An Empirical Reassessment of the Commodity Storage Model

Mario J. Miranda
and
Xiongwen Rui *

February 20, 1999

Abstract

We perform an empirical assessment of the rational expectations commodity storage model, with the specific aim of resolving the model's apparent inability to explain the high autocorrelation historically exhibited by primary commodity prices. In contrast to recent empirical tests of the model, we employ a classical supply of storage formulation to explain speculative and precautionary storage. To estimate the model, we develop a nested maximum likelihood estimation procedure that relies on numerical nonlinear functional equation methods to solve for the implied rational expectations equilibrium. Our results indicate that the rational expectations storage model can explain the autocorrelation exhibited by commodity prices extremely well, challenging recent findings to the contrary.

*The authors are Associate Professor and Postdoctoral Fellow, respectively, The Ohio State University, 2120 Fyffe Road, Columbus, OH, 53711, USA. Senior authorship is shared equally.

1 Introduction

In recent years, there has been a resurgence of interest in commodity price and storage dynamics following the development of the modern rational expectations commodity storage model (Williams and Wright). The central tenet of the modern storage model is that commodity price dynamics are governed mainly by the speculative and precautionary storage activity of rational commodity storers and processors. Modern storage theory can trace its origins to the classical "supply of storage" theory first introduced by Williams (1935) and later extended by Kaldor (1939) and Working (1949). Recent work on the rational expectations storage model includes papers by Salant (1983), Sheinkman and Schectman (1983), Wright and Williams (1982, 1984), and Miranda and Helmberger (1988).

Despite the intense interest in the rational expectations commodity storage model, few publications have been devoted to its empirical estimation and validation. Econometric estimation of the model has been hampered mainly by the absence of an analytically tractable closed form solution. Recently, however, Deaton and Larocque (1992, 1994), in a series of pathbreaking papers, employed numerical methods to estimate the rational expectations storage model and to assess its ability to explain the stylized facts of primary commodity price dynamics. Deaton and Larocque's analyses generated a series of interesting, but ultimately troubling, results: The rational expectations storage model, they concluded, can adequately explain the high volatility and positive skewness historically exhibited by primary commodity prices. The model, however, cannot adequately explain their high autocorrelation.

This paper is devoted to an empirical reassessment of the modern rational expectations commodity storage model, with the specific aim of resolving the model's apparent inability to explain the high autocorrelation historically exhibited by primary commodity prices. In so doing, we challenge Deaton and Larocque's assumption that storage costs arise solely from decay at a fixed rate. We argue that this formulation is unrealistic, largely because it fails to capture the coincidence of negative spot-futures price spreads and positive stock levels, a common occurrence in practice. As an alternative, we posit a classical supply of storage function, which admits negative intertemporal price spreads at positive stocks. With this formulation, we find that the rational expectations storage model can explain the high autocorrelation of commodity prices extremely well, reversing Deaton and Larocque's negative results.

In our empirical analysis, we encounter a series of methodological challenges that we address by developing novel econometric estimation techniques. The absence of reliable quantity data renders our rational expectations storage model unestimable by the generalized method of moments. As an alternative estimation strategy, we turn to full information maximum likelihood methods. In this context, however, we encounter the difficulty that rational expectations imposes a series of parametric restrictions that take the form of an infinite-dimensional nonlinear functional equation. To address this problem, we employ the orthogonal polynomial collocation method recently introduced to economists by Judd. The orthogonal collocation method replaces the functional equation with an approximating finite-dimensional system of nonlinear algebraic equations, converting the likelihood maximization problem into an unconstrained nonlinear programming problem with differentiable objective function.

Our paper is arranged as follows. In section 2, we present the modern rational expectations commodity storage model and discuss alternative supply of storage specifications. In section 3, we discuss how numerical functional equation methods can be used to derive the Markov probability transition rule followed by commodity prices under rational expectations. In section 4, we simulate the storage model under different parametric specifications and explore how the supply of storage affects autocorrelation in prices. In section 5, we discuss how to estimate the rational expectations storage model using nested full information maximum likelihood methods. In section 6, we empirically estimate the rational expectations storage model using historical price data for thirteen commodities and confirm that the model, when endowed with a classical supply of storage function, can satisfactorily explain high autocorrelation in commodity prices. We conclude the paper with a discussion of possible model extensions.

2 The Modern Storage Model

The centerpiece of the modern theory of storage is the competitive intertemporal arbitrage equation:

$$(1) \qquad \frac{1}{1+r}E_t p_{t+1} - p_t = c_t.$$

The intertemporal arbitrage equation asserts that, in equilibrium, expected appreciation in the commodity price p_t, discounted at the interest rate r, must equal the marginal cost of storage c_t. Dynamic equilibrium in the commodity market is enforced by competitive expected-profit-maximizing storers. Whenever expected appreciation exceeds the marginal cost of storage, the attendant profits motivate storers to increase their stockholdings until the equilibrium is restored. Conversely, whenever the cost of storage exceeds expected appreciation, the attendant loses motivate storers to decrease their stockholdings until the equilibrium is restored. Stockholding links supply and demand across time, inducing serial dependence in prices, even when production and consumption are serially independent.

The modern storage model is completed by the introduction of demand and production functions, a supply of storage function, a market clearing condition, and a theory of how price expectations are formed.[1] Denote available supply at the beginning of period t by s_t, quantity consumed in period t by q_t, stocks at the end of period t by x_t, and new production at the beginning of period t by y_t. In this paper, as in Deaton and Larocque's papers, we work with the simplest version of the storage model. Specifically, in addition to (1) we assume:

- available supply is the sum of initial stocks and new production:

$$s_t = x_{t-1} + y_t; \tag{2}$$

- available supply is either consumed or stored:

$$s_t = q_t + x_t; \tag{3}$$

- the market clearing price is a decreasing function $p(\cdot)$ of the quantity consumed:

$$p_t = p(q_t); \tag{4}$$

- the marginal cost of storage is a function $c(\cdot)$ of the quantity stored:

$$c_t = c(x_t); \tag{5}$$

- production y_t is exogenous, stochastic, and independently and identically distributed over time;
- and expectations are formed rationally in the sense of Muth.

Different versions of the modern rational expectations storage model assume different forms for the marginal cost, or "supply", of storage function $c(\cdot)$. Models that draw directly from the classical supply of storage literature assume that the marginal cost of storage comprises a marginal physical cost of storage and a marginal "convenience yield" (Williams; Kaldor; Working). According to classical supply of storage theory, the marginal convenience yield represents the amount processors are willing to pay to avoid the cost of revising their production schedules plus the option value of being in a position to take advantage of potential price increases (see in figure 1). If stock levels are high, the convenience yield is zero and the storage cost equals the physical storage cost, which is positive. As stock levels approach zero, however, the marginal convenience yield rises, eventually resulting in a negative storage cost. The classical supply of storage function has received strong empirical support over the years (e.g., Brennan; Fama and French).[2]

In contrast to classical supply of storage theory, Deaton and Larocque posit a simple storage technology in which storage costs arise exclusively from stock spoilage or decay at a fixed rate γ. Under this formulation, one unit of commodity stored in the current period will yield $(1-\gamma)$ units in the following period. This effectively implies that the marginal cost of storage is a constant multiple of the expected future price:

$$(6) \qquad c(x_t) = \frac{\gamma}{1+r} E_t p_{t+1}.$$

The expected price, however, falls with the level of stocks. Thus, under the constant decay assumption, the supply of storage function is decreasing in the stock level x_t. The constant decay supply of storage function is illustrated in figure 1. The main difference between the constant decay and classical supply of storage functions is evident: for low stock levels, the former is decreasing and positive, the latter is increasing and negative.

3 Deriving the Price Process

The commodity price p_t in the modern rational expectations storage model (1)-(5) follows a stationary first-order Markovian stochastic process. More specifically, price p_t in period t is a function of the lagged price p_{t-1} and the contemporaneous realization of an exogenous i.i.d. driving process, namely production y_t. To establish this fact and to uncover the Markov probability transition rule, let $\lambda(\cdot)$ denote the function that gives the equilibrium price p_t implied by the model for a given supply s_t. Given the equilibrium price function $\lambda(\cdot)$, the Markov probability transition rule that governs equilibrium market price dynamics can be derived as follows:

$$
\begin{aligned}
(7) \qquad p_t &= \lambda[s_t] \\
&= \lambda[x_{t-1} + y_t] \\
&= \lambda[s_{t-1} - q_{t-1} + y_t] \\
&= \lambda[\lambda^{-1}(p_{t-1}) - p^{-1}(p_{t-1}) + y_t].
\end{aligned}
$$

Here, we use the fact that $p_{t-1} = \lambda(s_{t-1}) = p(q_{t-1})$, where s_{t-1} is supply, q_{t-1} is consumption, and $x_{t-1} = s_{t-1} - q_{t-1}$ is carryout in period $t-1$.

To fully describe the Markov transition rule followed by the commodity price, one must derive the equilibrium price function $\lambda(\cdot)$. The equilibrium price function is characterized by a nonlinear functional equation, which stipulates that for every realizable supply s

$$(8) \quad \lambda(s) = p(s - x)$$

where x solves

$$(9) \quad \frac{1}{1+r} E_y \lambda(x + y) - p(s - x) = c(x).$$

Under mild regularity conditions, the functional equation (8)-(9) has an unique solution $\lambda(\cdot)$ (Scheinkman and Schectman). The equilibrium price function $\lambda(\cdot)$, however, does not generally possess a closed form.

Figure 2 illustrates the equilibrium price function $\lambda(\cdot)$ and its relation to the inverse consumption demand function $p(\cdot)$, which for the purposes of illustration is assumed to be linear. As can be seen in figure 2, for low initial supply, high prices discourage storage and supply is entirely consumed in the current period. Accordingly, the equilibrium price function $\lambda(\cdot)$ and the inverse demand function $p(\cdot)$ essentially coincide. For higher supplies, however, low prices encourage stockholding, raising prices above what they would otherwise have to be for consumers to completely clear the market of available supplies. Accordingly, the equilibrium price function $\lambda(\cdot)$ lies above the inverse demand function $p(\cdot)$.

Although an exact representation of the equilibrium price function typically cannot be derived analytically, an arbitrarily precise approximation can always be derived using numerical functional equation techniques. One particularly efficient technique for computing an approximation for the equilibrium price function is the Chebychev orthogonal collocation method. The efficacy of the Chebychev orthogonal collocation method is guaranteed by the Chebychev approximation theorems, which assert that a continuous function can be optimally approximated to any degree of accuracy by the polynomials that interpolate the function at the Chebychev nodes. Chebychev approximation is discussed in most numerical analysis textbooks (e.g., Atkinson; Press et al.). Chebychev orthogonal collocation and its application to nonlinear rational expectations models is extensively described and illustrated in Judd (1991, 1992) and Miranda.

The Chebychev orthogonal collocation method calls for the unknown equilibrium price function $\lambda(\cdot)$ to be approximated using a finite linear combination of the first $n+1$ Chebychev polynomials $\phi_0, \phi_1, \ldots, \phi_n$:

$$(10) \quad \lambda(s) \approx \sum_{j=0}^{n} a_j \phi_j(s).$$

In order to fix the $n+1$ coefficients a_j of the polynomial approximant, (8)-(9) are required to hold exactly, not at all possible supply points, but rather only at the $n+1$ Chebychev nodes $s_0, s_1, \ldots, s_n$ of the interval containing all the supply points.[3] To compute an approximate expectation in (9), Gaussian quadrature principles are used, effectively replacing the continuous distribution of production y with an m-point discrete distribution that assumes values $y_1, y_2, \ldots, y_m$ with probabilities $w_1, w_2, \ldots, w_m$, respectively.[4]

The practical value of the orthogonal collocation method is that it converts the original infinite-dimensional nonlinear functional equation problem into a finite-dimensional nonlinear equation problem that can be solved efficiently using standard numerical rootfinding techniques. Specifically, the approximation to $\lambda(\cdot)$ in (10) is obtained by solving for the $2n+2$ roots, a_i, x_i, $i = 0, 1, 2, \ldots, n$, of the $2n+2$ equations

$$(11) \quad f(a, x; \theta) = 0$$

$$(12) \quad g(a, x; \theta) = 0$$

where

$$(13)\quad f_i(a,x;\theta) \;=\; \sum_{j=0}^{n} a_j\phi_j(s_i) - p(s_i - x_i)$$

$$(14)\quad g_i(a,x;\theta) \;=\; \delta \sum_{k=1}^{m}\sum_{j=0}^{n} w_k a_j \phi_j(x_i + y_k) - p(s_i - x_i) - c(x_i).$$

Here, $\delta = 1/(1+r)$ is the discount rate and θ denotes the vector of model demand, supply, and storage function parameters.

The $2n+2$ roots of equations (11)-(12) can be computed relatively easily using successive approximation methods (Miranda 1994). The roots can also be computed by the less stable, but ultimately faster, Newton method. In the context of maximum likelihood estimation, the speed offered by Newton's method is highly desirable because the model must be re-solved every time the model parameters are perturbed by the hill-climbing routine seeking the maximum of the likelihood function. Using Newton's method to solve the model also has the added advantage that it generates information useful in computing the derivatives of the likelihood function, a feature that further accelerates and stabilizes the likelihood maximization procedure. Newton's method for solving equations (11)-(12) is discussed in greater detail in Appendix A.

4 Properties of the Price Process

Using Monte Carlo simulation methods, we now compare the behavior of commodity prices implied by the two competing versions of the rational expectations storage model: Deaton and Larocque's constant decay supply of storage model and the classical supply of storage model. In our analysis, we simulate the two models under alternative parametric specifications. For each specification, we derive the equilibrium price function $\lambda(\cdot)$ and simulate a representative price series of 100,000 years by sequentially generating random production levels y_t and applying the Markov transition rule (7). Using these price series, we compute estimates of the steady-state coefficient of variation, skewness, and first-order autocorrelation coefficient of prices implied by the different specifications.

Table 1 reports the results of simulations performed with the constant decay supply of storage model. In the first four cases, the demands are linear and production is normally distributed; in the last four, the demands are iso-elastic and production is log-normally distributed. In all cases, we assume a per-period interest rate of $r = 5\%$. The results in table 1 are essentially identical to those reported by Deaton and Larocque. With relatively elastic demand, there is little price volatility and limited storage; correlation and skewness are low, but positive. With relatively inelastic demand, price becomes more volatile, carryover becomes more prominent, and autocorrelation in prices becomes more significant as the intertemporal price link is strengthened. Higher storage costs (higher γ) discourages storage, weakening the intertemporal price link, raising price volatility, and reducing autocorrelation.

Table 1 also reveals the result that disappointed Deaton and Larocque: Although a broad range of parametric specifications is simulated, the price autocorrelation given a constant decay supply of storage never exceeds 0.47, although actual commodity prices typically exhibit substantially higher levels of autocorrelation.

Table 2 reports the results of simulations performed assuming a semi-log marginal cost of storage function that conforms to the classical supply of storage formulation. As seen in table 2, the classical supply of storage model behaves similarly to the constant decay model in some respects: relatively inelastic demand and lower storage costs imply higher price volatility and autocorrelation. However, the classical supply of storage model differs markedly from the constant decay model in one critical respect: the classical supply of storage model generates autocorrelations as high as 0.80, values that are comparable to those observed in practice.

The reason why price autocorrelation differs markedly between the two competing versions of the storage model can be explained with the aid of figure 3. Figure 3 shows the price expected in period $t + 1$ conditional on the price in period t. The relationship is computed by solving the functional equation (8)-(9) for the equilibrium price function λ and taking expectations with respect to production in equation (7) at different price levels p_t. The more linear the relationship, the higher price autocorrelation can be expected to be.

As seen in figure 3, low current prices imply abundant supplies, substantial stockholding, and a strong link between successive prices. Over this range, as the price rises, stocks fall and the expected future price rises concomitantly. This is true for both the constant decay and classical supply of storage models. The implications of the two storage models, however, diverge as prices rise. In the constant decay model, high prices induce a stockout, that is, a depletion of all stocks. Once a stockout occurs, the link between the current and expected future price is completely severed. Over this range, the expected future price remains constant, regardless of the current market price. In the classical supply of storage model, on the other hand, stocks never fall to zero and thus the link between the current and future price is never completely severed. In the classical supply of storage model, higher prices always induce some reduction in stocks and an increase in the expected future price. The constant decay model's inability to capture the high autocorrelation in prices is attributable to its implicit assumption of frequent market level stockouts — a phenomenon that is never observed in practice.

5 Estimation Method

The illustrative simulations presented in the preceding section demnostrated the rational expectation storage model's ability to generate realistically high price autocorrelations, provided the model is endowed with a classical supply of storage function. However, the question remains whether this version of the storage model can replicate the observed behavior of specific commodity price series. In this and the following section, we put the classical cost of storage version of the rational expectations model to the empirical test by estimating the model for thirteen commodities. In this section, we develop the method for estimating the model.

Two methods have been promoted for estimating nonlinear rational expectations models: the generalized method of moments (GMM) and the maximum likelihood method (Hansen and Singleton; Fair and Taylor; Miranda and Rui; Miranda and Glauber). GMM estimation is based on the model implication that ex-post errors in the intertemporal arbitrage condition at time $t + 1$,

$$(15) \quad z_{t+1} = \frac{1}{1+r} p_{t+1} - p_t - c(x_t),$$

must be uncorrelated with all information available at time t. GMM techniques, however, are not appropriate for estimating the rational expectations storage model. In order to perform GMM estimation, stock levels x_t must be observable. Unfortunately, reliable commodity stock data are often unavailable or, when available, are of questionable quality, making commodity prices the only reliably observed endogenous variable. Furthermore, even if reliable data were available, GMM estimation leaves critical model parameters unidentified, making it impossible to assess the storage model's predictions regarding the distribution of prices, particularly their autocorrelation.

The method of maximum likelihood, on the other hand, can explicitly treat the case of unobserved stocks and recovers sufficient model parameters to completely identify the Markov probability transition rule (7) that describes price dynamics.[5] In order to perform maximum likelihood estimation, we must introduce specific functional forms for the demand function, the cost of storage function, and the production probability density function. Following Deaton and Larocque, we posit a linear demand function

(16) $p(q) = a - bq$

and a production level that is normally distributed with mean μ and standard deviation σ. To capture the stylized facts of the classical supply of storage function, we introduce a semi-log cost of storage function

(17) $c(x) = \alpha + \beta \log(x)$.

We denote the vector of model parameters by $\theta = (a, b, \alpha, \beta, \mu, \sigma)$. In estimation, we fix the annual discount factor at $\delta = 1/1.05$, indicating an annual interest rate of $r = 5\%$.

Before developing our maximum likelihood estimation procedure, it must noted that the model is over-identified when quantity data are not explicitly observed. It can be shown that the storage model with parameter vector $\theta = (a, b, \alpha, \beta, \mu, \sigma)$ posited above implies the same Markov transition rule for prices (7) as a second storage model with parameter vector $\theta' = (a', b', \alpha', \beta, 0, 1)$ where $a' = a + b\mu$, $b' = \sigma b$, and $\alpha' = \alpha + \beta \log(\sigma)$.

The intuition is simple: if prices but not quantities are observed, then we are free to use any unit of measure for quantity that we wish — that is, we can shift and rescale quantities in our model arbitrarily. For convenience, we assume that $\mu = 0$ and $\sigma = 1$. Under this assumption, the remaining model parameters are just-identified.

Given that y_t is normally distributed with mean 0 and variance 1 by assumption, it follows from Markov transition rule (7) that the log-likelihood of any observed sequence of commodity prices $p_0, p_1, \ldots, p_T$ is

$$(18)\quad \mathcal{L}(\theta) = \frac{T}{2}\log(2\pi) - \sum_{t=1}^{T}\log|J_t| - \frac{1}{2}\sum_{t=1}^{T} y_t^2$$

where

$$(19)\quad y_t = \lambda^{-1}(p_t) - \lambda^{-1}(p_{t-1}) - p^{-1}(p_{t-1})$$

is the "fitted" or implied production level in period t and

$$(20)\quad J_t = \frac{\partial p_t}{\partial y_t}$$

is the Jacobian of the transformation $y_t \mapsto p_t$.

In order to evaluate the log-likelihood function $\mathcal{L}$ at a specified vector of model parameters θ, we must compute the production level y_t and Jacobian J_t implied by the parameters, for each observed price. To accomplish this, we use the orthogonal collocation methods developed in Section 3 to solve for the equilibrium price function $\lambda(\cdot)$ implied by the parameter vector θ. Given the implied equilibrium price function, we then apply standard rootfinding methods to (10) and simple Algebra to (17) to compute the implied supplies $s_t = \lambda^{-1}(p_t)$ and consumptions $q_t = p^{-1}(p_t)$ for every period t. We then compute the implied production levels y_t:

$$(21)\quad y_t = s_t - s_{t-1} + q_{t-1}$$

and Jacobians:

$$J_t = \frac{\partial p_t}{\partial y_t} = \frac{\partial p_t}{\partial s_t} = \sum_{j=0}^{n} a_j \phi_j'(s_t). \tag{22}$$

For the present study, we employ a quadratic hill-climbing routine to maximize the log-likelihood $\mathcal{L}$ function (Gill et al.). In order to accelerate convergence of the likelihood maximization routine, we compute the analytic derivatives of the log-likelihood function with respect to the model parameters θ. Computation of the derivatives of the log-likelihood function is discussed in Appendix B.

The nested functional equation algorithm was custom coded by the authors in standard Fortran 77, and was implemented using the Lahey Fortran Compiler 5.1 under MS-DOS 6.22 on a 90 megahertz Pentium personal computer.[6] In all instances, the equilibrium price function $\lambda(\cdot)$ was approximated using a 300th-order polynomial expressed as a linear combination of the Chebychev polynomials. A 3-point Gaussian-Hermite quadrature rule was used to approximate the distribution of yield shock y. The maximum likelihood algorithm was assumed to converge if the derivatives in all directions were less than 10^{-8}.

6 Estimation Results

In this section, we assess how well the modern rational expectations storage model can explain the stylized facts of commodity market price dynamics, particularly the autocorrelation exhibited by prices. Using the nested functional equation full information maximum likelihood methods developed in the preceding section, we estimated the classical cost of storage model for the same thirteen commodities examined by Deaton and Larocque: bananas, cocoa, coffee, copper, cotton, jute, maize, palm oil, rice, sugar, tea, tin, and wheat. We then used Monte Carlo simulation techniques to compute the autocorrelation implied by the models and compared the implied values to the historical values.

Commodity price data were obtained from the Commodity Division of World Bank, and consists of world prices for the thirteen commodities from 1900-1987, deflated by the U.S. consumer price index to 1980 constant dollars—these are the same data and deflation procedure as employed by Deaton and Larocque in their studies. In the present study, each deflated price series was further normalized to have a historical mean of one 1 by dividing by the sample average—this was done solely to allow easier comparisons of parameters estimates across commodity price series.

The parameter estimates for the thirteen commodities are given in table 3 with asymptotic t-statistic in parentheses. All 52 parameters were significant at the 0.01 level and possessed the expected signs: The cost of storage function is increasing $\beta > 0$, the inverse demand function is decreasing $b < 0$, and the intercept of the inverse demand function is positive $a > 0$.

Though the parameter estimates are of independent interest, our main goal is to assess the rational expectations storage model's ability to explain actual commodity prices behavior, particularly autocorrelation. Table 4 presents the first-order price autocorrelations implied by the classical supply of storage model and the price autocorrelation exhibited by the historical price series. For comparison, we present the autocorrelations predicted by Deaton and Larocque's constant decay supply of storage model.

As seen in table 4, there is a striking difference between the two competing models' ability to explain observed autocorrelation in commodity prices. The constant decay supply of storage model examined by Deaton and Larocque provides a very poor fit. In particular, the model consistently predicts autocorrelations that are substantially lower than those actually observed in practice. The classical supply of storage model, on the other hand, provides an excellent fit. The classical supply of storage model predicts autocorrelations that are as high as those observed in practice. Across the thirteen price series, the classical supply of storage model predicts the highest (lowest) autocorrelation for the commodities exhibiting the highest (lowest) historical autocorrelation. For ten of the thirteen commodities, the predicted and actual price autocorrelations differ by 0.04 or less; for six of the ten, the predicted and actual autocorrelations differ by less than 0.02.

7 Conclusion

The modern rational expectations commodity storage model explains the historical autocorrelation exhibited by commodity prices extremely well, provided that it is endowed with a classical supply of storage function. The closeness of the fit obtained in our empirical estimation is specially remarkable, given the conceptual simplicity of the model and the fact that only price data were used to estimate it. Our results refute Deaton and Larocque's earlier negative findings, which were obtained using a constant decay supply storage model. The discrepancy between our findings and Deaton and Larocque's findings are clearly attributable to differences in assumptions regarding the coincidence of negative intertemporal price spreads and positive stocks. Our results underscore the importance of accurately representing the supply of storage when modeling commodity price dynamics and suggest that further research be conducted on the microfoundations of the market-level supply of storage relationship.

The method that we developed to estimate the storage model provides a good example of how numerical functional equation methods can be used to compute full information maximum likelihood estimates of analytically insoluble nonlinear dynamic economic models. Our nested functional equation algorithm should be applicable, with only a few alterations, to the maximum likelihood estimation of other nonlinear rational expectations models and to dynamic decision models with continuous state and decision spaces.

Endnotes

1. Classical supply of storage theory focused on the form of the marginal cost of storage function, without addressing how prices and price expectations are formed. The modern theory of storage is characterized by the introduction of sufficient additional structure, including the assumption of rational expectations, to yield a complete, soluble partial equilibrium model of commodity market behavior.

2. Wright and Williams (1989) do not dispute the existence of a market-level supply of storage function like the one described in the classical storage literature. They do, however, dispute its microfoundations. They reject the convenience yield explanation, arguing instead that apparent negative storage costs are a consequence of spatial aggregation.

3. Deaton and Larocque (1994) used linear spline approximation, not orthogonal collocation, to solve the storage model. Both techniques are specific implementations of the Galerkin collocation method (Judd 1991). The former technique differs from the latter in that it uses

tent functions with finite support, not Chebychev polynomials, as basis functions ϕ_i in the approximation (10).

4. Guassian quadrature calls for the m probability mass points and the m probabilities to be chosen so that the discrete random variable has the same $2m$ lower moments as the continuous random variable being discretized. The quadrature rule exactly computes the expectation of a function of the original random variable, provided the function is a polynomial of exact degree $2m-1$ or less.

5. Deaton and Larocque use a pseudo maximum likelihood estimation method, which assumes that equilibrium prices are normally distributed. This approach imposes a distributional structure on prices that is inconsistent with that predicted by the underlying structural model and further ignores information about third and higher moments of the price distribution. Our maximum likelihood method, on the other hand, does not impose a contradictory distributional assumption on prices and uses all available moment information consistently with the underlying model specification.

6. After the model is solved the first time, the likelihood function and its derivatives can be subsequently computed in less than three seconds with each perturbation of the model parameters. The amount of time required to estimate the model varied across commodities, but never exceeded thirty minutes of computer time.

References

ATKINSON, K. E. (1989), *An Introduction to Numerical Analysis, 2nd Ed,* John Wiley & Sons: New York.

BRENNAN, M. J. (1958), “The Supply of Storage”, *American Economic Review,* 48, 50-72.

DEATON, A. and G. LAROCQUE (1992), “On the Behavior of Commodity Prices”, *Review of Economic Studies,* 59, 1-23.

DEATON, A. and G. LAROCQUE (1994), “Competitive Storage and Commodity Price Dynamics”, Unpublished paper, Princeton University.

FAMA, E. F. and K. R. FRENCH (1987), “Commodity Futures Prices: Some Evidence on Forecast Power, Premiums, and the Theory of Storage”, *Journal of Business,* 60, 55-73.

GILL, P. E., W. MURRAY, M. A. SAUNDERS, and M. H. WRIGHT (1986), “User’s Guide for NPSOL (Version 4.0): A Fortran Package for Nonlinear Programming”, Department of Operations Research, Stanford University, Technical Report SOL 86-2.

JUDD, K. L. (1991), *Numerical Methods in Economics,* Manuscript, Hoover Institution, Stanford University.

JUDD, K. L. (1992), “Projection Methods for Solving Aggregate Growth Models”, *Journal of Economic Theory,* 58, 410-452.

KALDOR, N. (1939-40), “Speculation and Economic Stability”, *Review of Economic Studies,* 7, 1-27.

MIRANDA, M. J. (1994), "Numerical Solution Strategies for the Nonlinear Rational Expectations Commodity Storage Model", Unpublished paper, The Ohio State University.

MIRANDA, M. J. and J. W. GLAUBER (1993), "Estimation of Dynamic Nonlinear Rational Expectations Models of Primary Commodity Markets with Private and Government Stockholding", *Review of Economics and Statistics,* 75.

MIRANDA, M. J. and P. G. HELMBERGER (1988), "The Effects of Price Band Buffer Stock Programs", *American Economic Review,* 78, 46-58.

MIRANDA, M. J. and X. RUI (1995), "Maximum Likelihood Estimation of Nonlinear Rational Expectations Asset Pricing Models", *Journal of Economic Dynamics and Control,* Forthcoming.

MUTH, J. F. (1961), "Rational Expectations and the Theory of Price Movements", *Econometrica,* 29, 315-35.

PRESS, W. H., S. A. TEUKOLSKY, W. T. VETTERLING, and B. P. FLANNERY (1992), *Numerical Recipes in FORTRAN: The Art of Scientific Computing, 2nd Ed.* Cambridge: Cambridge University Press.

SALANT, S. W. (1983), "The Vulnerability of Price Stabilization Schemes to Speculative Attack", *Journal of Political Economics,* 91, 1-38.

SCHEINKMAN, J. A. and J. SCHECHTMAN (1983), "A Simple Competitive Model with Production and Storage", *Review of Economic Studies,* 50, 427-41.

WILLIAMS, J. C. and B. D. WRIGHT (1991), *Storage and Commodity Markets,* New York: Cambridge University Press.

WILLIAMS, J. B. (1935), "Speculation and the Carryover", *Quarterly Journal of Economics,* 6.

WORKING, H. (1949), "The Theory of Price of Storage", *American Economic Review,* 39, 1254-62.

WRIGHT, B. D., and J. C. WILLIAMS (1984), "The Welfare Effects of the Introduction of Storage", *Quarterly Journal of Economics,* 99, 169-82.

WRIGHT, B. D., and J. C. WILLIAMS (1982), "The Economic Role of Commodity Storage", *Economic Journal,* 92, 596-616.

WRIGHT, B. D., and J. C. WILLIAMS (1989), "A Theory of Negative Prices for Storage", *Journal of Futures Markets,* 9, 1-13.

Appendix A: Solving the Storage Model Using Newton's Method

The $2n+2$ roots of equations (11)-(12) can be computed using Newton's method:

0. Initialization Step: Guess the values of the coefficients $a_0, a_1, \cdots, a_n$ and the equilibrium stock levels $x_0, x_1, \cdots, x_n$.

1. Linearization Step: At the incumbent coefficients $a_0, a_1, \cdots, a_n$ and stocks $x_0, x_1, \cdots, x_n$, compute the increments $\Delta a_0, \Delta a_1, \cdots, \Delta a_n$ and

$\Delta x_0, \Delta x_1, \cdots, \Delta x_n$ by solving the linear equation system

$$\left[\begin{array}{cc} \frac{\partial f}{\partial a} & \frac{\partial f}{\partial x} \\ \frac{\partial g}{\partial a} & \frac{\partial g}{\partial x} \end{array}\right]\left[\begin{array}{c} \Delta a \\ \Delta x \end{array}\right] = \left[\begin{array}{c} -f \\ -g \end{array}\right] \tag{23}$$

where

$$\begin{aligned} \frac{\partial f_i}{\partial a_j} &= \phi_j(s_i) \\ \frac{\partial g_i}{\partial a_j} &= \delta \sum_{k=1}^{m} w_k \phi_j(x_i + y_k) \\ \frac{\partial f_i}{\partial x_i} &= p'(s_i - x_i) \\ \frac{\partial g_i}{\partial x_i} &= \delta \sum_{k=1}^{m} w_k a_i \phi_j'(x_i + y_k) + p'(s_i - x_i) - c'(x_i) \\ \frac{\partial f_i}{\partial x_j} &= \frac{\partial g_i}{\partial x_j} = 0 \qquad \text{for } i \neq j. \end{aligned}$$

and f_i and g_i are defined in (13)-(14).

2. Update Step: Update the coefficients and stock levels by setting $a_j = a_j + \Delta a_j$ and $x_j = x_j + \Delta x_j$ for $j = 0, 1, \cdots, n$.

3. Convergence Check: If the norms of Δa and Δx are tolerably small, stop; otherwise, return to step 1.

Appendix B: Computing the Derivatives of the Likelihood Function

To compute the derivative of the log-likelihood function requires that we compute the derivatives of the random shock y_t and the Jacobian J_t with respect to θ:

$$\frac{\partial y_t}{\partial \theta} = \frac{\partial s_t}{\partial \theta} - \frac{\partial s_{t-1}}{\partial \theta} + \frac{\partial q_{t-1}}{\partial \theta} \tag{24}$$

and

$$(25)\quad \frac{\partial J_t}{\partial \theta} = \frac{\partial^2 p_t}{\partial \theta \partial s_t} + \frac{\partial^2 p_t}{\partial s_t^2}\frac{\partial s_t}{\partial \theta}$$

where

$$(26)\quad \frac{\partial s_t}{\partial \theta} = -\frac{\partial p_t}{\partial \theta} / \frac{\partial p_t}{\partial s_t}$$

and

$$(27)\quad \frac{\partial q_t}{\partial \theta} = \frac{\partial p^{-1}(p_t)}{\partial \theta} + \frac{\partial p^{-1}(p_t)}{\partial p_t} + \frac{\partial p_t}{\partial \theta}$$

Most of the terms in (24)-(27) can be computed analytically after deriving the implied availability levels s_t. The only terms that cannot be computed directly are the derivatives of p_t and $\frac{\partial p_t}{\partial s_t}$ with respect to θ. Differentiating (10) and (24) with respect to θ, we obtain

$$(28)\quad \frac{\partial p_t}{\partial \theta} = \frac{\partial \lambda}{\partial \theta}(s_t;\theta) = \sum_{j=0}^{n} \frac{\partial a_j}{\partial \theta}\phi_j(s_t)$$

and

$$(29)\quad \frac{\partial^2 p_t}{\partial \theta \partial s_t} = \frac{\partial^2 \lambda}{\partial \theta \partial s_t}(y_t;\theta) = \sum_{j=0}^{n} \frac{\partial a_j}{\partial \theta}\phi_j'(s_t).$$

Thus, computing the derivatives of the log-likelihood function $\mathcal{L}$ reduces to computing the derivatives of the Chebychev coefficients a_j with respect to θ.

The derivatives of the coefficients a_j with respect to θ are obtained by implicitly differentiating (15) with respect to θ, which gives rise to the following linear equation system

$$(30)\quad \begin{bmatrix} \frac{\partial f}{\partial a} & \frac{\partial f}{\partial x} \\ \frac{\partial g}{\partial a} & \frac{\partial g}{\partial x} \end{bmatrix} \begin{bmatrix} \partial a/\partial \theta \\ \partial x/\partial \theta \end{bmatrix} = \begin{bmatrix} -\frac{\partial f}{\partial \theta}(a,x;\theta) \\ -\frac{\partial g}{\partial \theta}(a,x;\theta). \end{bmatrix}$$

Most of the expense in solving linear equation system (30) comes from computing and factoring the matrix of partial derivatives of f and g with respect to a and x. This operation, however, is performed naturally during the Newton algorithm used to compute the coefficients a_j of the equilibrium price function $\lambda(\cdot)$ function and thus does not have to be repeated. Thus, the analytic derivatives of the likelihood function can be computed at little additional cost once the equilibrium price function $\lambda(\cdot)$ has been derived.

Table 1: Distribution of Prices for Constant Decay Supply of Storage Model.

	$p(q) = a - bq$ $y \sim N(1, 0.01)$				$p(q) = aq^{-b}$ $\log y \sim N(0, 0.01)$			
Parameters								
γ	0.05	0.00	0.05	0.00	0.05	0.00	0.05	0.00
a	2.00	2.00	6.00	6.00	1.00	1.00	1.00	1.00
b	1.00	1.00	5.00	5.00	1.00	1.00	5.00	5.00
Percent Variation	0.09	0.08	0.28	0.24	0.09	0.08	0.36	0.30
Autocorrelation	0.08	0.20	0.34	0.47	0.10	0.19	0.29	0.40
Skewness	0.47	0.86	1.63	2.01	0.67	1.00	3.08	3.64

Table 2: Distribution of Prices for Classical Supply of Storage Model.

	$p(q) = a - bq$ $c(x) = \alpha - 0.1 \log x$ $y \sim N(1, 0.01)$				$p(q) = aq^{-b}$ $c(x) = \alpha - 0.1 \log x$ $\log y \sim N(0, 0.01)$			
Parameters								
α	0.30	0.05	0.30	0.05	0.30	0.05	0.30	0.05
a	2.00	2.00	6.00	6.00	1.00	1.00	1.00	1.00
b	1.00	1.00	5.00	5.00	1.00	1.00	5.00	5.00
Percent Variation	0.08	0.05	0.30	0.16	0.08	0.05	0.36	0.17
Autocorrelation	0.20	0.60	0.41	0.80	0.19	0.60	0.33	0.80
Skewness	0.27	0.16	0.98	0.37	0.42	0.15	2.60	0.84

Table 3: Parameter Estimates for Classical Supply of Storage Model.*

Commodity	α	β	a	b
Bananas	-0.222	0.072	0.941	-1.518
	(0.013)	(0.005)	(0.056)	(0.173)
Cocoa	-0.150	0.112	0.948	-1.614
	(0.049)	(0.030)	(0.175)	(0.516)
Coffee	-0.175	0.116	1.003	-1.589
	(0.050)	(0.033)	(0.170)	(0.616)
Copper	-0.121	0.074	0.934	-1.013
	(0.033)	(0.018)	(0.109)	(0.260)
Cotton	-0.162	0.071	0.887	-2.772
	(0.007)	(0.004)	(0.041)	(0.261)
Jute	-0.126	0.121	0.976	-0.840
	(0.047)	(0.033)	(0.090)	(0.237)
Maize	-0.112	0.100	0.957	-1.034
	(0.046)	(0.030)	(0.112)	(0.289)
Palm Oil	-0.149	0.087	0.889	-1.736
	(0.049)	(0.024)	(0.191)	(0.704)
Rice	-0.169	0.083	0.771	-2.480
	(0.014)	(0.009)	(0.055)	(0.337)
Sugar	-0.091	0.101	0.924	-1.475
	(0.039)	(0.022)	(0.161)	(0.384)
Tea	-2.207	0.733	0.938	-0.834
	(4.302)	(1.102)	(0.093)	(0.306)
Tin	-0.147	0.070	1.334	-4.225
	(0.008)	(0.006)	(0.073)	(0.702)
Wheat	-0.155	0.082	0.918	-1.967
	(0.024)	(0.015)	(0.209)	(0.544)

*Standard error in parentheses.

Table 4: Actual Versus Predicted Autocorrelation under Alternative Supply of Storage Formulations.

Commodity	Actual	Classical Predicted	Constant Decay Predicted
Bananas	0.94	0.95	0.35
Cocoa	0.82	0.79	0.30
Coffee	0.80	0.80	0.24
Copper	0.84	0.81	0.39
Cotton	0.90	0.88	0.19
Jute	0.73	0.72	0.30
Maize	0.78	0.74	0.36
Palm Oil	0.75	0.83	0.42
Rice	0.85	0.87	0.22
Sugar	0.62	0.71	0.26
Tea	0.82	0.83	0.23
Tin	0.89	0.80	0.26
Wheat	0.88	0.85	0.26

Figure 1. Supply of Storage

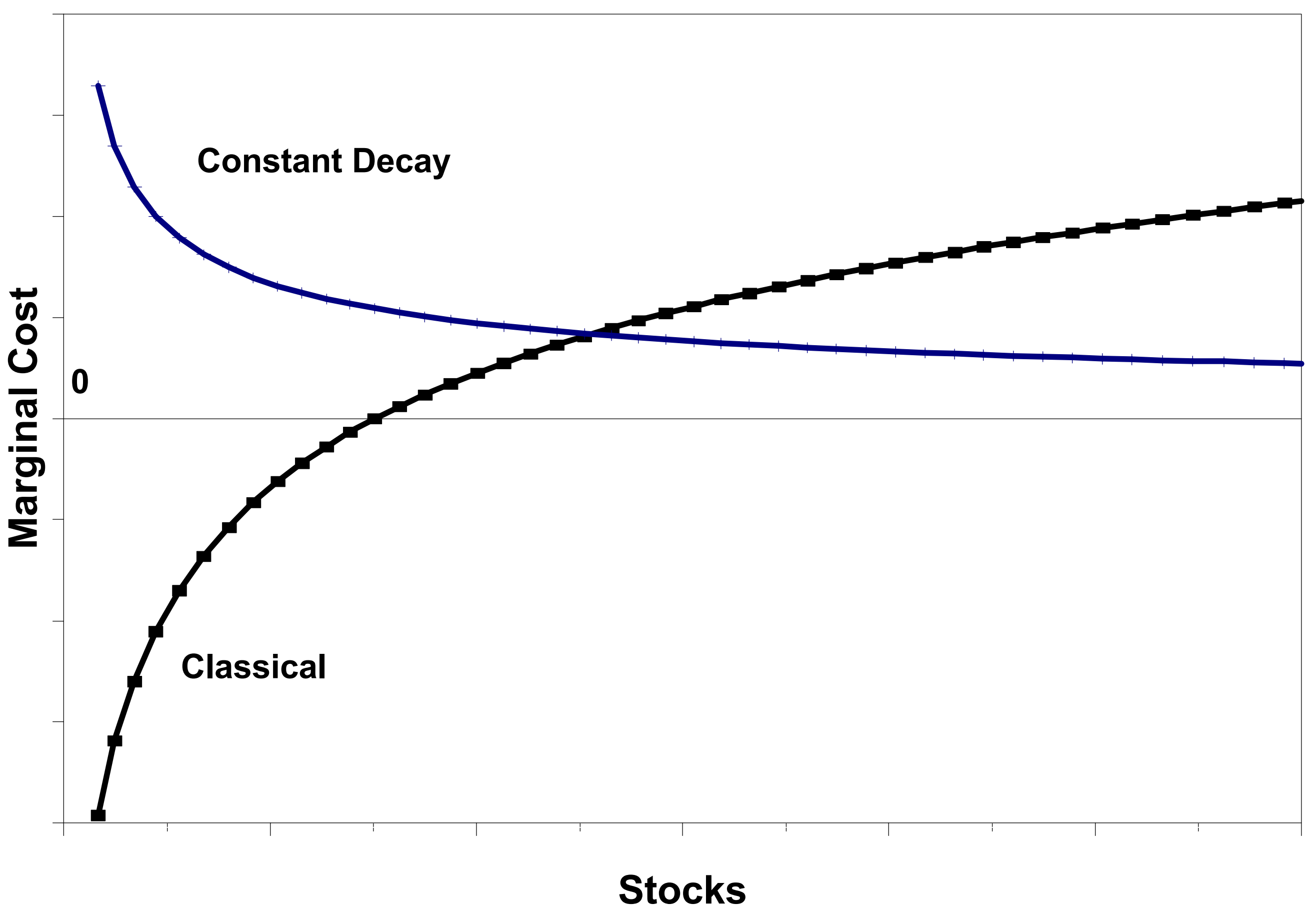

Figure 2. Equilibrium Price

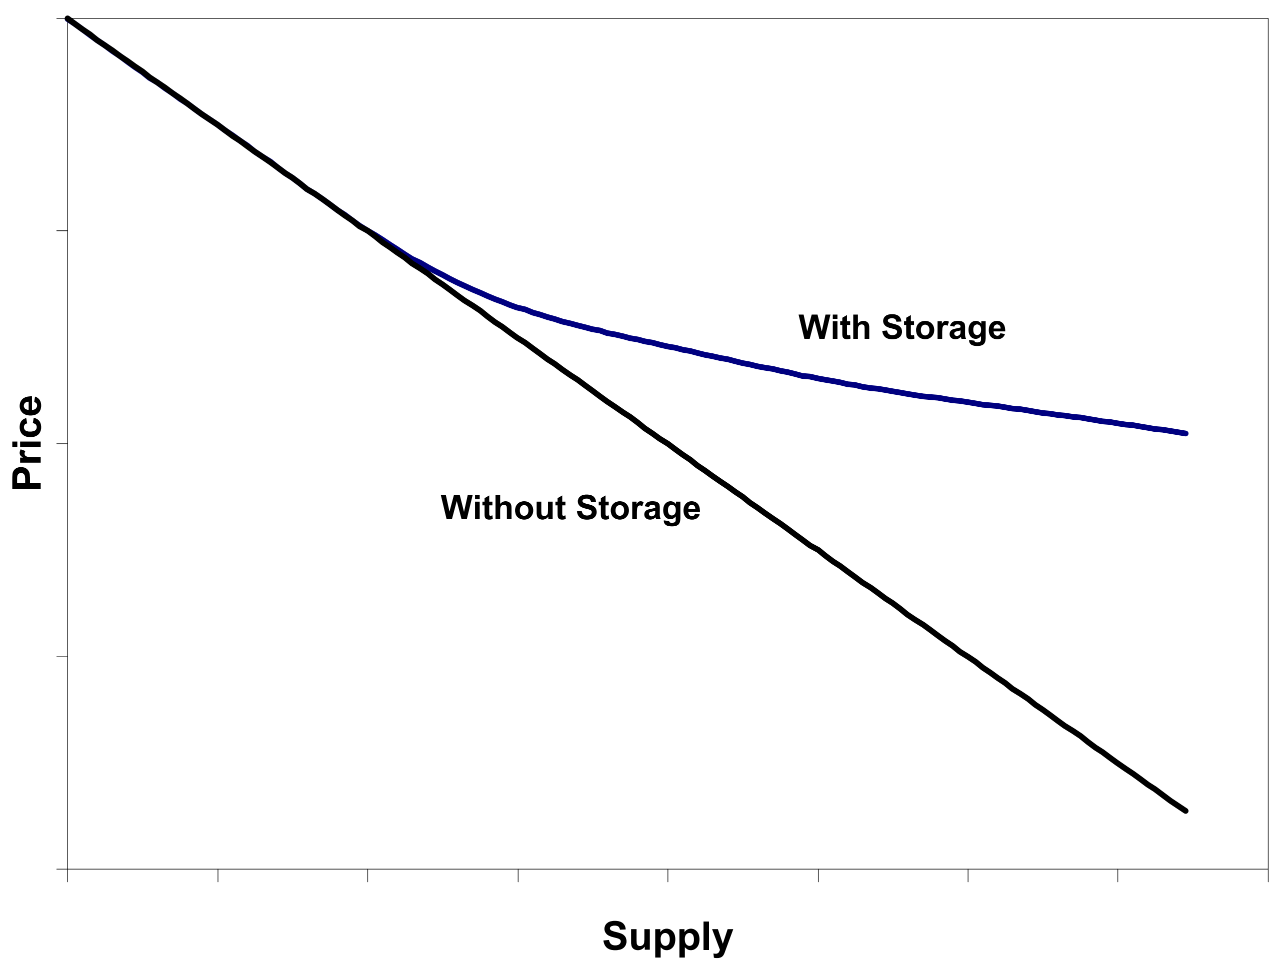

Figure 3. Markovian Price Transitions