

WORKING PAPERS

N° 17-776

March 2017

“Bayesian Estimation of the Storage Model using
Information on Quantities”

Christophe Gouel and Nicolas Legrand

Bayesian Estimation of the Storage Model using Information on Quantities*

Christophe Gouel^{†1,2} Nicolas Legrand^{‡3}

¹Économie Publique, INRA, AgroParisTech, Université Paris-Saclay, 78850, Thiverval-Grignon, France

²CEPII, 113 rue de Grenelle, 75007, Paris, France

³Toulouse School of Economics, 21 Allée de Brienne, 31000, Toulouse, France

January 20, 2017

Abstract

This paper presents a new strategy to estimate the rational expectations storage model. It uses information on prices and quantities – consumption and production – in contrast to previous approaches which use only prices. This additional information allows us to estimate a model with elastic supply, and to identify parameters such as supply and demand elasticities, which are left unidentified when using prices alone. The estimation relies on the Bayesian methods popularized in the literature on the estimation of DSGE models. It is carried out on a market representing the caloric aggregate of the four basic staples – maize, rice, soybeans, and wheat – from 1961 to 2006. The results show that to be consistent with the observed volatility of consumption, production, and price, elasticities have to be in the lower ranges of the elasticities in the literature, a result consistent with recent instrumental variable estimations on the same sample.

Keywords: Commodity price dynamics, storage, Bayesian inference.

JEL classification: C51, C52, Q11.

*We would like to thank Alain Ayong le Kama, Stéphane De Cara, Stéphane Lhuissier, and Fabien Tripier for their helpful suggestions. 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.

[†]christophe.gouel@grignon.inra.fr

[‡]nicolas.legrand@tse-fr.eu

1 Introduction

In research on commodity prices behavior's, the rational expectations storage model of [Gustafson \(1958\)](#) has become the workhorse model for prices formation and commodity market volatility studies. In the model, the dynamics of commodity prices is explained by real shocks to supply and demand influencing the decisions made by optimizing forward-looking agents. Although the theoretical foundations of the storage model are well grounded, its empirical validity continues to be challenged. In pathbreaking papers, [Deaton and Laroque \(1992, 1996\)](#) cast doubt on its relevance because of its inability to generate the observed high levels of autocorrelation in prices. This serial correlation puzzle has been partially solved in part through several refinements to improve the accuracy of the numerical methods ([Cafiero et al., 2011](#)), the development of a Maximum Likelihood estimator ([Cafiero et al., 2015](#)), the choice of the data used ([Guerra et al., 2015](#)), and management of the trend ([Gouel and Legrand, 2015](#)). However, all these previous estimations are based only on prices information, and ignore the data on quantities, although these would introduce some noise. Another problem with price-only estimations is that many parameters, such as those related to the supply process, are left unidentified. So these estimations have to assume an inelastic supply which is modeled by a unit normal random shock.

In this paper, we develop a new econometric approach to the storage model estimation that is able to use observables other than prices by borrowing from the recent literature on the estimation of DSGE models. This new approach allows us to estimate a richer storage model than the specifications previously estimated. For example, our storage model features two structural shocks to supply and demand, instead of only one to supply, and includes producers that react to expected price changes. Now that the literature has showed that the storage model is able to fit price dynamics, it is important to assess whether it is able to do so while also accommodating quantities behavior.

In the standard storage model, when observables other than prices are added, the model presents a stochastic singularity. With only one shock in the model, the combination of observables would not be consistent with the theory leading to a likelihood that would be zero with probability one. To avoid this problem, we follow the literature on the estimation of DSGE models by adding structural and measurement shocks ([Fernández-Villaverde, 2010](#)). Adding shocks of measurement errors may have important consequences for the estimation procedure. These additional shocks may make impossible to recover the state of the model from the observables and so to calculate the likelihood analytically. This situation would call for nonlinear filtering techniques, known as particle filtering. In this preliminary version of the paper, we focus on cases such that measurement errors do not prevent the observation of the state, which allows us to have likelihood functions that take analytic form. Because of this restriction, we have to assume that prices are always observed without noise so that a one-to-one mapping exists between prices and market availability. Only consumption and production are assumed to be observed with noise. Depending on the number of structural shocks and observed variables used, we estimate a sequence of model specifications allowing us to analyze how information from consumption and production contributes to parameters' identification. Our most simple specification, with prices as the only observables, leads to the same likelihood function as [Cafiero et al. \(2015\)](#). Starting from [Cafiero et al.](#) approach, we build more general likelihood functions in order to accommodate the additional observables. As for other rational expectations models, the likelihood functions are challenging to maximize and we estimate the model using Bayesian methods.

Because this is a new approach to estimating the storage model, we carry the estimations on a

sample for which we have good benchmark estimations to which our work can be compared: the global markets for calories coming from maize, rice, soybeans, and wheat from 1961 to 2006. [Roberts and Schlenker \(2013\)](#) use the storage model to motivate the choice of instrumental variables for reduced-form estimations of supply and demand on this market for calories. Aggregating staple foods into calories allows to account for about 75% percent of the global food production and to bypass many issues related to the close substitutability in production and consumption of these major crops, such as the risks of mixing own-price with cross-price elasticities. Even if our elasticities tend to be slightly lower, overall our results are consistent with [Roberts and Schlenker's](#) results. Our estimations also predict some stockouts over the estimation sample (2 to 5 depending on the specification), consistent with the rare price spikes observed in the sample.

The remainder of the paper is organized as follows. Section 2 describes the storage model and the solution method. Section 3 presents our data and some descriptive statistics. Section 4 presents the Bayesian estimation strategy, the chosen priors, and the model's posterior. Section 5 concludes the paper.

2 The Model

2.1 Model equations

Let consider a model of the world market for a storable staple (here the caloric aggregate of maize, rice, soybeans, and wheat). The storage model is based on [Wright and Williams \(1982\)](#) and is a partial equilibrium model featuring profit-maximizing producers, competitive storers with constant marginal storage costs which can transfer the good across time, and a demand function. We extend [Wright and Williams'](#) model to make it consistent with a balanced growth path driven by a deterministic increase in demand and a reduction in marginal production costs.

There are a few differences with respect to the storage models used for previous structural estimations ([Deaton and Laroque, 1992, 1996](#); [Cafiero et al., 2011, 2015](#)). First, in our case supply is elastic to prices while previous models tell us nothing about the agricultural producers' reactions to prices: producers are not represented explicitly and production consists only of an exogenous supply shock. Second, we add another structural shock to demand. Most storage models feature only one structural shock, usually a supply shock, which in fact can be interpreted as a net supply shock. When using information on consumption and production for the estimation, it is useful to have structural shocks corresponding to each observable. To avoid making the numerical method more complex, we assume that demand and supply shocks are additive, which allows us to sum shocks, available stocks and planned production, leading to a model with only one state variable.

Producers A representative producer makes its production decision and pays for inputs one period before bringing its output to the market. The production choice, made in period t , is denoted \tilde{Q}_t . Realized production differ from planned production by an exogenous additive shock, $\tilde{\eta}_{t+1}^Q$, occurring during the growing season (e.g., a weather disturbance). The producer's problem in period t can be written as

$$\max_{\tilde{Q}_t} (1+r)^{-1} E_t P_{t+1} \left(\tilde{Q}_t + \tilde{\eta}_{t+1}^Q \right) - \tilde{\Psi}_t \left(\tilde{Q}_t \right), \quad (1)$$

where r is the real interest rate which is assumed to be fixed, E_t is the expectation operator conditional on period t information, P_{t+1} is the price, and $\tilde{\Psi}_t(\cdot)$ is a differentiable and convex production cost function. The solution to this problem is the following first-order condition

$$(1+r)^{-1} E_t P_{t+1} = \tilde{\Psi}'_t(\tilde{Q}_t). \quad (2)$$

At each period, the producer rationally plants up to the point where the expected marginal benefit equals the marginal production cost.

Storers Competitive storers are risk-neutral. For storing an amount $\tilde{S}_t \geq 0$ from period t until $t+1$ they incur a physical cost of storage proportional to the stored quantities, $k\tilde{S}_t$, and an opportunity cost. Assuming rational expectations and taking account of the non-negativity constraint on storage yield the following arbitrage condition

$$(1+r)^{-1} E_t P_{t+1} - P_t - k \leq 0, = 0 \text{ if } \tilde{S}_t > 0. \quad (3)$$

Final demand Final demand for the good is the sum of a downward sloping demand function $\tilde{D}_t(P_t)$ and a demand shock denoted $\tilde{\eta}_t^D$.

Equilibrium The market clears when previous stocks plus production equal final demand plus demand for stocks:

$$\tilde{S}_{t-1} + \tilde{Q}_{t-1} + \tilde{\eta}_t^Q = \tilde{D}_t(P_t) + \tilde{\eta}_t^D + \tilde{S}_t. \quad (4)$$

Trends and steady-state growth Consumption and production of food grow steadily over time because of population increase, income growth, or technological progress. To account for this fact, we allow the demand function, the marginal cost function, and the structural shocks to have deterministic trends. Because of the additivity of the market clearing equation, to obtain steady-state growth requires the adoption of multiplicative trends. Detrended variables are expressed without “tilde” and related to their trending counterparts by the following relations

$$\tilde{D}_t(P_t) = (1+g)^t D(P_t), \quad (5)$$

$$\tilde{\Psi}'_t(\tilde{Q}_t) = \Psi'(Q_t), \quad (6)$$

$$\tilde{\eta}_t^Q = (1+g)^t \eta_t^Q, \quad (7)$$

$$\tilde{\eta}_t^D = (1+g)^t \eta_t^D, \quad (8)$$

where g is the assumed growth rate. Equation (6) imposes that the downward trend in marginal production costs perfectly offsets the trend in production, leaving prices stationary. Because of the additive nature of equation (3) determining the storage level, to be compatible with a steady-state growth any trend in prices will require marginal storage costs, k , either to be zero or to have the same trend as prices, two restrictions which are unlikely to hold. To avoid imposing these restrictions, we ignore potential trends in prices.¹ This set of trend assumptions induces a common multiplicative trend in all quantities: final

¹See Dvir and Rogoff (2014) for another storage model with trending quantities but without trending prices.

consumption, production and stocks. Using the convention $\tilde{S}_t = S_t (1 + g)^t$ and $\tilde{Q}_{t-1} = Q_{t-1} (1 + g)^t$, replacing trending quantities by their detrended counterparts in the market clearing equation leads to

$$\frac{S_{t-1}}{1 + g} + Q_{t-1} + \eta_t^Q = D(P_t) + \eta_t^D + S_t. \quad (9)$$

The division of S_{t-1} by $1 + g$ illustrates that, on average, stocks have to increase just to meet the pace of increase in production and demand, so detrended beginning stocks are discounted by $1 + g$ to maintain them at a level comparable to other detrended quantities.

Stationary equilibrium In equation (9), four variables are predetermined: stocks, planned production and the two shocks. They can be summed together in a single state variable, availability net of demand shock A_t :

$$A_t \equiv \frac{S_{t-1}}{1 + g} + Q_{t-1} + \eta_t^Q - \eta_t^D. \quad (10)$$

The capacity to combine all predetermined variables in one state variable is crucial to simplify numerical resolution of the model and this simplification is allowed by assuming an additive demand shock.

Equation (10) is the model's stationary transition equation. Applying previous transformations to the equilibrium equations leads to the following three stationary equilibrium equations:

$$(1 + r)^{-1} E_t P_{t+1} = \Psi'(Q_t), \quad (11)$$

$$(1 + r)^{-1} E_t P_{t+1} - P_t - k \leq 0, = 0 \text{ if } S_t > 0, \quad (12)$$

$$A_t = S_t + D(P_t). \quad (13)$$

We assume that the stationary demand function takes the following linear form

$$D(P_t) = \bar{D} \left(1 + \alpha^D \frac{P_t - \bar{P}}{\bar{P}} \right), \quad (14)$$

where \bar{D} is the steady-state demand (equal also to steady-state production since stocks are not held at the steady state), $\alpha^D < 0$ is the price elasticity of demand at the steady state, and \bar{P} is the steady-state price. Similarly, the stationary marginal cost function is assumed to be linear:

$$\Psi'(Q_t) = (1 + r)^{-1} \bar{P} \left(1 + \frac{1}{\alpha^S} \frac{Q_t - \bar{D}}{\bar{D}} \right), \quad (15)$$

where $\alpha^S > 0$ is the supply elasticity at the steady state. Because of the assumed specifications in deviation from the deterministic steady state, these demand and marginal cost functions depend only on parameters that can be directly interpreted, which is important for the prior elicitation of the Bayesian analysis.

η_t^Q and η_t^D , the stationary supply and demand disturbances, are assumed to be serially uncorrelated, independent of each other, and are normally distributed with zero mean and standard deviations σ_{η_Q} and σ_{η_D} .

We also estimate a version of the model with inelastic supply. In this case, equation (11) is dropped from the model and planned production is fixed at its steady-state level: $Q_t = \bar{D}$. Except for the demand shock and the growth rate g , this version of the storage model is the same as the one estimated in [Deaton](#)

and Laroque (1992).

2.2 Model solution

Equations (10)–(13) form a non-linear rational expectations system based on the variables A_t , Q_t , S_t , and P_t driven by the innovations $\eta_t = [\eta_t^Q, \eta_t^D]$. This system does not have a closed form solution and must be solved numerically to allow for a structural estimation. Solution of the rational expectations system takes the form of policy functions which describe control variables as functions of the contemporaneous state variable, net availability:

$$Q_t = \mathcal{Q}(A_t), \quad (16)$$

$$S_t = \mathcal{S}(A_t), \quad (17)$$

$$P_t = \mathcal{P}(A_t), \quad (18)$$

to which the transition equation (10) must be added.

Combining equations (10) to (13), one can see that the policy functions for all A_t have to satisfy:

$$\mathcal{P}(A_t) = \max \left[(1+r)^{-1} E_t \mathcal{P} \left(\mathcal{S}(A_t) + \mathcal{Q}(A_t) + \eta_{t+1}^Q - \eta_{t+1}^D \right) - k, D^{-1}(A_t) \right], \quad (19)$$

$$(1+r)^{-1} E_t \mathcal{P} \left(\mathcal{S}(A_t) + \mathcal{Q}(A_t) + \eta_{t+1}^Q - \eta_{t+1}^D \right) = \Psi'(\mathcal{Q}(A_t)). \quad (20)$$

Equation (19) makes clear the existence of two regimes. The first regime holds when speculators stockpile in the expectation of future prices covering 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 net availability in the market. Let us define P^* as the cutoff price above which there is no storage. This cutoff price is the solution to the following nonlinear equation

$$P^* = (1+r)^{-1} E_t \mathcal{P} \left(\mathcal{Q}(D(P^*)) + \eta_{t+1}^Q - \eta_{t+1}^D \right) - k, \quad (21)$$

which depends on the policy functions.

A variety of numerical techniques is available to solve rational expectations storage models and approximate the policy functions (Gouel, 2013). Previous estimations of the storage model relied on the fixed-point algorithm of Deaton and Laroque (1992). This has the advantage that it is very fast because it requires no rootfinding operation. However, it does not apply to models with elastic supply. The numerical algorithm we implement is the endogenous grid method proposed by Carroll (2006), adapted to the storage model in Gouel (2013). In the case of a model with inelastic supply, like Deaton and Laroque's algorithm it implies simple arithmetic operations only, but no rootfinding. With elastic supply, it is necessary to solve nonlinear equations, but this algorithm is still much faster than the alternatives. One advantage over the approach in Deaton and Laroque is that the endogenous grid method locates precisely the cutoff point of no storage even with a small number of grid points, whereas the the fixed-point algorithm requires a large number of grid points (Cafiero et al., 2011).

To implement the algorithm, the expectations terms in equations (11) and (12) are replaced by sums by discretizing the normal distribution of the net supply shocks $\eta^Q - \eta^D$ using a Gauss-Hermite quadrature with 11 nodes. We use a regular grid on storage of 50 nodes spanning the interval $[0, \bar{D}]$, where \bar{D} is fixed at 1, its central value in the estimations below. The algorithm stops when successive policy functions

cease to differ by more than a given tolerance threshold, set to 10 decimal places. See [Gouel \(2013\)](#) for more details on the implementation of the algorithm.

3 Overview of the grains market

3.1 Data

The data on world quantities for maize, rice, soybeans and wheat come from the Food and Agriculture Organization statistical database (FAOSTAT). Total demand is obtained by subtracting the variation in inventories from total production. Our sample starts in 1961, the first year available in FAOSTAT. It ends in 2006 because, after that year, the increasing importance of biofuels, supported strongly by US and EU policies, may have changed the behavior of the food markets (biofuel mandates create an additional inelastic demand for food products, while biofuel demand induced by tax credits may link food and oil markets).² Following the approach proposed by [Roberts and Schlenker \(2013\)](#), the four commodities are aggregated together into calories using the conversion ratios in [Williamson and Williamson \(1942\)](#). [Figure 1](#) depicts the patterns of the world production and consumption over the sample period. Although production appears more volatile than consumption, both production and consumption grow at the same steady linear rate over time. We need stationary variables for the estimation. Production and consumption are regressed on the same linear time trend. The residuals are centered on 2006 trend values and normalized to have unitary mean.

Regarding prices, a crucial issue is deciding what is the most relevant annual price series to use to estimate of the storage model. The series of prices used in [Deaton and Laroque \(1992\)](#) are formed by averaging prices over the calendar year, which might induce spurious correlations due to mixing together different marketing seasons. [Guerra et al. \(2015\)](#) compare the model fit for a storage model when estimated on different price series for maize and find that the price for the first month following the delivery is what works best. [Roberts and Schlenker \(2013\)](#) adopt the strategy of a single month per year, which is the approach used here. Nominal spot prices are taken from the World Bank pinksheets,³ which provide monthly prices by averaging the daily prices observed during each month. The annual prices is the price observed during the month following delivery in the main market (December for maize and wheat, and November for soybeans and rice). The resulting annual prices series are deflated by the US CPI and aggregated to a single caloric price index series similar to the calculation for quantities. The deflated price index exhibits a downward trend over the sample period ([figure 2](#)). To remove the trend, we follow [Roberts and Schlenker \(2013\)](#) and regress log prices over a restricted cubic spline with three knots which offers enough flexibility to account for small movements in the trend.⁴

3.2 Descriptive statistics

Next, we present some descriptive statistics for the detrended data and discuss their implications for the estimation of the storage model.

²See [Wright \(2014\)](#) and [de Gorter et al. \(2015\)](#) for more.

³We select the “Maize (US), no. 2, yellow”, “Rice (Thailand), 5% broken, f.o.b. Bangkok”, “Soybeans (US), c.i.f. Rotterdam”, and “Wheat (US), no. 1, hard red winter” series.

⁴The knots are located in 1963, 1984, and 2005 as suggested in [Roberts and Schlenker \(2013, Online appendix\)](#).

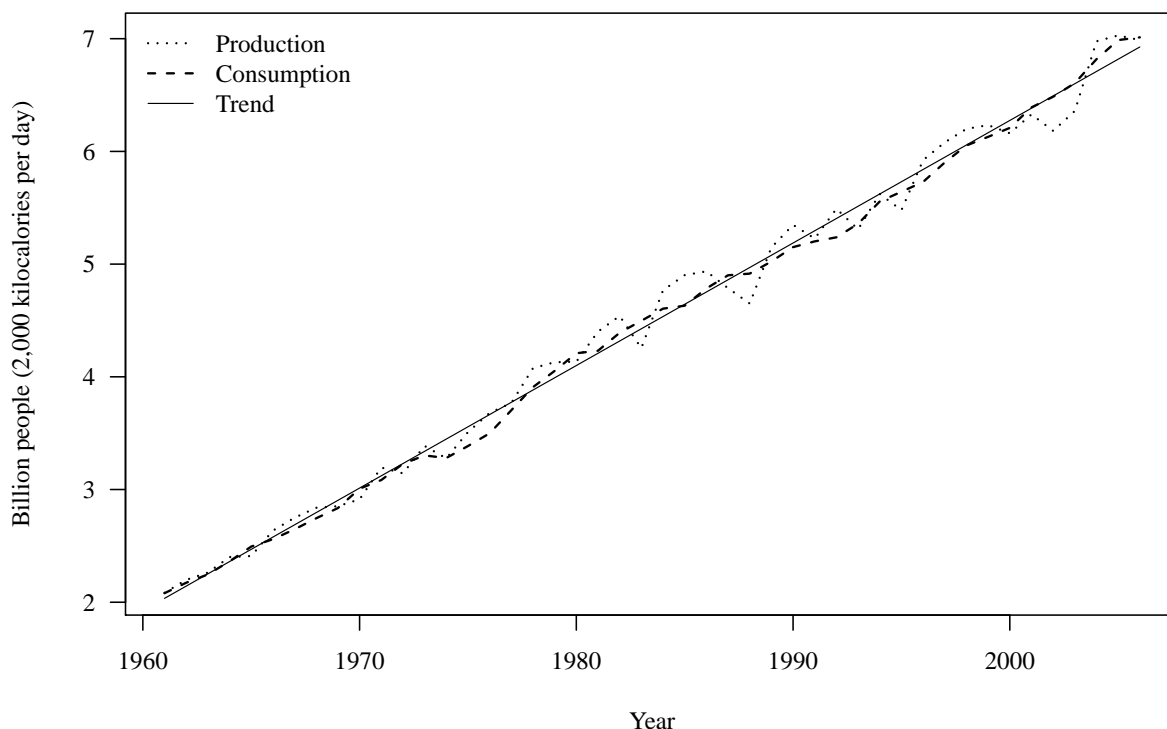


Figure 1: World caloric production and consumption, and their trend for 1961–2006. The y-axis is the number of people who hypothetically could be fed on a 2,000 kilocalories per day diet based on consuming only the four commodities.

The high levels of the correlation coefficients (all in excess of .85) between the calculated index of grains calories and the real price of each individual crop (table 1) are indicative of the large substitution possibilities between these crops. With the exceptions of soybeans and maize, the crops are correlated more to the aggregate than to any other crop. This supports use of an aggregated caloric indicator as a suitable measure of the state of the world food market. Estimation based on each crop individually would create the risk of mixing own-price and cross-price elasticities.

Table 1: Correlation coefficients of detrended real prices data, 1961–2006

Commodity	Grains	Maize	Rice	Soybeans	Wheat
Grains	1.000	0.897	0.972	0.856	0.950
Maize	0.897	1.000	0.795	0.886	0.873
Rice	0.972	0.795	1.000	0.774	0.875
Soybeans	0.856	0.886	0.774	1.000	0.813
Wheat	0.950	0.873	0.875	0.813	1.000

Notes: Prices are detrended using a restricted cubic splines using three knots. “Grains” refers to the caloric aggregate of maize, rice, soybeans, and wheat.

Table 2 reports the descriptive statistics of the data used to estimate the model. The first-order autocorrelation of prices is high at 0.615. Higher values of autocorrelation (mostly above 0.8) were observed on a different sample by Deaton and Laroque (1992), which led them to reject the storage model because of its inability to match these high levels. Competitive storage behavior is able to explain

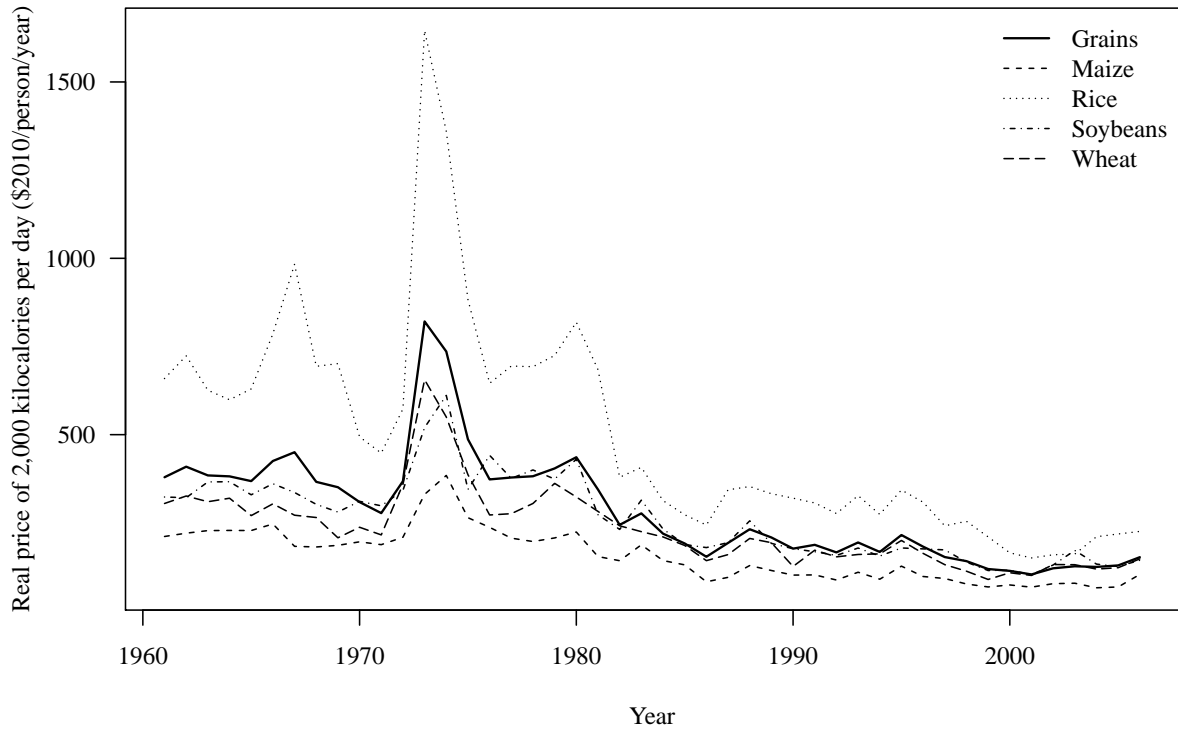


Figure 2: Real caloric prices over time. The y-axis is the annual cost of 2,000 kilocalories per day. The price series are taken from the World Bank pink sheets.

positively autocorrelated prices, but not very high levels. Hence, the importance of detrending the prices because the model is able to explain only cyclical fluctuations not long-run trends (Gouel and Legrand, 2015). Consumption also presents a high serial correlation at 0.73, which is consistent with the fact that, except in the case of structural and measurement shocks, consumption should follow price behavior. The first- and second-order autocorrelations are higher for consumption than for price, likely explained by autocorrelated demand shocks. We do not consider this possibility here, but it would make sense if demand shocks are not idiosyncratic but related to the business cycle. Production presents a small first-order autocorrelation of 0.163. With respect to the storage model, this small autocorrelation indicates that uncorrelated supply shocks are quantitatively more important than endogenous supply reactions. The absence of autocorrelation in production supports the relevance of storage to generate price persistence.

Table 2: Autocorrelation and coefficient of variation of detrended caloric data, 1961–2006

Variable	One-year a-c	Two-year a-c	Coefficient of variation
Production	0.163	-0.098	0.021
Consumption	0.730	0.418	0.011
Price	0.615	0.181	0.298

The coefficients of variation of production and consumption are quite low at 2.1% and 1.1% respectively. The fact that production is more volatile than consumption shows the importance of inter-annual storage. Without storage, changes in production levels would have to be matched each year by corre-

sponding changes in consumption levels. Price is much more volatile than consumption. To reproduce this would require a very inelastic demand curve.

Table 3 displays the correlation coefficients of production, consumption, and price. The correlations have the expected signs, with a negative correlation between consumption and price and a positive one between production and consumption. However, production and consumption are not perfectly correlated, which again points to the role of storage. Because of storage, in the model, consumption is not determined by production but by the sum of production and stocks. Therefore, production should affect consumption, but this effect is mitigated by the presence of stocks. The low correlation between consumption and price indicates that a lot of the movements observed for consumption stems not from movements along the demand curve, but from movements in the demand curve itself or from measurement errors. If it was due only to movements along the demand curve the correlation would be close to -1 . This low correlation between consumption and price exemplifies the standard econometric challenge in the estimation of demand elasticities: price variations are caused by supply and demand shocks and movements along the supply and demand curves, with the result that identification of a demand elasticity requires the price variations from exogenous supply shocks to be isolated. Some works have propose clever empirical strategies aimed at isolating clear exogenous supply shocks (Adjemian and Smith, 2012; Roberts and Schlenker, 2013). In this paper, the identification comes from the assumed structural form of the model, which allows to recover the various shocks from the observations.

Table 3: Correlation coefficients of detrended caloric data, 1961–2006

Variable	Production	Consumption	Price
Production	1.000	0.390	−0.122
Consumption	0.390	1.000	−0.121
Price	−0.122	−0.121	1.000

4 Estimation

In this section, we propose an econometric procedure to estimate the storage model using information on quantities. Thus far, structural estimations of the storage model have been undertaken assuming that: (i) supply is inelastic, (ii) only one shock, a supply shock, is generating prices dynamics, (iii) price is the only observable, and (iv) prices are observed without measurement errors. Under these assumptions, Cafiero et al. (2015) develop a Maximum Likelihood estimator. Starting from the same likelihood function, we gradually relax these four assumptions to provide a more comprehensive picture of how production, consumption and spot prices are connected by storage decisions. In this preliminary version of the paper, we maintain assumption (iv) that prices are observed without measurement errors. This assumption allows the likelihood to take an analytic form.

Once evaluated, the likelihood function can either be maximized over a vector of parameters θ following the frequentist approach, or according to Bayes' theorem, can be combined with a prior distribution of the model parameters to form the posterior distribution of the parameters. The posterior distribution then can be used for inference and model comparison. As is often the case when estimating dynamic models, the storage model likelihood is flat in some areas, and exhibits many local maxima. In this context, it is difficult to find the global maximum of the likelihood. This usually requires global search

algorithms which are costly to use when the model to be estimated, such as the storage model, requires a non-trivial length of time for its solution. To avoid these difficulties, we take a Bayesian approach by employing Markov chain Monte Carlo methods for sampling posterior parameters, an approach that has become the norm in macroeconomics for the estimation of DSGE models (Fernández-Villaverde et al., forthcoming).

We maximize the log-posterior function to estimate the mode of the posterior distribution. We use the mode of the posterior as the starting value for a random walk Metropolis-Hastings (RWMH) algorithm, which simulates a Markov chain whose stationary distribution corresponds to the posterior distribution of interest. For the random walk step, the initial variance-covariance matrix of the draws is taken to be the inverse of the outer product of the scores resulting from optimization of the log-posterior function. This matrix is then adjusted during the first 10^5 draws of the chain according to the strategy proposed by Vihola (2012).⁵ A sample of 10^7 draws from the posterior distribution is created from the RWMH algorithm, from which we discard the first half of the chain as a burn-in and retain 1 every 1,000 draws to thin the chain before testing for convergence. Convergence of the Markov chain is assessed graphically and using the convergence diagnostics of Geweke (1992). Both approaches confirm that the simulated chain is long enough to ensure convergence to the stationary distribution of the posterior and to deliver reliable moments estimates of the parameters.

4.1 The likelihood function

4.1.1 Generic expression

A variety of competing model specifications are estimated depending on assumptions about whether supply responds to prices variations, what variables are observable, and the number of structural shocks. For each assumption, different sets of parameters θ can be recovered from the data. Estimating several model specifications allows us to make the link with previous estimations of the storage model, to ease comparability of our results by progressively relaxing assumptions (i) to (iii), and to identify the contribution of each observable to the parameter estimations, given that supply and demand data do not deliver the same kind of information. In all the specifications, there are exactly as many shocks as there are observables to avoid stochastic singularity while maintaining a closed-form solution for the likelihood function.

Given the Markov structure of our model, a set of model parameters θ and a sample of observed variables of length T stacked in $Y_{1:T}^{\text{obs}} \equiv \{Y_1^{\text{obs}}, \dots, Y_T^{\text{obs}}\}$, the likelihood function can be expressed as

$$L(Y_{1:T}^{\text{obs}}|\theta) = \prod_{t=1}^T p(Y_t^{\text{obs}}|Y_{t-1}^{\text{obs}}; \theta). \quad (22)$$

Assumptions about structural measurement error shocks are made such that there are as many shocks as there are observables and such that the method of variable transformation can be applied. In this case, the conditional density of Y_t^{obs} can be written as

$$p(Y_t^{\text{obs}}|Y_{t-1}^{\text{obs}}; \theta) = p(V_t | J_t), \quad (23)$$

⁵This adjustment phase aims at reaching an acceptance rate closes to the optimal value of 23.4% (Robert and Casella, 2004). In all of our estimations the final acceptance rates lie in the interval 18.3–23.3%.

where V_t is a vector of the error terms which include structural and measurement error shocks, and $|J_t|$ is the determinant of the Jacobian of the mapping of $Y_t^{\text{obs}} \mapsto V_t$.

Under the assumption of centered normally distributed shocks, $p(V_t)$ is the density of a multivariate normal and the likelihood function of the state-space model takes the following general form:

$$L(Y_{1:T}^{\text{obs}}|\theta) = (2\pi)^{-\frac{nT}{2}} \prod_{t=1}^T |\Sigma|^{-1/2} \exp\left(-\frac{V_t' \Sigma^{-1} V_t}{2}\right) |J_t|, \quad (24)$$

where n is the number of observed variables stacked in Y^{obs} and Σ is the variance-covariance matrix of V .

A specific model is defined by its vector of error terms V_t which in turn, determines Σ and J_t . Given these three elements we can define the likelihood function.

4.1.2 Model-specific likelihood function

In the remainder of the paper, a model is denoted by letters and numbers respectively indicating the observables and the number of assumed structural shocks. For instance for the model PQ1, the observables are price and production, and there is one structural shock, a supply shock. For PDQ2, the observables are price, consumption, and production, and there are two structural shocks, supply and demand shocks.

For brevity, we demonstrate how to evaluate the likelihood function only for the two most comprehensive specifications with elastic supply and for which the vector of observed variables $Y^{\text{obs}} = (P^{\text{obs}}, D^{\text{obs}}, Q^{\text{obs}})'$ includes price, consumption, and production. These specifications nest the others.⁶

Model PDQ1 Let consider that prices fluctuations are driven solely by unexpected supply disturbances η^Q (i.e., $\eta^D = 0$). In contrast to consumption and production, prices are assumed to be observed without noise. So the vector of errors is $V_t = (\eta_t^Q, \varepsilon_t^D, \varepsilon_t^Q)'$, where ε^D and ε^Q are measurement errors on consumption and production. Following Cafiero et al. (2015), a production shock can be recovered from observation of two consecutive prices using equation (10). So the various elements of V_t are defined by

$$\eta_t^Q = \mathcal{P}^{-1}(P_t^{\text{obs}}) - \frac{\mathcal{S}(\mathcal{P}^{-1}(P_{t-1}^{\text{obs}}))}{1+g} - \mathcal{Q}(\mathcal{P}^{-1}(P_{t-1}^{\text{obs}})), \quad (25)$$

$$\varepsilon_t^D = D_t^{\text{obs}} - D(P_t^{\text{obs}}), \quad (26)$$

$$\varepsilon_t^Q = Q_t^{\text{obs}} - \mathcal{Q}(\mathcal{P}^{-1}(P_{t-1}^{\text{obs}})) - \eta_t^Q = Q_t^{\text{obs}} - \mathcal{P}^{-1}(P_t^{\text{obs}}) + \frac{\mathcal{S}(\mathcal{P}^{-1}(P_{t-1}^{\text{obs}}))}{1+g}. \quad (27)$$

The Jacobian J_t of V_t with respect to Y_t^{obs} is

$$J_t = \begin{pmatrix} \mathcal{P}^{-1'}(P_t^{\text{obs}}) & 0 & 0 \\ -D'(P_t^{\text{obs}}) & 1 & 0 \\ -\mathcal{P}^{-1'}(P_t^{\text{obs}}) & 0 & 1 \end{pmatrix}. \quad (28)$$

The determinant of J_t is equal to $|\mathcal{P}^{-1'}(P_t^{\text{obs}})|$.

We allow the measurement errors ε^D and ε^Q on the consumption and production observations to be correlated with the coefficient ρ^{DQ} , and we denote their standard deviation by σ_{ε^D} and σ_{ε^Q} . This leads to

⁶The other set-up are available upon request.

the following variance-covariance matrix of the vector of errors V

$$\Sigma = \begin{pmatrix} \sigma_{\eta^Q}^2 & 0 & 0 \\ 0 & \sigma_{\varepsilon^D}^2 & \rho^{DQ} \sigma_{\varepsilon^D} \sigma_{\varepsilon^Q} \\ 0 & \rho^{DQ} \sigma_{\varepsilon^D} \sigma_{\varepsilon^Q} & \sigma_{\varepsilon^Q}^2 \end{pmatrix}. \quad (29)$$

Assuming that only price is observable and that supply is inelastic (so $\mathcal{Q}(\mathcal{P}^{-1}(P_{t-1}^{\text{obs}})) = \bar{D}$) leads to the same likelihood function as [Caferio et al. \(2015\)](#).

Model PDQ2 Alternatively, let us consider now that price changes might result also from unexpected shifts in consumption. This augments the number of structural shocks to two. Since having a closed-form likelihood function requires as many shocks as observables, we have to remove one measurement error shock. So we assume that consumption is observed perfectly (i.e., $\varepsilon_t^D = 0$). This leads to the vector of errors $V_t = (\eta_t^Q, \eta_t^D, \varepsilon_t^Q)'$ with elements:

$$\eta_t^Q = \mathcal{P}^{-1}(P_t^{\text{obs}}) - \frac{\mathcal{S}(\mathcal{P}^{-1}(P_{t-1}^{\text{obs}}))}{1+g} - \mathcal{Q}(\mathcal{P}^{-1}(P_{t-1}^{\text{obs}})) + \eta_t^D, \quad (30)$$

$$= \mathcal{P}^{-1}(P_t^{\text{obs}}) - \frac{\mathcal{S}(\mathcal{P}^{-1}(P_{t-1}^{\text{obs}}))}{1+g} - \mathcal{Q}(\mathcal{P}^{-1}(P_{t-1}^{\text{obs}})) + D_t^{\text{obs}} - D(P_t^{\text{obs}}), \quad (31)$$

$$\eta_t^D = D_t^{\text{obs}} - D(P_t^{\text{obs}}), \quad (32)$$

$$\varepsilon_t^Q = Q_t^{\text{obs}} - \mathcal{P}^{-1}(P_t^{\text{obs}}) + \frac{\mathcal{S}(\mathcal{P}^{-1}(P_{t-1}^{\text{obs}}))}{1+g} - D_t^{\text{obs}} + D(P_t^{\text{obs}}). \quad (33)$$

and the Jacobian is:

$$J_t = \begin{pmatrix} \mathcal{P}^{-1'}(P_t^{\text{obs}}) - D'(P_t^{\text{obs}}) & 1 & 0 \\ -D'(P_t^{\text{obs}}) & 1 & 0 \\ D'(P_t^{\text{obs}}) - \mathcal{P}^{-1'}(P_t^{\text{obs}}) & -1 & 1 \end{pmatrix}. \quad (34)$$

As before, the determinant of J_t is equal to $|\mathcal{P}^{-1'}(P_t^{\text{obs}})|$.

Since the structural supply and demand shocks are assumed to be uncorrelated, the variance-covariance matrix of the vector of errors V is given by

$$\Sigma = \begin{pmatrix} \sigma_{\eta^Q}^2 & 0 & 0 \\ 0 & \sigma_{\eta^D}^2 & 0 \\ 0 & 0 & \sigma_{\varepsilon^Q}^2 \end{pmatrix}. \quad (35)$$

4.2 Prior distribution of the parameters

The model parameters are gathered into the vector $\theta \equiv (r, g, \rho^{DQ}, \sigma_{\varepsilon^D}, \sigma_{\varepsilon^Q}, \sigma_{\eta^D}, \sigma_{\eta^Q}, k, \bar{D}, \bar{P}, \alpha^D, \alpha^S)$. Before the estimation, two parameters are fixed: the real interest rate and the rate of growth of consumption and production. The real interest rate is always fixed in estimations of the storage model, since it would be very difficult to estimate it based only on observation of commodity prices. In the literature, it is fixed at either 2% or 5% ([Deaton and Laroque, 1992, 1996](#); [Caferio et al., 2011, 2015](#)). Given that our sample covers the period 1961 to 2006, it seems appropriate to fix the annual real interest rate at 2%. According to the assessment in [Barro and Sala-i Martin \(1990\)](#), for nine OECD countries for which historical data

are available the mean short-term interest rate was 1.87% over the period 1959–1989. Since the sharp rise to rates of about 5% in the 1980s, the world real interest rate started declining to reach an average yearly level of about 2% in the mid 2000s (International Monetary Fund, 2014, Chapter 3). The annual rate of growth of consumption and production is fixed at 2.6%, a value estimated by regressing the log of consumption and production on a time trend.

Apart from these calibrated parameters which will not be updated during the estimation procedure, the prior distributions for the other parameters are shown in table 4. For the standard errors of the shocks, we follow the usual practice in the estimation of DSGE models of using an inverse-gamma distribution which insures positive support, and without information on their size we use loose priors with mean and standard deviations equal to 0.05 for all. The correlation between measurement errors on consumption and production is assumed to be uniformly distributed from -1 to 1 . Data are detrended and centered on 1. Since the model is nonlinear, the steady state is different from the mean, but the location of the mean around 1 is nonetheless an indication that the steady state will be in this region. So the steady-state values of consumption and production are assumed to follow a gamma distribution with mean 1 and standard deviation of 0.32, high enough to cover a reasonable range of parameter values. The storage cost is assumed to follow a Gamma distribution with mean 0.25 and standard deviation 0.18 so that the prior extends from very low to prohibitive costs, covering World Bank and FAO (2012, Figure 2-4) values around 10%–15% of the price for wheat. For the elasticities, the abundant literature presents very diverse results with demand elasticities ranging from -0.03 (Roberts and Schlenker, 2013) to -1 (Adjemian and Smith, 2012), and similarly for supply elasticities. In light of this, elasticities are assumed to follow rather loose Gamma distributions with mean 0.5 and standard deviation 0.35, so that their domain covers all previous estimates.

Table 4: Marginal prior distribution for each model parameter

Parameter	Domain	Density	Hyperparameter (1)	Hyperparameter (2)
ρ^{DQ}	$(-1, 1)$	Uniform	-1	1
σ_{ε^D}	\mathbb{R}^+	Inverse Gamma	0.05	0.05
σ_{ε^Q}	\mathbb{R}^+	Inverse Gamma	0.05	0.05
σ_{η^D}	\mathbb{R}^+	Inverse Gamma	0.05	0.05
σ_{η^Q}	\mathbb{R}^+	Inverse Gamma	0.05	0.05
\bar{D}	\mathbb{R}^+	Gamma	10	0.1
\bar{P}	\mathbb{R}^+	Gamma	10	0.1
k	\mathbb{R}^+	Gamma	2	0.125
α^D	\mathbb{R}^-	Gamma (with transformation to have negative support)	2	0.25
α^S	\mathbb{R}^+	Gamma	2	0.25

Notes: Hyperparameter (1) and Hyperparameter (2) list the upper and lower bounds of the support for the Uniform distribution; the mean and the standard deviation for the Inverse Gamma distribution; and a and b for the Gamma distribution, where $p_G(x|a, b) \propto x^{a-1} \exp(-x/b)$.

4.3 Posterior estimates of the parameters

To maintain a continuum and ease comparability with past empirical storage model studies, we estimate several versions of the storage model from Deaton and Laroque’s simple model with inelastic supply

where only prices are observed, to more complete specifications with supply reactions and inferences involving consumption and production data in addition to prices.

4.3.1 Inelastic Supply

Here we fix the supply elasticity α^S to zero before estimation. The number of observed variables varies between the model specifications enabling analysis of what can be learned from the production and consumption observations, and eventually, which aspects of the global grains market dynamics can be captured by the model. For each model, the estimated mean and standard deviation of the parameter posterior distributions when supply is assumed inelastic are reported in table 5.⁷

Table 5: Bayesian estimates with inelastic supply

Parameter	P1	PD1	PQ1	PDQ1	PD2	PDQ2
ρ^{DQ}	0	0	0	0.6358 (0.1059)	0	0
$\sigma_{\varepsilon\varrho}$	0	0	0.0210 (0.0023)	0.0231 (0.0027)	0	0.0186 (0.0020)
$\sigma_{\varepsilon D}$	0	0.0143 (0.0016)	0	0.0147 (0.0016)	0	0
$\sigma_{\eta\varrho}$	0.0214	0.0264 (0.0095)	0.0179 (0.0028)	0.0189 (0.0029)	0.0333 (0.0100)	0.0219 (0.0027)
$\sigma_{\eta D}$	0	0	0	0	0.0151 (0.0018)	0.0153 (0.0018)
\bar{D}	1	0.9959 (0.0030)	1.0017 (0.0036)	0.9956 (0.0031)	0.9942 (0.0034)	0.9943 (0.0027)
\bar{P}	1.1483 (0.0980)	1.1274 (0.1598)	1.0905 (0.1073)	1.1196 (0.1129)	1.2066 (0.1648)	1.2087 (0.1136)
k	0.0595 (0.0208)	0.0373 (0.0148)	0.0518 (0.0198)	0.0508 (0.0183)	0.0453 (0.0196)	0.0698 (0.0228)
α^D	-0.0297 (0.0085)	-0.0160 (0.0043)	-0.0202 (0.0062)	-0.0197 (0.0052)	-0.0263 (0.0093)	-0.0297 (0.0072)
α^S	0	0	0	0	0	0
P^*	1.3953	2.0281	1.4264	1.5228	1.8866	1.4662
# Stockouts	5	2	3	2	2	3

Notes: The table shows the mean and in parenthesis the standard deviation of the posterior distribution. P^* and the number of stockouts are calculated for the parameters at the mean of the posterior distribution.

The first column, model P1, corresponds to the simplest specification and is the link with existing empirical storage model studies in which inference is made on prices data. In this case, the model is the same as in Deaton and Laroque (1992, 1996) and the likelihood function is the same as in Cafiero et al. (2015). In addition to assuming that supply is inelastic, inference resting on prices alone requires more parameters to be fixed before the estimation to allow the other parameters to be identified. Deaton and Laroque (1996, Proposition 1) show that when only prices are observed it is not possible to identify the demand function separately from the supply shocks. So we fix the steady-state value of quantities \bar{D} at unity, and the standard deviation of the supply innovation $\sigma_{\eta\varrho}$ at 0.0214, the standard deviation of the detrended production which amounts to assuming that all observed variations in production are explained

⁷Results based on the median, and the 5th and 95th percentiles of the posterior distribution are similar.

by structural shocks. The storage cost k is found to be equal to 5.2% of the steady-state price, a value much higher than the estimations in Cafiero et al. (2011, 2015) but close to those found by Gouel and Legrand (2015). The posterior mean of the demand elasticity, α^D , is -0.0296 . This is a very low demand elasticity but is in line with Roberts and Schlenker's (2013) estimates which are in the range of -0.014 to -0.066 for the same index of grains calories. Furthermore, it is consistent also with the fact that once aggregated, these staple products have very few, if any, substitution possibilities. Finally, the market is estimated to enter the stockout regime five times over the sample length of 46 years. This is in line with the intuition that stockouts are rare events corresponding to periods when prices spike. In our sample, the stockouts correspond (from the highest to the lowest detrended prices) to the peaks in grains prices in 1973 and 1974, to the price increase at the end of the sample in 2006, to 1980 and 1975.

In the second column, the model PD1 includes consumption as an observable. This allows the parameters of the supply process to be identified separately from those of the demand function, meaning that the parameters \bar{D} and $\sigma_{\eta\varrho}$ are now allowed to be estimated. \bar{D} is precisely estimated at close to 1, which is consistent with the fact that in this storage model the steady-state consumption should be close to the asymptotic mean consumption. The standard deviation of the structural supply shock is slightly above the volatility of production observed in the data (table 2). Compared to the specification P1, the demand function has a much steeper slope at -0.016 . Combined with a lower storage cost estimated at 3.31% of the steady-state price, the inferred cutoff price P^* over which inventories are sold out is 45% higher, which decreases the estimated number of stockouts now limited to 1973 and 1974. The standard deviation of the measurement error shock on consumption is of the same size as the volatility of consumption observed in the data, so the storage model explains little of the variation in consumption. This is expected given the low correlation observed between consumption and price.

In the specification PQ1, we use production as the observable rather than consumption, and assume only one structural shock. Compared to the specification PD1, it leads to a lower standard deviation of the posterior for $\sigma_{\eta\varrho}$, consistent with the fact that production is more directly informative than consumption about the production shock. Compare to PD1, demand is slightly more elastic to price, and storage is more costly (4.7% of the estimated steady-state price) but the differences are mostly all within the ranges of the standard deviations of the estimates, except for $\sigma_{\eta\varrho}$. Nonetheless, these differences lead to a lower cutoff price P^* implying the occurrence of a third stockout in 2006.

The specification PDQ1 includes all available observables. It leads to estimates that are close to the specifications PD1 and PQ1. This specification confirms that half of observed production fluctuations can be explained by structural shocks with the other half accounted for by measurement errors. As for PD1, the measurement error on consumption is substantial and would account for almost all observed consumption changes. There is a significant correlation coefficient of the measurement error on consumption and production: $\rho^{DQ} = 0.64$. This is consistent with the way these data are collected, with reasonably thorough surveys of production, and estimates for consumption that partly depend on production estimates.

We turn next to specifications PD2 and PDQ2 which feature two structural shocks: one on consumption and one on production. In these cases, a closed-form likelihood function with measurement errors on consumption is not possible. We assume that consumption is perfectly observed, and is the sum of a deterministic demand function and a structural shock. If we compare the specifications PD2 and PD1, what is captured in PD1 in the measurement error on consumption is accounted for now by the structural shock. Surprisingly, this leads also to a much larger structural production shock with a posterior mean for $\sigma_{\eta\varrho}$ of 0.0333, an estimate which far exceeds the observed volatility in production. The estimation of $\sigma_{\eta\varrho}$

is lower under the specification PDQ2 which uses production as an observable. Overall, the specifications allowing for two structural shocks lead to estimated net supply shocks ($\eta^Q - \eta^D$) larger in size than their counterparts with one shock (PD1 and PDQ1) which implies a more elastic demand function, closer to Roberts and Schlenker's estimates.

4.3.2 Elastic Supply

So far the literature on estimation of storage models has focused on models with inelastic supply because prices were the only observables; this does not permit estimation of demand and supply elasticities. By using consumption and production as additional observables we can estimate models with elastic supply.⁸

Theoretically, speculative storage and elastic production show interesting interactions.⁹ In the absence of storage, the model collapses to Muth's (1961) model where planned production is constant and equal to its steady-state value because the expected price is constant. As soon as speculative storage operates, the expected price and planned production varies. A high stock level weights on the expected prices, and by extension, on the expected profits from production, along with the level of effort invested in new production. Conversely, low stock levels lead to high expected prices and high levels of planned production, with an upper limit attained when stocks are null. Because production adjusts to the expected price and helps alleviate future scarcities, stock levels are less responsive to current availability. For various calibrations, Wright and Williams (1982) show that a model with elastic supply leads to lower price volatility than the same model with inelastic supply. This means that with elastic supply, lower stock levels will be required to fit the same price volatility in the data. This translates to the estimation (see table 6) of higher storage costs and lower cutoff prices than for the corresponding inelastic specifications. For the PDQ1 and PDQ2 specifications this leads to five stockouts, in 1973–5, 1980, and 2006, while in the PQ1 specification there is a sixth stockout in 1979.

For the same specifications, the estimations with elastic supply are similar to the estimations with inelastic supply for almost all parameters except the steady-state price and the storage cost; both of which increase significantly. Expressed as a percentage of the steady-state price, storage costs are of the order of 14%, a value consistent with the World Bank and FAO (2012). Supply elasticities are estimated to be around 0.05, so supply is twice as elastic as demand. These supply elasticities are in the lower range of the values in the literature. Compared to Roberts and Schlenker (2013), they are twice as low. Roberts and Schlenker (2013) estimate the supply elasticity as being between 0.085 and 0.116.

While most parameters are precisely estimated, the storage cost and supply elasticity are estimated with limited precision. The standard deviations of the posterior distribution of storage cost and supply elasticities are close to 43% and 53% of their respective means. Storage cost is the least precisely estimated parameter in the inelastic case, which can be explained by the fact that beyond very small values the likelihood function is quite flat with respect to storage cost. However, the standard deviation of k is much higher in the elastic case. This likely stems from the fact that storage and elastic production are substitutes for explaining price volatility. The observation of production should help choose between the two processes but given the size of the measurement error on production, its informativeness is probably limited which explains the lack of precision of the estimates.

⁸The sole observation of consumption is not informative enough to deliver sound estimates when the model has elastic supply. Hence, the PD1 and PD2 specifications are no longer considered in the subsequent estimations.

⁹See Wright and Williams (1982) for a thorough description of the effects of modeling supply responsiveness in the dynamics implied by the model.

Table 6: Bayesian estimates with elastic supply

Parameter	PQ1	PDQ1	PDQ2
ρ^{DQ}	0	0.6192 (0.1115)	0
$\sigma_{\varepsilon Q}$	0.0209 (0.0023)	0.0231 (0.0028)	0.0189 (0.0020)
$\sigma_{\varepsilon D}$	0	0.0145 (0.0016)	0
$\sigma_{\eta Q}$	0.0189 (0.0029)	0.0194 (0.0029)	0.0224 (0.0029)
$\sigma_{\eta D}$	0	0	0.0150 (0.0017)
\bar{D}	0.9978 (0.0038)	0.9921 (0.0032)	0.9903 (0.0029)
\bar{P}	1.3616 (0.1356)	1.4016 (0.1381)	1.4301 (0.1404)
k	0.2045 (0.0887)	0.1918 (0.0824)	0.1943 (0.0823)
α^D	-0.0218 (0.0065)	-0.0198 (0.0055)	-0.0300 (0.0077)
α^S	0.0560 (0.0290)	0.0507 (0.0270)	0.0492 (0.0285)
P^*	1.2811	1.3804	1.3713
# Stockouts	6	5	5

Notes: The table shows the mean and in parenthesis the standard deviation of the posterior distribution. P^* and the number of stockouts are calculated for the parameters at the mean of the posterior distribution.

Figure 3 summarizes the estimation results visually by plotting the prior distribution, and the posterior distributions of the inelastic and elastic PDQ2 models. It confirms that the posterior distributions are the same in the inelastic and elastic cases, except for \bar{P} and k . All parameters are estimated to be significantly different from zero. Except for the storage cost in the elastic case, the posterior distributions are different from the prior distributions. The similarity of the prior and posterior for storage cost if the model features elastic supply confirms a potential problem of identification.

5 Conclusion

This paper proposed a new estimation strategy for the rational expectations storage model. In addition to prices, it adds to the observable variables consumption and production. This allows identification parameters that are not identifiable if the inference rests on prices alone, and enables estimation of a richer storage model with elastic supply and two structural shocks. The estimation method takes its inspiration from recent developments in the estimation of DSGE models. It uses Bayesian econometric methods which are more suited to these types of dynamic models than frequentist methods, which are more likely to experience difficulty to find a global maximum. To ease comparability with the results in the literature, we started the estimations with a specification close to [Cafiero et al. \(2015\)](#), with inelastic supply and prices as the only observables. We introduced more observables gradually, and relaxed the assumption of inelastic supply. The variety of model specifications estimated allowed us to identify the contribution of

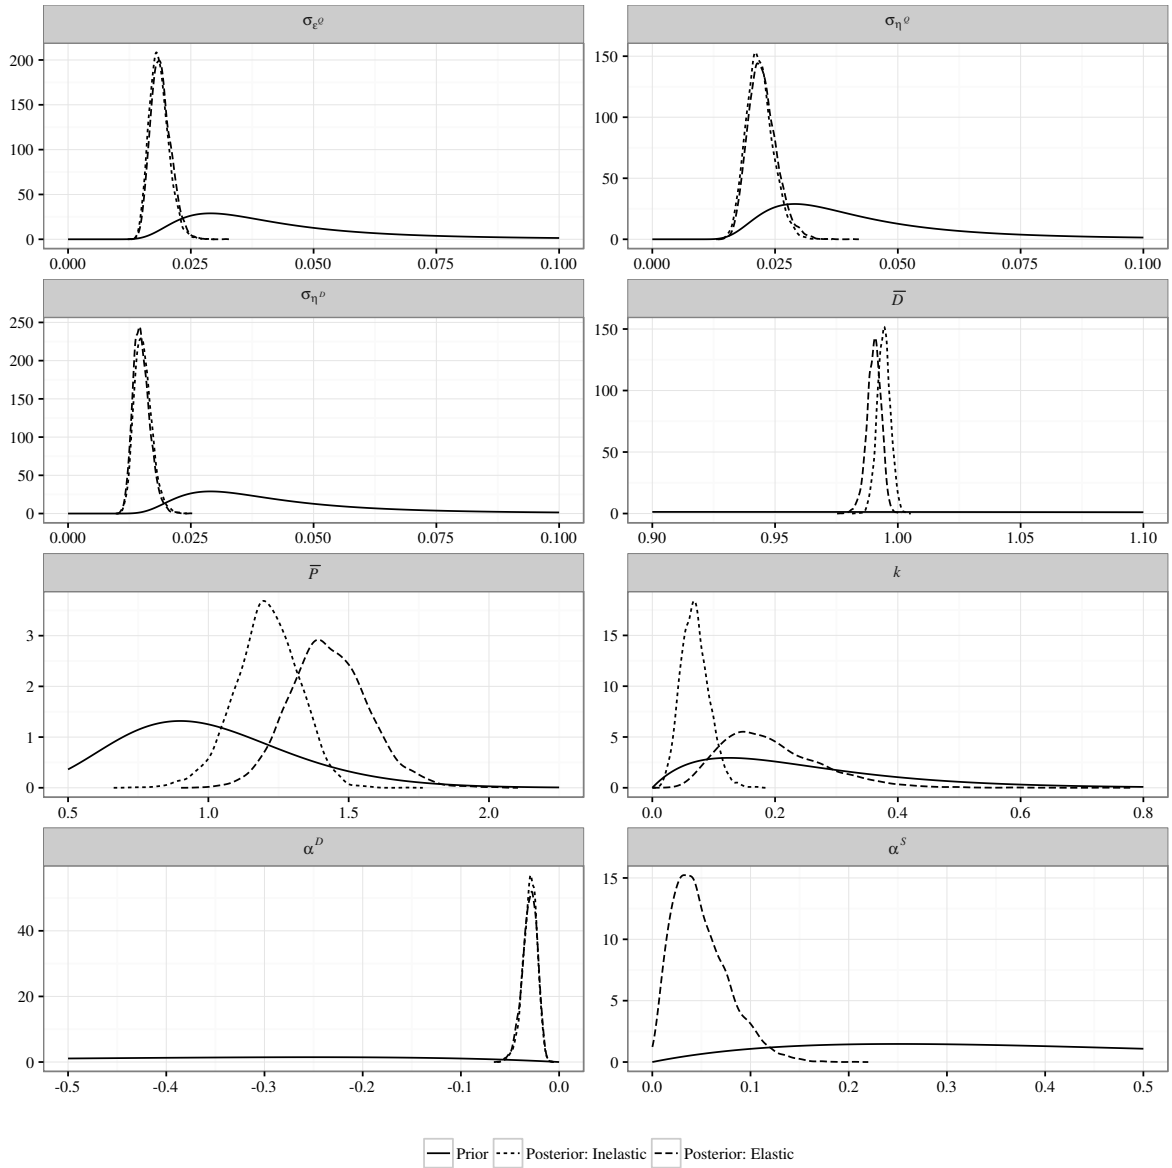


Figure 3: Estimated parameter distribution for the specification PDQ2

each set of observables to the results. In order to have a benchmark against which to compare our results, the model was estimated on the same set of data as in [Roberts and Schlenker \(2013\)](#): the market for the caloric aggregate of maize, rice, soybeans, and wheat from 1961 to 2006.

In our estimations, elasticities are in the lower range in the literature: -0.02 to -0.03 for the demand elasticity and 0.05 for the supply elasticity. These estimates are slightly lower than but consistent with [Roberts and Schlenker's](#) estimates, and other estimates based on the storage model (e.g., [Gouel and Legrand, 2015](#); [Guerra et al., 2015](#)). These low values were expected given the low volatility of observed consumption and production. Storage costs are estimated to be between 4% and 14% of the steady-state price. Storage was shown to be more costly when supply is elastic, because elastic supply and speculative storage have similar effects on price dynamics. The importance of storage decreases when supply is allowed to be elastic and the number of stockouts increases. These preliminary results tend to confirm

that our Bayesian approach is a reliable econometric method for structural estimations of the storage model.

In this preliminary version of the paper we made assumptions to permit the likelihood function of the storage model to take an analytic form. This involves severe restrictions: unlike quantities prices were assumed to be observed without noise, and it was not always possible to estimate measurement errors and structural shocks jointly. These restrictions will be lifted in the final version of the paper which will present estimations with more general specifications.

The new estimation methods for the storage model proposed in this paper constitute an important innovation; however one of their main benefits may be the new perspectives they open for this model. Deaton and Laroque (1992, 1996) rejected the storage model for its inability to fit the serial correlation in prices. This problem has been solved in recent works but with an estimation method which allows us to use all observables to fit the model, it is likely that the storage model's ability to fit the behavior of the observables will decrease with the number of observables. The dimensions where the model fails will open the way to new theoretical developments to improve the model fit. We have seen similar developments in the econometrics of DSGE models where extra elements have been added to the basic RBC model to improve its fit (Fernández-Villaverde, 2010). In the case of our estimations, the standard deviations of the measurement errors on consumption and production are large and could explain most the observed variations. So the current specification of the storage model is not able to account well for the dynamics of quantities, and should be extended with this shortcoming in mind.

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.
- Barro, R. J. and Sala-i Martin, X. (1990). World real interest rates. In O. J. Blanchard and S. Fischer (eds.) *NBER Macroeconomics Annual 1990*, volume 5, chapter 1, (pp. 15–74). Cambridge, MA: MIT Press.
- Cafiero, C., Bobenrieth, E. S. A., Bobenrieth, J. R. A. and Wright, B. D. (2011). 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.
- Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3), 312–320.
- 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.
- Dvir, E. and Rogoff, K. (2014). Demand effects and speculation in oil markets: Theory and evidence. *Journal of International Money and Finance*, 42, 113–128.
- Fernández-Villaverde, J. (2010). The econometrics of DSGE models. *SERIEs*, 1(1–2), 3–49.

- Fernández-Villaverde, J., Rubio-Ramírez, J. F. and Schorfheide, F. (forthcoming). Solution and estimation methods for DSGE models. In J. Taylor and H. Uhlig (eds.) *Handbook of Macroeconomics*, volume 2. Amsterdam: Elsevier.
- Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith (eds.) *Bayesian Statistics 4: Proceedings of the Fourth Valencia International Meeting*, New York: Oxford University Press, (pp. 169–193).
- de Gorter, H., Drabik, D. and Just, D. R. (2015). *The Economics of Biofuel Policies: Impacts on Price Volatility in Grain and Oilseed Markets*. New York: Palgrave macmillan.
- Gouel, C. (2013). Comparing numerical methods for solving the competitive storage model. *Computational Economics*, 41(2), 267–295.
- Gouel, C. and Legrand, N. (2015). *Estimating the Competitive Storage Model with Trending Commodity Prices*. Working Paper 2015-15, EconomiX.
- Guerra, E. A., Bobenrieth H., E. S. A., Bobenrieth H., J. R. A. and Cafiero, C. (2015). Empirical commodity storage model: the challenge of matching data and theory. *European Review of Agricultural Economics*, 42(4), 607–623.
- 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.
- International Monetary Fund (2014). *World Economic Outlook: Recovery strengthens, Remains uneven*. Washington, DC: International Monetary Fund.
- Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica*, 29(3), 315–335.
- Robert, C. and Casella, G. (2004). *Monte Carlo Statistical Methods*. New York: Springer-Verlag.
- 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.
- Vihola, M. (2012). Robust adaptive Metropolis algorithm with coerced acceptance rate. *Statistics and Computing*, 22(5), 997–1008.
- Williamson, L. and Williamson, P. (1942). What we eat. *Journal of Farm Economics*, 24(3), 698–703.
- World Bank and FAO (2012). *The Grain Chain: Food Security and Managing Wheat Imports in Arab Countries*. Washington, DC: World Bank.
- Wright, B. D. (2014). Global biofuels: Key to the puzzle of grain market behavior. *Journal of Economic Perspectives*, 28(1), 73–98.
- Wright, B. D. and Williams, J. C. (1982). The economic role of commodity storage. *The Economic Journal*, 92(367), 596–614.