

Economics Department
Working Papers in Economics

Boston College

Year 1997

Explaining the Persistence of Commodity
Prices

Serena Ng
Boston College,

Francisco Ruge-Murcia
Universite de Montreal,

Explaining the Persistence of Commodity Prices

Serena Ng*

Francisco J. Ruge-Murcia[†]

December 1997

Abstract

This paper extends the Competitive Storage Model by incorporating prominent features of the production process and financial markets. A major limitation of this basic model is that it cannot successfully explain the degree of serial correlation observed in actual data. The proposed extensions build on the observation that in order to generate a high degree of price persistence, a model must incorporate features such that agents are willing to hold stocks more often than predicted by the basic model. We therefore allow unique characteristics of the production and trading mechanisms to provide the required incentives. Specifically, the proposed models introduce (i) gestation lags in production with heteroskedastic supply shocks, (ii) multiperiod forward contracts, and (iii) a convenience return to inventory holding. The rational expectations solutions for twelve commodities are numerically solved. Simulations are then employed to assess the effects of the above extensions on the time series properties of commodity prices. Results indicate that each of the features above partially account for the persistence and occasional spikes observed in actual data. Evidence is presented that the precautionary demand for stocks might play a substantial role in the dynamics of commodity prices.

Keywords: commodity prices, persistence, speculative storage.

JEL Classification: B4, G1, Q1.

*Department of Economics, Boston College.

[†]Department of Economics and C.R.D.E., University of Montreal.

We have benefited from helpful comments by seminar participants at Johns Hopkins University. Jean-Philippe Laforte provided excellent research assistance. Financial support by the Social Science and Humanities Research Council of Canada (SSHRC) and the Fonds pour la Formation de Chercheurs et l'Aide à la Recherche du Québec (FCAR) is gratefully acknowledged. Correspondence: Serena Ng, Department of Economics, Boston College, Chestnut Hill, MA 02167, USA. E-Mail: serena.ng@bc.edu. Francisco J. Ruge-Murcia, Department of Economics, University of Montreal, C.P. 6128, succursale Centre-ville, Montreal (Quebec) H3C 3J7, Canada. E-Mail: rugemurf@ere.umontreal.ca.

1 Introduction

This paper studies the determination of commodity prices in a setup where (in addition to the producer and consumer) there is a rational, profit maximizing agent that can carry the good as inventory from the current to future periods. This framework has been employed by earlier researchers [see among others, Newbery and Stiglitz (1982), Williams and Wright (1991), and Deaton and Laroque (1992)] to examine the effect of speculative inventories on commodity prices. The models developed in this article also attribute inventory holdings a prominent role in determining the time series properties of the price process, but, in addition, they incorporate (i) gestation lags in production with heteroskedastic shocks, (ii) multi-period, overlapping forward contracts, and (iii) a convenience motive for inventory holding. The intention is that by extending the basic storage model to include other realistic aspects of the production process and the trading mechanism, the resulting specification would better capture the most relevant characteristics of the data.

Earlier research on this topic has documented two important features in the time series of commodity prices. First, prices are subject to occasional, dramatic increases or "spikes". Deaton and Laroque (1992) model these spikes as arising from stockouts. That is, in the absence of the smoothing effect of inventories, the price is solely determined by the available harvest and consumption demand. Second, prices exhibit a high degree of serial correlation. Chambers and Bailey (1996) and Deaton and Laroque (1996) seek to explain this feature by assuming serially correlated shocks to production but find that there is still substantial persistence left unaccounted for.

The specifications proposed in this article preserve stockouts as a plausible explanation of price spikes but allow news about future production to have effects on current prices that could replicate these pronounced price increases. More importantly, they explicitly model features of the production process and financial markets that might (in theory) explain the persistence of commodity prices. Loosely speaking, these features (production lags, contracts, and the convenience return) share the property that they induce profit-maximizing agents to hold stocks more often. Since there are fewer periods where the intertemporal price relationship is severed by stockouts, the persistence predicted by these models might be higher than in the basic storage model.¹

The study of the factors that determine the price of primary commodities is important for several reasons. First, many Less-Developed-Countries depend on the export of a small set of agricultural products (sometimes one) for most of their foreign currency earnings. Second, several countries spend a significant amount of resources on the regulation of their agricultural sectors through price support mechanisms and regulation boards. A good understanding of the determinants of the price

¹The model developed by Deaton and Laroque (1992) yields a high probability of stockouts and, consequently, predicts a far larger number of spikes than observed in actual data.

process is necessary to assess the effects of these government policies. Finally, the price behavior of industrial commodities (*e.g.*, copper, iron ore, oil, etc.) can have important implications for output and business fluctuations.

As noted above, three extensions are examined. First, we consider “time to build” as a source of persistence in commodity prices. On one hand, gestation lags in production might increase the possibility of stockouts because there are periods in which no new harvest is brought to the market. On the other hand, profit-maximizing speculators might consider the seasonality of output as an incentive to transfer some of the good to the period without production. Therefore, the overall effect of production lags might be to reduce the probability of stockouts. Gestation lags provide a convenient and reasonable way of introducing heteroskedasticity in harvest shocks that, under certain conditions, can also be associated with seasonality in consumption. Finally, from the point of view of understanding the formation of price expectations, gestation lags are interesting because they allow news about the incoming crop to convey information about the value of next period’s output and to influence the price of the commodity currently traded [see Lowry *et al.* (1987)].

Second, we suggest the presence of forward, multiperiod contracts as a plausible explanation of the serial correlation observed in prices. The intuition draws on earlier work by Fischer (1977) and Taylor (1979, 1980) where staggered labor contracts yield a significant degree of serial correlation in nominal wages. Multi-period contracts increase the predicted price persistence because they provide additional sources of supply in every period and unambiguously reduce the probability of stockout. In addition, the overlapping nature of the contracts means that even if the weather shocks are serially uncorrelated, their effect on the commodity price last longer than the contract period.

Finally, we allow for a demand for inventories other than for speculative purposes. This is achieved by relaxing the assumption of zero convenience yield. The convenience yield [Kaldor (1939)] is a catch-all term for the return accruing to consumers and producers by being able to use the stored commodity whenever desired. Since the convenience yield could partially compensate inventory holders for the expected loss when the basis is below carrying charges, the model predicts a smaller number of stockouts and a larger degree of serial correlation in prices.

An important characteristic of basic storage models is that the demand for speculative inventory will be greater than zero if and only if the future price expected by optimizing speculators is high enough to cover carrying costs. This non-negativity constraint on inventories implies that the equilibrium price is no longer a linear function.² In particular, the non-linearity takes the form of a kink at the price at which inventory demand becomes positive. Since the extensions proposed in

²The idea that inventories are bounded below by zero was originally discussed by Samuelson (1957) and Gustafson (1958). Muth (1961) obtains a linear rational expectations equation only under the assumption that speculative inventories can be either positive or negative

this article preserve the non-linear structure of the price process, we employ numerical routines to establish the rational expectation solution of the models.

Since the econometric estimation of the models above is non-trivial, it would seem reasonable that as a preliminary step in this research program, the competing specifications be assessed in a unified framework. To that effect, the models are calibrated for a set of 12 commodities and their empirical properties compared with the ones of the actual data.

The plan of the article is as follows. Section 2 briefly revises the basic storage model and introduces the main concepts employed throughout the paper. The reader already familiar with this model and the numerical techniques employed to solve non-linear rational expectations models could skip this part. This section also examines a benchmark case where serially correlated disturbances are incorporated into the competitive storage model. Sections 3, 4, and 5 respectively introduce the specifications with gestation lags, forward contracts, and a convenience yield. Section 6 presents the main empirical results and evaluates the different extensions. Finally, Section 7 concludes and suggests avenues of future research.

2 The Basic Storage Model

The rational expectations competitive storage model describes the optimal inventory decision on the part of risk-neutral speculators. This model is now briefly summarized. Let harvest, z_t , be given by

$$z_t = \bar{z} + u_t,$$

where \bar{z} is a constant, and u_t is a disturbance term assumed *i.i.d.*($0, \sigma^2$) with time invariant distribution and compact support.³ The random disturbance u_t is most intuitively interpreted as a weather shock. Denote by $I_{t-1,t}$ the quantity held as inventory from period $t-1$ to period t . Let r be the real interest rate and δ be a non-negative depreciation cost. That is, the storage technology yields $(1-\delta)$ units of good in period t for each unit stored in period $t-1$. It is further assumed that (i) the convenience yield is zero so that there is no demand for inventories other than for speculative purposes, and (ii) inventories are costly to hold, and hence $(1-\delta)/(1+r) < 1$. With the above notation, the quantity of commodity available at time t (denoted by x_t) can be written as

$$x_t = z_t + (1-\delta)I_{t-1,t}. \tag{1}$$

³This specification implicitly assumes that the commodity supply is inelastic with respect to the expected price. This postulate is solely made for tractability and does not affect the basic implications of the model.

Let p_t be the price of the commodity at time t and $D(p_t)$ be a deterministic demand function with inverse demand function denoted $P(x_t)$. Then, market clearing requires that

$$x_t = z_t + (1 - \delta)I_{t-1,t} = D(p_t) + I_{t,t+1}.$$

The demand for one-period-ahead inventories, $I_{t,t+1}$, is the result of intertemporal arbitrage by risk-neutral speculators. Formally, let $E_t p_{t+1}$ be the expectation of p_{t+1} conditional on information available at time t . The information set is assumed to contain p_t, z_t , and the distribution of supply shocks. Arbitrage implies that [see Scheinkman and Schechtman (1983)],

$$\begin{aligned} I_{t,t+1} &\geq 0 && \text{if } (1 - \delta)/(1 + r)E_t p_{t+1} \geq p_t, \\ I_{t,t+1} &= 0 && \text{otherwise.} \end{aligned} \tag{2}$$

The interpretation of this First Order Condition is straightforward. Profit maximizing stockholders will demand positive inventories if and only if the expected future price is high enough to cover carrying costs. Otherwise, the inventory demand will be zero.

Deaton and Laroque (1992) demonstrate the existence of an equilibrium for this model and show that prices follow an ergodic process with a non-zero probability of being in the stockout regime in finite time. Since inventories serve as intertemporal link between two periods, prices follow a linear first-order autoregressive process in the stockholding regime with a conditional variance that increases as the level of stocks diminishes. However, when stocks are not held, the market price is history independent since it is completely determined by the contemporaneous harvest. Accordingly, prices in the stockout regime follow a white noise process. The overall price process is a non-linear first-order Markov process with a kink at the price

$$p^* = (1 - \delta)/(1 + r)E_t p_{t+1},$$

that is constant under the assumption that shocks to harvests are *i.i.d.*

Due to the non-negativity constraint in inventories, the price process is non-linear. Thus, it not feasible to find an explicit close-form solution for the agents' price forecast. To address this problem, Deaton and Laroque (1992) obtain numerical solutions by iterating over the equilibrium price function until convergence. Other methods for solving non-linear rational expectations model are discussed in Taylor and Uhlig (1990). This article employs the procedure proposed by Williams and Wright (1991) [see also Den Haan and Marcet (1990)] that involves parameterizing the conditional price expectation in the terms of the model state variables. More precisely, the agents' price forecast is expressed as a finite, low-order polynomial of the state variables. For the basic storage model, there is only one state variable, namely the current level of inventories. Therefore, one

could represent the agents' forecast of the future price as a polynomial in the level of stocks with an associated set of coefficients.⁴ According, we postulate the function

$$E_t p_{t+1} = \psi_1(I_{t,t+1}), \quad (3)$$

where $\psi_1(\cdot)$ denotes a finite polynomial. Presumably, $\psi_1(\cdot)$ is a decreasing and convex function of the level of inventories so that the larger the current level of stocks, the larger the availability in the future period and the lower the expected price. Given a set of structural parameters and an initial grid of inventories and harvest shocks, it is possible to employ Ordinary Least Squares to obtain estimates of the polynomial coefficients and the Newton's method to adjust these coefficients until the rational expectations solution obtains.⁵

The basic speculative storage model can accurately replicate some of the features of the data. Specifically, the simulated model features the periodic price spikes that are occasionally observed in the actual time series [see Deaton and Laroque (1992, p. 14)]. On the other hand, reduced form estimations [Ng (1996)] and evaluations of the model based on estimation of the structural parameters [Deaton and Laroque (1992)] all find that the persistence in price data is much higher than predicted by the theory. This paper seeks to explain the persistence of commodity prices by explicitly modeling realistic features the financial markets and the productive process. It is argued at the theoretical and empirical levels that the extensions of the basic storage model proposed below can partially explain the high serial correlation in prices.

2.1 Adding Serially Correlated Supply/Demand Shocks

Before proceeding to the economic extensions of the competitive storage model, we will consider a useful benchmark that entails the mechanical introduction of serially correlated shocks in the supply or demand of the commodity. Deaton and Laroque (1996) and Chambers and Bailey (1996) have also examined this possibility for the case when weather shocks follow an AR(1) process with limited success in accounting for the serial correlation in prices. This section complements and extends their results by considering a MA(1) specification for the model shocks. As shown below, using MA(1) disturbances is considerably more tractable than postulating AR(1) errors and, for invertible specifications, it captures the serial correlation of production/demand shocks in a more parsimonious way than a higher-order autoregressive process.

Under this specification, the (inelastic) harvest in period t is

$$z_t = \bar{z} + u_t,$$

⁴Judd (1990) rigorously defends the use of polynomials to approximate non-linear functions.

⁵Williams and Wright (1991, pp. 81-90) describe the algorithm in detail for the cases with and without supply elasticity.

where, as before, \bar{z} is a constant, but now the weather disturbance u_t is assumed to follow the MA(1) process

$$u_t = e_t + \rho e_{t-1},$$

where e_t is *i.i.d.*($0, \sigma^2$) with time invariant distribution and compact support. In this simple extension of the basic storage model, the market clearing condition and the (profit-maximizing) arbitrage condition are still given by (2.1) and (2.2), respectively. Thus, prices still follow the non-linear process described above for the competitive storage model. However, as in Chambers and Bailey (1996), the price kink (namely, p^*) is no longer constant, instead it is a function of the current observed harvest (or weather shock). This result arises because with serially correlated disturbances, speculators' form their price forecasts using the information now contained in the current production shock.

An implication of the above observation is that for the model with MA(1) errors, the set of state variables should be enlarged to include the current observation of e_t . Thus, for the purpose of parameterizing the speculators' conditional expectation, we postulate the function

$$E_t p_{t+1} = \psi_2(I_{t,t+1}, e_t), \tag{4}$$

where the coefficients for ψ_2 are to be determined numerically as described in the previous section.

The assumption that shocks are serially correlated is a very simple and natural way to introduce persistence into the model. For example, in the case of tree crops (like coffee) damage to the plants as a result of bad weather can reduce not only current, but also future, production. In the case of industrial commodities, serial correlation in demand might arise as a result of business cycle fluctuations. However, in general, it is difficult to identify the origin of the serial correlation of prices without further assumptions on the structure of the model. As we will see in the next three sections, the persistence in commodity prices can be replicated even if the shocks are *i.i.d.*, once we enrich the model with more realistic features of production and preferences.

3 Gestation Lags and Heteroskedastic Supply Shocks

The production of commodities takes time. Cash crops such as cotton and tea are not immediately harvested after planting, and the eventual quantity of output depends on the sequence of (weather) shocks between planting and harvest. Likewise, industrial commodities such as iron and copper require time to be mined and refined. While the model presented below better describes the case of agricultural crops that constitute most of our sample, its implications are easily generalized to industrial commodities.

Gestation lags have several important implications. First, there might be periods with no production. For example, for agricultural commodities this might be due to the fact that harvesting

takes place in a particular season(s). *Ceteris paribus*, this limits availability to the market and increases the probability of stockout. However, profit-maximizing speculators have an incentive to transfer some of the good (as inventory) to the period without production and could therefore decrease the probability of stockouts. Thus, it is conceivable that the net effect of gestations lags is to reduce the number of periods where the intertemporal price link is severed and, consequently, to increase the estimated serial correlation.

Second, the periodicity in production permits the straightforward introduction of heteroskedasticity in supply shocks. Intuitively, the variance of weather shocks might differ across months or seasons. Under certain conditions, this heteroskedasticity can be also interpreted as modeling seasonal effects in *consumption*. Specifically, in the case when the inverse demand function is of the linear form $P(x) = a + bx$, the parameter b cannot be separately identified from the variance of the harvest [see Deaton and Laroque (1996), pp. 905-906]. Since a researcher using price data can only identify the ratio b/σ , changes in σ across seasons would be observationally equivalent to changes in the slope of the demand function b across seasons.

Third, while suppliers cannot revise their production schedules once the crop has been planted, speculators can observe shocks in between harvests. This information allows agents to update their forecast about the size of the incoming harvest and demand inventories consistent with their revised expectations of the one-period-ahead price. Thus, the trading of commodities takes place on a continuous basis in spite of the periodicity in the production process.

Finally, to the extent that news are allowed to have an effect on the prices, a model with gestation lags might generate large price increases without stockouts in the contemporaneous period. For example, consider the situation when a "bad" weather shock has affected the current crop. The anticipation of a smaller harvest and a high price in the incoming period prompts rational agents to increase their current demand for inventories to take advantage of the *ex-ante* profit opportunity. Thus, the current price rises. For a large enough negative shock, the price increase could be significant enough to replicate the large price spikes observed in commodity data.⁶

To formalize this idea, it is assumed that a production cycle takes two periods. More precisely, the crop grows during the odd periods while the harvest and subsequent planting take place in the even periods.⁷ Accordingly, harvest (denoted as before by z_t) is affected by the shocks during the growing and harvesting seasons. Let \bar{z} be mean output and u_j denote the *i.i.d.* shock to supply in period j . Assume that the statistical distribution generating the observations of u has a variance

⁶A recent example of this alternative explanation of price "spikes" was the large increase in current and future prices of coffee in June of 1994 following unusually low temperatures in Southern Brazil [see *The Economist*, 2 July 1994, p. 96].

⁷Lowry *et al.* (1987) develop a quarterly model for soybeans with storage within and across crop years. In their model, the grain is harvested only in one of the quarters.

that might differ according to whether the period is odd (σ_1) or even (σ_2).⁸ Then, the harvest in (the even) period t is

$$z_t = \bar{z} + u_{t-1} + u_t,$$

and the availability is

$$x_t = z_t + (1 - \delta)I_{t-1,t}.$$

Comparable expressions hold for the other even periods, namely $t + 2, t + 4$, etc. In the odd periods, when there is no harvest, the only source of availability is previously accumulated stocks. Specifically in period $t + 1$,

$$x_{t+1} = (1 - \delta)I_{t,t+1},$$

with similar expressions holding for the other odd periods, $t + 3, t + 5$, etc. Market clearing in even and odd periods requires that

$$\begin{aligned} t: \quad x_t &= z_t + (1 - \delta)I_{t-1,t} = D(p_t) + I_{t,t+1} \\ t+1: \quad x_{t+1} &= (1 - \delta)I_{t,t+1} = D(p_{t+1}) + I_{t+1,t+2}, \end{aligned}$$

where $D(p_t)$ is a deterministic demand function. Speculative stockholding is now governed by a pair of arbitrage conditions:

$$\begin{aligned} I_{t,t+1} &\geq 0 \quad \text{if} \quad (1 - \delta)/(1 + r)E_t p_{t+1} = p_t, \\ I_{t+1,t+2} &\geq 0 \quad \text{if} \quad (1 - \delta)/(1 + r)E_{t+1} p_{t+2} = p_{t+1}. \end{aligned} \tag{5}$$

Notice that the arbitrage conditions (3) imply (as in the basic model) that prices follow a linear first-order autoregressive process with constant coefficient $(1 + r)/(1 - \delta)$ in the stockholding regime and a white noise process in the stockout regime. However, once heteroskedasticity is allowed in production shocks, the threshold price p^* is no longer unique. Instead, there are two distinct p^* s associated with the odd and even periods respectively [see Chambers and Bailey (1996, p. 938)].

While gestation lags do not fundamentally alter the non-linear structure of the price process, they might still increase their serial correlation because both the periodic absence of new production and the information revealed by weather shocks have the effect of inducing speculators to hold inventories and might reduce the probability of stockouts. Since the intertemporal price link might be severed less frequently, the predicted price persistence might be larger than implied by the basic storage model.⁹

⁸Chambers and Bailey (1996, p. 940) show the existence and uniqueness of the equilibrium price function with periodic disturbances.

⁹Notice that the periodic arrival of the harvest might reduce the agents incentive to carry inventories from the odd to the even periods. Thus, gestation lags might also reduce the serial correlation in prices. Which one of the two effects dominates is an empirical matter to be examined in Section 6.

Although the interpretation of the above arbitrage conditions is similar to that of the basic storage model, there is an important difference in the information available to agents to construct their price forecasts. In the basic storage model, speculators always observe current period output. In the model with gestation lags, speculators only observe the current supply shock in the odd periods and the level of output in the even periods. Moreover, the shocks (or news) in the odd period reveal information about the size of the harvest and the price in the incoming period. Since agents are rational and make use of all information available, the information revealed by harvest shocks in the odd period is exploited by agents when making their inventory-holding decision for next period. More formally, there are two state variables in the odd periods without harvest (the current level of inventories and the current supply shock) and only one state variable in the even period with harvest (the level of inventories.) Thus,

$$\begin{aligned} t: \quad E_t p_{t+1} &= \psi_3(I_{t,t+1}), \\ t+1: \quad E_{t+1} p_{t+2} &= \psi_4(I_{t+1,t+2}, u_{t+1}). \end{aligned} \tag{6}$$

The coefficients for ψ_3 and ψ_4 are to be determined numerically as described above.

4 Multiperiod Forward Contracts

One institutional feature of trading in commodity markets is that contracts with different holding period are available. For example, the Chicago Board of Trade lists forward contracts for coffee for delivery as short as three months and as long as one year. This section examines the empirical implications of multiperiod contracts on commodity prices. Earlier research in macroeconomics [*e.g.*, Fischer (1977) and Taylor (1979, 1980)] shows that in labor markets, overlapping wage contracts increase the degree of rigidity of the endogenous variable. Thus, one might expect that overlapping forward contracts in commodity markets might increase the degree of serial correlation in prices *vis a vis* the one predicted by the basic storage model (where this feature is absent.)

To model multiperiod contracts, let $I_{t+i,t+j}$ denote inventories carried in period $t+i$ for delivery in period $t+j$. The basic storage model is the special case where $j-i=1$. For tractability, we concentrate on contracts with a maximum holding period of two periods. Thus, $j-i=1$ for one period contracts, and $j-i=2$ for two period contracts. In this setup, the availability in each period consists of the contemporaneous production and the inventories contracted in periods $t-1$ and $t-2$ to be delivered in the current period, that is,

$$x_t = z_t + (1-\delta)I_{t-1,t} + (1-\delta)^2 I_{t-2,t}.$$

As in the simple storage model, the harvest is assumed to be described by $z_t = \bar{z} + u_t$, where u_t is an *i.i.d* disturbance with mean zero and constant variance. Let $D(p_t)$ be a deterministic demand

function and note that market clearing requires

$$z_t + (1 - \delta)I_{t-1,t} + (1 - \delta)^2 I_{t-2,t} = D(p_t) + I_{t,t+1} + I_{t,t+2}.$$

In this case speculative stockholding must satisfy a pair of arbitrage conditions that hold simultaneously in every period to determine the demand for one-period and two-period ahead inventories. Specifically,

$$\begin{aligned} I_{t,t+1} &\geq 0 && \text{if } (1 - \delta)/(1 + r)E_t p_{t+1} = p_t, \\ I_{t,t+2} &\geq 0 && \text{if } [(1 - \delta)/(1 + r)]^2 E_t p_{t+2} = p_t. \end{aligned} \tag{7}$$

As before, speculators will carry inventories if and only if the expected capital gain suffices to cover storage costs. It is assumed that speculators holding two-period inventories are contractually prevented from rolling over their inventories. That is, once a stock holder agrees to deliver $I_{t,t+2}$ units of the good in period $t + 2$, she must physically hold the stocks during the contracted period and, consequently, is unable to modify her stockholding in light of new information available after her decision was made.¹⁰ Thus, the presence of overlapping contracts means that, even if the weather shocks are serially uncorrelated, their effect on the commodity price last longer than the contract period.

The persistent effects of weather shocks can also be seen by noting that in this set up there are two different threshold values p^* . More precisely, speculators will hold one-period contracts whenever the current price is below $p_1^* = (1 - \delta)/(1 + r)E_t p_{t+1}$ and two-period contracts whenever the price is lower than $p_2^* = [(1 - \delta)/(1 + r)]^2 E_t p_{t+2}$. As will be shown in Section 6, the fact that contracts of different maturity are subject to different arbitrage conditions makes it possible to decompose the level of stocks between those held for one period and two periods.

In this model with multiperiod contracts, the number of state variables is substantially larger than in the models considered above. The agents' forecast of the future price is a function of the level of all types of inventories currently held by speculators, and in theory, the one and two period ahead expectations should be respectively parameterized as $E_t p_{t+1} = \psi_5(I_{t-1,t+1}, I_{t,t+1}, I_{t,t+2})$, and $E_t p_{t+2} = \psi_6(I_{t-1,t+1}, I_{t,t+1}, I_{t,t+2})$. Notice however, that from the perspective of agents, $I_{t,t+1}$ and $I_{t,t+2}$ are based on and exactly contain the same information.. Thus, the inclusion of both arguments during the OLS estimation of the expectations polynomial yields colinearity among these two regressors. Therefore, for the calculation of the expectations coefficients, the following parameterization was employed

$$E_t p_{t+1} = \psi_5'(I_{t-1,t+1}, I_{t,t+1}), \tag{8}$$

¹⁰The assumption is equivalent to the one in Fischer (1977) where two-period labor contracts cannot be renegotiated at the end of the first term.

and

$$E_t p_{t+2} = \psi'_6(I_{t-1,t+1}, I_{t,t+2}), \quad (9)$$

where the coefficients ψ'_5 and ψ'_6 satisfy the assumption of rational expectations and are numerically estimated for each commodity in the sample.

Multi-period contracts might increase the predicted persistence in prices because they provide additional sources of supply in every period and unambiguously reduce the probability of stockout. The effect of overlapping contracts on the time series properties of prices will depend on the agent's willingness to carry multi-period inventories. We will show in Section 6 that in contrast to earlier literature on storage where contracts are absent and stockholders are free to roll-over their inventories, a model with overlapping contracts can partially explain the high serial correlation in prices.

5 The Convenience Yield

The arbitrage condition implied by the various models above all postulate that the demand for inventories will be positive only when the basis (*i.e.*, the difference between the future expected price and the current spot price) is enough to cover storage costs. As a direct consequence of this, the models admit zero as the lower bound for inventories. However, earlier literature on storage has pointed out that in certain instances individuals and firms seem to be willing to carry stocks even when the difference between the future expected price and the current price is less than the cost of storage.¹¹ Three possible explanations have been advanced to account for this apparent negative return to storage. First, Vance (1946) suggests that agents might unduly discount the future relative to the present. To support this view, Vance cites empirical evidence collected by Vaile (1944) that indicates that future prices are consistently lower at the beginning than at the close of the trading period for each contract considered.

A second explanation (that also accounts for Vaile's observation) is that futures prices might be downwardly biased as a result of a risk premium. This explanation is related to the theory of normal backwardation proposed by Keynes (1930) and postulates that since production must be planned well ahead of consumption, hedgers as a whole have a tendency to go short in futures and, consequently, a positive risk premium is required to persuade speculators to take the corresponding long position. Subsequent empirical research has sought to detect the presence and evaluate the magnitude of the expected risk premium with inconclusive results [see Telser (1958), Cootner (1960, 1967), Dusak (1973), Breeden (1980), and Fama and French (1987)].

¹¹This observation is not universally shared. Williams and Wright (1991) dismiss the empirical observation of stockholding below full carrying charges as an aggregation phenomenon [see also, Williams (1986) and Wright and Williams (1989)].

A third explanation for inventory holding below carrying charges is proposed by Kaldor (1939) who suggests that the return of storage might include a component of "convenience", that is, the possibility of using the stored materials whenever desired. In the case of a producer, this return could arise because inventories (*i*) reduce the need to revise production schedules, (*ii*) allow the supplier to take advantage of unexpectedly higher prices to quickly meet demand without increasing production, (*iii*) diminish the probability of having to turn down buyers, and (*iv*) reduce replenishment costs and time delays in delivery. In the case of a consumer, a convenience yield could arise if either the timing or the level of consumption are stochastic. Thus, whether the stockholders are players in the supply side or the demand side of the market, there is likely a convenience return that could partially offset the physical and financial costs of carrying stocks.

The proposition that inventory holders might derive a convenience yield from holding inventories is embedded in the classical supply of storage function examined empirically by Working (1934, 1949), Brennan (1958), Fama and French (1987, 1988), and Miranda and Rui (1996). Pindyck (1993) examines the extent to which the current price of a commodity can be explained by the present discounted sum of expected convenience returns. As a whole, this body of research appears to indicate that indeed there might be not only a speculative, but also a convenience motive for inventory holding. However, the implications of introducing convenience yield on the time series properties of commodity prices remains to be investigated. In particular, it would be interesting to examine how the introduction of convenience yield alter the degree of serial correlation in prices. Recall that the basic storage model predicts a large number of stockouts, and the weak intertemporal link served by inventories is responsible for the low degree of serial correlation in prices. Since the convenience yield could partially compensate inventory holders for the expected loss when the basis is below carrying charges, the model which allows for convenience yield should predict a smaller number of stockouts. One would therefore expect the demand for inventories for convenience purposes to strengthen the intertemporal link and hence predict a higher persistence in prices.

Since the sole determinant of the convenience yield is the level of inventories held, one could mathematically express the convenience return solely as a function of the level of stocks. Let $\phi(I_{t,t+1})$ be the convenience yield from holding inventories between period t and $t + 1$. Economic reasoning suggests that the marginal value of holding inventories should increase with scarcity, and diminishing marginal returns suggests a declining and concave relationship between stocks and the convenience yield. The latter arises from the fact that as inventories accumulate, marginal storage costs increase. Hence $\phi(I_{t,t+1})$ has the property that $\phi'(I_{t,t+1}) > 0$ and $\phi''(I_{t,t+1}) < 0$. For example, one could parameterize the (gross) marginal convenience yield as

$$\phi'(I_{t,t+1}) = \theta + (1 - \theta)g(I_{t,t+1}), \quad (10)$$

where

$$\theta = \frac{(1+r)(1-\varepsilon)}{(1-\delta)},$$

for an arbitrarily small ε and a normalizing constant $c \geq 0$, and

$$g(I_{t,t+1}) = \frac{I_{t,t+1}}{I_{t,t+1} + c}.$$

This specification for $\phi'(I_{t,t+1})$ nests the case with no convenience yield as a special case by setting $c = 0$, satisfies the requirement that the marginal convenience yield be a decreasing, convex function of the level of inventories, and insures that it will be appropriately bounded between $(1+r)/(1-\delta)$ when $I_{t,t+1} \rightarrow 0$ and 1 when $I_{t,t+1} \rightarrow \infty$. The property that $\phi'(I_{t,t+1})$ is bounded between $(1+r)/(1-\delta)$ and 1 is important because it guarantees the existence of the equilibrium price function [see Deaton and Laroque (1992)] and preserves the non-linearity of the price process. That is, there might be (negative) values of the intertemporal price spread that cannot be compensated by the convenience return, in which case, the non-negativity constraint on inventories binds. Miranda and Rui (1996) postulate a cost function that admits an unbounded marginal convenience yield. In this case, inventories are always strictly positive. Since the intertemporal price link created by stockholders is never severed by stockouts, this model predicts a high serial correlation in prices. However, it is unclear how a model without stockouts explains another fundamental feature of commodity prices, namely the presence of price spikes.

Following Gibson and Schwartz (1990), Schwartz (1997), and Routledge, Seppi, and Spatt (1997), we introduce the convenience yield multiplicatively. Then, stockholding is now governed by the arbitrage condition:

$$I_{t,t+1} \geq 0 \quad \text{if} \quad \phi'(I_{t,t+1})(1-\delta)/(1+r)E_t p_{t+1} = p_t, \quad (11)$$

As before, the stockholder will carry inventories if and only if the expected return covers the cost of storage *net* of the convenience yield, and prices will follow a white noise process in the stockout regime. However, a unique implication of the convenience yield is that prices will follow an AR(1) process with time-varying coefficient $(1+r)/[(1-\delta)\phi'(I_{t,t+1})]$ during the stockholding regime. Furthermore, the price above which inventory demand would be zero is $p^* = \phi'(I_{t,t+1})(1-\delta)/(1+r)E_t p_{t+1}$. Up the extent that the level of stocks is time-varying, p^* is no longer constant as in the basic model. Notice for a current price above p^* , even the convenience return will be insufficient to induce a demand for stocks.

In principle, the introduction of a convenience yield might allow us to decompose the demand for inventories in two components, namely speculative and precautionary demand. For speculative demand, $\phi'(I_{t,t+1}) = 1$ and the (gross) return is $(1-\delta)/(1+r)E_t p_{t+1}$. For precautionary demand,

$\phi'(I_{t,t+1}) \geq 1$ and the (gross) return is $\phi'(I_{t,t+1})(1 - \delta)/(1 + r)E_t p_{t+1}$. However, notice that the above parameterization of the marginal convenience yield implies that $\phi'(I_{t,t+1}) > 1$ for any non-negative level of inventories and, consequently, predicts that all stocks will be held for precautionary reasons. Therefore, we also consider an alternative, linear specification of $\phi'(I_{t,t+1})$ that satisfies the boundary conditions and yet permits the identification of speculative and precautionary inventory demand. Let

$$\phi'(I_{t,t+1}) = \theta + (1 - \theta)g(I_{t,t+1}), \quad (12)$$

where

$$\theta = \frac{(1 + r)}{(1 - \delta)},$$

$$g(I_{t,t+1}) = \min \left\{ \frac{\gamma}{\theta - 1} I_{t,t+1}, 1 \right\},$$

and $\gamma > 0$. In this case, the demand for inventories will have two kinks. The first one is the constant $p^{**} = (1 - \delta)/(1 + r)E_t p_{t+1}$, and speculative demand will be zero when $p > p^{**}$ (but precautionary demand will still be positive). The second one is the time-varying $p^* = \phi'(I_{t,t+1})(1 - \delta)/(1 + r)E_t p_{t+1}$, and when $p > p^*$, even the convenience return will be insufficient to induce a positive demand for inventories. Notice that since $\phi'(I_{t,t+1}) \geq 1$, $p^* \geq p^{**}$ for all t . Thus, whenever speculators are in the market (in the sense that they demand positive levels of stocks), precautionary holders might be in the market too, but their gross convenience return would be 1. For values of the current price strictly larger than p^{**} but smaller than p^* , only precautionary holders demand stocks. Finally, if the current price exceeds p^* , neither speculators nor precautionary holders demand inventories and a stockout occurs. While, it is clear that the existence of a second price kink is solely an artifact of the parameterization of the marginal convenience yield, this specification will prove useful below in assessing the relevance of the convenience yield in explaining the persistence of commodity prices.

The state variables for this model are the same as in the basic storage model, namely, the current level of inventories. Thus, $E_t p_{t+1}$ is parameterized by $\psi_7(I_{t,t+1})$ where the coefficients of the polynomial $\psi_7(\cdot)$ satisfy the rational expectations solution and are numerically estimated as in the models above.

6 Empirical Assessment

The specifications proposed above introduce features of the production process and financial markets that can (on theoretical grounds) explain the observed persistence in commodity prices. In order to assess the quantitative and qualitative implications of these extensions in a unified framework, we calibrated the models above using the parameter values estimated by Deaton and Laroque

(1996) for a set of agricultural and industrial commodities. The goal of the calibration exercise is to simulate a time series of commodity prices subject to the restrictions of the theoretical model and compare its properties, as summarized by a set of statistics, with the ones of the actual data. A statistic of primary interest is the first order autocorrelation of prices.

Before proceeding further, we examine whether the method of parameterized expectations used to obtain the rational expectations solution can replicate the results in Deaton and Laroque (1992, 1996). To this end, the basic storage model was first calibrated for the parameters values used in Deaton and Laroque (1992, p.11). In that article, the authors assume a linear inverse demand function of the form $P(x) = a + bx$ and examine the degree of serial correlation in prices for an arbitrary set of structural parameters (a, b, δ, r) . Note that given the identification proposition of Deaton and Laroque (1996), the parameter a cannot be separately identified from the mean of harvest, and the parameter b from the variance of harvest. Thus, variations in a across commodities can arise because of variations in autonomous demand, the mean of the shock, or both. With this interpretation in mind, the method of parameterized expectation yields an estimated first order autoregressive coefficient of 0.087 for the values $(a, b, \delta, r) = (200, -1.0, 0.05, .056)$. For the set $(a, b, \delta, r) = (600, -5.0, 0.0, .05)$, the estimated value is 0.441. Note that these statistics are very close to the values of 0.08 and 0.48 reported in Deaton and Laroque (1992).

Since our objective is to compare the properties of the extended models with those of the basic storage model, we also calibrated the latter using the estimates $(a, b, \text{ and } \delta)$ reported in Deaton and Laroque (1996, p. 911) for a set of 12 commodities. To be consistent with Deaton and Laroque, the rate of interest was fixed to $r = .05$. For each calibrated model, 5000 observation were simulated and the first order autoregressive coefficient was calculated. For this sample size the autoregressive coefficient has a standard error of 0.014. These results serve as our base case and are presented in Table 1. Results reported in Deaton and Laroque (1996) are also reproduced in Table 1 for convenience. Although this article uses the method of parameterized expectation to find the rational price forecast while Deaton and Laroque's uses value function iteration, the reader can verify that both solution methods yield coefficient estimates that are quantitatively quite close. The largest difference (in absolute value) is 0.057 for the case of copper and maize. The average difference is only 0.013. Thus, the low degree of serial correlation predicted by the basic model appears robust to the choice of solution methods. Thus, it seems reasonable to conclude, that up to the extent that our models (below) produce estimates of the autoregressive coefficients that differ from the ones of the basic storage model, this difference is not attributable to the use of an alternative procedure for modelling the agents' forecast.

The results for the extended models are reported in Tables 2, 3 and 4. In Table 2, a comparison of the serial correlation for the data (first column) and for the base case (second column) indicates

that the commodity prices simulated from the basic model have only about one-third the serial correlation found in the actual time series. Accounting for this discrepancy is the objective of our analysis.

Continuing in Table 2, results from the last three columns indicate that (as expected) naively introducing serially correlated shocks does increase the persistence predicted by the model. However, while for plausible values of the moving average coefficient ($\rho = 0.2$ and $\rho = 0.5$) the increase is substantial, the serial correlation is still below the one observed in actual data. Only, when the coefficient takes the more extreme value of $\rho = 0.8$, does the predicted serial correlation rises enough to be comparable with the empirical one for a small set of commodities.

Since the introduction of serially correlated shocks is only a crude way of handling the possible misspecification of the basic model, we now turn our attention to the economic models of inventory holding that are the main focus of this paper. It is hoped that this specifications might eliminate some (or all) of the sources of misspecification in the competitive storage model.

Consider first, the model with gestation lags. It is apparent from the results in Table 3, that production lags *alone* cannot account for the persistence observed in the time series of commodity prices. In the case of maize, sugar, tin and wheat, the degree of persistence decreases rather than increases. Interestingly, these tend to be commodities with a high estimate for δ . Less of what is stored is made available to the market when the depreciation rate is high, and this weakens the intertemporal link in prices.

However a model with gestation lags *and* heteroskedastic supply shocks does predict a higher price persistence than the basic model. Raising σ_2/σ_1 from 1 to 1.5 raises the calculated first order serial correlation by over 30 percent in the case of cocoa, and by as much as 70 percent in the case of wheat. An additional 7 to 30 percent increase is observed as σ_2/σ_1 is increased to 1.8.¹² Note also that the improvement tend to be larger when b is relatively small in absolute value. In spite of these results, it is apparent that this model does not generate enough price persistence to match the values observed in the actual series. The extension is particularly futile with commodities such as cotton, rice and wheat, whose depreciation rates are estimated to be high.

In order to further examine the role of heteroskedasticity in producing price persistence, we solved the model for cocoa with increasingly large values of σ_2/σ_1 . In addition to the values of already considered (1, 1.5, and 1.8), we also employed (2, 2.5, 3, and 5). The first-order serial correlation were estimated using simulated samples of 5000 observations. The calculated values are presented in Figure 1. These results clearly suggests that introducing heteroskedasticity in supply

¹²For the cases of maize and palm oil, the numerical procedure fails to find solutions that involve polynomials of order higher than 1 for the model with heteroskedastic disturbances. Since the linear approximation might not be as accurate as the parameterizations that include higher order terms, the results for these two commodities must be interpreted with caution

shocks has a substantial effect on the serial correlation predicted by the model, but the increase is marginal to nil once σ_2/σ_1 exceeds 2.5. Thus, even increasing σ_2/σ_1 to unrealistic values cannot replicate the persistence observed in the data.

The results for the model that incorporates overlapping contracts are reported in Table 4. Compared to the base case, the degree of serial correlation is increased from a modest 1 percent in the case of maize to an impressive 94 percent in the case of cotton. While these results are encouraging, the degree of persistence is still substantially smaller than the one estimated using actual the data.

As noted in Section 4, the introduction of contracts was conjectured to reduce the probability of stockouts by providing additional sources of availability. For example, even if in the current period (t) speculators are unwilling to carry one-period inventories, the decision taken in the preceding period ($t - 1$) to carry two-period inventories means that, should a bad harvest occur, a stockout next period (that is, $t + 1$) could be avoided. To examine whether this hypothesis is supported by the data, Figure 2 presents a subsample of 100 observations simulated using the model for cocoa. This figure suggests that while two-period contracts are demanded by speculators, they complement rather than replace one-period ahead contracts. Moreover, this type of inventories are held less frequently than the shorter term maturities. For example, while one-period contracts are held 96.3 percent of the times, two-period contracts are held only 29.3 percent of the periods and usually simultaneously with the shorter maturity contracts. This result underscores the role of inventories as a mechanism to transfer to goods from abundant to scare periods, and provide some insight of the role of contracts on the determination of commodity prices. However, they also indicate that contracts alone might not fully explain the high degree of serial correlation observed in the data.

The results of the simulations for the model with a convenience return are presented in the last four columns of Table 4. Consistent with the findings of the previous two extensions, serial correlation is higher for those commodities with a low depreciation rate. Though still lower than those observed in the data, it is apparent that the storage model with convenience yield produces the largest increase in serial correlation of the models considered. Indeed, with suitable choice of parameters (such as c and γ) both specifications of $\phi'(I_{t,t+1})$ are capable of increasing the serial correlation substantially. Recall that in the gestation lag model, even an unreasonable degree of heteroskedasticity cannot increase the serial correlation to levels close to those observed in the data. With the introduction of a convenience return using the smooth parameterization of ϕ' in (10), maize and palm oil now have serial correlation coefficients of over 0.6, which are much closer to the observed values of around 0.7 than the estimates obtained from the previous extensions. A further analysis of the simulations reveals that the introduction of a convenience-yield motive reduces significantly the probability of stockouts and the number of spikes observed in a finite

sample. With an additional reason for holding inventories, agents might carry stocks even if the expected return is lower than the carrying cost. Consequently, in a given sample, the number of observed stockouts diminishes and the number of periods in which inventory holding creates an intertemporal link in prices increases. Thus, the estimated serial correlation of prices rises.

With ϕ' defined in (12), the ability of the convenience yield to increase serial correlation in prices depends on the value of γ . For the case when $\gamma = 1$, the marginal convenience return rapidly decreases with inventories and hence does not significantly strengthen the intertemporal link in prices. Hence the estimates of serial correlation are numerically close to the values predicted by the basic model where the precautionary motive is absent. Notice for the case when $\gamma = .01$, the serial correlations predicted by the model are comparable to ones obtained with the smooth parameterization of ϕ' defined (10). Note also that the autoregressive coefficient implied by the two parameterizations of ϕ' are (different) time-varying functions of the level of stocks, accounting for the differences in the time series properties of prices implied by the two parameterizations.

In order to disentangle the demand for precautionary and speculative inventories and to further examine the effect of buffer-stockholding behavior on commodity prices, the fraction of inventories held for convenience purpose is calculated using the model with the piece-wise linear marginal convenience yield function described in (12) for various values of the coefficient γ . This parameter measures the reduction in the marginal convenience yield as a result of an unitary increase in the level of stocks and is smaller the stronger is the desire for holding buffer stocks. Notice that for all commodities, the price persistence is higher, the lower the value of γ . Furthermore, the last row of Table 4 reveals that the larger is γ , the lower the proportion of stocks held by speculators. The intuition for these results is straightforward, for lower values of γ the marginal convenience yield decreases more slowly with the level of stocks. Thus, there is a larger range of values of stocks for which the precautionary motive dominates the speculative motive for inventory holding. Up to the extent that agents are willing to hold stocks not only for the expected capital gain but also for the convenience return, the probability of observing stockouts decreases and the price persistence rises. The average proportion of precautionary stocks for all 12 commodities is smaller (22.6%) than in the two other cases considered (100% when $\gamma = .01$ and 84.2% when $\gamma = .1$).

The role of precautionary stockholding on the persistence of commodity prices is further documented in Figure 3. For this figure the model for cocoa was solved and simulated for a large range of values of γ . The estimates of serial correlation were based on samples of 5000 observations. This graph clearly suggests that price persistence is monotonically increasing in the proportion of stocks held for precautionary motives.

Finally, it is important to note that the model with a convenience return yields predictions of the price serial correlation that (i) are numerically similar to the ones obtained by mechanically

introducing MA(1) shocks (with $\rho = 0.2, 0.5$), but (ii) are based on more a more solid economic rationale, namely the precautionary nature of inventory demand.

7 Conclusions

The goal of this analysis was to generalize the basic storage model to reproduce the degree of serial correlation observed in actual data. The proposed specifications preserve important features of the basic model, namely the non-negativity constraint on inventories and the possibility of stockouts that account for occasional price spikes. However, they explicitly incorporate realistic aspects of the production process and financial markets that might potentially explain the persistence of commodity prices. The results are only partially successful in that the first two extensions are capable of reducing the discrepancy between the serial correlation in the data and the one predicted by theory from a factor of 3 to 2. The introduction of gestation lags with heteroskedastic shocks is somewhat more successful, but substantial persistence is still left unaccounted for.

The third specification considers the more general situation when agents can demand inventories for both precautionary and speculative reasons. This model predicts values of serial correlation that are closer to the actual ones, explains almost 65% of the observed price persistence, and unambiguously highlights the role of a convenience return in the dynamics of commodity prices. At the theoretical level, these results suggest that explicitly modeling the agents' risk aversion (that underlies the convenience yield) might produce a model of storage that better captures the features of the data. At the empirical level, these evidence indicates that a successful explanation for the time series properties of commodity prices might require the introduction of a non-speculative demand for stocks.

Nonetheless, other modifications to the basic model might be necessary. Including all the above features in a single model can be numerically challenging but might provide more accurate predictions than each separate specification on its own. More importantly, given the substantial differences in institutional and physical features across commodities and the idiosyncracies of their production and consumption processes, it could be argued that modeling these characteristics on a commodity by commodity basis may be necessary to obtain a better match between the model predictions and the data. The specifications proposed above could be useful in that respect, because they allow one to analyze commodities with a variety of trading and production mechanisms by suitably modelling the length and type of the contracts, the seasonality in production, and the heterogeneity of depreciation rates, price elasticity of demand, and spoilage costs.

Table 1. Comparison of Estimates of Serial Correlation

$$\text{Demand} = P(x) = a + bx$$

$$z \sim N(0, \sigma = 1)$$

Commodity	a	b	δ	D&L Procedure	Parameterized Expectations
Cocoa	.162	-.221	.116	0.298	0.352
Coffee	.263	-.158	.139	0.242	0.219
Copper	.545	-.326	.069	0.392	0.335
Cotton	.642	-.312	.169	0.192	0.173
Jute	.572	-.356	.096	0.302	0.289
Maize	.635	-.636	.059	0.356	0.413
Palm Oil	.461	-.429	.058	0.416	0.397
Rice	.598	-.336	.147	0.224	0.237
Sugar	.643	-.626	.177	0.264	0.266
Tea	.479	-.211	.123	0.230	0.213
Tin	.256	-.170	.148	0.256	0.238
Wheat	.723	-.394	.130	0.259	0.250

Notes: All simulations are based on a sample size of 5000 observations. The conditional price expectation was parameterized using a polynomial of order 3 in the current level of stocks.

Table 2. Estimates of Serial Correlation with MA(1) Shocks

$$\text{Demand} = P(x) = a + bx$$

$$z \sim N(0, \sigma)$$

Commodity	Actual	Basic Model	MA(1)	MA(1)	MA(1)
			Shocks	Shocks	Shocks
			$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$
Cocoa	0.834	0.352	0.447	0.547	0.609
Coffee	0.804	0.219	0.367	0.489	0.576
Copper	0.838	0.335	0.418	0.529	0.619
Cotton	0.884	0.173	0.328	0.468	0.564
Jute	0.713	0.289	0.391	0.522	0.589
Maize	0.756	0.413	0.481	0.584	0.644
Palm Oil	0.730	0.397	0.479	0.587	0.637
Rice	0.829	0.237	0.327	0.494	0.579
Sugar	0.621	0.266	0.364	0.504	0.583
Tea	0.778	0.213	0.335	0.478	0.571
Tin	0.895	0.238	0.377	0.521	0.567
Wheat	0.863	0.250	0.352	0.504	0.602

Notes: All simulations are based on a sample size of 5000 observations. The conditional price expectation was parameterized using polynomials of order 2.

Table 3. Estimates of Serial Correlation with Gestation Lags

$$\text{Demand} = P(x) = a + bx$$

$$z \sim N(0, \sigma)$$

Commodity	Actual	Basic Model	Gestation	Gestation	Gestation
			Lags	Lags	Lags
			$\sigma_2/\sigma_1 = 1$	$\sigma_2/\sigma_1 = 1.5$	$\sigma_2/\sigma_1 = 1.8$
Cocoa	0.834	0.352	0.353	0.467	0.511
Coffee	0.804	0.219	0.256	0.385	0.433
Copper	0.838	0.335	0.353	0.471	0.526
Cotton	0.884	0.173	0.196	0.278	0.365
Jute	0.713	0.289	0.298	0.417	0.486
Maize	0.756	0.413	0.385	0.604	0.620
Palm Oil	0.730	0.397	0.416	0.593	0.640
Rice	0.829	0.237	0.219	0.343	0.398
Sugar	0.621	0.266	0.239	0.355	0.427
Tea	0.778	0.213	0.218	0.361	0.428
Tin	0.895	0.238	0.226	0.380	0.428
Wheat	0.863	0.250	0.222	0.384	0.411

Notes: The standard deviation σ_1 was fixed to 1 while the standard deviation σ_2 was allowed to take the values 1, 1.5 and 1.8. All simulations are based on a sample size of 5000 observations. The conditional price expectation was parameterized using polynomials of order 2 (except for maize and palm oil where a polynomial of order 1 was employed for the cases $\sigma_2/\sigma_1 = 1.5$ and $\sigma_2/\sigma_1 = 1.8$).

Table 4. Estimates of Serial Correlation

$$\text{Demand} = P(x) = a + bx$$

$$z \sim N(0, \sigma)$$

Commodity	Actual	Basic Model	Overlapping Contracts	Conv.	Conv.	Conv.	Conv..
				Yield $c = 50$	Yield $\gamma = .01$	Yield $\gamma = .1$	Yield $\gamma = 1$
Cocoa	0.834	0.352	0.462	0.522	.537	.416	.357
Coffee	0.804	0.219	0.358	0.530	.477	.365	.231
Copper	0.838	0.335	0.394	0.608	.497	.375	.333
Cotton	0.884	0.173	0.337	0.473	.458	.341	.212
Jute	0.713	0.289	0.365	0.545	.530	.393	.333
Maize	0.756	0.413	0.418	0.623	.557	.482	.380
Palm Oil	0.730	0.397	0.438	0.625	.589	.432	.399
Rice	0.829	0.237	0.334	0.475	.441	.344	.254
Sugar	0.621	0.266	0.370	0.424	.456	.379	.279
Tea	0.778	0.213	0.302	0.509	.458	.352	.246
Tin	0.895	0.238	0.355	0.472	.486	.391	.281
Wheat	0.863	0.250	0.368	0.505	.508	.357	.270
Precautionary Stocks		0%	0%	100%	100%	84.2%	22.6%

Notes: The last row denotes the average of the percentage of precautionary stocks for all 12 commodities. All simulations are based on a sample size of 5000 observations. The conditional price expectation was parameterized using polynomials of order 2 (overlapping contracts) and 3 (convenience yield).

References

- [1] Breeden, D. T. (1980), Consumption Risks in Futures Markets, *Journal of Finance* **35**, pp. 503-520.
- [2] Brennan, M. J. (1958), The Supply of Storage, *American Economic Review* **47**, pp. 50-72.
- [3] Chambers, M. J. and Bailey, R. (1996), A Theory of Commodity Price Fluctuations, *Journal of Political Economy* **104**, pp. 924-957.
- [4] Cootner, P. H. (1967), Returns to Speculators: Telser vs. Keynes, *Journal of Political Economy* **68**, pp. 396-404.
- [5] Deaton, A. S. and Laroque, G. (1992), On the Behavior of Commodity Prices, *Review of Economic Studies* **89**, pp. 1-23.
- [6] Deaton, A. S. and Laroque, G. (1996), Competitive Storage and Commodity Price Dynamics, *Journal of Political Economy* **104**, pp. 896-923.
- [7] Den Haan, W. J. and Marcet, A. (1990), Solving the Stochastic Growth Model by Parameterizing Expectations, *Journal of Business and Economic Statistics* **8**, pp. 31-34.
- [8] Dusak, K. (1973), Futures Trading and Investor Returns: An Investigation of Commodity Market Risk Premiums, *Journal of Political Economy* **81**, pp. 1387-1406.
- [9] Fama, E. F. and French, K. R. (1987), Commodity Futures Prices: Some Evidence on Forecast Power, Premiums, and the Theory of Storage, *Journal of Business* **60**, pp. 55-73.
- [10] Fama, E. F. and French, K. R. (1988), Business Cycles and the Behavior of Metal Prices, *Journal of Finance* **43**, pp. 1075-1093.
- [11] Fischer, S. (1977), Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule, *Journal of Political Economy* **85**, pp. 191-205.
- [12] Gibson, R. and Schwartz, E. S. (1990), Stochastic Convenience Yield and the Pricing of Oil Contingent Claims, *Journal of Finance* **45**, pp. 959-976.
- [13] Gustafson, R. L. (1958), Carryover Levels for Grains: A Method for Determining Amounts that are Optimal under Specified Conditions, *USDA Technical Bulletin* **1178**.
- [14] Judd, K. L. (1990), Minimum Weighted Residuals Methods for Solving Dynamic Economic Models, Hoover Institution *Mimeo*.

- [15] Kaldor, N. (1939), Speculation and Economic Stability, *Review of Economic Studies* **7**, pp. 1-27.
- [16] Keynes, J. M. (1930), *A Treatise on Money, Volume II: The Applied Theory of Money*. London: Macmillan.
- [17] Lowry, M., Glauber J., Miranda, M., and Helmberger, P. (1987), Pricing and Storage of Field Crops: A Quarterly Model Applied to Soybeans, *American Journal of Agricultural Economics* November, pp. 740–749.
- [18] Miranda, M. J. and Rui, X. (1996), An Empirical Reassessment of the Commodity Storage Model. Ohio State University, *Mimeo*.
- [19] Muth, J. F. (1961), Rational Expectations and the Theory of Price Movements, *Econometrica*. 29, pp. 3-22.
- [20] Newbery, D. and Stiglitz, J. (1982), Optimal Commodity Stockpiling Rules, *Oxford Economic Papers* **34**, pp. 403–427.
- [21] Ng, S. (1996), Looking for Evidence of Speculative Storage in Commodity Markets, *Journal of Economic Dynamics and Control* 20, pp. 123-143.
- [22] Pindyck, R. S. (1993), The Present Value Model of Rational Commodity Pricing, *Economic Journal* **103**, pp. 511-530.
- [23] Routledge, B. R., Seppi, D. J., and Spatt, C. S. (1997), Equilibrium Forward Curves for Commodities, Graduate School of Industrial Administration, Carnegie Mellon University, *Mimeo*.
- [24] Samuelson, P. A. (1957), Intertemporal Price Equilibrium: A Prologue to the Theory of Speculation, *Weltwirtschaftliches Archiv* **79**, pp. 181-219.
- [25] Scheinkman, J. A. and Schechtman, J. (1983), A Simple Competitive Model with Production and Storage, *Review of Economic Studies* **50**, pp. 427-441.
- [26] Schwartz, E. S. (1997), The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging, *Journal of Finance* **52**, pp. 923-973.
- [27] Taylor, J. B. (1979), Staggered Wage Setting In a Macro Model, *American Economic Review* **69**, pp. 108–213.
- [28] Taylor, J. B. (1980), Aggregate Dynamics and Staggered Contracts, *Journal of Political Economy* **88**, pp. 1–23.

- [29] Taylor, J. B. and Uhlig, H. (1990), Solving Nonlinear Stochastic Growth Models: A Comparison of Alternative Solution Methods, *Journal of Business and Economic Statistics* **8**, pp. 1-17.
- [30] Telser, L. G. (1958), Futures Trading and the Storage of Cotton and Wheat, *Journal of Political Economy* **66**, pp.233-255.
- [31] Vaile, R. S. (1944), Cash and Future Prices of Corn, *Journal of Marketing* **9**, pp. 53-54.
- [32] Vance, L. L. (1946), Grain Market Forces in the Light of Inverse Carrying Charges, *Journal of Farm Economics* **28**, pp. 1036-1040.
- [33] Williams, J. (1986), *The Economic Fluctuations of Futures Markets*. New York: Cambridge University Press.
- [34] Williams, J. and Wright, B. (1991), *Storage and Commodity Markets*, Cambridge: Cambridge University Press.
- [35] Working, H. (1934), Price Relations Between May and New-Crop Wheat Futures at Chicago since 1885. Food Research Institute *Wheat Studies* **10** pp. 183-228.
- [36] Working, H. (1949), The Theory of Price Storage, *American Economic Review* **39**, pp. 1254-1262.
- [37] Wright, B. and Williams, J.. (1989), A Theory of Negative Prices for Storage, *Journal of Futures Markets* **9**, pp. 1-13.

Fig 1. Serial Correlation and Heteroskedastic Supply Shocks




Fig 2. Comparison of Inventory Holdings




Fig 3. Precautionary Storage and Serial Correlation

