

Estimating the Competitive Storage Model with Trending Commodity Prices*

Christophe Gouel[†]

Nicolas Legrand[‡]

May 29, 2015

Abstract

We present a method to estimate jointly the parameters of a standard commodity storage model and the parameters characterizing the trend in commodity prices. This procedure allows the influence of a possible trend to be removed without restricting the model specification, and allows model and trend selection based on statistical criteria. The trend is modeled deterministically using linear or cubic spline functions of time. The results show that storage models with trend are always preferred to models without trend. They yield more plausible estimates of the structural parameters, with storage costs and demand elasticities that are more consistent with the literature. They imply occasional stockouts, whereas without trend the estimated models predict no stockouts over the sample period for most commodities. Moreover, accounting for a trend in the estimation imply price moments closer to those observed in commodity prices. Our results support the empirical relevance of the speculative storage model, and show that storage model estimations should not neglect the possibility of long-run price trends.

Keywords: Commodity prices, non-linear dynamic models, storage, structural estimation, trend.

JEL classification: C51, C52, Q11.

*We would like to thank Fabien Tripier, Christophe Muller, Alain Ayong le Kama, and Stephane de Cara for their helpful suggestions. We are grateful to Carlo Cafiero for sending us its estimation code. This research was generously supported by the European Union's Seventh Framework Programme FP7/2007–2011 under Grant Agreements no 290693 FOODSECURE and by Total – Scientific Division. The authors only are responsible for any omissions or deficiencies. The European Union and European Commission are not accountable for the content of the paper.

[†]INRA, Économie Publique, and CEPII (christophe.gouel@grignon.inra.fr)

[‡]INRA, Économie Publique, and Climate Economics Chair (nicolas.legrand@grignon.inra.fr)

1 Introduction

Gustafson's (1958) commodity storage model is fundamental for explaining the annual behavior of commodity prices. It features forward-looking speculators that maximize profit by stockpiling a commodity based on the difference between the expected price and the current price. The source of volatility in the commodity storage model is the occurrence of unexpected supply shocks. The model has proven capable of reproducing many features of commodity prices such as sharp spikes, volatility clustering, positive skewness, and excess kurtosis (Deaton and Laroque, 1992). However, in early estimations of this model, Deaton and Laroque (1992, 1996) show that it could not explain the high degree of serial correlation observed in the price series. This finding was challenged. Cafiero et al. (2011b) show that using a finer grid to approximate the policy function and a different model specification, the storage model is able to generate the observed serial correlation for seven of the twelve commodities analyzed in Deaton and Laroque (1996). Since Cafiero et al. (2011b), several papers provide positive evidence for the role of storage arbitrage in price behavior (Bobenrieth et al., 2013; Cafiero et al., 2015; Guerra et al., forthcoming). However, if the model is estimated on untransformed real price indexes (as in Cafiero et al., 2011b), discretionary stocks are always strictly positive (i.e., there are no "stockouts") for most commodities over the sample interval. This result casts doubt on the appropriateness of using for estimation a nonlinear model with two regimes (with and without stocks) if, over long samples and for most commodities, the estimations imply that only one regime is active.

The absence of stockouts may be a consequence of the attempt to fit with the storage model a serial correlation that is artificially high, due to a possible non-stationarity in the price series. Commodity prices are unlikely to be stationary over long periods. Starting with the work by Prebisch (1950) and Singer (1950), a large literature has been devoted to characterizing the nature of this non-stationarity: whether trends are stochastic or deterministic, the existence of long-run cycles, or secular decline of commodity prices relative to those of manufactures (e.g., Grilli and Yang, 1988; Ardeni and Wright, 1992; Cuddington, 1992; Cashin and McDermott, 2002). Estimating the storage model, which features prices converging to a stationary distribution, with untransformed prices (as in Deaton and Laroque, 1992, 1996; Cafiero et al., 2011b) is likely to lead to biased parameter estimates if prices are non-stationary. The present article assesses the role of potential non-stationary price series in estimations of the storage model, and proposes an approach that statistically accounts for a trend in the price series.

How to estimate dynamic stochastic rational expectations models that are defined to be stationary around a steady state using non-stationary data is a very important question in the related literature on the estimation of DSGE models. In a recent paper, Canova (2014) summarizes the various strategies used in this literature. Most apply also to the storage model. Most DSGE models are estimated on transformed data in two steps. First, a statistical filter (linear detrending, Hodrick and Prescott filter, first-order differencing, or band-pass filter) is applied to the raw data, then the structural model is estimated using the transformed data. This approach has the attraction of convenience but is known to involve problems. The business cycle facts will depend on the choice of filter which is arbitrary (Harvey and Jaeger, 1993; Canova, 1998), due to lack of formal tests to select the most appropriate trend specification. For the storage model, the two-step approach is applied in Cafiero et al. (2011a), Bobenrieth et al. (2013, 2014), and Guerra et al. (forthcoming) where prices are detrended ex-ante using a log-linear trend.

Another approach involves the construction of a model that includes transitory and permanent shocks, the latter aimed at capturing non-cyclical fluctuations. The model is made stationary by scaling it by the

permanent shocks, and is fitted to the raw data. This approach has the appeal of theoretical consistency but introduces the risk of misspecification. Because it is not possible to make every model stationary for all possible specifications of permanent shocks, the model design and the nature of the shocks may be driven more by computational than economic motives. This approach is applied to the storage model in [Zeng \(2012\)](#). In [Zeng's](#) model, storers internalize the downward trend in commodity prices, and adjust their behavior accordingly. However, to obtain a stationary arbitrage equation while including a trend in prices, storage costs must either be zero or have the same trend as prices. Another issue related to this approach in the context of the storage model is that the structural trend included in the model may have to capture non-cyclical fluctuations which might be far from structural. The Boskin Report ([Boskin et al., 1996](#)) has initiated debate over a possible positive bias in the construction of the US CPI which is used in most works to deflate commodity prices, following [Deaton and Laroque \(1992\)](#). Potential bias in the deflator has numerous consequences for the literature on commodity prices dynamics. [Svedberg and Tilton \(2006\)](#) show that by adjusting the US CPI for the bias estimated by the Boskin Commission, the conclusion of a downward trend in the price of copper can be reversed. A related issue is the choice of price deflator. [Fernandez \(2012\)](#) shows that the conclusions related to the trend in commodity prices are not robust to the choice of price deflator. Because of the uncertainties related to the choice of an appropriate deflator and its potential accuracy, we prefer not to introduce the trend in the structural model, and pursue a different approach.

Here, we adopt the alternative method to estimate DSGE models using raw data proposed by [Canova \(2014\)](#).¹ In this approach, the econometrician defines a statistical model which is a combination of a DSGE and a reduced-form model; the reduced-form is aimed at capturing the component in the data that the structural model is unable to explain. This statistical model can be estimated using raw data, which leads to joint estimation of the structural and reduced-form parameters. Interestingly, this one-step approach allows us to select the most likely trend specification based on a statistical criterion for model selection. In this paper, we apply this approach by jointly estimating a storage model and a reduced-form trend that describes the non-cyclical component of price. The estimation procedure starts from the Maximum Likelihood estimator proposed for the storage model by [Cafiero et al. \(2015\)](#), and which was proved to have better small sample properties than [Deaton and Laroque's \(1996\)](#) Pseudo-Maximum Likelihood estimator. We extend the Maximum Likelihood estimator to account for a potential trend in prices and to exploit the information available from the first observation. This leads to the development of a new simulated unconditional Maximum Likelihood estimator. We consider only deterministic trend specifications because this assumption allows the likelihood to be expressed analytically.² As well as the case without trend, we consider a multiplicative trend, in which the logarithm of the trend can be linear as in [Cafiero et al. \(2011a\)](#), [Bobenrieth et al. \(2013, 2014\)](#), and [Guerra et al. \(forthcoming\)](#), or represented by a restricted cubic spline as in [Roberts and Schlenker \(2013\)](#). For the thirteen storable commodities considered in [Deaton and Laroque \(1992\)](#), there is a model with trend which presents a lower Akaike information criterion than the model without trend. Our estimates for the preferred models more closely replicate the key features of the data and allow for the occurrence of stockouts in line with the observed two-regime structure of long periods of stable prices interrupted by isolated spike episodes.

The remainder of the paper is organized as follows. Section 2 describes the competitive storage

¹See also [Ferroni \(2011\)](#) for an application.

²Stochastic trends would require a non-linear state-space approach and the use of particle filters ([Fernández-Villaverde and Rubio-Ramírez, 2007](#)), a promising but challenging approach for a model as non-linear as the storage model.

model discussed and estimated in [Deaton and Laroque \(1992, 1996\)](#) and [Cafiero et al. \(2011b\)](#). Section 3 presents the econometric procedure used to estimate the storage model with multiplicative deterministic trend, and describes how the unconditional maximum likelihood estimator is constructed. Section 4 presents the empirical results, and section 5 concludes.

2 The model

2.1 Model equations

We adopt the standard competitive storage model with no supply response, constant marginal storage cost, and no stock deterioration in line with [Cafiero et al. \(2011b\)](#). The exogenous supply is modeled by i.i.d. random production shocks ε_t following a normal distribution with mean μ and standard deviation σ truncated at five standard deviations. The demand for commodities consists of a demand for current consumption C_t associated with the inverse demand function $D^{-1}(C_t) = a + bC_t$, which is assumed to be linear with fixed parameters a and $b < 0$, and a speculative demand from competitive risk-neutral storers. Storers carry over $S_t \geq 0$ units of the commodity into the next period whenever they expect a positive return to storage over the interest and physical storage costs, and otherwise sell their past inventories. Assuming rational expectations and taking account of the non-negativity constraint on storage yields the following arbitrage condition:

$$\beta E_t P_{t+1} - P_t - k \leq 0, = 0 \text{ if } S_t > 0, \quad (1)$$

where $\beta = 1/(1+r)$ is the discount factor which is assumed to be fixed, $k \geq 0$ is the constant per unit physical cost of storage, P_t is the price, and E_t is the expectation operator conditional on period t information. In equilibrium, supply equals total demand such that

$$A_t = S_t + D(P_t), \quad (2)$$

where the amount on hand A_t at time t is the sum of the past inventories and the stochastic production ε_t written as

$$A_t \equiv S_{t-1} + \varepsilon_t, \quad (3)$$

with $A_t \in \mathbb{A} \equiv [-5\sigma, \infty)$.

Combined with the market clearing condition, the arbitrage condition (1) leads to two regimes in the price dynamics:

$$P_t = \max [\beta E_t P_{t+1} - k, D^{-1}(A_t)]. \quad (4)$$

The first regime holds when speculators stockpile expecting the future price to cover the full carrying costs and the purchasing cost. The second regime defines the stockout situation with empty inventories, where the market price is determined only by the final demand for consumption and the amount on hand in the market.

For this problem, a stationary rational expectations equilibrium is a price function $\mathcal{P} : \mathbb{A} \rightarrow \mathbb{R}$ which describes price as a function of contemporaneous availability. From equation (4), this price function satisfies for all A_t

$$\mathcal{P}(A_t) = \max [\beta E_t \mathcal{P}(S_t + \varepsilon_{t+1}) - k, D^{-1}(A_t)], \quad (5)$$

where, from (2), S_t is given by

$$S_t = A_t - D(\mathcal{P}(A_t)). \quad (6)$$

Building on [Deaton and Laroque \(1992\)](#), [Cafiero et al. \(2011b\)](#) prove that for this model there is a unique stationary rational expectations equilibrium \mathcal{P} in the class of continuous non-increasing functions.³ If we define $P^* \equiv \beta E \mathcal{P}(\varepsilon) - k$, the cutoff price for no storage, the price function has the following properties:

$$\mathcal{P}(A) = D^{-1}(A), \text{ for } A \leq D(P^*), \quad (7)$$

$$\mathcal{P}(A) > D^{-1}(A), \text{ for } A > D(P^*). \quad (8)$$

So P^* , which depends on the price function, defines the threshold between the two regimes. Prices above P^* are too high to make storage profitable, while for prices below P^* some stocks are carried over.

2.2 Numerical method

There is no closed-form solution for the equilibrium price function, which has to be approximated numerically. The numerical method follows the fixed-point approach proposed by [Deaton and Laroque \(1992\)](#). The equilibrium price function is approximated with a cubic spline over a grid of equally spaced availability points lying between -2 and 20 . The expectation term in equation (5) is replaced by a sum by discretizing the truncated normal distribution of the production shocks ε using a Gaussian quadrature calculated by the method of moments, with $N = 10$ nodes, where the production shocks and their associated probabilities are denoted ε^n and π^n . Then, using that production shocks are i.i.d., and combining (5) and (6) we have

$$\mathcal{P}(A) = \max \left[\beta \sum_{n=1}^N \pi^n \mathcal{P}(A - D(\mathcal{P}(A)) + \varepsilon^n) - k, D^{-1}(A) \right]. \quad (9)$$

The model is solved by iterating on this functional equation. Starting from a first guess for the price function, a price function applied on the right-hand side to all grid points leads by simple arithmetic operations to new values of the price function at the grid points on the left-hand side. The iterations stop when the Euclidean distance between two consecutive price functions at the interpolating nodes falls below a given tolerance threshold which we set to 10 decimal places.

[Cafiero et al. \(2011b\)](#) show that the estimation procedure is very sensitive to the accuracy of the model's numerical solution, which is determined mainly by the number of grid points used to approximate the policy function $\mathcal{P}(A)$. These authors show that [Deaton and Laroque \(1996\)](#) approximate the policy function on a grid that is too sparse to locate accurately the kink at the cutoff price P^* of empty stocks, which partly explains the inability of the storage model to generate the high serial correlation. Using a finer grid of 1,000 nodes, [Cafiero et al. \(2011b\)](#) obtain estimations of the parameters for which the storage model induces higher price autocorrelations. We retain their findings and choose a grid of 1,000 points.

³[Cafiero et al. \(2015\)](#) extend the proof to a model with free disposal and with a production support that may be unbounded. Free disposal has the advantage to prevent the realization of negative equilibrium prices, but increases significantly the time required to solve the model numerically preventing us from implementing it in this paper.

2.3 How can storage generate high serial correlation?

The debate over the empirical relevance of the storage model revolves around its ability to generate the high serial correlation observed in the data. Here, we explore the combination of parameters that allows the model to generate high serial correlation. Our storage model has six parameters, $\{a, b, k, r, \mu, \sigma\}$. In the remainder of this paper, we follow Deaton and Laroque (1996) and Cafiero et al. (2015) by fixing r at 5%. Deaton and Laroque (1996, Proposition 1) prove that it is not possible to identify separately the demand function and the distribution of supply shocks. So in this section, we set the mean and standard deviation of the harvest at 1 and 0.05.⁴ The mean price over the model asymptotic distribution is set to 1, which implies $a + b = 1$. Two degrees of freedom remain: storage cost and demand elasticity. We vary them to see how this affects the serial correlation. Given our assumptions, k can be interpreted as the ratio of storage costs with respect to the mean price, and the demand elasticity calculated at the mean price is simply equal to $1/b$.

To analyze the effect of storage cost on serial correlation, we set demand elasticity at -0.05 , corresponding to Roberts and Schlenker's (2013) best estimate of the elasticity of a caloric aggregate of the major crops. We vary storage costs between 0 and 20 percent of the mean price, and for every value of storage cost we solve the model and simulate it. We calculate the first-order autocorrelation for 100,000 series of 100 periods on the asymptotic distribution. As noted by Cafiero et al. (2011b), simulating the storage model generates time series with very volatile moments when the series length is around the number of observable annual prices (close to one hundred years). Therefore, it is not sufficient to compare the serial correlation of observable price to the average simulated first-order autocorrelation, we need also to compare it to the quantiles of the distribution of simulated first-order autocorrelation. The left panel of figure 1 displays the 5th, 50th, and 95th percentiles of the distribution of simulated first-order autocorrelation when we vary the storage cost. Serial correlation is a monotonically decreasing function of storage cost. This can be explained by the fact that the storage model displays two regimes. In one regime, there are positive stocks and prices are serially correlated. In the other, stocks are zero and prices are not serially correlated. The more time that is spent in the stockout regime, the lower will be the overall serial correlation generated by the model. Decreasing the storage cost makes storage more profitable, increases stock levels, thereby decreasing the likelihood of a stockout and increasing the serial correlation. With this calibration, even for a zero storage cost, the median first-order autocorrelation is well below the very high correlation observed in the price series (above 0.82 for all commodities except sugar). Even the 95th percentile is below 0.8.

This failure of the storage model to induce sufficient serial correlation in prices calls for a parameterization that is even more favorable to storage. This can be achieved by rotating the slope of the demand function around its mean. Indeed, in absence of inventories to buffer against short supply, the price adjustments are dictated only by the final demand for consumption. So the more inelastic the demand, the steeper the variations in prices and the greater the incentive to store. We set the storage cost at its zero lower bound and vary the demand elasticity between -0.4 and -0.005 (right panel of figure 1). Only for a very inelastic demand curve is the median of simulated first-order autocorrelation close to 0.8. The 95th percentile can be compatible with a first-order autocorrelation of 0.8 for a demand elasticity above -0.037 .

⁴A coefficient of variation of 5% for supply shocks is between what is observed for the commodities considered in this paper (see table A3 in the online appendix).




Figure 1: First-order autocorrelation implied by the storage model over 100 periods for several values of storage cost (with demand elasticity set at -0.05) and demand elasticity (with storage cost set at 0)

In the storage model with i.i.d. supply shocks, stockpiling is the sole source of time-dependency in prices, so only a model parameterization in which storage arbitrage is often active can generate high serial correlations. The model can generate high serial correlation only by decreasing the occurrence of stockouts which requires a parameterization of very low storage cost and very inelastic demand. So in the estimations that follow, we should expect that a storage model able to replicate the characteristics of the raw price series will be characterized by low storage costs and inelastic demand functions. Even with this combination, high first-order correlation is achieved only by the high percentiles of the asymptotic distribution.

3 Econometric procedure

In this section, following [Canova \(2014\)](#), we propose an econometric procedure to estimate the storage model and the trend in prices jointly. The idea behind this procedure is to capture in the trend the component of prices that cannot be accounted for by the storage model, in particular the non-cyclical fluctuations. As a result, the storage model has to account only for cyclical fluctuations in the observed data. We assume that observed prices, P_t^{obs} , can be decomposed into a multiplicative trend $\exp[\Gamma(t, \theta^\Gamma)]$ and a cyclical component denoted P_t^{sto} to be explained by the storage model:

$$P_t^{\text{obs}} = e^{\Gamma(t, \theta^\Gamma)} P_t^{\text{sto}}. \quad (10)$$

The vector of the parameters to be estimated θ can be split into two groups: the trend parameters θ^Γ , and the structural parameters of the storage model, θ^{sto} . In addition to the baseline case where any trend is ignored we consider three deterministic time trend specifications. In none of the trend specifications do we introduce an intercept because it would not be possible to identify separately the intercept of the trend from the intercept of the inverse demand function since both would be determined by the mean level of observed prices.

3.1 Trend specifications

No trend Our benchmark situation is where observed prices are assumed to be without trend. In this case, $\Gamma(t, \theta^\Gamma) = 0$, $P_t^{\text{obs}} = P_t^{\text{sto}}$, and θ^Γ is empty.

Linear trend Here we assume the trend is a deterministic linear time trend:

$$\Gamma(t, \theta^\Gamma) = g_1 t. \quad (11)$$

In this case $\theta^\Gamma = \{g_1\}$. For numerical stability, the time variable, t , is taken to vary between -1 and 1 .

Restricted cubic splines While the linear trend allows us to capture the overall long-run trend, it may not capture all the non-cyclical fluctuations that the storage model is unable to explain. It is often considered that trends in commodity prices might be non-constant (a feature captured, e.g., in [Arezki et al., 2014](#), by a piecewise linear trend with structural breaks). For a more flexible trend than in the linear case, we use restricted cubic splines. Cubic splines are piecewise cubic polynomials with continuous first and second derivatives. “Restricted” splines are splines that are constrained to be linear beyond the boundary knots which avoids a poor behavior in the tails, a feature common to polynomial trends. A restricted cubic spline with three knots has two parameters. With four knots, it has three parameters. So restricted cubic splines with three and four knots have the same degrees of freedom as quadratic and cubic polynomials but tend to be slightly more flexible. A spline with two knots would be the same as the linear trend above. Restricted cubic splines with three to five knots are also used in [Roberts and Schlenker \(2013\)](#) to capture trends in prices and quantities of agricultural commodities.

When represented by restricted cubic splines, the trend is expressed as

$$\Gamma(t, \theta^\Gamma) = \sum_{i=1}^I g_i B_i(t), \quad (12)$$

where I and $B_i(\cdot)$ are the degree of freedom and the basis functions of the spline,⁵ and g_i are the trend parameters to be estimated. The $B_i(\cdot)$ are functions of the knots, but once the knots are fixed the trend is linear in its parameters. Following the heuristics proposed in [Harrell \(2001\)](#), the knots for the cubic spline with three knots are located at the 10th, 50th, and 90th quantiles of the covariate, which correspond in our 1900–2011 sample to the years 1911, 1956, and 2000. The spline with four knots uses as knots the 5th, 35th, 65th, and 95th quantiles (1905, 1939, 1970, and 2006).

Since the knots are fixed before the estimation, only the slope parameters have to be estimated: $\theta^\Gamma = \{g_i\}_{i=1}^I$.

3.2 The likelihood estimator

Given that the price function \mathcal{P} is monotone decreasing ([Deaton and Laroque, 1992](#), Theorem 1, and [Cafiero et al., 2011b](#)), we can invert it to obtain the amount on hand from the price. Using the inverse of the price function, the cyclical component of prices, P^{sto} , follows a first-order Markov process with the

⁵For numerical stability, the splines are expressed in B-spline form and their basis matrices come from the command `ns` in the R package `splines`.

transition equation defined by equations (3) and (6) as

$$P_t^{\text{sto}} = \mathcal{P} \left(\mathcal{P}^{-1} (P_{t-1}^{\text{sto}}) - D (P_{t-1}^{\text{sto}}) + \varepsilon_t \right). \quad (13)$$

It is possible to link P_t^{sto} to observed prices, P_t^{obs} , using equation (10). So, given the price function \mathcal{P} , equations (10) and (13) define a mapping from the supply shocks ε_t to P_t^{obs} , conditional on P_{t-1}^{obs} and t .

Given a set of model parameters θ and a sample of observed prices of length T , noted $P_{1:T}^{\text{obs}} \equiv \{P_1^{\text{obs}}, \dots, P_T^{\text{obs}}\}$, and using the Markov structure of the problem, the likelihood function can be expressed as

$$L(\theta; P_{1:T}^{\text{obs}}) = f(P_1^{\text{obs}}; \theta) \prod_{t=2}^T f(P_t^{\text{obs}} | P_{t-1}^{\text{obs}}; \theta). \quad (14)$$

Using the mapping between observables and shocks, [Miranda and Rui \(1999\)](#) and [Cafiero et al. \(2015\)](#) obtain the conditional density $f(P_t^{\text{obs}} | P_{t-1}^{\text{obs}}; \theta)$ from the variable transformation method. Identification of the parameters of the demand function requires the parameters of the distribution of shocks to be set to arbitrary values. From here, the distribution of ε is assumed to be the unit normal distribution truncated at five standard deviations, with probability density function $f_\varepsilon(\varepsilon) = \phi(\varepsilon) / [\Phi(5) - \Phi(-5)]$ for $\varepsilon \in [-5, 5]$ and $f_\varepsilon(\varepsilon) = 0$ otherwise. We can write the conditional density of P_t^{obs} as:

$$f(P_t^{\text{obs}} | P_{t-1}^{\text{obs}}; \theta) = f_\varepsilon(\varepsilon_t) |J_t|, \quad (15)$$

where $|J_t|$ is the determinant of the Jacobian of the mapping $P_t^{\text{obs}} \mapsto \varepsilon_t$.

Based on equations (10) and (13), this mapping is

$$\begin{aligned} \varepsilon_t &= \mathcal{P}^{-1} (P_t^{\text{sto}}) - [\mathcal{P}^{-1} (P_{t-1}^{\text{sto}}) - D (P_{t-1}^{\text{sto}})], \\ &= \mathcal{P}^{-1} \left(e^{-\Gamma(t, \theta^\Gamma)} P_t^{\text{obs}} \right) - \left[\mathcal{P}^{-1} \left(e^{-\Gamma(t-1, \theta^\Gamma)} P_{t-1}^{\text{obs}} \right) - D \left(e^{-\Gamma(t-1, \theta^\Gamma)} P_{t-1}^{\text{obs}} \right) \right], \end{aligned} \quad (16)$$

which gives the expression of J_t :

$$J_t = e^{-\Gamma(t, \theta^\Gamma)} \mathcal{P}^{-1'} \left(e^{-\Gamma(t, \theta^\Gamma)} P_t^{\text{obs}} \right). \quad (17)$$

The probability of any element of $P_{1:T}^{\text{sto}}$ being equal to P^* is zero, so the derivative of \mathcal{P}^{-1} exists almost everywhere.

In this paper we extend the Conditional Maximum Likelihood Estimator pioneered by [Cafiero et al. \(2015\)](#) to its unconditional counterpart. The aim is to use all the available information from the first observation by accounting for the marginal density $f(P_1^{\text{obs}}; \theta)$ in equation (14). The marginal density $f(P_1^{\text{obs}}; \theta)$ can be expressed as the following integral over the steady-state distribution:

$$f(P_1^{\text{obs}}; \theta) = \int_{P_0} f(P_1^{\text{obs}} | P_0; \theta) f(P_0; \theta) dP_0. \quad (18)$$

This integral is intractable, because there is no closed-form solution for the steady-state distribution of the storage model. However, we can draw from the distribution with density $f(P_0; \theta)$, which is the unconditional probability density function of price in the storage model. Therefore, we can use Monte

Carlo integration to estimate $f(P_1^{\text{obs}}; \theta)$ by simulating the model on the stationary distribution:

$$f(P_1^{\text{obs}}; \theta) \approx M^{-1} \sum_{m=1}^M f(P_1^{\text{obs}} | P_0^{(m)}; \theta), \quad (19)$$

where $m = \{1, \dots, M\}$ indexes random draws from the unconditional price distribution. We set the number of draws M to 1,000,000 and obtain them by simulating 10,000 trajectories starting from the steady state for 120 periods and discarding the first 20 periods as burn-in periods. The random production shocks that generate the price simulations are drawn at the beginning of the estimation procedure and remain fixed throughout. Because of this simulation step, the Unconditional Maximum Likelihood Estimator falls within the class of simulated estimators.

As is well known in time-series econometrics (Hamilton, 1994, Ch. 5), if the sample size is sufficiently large the contribution of the first observation to the likelihood is negligible, while it is often much more complex to calculate the unconditional likelihood than the conditional likelihood. In the case of the storage model, it is true that the simulations necessary to calculate the marginal density make the likelihood evaluation much more costly. In Monte Carlo experiments designed following Michaelides and Ng (2000) and Cafiero et al. (2015), the unconditional likelihood performs only marginally better than the conditional likelihood.⁶ However, when using an actual sample, the unconditional likelihood has benefits which in our view outweigh its costs. For observed prices the conditional likelihood presents many local optima. The unconditional likelihood helps to select an optimum with an unconditional price distribution not too far from the price sample which may not be the case for a conditional likelihood. Indeed, in order to fit the high serial correlation of the data, the Conditional Maximum Likelihood Estimator often leads to parameter estimates for which the availabilities corresponding to observed prices (calculated using \mathcal{P}^{-1}) are set at very high levels which would correspond to very high stock levels. Observing large stock levels may have a high probability conditional on having large stocks in the previous period, but the unconditional probability of such a situation is very low. So, the Unconditional Maximum Likelihood Estimator helps to filter out some of these situations.

Based on all the previous elements, we can write the log-likelihood as

$$\begin{aligned} \log L(\theta; P_{1:T}^{\text{obs}}) = & -\frac{T}{2} \log 2\pi - T \log [\Phi(5) - \Phi(-5)] + \sum_{t=1}^T \left[-\Gamma(t, \theta^\Gamma) + \log \left| \mathcal{P}^{-1'} \left(e^{-\Gamma(t, \theta^\Gamma)} P_t^{\text{obs}} \right) \right| \right] \\ & - \sum_{t=2}^T \left(1_{|\varepsilon_t| \leq 5} \cdot \varepsilon_t^2 + 1_{|\varepsilon_t| > 5} \cdot \infty \right) / 2 + \log \left(M^{-1} \sum_{m=1}^M 1_{|\varepsilon^{(m)}| \leq 5} \cdot \exp \left(-\varepsilon^{(m)2} / 2 \right) \right), \quad (20) \end{aligned}$$

where 1 is the indicator function and $\varepsilon^{(m)} = \mathcal{P}^{-1} \left(\exp \left(-\Gamma(1, \theta^\Gamma) \right) P_1^{\text{obs}} \right) - S_0^{(m)}$. Given that the interest rate and the parameters of the distribution of production shocks have been fixed, there are three parameters that we need to estimate for the storage model $\theta^{\text{sto}} = \{a, b, k\}$, in addition to the parameters characterizing the trend, θ^Γ , defined above. From a set of parameters θ provided by the optimization algorithm, we calculate the detrended price P^{sto} and solve for the policy function of the storage model $\mathcal{P}(\cdot)$. Using this policy function, we can simulate the model to calculate the marginal probability and evaluate the likelihood for this set of parameters.

Evaluated on the observed prices, the above log-likelihood behaves badly. It displays many local

⁶Results available in section B of the online appendix.

optima. Gradient-based solvers and derivative-free local search methods converge only to local optima which are very sensitive to first guesses. Thus, we need to apply a global search algorithm to increase the likelihood of obtaining a global solution. Following a recent review of derivative-free algorithms (Rios and Sahinidis, 2013), and some tests on our problem, we choose a global solver, the particle swarm pattern search algorithm proposed by Vaz and Vicente (2007), and refine the solution it delivers with a local solver using a sequential quadratic programming approach (Nocedal and Wright, 2006, Ch. 18). Parameters are constrained to remain within bounds (this is required by the global solver). b is constrained to be strictly negative (≤ -0.001) and k to be positive or null. All the other bounds are chosen to be low enough or high enough to avoid their being binding. The particle swarm solver is initialized with 700 vectors of first-guess parameters, a combination of educated guesses, random draws, and vectors of previous solutions (e.g., the estimation without trend serves as a first guess for the linear trend). The tolerance for both optimization solvers is fixed at 10^{-6} .

Once a maximum has been identified, we estimate the asymptotic variance-covariance matrix of the parameters as the inverse of the outer product of the scores. If the highest log-likelihood is obtained with k constrained at its zero lower bound, calculating the variance-covariance matrix using the scores is inappropriate, and other methods such as bootstrap should be used. However, the number of estimations run in this paper prevents us from using the bootstrap method. For the estimates with k at zero, we do not report the standard errors for k but report the standard errors of the other parameters obtained using the inverse of the outer product of the scores by maintaining k at zero.

4 Estimation results

4.1 Data

Our data set is composed of the thirteen commodities analyzed in Deaton and Laroque (1992) (banana, cocoa, coffee, copper, cotton, jute, maize, palm oil, rice, sugar, tea, tin, and wheat). The original price series from Grilli and Yang (1988) is extended by Pfaffenzeller et al. (2007) and cover the period 1900 to 2011. The data were downloaded from Stephan Pfaffenzeller's personal website.⁷ The data are annual price indexes calculated by averaging the monthly price data provided by the World Bank Development Prospects Groups over the calendar year, and normalizing them with respect to the 1977–79 mean price. We deflated the nominal price indexes by the US CPI.

We use only price data to estimate the storage model but we also rely on shorter series of production data to illustrate the consequence of our estimations in terms of demand elasticities. The production data are from the FAOSTAT food balance sheets in the case of the agricultural commodities,⁸ from the British Geological Survey in the case of tin,⁹ and from the 2014 World Copper Factbook of the International Copper Study Group for copper.¹⁰ They cover the period 1961 to 2011. For each commodity, the logarithm of production is detrended by modeling the trend by a restricted cubic spline with five knots (as in Roberts and Schlenker, 2013).¹¹

⁷<http://www.stephan-pfaffenzeller.com>

⁸<http://faostat.fao.org/>

⁹<http://www.bgs.ac.uk/mineralsuk/statistics/home.html>

¹⁰<http://www.icsg.org/index.php/statistics/selected-data>

¹¹The knots are located at 1964, 1975, 1986, 1998, and 2008 as suggested by Harrell (2001).

4.2 Model selection

Joint estimation of the structural and the trend parameters allows us to select the preferred trend specification using model selection criteria. The various trends nest the specification without trend; however, the spline with four knots does not nest the spline with three knots, because the splines do not have same boundary knots. Thus, it is not possible to select the trend using a likelihood ratio test.¹² We select the preferred trend specification using the Akaike Information Criterion (AIC). The results are reported in table 1 which presents the values for the preferred models in boldface. For the model without trend the AIC is given in level, while for the three models with trend they are given in ratios to the AIC of the model without trend so that a value above unity means a lower AIC than the model without trend. The model without trend is never retained for any of the commodities. The model with a linear trend is preferred for copper, palm oil, and wheat. The model with a three-knot spline trend (RCS3) is preferred for coffee, cotton, jute, rice, and sugar. The model with a four-knot spline trend (RCS4) is preferred for banana, cocoa, maize, tea, and tin.

Table 1: Model Selection Using the Akaike Information Criterion

Commodity	No trend	Linear	RCS3	RCS4
Banana	-309.823	1.030	1.069	1.070
Cocoa	-417.014	0.998	1.010	1.025
Coffee	-377.997	0.995	1.005	1.000
Copper	-236.051	1.000	0.993	0.997
Cotton	-231.181	0.992	1.021	1.011
Jute	-184.036	1.015	1.056	1.045
Maize	-159.578	0.998	1.002	1.038
Palm oil	-219.800	1.042	1.040	1.034
Rice	-219.357	1.105	1.114	1.104
Sugar	-93.022	1.189	1.198	1.179
Tea	-275.557	1.011	1.010	1.014
Tin	-406.411	0.996	1.006	1.006
Wheat	-189.445	1.054	1.046	1.054

Note: For the “No trend” column, the AIC are given in levels while for the other columns they are reported in ratios to the “No trend” column so that a value above unity means a lower AIC than the model without trend. The preferred model for each commodity is in boldface.

4.3 No trend

The parameter estimates for the model without trend are given in table 2. In this setting without stock deterioration, estimates of a in the first column are directly interpretable as the ergodic means of the models. Thus, it would be reasonable to expect a to be not too far from the observed sample means reported in table 3, even if the sample mean is not the maximum likelihood estimator of a . However, this is not the case. With the exception of sugar, a is systematically higher than the sample mean, and for some commodities by a large margin. For example, a exceeds the sample mean by 255 percent for banana, 227 for cotton, 122 for rice, or 80 for wheat. This “bias” could be related to the presence of a

¹²Even if the models were all nested, a likelihood ratio test could not be applied to all commodities, because some estimations of k are constrained on their lower bound.

trend in the observables. A trend would generate a serial correlation higher than expected from storage alone. Estimating the ergodic mean of the model at above the sample mean implies that the sample is located in a region of larger than normal availabilities, and with large availabilities there are important stocks, and prices are more positively correlated than if availabilities are close to normal.

Table 2: Parameter Estimates Without Trend

Commodity	a	b	k	$\log L$	P^*	# Stockouts
Banana	1.9078 (0.6159)	-3.8117 (2.3678)	0.0011 (0.0037)	157.9117	3.6291	0
Cocoa	0.1826 (0.0383)	-0.8604 (0.0930)	0.0002 (0.0009)	211.5072	0.6149	0
Coffee	0.2571 (0.0259)	-0.8959 (0.0656)	0.0015 (0.0009)	191.9984	0.6876	0
Copper	0.6339 (0.0540)	-1.1740 (0.1382)	0.0046 (0.0034)	121.0254	1.1461	0
Cotton	1.7832 (0.7019)	-6.1122 (3.9400)	0.0029 (0.0027)	118.5906	4.7461	0
Jute	0.6934 (0.0970)	-1.7440 (0.2308)	0.0050 (0.0049)	95.0180	1.4927	0
Maize	0.7058 (0.0816)	-2.3422 (0.1738)	0.0009 (0.0030)	82.7891	1.8374	0
Palm oil	0.7311 (0.0802)	-1.5067 (0.1708)	0.0097 (0.0046)	112.8999	1.3887	1
Rice	1.1766 (0.3604)	-5.2418 (2.2061)	0 -	112.6784	3.8021	0
Sugar	0.5362 (0.0521)	-1.9907 (0.1123)	0.0050 (0.0026)	49.5109	1.4926	3
Tea	1.0839 (0.4672)	-3.6411 (2.2116)	0 -	140.7786	2.8518	0
Tin	0.3606 (0.0509)	-1.0371 (0.2022)	0.0023 (0.0015)	206.2057	0.8462	0
Wheat	1.0926 (0.1819)	-3.9100 (1.0561)	0 -	97.7225	3.0058	0

Note: Asymptotic standard errors in parentheses.

The limited number of stockouts confirms that the estimations localize all the samples in regions with large availability. With the exception of palm oil and sugar for which the respective number of periods without stocks over the sample are 1 and 3, commodity prices are always under their respective cutoff price P^* implying that inventories were carried over the whole 112 years of the sample. This feature is present also in the estimations in [Cafiero et al. \(2011b\)](#), where only sugar displays stockouts. In the model, stockouts occur when prices exceed P^* . Our estimations show that on average over all commodities, P^* exceeds its corresponding sample means by 4.5 times, making stockouts unlikely.

Figure 2 illustrates this in the case of wheat. Most observed prices are below the ergodic mean of the model, and all are far from P^* . If we compare the localization of the observed prices to the ergodic distribution of availability, we see that the observations are concentrated to the right of the distribution mode. The observations look more like an extreme sample in which the level of availability, and thus the

Table 3: Comparisons of Data Features and Predictions of the Model Without Trend

Commodity	Mean	One-year a-c	Two-year a-c	Coefficient of variation	Skewness	Excess kurtosis
Banana	0.54***	0.95***	0.90***	0.23***	-0.27***	-0.77***
Cocoa	0.17	0.86	0.71	0.60	1.24	1.64
Coffee	0.20	0.84	0.68	0.50	1.61	3.89
Copper	0.47•	0.83*	0.64	0.40	0.90*	0.56*
Cotton	0.55**	0.94***	0.85**	0.51	0.20***	-0.61**
Jute	0.52	0.84•	0.70	0.44•	0.41**	-0.23**
Maize	0.61	0.86	0.73	0.51	0.84•	1.22
Palm oil	0.46**	0.82*	0.65	0.60	2.43	11.84
Rice	0.53•	0.91•	0.79	0.50	0.42**	-0.39**
Sugar	0.61	0.70	0.51	0.69	1.62	3.45
Tea	0.45**	0.90*	0.81*	0.38*	-0.02***	-0.80***
Tin	0.20*	0.90**	0.78*	0.47•	1.48	2.84
Wheat	0.61•	0.91*	0.79•	0.47	0.82•	0.39•

Notes: Moments calculated on the observed prices P^{obs} . ***, **, *, and • indicate if the moments of observed prices are outside the 1%, 5%, 10%, and 20% two-sided percentiles of the simulated moments.

stock level are always high, than a typical sample from the distribution.

The small number of stockouts, and localization of the samples in regions with large availabilities question the empirical relevance of the storage model. This model is supposed to alternate between two regimes with relatively stable prices when there are stocks, and spikes during stockouts. If there are no stockouts, price spikes are explained by adverse production shocks only. They do not correspond to a much steeper part of the demand function and should be as likely as price troughs, which is not consistent with the stylized facts. Most commodity prices present a positive skewness (table 3) and there are few downward spikes to match the upward spikes (Deaton and Laroque, 1992).

Section 2.3 showed that the storage model can generate high serial correlation only with parameterization of very low storage cost and very inelastic demand. We next examine the magnitude of the estimated storage costs and demand elasticities. k/a is the ratio of storage cost to the ergodic mean price, so it is unit-free and directly interpretable. Because we assumed that supply shocks follow a truncated unit normal, it is not possible to recover the demand elasticity only from the demand function. To calculate the demand elasticity, we use Deaton and Laroque's (1996) Proposition 1 which shows that re-scaling the distribution of supply shocks to have mean and standard deviation μ and σ while adjusting the inverse demand function to $\tilde{D}^{-1}(C_t) = (a - b\mu/\sigma) + bC_t/\sigma$ does not affect the distribution of prices. Using this adjusted demand function, the demand elasticity evaluated at the model's ergodic mean price is given by $a\sigma/(b\mu)$.¹³ The coefficient of variation of the supply shocks, σ/μ , is obtained by calculating the standard deviation of the detrended logarithm of production, and is provided in the online appendix.¹⁴

¹³A similar approach is used in Guerra et al. (forthcoming).

¹⁴While reformulating the estimates as unit-free parameters is necessary to compare them between trend specifications, the expression as a demand elasticity relies on a literal interpretation of the storage model. In this specification of the storage model, additive demand shocks are equivalent to additive negative supply shocks, and thus the supply shocks in the model should be interpreted as net-supply shocks. If we assume that demand and supply shocks are uncorrelated, the elasticities calculated using only the coefficient of variation of supply will be downward biased. For staple food markets, where price volatility is often believed to arise from supply shocks the bias may be small, but it may be larger for commodities more subject to demand shocks such as metals.




Figure 2: Estimated storage model for wheat without trend. Lower panel: price function, $\mathcal{P}(A)$, with observed prices on it noted $+$. Upper panel: ergodic distribution of availability.

The estimated storage costs and price elasticities of demand are consistent with the results in section 2.3: both are low (table 4). The estimated storage costs are below 1 percent of the ergodic price for all commodities except palm oil, and are null for three commodities: rice, tea, and wheat. Information on storage costs to which we could compare these estimates are not readily available for all commodities. However, in a study of the grain chain in Middle East and North African countries, [World Bank and FAO \(2012, figure 2-4\)](#) report that the storage cost of wheat in the US was equal to US\$ 24.24 per ton in 2009, which would represent around 10% of the recent price of wheat. So, at least for cereals, the estimates of storage costs appear to be low.

In the model without trend, our implied price elasticities of consumption are comparable to those derived from other estimations of the storage model (see table 4 in [Guerra et al., forthcoming](#)). If we focus on the cereals, although these elasticities are plausible, in absolute value they appear to be in the low range of the elasticities in the literature (e.g., [Seale and Regmi, 2006](#); [Adjemian and Smith, 2012](#); [Roberts and Schlenker, 2013](#)). We would expect smaller elasticities than in most of the literature because most estimation methods do not account for the presence of stocks which tend to create a positive bias. Nevertheless, even [Roberts and Schlenker \(2013\)](#) who control for the effect of storage using instrumental variables find higher elasticities of demand, between -0.066 and -0.028 for aggregate calories from maize, rice, soybeans and wheat, and even higher values for each commodity individually.

Table 4: Estimated Values of Storage Costs and Demand Elasticities

Commodity	Storage costs: $100 \cdot k/a$				Price elasticity of demand: $a\sigma/(b\mu)$			
	No trend	Linear	RCS3	RCS4	No trend	Linear	RCS3	RCS4
Banana	0.06 (0.19)	0 –	0 –	0 –	–0.018 (0.013)	–0.033 (0.004)	–0.029 (0.009)	– 0.029 (0.008)
Cocoa	0.11 (0.51)	0.19 (0.54)	0.05 (0.23)	0.53 (0.63)	–0.015 (0.004)	–0.015 (0.004)	–0.015 (0.010)	– 0.027 (0.009)
Coffee	0.59 (0.35)	0.59 (0.35)	1.41 (1.01)	1.41 (1.07)	–0.022 (0.003)	–0.022 (0.003)	– 0.037 (0.010)	–0.037 (0.010)
Copper	0.72 (0.54)	1.09 (0.67)	0.99 (0.70)	0.07 (0.59)	–0.024 (0.004)	– 0.028 (0.004)	–0.028 (0.005)	–0.019 (0.005)
Cotton	0.16 (0.17)	0.19 (0.43)	0.34 (0.43)	0.34 (0.50)	–0.021 (0.016)	–0.022 (0.038)	– 0.029 (0.015)	–0.029 (0.020)
Jute	0.72 (0.71)	0.42 (0.96)	2.35 (0.91)	2.34 (0.99)	–0.045 (0.009)	–0.045 (0.008)	– 0.104 (0.009)	–0.104 (0.015)
Maize	0.13 (0.43)	1.07 (0.83)	1.34 (0.80)	3.25 (0.75)	–0.018 (0.002)	–0.030 (0.004)	–0.036 (0.005)	– 0.062 (0.005)
Palm oil	1.33 (0.64)	1.20 (0.91)	1.40 (0.88)	1.43 (0.98)	–0.023 (0.004)	– 0.027 (0.004)	–0.028 (0.004)	–0.028 (0.008)
Rice	0 –	0.28 (0.51)	0.30 (0.54)	0.33 (0.55)	–0.006 (0.003)	–0.014 (0.002)	– 0.016 (0.004)	–0.016 (0.005)
Sugar	0.94 (0.49)	2.27 (0.81)	2.43 (1.03)	2.35 (0.99)	–0.010 (0.001)	–0.020 (0.002)	– 0.019 (0.002)	–0.019 (0.003)
Tea	0 –	0.57 (0.36)	0.95 (0.51)	0.89 (0.46)	–0.006 (0.005)	–0.011 (0.002)	–0.014 (0.004)	– 0.014 (0.005)
Tin	0.63 (0.43)	0.26 (0.21)	0 –	0 –	–0.019 (0.005)	–0.015 (0.005)	–0.014 (0.003)	– 0.017 (0.002)
Wheat	0 –	1.11 (0.77)	1.18 (0.80)	0 –	–0.012 (0.004)	– 0.033 (0.005)	–0.034 (0.005)	–0.025 (0.005)

Notes: The preferred model values are in boldface and in parentheses the asymptotic standard errors obtained using the delta method.

4.4 Models with a time trend

The parameter estimates for the model with a linear trend are given in table 5. Since for numerical stability the time variable has been defined over the interval $[-1, 1]$, the trend coefficient g_1 is not directly interpretable. To make it interpretable, we report it also as annual growth rates in column G , where the standard deviation is calculated using the Delta method. For all commodities except tin, the annual growth rate is negative, which is consistent the Prebisch-Singer hypothesis of a long-run downward trend in commodity prices. If we exclude coffee and tin, the values range from an annual decline of 0.4 percent for copper to 1.94 percent for rice. The significance of many trend coefficients indicates that the model without trend is likely to be misspecified.

For the other trend specifications, since the trend and the structural parameters are not directly interpretable, the complete results are not displayed here. They are provided in the online appendix, along with a figure plotting the various trends with real prices. Rather than presenting detailed results, table 4 presents the parameters estimates expressed in a way that makes them comparable across trend specifications. The presence of a deterministic time trend in the model estimation can lead to large effects in terms of the key parameter estimates. Storage costs tend to increase when there is a trend, and also price elasticity of demand in absolute value. With the preferred model (in boldface) ten out of the thirteen

Table 5: Parameter Estimates with a Linear Trend

Commodity	G	g_1	a	b	k	$\log L$	# Stockouts
Banana	−0.0161 (0.0010)	−0.8909 (0.0550)	0.8770 (0.0434)	−0.9756 (0.1114)	0 −	163.4962	0
Cocoa	−0.0058 (0.0034)	−0.3215 (0.1899)	0.2036 (0.0501)	−0.9771 (0.1370)	0.0004 (0.0011)	212.1482	0
Coffee	−0.0000 (0.0022)	−0.0008 (0.1225)	0.2571 (0.0261)	−0.8963 (0.0747)	0.0015 (0.0009)	191.9992	0
Copper	−0.0040 (0.0009)	−0.2218 (0.0505)	0.6227 (0.0515)	−0.9837 (0.1082)	0.0068 (0.0042)	122.0631	0
Cotton	−0.0045 (0.0148)	−0.2476 (0.8197)	1.5846 (1.4137)	−5.3011 (8.1716)	0.0031 (0.0062)	118.6970	0
Jute	−0.0108 (0.0014)	−0.6014 (0.0784)	0.5357 (0.0856)	−1.3389 (0.1182)	0.0023 (0.0051)	97.4204	0
Maize	−0.0097 (0.0012)	−0.5386 (0.0671)	0.6957 (0.0694)	−1.3722 (0.1400)	0.0074 (0.0057)	83.6162	0
Palm oil	−0.0179 (0.0014)	−0.9921 (0.0763)	0.4775 (0.0491)	−0.8150 (0.0743)	0.0057 (0.0043)	118.5360	1
Rice	−0.0194 (0.0020)	−1.0755 (0.1119)	0.6286 (0.0647)	−1.2003 (0.1323)	0.0017 (0.0032)	125.2477	1
Sugar	−0.0107 (0.0005)	−0.5955 (0.0282)	0.5811 (0.0415)	−1.0836 (0.0725)	0.0132 (0.0046)	59.2928	5
Tea	−0.0134 (0.0034)	−0.7430 (0.1904)	0.7637 (0.0623)	−1.4313 (0.2095)	0.0043 (0.0027)	143.2308	0
Tin	0.0143 (0.0046)	0.7932 (0.2552)	0.4341 (0.0709)	−1.5516 (0.4729)	0.0011 (0.0009)	206.4246	0
Wheat	−0.0127 (0.0011)	−0.7033 (0.0600)	0.7133 (0.0531)	−0.9266 (0.1104)	0.0079 (0.0055)	103.8497	1

Notes: Asymptotic standard errors in parentheses. Column G is a transformation into an annual growth rate of the trend parameter g_1 .

commodities present higher storage costs than if there is no trend. With the exception of tin, all price elasticities are higher with the preferred model. The differences between the elasticities estimated with the models without trend and with the preferred models with trend are important. The elasticities double in the case of cocoa, coffee, jute, rice, sugar, and tea, and increase three-fold for maize and wheat.

For cereals, the elasticities of the preferred model although still low appear to be more consistent with the literature. Similarly, the increase in storage costs for cereals leads to more plausible parameters which nevertheless are low compared to some published figures.

For sugar, we can compare the parameter estimates to the Conditional Maximum Likelihood estimates of [Cafero et al. \(2015\)](#). They estimate their model on data from 1921 to 2009, because of a possible structural break between 1920 and 1921. On this subsample, the stationarity of deflated prices is more likely to hold (see figure [A1](#) in the online appendix) and the first-order correlation at 0.63 is lower than the 0.70 obtained over our extended sample. The purpose of our joint estimation approach is precisely to accommodate for such possible breaks. The preferred model for sugar is the model with a three-knot spline trend. The estimated trend is decreasing at the beginning of the sample and roughly constant after

1930 (figure A1 in the online appendix). We express the parameters in ratios to make them comparable across specifications. For $100 \cdot k/a$, our estimation is 2.43 (table A4 in the online appendix) and theirs is 2.23 (Cafiero et al., 2015, Table 3, ML setting $d = 0$). For a/b , our estimation is -0.52 and theirs is -0.49 . The estimations are very close, while they are very different if we use our estimation without trend ($100 \cdot k/a = 0.94$ and $a/b = -0.27$). This could indicate that for sugar our strategy succeeds in removing a source of non-stationarity in the original series.

4.5 Model predictions

Since one of the main critiques of the storage model is its inability to reproduce the observed serial correlation, one way to assess our new estimates is to compare the model predictions to the features of the data. For samples of the same length of the observables, a storage model generates moments that are highly volatile, so comparing the ergodic moments and observed moments would be inappropriate. We adopt the approach in Cafiero et al. (2011b): we match the observed moments to their corresponding percentiles in the ergodic distribution of the estimated model. The model predictions are consistent with the data when the percentiles are neither too low nor too high. We calculate the mean, first- and second-order correlations, coefficient of variation, skewness, and kurtosis for the observations. For the model with trend, the moments are calculated on the cyclical component of prices, P^{sto} , which is the component that the storage model is supposed to explain. We calculate the corresponding percentiles from 1,000,000 series of 112 periods from the asymptotic distribution.

The moments and their localization with respect to simulated percentiles are given in table 3 for the model without trend and in table 6 for the preferred model with trend. All moments from the observations are located within the implied empirical distribution of the models without and with trend. The corresponding percentiles are used to build symmetric confidence regions. As already noted by Cafiero et al. (2011b), if the model without trend is solved and estimated with a sufficiently precise grid, it is able to generate first-order correlation similar to the observations for several commodities (table 3). The storage model even appear to fail more often for reproducing the skewness and kurtosis, with observed moments which frequently are outside the 99 and 95 percent confidence intervals. It should be noted that the banana and tea price series present negative skewness, which makes them nearly impossible to reproduce with a storage model which on average generates positive skewness. When a commodity presents negative excess kurtosis, which is the case of banana, cotton, jute, rice and tea, it seems that the estimated model has difficulty reproducing it. Deaton and Laroque (1992, section 2.1) note that the storage model can produce negative excess kurtosis but only from a calibration with low price volatility and a limited role for storage, which does not correspond to our estimations.

We turn next to the predictions of the preferred model with trend (table 6). Compared to the model without trend, there is an improvement in the model predictions: many moments move inside smaller confidence intervals. However, the predictions for some commodities – banana, jute, and tea – do not improve. The storage model with the specifications we estimated appears to be unable to match the moments for these commodities. Due to the deterministic trend, serial correlation decreases and becomes more consistent with the model predictions. Except for the three commodities for which the storage model does not work and for maize, all the observed first- and second-order correlations fall inside the symmetric 95 percent confidence intervals. Similarly, if we exclude banana, jute and tea, then with the exception of 3 moments that are outside the 90 percent confidence intervals, all other moments are

Table 6: Comparisons of Data Features and Predictions of the Preferred Models

Commodity	Mean	One-year a-c	Two-year a-c	Coefficient of variation	Skewness	Excess kurtosis
Banana	0.66***	0.89***	0.81***	0.41	1.66	3.25
Cocoa	0.43	0.82	0.65	0.61	1.76	3.00
Coffee	0.21	0.78	0.59	0.47	1.86	5.97
Copper	0.47*	0.80*	0.58	0.38 [•]	0.81*	0.32**
Cotton	0.79**	0.66	0.43	0.45	3.38	19.96
Jute	0.62	0.69 [•]	0.40	0.31**	0.62***	0.17***
Maize	0.85	0.69**	0.46	0.36	1.47	2.79
Palm oil	0.42	0.72	0.47	0.39 [•]	0.99*	1.38 [•]
Rice	0.97 [•]	0.76	0.46	0.38	2.17	7.07
Sugar	0.93	0.62	0.35	0.67	3.00	13.73
Tea	0.67***	0.91***	0.85***	0.46	0.46**	0.01**
Tin	0.11	0.80	0.65	0.58	1.63	3.27
Wheat	0.57*	0.75 [•]	0.45	0.31*	1.26 [•]	2.17 [•]

Notes: The preferred model is chosen according to the AIC in table 1. Moments calculated on the cyclical component of prices, P^{sto} . ***, **, *, and [•] indicate if the moments of P^{sto} are outside the 1%, 5%, 10%, and 20% two-sided percentiles of the simulated moments.

consistent with the predictions of the estimated storage models.

A disturbing feature of the estimations without trend is the very small number (often zero) of implied stockouts. Considering the possibility of a trend increases the number of implied stockouts, which becomes positive for many commodities (table 7). However, even if we exclude banana, jute and tea, three commodities – cocoa, copper, and cotton – present zero stockouts.

Table 7: Number of Implied Stockouts over the Sample Interval

Commodity	No trend	Linear	RCS3	RCS4
Banana	0	0	0	0
Cocoa	0	0	0	0
Coffee	0	0	1	2
Copper	0	0	0	4
Cotton	0	0	0	0
Jute	0	0	4	8
Maize	0	0	3	12
Palm oil	1	1	1	1
Rice	0	1	1	1
Sugar	3	5	8	16
Tea	0	0	0	0
Tin	0	0	0	6
Wheat	0	1	1	0

Note: The preferred model values in boldface.

It is interesting to compare the number of implied stockouts to the probabilities of exceeding a given number of stockouts. The probabilities of stockouts are calculated in the same way as the price moments from 1,000,000 series of 112 periods from the asymptotic distribution. The results for the model without

trend and for the preferred model with trend for at least 1, 2, 3, 5, and 10 stockouts are given in table 8. Models with trend have much higher probabilities of stockouts than the model without trend. If there is a trend, the storage costs and the price elasticities are higher than without trend which discourages storage and leads to higher probabilities of stockouts. Apart from cocoa, cotton and tin, with the preferred model, samples without stockouts are highly unlikely. Even samples with only one stockout are fairly unlikely. The parameterization implied by the estimation of a model with trend is much more favorable to occasional stockouts and so provides a natural explanation for price spikes as periods where stocks are exhausted.

Table 8: Probabilities of at Least n Stockouts, in Samples of the Same Size as the Data

Commodity	No-trend model					Preferred model				
	$n = 1$	$n = 2$	$n = 3$	$n = 5$	$n = 10$	$n = 1$	$n = 2$	$n = 3$	$n = 5$	$n = 10$
Banana	0.91	0.81	0.70	0.48	0.13	0.98	0.94	0.89	0.73	0.30
Cocoa	0.63	0.47	0.35	0.18	0.02	0.86	0.74	0.62	0.39	0.09
Coffee	0.77	0.63	0.49	0.28	0.05	0.95	0.88	0.79	0.59	0.19
Copper	0.95	0.88	0.79	0.59	0.19	0.98	0.94	0.88	0.72	0.29
Cotton	0.75	0.60	0.47	0.26	0.05	0.87	0.75	0.63	0.41	0.10
Jute	0.88	0.77	0.65	0.43	0.10	1.00	0.99	0.99	0.95	0.68
Maize	0.76	0.61	0.48	0.27	0.05	1.00	1.00	1.00	0.99	0.86
Palm oil	0.95	0.88	0.79	0.59	0.19	0.97	0.92	0.86	0.68	0.26
Rice	0.64	0.49	0.36	0.19	0.02	0.95	0.88	0.79	0.59	0.19
Sugar	0.77	0.63	0.49	0.28	0.05	0.98	0.94	0.88	0.73	0.30
Tea	0.74	0.60	0.46	0.26	0.04	0.98	0.94	0.88	0.73	0.30
Tin	0.84	0.71	0.59	0.36	0.08	0.75	0.60	0.46	0.26	0.04
Wheat	0.72	0.57	0.44	0.24	0.04	0.99	0.97	0.93	0.82	0.41

Note: The preferred model is chosen according to the AIC of table 1.

5 Conclusion

Estimating the competitive storage model on untransformed series of commodity prices leads to very low demand elasticities and storage costs, which results in a prediction of very infrequent, and often zero, stockouts over the sample period. These results may stem from the presence of trends in prices, which could create statistical features difficult to explain by a storage model.

This article proposes a strategy inspired by [Canova \(2014\)](#) to estimate jointly the structural parameters of a storage model and the parameters characterizing the non-cyclical component of prices for which the storage model cannot account. For the non-cyclical component of prices, three deterministic time trends with increasing flexibility were tested and compared with the baseline model which ignores the possibility of a trend.

Our results show that storage models with trend are always preferred to models without trend, and the significance of the trend parameters indicates that the model without trend is likely to be misspecified. Accounting for a trend is quantitatively important for estimating the structural parameters. For most commodities, the storage model with a deterministic trend yields more plausible estimates of the structural parameters (e.g., higher storage costs and demand elasticities). It also increases the probability of observing stockouts, and more closely replicates the most salient features of the price data, including the high serial correlation which led [Deaton and Laroque \(1996\)](#) to question the relevance of the storage

model. For most commodities our results support the empirical relevance of the speculative storage model which is in line with the recent findings in Cafiero et al. (2011b, 2015) and prove that the joint estimation approach is a superior procedure to fit the storage model with the data. Future estimations of the storage model should no longer neglect the possibility of long-run trends in prices.

For banana, jute and tea, three of the commodities originally studied in Deaton and Laroque (1992), the storage model with or without a deterministic trend fails to reproduce the main features of the price dynamics. In the case of the specifications in this paper, the storage model is rejected as a relevant model to explain the price dynamics of these commodities. However, other specifications could be considered. For simplicity, this paper focused on deterministic time trends which allow an analytical likelihood to be characterized. The setup could be extended to other trend specifications. For example, only small changes would be required to account for deterministic trends with structural breaks. The inclusion of stochastic trends (e.g., ARIMA) would be more difficult. A stochastic trend would prevent a direct calculation of the likelihood and would require the use of filters for non-linear state-space models (see e.g., Fernández-Villaverde and Rubio-Ramírez, 2007, for use of the particle filter to estimate macroeconomic models).

References

- Adjemian, M. K. and Smith, A. (2012). Using USDA forecasts to estimate the price flexibility of demand for agricultural commodities. *American Journal of Agricultural Economics*, 94(4), 978–995.
- Ardeni, P. G. and Wright, B. (1992). The Prebisch-Singer hypothesis: A reappraisal independent of stationarity hypotheses. *The Economic Journal*, 102(413), 803–812.
- Arezki, R., Hadri, K., Loungani, P. and Rao, Y. (2014). Testing the Prebisch-Singer hypothesis since 1650: Evidence from panel techniques that allow for multiple breaks. *Journal of International Money and Finance*, 42, 208–223, Special issue: Understanding International Commodity Price Fluctuations.
- Bobenrieth, E. S. A., Bobenrieth, J. R. A., Wright, B. D. and Zeng, D. (2014). *How Biofuels Policies Boosted Grain Staple Prices: A Counterfactual Analysis*. Paper 5125, Agricultural and Applied Economics Association, 2014 Annual Meeting, July 27–29, 2014, Minneapolis, Minnesota.
- Bobenrieth, E. S. A., Wright, B. D. and Zeng, D. (2013). Stocks-to-use ratios as indicators of vulnerability to spikes in global cereal markets. *Agricultural Economics*, 44(S1), 43–52.
- Boskin, M. J., Dulberger, E. R., Gordon, R. J., Grilliches, Z. and Jorgenson, D. W. (1996). *Toward A More Accurate Measure Of The Cost Of Living*. Final report, US Senate Finance Committee.
- Cafiero, C., Bobenrieth, E. S. A. and Bobenrieth, J. R. A. (2011a). Storage arbitrage and commodity price volatility. In A. Prakash (ed.) *Safeguarding food security in volatile global markets*, chapter 15, (pp. 288–313). Rome: FAO.
- Cafiero, C., Bobenrieth, E. S. A., Bobenrieth, J. R. A. and Wright, B. D. (2011b). The empirical relevance of the competitive storage model. *Journal of Econometrics*, 162(1), 44–54.
- Cafiero, C., Bobenrieth, E. S. A., Bobenrieth, J. R. A. and Wright, B. D. (2015). Maximum Likelihood estimation of the standard commodity storage model: Evidence from sugar prices. *American Journal of Agricultural Economics*, 97(1), 122–136.

- Canova, F. (1998). Detrending and business cycle facts: A user's guide. *Journal of Monetary Economics*, 41(3), 533–540.
- Canova, F. (2014). Bridging DSGE models and the raw data. *Journal of Monetary Economics*, 67, 1–15.
- Cashin, P. and McDermott, C. J. (2002). The long-run behavior of commodity prices: Small trends and big variability. *IMF Staff Papers*, 49(2), 175–199.
- Cuddington, J. T. (1992). Long-run trends in 26 primary commodity prices: A disaggregated look at the Prebisch-Singer hypothesis. *Journal of Development Economics*, 39(2), 207–227.
- Deaton, A. and Laroque, G. (1992). On the behaviour of commodity prices. *Review of Economic Studies*, 59(1), 1–23.
- Deaton, A. and Laroque, G. (1996). Competitive storage and commodity price dynamics. *Journal of Political Economy*, 104(5), 896–923.
- Fernandez, V. (2012). Trends in real commodity prices: How real is real? *Resources Policy*, 37(1), 30–47.
- Fernández-Villaverde, J. and Rubio-Ramírez, J. F. (2007). Estimating macroeconomic models: A likelihood approach. *The Review of Economic Studies*, 74(4), 1059–1087.
- Ferroni, F. (2011). Trend agnostic one-step estimation of DSGE models. *The B.E. Journal of Macroeconomics*, 11(1), 1–36.
- Grilli, E. R. and Yang, M. C. (1988). Primary commodity prices, manufactured goods prices, and the terms of trade of developing countries: What the long run shows. *World Bank Economic Review*, 2(1), 1–47.
- Guerra, E. A., Bobenrieth, E. S. A., Bobenrieth, J. R. A. and Cafiero, C. (forthcoming). Empirical commodity storage model: the challenge of matching data and theory. *European Review of Agricultural Economics*.
- Gustafson, R. L. (1958). *Carryover Levels for Grains: A Method for Determining Amounts that are Optimal Under Specified Conditions*. Technical Bulletin 1178, US Dept. of Agriculture.
- Hamilton, J. D. (1994). *Time Series Analysis*. Princeton: Princeton University Press.
- Harrell, F. E. (2001). *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*. New York: Springer.
- Harvey, A. C. and Jaeger, A. (1993). Detrending, stylized facts and the business cycle. *Journal of Applied Econometrics*, 8(3), 231–247.
- Michaelides, A. and Ng, S. (2000). Estimating the rational expectations model of speculative storage: A Monte Carlo comparison of three simulation estimators. *Journal of Econometrics*, 96(2), 231–266.
- Miranda, M. J. and Rui, X. (1999). An empirical reassessment of the commodity storage model, mimeo, The Ohio State University.
- Nocedal, J. and Wright, S. J. (2006). *Numerical Optimization*. New York: Springer.
- Pfaffenzeller, S., Newbold, P. and Rayner, A. (2007). A short note on updating the Grilli and Yang commodity price index. *World Bank Economic Review*, 21(1), 151–163.
- Prebisch, R. (1950). *The economic development of Latin America and its principal problems*. New York: United Nations.
- Rios, L. M. and Sahinidis, N. V. (2013). Derivative-free optimization: a review of algorithms and

- comparison of software implementations. *Journal of Global Optimization*, 56(4), 1247–1293.
- Roberts, M. J. and Schlenker, W. (2013). Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate. *The American Economic Review*, 103(6), 2265–2295.
- Seale, J. L., Jr. and Regmi, A. (2006). Modeling international consumption patterns. *Review of Income and Wealth*, 52(4), 603–624.
- Singer, H. W. (1950). The distribution of gains between investing and borrowing countries. *The American Economic Review*, 40(2), 473–485.
- Svedberg, P. and Tilton, J. E. (2006). The *Real*, real price of nonrenewable resources: Copper 1870–2000. *World Development*, 34(3), 501–519.
- Vaz, A. I. F. and Vicente, L. N. (2007). A particle swarm pattern search method for bound constrained global optimization. *Journal of Global Optimization*, 39(2), 197–219.
- World Bank and FAO (2012). *The Grain Chain: Food Security and Managing Wheat Imports in Arab Countries*. Washington, DC.: World Bank.
- Zeng, D. (2012). Out of sight, out of mind? Estimating commodity price dynamics using detrended price, mimeo, UC Berkeley.

Online Appendix

A Computational details

The results were obtained using MATLAB R2014a on a PC with two quad-core processors Intel Xeon E5345 (2.33 GHz) with 32 GB of RAM running Ubuntu 12.04.5 64 bits. The maximization of the log-likelihood was done using Vaz and Vicente's (2007) free particle swarm pattern search software PSwarm version 2.1¹⁵ and the MATLAB function `fmincon` available in MATLAB Optimization Toolbox. The Monte-Carlo simulations have been performed using MATLAB's default random number generator with the seed set to 1. The Gaussian quadrature was calculated using a MATLAB's function from John Burkardt's website.¹⁶ For a unit normal distribution truncated at five standard deviations, the Gaussian quadrature with 10 nodes has nodes $\varepsilon^n = \{\pm 4.4576, \pm 3.3999, \pm 2.3838, \pm 1.4132, \pm 0.4684\}$ and weights $\pi^n = \{1.9834 \times 10^{-5}, 1.2876 \times 10^{-3}, 2.3048 \times 10^{-2}, 0.14029, 0.33536\}$.

B Small Sample Properties

In this section, we assess using Monte Carlo experiments the small sample properties of the simulated Unconditional Maximum Likelihood estimator (UML) we developed, and compare them to those of the Conditional Maximum Likelihood estimator (CML) proposed in Cafiero et al. (2015). From equation (20), the conditional log-likelihood without trend is obtained by removing the terms corresponding to the marginal likelihood:

$$\log L^C(\theta; P_{1:T}^{\text{obs}}) = -\frac{T-1}{2} \log 2\pi - (T-1) \log [\Phi(5) - \Phi(-5)] + \sum_{t=2}^T \log \left| \mathcal{P}^{-1'}(P_t^{\text{obs}}) \right| - \sum_{t=2}^T (1_{|\varepsilon_t| \leq 5} \cdot \varepsilon_t^2 + 1_{|\varepsilon_t| > 5} \cdot \infty) / 2. \quad (21)$$

Following Michaelides and Ng (2000) and Cafiero et al. (2015), we conduct four Monte Carlo experiments varying the parameterization and the length of the samples. The first set of parameters are $a = 1$, $b = -1$, and $k = 0.02$, which implies a storage cost of 2% of the mean price and, for supply shocks with a coefficient of variation of 5%, a demand elasticity of -0.05 in the range of the best estimates obtained by Roberts and Schlenker (2013) on a caloric aggregate of major crops, but slightly higher in absolute value than our estimated elasticities. The second parameterization only differs by the value of b , now equal to -2 . This rotation of the slope of the demand function around its mean halves the demand elasticity making this parameterization more favorable to storage, and closer to the values found in the article. For each set of parameters, we solve for the equilibrium price function on a grid of 1,000 points, and obtain 3,000 prices series of length $T = 50$ and $T = 100$ from the asymptotic distribution. The price series are obtained by the simulation of 3,000 trajectories starting from the steady-state availability and discarding the first 50 periods as burn-in periods.

The numerical methods follow what was described previously, but differ on two aspects. Firstly, to prevent the availability corresponding to the cutoff price to be below the lower bound of the grid of

¹⁵<http://www.mat.uc.pt/~lnv>

¹⁶http://people.sc.fsu.edu/~jburkardt/m_src/truncated_normal_rule/truncated_normal_rule.html

interpolation points, the lower bound is changed from being -2 to being -5 , the minimum availability. Secondly, since the log-likelihood optimization behaves better on simulated samples, we use a faster optimization algorithm: the generalized pattern search algorithm implemented by the MATLAB function `patternsearch` available in MATLAB Global Optimization Toolbox. The optimization starts from initial values randomly drawn in the range between 80% and 120% of the true values. If the optimization solver fails to converge for one of the two estimators, we discard the corresponding samples for both estimators. The results of valid estimates obtained on common samples are given in table A1 and A2.

Table A1: Comparison of Monte Carlo Experiment Results with Parameterization $a = 1$, $b = -1$, and $k = 0.02$

	UML			CML		
	a	b	k	a	b	k
$T = 50$						
Mean	0.9930	-0.9720	0.0193	0.9932	-0.9745	0.0194
Standard deviation	0.0650	0.1297	0.0059	0.0668	0.1344	0.0058
Bias	-0.0070 (0.70%)	0.0280 (2.80%)	-0.0007 (3.40%)	-0.0068 (0.68%)	0.0255 (2.55%)	-0.0006 (3.14%)
RMSE	0.0654 (6.54%)	0.1326 (13.26%)	0.0059 (29.67%)	0.0671 (6.71%)	0.1368 (13.68%)	0.0058 (29.24%)
$T = 100$						
Mean	0.9951	-0.9854	0.0196	0.9944	-0.9875	0.0194
Standard deviation	0.0510	0.1066	0.0053	0.0516	0.1070	0.0052
Bias	-0.0049 (0.49%)	0.0146 (1.46%)	-0.0004 (2.17%)	-0.0056 (0.56%)	0.0125 (1.25%)	-0.0006 (2.75%)
RMSE	0.0512 (5.12%)	0.1076 (10.76%)	0.0053 (26.70%)	0.0519 (5.19%)	0.1077 (10.77%)	0.0053 (26.26%)

Notes: The price samples for which one estimator does not converge are discarded. For $T = 50$, the total number of valid replications is 2,737 for UML and 2,733 for CML. For $T = 100$, it is 2,841 for UML and 2,846 for CML. The table reports the 2,593 and 2,764 valid estimates obtained on common samples for the short and long samples.

The results of the Monte Carlo experiments are similar to those obtained for the CML in Cafiero et al. (2015). They show that the two maximum likelihood estimators yield precise estimates of the parameters of the model, especially for a and b . The storage cost, k , is less precisely estimated with Root Mean Square Errors (RMSE) always above 26%. For all parameters, the bias is small, most of the RMSE coming from the standard deviation of the estimations.

For all parameterizations, the estimators perform better when the sample length increases. For the UML, a doubling of the sample length from 50 to 100 observations reduces the RMSE by 17% for both parameterizations. The CML benefits slightly more than the UML from an increase in the sample size. Indeed, they have similar RMSE for the long samples, but the UML performs better on the short samples. Regarding the influence of the parameterization, we observe that the parameterization more favorable to storage yields less precise estimates as all the RMSE of table A2 are higher than in table A1.

Table A2: Comparison of Monte Carlo Experiment Results with Parameterization $a = 1$, $b = -2$, and $k = 0.02$

	UML			CML		
	a	b	k	a	b	k
$T = 50$						
Mean	0.9793	-1.9639	0.0194	0.9772	-1.9695	0.0193
Standard deviation	0.1280	0.3202	0.0085	0.1351	0.3250	0.0085
Bias	-0.0207 (2.07%)	0.0361 (1.80%)	-0.0006 (3.14%)	-0.0228 (2.28%)	0.0305 (1.52%)	-0.0007 (3.69%)
RMSE	0.1296 (12.96%)	0.3222 (16.11%)	0.0085 (42.51%)	0.1370 (13.70%)	0.3265 (16.32%)	0.0085 (42.45%)
$T = 100$						
Mean	0.9848	-1.9838	0.0195	0.9825	-1.9860	0.0194
Standard deviation	0.1055	0.2592	0.0074	0.1090	0.2611	0.0076
Bias	-0.0152 (1.52%)	0.0162 (0.81%)	-0.0005 (2.56%)	-0.0175 (1.75%)	0.0140 (0.70%)	-0.0006 (3.14%)
RMSE	0.1065 (10.65%)	0.2597 (12.99%)	0.0074 (36.94%)	0.1104 (11.04%)	0.2615 (13.08%)	0.0076 (38.02%)

Notes: The price samples for which one estimator does not converge are discarded. For $T = 50$, the total number of valid replications is 2,820 for UML and 2,767 for CML. For $T = 100$, it is 2,899 for UML and 2,890 for CML. The table reports the 2,674 and 2,839 valid estimates obtained on common samples for the short and long samples.

C Production variation

Table A3: Production variation, 1961–2011

Commodity	Production CV (%)	Commodity	Production CV (%)
Banana	3.67	Palm oil	4.65
Cocoa	7.13	Rice	2.75
Coffee	7.68	Sugar	3.65
Copper	4.46	Tea	2.07
Cotton	7.20	Tin	5.52
Jute	11.35	Wheat	4.34
Maize	5.84		

Notes: The coefficients of variation (CV) are obtained by calculating the standard deviation of the detrended logarithm of observed production, modeling the trend using a restricted cubic splines with five knots.

D Parameters estimates

Table A4: Parameter estimates with 3-knot spline trend

Commodity	g_1	g_2	a	b	k	$\log L$	# Stockouts
Banana	-0.3323 (0.2804)	-1.2359 (0.2011)	1.3237 (0.2013)	-1.6729 (0.4397)	0 -	170.6770	0
Cocoa	-1.5599 (0.8803)	0.7076 (0.3317)	0.3679 (0.1595)	-1.7978 (0.9285)	0.0002 (0.0008)	215.6308	0
Coffee	0.0925 (0.3314)	-0.6813 (0.1557)	0.2522 (0.0455)	-0.5265 (0.1072)	0.0036 (0.0025)	194.9321	1
Copper	-0.4176 (0.1943)	-0.2464 (0.1424)	0.7348 (0.0789)	-1.1835 (0.1589)	0.0073 (0.0051)	122.1660	0
Cotton	-0.6422 (0.8240)	-2.0878 (0.3217)	1.5367 (0.4016)	-3.7632 (1.6051)	0.0053 (0.0065)	123.0208	0
Jute	-0.3974 (0.0756)	-0.9164 (0.0554)	0.6917 (0.0481)	-0.7573 (0.0401)	0.0163 (0.0062)	102.1980	4
Maize	-0.6588 (0.1101)	-0.8850 (0.0815)	0.9504 (0.0816)	-1.5377 (0.1432)	0.0128 (0.0075)	84.9322	3
Palm oil	-1.7887 (0.1181)	-1.2443 (0.0539)	1.0411 (0.0981)	-1.7009 (0.1865)	0.0146 (0.0091)	119.2428	1
Rice	-1.5405 (0.3035)	-1.4251 (0.1417)	1.2128 (0.1914)	-2.0786 (0.3684)	0.0037 (0.0066)	127.1550	1
Sugar	-1.2627 (0.1028)	-0.2394 (0.1652)	0.9701 (0.0809)	-1.8482 (0.1738)	0.0235 (0.0098)	60.7431	8
Tea	-0.9290 (0.3585)	-1.1775 (0.1357)	1.1089 (0.2085)	-1.6789 (0.4149)	0.0105 (0.0053)	144.1925	0
Tin	1.7957 (0.4707)	-0.7677 (0.2843)	0.1940 (0.0244)	-0.7567 (0.1185)	0 -	209.3322	0
Wheat	-1.2473 (0.1482)	-0.8441 (0.1517)	1.2069 (0.1069)	-1.5449 (0.2021)	0.0143 (0.0096)	104.0772	1

Note: Asymptotic standard errors in parenthesis.

Table A5: Parameter estimates with 4-knot spline trend

Commodity	g_1	g_2	g_3	a	b	k	$\log L$	# Stockouts
Banana	-0.5940 (0.2995)	-0.2663 (0.3033)	-1.5903 (0.2712)	1.2033 (0.1735)	-1.5061 (0.3602)	0 -	171.7633	0
Cocoa	0.2862 (0.2857)	-2.6193 (0.5337)	0.3417 (0.4282)	0.4498 (0.0981)	-1.1874 (0.2765)	0.0024 (0.0028)	219.7236	0
Coffee	0.1199 (0.1792)	0.0281 (0.3230)	-0.7565 (0.1351)	0.2339 (0.0495)	-0.4885 (0.0898)	0.0033 (0.0024)	194.9185	2
Copper	0.6214 (0.3452)	-0.8595 (0.3155)	-0.2444 (0.1968)	0.7416 (0.1155)	-1.7526 (0.3542)	0.0005 (0.0044)	123.6516	4
Cotton	-0.2545 (0.7621)	-1.1035 (0.8597)	-2.2689 (0.4857)	1.4300 (0.5264)	-3.5214 (2.0232)	0.0049 (0.0069)	122.8690	0
Jute	-0.0613 (0.1039)	-0.6035 (0.1805)	-0.9908 (0.0710)	0.6967 (0.0797)	-0.7614 (0.0715)	0.0163 (0.0066)	102.1845	8
Maize	0.0815 (0.1073)	-1.1899 (0.0887)	-1.1229 (0.0576)	0.9358 (0.0481)	-0.8819 (0.0508)	0.0304 (0.0068)	88.8597	12
Palm oil	-0.9208 (0.1402)	-1.6398 (0.3924)	-1.5053 (0.1190)	0.9971 (0.1744)	-1.6352 (0.3267)	0.0142 (0.0095)	119.6359	1
Rice	-0.8285 (0.2131)	-1.7697 (0.4977)	-1.6808 (0.1837)	1.2815 (0.3139)	-2.2062 (0.5310)	0.0042 (0.0070)	127.1070	1
Sugar	-0.6992 (0.1369)	-1.3979 (0.2635)	-0.3698 (0.2479)	1.1136 (0.1329)	-2.0999 (0.2176)	0.0262 (0.0106)	60.8484	16
Tea	-1.4922 (0.3475)	-0.4250 (0.4818)	-1.8818 (0.2257)	1.1959 (0.2308)	-1.7916 (0.4785)	0.0107 (0.0051)	145.6865	0
Tin	1.7411 (0.5157)	1.1081 (0.2943)	-0.2441 (0.1589)	0.1517 (0.0113)	-0.5043 (0.0414)	0 -	210.3823	6
Wheat	-0.1494 (0.1744)	-1.9057 (0.3087)	-0.9026 (0.1902)	1.3447 (0.1606)	-2.3618 (0.3405)	0 -	105.8230	0

Note: Asymptotic standard errors in parenthesis.

E Figures of price trends




Figure A1: Price trends

References

- Cafiero, C., Bobenrieth, E. S. A., Bobenrieth, J. R. A. and Wright, B. D. (2015). Maximum Likelihood estimation of the standard commodity storage model: Evidence from sugar prices. *American Journal of Agricultural Economics*, 97(1), 122–136.
- Michaelides, A. and Ng, S. (2000). Estimating the rational expectations model of speculative storage: A Monte Carlo comparison of three simulation estimators. *Journal of Econometrics*, 96(2), 231–266.
- Roberts, M. J. and Schlenker, W. (2013). Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate. *The American Economic Review*, 103(6), 2265–2295.
- Vaz, A. I. F. and Vicente, L. N. (2007). A particle swarm pattern search method for bound constrained global optimization. *Journal of Global Optimization*, 39(2), 197–219.