



Looking for evidence of speculative stockholding in commodity markets

Serena Ng

Department of Economics and C.R.D.E., University of Montreal, Montreal, Que. H3C 3J7, Canada

(Received May 1994; final version received December 1994)

Abstract

The theory of commodity price with speculative storage predicts that prices are a two-regime process depending on whether or not inventories are held. The price process is nonlinear in that it is nondifferentiable at some p^* which separates the data into a history independent regime and an autoregressive process. This paper looks for evidence of nonlinearity in the price data and tests the theory in the context of threshold autoregressive models under the assumption that shocks to harvest are i.i.d. While we find evidence for regime-specific behavior, we also find the degree of persistence in the stockout regime to be much stronger than that predicted by theory.

Key words: Commodity price; Speculative storage; Threshold models

JEL classification: C13; C22; G10

1. Introduction

Theory predicts that under the assumptions of rational expectations and i.i.d. shocks to harvest, inventories are the only link between the market for commodities in successive periods. Speculative inventories will be held only when prices are expected to rise by carrying cost. Prices are an autoregressive process when the demand for speculative inventories is positive but a linear function of white noise shocks when no inventory is held. Although the price series is a nonlinear

This is a substantially revised version of Chapter 2 of my thesis submitted to Princeton University. I would like to thank my advisors Angus Deaton and Pierre Perron for their guidance, Eric Ghysels and two referees for useful comments on an earlier draft. Financial support is provided by the Lynne and Harry Bradley Foundation (Princeton) and C.A.F.I.R. (U.M.).

first-order Markov process, persistence is derived solely from speculative stockholding. As well, large price hikes are possible in the event of a stockout.

Although the model is well accepted at the theoretical level, there is little empirical work which explicitly tests the speculative storage model under rational expectations. While the recent attempts of Deaton and Laroque (1992a,b) found some support for the theory, the authors also found aspects of the theory that are inconsistent with the data.

This paper is another attempt to fit the rational expectations variant of the commodity price model. The present study adopts a time series, nonstructural approach. Theory predicts that the price process for the stockholding and stockout regimes should have different stochastic properties. Our objective is to look for evidence of regime-specific behavior in the context of self-exciting threshold autoregressive models which explicitly allow the behavior of the price process to be regime-dependent. While such a reduced form analysis precludes the identification of structural parameters, they also avoid having to make specific assumptions about demand functions, for example, and provide a less restrictive platform for assessing the theory.

The rest of this paper is structured as follows. The key features of the commodity price model with speculative storage are summarized in the next subsection. This is followed by a brief review of the empirical literature. Section 2 tests for nonlinearity in thirteen commodity price series using a battery of statistics. Section 3 summarizes the basic structure of a threshold autoregressive model with one threshold. Section 4 presents the results based on a Kalman filter which allows the regime-specific parameters and the threshold value to be estimated jointly. An overview of the results and possible extensions to the analysis are discussed in the conclusion.

1.1. A simple commodity price model

Our goal is to analyze the behavior of prices for commodities such as copper, wheat, and cotton. The market for commodities is a well-developed one, with almost continuous trading occurring on a worldwide basis. Commodity prices do not exhibit obvious trends, but they are highly persistent and tend to be interrupted by large and positive spikes.

The demand for commodities can be derived from one of two sources. First, commodities can be used as input in the production process or they can be consumed. To the extent that consumption/production demand is closely tied to economic development, such nonspeculative demand is strictly positive but could be cyclical.

Second, commodities are also assets and therefore speculative demand may exist. An important feature of commodities is that most of them are storable, albeit at a nonzero cost. When there are profits to be made from storing a commodity, risk-neutral speculators will enter the market and drive up prices

until all profit opportunities are arbitrated away. However, when the holding cost exceeds expected capital gains, speculators will refrain from entering the market. Thus, unlike consumption demand, speculative demand can be zero but can also be quite volatile.

On the supply side, commodities can be made available to the market from one of two sources. They can come from new harvest or be supplied by speculative stockholders. In the absence of speculative storage, the market-clearing price will be a function of harvest, and hence the stochastic process of shocks to harvest. If harvest innovations are i.i.d. and consumption demand is linear and autonomous, then the price will be a linear function of contemporaneous harvest shocks. Allowing consumption demand to be stochastic will not change the simple dependency of price on current shocks as long as shocks to the linear demand function are also i.i.d. However, when there is speculative storage, the market-clearing price will depend on the availability of total supply (harvest plus inventories) relative to total demand (consumption cum speculative demand). It is the behavior of prices under the latter conditions that we wish to investigate.

The determination of commodity prices with speculative storage is a well-researched area. The idea that speculative stockholding is bounded from below by zero was discussed in Williams (1937) and Samuelson (1957), and a model incorporating rational expectations was formalized by Gustafson (1958). The original theory has since been much refined. See Deaton and Laroque (1992a) and the references therein for related work.

We are interested in testing the rational expectations version of the competitive storage model. Let r be the real interest rate and δ be the depreciation cost. The convenience yield is assumed to be zero so that there is no demand for inventories other than for speculative purposes. It is assumed that $(1 - \delta)/(1 + r) < 1$ which implies that speculative inventories, I_t , are costly to hold. If z_t is harvest, then total supply available to the market is given by

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

Let $P(x)$ be the inverse demand function corresponding to the demand function $D(p)$, where p_t is the price. The expectation of p_{t+1} given information on supply and demand available at time t is denoted $E_t p_{t+1}$. It is assumed that expectations are rational. The main predictions of the model can be summarized as follows:

Lemma 1. Assume that the regularity conditions on $D(p)$ given in Assumption 1 of Deaton and Laroque (1992a) hold and let shocks to harvest be i.i.d. Let $p^* = (1 - \delta)/(1 + r)E f(z)$ where $f(\cdot)$ is the equilibrium price function satisfying

$$f(x_t) = \max \left[\frac{(1 - \delta)}{(1 + r)} E_t f \{ z_{t+1} + (1 - \delta)I_t \}, P(x_t) \right]. \quad (1)$$

If $p_t \geq p^*$, then $I_t = 0$ and

$$p_{t+1} = f(z_{t+1}). \quad (2)$$

If $p_t < p^*$, then $I_t > 0$ and

$$p_{t+1} = \left(\frac{(1 - \delta)}{(1 + r)} \right)^{-1} p_t + \eta_{t+1}, \quad (3)$$

where $\eta_{t+1} = f(x_{t+1}) - p_t(1 + r)/(1 - \delta)$ and $E_t(\eta_{t+1}) = 0$.

A formal proof of the lemma is given in Deaton and Laroque (1992a) where it was shown that the price process is ergodic under the assumptions of their analysis. Since $\delta > 0$ by hypothesis, a stockout will occur and prices exceed p^* with probability one in finite time. The long-run distribution of prices will oscillate between the distribution corresponding to the case of no storage and the distribution corresponding to the case when inventories increase without bounds. The price series is therefore a renewal process with a stationary and invariant distribution.

The above model which emphasizes the importance of shocks to harvest is evidently more appropriate for agricultural than for industrial commodities. However, reinterpreting x above as the total (industrial and speculative) demand for an industrial commodity, one can construct a model such that (i) speculative demand is zero when industrial demand is high, in which case price equates the production of the commodity to noninventory demand, but (ii) speculative demand is positive when industrial demand is low, in which case price equates supply and total demand. The model will also have state-dependent features as summarized in the above lemma.¹

The essential features of the speculative storage model are as follows. First, since shocks to harvest are i.i.d., they reveal no information about future events. Thus, $E_t f(z_{t+1})$ is constant, and the expectation also holds unconditionally.

Second, expected profits of speculators are bounded from below by zero since they have the option not to enter the market. There will exist a critical value, p^* , such that no inventory will be held if the current price exceeds that critical level. The variable p^* is the price (after discounting) that is expected to prevail if harvests were the only source of supply. When the current price is too high and expected capital gains are low or negative, speculators will not enter the market. When no inventory is brought to the market, the price will be a sole function of

¹ See Gilbert (1993) for a structural model for industrial commodities.

harvest. Thus, the price following the period when no inventory is held is history-independent and is given by (2).

Third, when speculative demand is positive, total supply in period $t + 1$ will be the sum of harvest in $t + 1$ plus inventories held over from period t . The price in period t and $t + 1$ are linked to the extent that speculative demand in period t depends on the price that clears the market at time t . Thus, the supply of inventories in $t + 1$ will also depend on the price at time t . In that case, the evolution of price is given by (3), a first-order Markov process with an autoregressive coefficient $[(1 - \delta)/(1 + r)]^{-1}$ that is greater than one by assumption. If inventories are to be held, the price must be expected to rise at the rate of carrying cost.

Fourth, the model has specific statistical implications. The conditional mean of prices is

$$E(p_{t+1} | p_t) = \min(p^*, p_t) \left[\frac{(1 - \delta)}{(1 + r)} \right]^{-1}. \quad (4)$$

The price function is nondifferentiable at p^* . This induces a kink, and therefore a source of nonlinearity in the conditional mean of the price process. Furthermore, it can be seen from the definition of η_{t+1} and the i.i.d. assumption on z_{t+1} that the conditional variance of prices is heteroskedastic in p_t when stocks are held but homoskedastic otherwise. Intuitively, fewer inventories are being stored at higher level of prices, and the increasing probability of being unable to satisfy demand causes the conditional variance of p_{t+1} to increase with p_t . In consequence, the overall price process is also conditionally heteroskedastic.² More generally, the higher-order conditional moments of p_t also have non-Gaussian properties.

1.2. Empirical evidence

There exist many attempts in the literature to test the commodity price model assuming myopic expectations and using *ad hoc* distributed lags to track the serial correlation in the data. Ghosh, Gilbert, and Hughes Hallett (1987) provide a review of work along these lines. The rational expectations version of the model with speculative stockholding is conceptually more appealing, but it is numerically more complex. The problem is that there is no simple closed form for the equilibrium price function. Any test of the structural model has to rely on numerical approximations for $f(x)$, as in Newbery and Stiglitz (1982) and Williams and Wright (1991), or exploit the idea that $f(x)$ has an invariant

² The policy implications of conditional heteroskedasticity in terms of risk management is discussed in Beven, Collier, and Gunning (1987).

distribution, as in Deaton and Laroque (1992a), and then simulate the model to examine its reasonableness. Such modeling strategies also rely on assumptions about demand functions, the variance of the shocks, and the carrying cost. As Deaton and Laroque (1992a) have shown, there are qualitative differences in the simulation results for the linear and the iso-elastic demand case, with the latter generating more skewness and kurtosis and a higher frequency of stockout than the former. Furthermore, the data simulated under the maintained structural assumptions exhibit a much lower degree of serial correlation than the actual data.

An alternative to testing the model on data simulated from a given set of parameters is to obtain the full-information estimates of those parameters. To do so, it would be necessary to solve out $f(x)$ for each observation. Although numerical methods can be used to resolve this problem once a tractable demand function is assumed, the procedure is computationally costly. An attempt along these lines was undertaken by Deaton and Laroque (1992b) albeit with mixed success. While their storage model outperformed a model which implies prices are i.i.d. draws from a normal distribution, the structural model was also found to be beaten by a simple first-order autoregressive model. However, it is not clear whether the problem lies in the structural assumptions (such as linearity of the demand function) or with the theory itself.

There exists yet another way of testing the theory. The idea is not to impose assumptions about the demand function or solve for $f(x)$ but to test for consistency between the data and the reduced form model while maintaining only that there are two regimes governing the behavior of prices if the theory was true. For example, one can test the overidentifying restrictions imposed on the first moment condition (4) such as in Deaton and Laroque (1992a). Although this is a weaker test of the theory, reduced form estimations require less restrictive assumptions and are easier to implement.

The present analysis is in the same spirit as Deaton and Laroque (1992a) but has a broader scope. One cannot tell from testing the first moment condition alone whether it is the stockholding or the stockout part of the theory that is inadequate. Our analysis supplements this information by explicitly modeling the behavior of prices in both regimes. Such regime-specific information not only allows us to examine previously untested aspects of the theory, it also isolates features in the data not captured by the basic competitive storage model. Before proceeding to such an analysis, we shall perform a basic test of the theory, namely, that commodity prices are nonlinear processes.

2. Testing for nonlinearity

Theory suggests that the price of a commodity should have a conditional mean and variance that depend on whether or not inventories are held, and the

kink at p^* in turn implies a specific form of nonlinearity on the price process.³ To test for nonlinearity without prejudging its nature, we apply six tests to the data: the Neural-Network, Macleod–Li, Keenan, RESET, Tsay, and ARCH. Thus, given a series y_t , and $x_t = (y_{t-1} \dots y_{t-p})$, the null hypothesis of linearity is $y_t = x_t'\theta + e_t$.

The neural network statistic tests the null hypothesis against the alternative that $y_t = x_t'\theta + \sum_{j=1}^q \beta_j \psi(\gamma_j x_t) + e_t$. The idea is to extract nonlinearity using an augmented single hidden layer network as modeled by the logistic function ψ and random variables γ drawn from a uniform distribution. It can detect general forms of nonlinearity and is found to have good size and power in Lee, White, and Granger (1993).

The Macleod–Li statistic is based on the correlation coefficients between y_t^2 and its lags since for a linear stationary process, $\text{corr}(y_t^2, y_{t-k}^2) = \text{corr}(y_t, y_{t-k})^2$. The statistic will pick up nonlinearity in the conditional mean but will also detect nonlinearity in variance such as due to ARCH effects. The ARCH test is the by now familiar Lagrange multiplier test due to Engle (1982). It specifically tests for nonlinearity in the conditional variance.

The Keenan, RESET, and Tsay tests are F tests which detect departures from linearity in mean. Denoting $f_t = x_t'\hat{\theta}$, the Keenan test checks to see if f_t^2 has additional forecasting ability for y_t . Both the RESET and the Tsay tests are generalizations of this idea. The former checks to see if higher orders of f_t have additional forecasting power while the latter looks for the significance of cross-product terms such as $y_{t-j}y_{t-k}$, $j, k = 1 \dots p$, in the null model. Formal definitions of all six statistics are given in Lee et al. (1993).

The regressions used to produce the neural network and the RESET tests are susceptible to collinearity in the regressors. In practice, the problem is circumvented by using the main principal components (but not the largest) instead of all the regressors. We let $q = 15$ and $q^* = 2$ (principal components) in the neural network test, while the RESET test is based on one principal component. The Macleod–Li statistic has 15 degrees of freedom. The data are prewhitened, or the lag order chosen, where necessary, using the Akaike information criteria. The tests are robust to the choice of these parameters. Since the neural network test is based on randomly generated uniform variates, it is repeated several times with the same parameterization to ensure stability of the results. We therefore report three sets of the neural network test for reference. The statistics are programmed in Gauss and the exact size of each test is simulated to ensure that it is close to the nominal size.

The tests are applied to the thirteen commodity prices used in Deaton and Laroque (1992a). The data are indices of average commodity prices for each

³ The model predicts nonlinearity and conditional non-Gaussianity, but since features of nonnormality are induced by nonlinearity, we focus on identifying nonlinearity in the data.

Table 1
Tests for nonlinearity

Commodity	Neural $\chi^2(2)$	Keenan $F(1, 84)$	Reset $\chi^2(1)$	Tsay $F(1, 85)$	Arch(1) $\chi^2(1)$	TAR $F(1, 85)$
Bananas	0.50 2.25 5.31**	0.0005	0.03	0.0005	0.50	0.63
Cocoa	2.65 1.81 2.60	1.956	1.78	1.98	0.89	0.80
Coffee	2.83 9.96* 10.46*	4.78*	5.44*	4.83*	23.83*	0.82
Copper	8.60* 7.28* 12.10*	3.98*	4.06*	4.03*	1.52	1.45
Cotton	4.87** 5.85* 6.19*	1.08	0.82	1.92	0.53	1.14
Jute	7.40* 8.44* 0.88	7.47*	9.39*	7.56*	4.17*	2.68**
Maize	3.68 2.23 4.47**	2.98	1.64	3.02**	0.008	1.49
Palm oil	16.91* 10.98* 13.73*	18.26*	31.93*	18.48*	24.42*	2.59**
Rice	0.23 4.50 1.25	0.99	2.07	2.79*	5.30*	0.50
Sugar	14.70* 15.29* 16.71*	16.65*	5.05*	16.85*	1.66	5.18**
Tea	2.74 1.50 2.16	0.01	0.28	0.01	1.13	0.04
Tin	0.47 0.71 0.23	0.11	1.09	0.11	3.01**	0.29
Wheat	0.98 2.79 3.06	0.85	1.10	0.86	0.08	0.14

* (**) Denotes significance at the 5 (10) percent level.

calendar year, prepared by the World Bank Commodities Division, and are deflated by the U.S. consumer price index. The results are reported in Table 1. Tests which are statistically significant at the five and the ten percent levels are marked with one and two asterisks respectively. We cannot reject the null hypothesis of linearity at all in three series, namely, cocoa, tea, and wheat. Very weak evidence of nonlinearity can be detected in bananas, cotton, maize, rice, and tin. However, there is convincing evidence of nonlinearity in the price of coffee, copper, jute, palm oil, and sugar.

With the exception of the ARCH test, the statistics used essentially test the null hypothesis of linearity against an unspecified form of nonlinearity and are therefore portmanteau tests. Together with the small sample size on hand, one might expect the power of the tests to be low. Yet, our results still show signs of nonlinearity in many of the commodity prices examined. As nonlinearity is a prerequisite of the data if the theory of speculative stockholding was to hold, we therefore proceed to test the specific implications of the theory.

3. Threshold autoregressive models

A distinct prediction of the speculative storage model is that commodity prices are a nonlinear, two-regime, process with the switch in regime occurring at $p^* = (1 - \delta)/(1 + r)Ef(z)$. Although there exist many nonlinear time series models in the literature – see, for example, Priestley (1988) – it seems appropriate to work with a model that exploits the two-regime nature of the price function. This suggests using a Self Exciting Threshold Autoregressive model due to Tong (1978) and further developed in Tong and Lim (1980). A SETAR(r, q, c, d) model is summarized by a set of autoregressions:

$$p_t = \sum_{i=1}^q a_{ij}p_{t-i} + e_{jt} \quad \text{if} \quad c_{j-1} \leq p_{t-d} < c_j, \quad j = 1, \dots, r, \quad (5)$$

where $e_{jt} \sim \text{i.i.d.}(0, \sigma_j^2)$, c_j is the threshold characterizing regime j , p_{t-d} is the threshold variable, q is the order of autoregression, and r is the number of thresholds identifying the $r + 1$ regimes.⁴ Readers are referred to Tong (1990) for a systematic analysis of the SETAR model.

⁴The term 'regime switching' is used in a casual sense and is to be distinguished from the type of regime switching models such as developed by Hamilton (1989). In self-exciting threshold models, regime switching is determined by the value of the variable's own past, and a shift in regime is certain once the threshold level is reached. In Hamilton's model, regime shifting is probabilistic as determined by the fundamentals (possibly variables other than the own lag) underlying the Markov transition matrix.

Under the threshold principle, the parameters of the model are allowed to vary according to past values of p_t or values associated to lags of p_t . Hence, the terminology ‘self-exciting’. Within each regime, p_t is a linear autoregressive process, but p_t itself is the sum of r Gaussian processes and therefore exhibits non-Gaussian and nonlinear behavior. The threshold model given by (5) can therefore be regarded as a ‘piecewise linear approximation’ to a general, first-order, nonlinear model of the form $p_t = f(p_{t-1}) + e_t$.

While a general SETAR model has many parameters, namely, r , q , c , and d , the empirical model that we will use to analyze commodity prices is simpler than that written down in (5). This is because theory tells us what the values of r , q , and d should be. Rational expectations and white noise shocks dictate that d , the delay parameter, and q , the order of autoregression, be one. Since there are only two regimes possible, r is unity. Furthermore, since shocks are i.i.d., the threshold, c , is p^* , a constant. The model of interest is therefore a SETAR(1, 1, p^* , 1) with p^* as the threshold. If $p_{t-1} < p^*$, inventories are held and prices are intertemporally related. When $p_{t-1} \geq p^*$, p_t is entirely determined by harvest in period t . The empirical model which will be used to test the theory is

$$\begin{aligned}
 p_t &= a_1 + b_1 p_{t-1} + e_{1t} & \text{if } p_{t-1} \leq p^*, \\
 p_t &= a_2 + b_2 p_{t-1} + e_{2t} & \text{otherwise.}
 \end{aligned}
 \tag{6}$$

Under the null hypothesis, prices should be expected to rise in the stockholding regime but be history-independent in the stockout regime. This implies $b_1 > 1$ and $b_2 = 0$. Furthermore, prices should be higher and more variable in the stockout regime because of the absence of inventories to buffer excess demand. These are testable implications regardless of the assumptions about demand functions. The only maintained assumptions are that shocks to harvest are i.i.d. and that there are two regimes underlying the data.

It is interesting to note that (6) can be used to test if commodity markets are efficient in the sense that the conditional expectation of p_t from the two autoregressions in (6) should converge as $p_{t-1} \rightarrow p^*$. This ensures no arbitrage opportunities arising from jumps in the expectation of prices. A model with market efficiency imposed should therefore satisfy $a_1 + b_1 p^* = a_2 + b_2 p^*$. Note, however, that market efficiency so defined does not imply and is not implied by rational speculative stockholding which requires $b_1 > 1$ and $b_2 = 0$.

Tsay (1989) proposed a simple test of linearity against the alternative of threshold nonlinearity using a rearranged autoregression with the data ordered according to the threshold variable. The idea is that the orthogonality between the predictive residuals and the regressors in the recursive ordered autoregression will be destroyed if there is a regime change. We compute the test statistic for the thirteen series setting q , r , and d to unity. The results are reported in the last column of Table 1. The test finds evidence for threshold nonlinearity in three

series only, namely, jute, palm oil, and sugar. While such results are somewhat discouraging, it could be the case that the test has low power because of the small number of observations. To further investigate the properties of the data, we proceed to estimate the parameters for the two regimes.

4. Estimation issues

In a stationary SETAR context, the state is always known at time t since the value of p_{t-1} determines the regime the series will belong each period, and hence the process driving the series in that period. Given this structure, normality of the error process, and stationarity of the regime-specific autoregressions, least squares should produce estimates of the autoregressive parameters that are \sqrt{T} -consistent and asymptotically normal. This result was formally proved in Chan (1993).

In practice, estimation of the SETAR parameters is quite involved especially when the threshold value is unknown because there are many free parameters.⁵ Although it is not necessary to estimate r , d , and q in our case because they are already pinned down by theory, our empirical model also has two nonstandard features. The first arises from the fact that the errors are conditionally heteroskedastic in the regime in which stocks are held. The second relates to the fact that the autoregressive coefficient in the stockholding regime exceeds one under the null hypothesis.

As for the first problem, omitting conditional heteroskedasticity should be an issue of efficiency and not consistency, and one would expect approaches developed for homoskedastic errors to continue to be valid when the errors are conditionally heteroskedastic.⁶ As for the second problem, little is known about the properties of the least squares estimator when the root of the autoregression is unstable in one or more regimes. Pham, Chan, and Tong (1991) show in the case of two regimes that the least squares estimator is strongly consistent when one and/or both regimes have a unit root with d and r known. Results for more general forms of nonergodicity are unavailable, but Pham et al. (1991) noted that 'recurrence' might be necessary for strong consistency to hold when there is nonstationarity of other forms in some of the regimes.

⁵ Tong and Lim (1980) proposed an iterative search over r , q , c , and d for the best model. Tsay (1989) proposed a graphical method that is computationally less demanding.

⁶ Commodity prices are conditionally heteroskedastic in its own lag and cannot, strictly speaking, be parameterized as an ARCH or a GARCH process. We cannot therefore appeal to the results of Gouriéroux and Monfort (1992) for qualitative threshold models (QTARCH) and Zakoïan (1994) for threshold ARCH and GARCH models.

The statistical result that is ultimately required for our analysis is consistency of the estimator with unknown thresholds and with possible nonstationarity and conditional heteroskedasticity in one of the two regimes. The issue of whether the estimator converges at the usual rate of \sqrt{T} or faster is also an issue to be resolved. The ambiguity arises because commodity prices are renewal processes which are ergodic by assumption. Since prices in the nonstationary regime are bounded from above by p^* , the sample moments of the process in that regime will not accumulate as fast as those of a strictly nonstationary process. Therefore, although there is nonstationarity in one regime, results relevant for nonergodic and nonstationary processes would seem overly stringent. A formal statistical analysis would be beyond the scope of this paper, but we will use simulations to examine the properties of the quasi-maximum-likelihood estimator under the assumptions of the speculative storage model. This numerical analysis is carried out using a Kalman filter framework, to which we now turn.

4.1. The Kalman filter

In this section, we use a Kalman filter to set up the two-regime autoregressions. The threshold value is a parameter to be estimated by maximizing the likelihood function of the threshold model. The approach is more mechanical than the graphical approach of Tsay (1989) and seems appropriate for modeling a battery of series such as in this study. We will cast the two-regime threshold model in state space form and will follow the analysis of Harvey (1981) closely. Let $\alpha_t = (1, p_t)$ be the state vector and $Z = [0 \ 1]$ be a fixed matrix. The measurement and transition equations are

$$p_t = Z\alpha_t,$$

$$\alpha_t = T_{t|t-1}\alpha_{t-1} + e_t.$$

The transitional matrix at time t depends on the state in $t - 1$, hence denoted $T_{t|t-1}$. More precisely, if $p_{t-1} \leq p^*$, the transition equation is

$$\begin{bmatrix} 1 \\ p_t \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ a_1 & b_1 \end{bmatrix} \begin{bmatrix} 1 \\ p_{t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ e_{1t} \end{bmatrix}, \quad Ee'_{1t}e_{1t} = \sigma_1^2.$$

If $p_{t-1} > p^*$, the transition equation is

$$\begin{bmatrix} 1 \\ p_t \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ a_2 & b_2 \end{bmatrix} \begin{bmatrix} 1 \\ p_{t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ e_{1t} \end{bmatrix}, \quad Ee'_{2t}e_{2t} = \sigma_2^2.$$

Let $\hat{\alpha}_{t-1}$ be the minimum mean-squared estimator for α given all the information up to $t - 1$ with covariance $\sigma_{t-1}^2 V_{t-1}$, where σ_{t-1}^2 is either σ_1^2 or σ_2^2 depending on the regime. Then the prediction equations given information at $t - 1$ are

$$\hat{\alpha}_{t|t-1} = T_{t|t-1} \hat{\alpha}_{t-1},$$

$$V_{t|t-1} = T_{t|t-1} V_{t-1} T'_{t|t-1} + \sigma_{t-1}^2.$$

The updating rule is given by

$$\hat{\alpha}_t = \hat{\alpha}_{t|t-1} + V_{t|t-1} z_t (p_t - z'_t \hat{\alpha}_{t|t-1}) / f_t,$$

$$v_t = p_t - z'_t \hat{\alpha}_{t|t-1},$$

$$f_t = z'_t V_{t|t-1} z_t.$$

The objective is to maximize the pseudo-log-likelihood function

$$L = -\frac{1}{2} \log 2\pi - \frac{1}{2} \sum_{t=1}^T \log \sigma_t^2 f_t - \frac{1}{2} \sum_{t=1}^T \log v_t^2 / (\sigma_t^2 f_t^2). \quad (7)$$

The threshold, p^* , enters the objective function since the transition matrix in each period depends on the state, and hence on the value of p^* . We have seven parameters to estimate, namely, $a_1, b_1, \sigma_1^2, a_2, b_2, \sigma_2^2$, and p^* .

A key feature of the commodity price model is that there is a kink in the price process at the threshold value. Accordingly, the likelihood function is nondifferentiable at the kink point. The approach taken here is to use the downhill simplex method of Nelder–Mead which is derivatives-free. Briefly, if there are n parameters to be estimated, the algorithm requires as input $n + 1$ sets of initial values to form a simplex. The algorithm seeks to move a point in the simplex towards the direction where the objective function is minimized (in this case, the negative of L). The algorithm is repeated several times with updated starting values to ensure that a global maximum is achieved. The routine is taken from Press, Teukolsky, Vetterling, and Flannery (1992) and recoded in Gauss.

Since the algorithm is derivatives-free, it does not compute the gradient or the Hessian. In consequence the standard errors for the estimates have to be computed by alternative means. As a matter of practice, this is achieved as follows. Given a set of initial values, the Nelder–Mead algorithm provides an estimate of p^* . This allows us to separate the sample according to \hat{p}^* and construct the regime-specific moment matrix of regressors and regression

residuals. These are then used to construct the heteroskedastic-consistent standard errors of White (1980).⁷

4.2. A Monte Carlo experiment

As noted earlier, the theoretical properties of the quasi-maximum-likelihood estimator for a model with features such as ours are not known. We therefore perform a Monte Carlo experiment to assess the adequacy of the estimator. To this end, we simulated data using the model:

$$p_t = 0.1 + 1.02p_{t-1} + \varepsilon_{1t} \quad \text{if } p_{t-1} \leq p^* = 0.7,$$

$$\varepsilon_{1t} = h_t e_{1t}, \quad h_t^2 = 0.1 + 0.4\varepsilon_{1,t-1}^2,$$

$$p_t = 0.812 + e_{2t} \quad \text{otherwise,}$$

$$e_{1t}, e_{2t} \sim N(0, 1).$$

This is a two-regime threshold model with an explosive root and ARCH effects in one regime. It is set up to broadly mimic the theoretical properties of commodity prices. The price process is continuous at p^* by construction. We set the sample size to 88, the length of the commodity price series used in the empirical analysis, and estimated the model 100 times using the Nelder–Mead algorithm to maximize the quasi-maximum-likelihood function described in the previous section. The mean and standard error of the estimates in the 100 simulations are as follows:

	a_1	b_1	a_2	b_2	σ_1	σ_2	p^*
Mean	0.1076	0.9638	0.7255	0.0417	0.3779	1.0252	0.7263
S.D.	0.0722	0.0914	0.3058	0.0391	0.0604	0.3345	0.1131

The true values of the parameters in the autoregressions are within one standard error of the estimates. Of special interest is that p^* is estimated with high precision even with the small sample size on hand. It has been shown in

⁷ This is as though we run an autoregression for each regime using p^* to split the sample. The least squares estimates of the two autoregressions provide a check for the reasonableness of the quasi-maximum-likelihood estimates. By Theorem 2.1 of Tsay (1989), given the threshold value, the least squares estimator should converge almost surely to the true, regime-specific, parameters. Thus, conditional on p^* , the quasi-maximum-likelihood estimates should be close to those of the conditional least squares estimates because the sample is divided on the basis of the same threshold value. However, the conditional least squares approach takes p^* as given but the maximum-likelihood approach treats p^* as parametric. In practice, the two sets of estimates are very similar.

Chan (1993) that for stationary ergodic cases, the least squares estimator for p^* is super (order T) consistent. The intuition for that result is that given a series ordered according to p_{t-a} , p^* can be identified from the observations in the neighborhood of p^* , the location of the discontinuity of the autoregressive function. Precise estimates can therefore be obtained without using data on p_t over a large range. This allows p^* to be consistently estimated at an accelerated rate. It appears that this fast rate of convergence will continue under the stochastic assumptions of our model, although this conjecture needs to be formalized.

The regression model used in the Monte Carlo analysis has ignored conditional heteroskedasticity in one of the regimes and the estimated variances resemble the unconditional variance of the respective regimes. The issue of interest is the properties of \hat{b}_1 , \hat{b}_2 , and \hat{p}^* when conditional heteroskedasticity is ignored, and our results suggest that these parameters are still estimated with high precision. The result that the gain in modeling conditional heteroscedasticity is one of efficiency rather than consistency appears to carry over from piecewise linear to threshold models.⁸

There is an explanation for our seemingly favorable Monte Carlo results. If the autoregressive coefficient was outside the unit circle in the stockout rather than the stockholding regime, then prices would grow forever and would exceed p^* with probability one. Ergodicity would be violated. But such is not the case here. Nonstationarity occurs in the autoregression for the regime in which stocks are held. Since carrying costs are positive, prices will rise (in expectations) and will eventually hit p^* . Thus, there must exist periods when no speculative inventory is held. The market-clearing mechanism is such that prices are prevented from moving in one direction forever. This makes commodity price a renewal and an ergodic process with a nonzero probability of falling into one of the two regimes. For this reason, although there is nonstationarity in one regime, the empirical properties of the quasi-maximum-likelihood estimator behaves as though the price process is stationary and ergodic.

4.3. *Estimates for the commodity prices*

The quasi-maximum-likelihood estimates for the model without market efficiency imposed are reported in Table 2. Evidence of regime-specific behavior can be seen from the fact that the conditional mean in the stockholding regime is uniformly lower than in the stockout regime. With the exceptions of cocoa, rice, and tea, the variances are also higher in the stockout regime.

Using the estimated values of p^* to separate the sample, the highest incidence of stockout is found in the case of cotton (twenty-five percent) and the lowest in

⁸ Generalizing the model to handle GARCH effects within regimes is, in theory, straightforward. In practice, this additional source of nonlinearity poses convergence problems at the optimization level.

Table 2
 Estimation of threshold autoregressive model
 $p_t = a_1 + b_1 p_{t-1} + e_{1t}$, if $p_{t-1} < p^*$, $p_t = a_2 + b_2 p_{t-1} + e_{2t}$ otherwise.

Commodity (like)	a_1	b_1	a_2	b_2	\hat{p}_1	σ_1	p_2	σ_2	p^*
	s.e.	s.e.	s.e.	s.e.	stockin	stockin	stockout	stockout	% stockout
Banana	-0.0003	1.0017	0.2467	0.6279	0.561	0.038	0.703	0.035	0.7036
232.4	(0.030)	(0.052)	(1.507)	(2.048)			10.3		
Cocoa	-0.0005	1.0270	-0.1603	1.1202	0.182	0.050	0.335	0.020	0.3431
-202.8	(0.015)	(0.102)	(0.280)	(0.556)			8.0		
Coffee	0.031	0.9965	0.5087	-0.1781	0.209	0.050	0.423	0.119	0.4060
-198.7	(0.017)	(0.093)	(0.664)	(1.154)			8.0		
Copper	-0.0005	1.0272	0.7257	0.0000	0.441	0.063	0.725	0.192	0.6619
-180.3	(0.029)	(0.070)	(0.861)	(1.013)			16.0		
Cotton	0.0151	0.9840	0.2576	0.6832	0.562	0.080	0.891	0.147	0.7866
-152.5	(0.037)	(0.063)	(0.477)	(0.519)			25.2		
Jute	0.0000	1.0350	0.8188	-0.0106	0.552	0.114	0.809	0.182	0.8007
-127.8	(0.055)	(0.108)	(0.757)	(0.754)			18.3		
Maize	-0.0000	1.0309	0.1892	0.6896	0.665	0.110	1.019	0.367	0.9884
-124.1	(0.051)	(0.081)	(1.070)	(0.714)			13.7		
Palm	-0.0007	1.0040	0.6385	0.3569	0.496	0.089	1.054	0.508	0.8128
-144.2	(0.040)	(0.079)	(1.365)	(1.252)			8.0		
Rice	0.0008	1.0030	-0.8412	1.4384	0.608	0.109	0.851	0.078	1.133
-139.7	(0.038)	(0.062)	(3.820)	(3.039)			3.4		
Sugar	0.0136	1.0701	1.1115	-0.0000	0.580	0.271	1.115	0.376	1.04
-63.2	(0.091)	(0.217)	(0.707)	(0.508)			22.9		
Tea	0.0030	1.0049	-0.0459	0.9900	0.473	0.080	0.635	0.055	0.6032
-172.4	(0.065)	(0.135)	(0.179)	(0.257)			22.9		
Tin	0.0114	0.9620	-0.0457	1.0168	0.205	0.039	0.451	0.058	0.4241
-223.0	(0.020)	(0.113)	(0.948)	(2.019)			6.8		
Wheat	0.1007	0.8413	0.7914	0.4145	0.666	0.120	1.393	0.223	1.194
-127.8	(0.037)	(0.057)	(6.341)	(4.146)			3.4		

Table 3
 Estimation of threshold autoregressive model
 $p_t = a_1 + b_1 p_{t-1} + e_{1t}$ if $p_{t-1} < p^*$,
 $p_t = a_2 + b_2 p_{t-1} + e_{2t}$ otherwise.
 Restriction: $a_1 + b_1 p^* = a_2 + b_2 p^*$.

Commodity (like)	a_1 s.e.	b_1 s.e.	a_2 s.e.	b_2 s.e.	\bar{p}_1		\bar{p}_2		p^* % stockout
					stockin	stockout	stockin	stockout	
Banana 232.2	0.0004 (0.039)	1.0019 (0.070)	0.4364 (0.393)	0.3678 (0.555)	0.554	0.040	0.699	0.026	0.6876
Cocoa -191.5	-0.0007 (0.014)	0.9625 (0.092)	0.1690 (0.900)	0.5707 (1.564)	0.188	0.060	0.465	0.044	0.4331
Coffee -198.6	-0.0064 (0.017)	1.0130 (0.092)	0.3703 (0.570)	-0.1114 (0.881)	0.209	0.050	0.423	0.117	0.4114
Copper -179.9	-0.0007 (0.029)	1.0324 (0.073)	0.6845 (0.886)	-0.0013 (1.076)	0.441	0.063	0.725	0.188	0.6651
Cotton -152.5	0.0145 (0.037)	0.9856 (0.063)	0.2393 (0.477)	0.7004 (0.518)	0.562	0.080	0.891	0.149	0.7881
Jute -127.8	0.0000 (0.055)	1.0339 (0.108)	0.8703 (0.743)	-0.0619 (0.746)	0.552	0.114	0.809	0.182	0.7942
Maize -123.6	0.0034 (0.051)	1.0236 (0.081)	0.6293 (1.234)	0.3898 (0.873)	0.665	0.110	1.019	0.382	0.9884
Palm -144.1	-0.0009 (0.040)	1.0049 (0.079)	0.4273 (1.401)	0.4780 (1.353)	0.496	0.089	1.054	0.515	0.8128
Rice -135.7	0.0455 (0.039)	0.9326 (0.065)	-1.2862 (24.38)	-1.1620 (20.87)	0.608	0.110	0.851	0.358	1.133
Sugar -63.2	-0.0069 (0.091)	1.0984 (0.216)	1.1229 (0.708)	-0.0000 (0.509)	0.580	0.271	1.111	0.378	1.02
Tea -168.0	0.0003 (0.069)	0.9684 (0.144)	-0.0112 (0.222)	0.9503 (0.317)	0.473	0.080	0.635	0.064	0.6032
Tin -221.0	0.0020 (0.021)	0.9886 (0.114)	-0.1099 (1.545)	1.2500 (3.265)	0.205	0.040	0.451	0.081	0.4241
Wheat -127.0	0.0701 (0.037)	0.8870 (0.058)	0.8042 (9.185)	0.2725 (6.006)	0.666	0.122	1.393	0.298	1.194

the case of rice and wheat (three percent), but there is evidence of infrequent stockouts in every commodity. The estimated values of p^* are in line with those reported in Deaton and Laroque (1992a) with the exception of cotton and palm oil. In their analysis, palm oil is estimated to have the lowest (one percent) incidence of stockout. Our estimated value of p^* is less extreme and suggests stockouts occurring in the palm oil market eight percent of the time. For the case of cotton, our estimate of p^* is lower than theirs and implies stockouts happening twenty-five percent of the time. By contrast, Deaton and Laroque (1992a) reports a more modest eight percent incidence of stockout. Thus, cases when the results of the two analyses do not accord pertain to situations when the frequency of stockouts is estimated to be either extremely high or extremely low.

An implication of the speculative storage model is that speculators expect prices to increase since carrying costs are positive. The point estimate of b_1 exceeds one in nine cases with the estimated coefficient in the remaining four cases within one standard deviation of unity. If stockouts are indeed infrequent events, the number of observations in the stockout regime will be small, and the standard errors for the estimates in that regime will likely be large. The parameters in the stockout regime will not be estimated with as much precision as those of the stockholding regime. We therefore also consider the absolute size of the estimates in assessing whether the stockout regime has zero degree of persistence with $\hat{b}_2 = 0$. The estimates suggest that in nine of the thirteen cases, strong persistence is found in the stockout regime. This is one aspect of theory most rejected by the data.

The empirical properties of commodity prices can therefore be subdivided into three categories depending on the values of \hat{b}_1 and \hat{b}_2 . For coffee, copper, jute, and sugar, \hat{b}_1 is close to or exceed one and \hat{b}_2 is statistically and numerically insignificant. All aspects of the theory are therefore supported. For banana, cotton, maize, palm, and wheat, \hat{b}_2 is less than one and in some cases statistically insignificant, but the validity of the theory is still of some concern because the numerically large estimates for b_2 imply nonnegligible persistence in the stockout regime. As for cocoa, rice, tea, and tin, the values of \hat{b}_1 and \hat{b}_2 are close to or exceed one. This is incompatible with an assumption underlying our Monte Carlo experiment and the estimates should therefore be interpreted with some caution. However, taking the estimates at face value would imply that these four price processes are not ergodic and cannot therefore be renewal processes as predicted by theory.

Turning now to the results reported in Table 3 with market efficiency imposed, the constrained estimates generally suggest a lower degree of persistence in the stockout regime than in the model without market efficiency imposed. The data for coffee, copper, jute, and sugar continue to be the only commodities that show decisive support for the speculative storage model. The efficient market hypothesis is rejected in four case: cocoa, rice, tea, and tin. Of

these, the price of tea and tin are also strongly persistent in the stockout regime. This is inconsistent with rational storage behavior.

Together with the results of the nonlinearity tests, coffee, copper, jute, and sugar appear to be the only commodities supporting all aspects of the theory. The remaining prices either exhibit weak or no evidence of nonlinearity, or the estimates of \hat{b}_2 are more persistent than predicted by theory. Interestingly, cases when evidence for nonlinearity and market efficiency is weakest are also cases when \hat{b}_2 is both numerically and statistically significant and not too different from \hat{b}_1 . Indeed, a key distinction between the regimes would be lost if the two regimes have the same degree of persistence. Unless the intercepts in the two autoregressions are sufficiently different, the processes defining the two regimes would be approximately identical, in which case, the price process can be viewed as piecewise linear. The failure to reject linearity in cocoa, tea, and tin reinforces the similarities in the stochastic properties of the two regimes.

5. Concluding comments

The assumption of i.i.d. shocks to harvest imposes a strong identification restriction on the behavior of prices across regimes. Prices should be autocorrelated in one regime but not the other. The finding that in eight of the thirteen cases, the autoregressive coefficient in the stockout regime is numerically and/or statistically significant suggests that prices are serially correlated whether or not stocks are held. Indeed, the degree of persistence found in the stockout regime is rather strong; the autoregressive coefficient exceeds 0.8 in many cases. Given that the data exhibit such high level of serial correlation, a model which constrains the stockout regime to have zero persistence is clearly inappropriate. It is therefore not surprising that Deaton and Laroque (1992b) found a one-regime AR(1) model to track the data better than a two-regime model with no persistence in one regime.

Serial correlation in commodity prices can arise for a variety of reasons. One explanation in the case of tree crops such as cotton and tea is the long gestation lag between planting and harvest. A supply shock can lead to prolonged periods of excess demand which in turn induces autocorrelation in prices. In the same vein, cyclical but serially correlated shocks to the demand for industrial commodities such as tin will also generate a similar persistent effect. In these cases, the assumption that shocks to commodity markets are i.i.d. will not be appropriate.

This paper suggests further work in two areas. The first is a rigorous statistical analysis to formalize the properties of our threshold model. This would require integrating results for nonstationary variables and conditional heteroskedasticity with those of SETAR models. The results will be useful not just for analyzing commodity prices, but more generally when threshold modeling of

nonstationary and conditional heteroskedastic data is appropriate. The second line of research is to generalize the speculative storage model to allow for serial correlation, establish the theoretical properties, and test the model predictions. On this, some progress has been made.

The theoretical properties of commodity prices when shocks are time-dependent is the subject of research in a recent paper by Chambers and Bailey (1993). A conclusion from that analysis is that while p_t will still be ergodic, p^* will no longer be a constant once the i.i.d. assumption is relaxed. From the standpoint of hypothesis testing, the time-varying nature of p^* raises an econometric problem since prices in both regimes will be serially correlated whether or not speculative inventories are held. The feature which allows us to identify the two regimes in the i.i.d. case is no longer valid. One would need to isolate the effects of serially correlated shocks from persistence due to speculative stockholding to test the competitive storage model. This issue is currently under investigation by the author.

References

- Bevan, D.L., P. Collier, and J. Gunning, 1987, Consequences of a commodity boom in a controlled economy: Accumulation and redistribution in Kenya 1975–83, *World Bank Economic Review* 1, 489–514.
- Chambers, M.J. and R. Bailey, 1993, Commodity prices subject to time dependent and heterogeneous disturbances, *Mimeo.* (University of Essex, Essex).
- Chan, K.S., 1993, Consistency and limiting distribution of the least squares estimates of a threshold autoregressive model, *Annals of Statistics* 32, 520–533.
- Deaton, A.S. and G. Laroque, 1992a, On the behavior of commodity prices, *Review of Economic Studies* 89, 1–23.
- Deaton, A.S. and G. Laroque, 1992b, Estimating the commodity price model, *RPDS working paper* 154 (Princeton University, Princeton, NJ).
- Engle, R., 1982, Autoregressive conditional heteroskedasticity with estimates of variance of U.K. inflation, *Econometrica* 50, 987–1007.
- Ghosh, S., V. Gilbert, and A. Hughes Hallett, 1987, *Stabilizing speculative commodity markets* (Clarendon, Oxford).
- Gilbert, C.L., 1993, *Modelling market fundamentals: A model of the aluminum market*, Working paper series of the network in financial markets 30 (European Science Foundation, London).
- Gourieroux, C. and A. Monfort, 1992, Qualitative threshold ARCH models, *Journal of Econometrics* 52, 159–199.
- Gustafson, R.L., 1958, Carryover levels for grains, *Technical bulletin* 1178 (U.S. Department of Agriculture, Washington, DC).
- Hamilton, J., 1989, A new approach to the economic analysis of non-stationary time series and the business cycle, *Econometrica* 57, 357–384.
- Lee, T.H., H. White, and C.W.J. Granger, 1993, Testing for neglected nonlinearity in time series models, *Journal of Econometrics* 56, 269–290.
- Newbery, D. and J. Stiglitz, 1982, Optimal commodity stockpiling rules, *Oxford Economic Papers* 34, 403–427.

- Pham, D.T., K. Chan, and H. Tong, 1991, Strong consistency of the least squares estimator for a non-stationary threshold autoregressive model, *Statistica Sinica* 1, 361–369.
- Press, W.H., S. Teukolsky, W. Vetterling, and B. Flannery, 1992, *Numerical recipes in C*, 2nd ed. (Cambridge University Press, Cambridge).
- Samuelson, P., 1957, Intertemporal price equilibrium: A prologue to the theory of speculation, *Weltwirtschaftliches Archiv* 79, 181–219.
- Tong, H., 1990, *Non-linear time series: A dynamical system approach* (Oxford Science Publications, Oxford).
- Tong, H. and K. Lim, 1980, Threshold autoregression, limit cycles and cyclical data, *Journal of the Royal Statistical Society B* 42, 245–292.
- Tsay, R., 1989, Testing and modeling threshold autoregressive processes, *Journal of the American Statistical Association* 84, 231–240.
- White, H., 1980, A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* 48, 817–838.
- Williams, J. and B. Wright, 1991, *Storage and commodity markets* (Cambridge University Press, Cambridge).
- Williams, J.B., 1937, Speculation and the carryover, *Quarterly Journal of Economics* 50, 436–455.
- Zakoïan, J., 1994, Threshold heteroskedastic models, *Journal of Economic Dynamics and Control* 18, 931–955.