, these findings strongly reject the hypothesis that autocorrelated demand explains the close co-movements in commodity spot prices and old crop and new crop futures prices. They therefore imply that some other linkage between prices among seasons is required to generate the observed price behavior. These preliminary results suggest that intertemporal substitution is a plausible candidate for this linkage.

5 Summary and Conclusions

The high degree of autocorrelation in commodity price time series has defied explanation by the standard storage model. Storage alone induces some autocorrelation into commodity prices, but not nearly as much as observed in actual data. One proposed solution to this problem has been to posit highly autocorrelated demand shocks. Existing empirical evidence (Deaton and Laroque, 1992, 1995, 1996) shows that a very high demand autocorrelation is required to generate the autocorrelations observed in actual commodity prices.

Although high demand autocorrelation can reconcile the standard model's predictions with the behavior of low frequency (e.g., annual) commodity price data, this article demonstrates that high demand autocorrelation cannot address the fundamental failings of the standard storage model in high frequency futures price data. Specifically, high demand autocorrelation in the standard model explains neither the high correlations between old crop and new crop futures prices for seasonally produced commodities nor the responsiveness of old crop futures prices to information about the expected size of the harvest.

In contrast, a modified storage model can explain these salient features of futures prices. Specifically, introducing intertemporal substitution in consumption to the standard model rectifies its failure to explain the behavior of futures prices at daily frequencies. A storage model with intertemporal substitution can produce old crop-new crop correlations and old crop price responses to harvest shocks of the magnitude observed in actual data for a variety of seasonally produced commodities. Thus, intertemporal substitution, rather than demand autocorrelation, may be the missing link between commodity prices over time.

The intuition behind these various results is straightforward. It is typically inefficient to carry inventories of a seasonally produced commodity from before the harvest to the harvest period. Thus, in normal circumstances storage cannot link old crop and new crop futures prices. This implies that news about the expected harvest should have little effect on old crop futures prices in the absence of intertemporal substitution. Intertemporal substitution provides a channel by which shocks to the expected harvest influence old crop prices. Intertemporal substitution therefore induces additional correlation between old crop and new crop prices. It also causes news about the expected harvest to influence old crop prices. Thus, storage plus intertemporal substitution can explain salient features of commodity prices at both high and low frequencies. Storage likely plays a secondary role in linking new crop and old crop futures prices and linking prices over time. This article suggests that intertemporal substitution plays the primary role in forging these links.

The approach taken in this paper to resolve the empirical deficiencies of commodity storage models of price dynamics differs from the efforts of Miranda and Rui (1999) to achieve the same objective. Miranda-Rui rely on the concept of the convenience yield in their attempt to reconcile the predictions of the storage model with the data. Although they find that the convenience yield model can generate price autocorrelations similar to those found in practice for a wide variety of commodities, their work examines prices at low (annual) frequencies. Moreover, they do not examine the implications of the convenience yield for the behavior of futures prices for seasonal commodities. Indeed, given the nature of convenience yield hypothesis it is difficult to see how it can generate such implications. That hypothesis lacks any rigorous microfoundations. In contrast, the explanation advanced herein advances a credible microeconomics-based explanation of the behavior of commodity spot and futures prices.

Future research should concentrate on estimating econometric models of commodity prices that incorporate intertemporal substitution as well as storage, although it must be recognized that such an endeavor faces a variety of serious challenges. First, the implication from the simple storage model that there is some "cutoff" price above which storage is zero and below which storage is positive does not hold in this seasonal model with autocorrelated demand and supply shocks. Thus, limited information methods such as GMM (along the lines of Chambers-Bailey (1996) and Deaton-Laroque (1992)) cannot be used to test the implications of the seasonal model with supply shocks. Second, whereas the simple storage model with autocorrelated net demand shocks requires estimation of a model that involves only a single pricing functional and two latent state variables, the seasonal model requires estimation of several pricing functionals with three latent state variables. The technical difficulties involved with the simpler model are trying enough (see Deaton-Laroque, 1995); those confronting the estimation of the structural parameters in the seasonal model are substantially greater.

Nonetheless, the evidence presented in this article does suffice to show that (1) the basic storage model augmented with autocorrelated demand shocks cannot explain key features of the behavior of old crop and new crop futures prices, and (2) intertemporal substitution can explain these same features. Thus, the theory and evidence presented herein advance our understanding of the forces driving commodity prices and provide a foundation for further research on this subject. In particular, the theory and evidence demonstrate that researchers can use high frequency futures price data to formulate and test sharper hypotheses that can discriminate between models in ways that are impossible when relying on low frequency spot price data alone.

Table 1
Panel 1
New Crop-Old Crop Futures Contract Information

COMMODITY	SYMBOL	OLD CROP MONTH	NEW CROP MONTH	PERIOD
CANOLA	CA	JUNE	NOVEMBER	MAY-JUNE
CORN	CN	MAY	DECEMBER	APRIL-MAY
COTTON	CT	MAY	OCTOBER	APRIL-MAY
OATS	OA	MAY	DECEMBER	APRIL-MAY
SOYBEANS	SY	MAY	NOVEMBER	APRIL-MAY
HRW WHEAT	HRWW	MARCH	JULY	FEBRUARY-MARCH
SRW WHEAT	SRWW	MARCH	JULY	FEBRUARY-MARCH
HS WHEAT	HSW	MAY	DECEMBER	APRIL-MAY

Table 1
Panel 2
New Crop-Old Crop Correlations

YEAR	CA	CN	CT	OA	SY	SRWW	HRWW	HSW
1981	.8756	.6481	.6580	.6310	.9297	.5847	.9455	.7835
1982	.9428	.8586	.6867	.8242	.9055	.8999	.7374	.8594
1983	.7324	.7099	.8839	.8517	.9787	.9444	.6950	.6519
1984	.5023	.7603	.7755	.6257	.6542	.8221	.7568	.7842
1985	.6799	.3736	.5816	.5985	.8142	.7059	.6200	.7219
1986	.9755	.5920	.2281	.9311	.8340	.5510	.4192	.8074
1987	.9476	.8281	.8253	.5546	.9175	.7334	.7321	.8565
1988	.9810	.8600	.5722	.8419	.9174	.8843	.8543	.8302
1989	.9815	.7526	.7267	.9402	.8213	.6844	.6484	.8659
1990	.9042	.7482	.7203	.9713	.9625	.5517	.6933	.6451
1991	.7739	.8600	.7089	.9380	.9540	.9370	.9531	.9191
1992	.9712	.8770	.9102	.9583	.9182	.8007	.8270	.8399
1993	.6233	.9085	.9211	.9139	.9281	.6565	.7053	.4875
1994	.2888	.8273	.7215	.9450	.9438	.5896	.7131	.5747
1995	.8318	.8880	.5838	.9272	.9715	.7605	.6638	.5302
1996	.9005	.4638	.6913	.8364	.9146	.4356	.6886	.6973
1997	.6301	.6773	.6297	.9054	.6831	.4589	.6456	.9020

Table 2
Slope Coefficients in Regression of
Old Crop Futures Price Change on New Crop Futures Price Change

Commodity	Announcement Day	Slope Coefficient	R^2
Corn	Yes	1.00	.863
Corn	No	.922	.832
Soybeans	Yes	.998	.979
Soybeans	No	1.007	.935
Wheat	Yes	.974	.855
Wheat	No	.938	.895

Table 3
Old Crop and New Crop Futures Price Change Variances
On USDA Announcement Days and All Other Days
in June-September

Commodity	Announcement Day Variance	Other Day Variance	$F-$statistic
Old Crop CN	17.71	8.32	2.13
New Crop CN	15.28	8.12	1.88
Old Crop SY	111.52	62.32	1.79
New Crop SY	103.57	57.64	1.80
Old Crop HSW	9.10	7.76	1.19
New Crop HSW	10.18	7.32	1.39

Note: Variances in cents squared. All $F-$statistics significant at the .995 level for 73 and 1165 degrees of freedom.

A Supply and Demand Variances in the Net Demand Setup

Deaton and Laroque invoke a high net demand autocorrelation to explain the high autocorrelation in commodity prices. This appendix demonstrates that this explanation requires demand shock variances to be far larger than supply shock variances. This is not plausible for the commodities I study.

Virtually all articles in the storage literature specify a "net demand shock" that is the difference between a demand shock and a supply shock. Consider the net demand $X_t = \eta_t - Z_t$ where η_t is demand at t and Z_t is production at t. Scale the variables such that $E(\eta_t) = 0$ and $E(Z_t) = 0$. For non-tree crops (e.g., soybeans or corn) it is plausible that Z_t exhibits low serial correlation. For purposes of exposition, assume that harvest autocorrelation is zero, that is $E(Z_t Z_{t-1}) = 0$. It is also plausible that $E(Z_{t-1}\eta_t) \approx 0$. Under these assumptions, the autocorrelation of X_t is:

$$\rho_X = \frac{E(\eta_t \eta_{t-1})}{E(\eta_t^2) + E(Z_t^2)}$$

If the autocorrelation in the demand shock is ρ_η, this implies:

$$\rho_X = \frac{\rho_\eta E(\eta_t^2)}{E(\eta_t^2) + E(Z_t^2)}$$

Even if the demand shock follows a random walk (i.e., $\rho_\eta = 1$), this implies that the variance of the demand shock–$E(\eta_t^2)$–must be 9 times as large as the variance of the supply shock-$E(Z_t^2)$–to generate a $\rho_X = .9$ (the magnitude required to explain price autocorrelations in the Deaton-Laroque (1995, 1996) framework). Thus, if supply shocks are serially uncorrelated (or only weakly serially correlated), the Deaton-Laroque explanation of high price autocorrelations requires supply shocks to be unimportant. This is implausible for the field crops studied here; the attention of traders to weather information and supply forecasts clearly indicate that supply shocks are of crucial importance in determining crop futures prices. This further highlights the necessity of identifying some other source of intertemporal linkage.

References

[1] Chambers, Marcus, and Roy Bailey. 1996. "A Theory of Commodity Price Fluctuations." 104 *Journal of Political Economy*: 924-957.

[2] Deaton, Angus, and Guy Laroque. 1992. "On the Behavior of Commodity Prices." 59 *Review of Economic Studies*: 1-23.

[3] Deaton, Angus, and Guy Laroque. 1996. "Competitive Storage and Commodity Price dynamics." 104 *Journal of Political Economy*: 896-923.

[4] Deaton, Angus and Guy Laroque. 1995. "Estimating a Nonlinear Rational Expectations Model with Unobservable State Variables." 10 *Journal of Applied Econometrics*: S9-S40.

[5] Miranda, Mario and Xiongwen Rui. 1999. "An Empirical Reassessment of the Commodity Storage Model." Working paper, Ohio State University.

[6] Routledge, Brian, Duane Seppi, and Chester Spatt. 1997. "Equilibrium Forward Curves for Commodities." Working paper, Carnegie Mellon University.

[7] Scheinkman, Jose, and Jack Schectman. 1983. "A Simple Competitive Model of Production with Storage." 50 *Review of Economic Studies*: 427-441.

[8] Williams, Jeffrey, and Brian Wright. 1991. *Storage and Commodity Prices.* Cambridge: Cambridge University Press.