ESTIMATING THE COMMODITY PRICE MODEL

Discussion Paper #154

by

Angus Deaton*

and

Guy Laroque**

*Research Program in Development Studies
Center of International Studies
Woodrow Wilson School
Princeton University
Princeton, NJ 08544

**INSEE
Paris, France

Research Program in Development Studies
Center of International Studies
Woodrow Wilson School
Princeton University
Princeton, NJ 08544-1022

May and November 1991

0 Introduction

This paper is concerned with the estimation and testing of a model of price formation for primary commodities, where supply and demand shocks are mediated by speculative storage. In the version that we estimate, the model supposes that, in the absence of storage, price will be a linear function of an independent and identically distributed shock. The addition of risk-neutral and profit maximizing stockholders means that expected future prices cannot be greater than current prices by more than the cost of holding inventories into the future, and this condition induces a relationship between prices between periods in which stocks are held. The theoretical properties of this model are by now relatively familiar from the contributions of Samuelson (1971) and Newbery and Stiglitz (1981), as well as the excellent recent monograph by Williams and Wright (1991). In Deaton and Laroque (1992), we made a first attempt to confront the model with actual commodity prices using annual data on thirteen commodities from 1990 to 1987. We simulated the model in an attempt to reproduce some of the stylized facts of commodity price behavior, and we tested some of its implications, for example about the relationship between current and expected future prices. Our results were encouraging, at least in some respects. Simulated data reproduce the pattern of 'doldrums' interrupted by upward (not downward) 'spikes' that is characteristic of many actual commodity prices. Furthermore, our econometric tests could not reject the implication of the model that, below a fixed cutoff, one period ahead price expectations are current prices multiplied by a factor greater than unity, while above the cutoff, expectations are constant. However, we could not find parameter values where the associated simulations reproduced the high levels of positive autocorrelation which are displayed by the actual series.

In the current paper, we make a more serious attempt to fit the model to the data, and to estimate the parameters. We did not do so in the previous paper because estimation of the model is technically difficult, and we did not know whether we could generate useful results. However, we have overcome at least some of the problems, and it is now clear that the simple model cannot account for the evidence. The autocorrelation properties of the data are quite different from those that are generated by the model, so that, for example, one period ahead price expectations are in most cases better matched by a simple first-order linear autoregression. Further, if the underlying shocks to the process are independent and identically distributed over time, speculators will break even by storing a commodity

whenever its price is below some constant critical level. In practice, such a strategy would have generated considerable losses for all reasonable assumptions about the costs of storage. If the autocorrelation in commodity prices is to be explained by the actions of stockholders without autocorrelation in supply or demand shocks, then prices must typically be expected to rise by the carrying costs of stocks. In practice, prices do not rise by this amount on average, so that stockholding is typically unprofitable unless speculators have more information than is allowed by the model. It seems that commodity price autocorrelation must be explained at least in part by autocorrelation in the underlying processes of supply and demand; speculative storage by itself is not sufficient.

We also have a methodological aim in writing this paper. The technical problems of estimation come from the fact that there is no analytical expression for the stochastic process that is implied by the model for the observed commodity price. Instead, prices are supposed to follow a nonlinear first-order Markov process, and the nonlinear function is defined only implicitly by a functional equation involving the underlying parameters of the problem. Such situations are frequently encountered when processes are the solutions to dynamic programs, or can be interpreted as such, for example in nonlinear rational expectations macro models. There is now a considerable body of experience in approximating the solutions to the functional equations, solutions that can be used to generate stochastic simulations. There is a smaller, but rapidly developing literature that goes one step further, and estimates parameters by matching the properties of the numerical solutions to those of the actual data. The problem we examine in this paper is a very simple example of this general methodology, and as such, provides useful lessons about at least some of the many pitfalls. For example, approximation techniques that work well for simulation can be quite unsuitable for estimation. Very simple nonlinearities, in this case due to non-negativity requirements, can generate multiple discontinuities in likelihood functions, as can ill-chosen approximation techniques. Moreover, fundamental questions, like model identification, can become obscure and difficult. While fast computation is necessary to solve these problems, thoughtless application of computing power is most unlikely to lead to correct or convincing answers.

The paper is arranged as follows. Section 1 summarizes the simple storage model in the case where the driving variable (the harvest) is independently and identically distributed over time. We introduce a new result on the way in which the rational expectations equilibrium

is affected by different assumptions about the harvest in the case where consumption is a linear function of price. In particular, we show that the amount of storage is proportional to the standard deviation of the harvest, and is unaffected by its mean. This result is of interest in its own right, and provides great simplifications in the econometric application. Section 2 is concerned with the econometric procedures, and with the detail of how the computations were organized. This section is entirely methodological, and can be skipped by readers who are primarily interested in commodity prices, and in the ability of the storage model to explain them. Section 3 contains the empirical results.

1. A simple model of a commodity price

The basic model of commodity price is one in which there are profit-maximizing risk-neutral stockholders who modify what would otherwise be the outcome of a simple process of supply and demand. Suppose that, in the absence of stocks, there is a price p_t and a stochastic shock z_t, such that

$$p_t = P(z_t) = a + bz_t. \tag{1}$$

In the simplest example, z_t would be an inelastically supplied harvest, and (1) would be the (inverse) demand function; with reference to this interpretation, we shall often refer to z_t as the "harvest." More generally, both supply and demand could be linear functions of current price, with shocks in both schedules, so that z_t would be an amalgam of shocks to both supply and demand. Note that if, to specialize further, z_t is an independently and identically distributed normal shock with mean μ and variance σ^2, price will be distributed as $N(a+b\mu, b^2\sigma^2)$, so that, if only price is observed, a is not identified separately from μ, nor b from σ.

At the end of period t, stockholders hold a non-negative amount of inventory I_t; the market cannot hold negative stocks, since commodities cannot be consumed before they exist. There is a constant real rate of interest r, and stocks physically deteriorate at rate δ, so that I_t units of stocks yields $(1-\delta)I_t$ units one period later. If prices are expected to fall, or to rise by less than carrying costs, inventories are zero, while, for positive inventories, the expected

price one period ahead must be the current price plus carrying costs. Hence, in the presence of speculators, price must satisfy

$$p_t = \max\left(\frac{1-\delta}{1+r}E_t p_{t+1}, P\left(z_t+(1+\delta)I_{t-1}\right)\right). \tag{2}$$

The first term on the right hand side is the expected value of holding one unit of the good into the next period, after allowing for deterioration and interest costs; if inventories are being held, this must equal the current price. The second term is the value of the current price if no stocks are taken into $t+1$, with the current harvest z_t and surviving inventories $(1-\delta)I_{t-1}$ sold for what they can fetch. If this latter price is larger than the former, speculators will be unwilling to hold inventories, and it is the latter that will set the market price. However, if selling everything would drive the price lower than the future price less costs, stocks will be held, and the price arbitraged up to the first term in the square brackets.

The evolution of prices governed by (2) depends on the nature of the stochastic process z_t, and in particular on how successive demand and supply shocks are linked over time. The simplest case, and the one on which much of the literature has focussed, is that in which z_t is independently and identically distributed. Again, the easiest case to think of is where the commodity is an agricultural crop, the shocks are supply shocks and are driven by the weather. Of course, even with an i.i.d. z_t process, equation (2) will generate autocorrelation in the prices, and our main topic of interest is whether stockholding alone is sufficient to explain the actual degree of autocorrelation in commodity prices, even if the forcing variable z_t is not autocorrelated.

Suppose then that z_t is i.i.d., and define the state variable x_t, the "amount on hand" by

$$x_t = (1-\delta)I_{t-1} + z_t. \tag{3}$$

Under appropriate assumptions, the most important of which are that $r+\delta>0$ and that z_t has compact support, see Deaton and Laroque (1992, Theorem 1) there is a rational expectations equilibrium where the price p_t is a function of the amount on hand

$$p_t = f(x_t) \tag{4}$$

where $f(x)$ is the unique non-negative non-increasing solution to the functional equation

$$f(x) = \max\left(\frac{1-\delta}{1+r}E_z f\left[z + (1-\delta)\left(x - P^{-1}f(x)\right)\right], P(x)\right). \tag{5}$$

Note that $f(x)$ depends on the fundamental "parameters" of the problem, the depreciation rate δ, the interest rate r, the demand parameters a and b, and the probability distribution function of the harvest process z, with respect to which the expectation in (5) is taken. Once the function is calculated, the price is given by (4), and the state variable x_t is updated according to

$$x_{t+1} = (1-\delta)\left(x_t - P^{-1}\left(f(x_t)\right)\right) + z_{t+1} \tag{6}$$

so that inventories in t are the amount on hand less consumption, given by inserting price $f(x_t)$ into the demand function $P^{-1}(p)$. To get next period's amount on hand, they are multiplied by the depreciation factor $(1-\delta)$, and augmented by the new harvest z_{t+1}.

In Deaton and Laroque (1992), we provide a number of results characterizing the properties of the stationary rational expectations equilibrium defined by equations (4), (5), and (6). Some of these are required later in this paper, and are restated here. Associated with the price function $f(x_t)$ is a stocking or inventory function $I(x_t)$, defined by

$$I(x_t) = x_t - P^{-1}\left(f(x_t)\right) \tag{7}$$

so that, given the amount on hand x_t, inventories carried into the next period are the amount on hand less consumption, which is given by applying the demand function $P^{-1}(p_t)$ to the current price $f(x_t)$. $I(x_t)$ has the property that there exists a critical amount on hand x^* and associated price $p^* = f(x^*)$, such that for $x \le x^*$, and hence $p_t \ge p^*$, $I(x_t) = 0$, so that no stocks are carried forward, and $f(x_t) = P(x_t)$. If the amount on hand is greater than x^*, it is optimal to store some of it, so that $I(x_t) > 0$ and $f(x_t) > P(x_t)$, because speculative inventories are a source of demand over and above consumption. Neither $I(x_t)$ nor $f(x_t)$ are differentiable at $x = x^*$, where $p = p^*$. Since inventories are always positive when $p_t < p^*$, and because inventories can be positive only if the expected price next period is sufficient to cover the holding costs, we have immediately that for $p_t < p^*$,

$$E(p_{t+1}|p_t) = \frac{1+r}{1-\delta} p_t \tag{8a}$$

while for $p_t > p^*$,

$$E(p_{t+1}|p_t) = E_t f(z) = \text{constant.} \tag{8b}$$

Equations (8a) and (8b) characterize the autoregression function of prices when the simple storage model is correct, and will be useful in interpreting the results in Section 3.

The analysis in this paper assumes that the inverse demand function is linear, equation (1), in which case it is possible to be a good deal more specific about the nature of the functions $f(x_t)$ and $I(x_t)$. To explain these new results, it is useful to begin from the case where there is no storage, and where the harvests are i.i.d. as $N(0,1)$. Given the linear price function (1), prices will also be i.i.d. normal, with mean a and variance b^2. Imagine a second economy, with an identical history of harvest shocks, except that the mean and variance are now μ and σ^2. With the same inverse demand function (1), prices in this second economy will be the same as prices in the first economy, but with the mean shifted to $a+b\mu$ and the variance $(\sigma b)^2$. If, instead of (1), the second economy's inverse demand function had been

$$\tilde{P}(x) = (a - \frac{b\mu}{\sigma}) + \frac{b}{\sigma} x \tag{9}$$

then the differences in means and variances would be offset, and the sequence of prices would be identical in the two economies. If only prices are observed, it is also clear that the four parameters a, b, μ, and σ cannot be separately identified, and that, in this case, there would be no loss in generality in assuming that the harvests have mean zero and unit variance.

Perhaps more surprisingly, the same result holds in the presence of optimal speculative storage, although the reasoning is a good deal more complex. Optimal storage in the model provides a stationary buffer stock that shields consumption against variations in the harvest, transferring the commodity from periods when it is less valued to periods when it is more valued. With linear demand functions, it is reasonable to conjecture that an increase in the mean of the harvest will simply reduce prices by the appropriate amount, without changing storage behavior. Inventories follow a stationary process, and the increase in mean supply

must eventually be passed through into the same increase in consumption. In contrast, changes in the variability of the harvest will produce more opportunities for speculative storage, since the spread between today's price and tomorrow's expected price that is required to induce storage will now occur more often. Once again, it might reasonably be conjectured that the amount of storage is proportional to the standard deviation of the harvest, and we show in the Appendix that these conjectures are correct when the demand function is linear. From this, it is easy to see that the result of the previous paragraph continues to hold in the presence of storage. At all times, the second economy will have σ times as much storage as does the first economy, so that in any period t, if the amount on hand in the second economy is $\tilde{x}_t$, we have

$$\tilde{x}_t = (1-\delta)\tilde{I}_{t-1} + \tilde{z}_t = \sigma(1-\delta)(I_{t-1} + z_t) + \mu = \sigma x_t + \mu \tag{10}$$

where the quantities without tildes refer to the first, baseline economy. Hence, just as the harvest is subject to a linear transformation in the no-storage case, so is the amount on hand in the storage case. The effect on prices is the same.

We state the result formally:

Proposition 1 *(identification when only prices are observed):*

Consider an economy E, with real interest rate r and inventory depreciation parameter δ*, where the harvest shocks z have zero mean and unit variance, and where the inverse demand function P(x) is linear, as in equation (1). Any other economy* $\tilde{E}$*, with real interest rate r and inventory depreciation parameter* δ*, with harvest shocks* $\tilde{z} = \mu + \sigma z$ *and inverse demand function (9) has the same rational expectations price process as the base economy E.*

Proof: See Appendix.

This result is of some interest in its own right, since it shows very clearly how the variability of harvests affects the amount stored, albeit in the very simple case where the demand function is linear. (Note however, as can be confirmed from the proof in the

Appendix, the result does not require that the harvests be normally distributed.) The result also has applications beyond the present context. For example, in a model of consumption where assets are required to be nonnegative, and where preferences are additively separable and quadratic (as in the permanent income hypothesis), and where households are sufficiently impatient so that the rate of time preference exceeds the constant real rate of interest, then different households with i.i.d. income processes differing only in mean and variance will have savings and asset processes that differ only in scale, the scale being proportional to the standard deviation of income. For the purposes of the current work, the result is important because it prevents us attempting to estimate an underidentified model, and instead allows us to reduce the number of parameters to be estimated. Clearly, we can arbitrarily set the mean and standard deviation of the harvest process to be one and zero respectively, and estimate only the four parameters a, b, r, and δ rather than the original six. This should be regarded as an arbitrary normalization without economic significance; we do not suppose that harvest can be negative. Nor do we need to take any action to guarantee that prices should be positive. In this model, and *in the absence of storage*, there is nothing to stop prices being negative. However, with speculative stockholding, there is always a rational expectations equilibrium in which the function $f(x)$ is non-negative and our algorithms will always generate non-negative prices. The reduction in the dimensionality of the estimation problem is an important advantage of the linear demand function. It is clear that, in general, we cannot expect the distribution of prices alone to identify both the harvest process and the demand function, and our specification recognizes this general lack of identification in an explicit and transparent manner.

Equations (4), (5), and (6) provide a complete description of the model that is to be fitted to the data. Given a set of parameters, and a distribution function for z, the function $f(x)$ is found that satisfies (5). A starting value for x_0 is selected, for example μ, the mean of z, which gives an initial price according to (4). The updating rule (6) gives x_1, and thus p_1 from (4), and so on. These predicted prices can be compared with the actuals, and the parameters suitably adjusted. Since the function $f(x)$ is monotone, Deaton and Laroque, Theorem 1, it can be inverted, so that (4) and (6) can be combined to give

$$p_{t+1} = f\left((1-\delta)\left(f^{-1}(p_t) - P^{-1}(p_t)\right) + z_{t+1}\right) \tag{11}$$

which shows that the price follows a nonlinear first-order Markov or autoregressive process. Equation (11) completely defines the distribution of p_{t+1} conditional on p_t, and thus provides a likelihood function that could, in principle, be used for estimation.

2 *Econometric and computational issues*

In practice, the conceptual scheme of the previous paragraph is adhered to quite closely in what follows. However, there are a large number of practical questions that have to be resolved, and it is to these that we now turn. We have already dealt with the first practical issue, the fact that the full six parameter model is not identified, and shown how to use Proposition 1 to reduce the parameter space to four parameters. Second, we must find a procedure for approximating the function $f(x)$, preferably one that is computationally fast, since it will be used at least once for each likelihood evaluation, and many more times for calculating derivatives. Third, we need an estimation method that is well-behaved for the problem at hand. As we shall see, maximum likelihood is not a good choice. The second and third issues are not independent, since the robustness of the estimation procedure can be seriously affected by the way in which the function $f(x)$ is approximated.

We require a numerical approximation to the function which, to some prespecified degree of accuracy and over some range, satisfies equation (5). One technique that has been commonly employed in the literature is the use of polynomial approximations, see for example the symposium edited by Taylor and Uhlig (1990), and Judd (1991), and in the current context, Williams and Wright (1991, Chapter 2). For example, suppose that we have some guess for the range of the state variable x, and that we select some set of points $(x_1, x_2, \ldots, x_G)$ over that range. Given some low degree polynomial approximation for $f(x)$, the parameters can be chosen to minimize the distance between the left-hand and right-hand sides of (5) at the pre-specified points. Higher accuracy can be obtained by increasing the number of points and the degree of the polynomial. We can also exploit known particular features of the problem, for example that for x less than some x^*, $f(x) = P(x)$, so that if we make x^* a parameter, we can fit the polynomial approximation only for $x > x^*$. Williams and Wright's procedure is of this form, and has the desirable feature that the polynomial is not required to fit the kink at x^*. In our experience, algorithms such as these are fast and accurate,

and they are relatively easy to compute. However, we encountered difficulties in using polynomial approximations for the estimation. First, the function $f(x)$ is non-negative and convex, properties that cannot easily be imposed on polynomials. Second, and even more serious, as equation (11) makes clear, fitting the data requires not only $f(x)$, but also its inverse $f^{-1}(x)$, the existence of which is guaranteed by the monotonicity of a function satisfying (5), again see Deaton and Laroque, Theorem 1. We found no convenient automatic way of guaranteeing the monotonicity of the polynomial approximation. Of course, these are not issues for model simulation, where the inversion is not required, and the polynomial approximations appear to work well.

An alternative to polynomial approximation is to calculate the function over a grid of points, the finer the grid the more accurate the approximation. The idea is to start from some suitable non-negative, convex, monotone non-increasing function $f_0(x)$, for example,

$$f_0(x) = \max\left(P(x), 0\right) \tag{12}$$

evaluated over a grid of points xx, say, where $xx = (x_1, x_2, \ldots, x_G)$. The function is then updated for each of the G grid points according to, c.f. (5):

$$f_{n+1}(x) = \max\left(\frac{1-\delta}{1+r}E_z f_n\left[z + (1-\delta)\left(x - P^{-1}\left(f_n(x)\right)\right)\right], P(x)\right) \tag{13}$$

Note that the updating rule requires the evaluation of points between the grid points, so that an interpolation rule is required. In our earlier work, we used a relatively fine grid (typically 100 points) together with linear extrapolation between points, a technique that worked well for simulation. However, the fineness of the grid comes at a computational price, and the linear interpolation means that the approximating function is non-differentiable at each of the grid points. These non-differentiabilities generate serious problems when we come to estimate the parameters, problems that are discussed in more detail below. We therefore replaced linear interpolation by a standard cubic spline procedure, as for example in Press et al (1989, pp. 86–89), a method that is rapidly computed, especially over a uniformly spaced grid. By construction, cubic splines are twice differentiable at the knots, or grid points, and the

interpolating cubic polynomials between the knots are required to be accurate only over a small range.

If the innermost $f_n(.)$ in (13) were to be replaced by $f_{n+1}(.)$, the iteration would be a contraction, and convergence would be guaranteed. However, such a procedure would require an iterative calculation for each new n, which would greatly increase computational time. And although the iteration defined by (13) is not generally a contraction, the procedure always seems to converge in practice, usually requiring about twenty iterations for the maximum distance between successive functions to be less than 10^{-5}. Again, if we could use the (infeasible) contraction, and if the splines were sufficiently accurate, the iterations would preserve non-negativity, monotonicity, and convexity. The feasible procedure, although not guaranteeing these properties, always delivered them in our calculations. Furthermore, the inversion of a function on a grid is straightforwardly accomplished by reversing the roles of the grid and the function, and using cubic splines again for interpolation, now on the non-uniform grid generated by the function values.

Given a method for calculating the function and its inverse, we can use equation (11) to generate a likelihood function for maximization. We use a normal distribution for the harvest process z_t, so that (11) generates the conditional density of p_t given p_{t-1}, $H(p_t \mid p_{t-1})$, say, with a log likelihood function

$$\ln L = \sum_{2}^{T} \ln H(p_{t+1} \mid p_t). \tag{14}$$

However, the maximization of (14) presents serious difficulties. Since the random variable z_t appears inside the function $f(.)$ in (11), the conditional density $H(p_t \mid p_{t-1})$ contains the Jacobian of the transformation, which requires the first derivative of $f(.)$. The problem here is not numerical, since given the splines the derivatives can readily be calculated, but that, as we have already seen, the function is not differentiable at the point where speculators begin to hold stocks. At this point, the Jacobian and the likelihood function will be discontinuous. While discontinuities are not necessarily incompatible with consistency of maximum likelihood estimates, there are obvious computational difficulties in dealing with such a function, or in locating a maximum which may be at a point of discontinuity.

To avoid these problems, we used (second-order) pseudo-maximum likelihood (PML) rather than full maximum likelihood. From equation (11), given the function $f(.)$ and its inverse, numerical integration can be used to calculate the mean and variance of p_t conditional on p_{t-1}. We write these:

$$\begin{aligned} m(p_{t-1}) &= E(p_t \mid p_{t-1}) \\ s(p_{t-1}) &= V(p_t \mid p_{t-1}) \end{aligned} \tag{15}$$

and then the (log) pseudo likelihood function

$$2\ln L = \sum_{2}^{T} \ln l_t = -(T-1)\ln(2\pi) - \sum_{2}^{T} \ln s(p_{t-1}) + \sum_{2}^{T} \frac{(p_{t-1} - m(p_{t-1}))^2}{s(p_{t-1})} \tag{16}$$

Equation (16) would be the exact likelihood function if, conditional on lagged price, price were normally but heteroscedastically distributed. If the theory is correct, this will not be the case, but the estimates will nevertheless be consistent, see Gourieroux, Monfort, and Trognon (1984). The Jacobian is not required by (16) and the pseudo-likelihood is continuous in the parameters. Of course, the non-differentiability of $f(.)$ carries through to (16); both functions m and s are differentiable almost everywhere, and the results of Laroque and Salanie (1991) guarantee the consistency and asymptotic normality of the estimates. The variance covariance matrix is calculated in the standard way: for the k-vector of parameters β, define the $(T\text{-}1)$ by k matrix G and the k by k matrix J by

$$g_{ti} = \frac{\partial \ln l_t}{\partial \beta_i}; \quad J_{ij} = \frac{\partial^2 \ln L}{\partial \beta_i \partial \beta_j}. \tag{17}$$

then the asymptotic variance covariance matrix of the parameters is given by

$$V = J^{-1}(T^{-1}G'G)J^{-1}. \tag{18}$$

In spite of the advantages, there is a real compromise in substituting the pseudo likelihood for the true likelihood. The non-negativity of inventories which is responsible for the kink in the price function also implies that the distribution of prices is asymmetric, since stockholders can always stop prices from going down by accumulating stock, but they cannot

stop prices from going up once their inventories have been exhausted. As a result, the conditional distribution of prices, $H(p_t \mid p_{t-1})$ displays both skewness and kurtosis. Pseudo maximum likelihood estimation pays attention only to the first two conditional moments, and ignores any correspondence (or lack of correspondence) between the theory and the model in higher moments. This is likely to unfairly penalize the storage model since the skewness that it predicts is clearly present in the data, see Deaton and Laroque (1992, Table 1). We have not been able to devise an extension to pseudo maximum likelihood that uses higher order moments. A generalized methods of moments estimator is clearly possible, but it is unclear to us how to combine the contributions from, say, the first four conditional moments.

To complete this methodological section, it remains to note only a few other points. We made no attempt to estimate the real interest rate r; instead it is fixed at 5% in all the calculations. Note that the role of the interest rate is to define, along with δ, the cost of storage. Although δ appears separately in the equation controlling the evolution of inventories, we are doubtful of our ability to identify the two parameters in practice. Real interest rates are typically not well estimated in empirical models of intertemporal allocation. We plan to do better than this in future work, not by estimating the parameter, but by reformulating the model to allow variable interest rates, and looking at the role of such variations in explaining the correlations among different commodity prices.

We also chose to replace the normal distribution for the harvest z_t by a discrete distribution; with equal probability 0.1 the harvest takes on the ten possible values (± 1.755, ± 1.045, ± 0.677, ± 0.386, ± 0.126). These values are the conditional means of ten equi-probable divisions of the standard normal distribution. Using them in place of the normal itself means that (costly) numerical integration is replaced by the simple operation of taking the mean of ten values, both when calculating the expectation in the iteration (13), and in calculating the conditional means and variances (15) from (11). In the few cases where we calculated both, the functions $f(.)$ calculated from the normal using numerical integration were not distinguishable from those calculated using the discrete distribution.

The selection of the grid xx over which the function $f(x)$ is evaluated was a matter of some delicacy. The cubic splines provide approximations only within the grid, and although it would have been possible to extrapolate beyond the end points in either direction, we took considerable pains to avoid this happening. While it is desirable to have the smallest possible

Table I: Grid limits and sizes for each commodity

	lower limit	upper limit	grid size
bananas	-5	20	10
cocoa	-5	20	10
coffee	-6	20	10
copper	-5	50	20
cotton	-5	10	10
jute	-5	15	10
maize	-5	30	10
palm-oil	-10	40	20
rice	-5	10	10
sugar	-5	10	10
tea	-5	15	10
tin	-5	40	20
wheat	-3	50	20

range of x, so that a given number of grid points gives a finer grid and greater accuracy, the range must be large enough so that, for any relevant parameter values, the repeated application of the iteration (13) never generates values of the amount on hand greater than the maximum value in xx. The highest value of x must be less than the value that would be reached if the harvest indefinitely remained at its maximum value, and if nothing were consumed, so that the amount in storage would eventually reach $1.755/\delta$. At the other extreme, and since inventories cannot be negative, the smallest value of x is the minimum harvest, i.e. -1.755, so that the minimum of xx cannot be greater than this. The use of these values as end points of the grid guarantees that the calculation of $f(x)$ will never generate values off the grid. However, there is still no guarantee that, for any given parameter values, one or more of the actual data points will not lie outside the range $[f(xx_G), f(xx_0)]$ in which case we have no means of calculating $f^{-1}(p)$ in (11). In these cases, we have found no better solution than stopping the algorithm, resetting the grid, and beginning again. Changing the grid within the estimation procedure is hazardous because grid changes alter likelihood values even without a change in parameter values, so that control of the algorithm becomes

extremely difficult. One possibility would be to add points to a pre-existing grid, leaving previous grid points intact, and although this should not alter the likelihood by much, the additional points would affect the cubic splines for the original points, and the likelihood evaluations would be affected. In practice, we were guided by the maxima and minima discussed above, but proceeded by trial and error. Table 1 lists lower and upper limits that were used for each of the commodities; we make no claims that persistent search might not generate tighter bounds.

Table 1 also shows the number of grid points for each good. In all cases, we began with ten, and increased it if the algorithm showed difficulty in finding a flat maximum of the likelihood, the idea being that a finer grid would allow a better approximation to the likelihood function and (perhaps) allow easier convergence. The fineness of the grid over x varied from commodity to commodity, but 10 or 20 points were generally sufficient to generate a well-defined maximum.

Finally, two of the parameters, b and δ, must satisfy restrictions for the model to make sense. In particular, b must be negative, and $\delta > -r$ for the model to be coherent and to guarantee that there exists a non-negative price function $f(x)$. It is not possible to estimate without these restrictions and check whether they are satisfied *ex post*; the calculations cannot be carried through for values outside the correct ranges. To guarantee the inequalities, we estimate the three parameter vector θ defined by (recall that r is set at 0.05):

$$a = \theta_1 \quad b = -e^{\theta_2} \quad \delta = -0.05 + e^{\theta_3} \tag{19}$$

As is perhaps apparent by this point, these calculations were not straightforward, although once we had implemented the procedures described above, the estimation proceeded without difficulty using a version of the Berndt, Hall, Hall and Hausman (1974) algorithm with numerical derivatives. Our early experiments using maximum likelihood and linear interpolation frequently showed multiple maxima; our final choice of technique has largely eliminated this problem, as we shall see in the Monte Carlo experiments reported below. All the calculations were carried out using GAUSS on a 386-series PC, although the algorithms were programmed directly; no use was made of the maximization routines in GAUSS.

3 Empirical results

The data used in this study come from the commodity division of the World Bank, and are the prices of thirteen commodities from 1900 to 1987; they are identical to the data used in Deaton and Laroque (1992). The commodities are bananas, cocoa, coffee, copper, cotton, jute, maize, palm-oil, rice, sugar, tea, tin, and wheat. The price data are indexes of annual averages prices divided by an index of the U.S. consumer price index. Better than the U.S. price would be an index of consumer prices in industrialized countries, but no such index is available over the period. The degree to which the storage model can be expected to apply to these commodities clearly varies a good deal from case to case. For many commodities, there have been more or less successful attempts to influence the world price through the formation of cartels, usually by the producers, but occasionally, as with coffee, by producers and consumers acting in concert. The history of these schemes is discussed in Gilbert (1988); all have tended to break down eventually, although some have clearly had effects on prices in the short-term if not in the long. Even so, we believe that it is a useful exercise to assess the extent to which the simple storage model is consistent with actual outcomes. If the model fits the data, one interpretation might be that the cartels have been successful only when the administered price was close to what the market would have set, while if the model fails, the actions of the cartels can be examined for an alternative explanation. Of course, there are other reasons, apart from cartels, why the model may not fit the data. On the demand side, there is no *a priori* reason for shocks to be uncorrelated over time. Indeed, over the last fifteen years, random walk models of aggregate consumption have been widely canvassed, and it would be plausible to find some autocorrelation in the demand for commodities, even if stationarity is maintained. On the supply side, metals like copper and tin do not have annual 'harvests' and it is at best unclear whether the analogy with agricultural crops is valid. Even among the genuine agricultural crops, bananas, cocoa, coffee and palm-oil are tree-crops, where it is hard to justify the supposition that harvest shocks are independent over time. Trees that are destroyed by fire and frost take several seasons to replace, so that even although the weather itself may be uncorrelated, its effects on price will not be. We regard this topic as an important one for further research. The model of Section 1 can be extended to handle persistence in shocks, but the econometric procedures are a good deal more

Table II: Parameter estimates for thirteen commodities

	a	*s.e.*	*b*	*s.e.*	δ	*s.e.*
bananas	0.5845	0.0167	-0.1363	0.0146	0.0026	0.0192
cocoa	0.1967	0.0144	-0.1562	0.0121	0.1206	0.0282
coffee	0.2634	0.0215	-0.1581	0.0254	0.1394	0.0191
copper	0.5216	0.0382	-0.3272	0.0570	0.0412	0.0266
cotton	0.6417	0.0410	-0.3119	0.0366	0.1688	0.0283
jute	0.5724	0.0268	-0.3560	0.0577	0.0956	0.0528
maize	0.6347	0.0428	-0.6357	0.1460	0.0589	0.0330
palm oil	0.4613	0.0524	-0.4292	0.0602	0.0575	0.0297
rice	0.5976	0.0270	-0.3357	0.0296	0.1471	0.0394
sugar	0.6428	0.0458	-0.6257	0.0649	0.1771	0.0329
tea	0.4788	0.0170	-0.2106	0.0248	0.1232	0.0328
tin	0.2557	0.0441	-0.1702	0.0490	0.1484	0.0513
wheat	0.7233	0.0357	-0.3938	0.0338	0.1302	0.0319

difficult, and we believe that the best way forward is to try to use data on harvests or stocks or on futures prices to complement the spot price data to better identify the parameters of such a model.

The parameter estimates for the thirteen goods are given in Table 2. Apart from bananas, where we still have reason to doubt the estimates, we believe that the computations leading to these estimates can be trusted. For bananas, the estimate of δ is so low that very large potential stocks must be accommodated in the calculations, and the resulting grid for the amount on hand is both long and coarse, and we have little confidence that we have actually maximized the pseudo-likelihood value. We have not given this problem much priority in our research because it is clear from the start that the simple storage model has no hope of explaining the behavior of the price of bananas. Over the data period, 1900–1987, the banana price has a first-order autocorrelation coefficient of 0.91, the highest among the thirteen commodities. To explain this autocorrelation, the model can only rely on the existence of banana inventories carried over from one harvest year to the next, a most implausible mechanism given the perishability of the commodity. The explanation of banana prices must

be sought elsewhere, most promisingly in the behavior of the oligopolistic cartel that controls a substantial fraction of output.

For all the commodities in the Table, there was no great difficulty in obtaining the parameter estimates *once we had come to understand the various issues raised in Section 1.* We used the gradient of the objective function as a convergence criterion, and the calculations were stopped when the elasticity of the criterion with respect to each parameter was less than 1.E-04. Based on the experience of earlier attempts with inferior approximation techniques, we have some concern about the presence of multiple maxima. For three of the commodities, maize, rice, and sugar, we have run Monte Carlo experiments in an attempt to assess the danger. For each good, we ran 100 experiments starting from the values of the θ parameters corresponding to the estimates in Table II. In each experiment, the starting values were obtained by multiplying the θ's by random numbers drawn from a lognormal distribution, so that the standard deviation of the log perturbation was 5 percent. For maize, 99 of the 100 experiments converged to the values in the table, and one converged to different estimates, (0.6362, –0.6320, 0.0611), and a log pseudo-likelihood of 32.051, as opposed to 32.056 for the estimates in the table and for the other 99 simulations. This is the only case of a multiple maximum we have found, in the maize experiments, or in those for the other two commodities. In these experiments, the modal number of iterations to convergence was 17, with a minimum of 9 and a maximum of 24. Of the 100 experiments for rice, one failed to produce estimates because of an inadequately large grid, and the other 99 converged to the estimates shown in Table II. In this case, more iterations were required, with a minimum of 28, a maximum of 60, and a mode of 40; 79 of the 99 cases converged in 35–44 iterations. For sugar, there were 7 cases where the starting values led to grids that were too small, while the other 93 converged to the same values. The minimum number of iterations was 11, the maximum 26, and the mode 18. [In the next version of this paper, we hope to report results of similar experiments for all the commodities, with the grid problems removed, and with the starting values further from the (provisional) final values.]

Suppose then that, apart from bananas, we accept the estimates in Table II. The estimates of δ in the final column are directly interpretable. The a and b parameters are best thought of in terms of the distribution of prices that would have been the case in the absence of storage, which, according to the model, is $N(a,b^2)$. Table III shows the actual means and

Table III: Means and standard deviations of commodity prices

	mean price	standard deviation	ratio standard deviations, post to pre-storage
bananas	0.58	0.10	0.72
cocoa	0.20	0.11	0.69
coffee	0.23	0.10	0.65
copper	0.49	0.19	0.57
cotton	0.65	0.22	0.72
jute	0.60	0.19	0.55
maize	0.71	0.27	0.43
palm oil	0.54	0.26	0.60
rice	0.62	0.23	0.67
sugar	0.71	0.43	0.68
tea	0.51	0.13	0.63
tin	0.22	0.09	0.54
wheat	0.69	0.26	0.67

Note: Columns 1 and 2 are the means and standard deviations of the deflated commodity prices (index numbers) from 1900–1987. The final column is the ratio of column 2 to the b parameters shown in Table II.

standard deviations of the prices, together with, in the final column, the ratio of the actual standard deviation to the absolute value of the b parameter in Table II. If the model were exactly correct, this ratio would show how much the standard deviation is reduced below what would be its hypothetical value in the absence of storage. This reduction is typically substantial, even for those commodities where the actual variability is very high. For example, the price of sugar has a coefficient of variation of 0.60 from 1900–1987, but according to the model, if there had been no storage, the coefficient would have been 0.88. The a parameters, which are the mean prices in the absence of storage, are close to the actual sample means. If the inventory depreciation rate δ were zero, storage would not change the mean price because the demand functions are linear; what is grown is eventually sold, and the effect of an additional unit of supply on price is the same whenever that unit is sold. With

Table IV: Log likelihood values for storage model, constant, and AR(1)

	storage model	i.i.d. normal	AR(1)
bananas	113.92	80.51	162.38
cocoa	120.62	72.06	125.48
coffee	111.01	76.32	120.24
copper	70.51	23.30	82.03
cotton	29.83	7.33	75.04
jute	44.79	19.58	50.75
maize	32.06	-10.37	27.34
palm oil	22.20	-5.44	27.91
rice	26.00	6.65	61.72
sugar	-10.69	-49.77	-27.34
tea	69.28	53.89	102.03
tin	108.93	85.43	152.63
wheat	24.64	-7.46	53.36

positive depreciation, storage destroys some of the commodity, and the mean price must rise. Comparing Tables II and III, we see that the sample mean is greater than the estimate of a for eight of the thirteen commodities, and that it is essentially identical in the case of bananas where the depreciation rate is essentially zero. For the remaining four goods, the difference has the wrong sign, but this could reflect sampling variability of the parameter estimate and the mean, as well as possible inadequacies in the model.

Table IV lists some measures of how well the model fits the data. Three likelihood or pseudo-likelihood values are shown. The first is the maximized pseudo log likelihood for the storage models and corresponds to the parameter estimates shown in Table II. The second column gives the absolute minimum standard that the model must beat, the log likelihood corresponding to the (clearly false) hypothesis that the prices are identical independent drawings from a normal distribution. The third and final column is the log likelihood of an AR(1) fitted to the data by ordinary least squares. The storage model comfortably beats the (absurd) i.i.d. model, but is beaten by the simple AR(1), except for the cases of maize and sugar. The simple autoregressive model has no theoretical support, but it is perhaps the first

model that any time series analyst would try, and the failure of the storage model to fit even this well is a serious condemnation of its ability to account for the data. Nevertheless, there are important respects in which the comparison in Table IV is unfair to the storage model. The AR(1) has no obvious theoretical justification, so that even if it fits better, it offers no coherent account of what is going on. It is not constrained to be stationary, nor can it reproduce the skewness and kurtosis in the data, which *are* generated by the storage model, but which are neither recognized nor rewarded by the pseudo-likelihood function. All this must be recognized, but it does little to gainsay the basic fact of Table IV, that the storage model cannot account for the degree of first-order autocorrelation displayed by the commodity prices.

Table V makes the same point more directly by comparing the actual and predicted autocorrelation coefficients for each commodity. The actual autocorrelations are calculated directly from the data; they are the first and second order (total) autocorrelation coefficients, and are reproduced from Table 1 of Deaton and Laroque (1992). The theoretical autocorrelations are calculated using a technique described more fully in the same paper. The basic idea is to use the parameter estimates in Table II to simulate the commodity model, and to use the simulated prices to calculate autocorrelations. In fact, such a procedure is not very precise because a very large number of simulations are required to obtain accurate estimates of the population parameters. The nonlinear Markov processes (11) generate time series with rare but very sharp spikes, and very long runs are required to get enough such spikes to permit the accurate calculation of population statistics. Instead, we divide the range of possible prices into 100 discrete points, and use the empirical estimates of the functions in (11) to calculate a transition matrix from the price in period t to the price in period $t+1$, with each point on the price grid capturing the points closest to it. This transition matrix has a unit eigenvalue, and the corresponding eigenvector gives the limiting probabilities of each price in the invariant stationary distribution. These probabilities, together with the transition probabilities define the limiting joint distribution of p_{t+1} and p_t, from which autocovariances, and thus autocorrelations, can be directly calculated. This calculation is harder to describe than to perform, and readily yields estimates that can be replicated only by simulations of several hundred thousand periods.

Table V: Theoretical and actual autocorrelation coefficients

	theoretical autocorrelations		actual autocorrelations	
	ρ_1	ρ_2	ρ_1	ρ_2
bananas	0.347	0.178	0.908	0.818
cocoa	0.298	0.142	0.834	0.660
coffee	0.242	0.099	0.804	0.621
copper	0.392	0.226	0.838	0.638
cotton	0.192	0.064	0.884	0.679
jute	0.302	0.145	0.713	0.448
maize	0.356	0.215	0.756	0.532
palm oil	0.416	0.247	0.730	0.478
rice	0.224	0.087	0.829	0.610
sugar	0.264	0.116	0.621	0.390
tea	0.230	0.086	0.778	0.593
tin	0.256	0.106	0.895	0.757
wheat	0.259	0.105	0.863	0.680

The two sets of figures in Table V bear little relationship to one another. The estimated storage models cannot reproduce the high levels of autocorrelation in the actual data. Indeed the predicted first-order autocorrelation coefficients are typically less than a half of the actuals, and the discrepancy is, if anything, larger for the two year autocorrelations. In the simulations reported in Deaton and Laroque (1992), we noted that, although the storage model could reproduce several features of the data, we were unable to find parameter values that generated autocorrelation coefficients as high as those in the data. Of course, it is impossible to run simulations for all possible parameters, so we were unable to conclude that the storage model was incapable of explaining highly autocorrelated prices. Table V provides the missing evidence. The best fitting storage models can account for year to year price autocorrelations of 0.3 or 0.4, but they cannot account for the actual autocorrelations, which are never lower than 0.6.

Figures 1 and 2 on the next two pages show how the storage model fails for two commodities, cotton and tea. Much the same phenomena are apparent for the other

commodities, but these two graphs show it most clearly. Each figure shows the actual price together with the one step ahead predictions from an AR(1) and from the storage models with parameters as estimated in Table II. The AR(1) fits the data badly, and cannot match the fact that prices come off their peaks quickly, so that its prediction is clearly not well synchronized with the actual data. Clearly, the data contain more complex dynamics than are allowed for by this simple specification. However, even the AR(1) does much better than the predictions of the storage model, which do not track the data at all whenever prices are high. In both graphs, the storage model predictions have upper limits, and are constant at that upper limit, so that there are periods of constant price predictions that are clearly contradicted by the most elementary inspection of the data, or indeed by a simple and poorly fitting AR(1). This failure of the model is not something that is specific to these two commodities, or to the peculiarities of the parameter estimates, but is fundamental to the storage model when harvests are independently and identically distributed. With optimal storage, the autoregression function of prices is given by (8), which can be rewritten as

$$E(p_t|p_{t-1}) = \frac{1+r}{1-\delta}\min(p^*, p_{t-1}) \tag{20}$$

for constant $p^* = E(f(z))$. When prices are above the critical value p^*, and because with i.i.d. harvests high prices are no indication of poor z's in the future, speculators will not choose to hold inventories, because next year's price is not expected to be high enough to justify the interest and depreciation costs of carrying inventories into the next period. But with no serial correlation in either demand or supply behavior, inventories are the only link between the present and the future, so the expectation of next period's price is a constant p^*, which is simply the expectation with respect to the harvest of $f(z)$, the price when there are no inventories carried forward. Hence, as in (16), the model only predicts autocorrelation when prices are below p^*; above that level, expected prices are constant. This prediction is clearly falsified by the behavior of the actual prices, where there is clear evidence of positive autocorrelation even when prices are high. The AR(1), with all its manifest faults, captures this 'high price' autocorrelation better than does the storage model, and it is this that accounts for its superior performance in the likelihood comparisons in Table III.

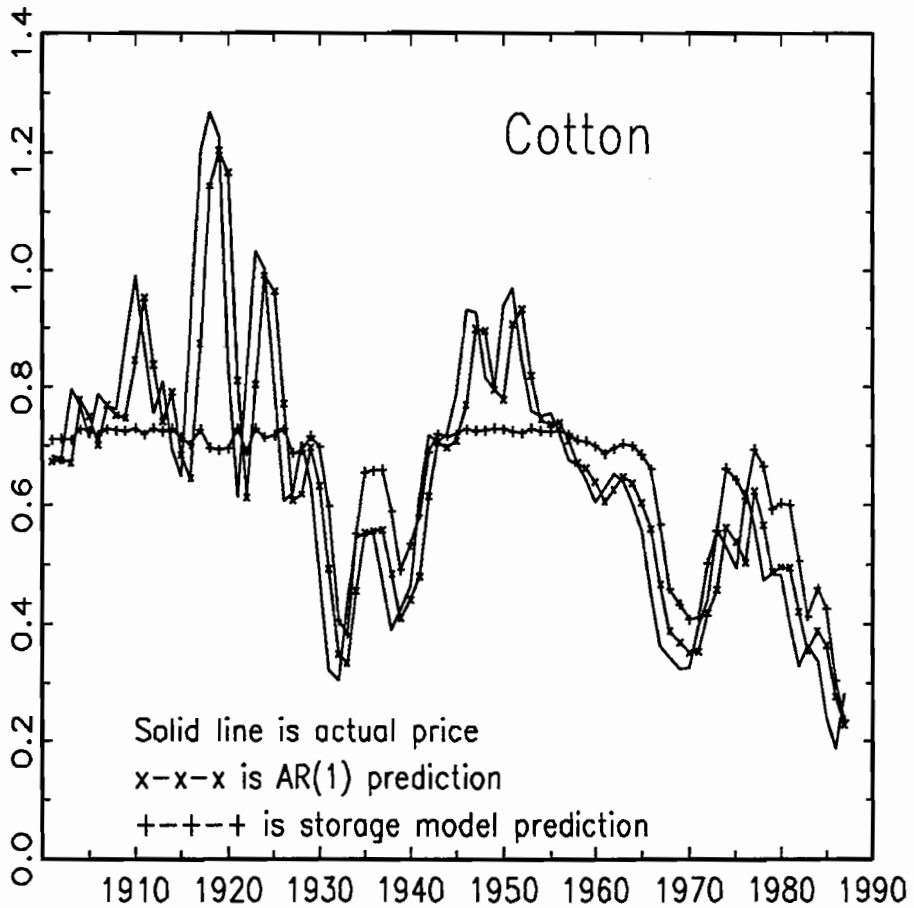

Figure 1: Actual and one period ahead expectations of cotton prices

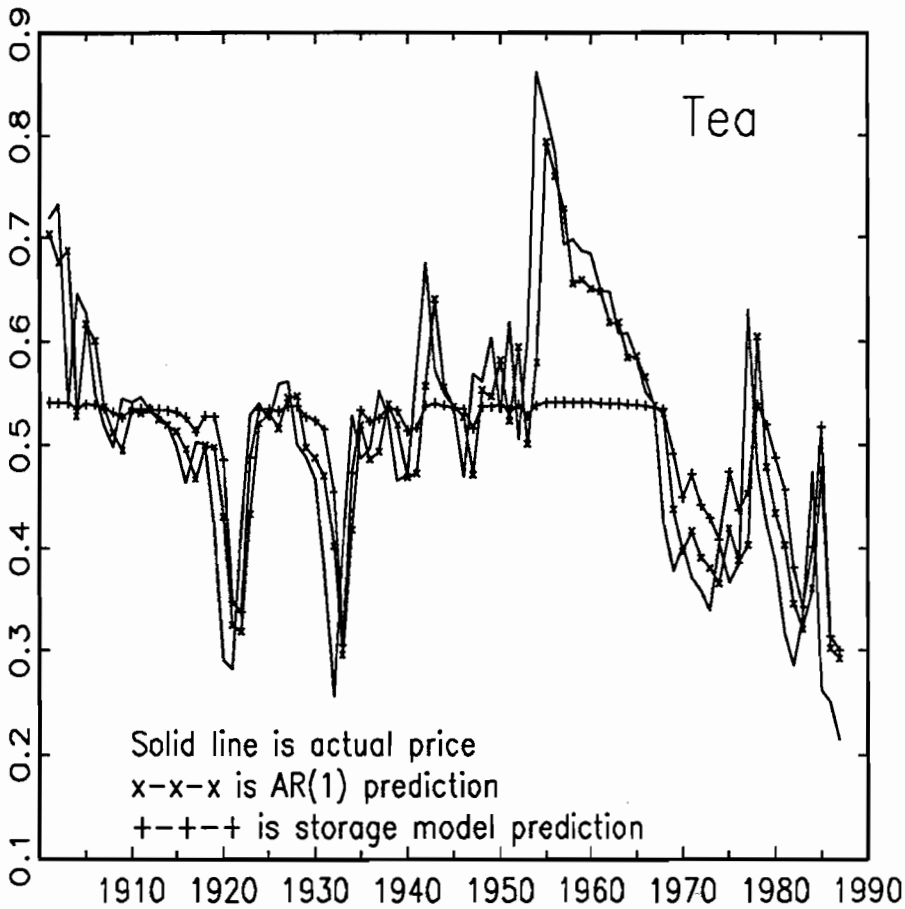

Figure 2: Actual and one period ahead expectations of tea prices

In the speculative storage model with i.i.d. harvests, it is the behavior of profit maximizing risk-neutral speculators that generates autocorrelation in prices. If in the absence of storage prices would be abnormally low, the demand for speculative storage will bid up the price, and the release of stocks when prices are high will moderate what would otherwise be even higher prices. Adjacent prices are therefore brought into closer conformity than would be the case without the speculators. If this mechanism is to account for the high levels of autocorrelation in the data, stocks must normally be held, and on average it must be profitable (or at least not loss-making) for speculators to hold them. In fact, given the actual data, it is hard to simultaneously reconcile profitable speculation and infrequent stockouts. Consider a speculator whose strategy is to buy one unit of the commodity whenever its price is below some critical level, and to resell it in the next period. When harvests are i.i.d., we know that the expected profit associated with this strategy is zero whenever the critical price is below p^*, and becomes negative when the critical price is above p^*. If the model is true, prices can be expected to rise to cover the cost of holding inventories whenever prices are below p^*, so that if we calculate average actual profits for a range of different cutoffs, profits should be zero for low prices, and negative as soon as our cutoff is higher than p^*. If the autocorrelation is to be explained, the losses should not start until the cutoff is at a relatively high point in the actual distribution of prices, since losses are not compatible with stocks, and autocorrelation is not compatible with stockouts. In fact, if the cutoff price is set at the sample median price for each commodity, there will be stockouts at half the sample points, and yet, if the total cost of storage is no less than 5% per annum, speculators would have made losses over the period as a whole on all of the commodities except cocoa, coffee, and maize. If the cutoff is at the upper quartile, with 25% stockouts, speculators would have made losses on all of the commodities. Such statistics give further emphasis to the evidence in Figures I and II. It is simply not plausible that the behavior of these commodity prices can be explained in a model where the only source of serial correlation is the storage activity of risk-neutral speculators.

The storage model with i.i.d. harvests is the leading model of commodity prices in the literature, and has been since the original contribution of Gustafson in 1958. The results of this paper make it hard for us to believe that it can offer a coherent account of the data, at least not without serious modification. If actual prices are to be explained, we require some

other source for the high degree of serial correlation displayed by the actual prices. As always, there are two possibilities, supply and demand. For some crops, particularly for perennial tree crops, there is undoubtedly autocorrelation in the harvest shocks. Damage to trees, although ultimately repairable, is clearly capable of generating supply shocks that last through several periods. However, it is much less clear that for annuals such as cotton, jute, maize, rice, sugar, and wheat, that the technology of production can be expected to display enough serial dependence to account for the very large gap between actual and predicted autocorrelations in Table V. Not does inspection of that Table suggest that the problems are exclusively with metals and tree-crops, or that they are worse for those goods than for the annuals. We suspect that the most plausible source of autocorrelation is on the demand side. Autocorrelation can be much more readily explained if there is a demand for stocks even when prices are expected to fall, a demand that might be provided by the activities of processors or manufacturers, for whom the commodity is a factor of production, and who will pay to avoid the possibility that they might run out of stock. Modelling such activities in a satisfactory way is no easy task, but there would seem to be no way of avoiding it if the data are to be explained.

3. List of Works Cited

Berndt, Ernst K., Bronwyn H. Hall, Robert E. Hall, and Jerry A. Hausman, 1974, 'Estimation and inference in nonlinear structural models,' *Annals of Economic and Social Measurement*, **3**, 653–666.

Deaton, Angus and Guy Laroque, 1992, 'On the behavior of commodity prices,' *Review of Economic Studies*, forthcoming.

Gilbert, Christopher L., 1987, 'International commodity agreements: design and performance,' *World Development*, **15**, 591–616.

Gourieroux, Christian, Alain Monfort, and A. Trognon, 1984, 'Pseudo maximum likelihood methods: theory,' *Econometrica*, **52**, 681–700.

Judd, Kenneth L., 1990, 'Minimum weighted residual methods for solving dynamic economic models,' Stanford University, Hoover Institution, March.

Laroque, Guy and Bernard Salanié, 1990, 'Estimating the canonical disequilibrium model: asymptotic theory and finite sample properties,' Discussion paper # 9005, INSEE, Paris.

Newbery, David M.G., and Joseph E. Stiglitz, 1981, *The theory of commodity price stabilization: a study in the economics of risk*, Oxford. Oxford University Press.

Press, William H., Brian P. Flannery, Saul E. Teukolsky, and William T. Vettering, *Numerical recipes: the art of scientific computing*, Cambrudge. Cambridge University Press.

Samuelson, Paul A., 1971, 'Stochastic speculative price,' *Proceedings of the National Academy of Science*, **68**, 335–337.

Taylor, John B., and Harald Uhlig, 1990, 'Solving nonlinear stochastic growth models: a comparison of alternative solution methods,' *Journal of Business and Economic Statistics*, **8**, 1–17.

Williams, Jeffrey C., and Brian D. Wright, 1991, *Storage and commodity markets*, Cambridge. Cambridge University Press.

Appendix: Proof of Proposition 1

Let $f(x)$ be the rational expectations equilibrium associated with the base economy, in which harvests have zero mean and unit variance, and in which the inverse demand function is linear, equation (1). Denote $I(x)$ the amount of inventories carried forward when the amount on hand is x, and given by equation (7). By definition, $f(x)$ is the unique non-negative non-decreasing solution of the functional equation

$$f(x) = \max\left[\frac{1-\delta}{1+r}E_z f\big(z + (1-\delta)I(x)\big), P(x)\right]$$

see (5) and (7) above, or Deaton and Laroque (1992) for the derivation and a detailed analysis of this equation.

We shall prove that $\tilde{f}(x)$, the rational expectations equilibrium price function associated with the economy $\tilde{E}$, satisfies:

$$\tilde{f}(\mu+\sigma x) = f(x) \tag{A.1}$$

and that the corresponding inventory function for the economy $\tilde{E}$, $\tilde{I}(x)$, satisfies:

$$\tilde{I}(\mu+\sigma x) = \sigma I(x). \tag{A.2}$$

Now since the evolution of the state variable is governed by (3):

$$\tilde{x}_{t+1} = (1-\delta)\tilde{I}(\tilde{x}_t) + \tilde{z}_t.$$

So that letting $\tilde{x}_{t+1} = \mu+\sigma x_{t+1}$, we see from (A.2) that:

$$x_{t+1} = (1-\delta)I(x_t) + z_{t+1}$$

and x_{t+1} is the amount on hand in the base economy. It then follows from (A.1) and (A.2) that the price processes of the two economies are identical, and the inventory processes are proportional with a factor of proportionality equal to σ.

We complete the proof by showing that (A.1) and (A.2) are in fact correct. By definition:

$$\tilde{f}(\tilde{x}) = \max\left[\frac{1-\delta}{1+r}E_z\tilde{f}\left(\tilde{z}+(1-\delta)\tilde{I}(\tilde{x})\right), \tilde{P}(\tilde{x})\right]. \tag{A.3}$$

We let $\tilde{x} = \mu+\sigma x$. First, from (9),

$$\tilde{P}(\tilde{x}) = \left(a-\frac{b\mu}{\sigma}\right)+\frac{b}{\sigma}(\mu+\sigma x) = P(x)$$

Second:

$$\begin{aligned}\tilde{I}(\tilde{x}) &= \mu+\sigma x-\tilde{P}^{-1}\left(\tilde{f}(\tilde{x})\right)\\ &= \mu+\sigma x-\frac{\tilde{f}(\tilde{x})-a+b\mu/\sigma}{b/\sigma}\\ &= -\sigma\left[x-P^{-1}(\tilde{f}(\tilde{x})\right].\end{aligned}$$

Substituting these expressions for $\tilde{P}(\tilde{x})$ and $\tilde{I}(\tilde{x})$ into the right hand side of (A.3), we obtain

$$\tilde{f}(\tilde{x}) = \max\left[\frac{1-\delta}{1+r}E_z\tilde{f}\left(\mu+\sigma z+(1-\delta)\sigma\left(x-P^{-1}(\tilde{f}(\tilde{x}))\right)\right), P(x)\right]$$

so that, since $x = (\tilde{x}-\mu)/\sigma$, we have,

$$\tilde{f}(\mu+\sigma x) = \max\left[\frac{1-\delta}{1+r}E_z\tilde{f}\left(\mu+\sigma\left[z+(1-\delta)\left(x-P^{-1}(\tilde{f}(\mu+\sigma x))\right)\right]\right), P(x)\right].$$

But this is exactly the same functional equation that defines $f(x)$, and since that functional equation has a unique solution, we have immediately that

$$\tilde{f}(\mu+\sigma x) = f(x).$$

Recent RPDS Working Papers

100. "Environmental Pollution and Optimum Taxation Under Market Distortions in Developing Countries," by Gill C. Lim and Joon Koo Lee (October, 1981), pp. 18.

101. "Decomposing Multipliers for Economic Systems Involving Regional and World Trade," by Jeffery I. Round (December, 1981), pp. 26.

102. "Environmental Policies in Developing Countries: A Case of International Movements of Polluting Industries," by Joon Koo Lee and Gill Chin Lim (March, 1982), pp. 19.

103. "Social Welfare Underpinnings of Urban Bias and Unemployment," by Richard J. Arnott and Mark Gersovitz (April, 1982), pp. 32.

104. "A Preliminary Exploration of the Links Between Long-Term Environmental Management and Economic Development," by H. Jeffrey Leonard (June, 1982), pp. 71.

105. "Production Relations in Agriculture," by Hans P. Binswanger and Mark R. Rosenzweig (June, 1982), pp. 56.

106. "Savings, Nutrition and Development," by Mark Gersovitz (June, 1982), pp. 24.

107. "Riverains and Lacustrines: Toward International Cooperation in the Nile Basin," by John Waterbury (September, 1982), pp. 150.

108. "Trade, Capital Mobility and Sovereign Immunity," by Mark Gersovitz (September, 1983), pp. 27.

109. "Macroeconomic Policy Under Currency Inconvertibility," by Jorge Braga de Macedo (December, 1983), pp. 33.

110. "Agro-Industrial Processing and Agricultural Pricing Under Uncertainty," by Mark Gersovitz (December, 1983), pp. 43.

111. "The Demand for Personal Travel in Developing Countries: Pricing and Policy Analysis," by Angus Deaton (April, 1984), pp. 39.

112. "Econometric Issues for Tax Design in Developing Countries," by Angus Deaton (January, 1984), pp. 31.

113. "Optimal Pricing of Publicly Supplied Goods and Services," by Michael L. Katz (April, 1984), pp. 28.

114. "The Effects of Domestic Taxes on Foreign Private Investment," by Mark Gersovitz (April, 1984), pp. 41.

115. "Implementation of Environmental Impact Assessment in Developing Countries," by Gill Chin Lim (July, 1984), pp. 44.

116. "Issues in the Methodology of Multimarket Analysis of Agricultural Pricing Policies," by Angus Deaton, July 1984, pp. 22.

117. "Small Countries in Monetary Unions: The Choice of Senegal," by Jorge Braga de Macedo, March 1985, pp. 77.

118. "Small Countries in Monetary Unions: A Two-Tier Model," by Jorge Braga de Macedo, March 1985, pp. 56.

119. "Does it Make Sense to Collect Panel Data for Developing Countries?" by Orley Ashenfelter, Angus Deaton, Gary Solon, May 1985, pp. 52.

120. "Land Policy in Developing Countries: A Conceptual Framework and A Case Study," by Gill-Chin Lim, June 1985, pp. 28.

121. "The New Development Economics," by Joseph E. Stiglitz, July 1985, pp. 28.

122. "The Influence of Household Composition of Household Expenditure Patterns, Theory and Spanish Evidence," by Angus Deaton, Javier Ruiz-Castillo, and Duncan Thomas, September 1985, pp. 53.

123. "Optimally Uniform Commodity Taxes, Taste Differences and Lump-sum Grants," Angus S. Deaton, Princeton University and Nicholas Stern, Warwick University, June 1985, pp. 10.

124. "The Pure Theory of Country Risk," Jonathan Eaton, University of Virginia, Mark Gersovitz and Joseph E. Stiglitz, Princeton University, October 1985, pp. 58.

125. "Poverty: Measurement, Alleviation and the Impact of Macroeconomic Adjustment," by S. M. Ravi Kanbur, Department of Economics, University of Essex and Woodrow Wilson School, Princeton University, November 1985.

126. "Consumers' Expenditure," by Angus Deaton, Department of Economics, Princeton University, July 1986, pp. 63.

127. "Quality, Quantity, and Spatial Variation of Price," by Angus Deaton, Living Standards Unit, The World Bank, and Research Program in Development Studies, Princeton University, November, 1986, pp. 49.

128. "Estimation of Own and Cross-price Elasticities from Household Survey Data," Angus Deaton, Living Standards Unit, The World Bank, and Research Program in Development Studies, Princeton University, January, 1987, pp. 33.

129. "Food Subsidies and Poverty Alleviation", by Timothy Besley, All Souls College, Oxford, U.K. & Institute of Economics and Statistics, University of Oxford, U.K. and Ravi Kanbur, Department of Economics, University of Warwick, Coventry, U.K. & Woodrow Wilson School, Princeton University, USA, March, 1987, pp. 38.

130. "The allocation of goods within the household: adults, children, and gender," Angus Deaton, Living Standards Unit, The World Bank and the Research Program in Development Studies, Princeton University, June, 1987.

131. "Saving and Development", by Mark Gersovitz, Woodrow Wilson School, Princeton University, July 1987, pp. 80.

132. "Country Risk and the Organization of International Capital Transfer", by Jonathan Eaton, University of Virginia and NBER, and Mark Gersovitz, Woodrow Wilson School, Princeton University, May 1987, pp. 38.

133. "Transportation, State Marketing, and the Taxation of the Agricultural Hinterland", by Mark Gersovitz, Woodrow Wilson School, Princeton University, September 1987, pp. 33.

134. "Household surveys and policy reform: cocoa and coffee in the Cote d'Ivoire", Angus Deaton and Dwayne Benjamin, Welfare and Human Resources Division, The World Bank and the Research Program in Development Studies, Princeton University, November 1987.

135. "On the Use of Spatial Autoregressive Models in Demand Analysis", by Anne C. Case, Economics Department, Princeton University, October 1987.

136. "Agricultural pricing policies and demand patterns in Thailand", by Angus Deaton, Research Program in Development Studies, Woodrow Wilson School, Princeton University, February 1988, Revised, March 1988.

137. "Household Savings in Thailand: Responses to Income Shocks", by Christina H. Paxson, Research Program in Development Studies, Woodrow Wilson School, Princeton University, February 1988.

138. "Price elasticities from survey data: Extensions and Indonesian results", by Angus Deaton, Research Program in Development Studies, Woodrow Wilson School, Princeton University, May 1988.

139. "New approaches to household survey data from developing countries", by Angus Deaton, Research Program in Development Studies, Woodrow Wilson School, Princeton University, July 1988.

140. "Household Composition and Labor Demand: A Test of Rural Labor Market Efficiency", by Dwayne Benjamin, Research Program in Development Studies, Woodrow Wilson School, Princeton University, July 1988, Revised November 1988.

141. "A Shadow Wage Approach to Estimating a Labor Supply Model for Peasant Households: An Application to the Peruvian Sierra" by Hanan Jacoby, Research Program in Development Studies, Center of International Studies, Woodrow Wilson School, Princeton University, February 1988, Revised November 1988.

142. "Workfare vs. Welfare: Incentive Arguments for Work Requirements in Poverty Alleviation Programs" by Timothy Besley, Research Program in Development Studies, Woodrow Wilson School, Princeton University and Stephen Coate, John F. Kennedy School of Government, Harvard University, October 1988.

143. "General Equilibrium with Parallel Markets for Goods and Foreign Exchange: Theory and Application to Ghana" by Jean-Paul Azam, CERDI, University of Clermont-Ferrand and Tim Besley, Research Program in Development Studies, Woodrow Wilson School, Princeton University, Revised December 1988.

144. "Saving in Developing Countries: Theory and Review" by Angus Deaton, Research Program in Development Studies, Woodrow Wilson School, Princeton University, March 1989.

145. "On the Behavior of Commodity Prices" by Angus Deaton, Research Program in Development Studies, Woodrow Wilson School, Princeton University and Guy Laroque, INSEE, Paris, July 1989.

146. "Targeting Taxes and Transfers: Administrative Costs and Policy Design in Developing Economies" by Timothy Besley, Research Program in Development Studies, Woodrow Wilson School, Princeton University, Revised October 1989.

147. "Gender Effects in Indian Consumption Patterns" by Shankar Subramanian, Indira Gandhi Institute of Development Research, Bombay and Angus Deaton, Research Program in Development Studies, Woodrow Wilson School, Princeton University, April 1990.

148. "Patterns of Aging in Thailand and Côte d'Ivoire" by Angus Deaton and Christina H. Paxson, Research Program in Development Studies, Woodrow Wilson School, Princeton University, March 1990, Revised May 1990.

149. "The Economics of Rotating Savings and Credit Associations" by Timothy Besley, Research Program in Development Studies, Princeton University, Stephen Coate, Department of Economics, University of Pennsylvania, and Glenn Loury, John F. Kennedy School of Government, Harvard University, May 1990.

150. "Consumption and Income Seasonality in Thailand" by Christina H. Paxson, Research Program in Development Studies, Center of International Studies, Woodrow Wilson School, Princeton University, January 1991.

151. "Demand Analysis for Tax Reform in Pakistan" by Angus Deaton and Franque Grimard, Research Program in Development Studies, Center of International Studies, Woodrow Wilson School, Princeton University, February 1991.

152. "Group Lending, Repayment Incentives and Social Collateral" by Timothy Besley, Research Program in Development Studies, Center of International Studies, Woodrow Wilson School, Princeton University and Stephen Coate, Department of Economics, University of Pennsylvania, June 1991.

153. "Household Saving in LDC's: Credit Markets, Insurance, and Welfare" by Angus Deaton, Research Program in Development Studies, Center of International Studies, Woodrow Wilson School, Princeton University, April 1991, revised August, 1991.

154. "Estimating the Commodity Price Model" by Angus Deaton, Research Program in Development Studies, Center of International Studies, Woodrow Wilson School, Princeton University, and Guy Laroque, INSEE, Paris, May and November 1991.