



HAL
open science

The Role of Storage in Commodity Markets: Indirect Inference Based on Grains Data

Nicolas Legrand, Christophe Gouel

► **To cite this version:**

Nicolas Legrand, Christophe Gouel. The Role of Storage in Commodity Markets: Indirect Inference Based on Grains Data. 2022. hal-03809825

HAL Id: hal-03809825

<https://hal.inrae.fr/hal-03809825>

Preprint submitted on 10 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Role of Storage in Commodity Markets: Indirect Inference Based on Grains Data

Christophe Gouel & Nicolas Legrand

Highlights

- Understanding the drivers of commodity prices dynamics is crucial. Unfortunately the central economic model for representing commodity prices, the competitive storage model, is not yet empirically validated.
- In this work, we develop a rich storage model with four demand and supply shocks, elastic supply, and long-run trends and estimate it structurally on a caloric aggregate of the four most important grains.
- Our estimated model is consistent with most of the moments in the data, validating the empirical relevance of the storage model.
- The estimated model shows that speculative storage while crucial cannot explain alone the persistence of grain price. It explains 42% of the price autocorrelation, the rest being accounted for by the price trend, the persistence of demand shocks, and the presence of news shocks about production.



Abstract

Understanding commodity prices dynamics is of crucial importance for assessing the persistence of cost-push costs or for countries dependent on commodity exports. Unfortunately, despite decades of research, the workhorse theoretical model in the field, the rational expectations storage model, is yet to be empirically validated. This paper provides the first full empirical test of the storage model. We first build a new storage model featuring a supply response, long-run demand and cost trends, and four structural shocks. We then develop a flexible empirical approach which relies on the indirect inference method and exploits the joint dynamics of prices and quantities unlike previous estimations which only use price information. The information contained in quantities is essential to relax restrictive identifying assumptions and empirically assess the overall consistency of the model's new features. Finally, we carry out a structural estimation on the aggregate index of the world most important staple food products: maize, rice, soybeans, and wheat. The results show that our extended storage model is consistent with most of the moments in the data, including the high price autocorrelation of which up to 42% can be explained by the transfer of inventories over time. They also show that, although for these commodities supply shocks are the main drivers of market dynamics, over the past 60 years all price spikes have been associated with large positive demand shocks.

Keywords

Commodity Price Dynamics, Indirect Inference, Monte Carlo Analysis, Storage.

JEL

C51, C52, Q11

Working Paper

CEPII

CEPII (Centre d'Etudes Prospectives et d'Informations Internationales) is a French institute dedicated to producing independent, policy-oriented economic research helpful to understand the international economic environment and challenges in the areas of trade policy, competitiveness, macroeconomics, international finance and growth.

CEPII Working Paper
Contributing to research in international economics

© CEPII, PARIS, 2021

All rights reserved. Opinions expressed in this publication are those of the author(s) alone.

Editorial Director:
Sébastien Jean
Production: Laure Boivin
Published on 21.07.22

No ISSN: 1293-2574

CEPII
20, avenue de Ségur
TSA 10726
75334 Paris Cedex 07
+33 1 53 68 55 00
www.cepii.fr
Press contact: presse@cepii.fr

RESEARCH AND EXPERTISE
ON THE WORLD ECONOMY



The Role of Storage in Commodity Markets: Indirect Inference Based on Grains Data¹

Christophe Gouel* and Nicolas Legrand†

1. Introduction

Many key policy decisions depend on a good understanding of commodity price dynamics. This is the case for example of how monetary policy should react to commodity-driven cost-push shocks or how commodity-exporters should deal with the windfalls from commodity price spikes. These questions crucially depend on the persistence of commodity prices. Unfortunately, the workhorse model supposed to explain this persistence, the rational expectations storage model which is used in many applied and policy works (Gouel, 2013b; Porteous, 2019; Steinwender, 2018), is far from being empirically validated. Its first estimations by Deaton and Laroque (1992, 1996) show that the model is able to account qualitatively for many of the stylized facts: commodity prices have a high serial correlation, are highly volatile, their distribution is positively skewed, and they display a succession of doldrums, boom and bust cycles. By buying low and selling high, the speculators tend to smooth shocks, prevent price troughs, and link together current and expected prices which induces autocorrelation in prices. However, Deaton and Laroque (1996) found that the storage model while able to generate autocorrelation in prices is not able to match the level observed in the data. Subsequent work offered some solutions to raise the persistence induced by the model and better match this central feature of the data (Cafiero et al., 2011, 2015; Gouel and Legrand, 2017), but all these studies build on Deaton and Laroque's approach where the model is estimated only on prices, which requires restrictive identifying assumptions, prevents estimation of all the model parameters, and have typically involved an all-or-nothing approach to the autocorrelation issue. Indeed, by forcing a simple storage model to explain all the observed price persistence by the transfer of inventories over time, either shocks need to be autocorrelated such that all the persistence comes from the shocks themselves and storage is irrelevant (Deaton and Laroque, 1996); or shocks are assumed i.i.d. and all the persistence is generated by storage which requires negligible storage costs (Cafiero et al., 2011). This seriously limits the model's usefulness in studying how price fluctuations are driven by the underlying shifts in supply and demand, and in assessing the respective importance of these shocks.

¹The authors are grateful to Alain Carpentier, Dan Cao, Joe Janzen, and John Rust; and to seminar audiences at the RCEA Time Series Workshop, Georgetown University, NCCC-134 2021 conference, Penn State, SMART, University of California Davis, and University of Illinois Urbana-Champaign for helpful comments. This work was supported by the INRAE ECOSOCIO department.

*INRAE and CEPII: christophe.gouel@inrae.fr

†INRAE: nicolas.legrand@inrae.fr

In this work, we build and estimate a rich rational expectations storage model with the purpose of examining its empirical validity beyond the ability to fit price dynamics. To do this, we depart from the standard model set-up estimated so far. Specifically, we first extend the simple storage model to include: (i) a supply response, (ii) long-run trends in prices and quantities, (iii) a persistent demand shock, and (iv) three supply shocks with different timings. Next, we show how to exploit the information contained in the joint dynamics of quantity and prices to identify all the structural parameters of the model. Last, we take our richer storage model to five time series of the global grains market represented by an aggregate index of the world most important staple food products.² We find that our model is able to generate the observed high price autocorrelation and the transfer of inventories over time explains 42% of it, the rest being explained by the other model features. We also show that, once fully-specified, the storage model is able to fit quite well the main moments of the global food market. Importantly, with our econometric strategy that exploits the joint dynamics of price and quantity, we can empirically assess the overall consistency of the model's combined extensions while identifying more formally which ones help to match the moments in the data.

The construction of our model was guided by the following considerations. Estimating a supply and demand model presents the usual problem of simultaneity bias with equilibrium price and quantity that are jointly determined. Correct identification in this setting requires accounting for unobservable shifts in each curve. Considering this, we build on the recent innovation in this literature by Roberts and Schlenker (2013) who use the storage theory to find an appropriate instrument to estimate supply elasticities in storable commodity markets. While storage theory inspires their econometric strategy to identify demand and supply elasticities, they do not develop a storage model consistent with their strategy. In contrast, we introduce in our model various demand and supply shocks, with heterogeneous timing, guided by both the theoretical structure implicit behind Roberts and Schlenker's instrumental variable estimation approach and by the moments in the data.

Despite the richness of our model compared to most models in the storage literature, it remains quite stylized, and particularly compared to the number of observables. More precisely, with only four shocks driving the fluctuations of five observables, the model presents a stochastic singularity. This is an obstacle to a likelihood-based estimation since, by construction, the model could not be expected to account for the richness of the data. To deal with this stochastic singularity issue, we adopt an estimation approach that can be applied despite the singularity, which allows us to choose the dimensions of the data to match, and which remains fully transparent with respect to the factors driving the estimation. This approach is the indirect inference proposed in Gourieroux et al. (1993) and Smith (1993). It is a simulated moment-based method in which the model is estimated by targeting parameter estimates from an auxiliary model. Put simply, indirect inference is based on the use of an auxiliary model as a statistical model which provides a rich description of the features in the data. This auxiliary model, which here is the supply

²These are maize, rice, soybeans, and wheat.

and demand model of Roberts and Schlenker (2013), is estimated on both the true data and on simulated data from the structural model, and the structural model parameters are adjusted to minimize the distance between both sets of estimates from the auxiliary model. This approach allows us to exploit an econometric literature where intuitions about which moment is driving a parameter estimation are more explicit than full-information techniques. Also interesting with this approach is that it can be applied in the absence of information about stocks which are generally not available or too noisy to be of use.

We apply this indirect inference approach on the data used by Roberts and Schlenker (2013), which includes five observed variables: price, expected price, demand, production, and yield shock. This allows us to estimate all the parameters of the model. Using these estimates, we present four sets of results. First, a credible solution to the autocorrelation puzzle can be found by accounting for sufficient features of the international grains market. Using our benchmark estimations, we show that 42% of the observed one-year autocorrelation can be explained by storage, a third by a long-run trend in prices, 20% by autocorrelated demand shocks, and the final 5% by the shocks on supply.

Second, we evaluate the ability of our extended storage model to capture the empirical time series properties of both price and quantity data. As is typically done in the literature of dynamic stochastic general equilibrium (DSGE) models, we assess the performance of the estimated storage model by comparing the covariances based on model simulations and those based on observations. Generally, the covariances are similar for simulations and observations, suggesting that the model is able to mimic the main moments in the data. Interestingly, our results raise a new puzzle: the model proves unable to match the correlations between price and quantities, consumption as well as production, which are much lower in the data than in the model.

Third, we use the estimated model to analyze the sources of commodity price movements in the global market of grains. In this respect, we show that for these commodities supply shocks are the main drivers of market dynamics with an aggregated standard deviation 30% larger than for the aggregate demand shocks. However, over the past six decades, all price spikes have been associated with large positive demand shocks.

Fourth, we revisit Wright and Williams's (1984) study of the welfare effects of the introduction of storage. Unlike them and the rest of the literature which use various calibrations in the absence of credible estimations of the model parameters, our welfare analysis relies on a fully estimated version of the model. We document substantial distributional effects with a 4% increase in consumer surplus following the introduction of storage. However, because this change is related mostly to a reduction in the mean price, the corresponding decrease in the producer profit leads to modest welfare effects overall.

Our work relates to three strands of research. The first strand studies the theoretical and empirical properties of the storage model. Our storage model builds on earlier studies that introduce similar features separately. For example, Wright and Williams's (1982) competitive

storage model includes an elastic supply. Autocorrelated shocks were introduced by Chambers and Bailey (1996), Deaton and Laroque (1996), and Routledge et al. (2000). Production shocks with different timings have been used in several papers (e.g., Lowry et al., 1987; Osborne, 2004; Gouel, 2020). Dvir and Rogoff (2014) develop a storage model with trending quantities and Bobenrieth et al. (2021) introduce a supply trend which in turn generates quantity and price trends. Relative to this literature, our use of information on both price and quantities enable us to disentangle the effects of the core storage theory from the set of auxiliary assumptions needed for inference. Indeed, along three dimensions—the persistence of the demand shock, the supply elasticity, and the size and cross-correlation of the supply shocks—the dynamics of quantities play a critical role because price data alone cannot identify any of them.

The second strand is a literature that uses structural vector autoregressions (VAR) to identify supply and demand shocks in commodity markets (e.g., Kilian and Murphy, 2014; Carter et al., 2017; Baumeister and Hamilton, 2019). Our paper provides one of the first fully structural approach in the commodity price literature allowing to identify the various shocks in a theoretically consistent way (another paper doing it with a structural model, but without storage and for the oil market, is Bornstein et al., forthcoming).

Last, our approach bridges two literatures: the literature on the estimation of the storage model and the literature on the estimation of DSGE models, which conceptually and numerically are close to the storage model. The estimation of storage models has been so far restricted to small models too stylized to capture the richness of these markets. This was also the case for DSGE models up to the contributions of Smets and Wouters (2003, 2007), who show how to build and estimate DSGE model with rich stochastic structures. We follow Smets and Wouters by adding a rich set of structural shocks to the storage model. Our estimation approach borrows also from the DSGE literature where indirect inference is commonly applied.³ In this literature, the auxiliary model is often structural VARs and the estimations depend on targeting the impulse responses (e.g., Rotemberg and Woodford, 1997; Christiano et al., 2005; Ruge-Murcia, 2020). In our case, we show that a system of linear equations based on the instrumental variable model in Roberts and Schlenker (2013) is enough to capture all the dynamic relationships of interest (as in Guvenen and Smith, 2014), and in particular the strong nonlinearities.⁴ However, a VAR should also work since Ghanem and Smith (2022) adapted in a VAR Roberts and Schlenker's model, which provides the basis for our auxiliary model.⁵

The rest of the paper is as follows. Section 2 describes the storage model. Section 3 presents

³This paper is not the first to estimate a storage model by indirect inference. Michaelides and Ng (2000) employed this approach in a Monte Carlo comparison of simulation estimators. However, as Michaelides and Ng (2000) followed Deaton and Laroque by estimating the model only on prices, the various auxiliary models they consider are all based on univariate time-series models.

⁴Since commodities cannot be consumed before being produced, there is a nonnegativity constraint on inventories. This zero lower bound on storage introduces an essential nonlinearity which carries through into nonlinearity of the predicted commodity price series.

⁵See also Carter et al. (2017) for a VAR estimation of a storage model applied to the maize market.

the econometric strategy which starts by deriving the instrumental variable approach consistent with the model followed by the indirect inference approach. The short-sample properties of these estimation approaches are studied using Monte Carlo simulations in section 4. Section 5 describes the data and gives descriptive statistics. Section 6 discusses the estimation results and assesses the model fit on moments not included in the estimation. Based on the model estimated, section 7 analyzes the role of storage in price dynamics and welfare, and studies the contribution of the various shocks to the market dynamics. Section 8 concludes the paper.

2. The model

This section presents the storage model to be estimated. Although the storage model is used to explain short-run dynamics in commodity markets, long-run dynamics can potentially affect short-run incentives and should not be neglected in the model. Consumption and production of food increase over time due to rising population numbers, income growth, and technological progress. There is a large literature analyzing the nature of the long-run trends in commodity prices (see section 5.2). To account for these long-run dynamics, we allow both the demand and marginal cost functions to have trends, which in turn translate into quantity and price trends. However, for simulation purposes, the storage model must be a stationary model. Therefore, we first present the storage model with trends, and second we express it in terms of the detrended variables, which shows how the trends affect agents' incentives.

2.1. Nonstationary model

Producers A representative producer makes its production decision and pays for inputs one period before bringing its output to the market. The production choice represented by the acreage is made in period t and denoted H_t . The producer decision is affected by two shocks: η_t , a planting-time yield shock, and ω_t , a cost shock. The planting-time yield shock represents the component of yield shock that is observable by the producer when planting, for example related to the field-conditions during planting, the groundwater level, and the seasonal weather forecasts. Roberts and Schlenker (2013) take also the example of the soybean rust which is observable from the previous growing season. The cost shock aggregates a variety shocks, for example related to fertilizers, seeds, labor, and fuel. Realized production differs from planned production because of an unpredictable harvest-time yield disturbance denoted ϵ_{t+1} . The shocks are normal with zero mean and no autocorrelation, and their respective variances are σ_η^2 , σ_ω^2 , and σ_ϵ^2 .

Although in reality, planting-time and harvest-time yield shocks may be correlated, because of the rational expectations assumption there is no need to introduce in the model a correlation between η_t and ϵ_{t+1} . If producers are efficient forecasters (in the sense of Nordhaus, 1987), they will account for the existing correlation and their forecasting errors should be independent of the observables at period t . In other words, ϵ_{t+1} can be interpreted as the yield forecast error

at planting time, which because of rational expectations must be uncorrelated to any period- t variable.

We cannot exclude the possibility of a correlation between the two planting-time shocks, η_t and ω_t , since a year with low yield prospects, for example, could be associated also with higher marginal costs to achieve the same level of production. Therefore, we assume they are correlated with a coefficient $\rho_{\eta,\omega} \in (-1, 1)$.

The producer's problem in period t can be written as

$$\max_{H_t \geq 0} \beta E_t (P_{t+1} H_t e^{\eta_t + \epsilon_{t+1}}) - \Gamma_t (H_t) e^{\omega_t + g_p t}, \quad (1)$$

where $0 < \beta < 1$ is the annual discount factor which is assumed to be fixed, E_t is the expectation operator conditional on period t information, P_{t+1} is the price, $\Gamma_t(\cdot)$ is a nonstationary, differentiable, and convex production cost function, and g_p is the price trend which appears as a production cost trend. The solution to this problem is given by the following first-order condition

$$\beta e^{\eta_t} E_t (P_{t+1} e^{\epsilon_{t+1}}) = \Gamma_t' (H_t) e^{\omega_t + g_p t}. \quad (2)$$

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

From an econometrics perspective, we assume that only the combined yield shock is observable and that it is not possible to observe η_t and ϵ_{t+1} separately. We therefore introduce $\psi_{t+1} = \eta_t + \epsilon_{t+1}$ as the observable yield shock. Final production $Q_{t+1} = H_t \exp(\psi_{t+1})$, is also observable in publicly available statistics. Note that assuming a multiplicative cost shock separable from the other costs implies that this shock can be moved to the left-hand side of equation (2) where it would play the same role in final production as the planting-time yield shock, the only difference being that the yield shock is observable with noise ex-post in ψ_{t+1} but not the cost shock. Since ω_t can be moved to the left-hand side, this means it might capture also some incentive shocks (e.g., because of changes to agricultural and trade policies or because of price changes in competing crops).

Storers For the storage sector, we assume free entry, competitive behavior, and risk-neutrality. To store an amount $X_t \geq 0$ from period t until $t + 1$ competitive storers incur several costs. They incur an opportunity cost because they have to buy one period before being able to sell. Following most of the storage literature (Gustafson, 1958; Steinwender, 2018; Wright and Williams, 1982, 1984), we assume that storers incur a physical cost of storage proportional to the stored quantity, $k \bar{P}_t X_t$, where \bar{P}_t is the price on the growth path (i.e., in the absence of shocks) and $k \geq 0$ is the per-unit physical storage cost expressed as a percentage of this price. To be compatible with a model that ultimately could be expressed in terms of stationary variables, the per-unit storage cost must be assumed either to be null (the assumption adopted

in Bobenrieth et al., 2021) or as adopted here to follow the same trend as the price. Finally, following Deaton and Laroque (1992, 1996), we assume that because of deterioration stocks shrink by a proportion $\delta \in [0, 1)$ every period.⁶ Storage technologies impose a trade-off between physical storage cost and deterioration. Cold storage in dedicated facilities could result in almost no shrinkage but high costs, while piling bags under a tarp would involve limited costs but high shrinkage. Representing both types of costs allows to estimate their share in global storage costs.

Under this structure of costs and the assumption of rational expectations, the representative storer maximizes its expected profit,

$$\max_{X_t \geq 0} E_t \{ [\beta(1 - \delta) P_{t+1} - P_t - k\bar{P}_t] X_t \}, \quad (3)$$

which taking account of the non-negativity constraint on storage yields the following arbitrage condition

$$\beta(1 - \delta) E_t P_{t+1} - P_t - k\bar{P}_t \leq 0, \quad = 0 \text{ if } X_t > 0. \quad (4)$$

When the expected price is too low to cover the purchase and storage costs (i.e., $\beta(1 - \delta) E_t P_{t+1} \leq P_t + k\bar{P}_t$), no stocks are held. Conversely, when the expected price covers the purchase and storage costs, stocks are acquired up to the level where the expected marginal profit is null: $\beta(1 - \delta) E_t P_{t+1} = P_t + k\bar{P}_t$, which involves an intertemporal relationship between current and expected prices.

Total marginal storage costs equal $k\bar{P}_t - [\beta(1 - \delta) - 1]P_{t+1}$, which shows that a key difference between per-unit storage costs and the two other costs lies in the fact that opportunity and deterioration are storage costs that rise with price level.

Final demand Final demand for the good is the product of a downward sloping demand function $D_t(P_t)$ with a demand shock, $\exp(\mu_t)$, where μ_t follows a first-order autoregressive process with autocorrelation $\rho_\mu \in [0, 1)$ and innovation $v \sim \mathcal{N}(0, \sigma_v^2)$:

$$\mu_{t+1} = \rho_\mu \mu_t + v_{t+1}. \quad (5)$$

Equilibrium The market clears when the sum of previous remaining stocks and production equals the final demand for immediate consumption plus the speculative demand for stocks:

$$(1 - \delta) X_{t-1} + H_{t-1} e^{\eta_{t-1} + \epsilon_t} = D_t(P_t) e^{\mu_t} + X_t. \quad (6)$$

⁶We do not consider the possibility of an upper bound on storage capacities (Oglend and Kleppe, 2017) because, contrary to oil and gas, grains can be stored outside dedicated facilities.

2.2. Stationary model

Detrended variables and functions are denoted in lower case and relate to their trending counterparts based on the following relations

$$P_{t+1} = p_{t+1} e^{g_p t}, \quad (7)$$

$$D_t(P_t) = e^{g_q t} d(p_t), \quad (8)$$

$$\Gamma'_t(H_t) = \gamma'(h_t), \quad (9)$$

where g_q is the assumed rate of growth of quantities. In equation (7), the fact that the price trend in t is applied to the price in $t + 1$ comes from equation (1), where the price trend enters through the cost to produce quantities, which in turn will determine the prices in the next period.

For reasons of market equilibrium, all quantities—final consumption, production, and stocks—must share the same multiplicative trend, so that any discrepancy between the demand and the cost trend will emerge as a price trend.⁷ Defining detrended stocks and acreage using $X_t = x_t \exp(g_q t)$ and $H_{t-1} = h_{t-1} \exp(g_q t)$, we can replace the trending quantities by their detrended counterparts in the above market clearing equation (6):

$$(1 - \delta) x_{t-1} e^{-g_q} + h_{t-1} e^{\eta_{t-1} + \epsilon_t} = d(p_t) e^{\mu_t} + x_t. \quad (10)$$

The multiplication of x_{t-1} by $\exp(-g_q)$ shows that, on average, stocks have to increase just to keep pace with the increased production and demand (for $g_q > 0$), so the detrended past stocks are discounted to maintain them at a level comparable to other detrended quantities.

Similarly, since $\bar{P}_{t+1} = \bar{p} \exp(g_p t)$ where \bar{p} is the steady-state price, the storage non-arbitrage equation (4) can be expressed with detrended variables as

$$\beta (1 - \delta) e^{g_p} E_t p_{t+1} - p_t - k \bar{p} \leq 0, \quad = 0 \text{ if } x_t > 0. \quad (11)$$

The presence of $\exp(g_p)$ in the equation shows that in the stationary model, the price trend is equivalent to adjusting the opportunity cost of storage. Intuitively, a negative price trend—as empirically found in section 5—raises the opportunity cost because, since prices tend to decrease over time, a higher expected price is required to maintain the same level of stocks. Associated with the price trend, the condition $g_p < -\log[\beta(1 - \delta)]$ ensures that inventories are costly and is a necessary condition for the existence of a stationary rational expectations equilibrium,⁸ which is always satisfied for decreasing trends.

⁷Assigning exclusively the quantity trend to demand and the price trend to cost is done here for analytical convenience. Without additional information about the drivers of the trends, we have no basis for doing the assignment, so we opted for the approach requiring the fewest mathematical notations. Given that the econometric analysis is done in deviation from trend, the origin of trends is irrelevant for the results.

⁸It corresponds to the assumption 2 of Deaton and Laroque (1992).

In equation (10), five variables are predetermined: stocks, acreage, and the three shocks. Four of these variables are combined in a single state variable, total available supply s_t , as follows

$$s_t \equiv (1 - \delta) x_{t-1} e^{-g_q} + h_{t-1} e^{\eta_{t-1} + \epsilon_t}. \quad (12)$$

Applying previous transformations to the equilibrium equations leads to the following system of three stationary equilibrium equations associated with three equilibrium variables:

$$h_t : \beta e^{\eta_t - \omega_t} E_t (p_{t+1} e^{\epsilon_{t+1}}) = \gamma' (h_t), \quad (13)$$

$$x_t : \beta (1 - \delta) e^{g_p} E_t p_{t+1} - p_t - k\bar{p} \leq 0, \quad = 0 \text{ if } x_t > 0, \quad (14)$$

$$p_t : s_t = x_t + d(p_t) e^{\mu_t}. \quad (15)$$

It can be seen that, in the stationary model, while the price trend is equivalent to a change in the opportunity cost of storage, the quantity trend does not directly affect the incentives. However, it affects them indirectly through its scaling of past stocks in equation (12). One unit of stocks is less valuable with a positive quantity trend than the same unit without any quantity trend. So a positive quantity trend is equivalent to an increase in the opportunity cost of storage, albeit a one harder to quantify than that coming from the price trend.

2.3. Functional forms

We assume that the stationary demand function takes an isoelastic form such that

$$d(p_t) = \bar{d} \left(\frac{p_t}{\bar{p}} \right)^{\alpha_D}, \quad (16)$$

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

$$\gamma' (h_t) = \beta \bar{p} \left(\frac{h_t}{\bar{d}} \right)^{1/\alpha_S}, \quad (17)$$

where $\alpha_S > 0$ is the supply elasticity. Because of the assumed specifications with variables expressed in deviation from the deterministic steady state, these demand and marginal cost functions depend only on parameters that can be interpreted directly.

Under these assumptions, the four model equations can be expressed as

$$\frac{s_t}{\bar{d}} = (1 - \delta) \frac{x_{t-1}}{\bar{d}} e^{-g_q} + \frac{h_{t-1}}{\bar{d}} e^{\eta_{t-1} + \epsilon_t}, \quad (18)$$

$$\frac{h_t}{\bar{d}} = \left[e^{\eta_t - \omega_t} E_t \left(\frac{p_{t+1}}{\bar{p}} e^{\epsilon_{t+1}} \right) \right]^{\alpha_S}, \quad (19)$$

$$\beta (1 - \delta) e^{g_p} E_t \left(\frac{p_{t+1}}{\bar{p}} \right) - \frac{p_t}{\bar{p}} - k \leq 0, = 0 \text{ if } \frac{x_t}{\bar{d}} > 0, \quad (20)$$

$$\frac{s_t}{\bar{d}} = \frac{x_t}{\bar{d}} + \left(\frac{p_t}{\bar{p}} \right)^{\alpha_D} e^{\mu_t}. \quad (21)$$

From these equations, we see that the only effect of the steady-states quantity (\bar{d}) and price (\bar{p}) is that they scale the mean value of the variables. Once normalized by their mean, all model moments should be identical whatever the choice of these parameters.

Note that these assumed functional forms and the stochastic assumptions imply $E[d^{-1}(\bar{d} \exp(\psi - \mu))] < \infty$, which rules out bubble models such as Bobenrieth et al. (2002).

2.4. Model solution

Equations (5) and (18)–(21) represent a nonlinear rational expectations system based on the variables μ_t , s_t , h_t , x_t , and p_t driven by the innovations $\{\eta_t, \omega_t, \epsilon_t, \nu_t\}$. This system does not have a closed form solution and must be solved numerically to allow for a structural estimation. The solution to the rational expectations system takes the form of policy functions which describe the control variables as functions of the contemporaneous state variables. Different definitions of the state variables can be employed. Given that for the numerical solution we use a projection method, it is important to limit the number of state variables. So far only some of the predetermined variables have been combined in the availability, but a further reduction in the dimensionality of the problem can be achieved.

Instead of working with the acreage h_t , we can work with $q_{t+1}^e = E_t q_{t+1} \exp(-\sigma_\epsilon^2/2) = h_t \exp(\eta_t)$, which is the expected production corrected for the mean harvest-time shock and which is given by

$$q_{t+1}^e = \bar{d} e^{\eta_t} \left[e^{\eta_t - \omega_t} E_t \left(\frac{p_{t+1}}{\bar{p}} e^{\epsilon_{t+1}} \right) \right]^{\alpha_S}. \quad (22)$$

In this case, the transition equation is defined as

$$s_{t+1} = (1 - \delta) x_t e^{-g_q} + q_{t+1}^e e^{\epsilon_{t+1}}. \quad (23)$$

We combine the two planting-time shocks that appear in equation (22) to form the aggregate planting-time shock $\varphi_t \equiv (1 + \alpha_S)\eta_t - \alpha_S\omega_t$. This allows for a further simplification of the supply equation:

$$q_{t+1}^e = \bar{d} e^{\varphi_t} \left[E_t \left(\frac{p_{t+1}}{\bar{p}} e^{\epsilon_{t+1}} \right) \right]^{\alpha_S}. \quad (24)$$

We can see also that in the absence of demand for stock, the market clearing equation (15) collapses to $s_t = d(p_t) e^{\mu_t}$. This simplification implies that, in this situation, the availability and the demand shock can be combined into a variable that we define as net availability, $\tilde{s}_t \equiv s_t e^{-\mu_t}$, i.e., availability in the market corrected for the demand shock.

From the above, it is clear that the minimum set of state variables can be defined as $\{\tilde{s}_t, \varphi_t, \mu_t\}$. We therefore define the policy functions on the set of state variables $\{\tilde{s}_t, \varphi_t, \mu_t\}$:

$$q_{t+1}^e/\bar{d} = \mathcal{Q}(\tilde{s}_t, \varphi_t, \mu_t), \quad (25)$$

$$x_t/\bar{d} = \mathcal{X}(\tilde{s}_t, \varphi_t, \mu_t), \quad (26)$$

$$p_t/\bar{p} = \mathcal{P}(\tilde{s}_t, \varphi_t, \mu_t). \quad (27)$$

To simplify the succeeding expressions, the policy functions are expressed as the variables divided by the steady-state values. Combining the equations defining the model shows that the policy functions for all $\{\tilde{s}_t, \varphi_t, \mu_t\}$ have to satisfy:

$$\mathcal{P}(\tilde{s}_t, \varphi_t, \mu_t) = \max \left\{ (\tilde{s}_t/\bar{d})^{1/\alpha_D}, \right. \\ \left. \beta(1-\delta) e^{g_D} E_t \left[\mathcal{P} \left([(1-\delta) \mathcal{X}(\tilde{s}_t, \varphi_t, \mu_t) e^{-g_q} + \mathcal{Q}(\tilde{s}_t, \varphi_t, \mu_t) e^{\epsilon_{t+1}}] e^{-\mu_{t+1}}, \varphi_{t+1}, \mu_{t+1} \right) \right] - k, \right\}, \quad (28)$$

$$e^{\varphi_t} \left\{ E_t \left[\mathcal{P} \left([(1-\delta) \mathcal{X}(\tilde{s}_t, \varphi_t, \mu_t) e^{-g_q} + \mathcal{Q}(\tilde{s}_t, \varphi_t, \mu_t) e^{\epsilon_{t+1}}] e^{-\mu_{t+1}}, \varphi_{t+1}, \mu_{t+1} \right) e^{\epsilon_{t+1}} \right] \right\}^{\alpha_S} \\ = \mathcal{Q}(\tilde{s}_t, \varphi_t, \mu_t). \quad (29)$$

Equation (28) reveals that two regimes exist. The first regime holds when speculators stockpile in the expectation of future prices covering the full carrying and purchasing costs. The second regime refers to the stockout situation with empty inventories, where the market price is determined only by the final demand for consumption. In the absence of stocks, the equation collapses to $\mathcal{P}(\tilde{s}_t, \varphi_t, \mu_t) = (\tilde{s}_t/\bar{d})^{1/\alpha_D}$, which shows that in this case the only relevant state variable for price determination is net availability. However, the other two state variables determine the production level given that production is based on forward-looking behavior affected by shocks observable at planting time.

This model has no closed-form solution which means its solution must be approximated numerically. Cafiero et al. (2011) show that the precision of the numerical solution is important in the context of estimating a storage model involving simulations; lack of precision could bias the estimates. Thus, we need to balance the need for a solution that is both precise and fast, because the model must be solved at each iteration of the estimation procedure. In Appendix, section A, we propose a new solution method to the storage model based on recent developments in the literature (Maliar and Maliar, 2014) which satisfies this trade-off.

3. Econometric procedure

Not all of the storage model variables are observable. For example, stock levels are available from the United States Department of Agriculture (USDA) statistics but for many countries they are based on USDA estimates in the absence of official statistics.⁹ In this paper, we use the five observable variables proposed by Roberts and Schlenker (2013): price, expected price, consumption, production, and yield shock: $[p_t, E_t p_{t+1}, c_t, q_t, \psi_t]$.

The unknown parameters to be estimated are gathered in the n -vector $\theta \in \Theta$. Our storage model includes fifteen parameters, ten of which are estimated in combination. The other five parameters are fixed or are estimated separately from the procedure described below. As already mentioned, the only role played by the steady-state quantity and price values is to scale the averages of the model variables, hence without loss of generality they are fixed to 1. It is well-known that it is difficult to identify the real discount factor, and especially in short samples involving annual data. Therefore, in structural estimations of storage models it tends to be kept constant. We fix the annual real interest rate at 2%, the value commonly used in the storage literature. It is in line also with Barro and Sala-i-Martin (1990) who derive a mean short-term interest rate of 1.87% for the period 1959–89 for nine OECD countries for which historical data are available. Following the sharp rise to rates of about 5% in the 1980s, the world real interest rate began to decline and reached an average yearly level of about 2% in the mid 2000s (IMF, 2014, Chapter 3). Annual rates of growth of quantities and prices, g_q and g_p , are characterized by the trending behavior of the data (discussed in section 5.2).

Below, we present two estimation strategies. The first is an instrumental variable approach which is in line with Roberts and Schlenker (2013) with the difference that we can derive the equations to estimate from the storage model equations whereas Roberts and Schlenker (2013) had to rely on intuitions from a storage model to propose their estimation strategy. This approach allows us to estimate directly four parameters (α_S , α_D , ρ_μ , and σ_v) but leaves five parameters unidentified. The second strategy is the indirect inference approach. It relies on the supply and demand model from the instrumental variables approach, which is used to build an auxiliary model and enables identification of all the parameters.

3.1. Instrumental variables approach

To ease the notations, our instrumental variables approach is presented with stationary variables. However, the estimations on the observations are based on trending variables. To account for the trends in the variables, flexible trends are added to each equation following Roberts and Schlenker (2013).

⁹The measurement error related to USDA stock levels can be large due to frequent data revisions. E.g., in May 2001 and November 2015, the USDA raised Chinese grain stocks by 164 million tons or by more than 10% of 2001 global production, and Chinese maize stocks by 23.8 million tons or nearly 2.5% of 2015 global production of maize.

3.1.1. Production

Expressed in logarithm, the supply equation (19) is

$$\log q_t = \log (h_{t-1} e^{\psi_t}) = \log (\bar{d}/\bar{p}^{\alpha_S}) + \alpha_S (\eta_{t-1} - \omega_{t-1}) + \alpha_S \log (E_{t-1} (p_t e^{\epsilon_t})) + \psi_t. \quad (30)$$

In this equation, $\eta_{t-1} - \omega_{t-1}$ and $E_{t-1}(p_t e^{\epsilon_t})$ are not observable. However, it is possible to use the expected price $E_{t-1} p_t$ to proxy for the true producer price incentives, which leads to the following estimation equation

$$\log q_t = a_q + b_q \log (E_{t-1} p_t) + c_q \psi_t + u_{q,t}. \quad (31)$$

Since the planting-time shocks are present in the residuals, $u_{q,t}$, and are correlated with the expected price, an ordinary least square (OLS) estimation would suffer from an omitted variable bias. Therefore, following Roberts and Schlenker (2013), we instrument the expected price by the lagged yield shocks ψ_{t-1} . Under the model assumptions, lagged yield shocks are a valid instrument because storage implies that past yield shocks have contemporaneous effects on prices through the availability in the market, and they are not correlated with the planting-time shocks and thus with the residuals. The first-stage equation is

$$\log (E_{t-1} p_t) = a_{E_p} + b_{E_p} \psi_{t-1} + c_{E_p} \psi_t + u_{E_p,t}. \quad (32)$$

This supply-side estimation strategy deserves a few comments. First, substituting the expected price $E_{t-1} p_t$ for the producer incentive price $E_{t-1} (p_t e^{\epsilon_t})$ could potentially create a bias because the former does not include the correlation between the harvest-time yield shock and the price. This implies that b_q will not be a consistent estimator of α_S with the size of the bias depending on the conditional covariance between p_t and ϵ_t . Following the analysis in Gouel (2020, Appendix), this bias is likely to be small for typical parameter values. The Monte Carlo analysis that follows sheds light on this issue.

Second, though this regression allows us to estimate only the supply elasticity, it provides indirect information on the other parameters. Specifically, the estimation of c_q provides information on a combination of the other supply parameters. Neglecting the previously mentioned bias and assuming that $b_q \log (E_{t-1} p_t) = \alpha_S \log (E_{t-1} p_t e^{\epsilon_t})$, we can write

$$\log q_t - b_q \log (E_{t-1} p_t) = a_q + c_q \psi_t + u_{q,t} = \log (\bar{d}/\bar{p}^{\alpha_S}) + \alpha_S (\eta_{t-1} - \omega_{t-1}) + \psi_t. \quad (33)$$

A standard OLS estimator formula gives c_q as a function of the model's parameters:

$$c_q = \frac{\text{cov} (\log q_t - \alpha_S \log (E_{t-1} p_t), \psi_t)}{\text{var} \psi_t} \quad (34)$$

$$= \frac{\text{cov} (\alpha_S (\eta_{t-1} - \omega_{t-1}) + \psi_t, \psi_t)}{\text{var} \psi_t} \quad (35)$$

$$= 1 + \alpha_S \sigma_\eta \frac{\sigma_\eta - \rho_{\eta,\omega} \sigma_\omega}{\sigma_\psi^2}. \quad (36)$$

This (omitted variable bias) formula implies that, if $\rho_{\eta,\omega} \geq 0$, then $c_q \leq 1 + \alpha_S \sigma_\eta^2 / \sigma_\psi^2 \leq 1 + \alpha_S$. It turns out that c_q can exceed $1 + \alpha_S$ only if $\rho_{\eta,\omega} < 0$, an implication that will be useful later to make the link between the 2SLS and the indirect inference estimates.

Similarly, the residuals can be used to obtain a measure of the total supply shock, which we denote ϑ . As for c_q , we can reorganize equation (31) to get

$$\log q_t - \log(\bar{d}/\bar{p}^{\alpha_S}) - b_q \log(E_{t-1} p_t) = c_q \psi_t + u_{q,t} = \alpha_S (\eta_{t-1} - \omega_{t-1}) + \psi_t = \vartheta_t. \quad (37)$$

Thus, although c_q and $u_{q,t}$ cannot be used to directly identify any structural parameter, they provide information when used in the subsequent indirect inference approach.

Third, Hendricks et al. (2015) note that the observable yield shock ψ_t is likely correlated with the planting-time shocks, η_{t-1} and ω_{t-1} (by construction in our model), and hence including it as a control variable mitigates the omitted variable bias. In this context, there is a tradeoff between the consistency of a 2SLS estimate and the higher precision of an OLS estimate. Based on our structural model and the Monte Carlo experiment, we contribute to this debate on whether instrumental variables are actually useful for estimating supply elasticity.

Fourth, although the exclusion restriction holds under the model assumptions, it may not hold in reality. One concern arises if yields are serially correlated. As shown in section 5.3, they display little autocorrelation. To address other concerns regarding the endogeneity of yields, Roberts and Schlenker (2013) replace yield shocks with weather variables, for which endogeneity is less of a concern, and find similar but not as significant results, because weather variables are weaker instruments.

3.1.2. Consumption

From equation (16), logged consumption (denoted c_t), is given by

$$\log c_t = \log(d(p_t) e^{\mu_t}) = \log(\bar{d}/\bar{p}^{\alpha_D}) + \alpha_D \log p_t + \mu_t. \quad (38)$$

By calculating $\log c_t - \rho_\mu \log c_{t-1}$ and using equation (5), we can recover the innovation v_t in the demand equation:

$$\log c_t = (1 - \rho_\mu) \log(\bar{d}/\bar{p}^{\alpha_D}) + \alpha_D \log p_t - \alpha_D \rho_\mu \log p_{t-1} + \rho_\mu \log c_{t-1} + v_t. \quad (39)$$

The fact that v_t is unobservable but correlated with p_t implies that an OLS estimation of equation (39) would again lead to an omitted variable bias. We solve this by instrumenting prices with the yield shocks. Thus, the estimation equation is

$$\log c_t = a_c + b_c \log p_t + c_c \log p_{t-1} + d_c \log c_{t-1} + u_{c,t}, \quad (40)$$

with the associated first stage

$$\log p_t = a_p + b_p \psi_t + c_p \log p_{t-1} + d_p \log c_{t-1} + u_{p,t}. \quad (41)$$

Note that this approach identifies all the demand-side parameters: α_D and ρ_μ in the equation, and σ_v as the standard deviation of the residuals, $u_{c,t}$. This approach differs slightly from that in Roberts and Schlenker (2013) where equation (38) is estimated directly using

$$\log p_t = a_p + b_p \psi_t + u_{p,t} \quad (42)$$

as first stage, since Roberts and Schlenker's focus is on the demand elasticity and not the other parameters. These two approaches are asymptotically equivalent in terms of estimating the demand elasticity.

Since equation (40) includes a lagged dependent variable, a condition for d_c to be consistently estimated is the absence of serial correlation in the residuals, which will be tested using the test proposed by Cumby and Huizinga (1992) which is valid for models that have endogenous regressors. Even in the absence of serial correlation in the residuals, standard estimators of autoregressive models are biased in finite sample. We correct for the finite sample bias using Orcutt and Winokur's (1969) formula: $\hat{\rho}_\mu = (1 + T \hat{d}_c)/(T - 3)$, where T is the sample length.

3.2. Indirect inference approach

Indirect inference requires selection of an auxiliary model. Here, we use the supply and demand model presented above, with some adjustments. The auxiliary model consists of the following system of equations:

$$\log q_t = a_q + b_q \log (E_{t-1} p_t) + c_q \psi_t + u_{q,t}, \quad (43)$$

$$\log (E_{t-1} p_t) = a_{E_p} + b_{E_p} \psi_{t-1} + c_{E_p} \psi_t + u_{E_p,t}, \quad (44)$$

$$\log c_t = a_c + b_c \log p_t + c_c \log p_{t-1} + d_c \log c_{t-1} + u_{c,t}, \quad (45)$$

$$\log p_t = a_p + b_p \psi_t + c_p \log p_{t-1} + d_p \log c_{t-1} + u_{p,t}, \quad (46)$$

$$\psi_t = a_\psi + u_{\psi,t}. \quad (47)$$

The model includes both the first and second-stage equations presented previously, and equation (47) which is included to ensure that the model also fits the standard deviation of yields, an aggregate shock we are able to observe.

The discussion in the above section might suggest that we should estimate the supply and demand equations (43) and (45) using 2SLS since this approach would lead to the lowest biases in the elasticities estimated in the auxiliary model. However, this is not the best option, since use of the indirect inference means that the supply and demand elasticity estimates will not be equal to b_q and b_c . The indirect inference combines the various moments and produces estimates which are the most consistent with the theoretical structure and the moments. Through the lens of the omitted variable bias formula, the theoretical structure of the model imposes a clear mapping between b_q estimated by OLS and the model parameters (and similarly for b_c). As a result, employing equations estimated using OLS in the auxiliary model provides similar

information to equations estimated using 2SLS, and has the advantage of being more precise, since the precision of the equations estimated with 2SLS is dependent on the correlation between the endogenous regressors and the instruments.

Hence, our benchmark auxiliary model is based on the system (43)–(47) estimated by OLS. However, we retain the first-stage equations in the system because they contain information not provided in the other equations. For robustness, we also use the supply and demand model estimated by 2SLS as an additional auxiliary model. See Li (2010) and Guvenen and Smith (2014) for two other papers that rely on linear equations estimated by OLS as the auxiliary model in an indirect inference setting.¹⁰ Using the selected auxiliary model, we can define the objective using a subset of the model parameters which excludes the intercepts, since these are informative only about the steady-state values which we normalize to unity: $\zeta = [b_q, c_q, \sigma_{u_q}, b_{E_p}, c_{E_p}, \sigma_{u_{E_p}}, b_c, c_c, d_c, \sigma_{u_c}, b_p, c_p, d_p, \sigma_{u_p}, \sigma_{u_\psi}]$.

This auxiliary model has two important benefits. First, since it involves only linear regressions, it is trivial to estimate, and avoids the indirect inference procedure being burdened by a computationally costly auxiliary model. Second, it is quite transparent regarding the relationships between the auxiliary model and the storage model parameters. b_q is asymptotically equal to the supply elasticity plus the omitted variable bias. From equation (36), c_q and similarly σ_{u_q} are both nonlinear combinations of α_S , σ_ϵ , σ_η , σ_ω , and $\rho_{\epsilon,\omega}$. From Hendricks et al. (2015), c_{E_p} is related to the predictability of the yield shocks, and thus to σ_η . In equation (45), b_c consists of the demand elasticity plus the omitted variable bias which is related to ρ_μ and σ_v , themselves informed by d_c and σ_{u_c} . In equation (46), c_p is linked to the first-order autocorrelation of $\log p$, which conditional on the other parameters, depends directly on the storage costs δ and k . More precisely, lower storage costs imply more storage and hence a higher price autocorrelation (Gouel and Legrand, 2017, Figure 2), and vice versa. In equation (47), $\sigma_{u_\psi}^2 = \sigma_\epsilon^2 + \sigma_\eta^2$. Finally, the inclusion of $\sigma_{u_{E_p}}$ and σ_{u_p} is almost equivalent to including the standard deviations of the price and the expected price in the objective, and ensures that the estimated model will also fit these targets.

We use ζ_T to denote the 15×1 vector of the auxiliary model estimates from the observations of length $T + 1$, while $\zeta_T^i(\theta)$ denotes the counterpart of ζ_T estimated on artificial data generated by the storage model for a given set of parameters θ . We simulate $\tau \geq 1$ samples of size $T + 1 + t^{\text{burn}}$. The first $t^{\text{burn}} = 50$ simulations are dropped as burn-in periods to remove the influence of the initial state. The final $T + 1$ simulations are used for the estimations, but the first is dropped due to the lagged variables appearing in the auxiliary model. The indirect inference estimator then is

$$\hat{\theta} = \arg \min_{\theta \in \Theta} \left[\hat{\zeta}_T - \frac{1}{\tau} \sum_{i=1}^{\tau} \hat{\zeta}_T^i(\theta) \right]' W \left[\hat{\zeta}_T - \frac{1}{\tau} \sum_{i=1}^{\tau} \hat{\zeta}_T^i(\theta) \right], \quad (48)$$

¹⁰See also Simonovska and Waugh (2014) for an estimation approach in which a biased auxiliary model is used to obtain an unbiased simulated estimator.

where W is a 15×15 symmetric nonnegative definite weighting matrix. This estimator minimizes the weighted distance between the auxiliary model parameters estimated using actual data, and those estimated using data simulated from our structural storage model.

At every step of the minimization, a new set of parameters θ is proposed. For this new θ , a numerical solution of the storage model is computed using the algorithm proposed in Appendix A. The resulting policy functions are used to simulate the model starting from the deterministic steady state and using random shocks drawn at the beginning of the estimation procedure and kept fixed throughout.

In line with [Gourieroux et al. \(1993\)](#), and assuming that W is the optimal weighting matrix, the variance-covariance matrix for the parameter estimates converges asymptotically to

$$\left(1 + \frac{1}{\tau}\right) (J'WJ)^{-1}, \quad (49)$$

where $J = (1/\tau) \sum_{i=1}^{\tau} E[\partial \hat{\zeta}_T^i(\theta) / \partial \theta]$ is a $15 \times n$ full rank matrix, evaluated by central difference at $\theta = \hat{\theta}$. The optimal weighting matrix is the inverse of the variance-covariance matrix of the estimate of ζ_T . We calculate this using the formulas for standard errors robust to heteroskedasticity for the standard regression parameters $(b_q, c_q, b_{E,p}, c_{E,p}, b_c, c_c, d_c, b_p, c_p, d_p)$, and using the following formulas for the standard deviations $(\sigma_{u_q}, \sigma_{u_{E,p}}, \sigma_{u_c}, \sigma_{u_p}, \sigma_{u_\psi})$:

$$\text{var}(\sigma^{\text{OLS}}) = \frac{(\sigma^{\text{OLS}})^2}{2(T-l)} \quad \text{and} \quad \text{var}(\sigma^{\text{2SLS}}) = \frac{(\sigma^{\text{2SLS}})^2}{2(T-l)R_p^2}, \quad (50)$$

where $T-l$ is the degree of freedom of the corresponding regression, and in the case of the residuals from the second stage of the 2SLS, R_p^2 is the partial R^2 from the first stage where the endogenous variables and the instruments have both been regressed on the exogenous variables in a first step ([Bound et al., 1995](#)). This gives a diagonal weighting matrix, a common simplification in the indirect inference literature (see, e.g., [Christiano et al., 2005](#); [Ruge-Murcia, 2020](#)).

There are more parameters included in the auxiliary model than parameters to be estimated in the storage model which means that there are overidentification restrictions, which will be tested using the statistics

$$\frac{T\tau}{1+\tau} \min_{\theta \in \Theta} \left[\hat{\zeta}_T - \frac{1}{\tau} \sum_{i=1}^{\tau} \hat{\zeta}_T^i(\theta) \right]' W \left[\hat{\zeta}_T - \frac{1}{\tau} \sum_{i=1}^{\tau} \hat{\zeta}_T^i(\theta) \right], \quad (51)$$

which follows asymptotically a χ^2 distribution with $15-n$ degrees of freedom ([Gourieroux et al., 1993](#)).

Since it is costly to evaluate the objective in equation (48), because it requires a new solution and additional simulations of the storage model for each updated set of parameters, and in

the absence of analytical derivatives, we employ for minimization a derivative-free algorithm, BOBYQA (Powell, 2009). We also use bounds to avoid exploration of parameter values outside their domain of definition but also those that would make it difficult to solve the model (see table A2). Furthermore, to limit the risk of finding only a local optimum, the optimization algorithm starts from 500 different initial values of θ , with the exception of the Monte Carlo experiments in the next section which uses a unique starting point. Finally, although it is costly to solve for the rational expectations equilibrium of the model, it is less costly to simulate from it. We therefore choose $\tau = 200$ to minimize the simulation-related uncertainty in the estimates.

4. Monte Carlo experiments

Except for Michaelides and Ng (2000), there is no example of using indirect inference to estimate the storage model, and this work involved a much simpler storage model than ours, as well as different auxiliary models. Therefore, in this section we employ a Monte Carlo analysis to study the small-sample properties of this estimator and gauge the ability of our selected auxiliary model to reveal the true structural parameters. Since Roberts and Schlenker (2013)'s supply and demand model allows direct estimation of some of the model parameters and forms the basis of our auxiliary model, we include it in the Monte Carlo analysis.

All the experiments are based on 500 replications and use the same sample size $T = 56$ as used subsequently on observations (the results for longer samples are provided in the Appendix). The model parameters chosen for the experiments are based on the estimates in section 6, except for σ_ω and δ for which different values are chosen to illustrate some of the difficulties that can be encountered. The parameter values used are $\beta = 0.98$, $g_q = 2.5\%$, $g_p = -2\%$, $\rho_\mu = 0.5$, $\rho_{\eta,\omega} = -0.4$, $\sigma_\eta = 1.5\%$, $\sigma_\epsilon = 2\%$, $\sigma_v = 1.6\%$, $k = 3\%$, $\alpha_D = -0.07$, and $\alpha_S = 0.08$. Since the cost shock, ω , is a crucial and unobserved determinant of the omitted variable bias in the supply equation, we run the Monte Carlo experiments for three values of its standard deviation: $\sigma_\omega = \{0.05, 0.1, 0.2\}$, the latter being close to the value estimated in section 6. δ is estimated to be zero, but we use here $\delta = 2\%$ in order to test whether our estimation strategy can estimate precisely both types of storage cost. For the indirect inference, the optimization for each replication starts from a different vector θ with values drawn randomly from the 80% and 120% range of the true values.

The results of the OLS and 2SLS approaches are reported in table 1 panels A and B, and the results for the indirect inference approach with an auxiliary model based only on OLS regressions are presented in table 2. The results for the longer samples and the indirect inference based on 2SLS regressions for the auxiliary model are contained in Appendix tables A3–A5. These tables show that, for the parameters that are common to both methods, the indirect inference approach is more precise than either the OLS or 2SLS approaches, as evidenced by the lower root mean squared errors (RMSEs) obtained in either small or large samples. Note also that estimates of the demand and supply elasticities exploiting the indirect inference approach are not biased in either the small or the large samples which contrasts with the OLS estimates on

which they are based.¹¹ This confirms that in the context of indirect inference the elasticities are not just set equal to their OLS counterparts. More precisely, the approach relies on the information derived from b_c^{OLS} and c_q^{OLS} in combination with the other parameters, and delivers unbiased and consistent elasticity estimates.¹²

Nevertheless, using the indirect inference approach, three parameters are difficult to estimate: the storage costs (δ and k), and the correlation between the planting-time shocks ($\rho_{\eta,\omega}$). For the storage costs, one difficulty comes from the inability of indirect inference to estimate separately δ from k . Both storage costs are identified from the same moments (a point proved on estimations made on observations in section 6.4), making their separate estimation challenging. If we consider instead the total storage at steady state (excluding opportunity costs), $k + \delta$, then the RMSE is much lower at 55%. Although still high, this value is similar to what is obtained in a Monte Carlo experiment with parameters $\delta = 0$ and $k = 5\%$ (in which case all values are close to those in table 2 except for k with mean 5.07%, standard deviation 2.68%, RMSE 53.70%, and asymptotic standard error 1.91%). The limited precision of the estimation of total storage costs could stem from the fact that they are identified only indirectly, in part through their effect on the autocorrelation and volatility of prices. The full information approaches in Cafiero et al. (2015) and Gouel and Legrand (2017) provide lower RMSE for their storage cost parameter. However, in our context these approaches are not feasible given that they require observability of the planting-time shocks.

The parameter $\rho_{\eta,\omega}$ is estimated based on its effect on the auxiliary parameter c_q^{OLS} (see equation (36)). However, what matters for estimating $\rho_{\eta,\omega}$ is $c_q^{\text{OLS}} - 1 = \alpha_S \sigma_\eta (\sigma_\eta - \rho_{\eta,\omega} \sigma_\omega) / \sigma_\psi^2$ and this is not precisely estimated in the auxiliary model (table 1). The estimates of $\rho_{\eta,\omega}$ will be affected not only by the uncertainty related to the estimates of $c_q^{\text{OLS}} - 1$, but also by the uncertainty related to the other parameter estimates which explains its high RMSE. However, the challenges related to estimating $\rho_{\eta,\omega}$ are of secondary importance. In section 2.4, this parameter is absent from the rational expectations problem expressed in compact form. The equilibrium price, expected price, demand, and production depend not on the specific value of the shock ω but rather on the aggregate shock φ . In a Monte Carlo experimental setting, it is possible to calculate the RMSE for σ_φ . At 18%—for $\sigma_\omega = 20\%$ —this is similar to the RMSE for the other shocks. Overall, the empirical method seems to be appropriate for estimating the volatility of all the shocks, but the various errors will be compounded in $\rho_{\eta,\omega}$ which is difficult to estimate, although without consequences for the rest of the model.

Tables 1–2 and A3–A5 show that both approaches have good asymptotic properties. The RMSE and their two components vanish “asymptotically”—i.e., as the sample length increases from 56, to 100, 200, and 1000—showing the consistency of both estimators (apart from a small bias in

¹¹In the case of the supply elasticity, the inconsistency caused by using the expected price to substitute for the true incentive price can be evaluated employing an OLS regression to estimate equation (30) where $E_{t-1}(p_t \exp(\epsilon_t))$ is replaced by $E_{t-1} p_t$. At -1.7% , this bias is small under these parameters.

¹²This also applies in the case of an auxiliary model based on 2SLS regressions (see table A5).

Table 1 – Monte Carlo experiment with OLS and 2SLS estimations of the supply and demand equations

	ρ_μ	$c_q - 1$	σ_ψ (%)	σ_θ (%)	σ_ν (%)	α_D	α_S
<i>Panel A. OLS</i>							
$\sigma_\omega = 5\%$							
Mean	0.36	0.049	2.49	2.64	1.28	-0.021	0.067
St. dev.	0.13	0.023	0.23	0.25	0.14	0.011	0.005
RMSE (%)	37.93	43.240	9.40	9.65	21.70	71.702	17.712
SE	0.14	0.024	0.24		0.13	0.011	0.005
$\sigma_\omega = 10\%$							
Mean	0.37	0.065	2.49	2.75	1.29	-0.022	0.053
St. dev.	0.13	0.043	0.23	0.27	0.14	0.011	0.009
RMSE (%)	36.57	55.737	9.40	10.37	20.97	69.774	35.730
SE	0.14	0.045	0.24		0.13	0.011	0.009
$\sigma_\omega = 20\%$							
Mean	0.39	0.077	2.49	3.00	1.33	-0.026	0.018
St. dev.	0.13	0.078	0.23	0.31	0.14	0.010	0.015
RMSE (%)	33.51	72.009	9.40	14.01	19.25	65.089	79.423
SE	0.13	0.083	0.24		0.13	0.010	0.014
<i>Panel B. 2SLS</i>							
$\sigma_\omega = 5\%$, {Supply: $E(F) = 14$, $E(p\text{-value}) = 0.34$ }, {Demand: $E(F) = 24$, $E(p\text{-value}) = 0.00$ }							
Mean	0.50	0.069	2.49	2.70	1.65	-0.072	0.080
St. dev.	0.18	0.032	0.23	0.26	0.35	0.024	0.015
RMSE (%)	36.02	47.704	9.40	9.76	21.92	34.804	18.332
SE	0.19	0.035	0.24		0.28	0.021	0.016
$\sigma_\omega = 10\%$, {Supply: $E(F) = 14$, $E(p\text{-value}) = 0.28$ }, {Demand: $E(F) = 23$, $E(p\text{-value}) = 0.01$ }							
Mean	0.49	0.111	2.49	2.90	1.64	-0.072	0.082
St. dev.	0.18	0.064	0.23	0.32	0.34	0.024	0.028
RMSE (%)	35.17	60.495	9.40	11.05	21.61	34.299	35.253
SE	0.19	0.069	0.24		0.28	0.021	0.029
$\sigma_\omega = 20\%$, {Supply: $E(F) = 12$, $E(p\text{-value}) = 0.17$ }, {Demand: $E(F) = 22$, $E(p\text{-value}) = 0.02$ }							
Mean	0.49	0.202	2.49	3.46	1.65	-0.072	0.089
St. dev.	0.17	0.147	0.23	0.69	0.34	0.024	0.066
RMSE (%)	33.51	81.341	9.40	20.83	21.59	34.297	82.679
SE	0.18	0.154	0.24		0.29	0.022	0.064

Notes: Monte Carlo experiment based on 500 replications, with a sample size $T = 56$. True values: $\rho_\mu = 0.5$, $\sigma_\psi = 2.5\%$, $\sigma_\nu = 1.6\%$, $\alpha_D = -0.07$, and $\alpha_S = 0.08$. The values of c_q and σ_θ vary with σ_ω as follows $c_q = \{1.067, 1.106, 1.182\}$ and $\sigma_\theta = \{2.70, 2.88, 3.36\}$ corresponding to $\sigma_\omega = \{0.05, 0.1, 0.2\}$. The mean and standard deviations are respectively the average and standard deviations of the empirical parameter distribution. They are combined to calculate the RMSE expressed as a percentage of the true parameter value. SE is standard errors and represents the average of the standard errors robust to heteroskedasticity. $E(F)$ is the average first-stage F -statistics. $E(p\text{-value})$ is the average p -value for the Hausman test of endogeneity. ρ_μ in panel B is bias adjusted (Orcutt and Winokur, 1969)

the supply elasticity discussed above).

Table 2 – Monte Carlo experiment with indirect inference approach (auxiliary model based on OLS regressions)

	ρ_μ	$\rho_{\eta,\omega}$	σ_ω (%)	σ_η (%)	σ_ϵ (%)	σ_ν (%)	δ (%)	k (%)	α_D	α_S
$\sigma_\omega = 5\%$	OID: 0.043									
Mean	0.50	-0.45	5.05	1.47	1.98	1.61	1.97	3.08	-0.071	0.080
St. dev.	0.11	0.31	0.66	0.33	0.28	0.26	1.34	2.27	0.016	0.008
RMSE (%)	22.19	78.65	13.31	22.14	14.13	16.26	67.26	75.82	22.380	9.539
ASE	0.09	0.39	0.67	0.36	0.30	0.26	19.33	18.19	0.020	0.008
$\sigma_\omega = 10\%$	OID: 0.049									
Mean	0.50	-0.44	10.22	1.46	1.98	1.62	2.00	3.08	-0.071	0.081
St. dev.	0.11	0.29	1.63	0.35	0.29	0.26	1.47	2.26	0.016	0.014
RMSE (%)	22.67	72.08	16.41	23.31	14.54	16.08	73.66	75.45	22.346	17.661
ASE	0.09	0.34	1.56	0.38	0.32	0.26	18.58	17.58	0.020	0.014
$\sigma_\omega = 20\%$	OID: 0.038									
Mean	0.50	-0.45	21.01	1.46	1.98	1.62	2.00	3.17	-0.072	0.083
St. dev.	0.11	0.28	5.60	0.38	0.31	0.24	1.62	2.28	0.015	0.026
RMSE (%)	22.94	70.52	28.44	25.25	15.53	15.25	81.17	76.20	21.506	32.172
ASE	0.10	0.33	5.23	0.41	0.34	0.26	18.51	17.09	0.020	0.025

Notes: Monte Carlo experiment based on 500 replications, with a sample size $T = 56$. True values: $\rho_\mu = 0.5$, $\rho_{\eta,\omega} = -0.4$, $\sigma_\eta = 1.5\%$, $\sigma_\epsilon = 2\%$, $\sigma_\nu = 1.6\%$, $\delta = 2\%$, $k = 3\%$, $\alpha_D = -0.07$, and $\alpha_S = 0.08$. The mean and standard deviations are respectively the average and standard deviations of the empirical parameter distribution. They are combined to calculate the RMSE expressed as a percentage of the true parameter value. ASE means asymptotic standard errors, based on equation (49), and represents the average standard errors calculated at the solutions. OID is the empirical size of the chi-square test of overidentifying restrictions.

The standard errors (rows SE in table 1 and asymptotic standard errors (ASE) in table 2) are similar to the standard deviations of the Monte Carlo estimates showing that for both methods the standard errors are consistent with the standard deviations in the population. The only exception is the storage costs for which the asymptotic standard errors are one order of magnitude above the population standard deviations. This is only true when both storage costs are included in the model. In a model with only per-unit storage costs (with $k = 5\%$), standard deviations and asymptotic standard errors are similar. The comparability of standard deviations and standard errors is an important result for two reasons. Reliable auxiliary model standard errors matter because in the indirect inference approach they directly enter the weighting matrix. Also, consistent indirect inference standard errors in the Monte Carlo analysis suggests that the asymptotic formula we apply has a limited small-sample bias (tables A4 and A5 show that with longer samples the biases are negligible). The empirical size of the OID statistic is close to but below its 5% critical value. This means that this test statistic is biased against rejecting the model identification restrictions. This problem is exacerbated in longer samples (table A4).¹³

¹³This size-distortion issue related to specification tests is acknowledged in the literature (see, e.g., Ruge-Murcia, 2007; Michaelides and Ng, 2000).

We also ran Monte Carlo estimations for gradually increasing sizes of σ_ω to analyze its role in the parameter estimations. Table 1 panels A and B show that an increase from 5 to 20% in σ_ω affects only the OLS and 2SLS performances for the supply-side parameters estimates. Varying σ_ω fleshes out the trade-off between consistency and precision in the supply elasticity estimates highlighted by Hendricks et al. (2015). What is gained in terms of reduced bias from using 2SLS is lost through higher volatility of the estimates, resulting in similar but lower RMSE for the OLS compared to the 2SLS. This is because a higher σ_ω implies a larger omitted variable bias but it also makes the lagged yield shocks a weaker instrument because their role in explaining price changes declines as the variance of cost shocks increases. For this choice of storage model parameters, deciding between estimating supply using OLS or 2SLS is difficult given that both approaches have some limitations. However, in the present context, as documented in table 2, the indirect inference approach is much more robust to σ_ω with RMSEs which deteriorate less as this parameter increases.

In addition to the Monte Carlo results on the parameters, table 1 displays some results about the 2SLS diagnostics statistics: the first-stage F -statistics and the p -value for the Hausman test of endogeneity. The F -statistics show that the instrument is much weaker for the supply equation than for the demand equation with a mean value closer to 10, and deteriorating with σ_ω . This is consistent with the idea that a lagged yield, used to instrument the expected price in the supply equation, is a worse predictor of price than current yield is which is something that will also be found on observations thereafter. The Hausman test shows that the null of exogenous expected prices is not rejected on average, in a context where we know that these prices are endogenous and should be instrumented. This result disappears in longer samples (table A3) and is explained by the important standard errors that the supply elasticity displays in small samples.

Finally, before deciding about the most appropriate auxiliary model, we rely on the Monte Carlo estimations to investigate the effect of substituting the auxiliary model based on OLS estimates of the demand and supply equations by the 2SLS estimates. Appendix Table A5 reports the Monte Carlo results using the parameters estimated by 2SLS. The two indirect inference approaches have similar performance, apart from α_S , σ_v , and σ_ω which are estimated with much higher precision in the OLS-based model; loss of precision is associated with the instrumentation. These results support our choice to use the OLS regression based auxiliary model as the baseline and to use the auxiliary model based on the 2SLS regression as a robustness check.

5. Overview of the grains market

With some small modifications, our data series is constructed following Roberts and Schlenker (2013) but for completeness we present all the different choices along with the descriptive statistics.

5.1. Data

The observations include five annual time series—price, expected price, consumption, production, and yield shock—for a caloric aggregate of the four basic staples: maize, rice, soybeans, and wheat. Information on quantities come from the Food and Agriculture Organization statistical database (FAO, 2020) with data for 1961–2017 on production, stock variations, yield and area harvested. Consumption is obtained by subtracting stock variations from total production. Following Roberts and Schlenker (2013), the four commodities are aggregated into calories using the conversion ratios in Williamson and Williamson (1942).

For country i , crop l , and κ_l the caloric content of a ton of crop l , the global annual yield shocks Ψ_t are computed according to the approach proposed by Hendricks et al. (2015):

$$\Psi_t = \frac{\sum_l \sum_i A_{lit} \kappa_l Y_{lit}}{\sum_l \sum_i A_{lit} \kappa_l \hat{Y}_{lit}} = \sum_l \sum_i \rho_{lit} \Psi_{lit}, \quad (52)$$

where A_{lit} is the harvested area, Y_{lit} is the yield, \hat{Y}_{lit} is the trend yield, and

$$\rho_{lit} = \frac{A_{lit} \kappa_l \hat{Y}_{lit}}{\sum_{l'} \sum_{i'} A_{l'i't} \kappa_{l'} \hat{Y}_{l'i't}} \quad (53)$$

is the weight of the country-crop shocks in the aggregate shock. Yields are decomposed multiplicatively into a trend yield and a yield shock: $Y_{lit} = \hat{Y}_{lit} \Psi_{lit}$. The trend yield is obtained from the model prediction regressing the logarithm of yield over 4-knot natural cubic spline with the corresponding observation deleted. The trend yield model has to be run separately for each country, crop, and year. The prediction is corrected for the transformation bias introduced by the logarithm using the residual variance of the trend yield model. All countries are included in the calculation but the smallest contributing less than 0.5% to a crop's world production are aggregated.

This data construction implies that the yield shock in the model corresponds to the logarithm of the yield shock calculated here, $\psi_t = \log \Psi_t$, and the acreage in the model corresponds in the data to $H_{t-1} = Q_t / \Psi_t = \sum_l \sum_i A_{lit} \kappa_l \hat{Y}_{lit}$. Following the discussion in Hendricks et al. (2015), this definition has implications for the interpretation of the supply elasticity as represented in the model. The model supply elasticity combines an acreage elasticity and an average trend yield effect related to changes in the composition of growing areas across countries associated with price changes. Hendricks et al. (2015) argue that to avoid this composition effect the supply elasticity should be estimated based only on acreages. In the present context of a market model, it is the total supply elasticity that matters since this determines the price.

There are several sources of price information but it is important to choose the prices that are the most consistent with the model. For example, the annual prices in Deaton and Laroque (1992) are from the World Bank and are obtained by averaging prices over the calendar year, which can

induce spurious correlations due to mixing different marketing seasons (Guerra et al., 2015). The model includes two prices: the current price P_t , which is the price received by the farmers at harvest time and paid by consumers, and the expected price $E_{t-1} P_t$, which corresponds to the farmers' rational expectations at planting time about the price P_t they will receive at harvest time. Since Gardner (1976), it is common to use futures prices in place of the unobservable expected price. This is a valid approach if futures prices are unbiased predictor of spot prices, which is not true for all commodities but is true for the commodity prices studied here according to Chinn and Coibion (2014).¹⁴ Given the annual time-frame of the model, we take futures contracts with a one-year horizon. For consistency, P_t is the corresponding futures contract at delivery. Following Roberts and Schlenker (2013), we use prices from the Chicago Board of Trade futures for the main month following each crop harvest (i.e., December for maize and wheat, November for rice and soybeans).¹⁵ Monthly prices are obtained by averaging the daily prices observed during each month. Futures prices for rice started trading in 1986. Due to lack of data, we exclude rice from our calculation of the price index (which is in line with Roberts and Schlenker, 2013). Futures prices are deflated by the US CPI and aggregated into a single caloric price index series using the caloric weights, ρ_{lit} , derived in equation (53):

$$P_t = \frac{\sum_{l \neq \text{rice}} (\sum_i \rho_{lit}) P_{lt|t}/\kappa_l}{\sum_{l \neq \text{rice}} \sum_i \rho_{lit}} \quad \text{and} \quad E_{t-1} P_t = \frac{\sum_{l \neq \text{rice}} (\sum_i \rho_{lit}) P_{lt|t-1}/\kappa_l}{\sum_{l \neq \text{rice}} \sum_i \rho_{lit}}, \quad (54)$$

where $P_{lt|t-n}$ denotes the real crop- l futures price at time $t - n$ for delivery at time t .

5.2. Non-stationarity

Figures 1 and 2 plot the constructed production, consumption, and price series used for inferences thereafter. In line with the model trend assumptions, these series do not appear stationary. There is a large literature on the nature of trends in commodity prices which was motivated by the Prebisch-Singer hypothesis of a secular deterioration in primary commodity prices relative to the prices of manufactured goods (e.g., Ghoshray, 2010; Lee et al., 2006). An important take-away from this literature is that, over long periods, it is necessary to account for possible breaks in deterministic trends to avoid spurious rejection of the assumption of a deterministic trend.¹⁶ We test for stationarity using the endogenous two-break Lagrange Multiplier (LM) unit root test developed by Lee and Strazicich (2003, 2013) and Lee et al. (2006). The LM tests allow for one or two structural breaks with or without a linear or quadratic deterministic trend under both the null and alternative hypotheses.

¹⁴However the lack of convergence for several grain futures have partly altered this property during the period 2005–10 (Garcia et al., 2015).

¹⁵At the beginning of the series, not all futures contracts extended one year in advance. In these cases, we use the average price for the first month the contract was traded.

¹⁶It is well-known that omitting possible structural breaks can lead to a bias resulting in retention of the unit root null hypothesis when it should be rejected (Perron, 1989; DeJong et al., 1992; Zivot and Andrews, 1992).

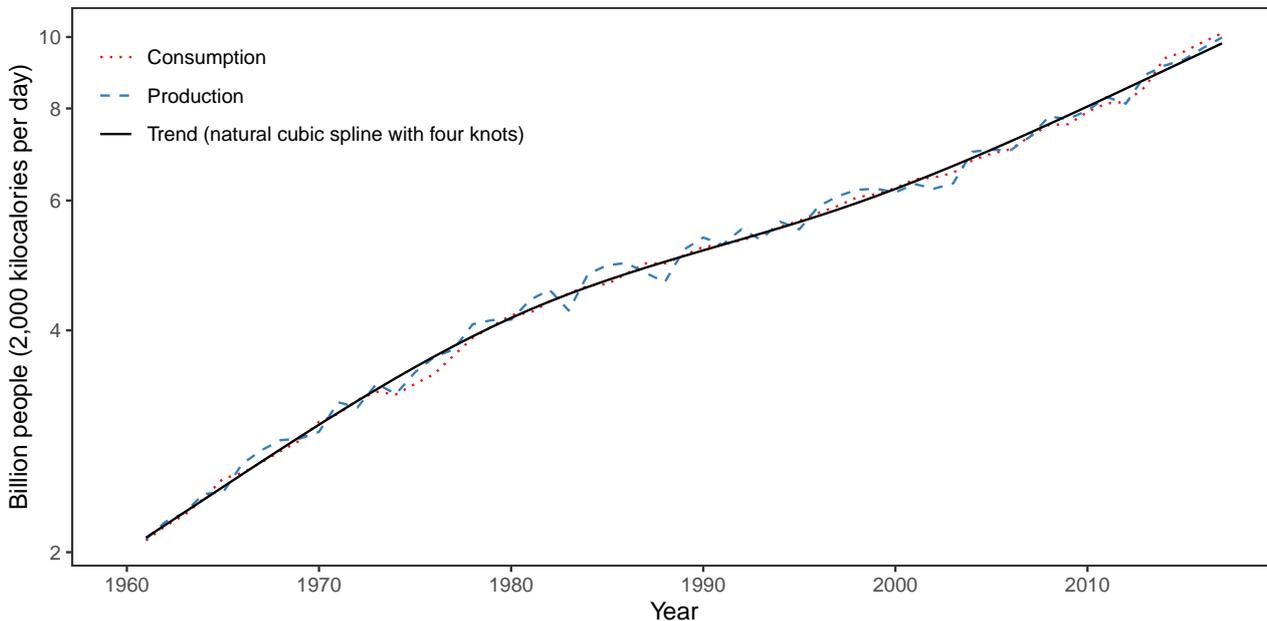


Figure 1 – World caloric production and consumption, and their trend for 1961–2017. The y-axis is the number of people that hypothetically could be fed 2,000 kilocalories per day diet based on consumption of only the four commodities.

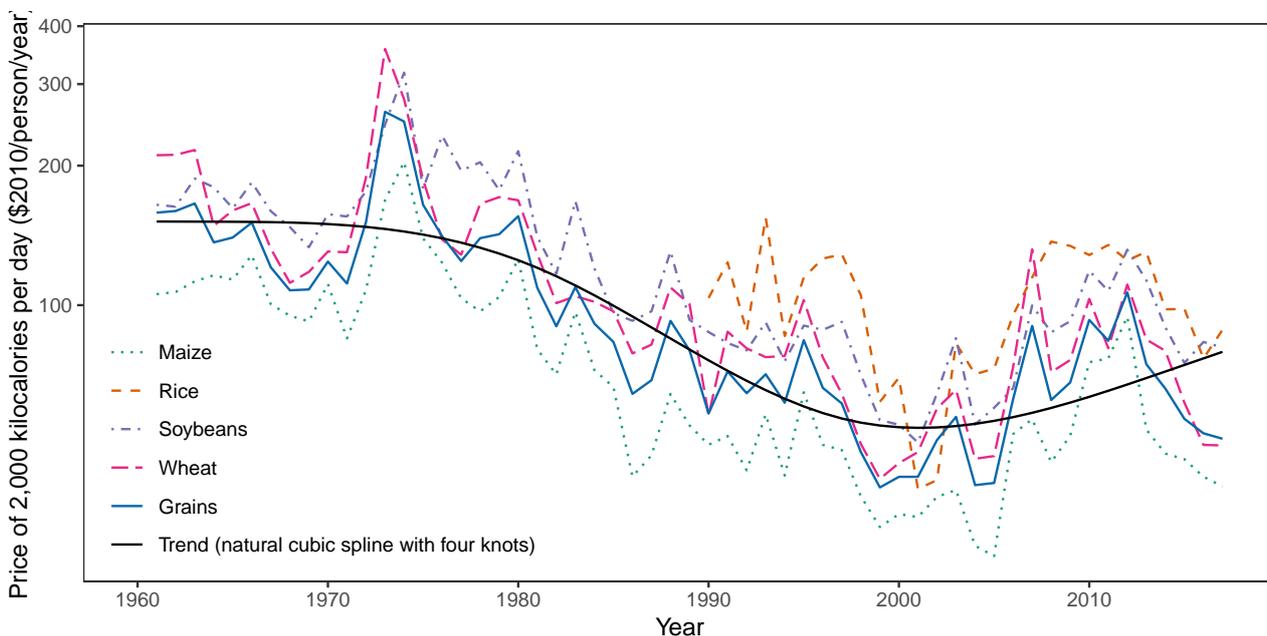


Figure 2 – Real caloric prices at delivery. The y-axis is the annual cost of 2,000 kilocalories per day.

Although our econometric models call for variables in logarithms, it is well known that unit root

tests are highly sensitive to data transformation which is likely also to transform the underlying trends (Corradi and Swanson, 2006). For example, in levels, quantities exhibit a nearly linear trend up to the mid 2000s, but this is less evident in logarithm. We therefore apply the variable tests in levels (results are reported in Appendix table A6, panel A). The null of difference stationarity is rejected for all the variables with one break, two breaks, or in both specifications using the bootstrap critical values given by Lee et al. (2006).¹⁷ More precisely, for production and consumption, the unit root assumption is rejected at the 5% level of significance with two structural breaks in 1982 and 2000, and 1984 and 2007. Regarding the spot and expected prices, the two-break LM test with a quadratic trend rejects the null at the 5% level with a single estimated break occurring in 1979 and 1980.¹⁸

These tests support our deterministic trends modeling choice. However, there is a mismatch between the log-linear trends assumed in the model and the flexibility needed to make the data stationary. This difference is common in macroeconomic models; a trend consistent with a growth path may not be sufficiently flexible to stationarize the data. Various solutions to the problem have been explored; all involve tradeoffs related to consistency between the theoretical and empirical models (see the discussion in Canova, 2014). In our case, the consequences of this mismatch are likely to be small for two reasons. First, the quantitative effect of the trend g_q on the variables of interest is quite limited (see results of table 11 in section 7.1). Second, the deviations from the log-linear trends are small meaning that the theoretical model accounts well for the first-order effects of trends. So, the variations around the linear trend captured by our more flexible specification are likely of small quantitative importance.

Since our econometric models use variables in logarithms, we need log-detrended variables. To be consistent with Roberts and Schlenker's empirical approach, we adopt their natural cubic spline specification to model the trend and consider three levels of flexibility, with three to five knots.¹⁹ We confirmed the stationarity of the detrended variables by running the usual augmented Dickey-Fuller (ADF), Phillips-Perron (PP) and Kwiatkowski-Phillips-Schmidt-Shin (KPSS) unit root tests. Results are reported in table A6, panel B, with an increasing degree of flexibility going from the top to the bottom of each column. We find, with the exception of the variables detrended with three knots, the remainder are stationary at the 1% level of significance.²⁰ In other words, a natural cubic spline with three knots—i.e., flexibility equivalent to a quadratic trend—is not sufficiently flexible to make both price and quantity data stationary. Since the four-knot spline involves the minimum flexibility needed to make the data stationary, this is our preferred trend specification; as a robustness check we test for more and less flexible

¹⁷Based on 5,000 replications of sample sizes $T = 100$.

¹⁸It is interesting that if we assume two breaks for prices, the dates correspond to two food crises after which food prices settled at higher average levels. This applies also to consumption in relation to a regime change in 2007 which followed the implementation of the biofuels mandates in Europe and the United States (Wright, 2014).

¹⁹Unless indicated otherwise, when natural cubic splines are used, their knots are located according to the percentiles method suggested in Harrell (2001): 1967, 1989, 2011 for 3 knots; 1964, 1981, 1997, 2014 for 4 knots; and 1964, 1976, 1989, 2002, 2014 for 5 knots.

²⁰Recall that in the KPSS test the null is a trend-stationary series.

trends.

Finally, to simulate the storage model requires trend parameters g_q and g_p . In contrast to the other parameters, these are estimated separately and before applying the indirect inference. In the theoretical model, consumption and production, and the demand and supply prices show common trends respectively denoted g_q and g_p . We estimate $g_q = 2.54\%$ by regressing the logged quantities (consumption and production) on a common linear trend and similarly with the logged prices to estimate $g_p = -2.03\%$.

5.3. Descriptive statistics

In this section we present some descriptive statistics for the detrended data and discuss their implications for the estimation of the storage model.

Table 3 contains the correlation between the detrended real prices at delivery. It shows that crop prices are strongly correlated with one another, and all but rice have a correlation with the grains index in excess of 0.88. These high correlations are indicative of the large substitution possibilities between these basic staples. We observe that with the exception of the correlation between rice and soybeans, crop prices are correlated more strongly to the grain index than to the prices of any of the other crops. These high correlations support use of an aggregated caloric index to measure the state of the world grain market. In addition to the issues involved in solving and estimating a multi-crop storage model, an estimation based on the separate crops considered would risk mixing own-price and cross-price elasticities.

Table 3 – Correlation coefficients of detrended real prices at delivery, 1961–2017 (except rice, 1986–2017)

Commodity	Maize	Rice	Soybeans	Wheat
Maize				
Rice	0.662			
Soybeans	0.858	0.772		
Wheat	0.790	0.611	0.776	
Grains	0.923	0.688	0.887	0.959

Notes: Prices are detrended using a natural cubic spline using four knots. "Grains" includes the caloric aggregate of maize, soybeans, and wheat.

Table 4 reports the autocorrelations and standard deviations in the data used to estimate the model. The first-order autocorrelations of spot and futures prices are both greater than 0.57. It was the inability of the storage model to match these high serial correlation levels in prices for a range of storable commodities that originally led Deaton and Laroque (1992, 1996) to reject the storage model. Consumption persistence is also substantial with a first order autocorrelation coefficient of 0.64 which suggests the inclusion in the model of a persistent demand shock.

Production and yield shocks have small and insignificant autocorrelation in line with our model assumption of supply shocks without serial correlation.

Table 4 – Autocorrelation and standard deviation of log detrended caloric data, 1961–2017

Variable	One-year autocorrelation	Two-year autocorrelation	Standard deviation
Demand price ($\log(p_t)$)	0.576	0.167	0.236
Supply price ($\log(E_t p_{t+1})$)	0.652	0.236	0.192
Consumption ($\log(c_t)$)	0.642	0.302	0.019
Production ($\log(q_t)$)	0.042	−0.095	0.028
Yield shock (ψ_t)	0.148	0.050	0.023

The pattern of the standard deviations is coherent with a storage model with small elasticities. The coefficient of variation of quantities is one order of magnitude lower than the coefficient of variation of prices. Consumption volatility is lower than production volatility, which is consistent with a smoothing by storage associated with larger supply than demand shocks. Put simply, without storage, yearly changes in production levels would have to be matched by corresponding variations in consumption levels. The standard deviation of the yield shock accounts for 82% of that of production, suggesting the importance of these shocks for the variations in production. Finally, the lower volatility of the expected compared to the spot price is as predicted and is consistent with the “Samuelson effect”: decreasing futures price volatility based on the contract horizon.

Table 5 displays the correlation coefficients of all the detrended variables in logarithm. The correlations with obvious counterparts in the model have the expected signs. Current and expected prices are strongly correlated, consistent with equation (4) in the presence of frequent stocks. The fact that production and consumption are not perfectly correlated is another indication of the role played by storage. The observed negative correlation between consumption and price suggests that the changes in consumption stem from movements along the demand curve and from shifts in the demand curve. Were they due only to changes along the demand curve the correlation would be close to -1 .

6. Estimation

6.1. Structural parameters

Before analyzing the results obtained by indirect inference in section 6.1.2, we report the 2SLS and OLS estimates of the supply and demand equations. These estimates provide direct values for some parameters (α_D , α_S , σ_v , ρ_μ , and σ_ϑ), and indirect information about the others.

Table 5 – Correlation coefficients of log detrended caloric data, 1961–2017

Variable	Demand price ($\log(p_t)$)	Supply price ($\log(E_t p_{t+1})$)	Consumption ($\log(c_t)$)	Production ($\log(q_t)$)
Demand price ($\log(p_t)$)				
Supply price ($\log(E_t p_{t+1})$)	0.935			
Consumption ($\log(c_t)$)	−0.488	−0.451		
Production ($\log(q_t)$)	−0.406	−0.270	0.395	
Yield shock (ψ_t)	−0.532	−0.498	0.527	0.775

6.1.1. Instrumental variable estimations

Supply and demand equations can be estimated on raw trending data. Following Roberts and Schlenker, we augment all the first and second stage equations with trend variables generated by natural cubic splines with three to five knots. Tables 6 and 7 present the supply and demand estimates. To enable comparison with Roberts and Schlenker (2013), we replicate these estimates in Appendix (Tables A7 and A8) for a shorter sample (1962–2007) which corresponds to the sample length they used. The Appendix tables have some minor differences with the Table 1 in Roberts and Schlenker. These are due to two deviations from their approach: a slightly different procedure to construct the yield shock (in line with Hendricks et al., 2015), and the detrending of yields using a 4-knot spline rather than a 3-knot spline which is more consistent with our longer sample.

Table 6 reports the estimations of the supply equation. For the 2SLS estimates, the Cumby-Huizinga test rejects the hypothesis of residuals without serial correlation. We nevertheless report standard errors and diagnostic tests that are robust only to heteroskedasticity. Not only this is the most conservative choice in this particular setting but it also allows us to use the same type of standard errors for the supply and demand equations as well as for the weighting matrix of the indirect inference approach. 2SLS estimates of the supply elasticity are around 0.08, slightly lower than the values obtained by Roberts and Schlenker (2013). However, comparison with table A7 shows that the difference is entirely explained by our longer sample. The c_q estimates are always above $1 + \alpha_S$ (although not significantly). According to the discussion in section 3.1.1, this indicates a negative correlation between the two planting-time shocks (η and ω). The estimations using four and five knots are similar but present small differences with the estimations using three knots which is in line with the previous stationarity test results. Consistent with Hendricks et al.'s (2015) insights, the OLS and 2SLS supply elasticity estimates show only small and insignificant differences indicating that using the yield shock as a control variable helps to mitigate the omitted variable bias. This is further confirmed by the Hausman test which fails to reject the null of exogenous expected prices. However, the Monte Carlo analysis shows that this result was to be expected in such short samples even with endogenous prices. Therefore, we do not follow the Hausman test and for the comparisons that will follow

Table 6 – Supply equation estimation

	(1)	(2)	(3)
<i>Panel A. 2SLS</i>			
Supply elasticity b_q	0.088** (0.038)	0.075*** (0.026)	0.082*** (0.026)
Shock c_q	1.153*** (0.194)	1.154*** (0.141)	1.137*** (0.150)
<i>Panel B. First stage</i>			
Lagged shock b_{E_p}	-4.045*** (1.474)	-3.783*** (0.991)	-3.821*** (0.993)
Shock c_{E_p}	-2.470 (1.927)	-2.382* (1.382)	-2.343* (1.334)
<i>Panel C. OLS</i>			
Supply elasticity b_q	0.135*** (0.014)	0.058*** (0.013)	0.061*** (0.012)
Shock c_q	1.298*** (0.154)	1.103*** (0.099)	1.078*** (0.107)
$\sigma_{u_q^{2SLS}}$	0.028	0.015	0.015
$\sigma_{\vartheta^{2SLS}}$	0.038	0.031	0.030
$\sigma_{u_{E_p}}$	0.228	0.165	0.166
$\sigma_{u_q^{OLS}}$	0.026	0.015	0.015
$\sigma_{\vartheta^{OLS}}$	0.039	0.030	0.029
First stage F -stat	7.531	14.567	14.811
p -value for Hausman test	0.172	0.414	0.302
p -value for Cumby-Huizinga test (panel A)	0.000	0.004	0.004
Observations	56	56	56
Spline knots	3	4	5

Notes: Standard errors robust to heteroskedasticity in parenthesis. ***, **, and * indicate significance at the 99%, 95%, and 90% levels, respectively.

our benchmark estimate is the 2SLS with four knots. For this specification, total supply shocks have a standard deviation σ_{ϑ} equal to 0.031, a value slightly above the standard deviation of production in table 4.

Table 7 presents the estimation results of the demand equation. The demand elasticity estimates are higher in absolute values than in Roberts and Schlenker (2013), which again seems to result from using a longer sample (see table A8). We use equation (40) to estimate both the demand elasticity and autocorrelation of the demand shock. This contrasts with Roberts and Schlenker (2013) who use equation (38) which identifies only the demand elasticity. By comparing the results in panels A and D, we see that the estimates do not differ significantly between these

Table 7 – Demand equation estimation

	(1)	(2)	(3)
<i>Panel A. 2SLS</i>			
Demand elasticity b_c	-0.051* (0.028)	-0.065** (0.026)	-0.060** (0.027)
Lagged price c_c	0.041** (0.016)	0.019 (0.014)	0.014 (0.014)
Lagged demand d_c	1.054*** (0.070)	0.535*** (0.159)	0.442** (0.203)
<i>Panel B. First stage</i>			
Shock b_p	-4.287*** (0.882)	-4.112*** (0.937)	-4.014*** (1.056)
Lagged price c_p	0.569*** (0.087)	0.486*** (0.105)	0.498*** (0.111)
Lagged demand d_p	1.446* (0.745)	-0.130 (1.690)	0.523 (2.012)
<i>Panel C. OLS</i>			
Demand elasticity b_c	-0.012 (0.010)	-0.021** (0.010)	-0.018* (0.010)
Lagged price c_c	0.015 (0.010)	-0.005 (0.011)	-0.010 (0.011)
Lagged demand d_c	0.949*** (0.044)	0.547*** (0.118)	0.413** (0.162)
<i>Panel D. 2SLS using Roberts and Schlenker's approach (eqs. (38) for 2nd stage and (42) for 1st)</i>			
Demand elasticity b_c	-0.069 (0.049)	-0.079*** (0.023)	-0.066*** (0.023)
$\sigma_{u_c^{2SLS}}$	0.018	0.016	0.016
σ_{u_p}	0.180	0.180	0.180
$\sigma_{u_c^{OLS}}$	0.016	0.014	0.013
$\sigma_{u_c^{2SLS, RS}}$	0.049	0.020	0.017
$\sigma_{\mu^{2SLS}}$		0.019	0.018
First stage F -stat (panel A)	23.627	19.252	14.443
p -value for Hausman test (panel A)	0.137	0.043	0.054
p -value for Cumby-Huizinga test (panel A)	0.851	0.199	0.057
First stage F -stat (panel D)	16.668	27.501	22.935
p -value for Hausman test (panel D)	0.000	0.029	0.052
p -value for Cumby-Huizinga test (panel D)	0.000	0.014	0.045
Observations	56	56	56
Spline knots	3	4	5

Notes: Standard errors robust to heteroskedasticity in parenthesis, except for panel D where they are also robust to autocorrelation. The lagged demand estimates in panel A are bias adjusted (Orcutt and Winokur, 1969). ***, **, and * indicate significance at the 99%, 95%, and 90% levels, respectively.

two approaches.²¹ The Cumby-Huizinga test cannot reject the hypothesis of absence of serial

²¹Monte Carlo simulations (not reported here) show that using equation (40) instead of equation (38) leads to

correlation in the residuals for equation (40) (but not for equation (38)), which is a necessary condition for the consistent estimation of autoregressive terms. Estimates of the autocorrelation of the demand shocks differ depending on the number of knots. ρ_μ estimated along with a 3-knot spline is not statistically different from 1 indicating a non-stationary demand, which confirms the results in section 5.2 which shows that a 3-knot spline is not sufficiently flexible to obtain stationary series. A higher number of knots reduces ρ_μ by reducing the autocorrelation in the data, but at 0.53 (0.16) and 0.44 (0.20) for four and five knots the estimates are similar. The last parameter which can be identified from the demand estimation is the standard error of the demand shock. Using 4- and 5-knot splines, σ_v (estimated by $\sigma_{u_c^{2SLS}}$) is about 0.016, which is slightly lower than the volatility of consumption observed in the raw data reported in table 4.

With the exception of the supply equation with three knots all first-stage F -statistics exceed the standard threshold of 10. For the first-stage of supply, the coefficient of contemporaneous yield shock is negative which is consistent with a positive supply shock decreasing the prices but barely significant, indicating the limited predictability of yield shocks. The coefficient of the lagged yield shock is negative and significant because a lagged positive supply shock increases current availability through its effect on storage and thus depresses prices. Similarly, the supply shock in the first-stage of the demand equation is significantly negative.

Were the residuals of the demand and supply equations correlated, a more efficient strategy would be a three-stage least squares (3SLS). For the three degrees of flexibility considered, the correlation between the residuals is small at 0.16, -0.09 , and -0.09 . This low correlation means that the 2SLS and 3SLS results are very similar and thus the latter are not reported here. Since the standard deviation of the residuals of the supply equation σ_{u_q} can be expressed as a function of the various supply shocks, the lack of correlation between the residuals supports our assumption of no correlation between demand innovations v_t and supply shocks.

6.1.2. Indirect inference estimations

We followed Roberts and Schlenker by presenting the instrumental variable results for natural cubic spline trends with three to five knots. However, both the unit-root tests and the estimates from table 7 suggest that 3-knot spline estimations could be problematic since the trend is not sufficiently flexible to stationarize the series. Moreover, a 3-knot spline creates numerical problems in the indirect inference approach because the storage model is difficult to solve for values of ρ_μ close to 1. Hence, in the following indirect inference approach, we vary the number of knots only between four and five. The estimation results using the auxiliary model based on OLS regressions are presented in table 8.

Most parameters are estimated precisely for both trend specifications despite the rather short

slightly smaller RMSE, consistent with the fact that more spherical residuals should make the estimator more efficient.

Table 8 – Estimation results for the indirect inference approach (auxiliary model based on OLS regressions)

	4-knot spline		5-knot spline	
	Estimate	Standard error	Estimate	Standard error
ρ_μ	0.702	(0.068)	0.681	(0.079)
$\rho_{\eta,\omega}$	-0.442	(0.307)	-0.370	(0.331)
σ_ω	0.188	(0.031)	0.185	(0.030)
σ_η	0.014	(0.006)	0.014	(0.006)
σ_ϵ	0.020	(0.005)	0.020	(0.005)
σ_ν	0.019	(0.003)	0.018	(0.003)
δ	0		0	
k	0.037	(0.014)	0.034	(0.012)
α_D	-0.068	(0.019)	-0.059	(0.019)
α_S	0.086	(0.016)	0.085	(0.016)
σ_φ	0.027	(0.005)	0.026	(0.005)
σ_ψ	0.025	(0.002)	0.025	(0.002)
σ_μ	0.026	(0.005)	0.024	(0.005)
σ_ϑ	0.034	(0.004)	0.033	(0.004)
OID p -value	0.084		0.097	

Notes: $\sigma_\varphi = \sqrt{(1 + \alpha_S)^2 \sigma_\eta^2 + (\alpha_S \sigma_\omega)^2 - 2\rho_{\eta,\omega} \alpha_S (1 + \alpha_S) \sigma_\eta \sigma_\omega}$, $\sigma_\psi = \sqrt{\sigma_\eta^2 + \sigma_\epsilon^2}$, $\sigma_\mu = \sigma_\nu / \sqrt{1 - \rho_\mu^2}$, and $\sigma_\vartheta \equiv \sqrt{\sigma_\epsilon^2 + \sigma_\varphi^2}$. The standard errors of σ_φ , σ_ψ , σ_μ , and σ_ϑ are calculated using the Delta method.

sample size.²² The exceptions are the correlation between the planting-time shocks ($\rho_{\eta,\omega}$), and to a lesser extent the per-unit storage cost (k). The fact that $\rho_{\eta,\omega}$ is not precisely estimated is not surprising, given the large RMSE values obtained in table 2 for the Monte Carlo analysis, and given the lack of precision in table 6 of the estimates of $c_q - 1$ from which $\rho_{\eta,\omega}$ is derived.

The parameters estimated using both methods (i.e., ρ_μ , σ_ν , α_D , α_S , and σ_ϑ), do not differ significantly across methods but precision is greater with indirect inference as suggested by the Monte Carlo studies. Although not significantly different from the 2SLS estimates, the indirect inference estimate of ρ_μ is sufficiently higher to be a potential concern and could indicate some misspecification of the model on the demand side. This is confirmed later by the limited fit of some demand-related moments.

The volatility of the cost shock σ_ω is about 19% which is an order of magnitude larger than the estimates of the other shocks. However, the cost shock has no direct effect on quantities.

²²Since δ is estimated at its lower bound, it is not possible to calculate its standard error, and the model is too costly to optimize to do it by bootstrap.

Making it comparable to the other shocks requires its multiplication by α_S which produces 1.6% with four and five knots that is a contribution similar to the planting-time yield shock $((1 + \alpha_S)\sigma_\eta)$. In the Monte Carlo analysis, such a large cost shock would make the 2SLS estimation of the supply equation very imprecise because the lagged yield shock would be a weak instrument, and could also create a wide gap between the OLS and the 2SLS estimates. This is not fully consistent with the results in table 6 where the OLS and 2SLS estimates are similar, indicating possible overestimation of σ_ω . The planting-time shocks η and ω can be aggregated in the shock φ . The standard deviation of φ exceeds the standard deviation of harvest-time yield shock σ_ϵ , which indicates that the majority of supply shocks is known before deciding to produce. Finally, these three supply shocks can be aggregated together. The last row in table 8 shows that the standard deviation of the resulting total supply shock ϑ is about 30% larger than the standard deviation of the demand shock, μ .²³

Storage costs are estimated to be only composed of per-unit cost without any shrinkage with δ estimated at its lower bound of 0, which confirms a similar result in Cafiero et al. (2011). However, the Monte Carlo analysis has shown that while the model is able to recover the average storage cost, it may not be able to distinguish between per-unit cost and shrinkage. The physical per-unit storage cost (k) is estimated a 3.7% of the steady-state price with four knots. By combining the opportunity costs related to the interest rate and the price trend, we obtain an estimated total annual storage cost of around 7.6% at the steady state $(k + 1 - \beta(1 - \delta) e^{g_p})$. Note that estimating the model without a price trend—i.e., by setting $g_p = 0$ —barely changes the parameter estimates apart from the storage cost which increases by 2% which is exactly the opportunity cost implied by the downward price trend. The cost created by the positive quantity trend also contributes to higher storage costs but cannot be characterized analytically and so is ignored in this discussion.

Overall, these results suggest that our indirect inference approach returns fairly precise parameter estimates which are reasonably consistent with the 2SLS estimates. Since the differences across trend specifications are small, all the subsequent analyses are based on the estimation using the 4-knot spline, our preferred trend specification.

6.2. Inspecting the auxiliary model

The overidentification test cannot reject the model specification at the 5% threshold level, with the caveat that the test is biased against rejection of the null (according to the Monte Carlo experiment in table 2).²⁴ It is nonetheless interesting to check also the similarity between the

²³The VAR literature on oil price fluctuations has not settled on the respective role of demand and supply shocks with Kilian (2009) finding that oil price fluctuations have been mainly driven by demand shocks, Baumeister and Hamilton (2019) by supply shocks, and Caldara et al. (2019) finding equally important roles.

²⁴The overidentification test is not defined when a parameter is at a bound, so it is calculated by assuming that the model was estimated with the restriction $\delta = 0$, which implies that the statistics in equation (51) follows a $\chi^2(6)$.

estimates of the auxiliary model parameters based both on observations and simulations.

Table 9 reports the auxiliary parameters obtained respectively from the actual and the simulated data along with their standard errors estimated on the observations. Note that the standard errors column corresponds to the inverse of the square root of the diagonal of the weighting matrix, W . For each parameter we can calculate a t -statistic of equality of the coefficients and test for consistency of the auxiliary model (Gourieroux et al., 1993, Appendix 3). Apart from the OLS-estimated parameter, d_p from equation (41), we cannot reject the null of equality between the estimates based on observations and those based on simulations from the structurally estimated model. Although some parameters differ a lot between the two columns (e.g., $b_{E,p}$ or $c_{E,p}$), they are estimated imprecisely in the auxiliary model, and thus were given a small weight in the objective function which the indirect inference procedure minimizes.

Table 9 – Coefficients of the OLS auxiliary model: estimation based on observations versus based on simulations

Coefficient	Observations		Model
	Estimate	Standard error	Estimate
b_q	0.058	0.013	0.048
c_q	1.103	0.099	1.148
σ_{u_q}	0.015	0.001	0.015
b_c	-0.021	0.010	-0.007
c_c	-0.005	0.011	0.011
d_c	0.547	0.118	0.534
σ_{u_c}	0.014	0.001	0.014
$b_{E,p}$	-2.382	1.382	-1.687
$c_{E,p}$	-3.783	0.991	-2.303
$\sigma_{u_{E,p}}$	0.165	0.016	0.160
b_p	-4.112	0.937	-4.445
c_p	0.486	0.105	0.456
d_p	-0.130	1.690	2.881*
σ_{u_p}	0.180	0.018	0.180
$\sigma_{u_{\psi}}$	0.023	0.002	0.025
b_q^{2SLS}	0.075	0.021	0.086
b_c^{2SLS}	-0.065	0.026	-0.068

Notes: Standard errors robust to heteroskedasticity for the parameters and based on equation (50) for the standard deviations. The lower panel presents the parameters estimated by 2SLS not present in the auxiliary model used for the estimation. * indicates significant difference between the estimates based on observations and those based on simulations at the 90% level.

Although the auxiliary model used here involves only OLS estimations, it is useful to compare

also the fit with the supply and demand elasticities estimated by 2SLS (see lower panel in table 9). The fit is very good for both elasticities. However, we can note that the model tends to overestimate b_q^{2SLS} and to underestimate b_q estimated by OLS (albeit insignificantly for both). Since the difference between the supply elasticities estimated by OLS and 2SLS is supposed to increase with σ_ω , this difference could confirm the possible overestimation of σ_ω highlighted above.

These results suggest an overall good fit of the auxiliary model between observations and simulations, with the exception of one demand-side parameter.

6.3. Inspecting the model fit on other moments

We next assess the performance of the estimated storage model by comparing the variances and covariances based on model simulations and those based on observations (as typically done following the estimation of DSGE models, e.g., Smets and Wouters, 2003). Recall that so far the empirical performance of estimated storage models was judged based only on their ability to replicate price-based moments given that only prices were used for the estimations. By focusing on second-order moments calculated up to one lag for each of our 5 observables, our empirical setting now allows evaluation of the model fit over 40 moments. The results of this exercise are presented in table 10 which includes all the moments calculated on the detrended observations, their standard deviation calculated by bootstrap, the corresponding moments from the simulated model, and an indication of whether the simulated moment lies within the bootstrap confidence intervals of the observed moment. Note that some of these moments were included in the auxiliary model—either directly (σ_ψ as σ_{u_ψ}) or indirectly ($\phi_{\ln p}(1)$ as c_p)—but many others were not and therefore constitute a good test of the model's overall quantitative performance. The majority of the moments are similar for observations and simulations, indicating that our extended storage model is generally able to replicate the main dynamics in the data. This applies in particular to the first-order autocorrelation of price, the subject of long-standing debates since Deaton and Laroque (1992).²⁵

However, it can be seen that the storage model fails to match some moments (14 lie outside the 10% bootstrap confidence interval including 11 outside the 5% confidence interval). These moments mostly relate to two aspects. Six moments are related to consumption and its (lagged) covariance with current and expected prices. In particular, the model fit related to the negative correlation between consumption and spot prices is problematic: $\text{cor}(\ln p_t, \ln c_t) = -0.49$ on observations but 0.08 on simulations. Logically, given the strong autocorrelation of both prices and consumption combined with the strong correlation between current and expected prices, this issue persists with a lag as well as if we consider expected instead of current prices. A similar problem arises for four moments related to production and its (lagged) covariance with current and expected prices.

²⁵However, this is not surprising since this moment was included in the objective function through the parameter c_p .

Table 10 – Comparison of actual and model-based second-order moments

Moment	Observed	Standard deviation	Simulated
$\sigma_{\ln p}$	0.236	0.023	0.262
$\sigma_{\ln c}$	0.019	0.002	0.018
$\sigma_{\ln q}$	0.028	0.002	0.031
$\sigma_{\ln E p}$	0.193	0.018	0.180
σ_{ψ}	0.024	0.002	0.025
$\phi_{\ln p}(1)$	0.576	0.110	0.559
$\phi_{\ln c}(1)$	0.642	0.146	0.568
$\phi_{\ln q}(1)$	0.042	0.140	-0.011
$\phi_{\ln E p}(1)$	0.652	0.116	0.607
$\phi_{\psi}(1)$	0.146	0.142	0.001
$\phi_{\ln p, \ln c}(0)$	-0.488	0.102	0.083***
$\phi_{\ln p, \ln q}(0)$	-0.406	0.103	-0.183**
$\phi_{\ln p, \ln E p}(0)$	0.939	0.017	0.871***
$\phi_{\ln p, \psi}(0)$	-0.534	0.118	-0.454
$\phi_{\ln c, \ln q}(0)$	0.395	0.109	0.590*
$\phi_{\ln c, \ln E p}(0)$	-0.452	0.106	0.283***
$\phi_{\ln c, \ln \psi}(0)$	0.529	0.116	0.463
$\phi_{\ln q, \ln E p}(0)$	-0.271	0.115	-0.025**
$\phi_{\ln q, \psi}(0)$	0.775	0.050	0.831
$\phi_{\ln E p, \psi}(0)$	-0.500	0.118	-0.292
$\phi_{\ln p, \ln c}(1)$	-0.469	0.125	0.191***
$\phi_{\ln p, \ln q}(1)$	0.104	0.156	-0.015
$\phi_{\ln p, \ln E p}(1)$	0.643	0.069	0.627
$\phi_{\ln p, \psi}(1)$	-0.274	0.142	-0.183
$\phi_{\ln c, \ln p}(1)$	-0.326	0.109	0.205***
$\phi_{\ln c, \ln q}(1)$	0.184	0.110	0.299
$\phi_{\ln c, \ln E p}(1)$	-0.300	0.118	0.181***
$\phi_{\ln c, \psi}(1)$	0.304	0.127	0.187
$\phi_{\ln q, \ln p}(1)$	-0.257	0.110	0.216***
$\phi_{\ln q, \ln c}(1)$	0.323	0.110	0.352
$\phi_{\ln q, \ln E p}(1)$	-0.212	0.116	0.092**
$\phi_{\ln q, \psi}(1)$	0.067	0.134	-0.143*
$\phi_{\ln E p, \ln p}(1)$	0.566	0.094	0.534
$\phi_{\ln E p, \ln c}(1)$	-0.508	0.116	0.293***
$\phi_{\ln E p, \ln q}(1)$	0.070	0.147	0.043
$\phi_{\ln E p, \psi}(1)$	-0.358	0.129	-0.138*
$\phi_{\ln \psi, \ln p}(1)$	-0.162	0.108	-0.120
$\phi_{\ln \psi, \ln c}(1)$	0.334	0.127	0.123
$\phi_{\ln \psi, \ln q}(1)$	-0.115	0.122	0.002
$\phi_{\ln \psi, \ln E p}(1)$	-0.203	0.115	-0.226

Notes: Moments calculated over 100,000 sample observations from the asymptotic distribution simulated with a storage model calibrated with the indirect inference estimates with a 4-knot spline from table 8. $\phi(1)$ denotes first-order serial correlation and $\phi_{i,j}(l) = \text{cor}(i_{t-l}, j_t)$ denotes l^{th} -order correlation between variable i and j . Statistics involving $E p$ refer to $E_t p_{t+1}$, e.g., $\phi_{\ln p, \ln E p}(0) = \text{cor}(\ln p_t, \ln E_t p_{t+1})$. Standard deviation calculated by bootstrapping the dataset of detrended variables using 5,000 bootstrap replicas. ***, **, and * indicate that the simulated moment is outside the 99%, 95%, and 90% bootstrap confidence interval (adjusted bootstrap percentile method), respectively.

The correlation between consumption and price is governed in the model by the demand elasticity and the relative size of the supply and demand shocks. Indeed, in the absence of demand shocks the correlation would be -1 . The higher the variance of demand shocks, the higher the correlation which can even turn positive for demand shocks with a sufficiently large variance. The indirect inference estimations lead to higher demand shock autocorrelation and larger variance of demand shocks compared to those obtained using 2SLS. These differences between 2SLS and

indirect inference could contribute to explaining the difficulty related to fitting the consumption-price correlation and confirm a likely model misspecification on the demand side.

Similar mechanisms apply to the correlation between production and prices, which is governed by the supply elasticity and the relative size of demand and supply shocks. Then again, without supply shocks and a positive supply elasticity, production and prices would be positively correlated as production would increase with demand shocks. At the other extreme, without demand shocks and an inelastic supply, the correlation would be negative as supply shocks would depress prices. Hence, the inability to match the negative correlation between production and price could also come from demand shocks too large relative to supply shocks, which would be consistent with the previous problem.

6.4. Sensitivity analyses

In this sensitivity analysis, we discuss the role of the auxiliary model, the storage cost, and the data. Appendix C displays all the tables.

The main results are based on an auxiliary model in which all equations are estimated by OLS, so involving biased parameters. We now compare with the results obtained when the supply and demand equations of the auxiliary model are estimated by 2SLS. The latter estimates are available in table A9 in the Appendix. Comparing tables 8 and A9, most of the parameter estimates are not significantly different, but the elasticities deviate more from the 2SLS benchmark in the case of the auxiliary model estimated by 2SLS. Corroborating the Monte Carlo results and the intuition that the moments from 2SLS estimates are noisier, the indirect inference based on the 2SLS supply and demand equations delivers several estimates that are much less precise.

Storage costs are challenging to estimate with our approach. The main results conclude on the absence of shrinkage, but the Monte Carlo analysis shows that the precision of such a conclusion is limited. To assess the robustness of the results to the type of storage costs, we present in the first column of table A10 the results of an estimation imposing zero per-unit storage costs (i.e., $k = 0$) and only shrinkage as in Deaton and Laroque (1992, 1996). This constraint barely affects any estimate while the shrinkage rate δ is found almost equal to the per-unit cost. The only difference between the two models is that the unrestricted model presents a smaller value of the objective (which given equation (51) can be seen in its higher p -value for the OID test).

For the data, we have followed Roberts and Schlenker (2013) and considered the four most important crops for quantities but excluded rice of the price index because of the short sample of rice price futures. This could be a concern if the rice market behaves differently from the other markets. To verify this, we estimate the model on data from which the rice sector has been removed altogether. Similar results are obtained (table A10) except that all shocks and elasticities are higher in absolute values (albeit not significantly). This could be explained by the fact that rice consumption and production are more stable than for the other crops, because of

its almost exclusive use for food consumption and its large share of irrigated production which limits production shocks.

In addition to FAOSTAT, it is possible to obtain almost-global information about quantities from the USDA. The USDA Production, Supply and Distribution (PSD) database (USDA, 2020) provides information about a smaller sample of countries, which excludes some countries with minor contribution to the global food balance. Although it also allows using a longer sample, for comparability with FAOSTAT data, we maintain the same 1961–2017 sample. 2SLS and indirect inference results with USDA-PSD are available in table A10. Results based on USDA data are extremely similar to those based on FAOSTAT data whatever the estimator. One noticeable difference is the more elastic demand. This difference could be due to differences in stock changes data, which would appear here as a difference in consumption, given that production is less susceptible to measurement errors. However, this different estimate does not help improve the model fit as studied in table 10.

We have carried out our estimations on a caloric aggregate following Roberts and Schlenker (2013), because it provided us with a 2SLS benchmark to compare our indirect inference estimates. However, aggregating commodities may create bias. In table A11, we present the crop by crop results of 2SLS and indirect inference estimates for maize, soybeans, and wheat (the sample for rice is too short to obtain reliable results). At the commodity level, only the elasticities of maize are significant when estimated by 2SLS, but they are all precisely estimated with indirect inference, except for the demand elasticity of soybeans. Demand and supply are more elastic (except for the supply of wheat) at the commodity level, consistent with the idea that these crops are substitutes. The shocks tend also to be much larger since they are no longer smoothed by the aggregation. The conclusion that supply shocks are larger than demand shocks remain, but with a smaller difference between them. This could be explained by the fact that at the commodity level a supply shock for one commodity can become a demand shock for another: for example a bad wheat harvest could create increased demand for maize for feed.

Finally, to be transparent regarding the identification, we report Andrews et al.'s (2017) measure of sensitivity of the estimates to the estimating moments in tables A12 and A13. This measure is calculated as $\Lambda = -(J'WJ)^{-1}J'W$ and describes how estimated parameters change with the moments. To normalize this measure as elasticities of changes in parameters with respect to moments, we display $\text{diag}(\hat{\theta}^{-1})\Lambda \text{diag}(\hat{\zeta}_T)$, where diag transforms a vector in a diagonal matrix. Table A12 presents this measure for our benchmark estimation. It confirms the intuitions about identification laid out in section 3.2, but it also shows that for most parameters, identification actually comes from a combination of moments. For example, the four moments to which the supply elasticity is the most sensitive are the three moments associated with the supply equation (b_q , c_q , and σ_{u_q}), with the expected signs, as well as σ_{E_p} because more volatile expected prices would imply a lower supply elasticity. Table A13 presents this measure of sensitivity for the estimation without per-unit storage costs. Comparing the two tables shows that k and δ present very similar sensitivity to the observed moments which explains the difficulty to estimate them

separately.

7. Applications

Having demonstrated that, apart from the demand-related misspecification mentioned above, our rich storage model shows a reasonable fit with the data of the global grains market, we can use it to address various questions linked to the role of speculative storage in the formation and behavior of commodity prices in the world market. In particular, how do the different model components interact with one another and drive the implied dynamics? What are the relative contributions of the various supply and demand structural shocks to price and quantity developments in the global grains market? What are the expected welfare effects of speculative demand for storage? These issues are studied in turn in the succeeding subsections.

7.1. The role of storage in market dynamics

The introduction of many new features in our storage model calls for investigation of their respective contributions to the price and quantity dynamics generated by the model. In this section, we explore the role of storage in the movement of prices based on the alternative exclusion of the various model features. For reasons of space, we restrict the discussion to six moments of interest: price autocorrelation which since Deaton and Laroque (1992) is the benchmark metric used to assess the performance of the storage model, price, consumption as well as production volatilities, and the correlation between price and consumption, and price and production. $\phi_{\ln p, \ln c}(0)$ and $\phi_{\ln p, \ln q}(0)$ are of particular interest because in the previous section we showed that the model struggles to match these moments; thus, it is helpful to examine which model characteristics is driving their behavior. Table 11 reports the results of this exercise as well as the same moments calculated for comparison on the raw and detrended data.

Switching off the model features one at a time allows us to quantify their respective contribution to price persistence. The trend captured by the 4-knot spline explains one third of the 0.87 one-year autocorrelation in the raw data. Regarding the remaining serial correlation explained by the benchmark model, the simulations of the various models show that the three features which matter most for this moment are the autocorrelation coefficient of the demand innovations (model 2), the presence of planting-time shocks (model 11), and the smoothing effect of storage (model 15). Because of their interactions, turning off each feature leads to contributions that sum to more than 100% and so we normalize each contribution by the total. Demand shock persistence explains 20% of the price autocorrelation, planting-time shocks account for 5%, and storage accounts for the remaining 42%. So storage, while key to induce price persistence, explains less than half of the actual serial correlation, which means that the other model features matter too. This result contrasts with Deaton and Laroque's (1996) estimation results for a model with autocorrelated supply shocks. Indeed, they found that almost all the serial correlation in prices was attributable to shock persistence not speculative storage. The difference with our

Table 11 – Role of model assumptions in price and quantity dynamics

Data or model	$\phi_{\ln p}(1)$	$\sigma_{\ln p}$	$\sigma_{\ln c}$	$\sigma_{\ln q}$	$\phi_{\ln p, \ln c}(0)$	$\phi_{\ln p, \ln q}(0)$
Trending data	0.87	0.46	–	–	–	–
Detrended data	0.58	0.24	0.019	0.028	–0.49	–0.41
1. Benchmark	0.56	0.26	0.018	0.031	0.09	–0.18
2. $\rho_\mu = 0$	0.38	0.21	0.017	0.029	–0.33	–0.50
3. $\rho_\mu = 0, \sigma_v = \sigma_\mu$	0.38	0.23	0.022	0.030	–0.04	–0.38
4. $\alpha_S = 0$	0.65	0.30	0.014	0.024	0.12	–0.16
5. $g_q = 0$	0.56	0.26	0.017	0.031	0.08	–0.17
6. $g_p = 0$	0.60	0.24	0.018	0.032	0.19	–0.14
7. $k = 0.018$	0.60	0.24	0.018	0.032	0.19	–0.14
8. $\sigma_\eta = 0$	0.53	0.25	0.017	0.027	0.19	–0.12
9. $\sigma_\omega = 0$	0.54	0.25	0.016	0.026	0.20	–0.11
10. $\sigma_\eta = 0, \sigma_\epsilon = \sigma_\psi$	0.52	0.26	0.017	0.031	0.09	–0.19
11. $\sigma_\omega = \sigma_\eta = 0, \sigma_\epsilon = \sigma_\psi$	0.51	0.26	0.017	0.028	0.15	–0.16
12. $\rho_\mu = 0, \sigma_v = \sigma_\mu, \alpha_S = 0$	0.24	0.20	0.022	0.027	0.06	–0.36
13. $\rho_\mu = 0, \sigma_v = \sigma_\mu, \alpha_S = 0, \sigma_\eta = 0, \sigma_\epsilon = \sigma_\psi, g_q = 0$	0.25	0.20	0.022	0.028	0.00	–0.40
14. $\rho_\mu = 0.535, \sigma_v = 0.016$	0.47	0.23	0.015	0.030	–0.25	–0.39
15. $k = \infty$	0.16	0.45	0.025	0.025	–0.58	–0.58

Notes: Moments calculated over 100,000 sample observations from the asymptotic distribution simulated with models calibrated with the indirect inference estimates with 4-knot spline from table 8, except for the parameter values indicated in the first column.

results lies in our use of quantities as observables: this ensures that any shock autocorrelation must be compatible with the quantity dynamics, which is not the case if we only use information contained in prices. Planting-time shocks contribute to price persistence by linking periods. More precisely, shocks at planting time affect production and therefore the prices in the next period, but since they are immediately observed they also affect current prices because of the intertemporal link created by storage. The presence of a supply response has an ambiguous effect on price autocorrelation, and is excluded from the above decomposition. If we compare the benchmark setup with model 4, we can see that an elastic supply decreases price serial correlation. On the other hand, in the absence of an autoregressive exogenous demand process—i.e., comparing models 3 and 12—a supply response increases price persistence.

The simulations of the estimated model raise a new puzzle about the inability of the model to match the price–consumption correlation. This moment is explained by the respective roles of the demand and supply shocks in driving price movements, combined with the demand elasticity. At the extreme without demand shocks, the correlation would be -1 . Therefore, removing planting-time shocks (models 8–11) or the supply response (model 4) would only decrease the role of supply shocks and exacerbate the problems related to this moment. Some improvement can be achieved by removing the persistence of the demand shock (models 2–3) or increasing the storage cost (model 15), but both lead to a lower fit of the price autocorrelation. The indirect inference approach overestimates ρ_μ by 0.168 and σ_v by 0.003 compared with the 2SLS

approach. Comparing models 2 and 3 with the benchmark shows that overestimation of ρ_μ would contribute only a little to solving this puzzle. However, setting the size of the demand shock equal to its 2SLS estimate level, in addition to ρ_μ (model 14), would bring the simulated moment closer to the observed moment, inside the 99% bootstrap confidence interval but outside the 95% interval. In other words, the covariance mismatch between consumption and price might be due in part to the overstatement of both the persistence and variance of the demand shocks.

The inability of the estimated model to match the price-production correlation also seems to be related to the demand-side estimates. Then again, setting ρ_μ and σ_v to their 2SLS values is enough to obtain a perfect fit of this moment.

Price volatility is well explained by the model if we remove the large share of this volatility caused by the trend (as shown for two other commodities in Bobenrieth et al., 2021). Storage explains the order of magnitude of the price fluctuations. Indeed, without storage, the price volatility implied by our model would be 73% higher (model 15). The other model components contribute much less but in the expected direction. For example, the autocorrelation of the demand innovations reduces the ability of storage to smooth these shocks. Indeed, compared with the benchmark model 1, the price variance is lower in model 3 when shock to consumption demand μ_t collapses to an i.i.d. normal error term. Thus, speculative storage can smooth transitory shocks but is less efficient in the case of persistent disturbances.

Overall, the effects of the various model features on consumption and production volatility have the expected signs. We next discuss the effects of the model variants not considered so far. In model 5, the positive trend on quantities g_q is removed. As discussed in section 2.2 this boils down to decreasing storage costs which slightly increases price persistence. In model 6, the negative trend on price g_p is removed. Because the price trend directly affects the storers' incentives, for a value similar to g_q it has a stronger impact. Comparing models 6 and 7 shows that its impact is very similar to the effect of a corresponding decrease in the per-unit storage cost k (i.e., decreasing it by $\beta(1 - \exp g_p)$). The specification of model 13 is the closest to the model estimated in Deaton and Laroque (1992): it includes neither persistent shocks nor planting-time disturbances, and includes an inelastic supply. In this specification, price autocorrelation decreases significantly from 0.56 to 0.25 which suggests that to match the true persistence of prices, estimation of the simpler version of the storage model considered in the literature so far would require lower storage costs.

7.2. Historical decomposition

The model can be used to perform a historical decomposition, i.e., to extract the various shocks from the series. This does not require indirect inference estimation per se. The linear regressions estimates would be sufficient as long as the residuals are given a structural interpretation as proposed in section 3.1. Figure 3 depicts the shocks that are identified along with the log deviations of their price and quantity trends. Our preferred estimates are from section 6.1.2:

the indirect inference estimates with 4-knot spline.²⁶

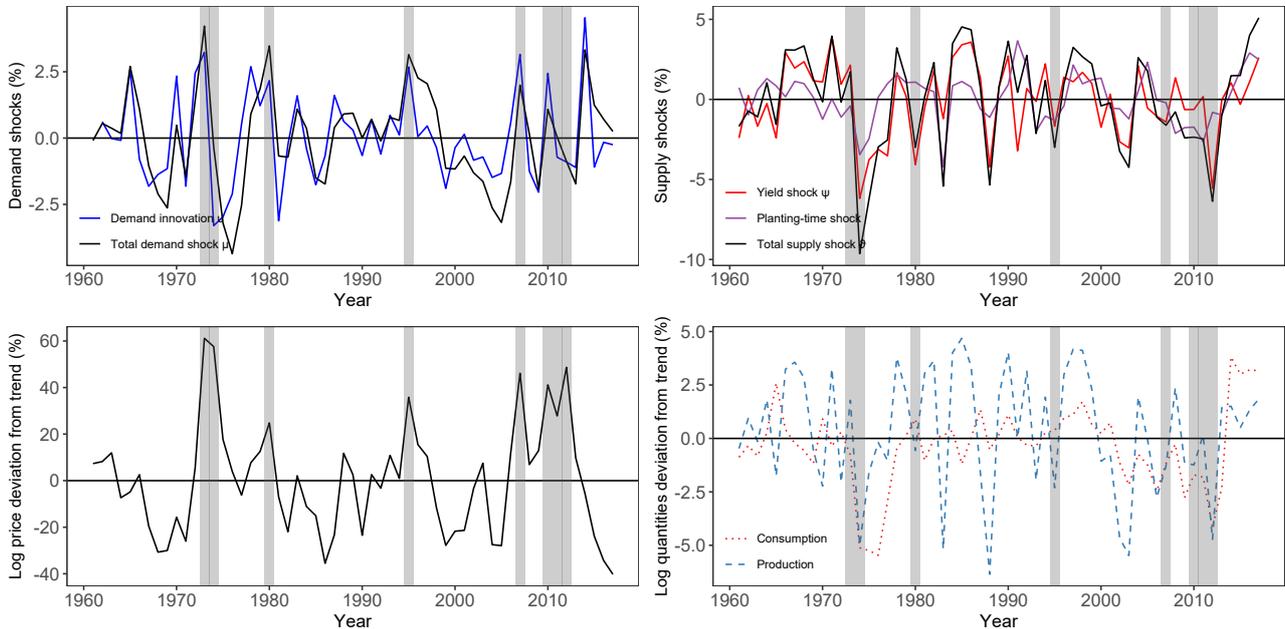


Figure 3 – Historical decomposition of the world price, production and consumption of grains into the various shocks. Gray areas denote price spike periods defined as log deviations from the trend greater than one standard deviation, 23.6%. The planting-time supply shock in purple corresponds to $\alpha_S(\eta_{t-1} - \omega_{t-1})$.

This decomposition helps to explain the market movements through our structural model lens. However, there is one missing piece which is stock levels, though as argued earlier the related statistics are unreliable at the global level. In the absence of storage, the effects of the shocks are not linked over time. Still, a couple of observations are warranted.

First, in line with the estimated standard deviations of the shocks, supply disturbances are larger than demand disturbances. However, all the price spikes are associated with large positive demand shocks. This applies also to the recent price spikes of 2007 and 2010–2, when the demand shocks took the form of biofuels mandates (see e.g., Roberts and Schlenker, 2013; Wright, 2014).

Second, there are seven years when total supply shocks ϑ are one standard deviation below the mean ($< -3.3\%$): 1974–5, 1983, 1988, 2002–3, and 2012, but in these seven years only two (1974 and 2012) correspond to price spikes. In all the other years, prices are close to their trends. This demonstrates the importance of storage to buffer against supply shortages. In the absence of inventories, a -3.3% supply shock would lead to a 63% price increase because inelastic consumption would have to respond one-to-one to the supply shortfall.

²⁶Figure 3 would nonetheless be very similar if created using instead the 2SLS estimates.

7.3. The welfare effect of private storage

In this section, we assess the welfare effect of storage in agricultural commodity markets. This issue is studied in depth in Wright and Williams (1984) but in the absence of credible estimations of the model parameters Wright and Williams (1984) use various calibrations.²⁷ In addition, they consider a simpler version of the model with only a harvest-time supply shock. However, they consider a more general inverse demand function than the simple constant elasticity function assumed here. Our structural estimates allow us to revisit this issue. In our model, welfare is defined as the sum of agents'—consumers, storers, and producers—surpluses. Produced and consumed quantities follow a trend, so that corresponding surpluses increase at the same rate. This implies that if the discount factor β is inferior to $\exp(-g_q)$, then intertemporal welfare will be diverging. This is the case here; to avoid this problem, we calculate welfare assuming no trends in either quantities or prices.

Using the notations from the detrended model, instantaneous welfare can be defined, up to an integration constant, as the sum of the following three surpluses:

$$w_t = \underbrace{-\bar{d} \frac{p_t^{1+\alpha_D}}{1+\alpha_D} \bar{p}^{-\alpha_D} e^{\mu t}}_{\text{Consumer surplus}} + \underbrace{(1-\delta) p_t x_{t-1} - (p_t + k\bar{p}) x_t}_{\text{Storer profit}} + \underbrace{p_t h_{t-1} e^{\eta_{t-1} + \epsilon_t} - \gamma(h_t) e^{\omega t}}_{\text{Producer profit}}. \quad (55)$$

We introduce storer profit because it is useful for the subsequent decomposition but due to the assumption of constant marginal storage cost, storers operate at zero profit (in expectations) so their average profit is zero. Dividing instantaneous welfare by the steady-state value of consumption $\bar{p}\bar{d}$ and using equation (12) for simplification, we can derive a unit-free expression of instantaneous welfare:

$$\frac{w_t}{\bar{p}\bar{d}} = \underbrace{-\frac{(p_t/\bar{p})^{1+\alpha_D}}{1+\alpha_D} e^{\mu t} + \frac{p_t S_t - x_t}{\bar{p} \bar{d}}}_{\text{Consumer efficiency gains}} \underbrace{-k \frac{x_t}{\bar{d}}}_{\text{Storage costs}} \underbrace{-\beta \frac{(h_t/\bar{d})^{1+1/\alpha_S}}{1+1/\alpha_S} e^{\omega t}}_{\text{Production costs}}. \quad (56)$$

In this expression, the terms are reorganized to provide a different decomposition. Since in such models one of the main welfare effects of storage is transfer between consumers and producers caused by a change in the mean price (Wright and Williams, 1984), it is useful to focus on efficiency. To do this, we correct the consumer surplus using the consumption value.

From the instantaneous welfare, we can calculate the intertemporal welfare normalized to an annual value by

$$W_t = (1 - \beta) w_t / (\bar{p}\bar{d}) + \beta E_t W_{t+1}. \quad (57)$$

²⁷In the absence of structural estimates, all past welfare applications of the storage model rely largely on calibrations (e.g., Gouel, 2013b), or a combination of estimation and calibration as in Steinwender (2018), Porteous (2019), and Gouel (2020).

Equation (56) can be evaluated over any state variables using the policy functions defined in section 2.4. Equation (57) is a Bellman equation evaluated using value function iterations. The resulting welfare is a function of the state variables. This welfare function is applied to the simulated observations to recover the expected welfare over the asymptotic distribution.

These welfare effects are presented in table 12 along with the two decompositions. It shows that with an increase of 0.31% in annual steady-state consumption, the overall welfare effects are modest. However, the distributional effects are large and show a 4.01% increase in consumer surplus and a corresponding decrease in producer profit. This change is related mostly to the change in the mean price: the presence of storage reduces the mean price by 3.68% compared to the situation without storage. Our assumption of a constant elasticity demand function means that a mean quantity preserving reduction in the consumption dispersion leads to a mean price decrease which explains the distributional effects. By abstracting from this transfer, the decomposition in the last three columns of the table displays only the efficiency changes and shows that the gains are shared more equally between consumer efficiency gains and production costs.

Table 12 – Welfare effects of introducing storage (expressed as a percentage of the steady-state consumption, $\bar{p}\bar{d}$)

Total	Consumer surplus	Producer profit	Consumer efficiency gains	Storage costs	Production costs
0.31	4.01	-3.70	0.27	-0.11	0.15

Notes: Calculated over 100,000 sample observations from the asymptotic distributions simulated with models calibrated with the indirect inference estimates with the 4-knot spline from table 8 except for the model without storage where we impose $k = \infty$.

The small size of the overall effects is related to the choice of a setting without market failures where risks do not matter. So the total effects are equal to the benefits derived from arbitrage: transferring the commodities from periods of low values to periods of high values. With risk averse agents (as in Gouel, 2013b), the welfare effects would be larger. Finally, note that this is only an assessment of the long-run welfare difference from introducing storage. It ignores any temporary welfare changes due to the transition between the steady-state distributions.

8. Conclusions

This paper proposes a new empirical strategy to estimate the rational expectations storage model. It requires five observables (current price, expected price, production, consumption, and supply shock) and reliance on a simple linear supply and demand model as the auxiliary model in an indirect inference approach. Including quantities as well as prices within the set of observables is crucial because it allows estimation of all the model parameters which is important to empirically validate the model and run counterfactual simulations for policy applications. Although the key role of storage for mediating the dynamics of commodity prices has long been

acknowledged and has been exploited widely in finance and economics, so far a full empirical validation of the storage model has not been carried out. To apply our approach, we chose the empirical setting of the global grains market following Roberts and Schlenker (2013), who use an instrumental variable strategy motivated by storage theory. While they estimate only a subset of the structural parameters, their strategy provides a good benchmark for comparing our indirect inference estimates. We also used their estimating equations to choose our auxiliary model.

Our results show that the long-standing price autocorrelation puzzle highlighted by Deaton and Laroque (1992, 1996) can be solved convincingly by accounting for sufficient features of the market for grains, such as (in decreasing order of importance): storage, a long-run price trend, autocorrelated demand shocks, and producers' incentive shocks associated with an elastic supply. We used our estimated model to quantify the relative size and contribution of the various structural disturbances to the boom and bust episodes recorded over the past 60 years. We found that total supply shocks are 30% larger than demand shocks, but that all price spikes have been associated with large positive demand shocks.

While our estimated storage model is able to rationalize many of the observed moments, it fails to reproduce the observed levels of the negative correlation between price and quantities. Finding a solution to this issue will be critical to estimate the model using full-information likelihood techniques which are likely to be more sensitive to such misspecification. Here, we can only speculate about possible sources of misspecification in our approach. A first is the aggregation of different commodities, which may introduce aggregation bias. A second is the deterministic arbitrage relationship assumed for storage which creates a stochastic singularity between price and expected price. This arbitrage equation is standard in the storage literature but there are alternatives that include a shock to the cost of storage such as in Knittel and Pindyck (2016). A third possible source of misspecification is the assumption that all wedges between quantities and prices are accounted for by structural shocks. This could be avoided by assuming the presence of measurement errors as is commonly assumed when estimating DSGE models (Canova, 2014). Despite these limits, this paper has proved that a simple storage model is able to capture the most important dynamic features of a global commodity market.

While the present paper follows Roberts and Schlenker (2013) and focuses on the grains market, our empirical methodology is fully applicable to the oil market or to any other storable commodity as long as there is an observable demand or supply shock (e.g., a demand shock based on freight rates as suggested by Kilian, 2009). This development could also help link the rational expectations storage literature to the estimation of VARs for commodity prices (e.g., Kilian and Murphy, 2014; Baumeister and Hamilton, 2019). Unlike the macroeconomic literature where the interaction between the VAR and DSGE modeling is fruitful, in research on commodity price dynamics the storage model has so far not been considered a relevant empirical model.

References

- Andrews, I., Gentzkow, M. and Shapiro, J. M. (2017). Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics*, 132(4), 1553–1592.
- 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.
- Baumeister, C. and Hamilton, J. D. (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *The American Economic Review*, 109(5), 1873–1910.
- Bobenrieth, E. S. A., Bobenrieth, J. R. A., Guerra, E. A., Wright, B. D. and Zeng, D. (2021). Putting the empirical commodity storage model back on track: Crucial implications of a “negligible” trend. *American Journal of Agricultural Economics*, 103(3), 1034–1057.
- Bobenrieth, E. S. A., Bobenrieth, J. R. A. and Wright, B. D. (2002). A commodity price process with a unique continuous invariant distribution having infinite mean. *Econometrica*, 70(3), 1213–1219.
- Bornstein, G., Krusell, P. and Rebelo, S. (forthcoming). A world equilibrium model of the oil market. *The Review of Economic Studies*.
- Bound, J., Jaeger, D. A. and Baker, R. M. (1995). Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak. *Journal of the American Statistical Association*, 90(430), 443–450.
- 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.
- Caldara, D., Cavallo, M. and Iacoviello, M. (2019). Oil price elasticities and oil price fluctuations. *Journal of Monetary Economics*, 103, 1–20.
- Canova, F. (2014). Bridging DSGE models and the raw data. *Journal of Monetary Economics*, 67, 1–15.
- Carter, C. A., Rausser, G. C. and Smith, A. (2017). Commodity storage and the market effects of biofuel policies. *American Journal of Agricultural Economics*, 99(4), 1027–1055.
- Chambers, M. J. and Bailey, R. E. (1996). A theory of commodity price fluctuations. *Journal of Political Economy*, 104(5), 924–957.
- Chinn, M. D. and Coibion, O. (2014). The predictive content of commodity futures. *Journal of Futures Markets*, 34(7), 607–636.
- Christiano, L., Eichenbaum, M. and Evans, C. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1), 1–45.

- Corradi, V. and Swanson, N. R. (2006). The effect of data transformation on common cycle, cointegration, and unit root tests: Monte Carlo results and a simple test. *Journal of Econometrics*, 132(1), 195–229, common Features.
- Cumby, R. E. and Huizinga, J. (1992). Testing the autocorrelation structure of disturbances in ordinary least squares and instrumental variables regressions. *Econometrica*, 60(1), 185.
- Deaton, A. and Laroque, G. (1992). On the behaviour of commodity prices. *The 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.
- DeJong, D. N., Nankervis, J. C., Savin, N. E. and Whiteman, C. H. (1992). Integration versus trend stationary in time series. *Econometrica*, 60(2), 423–433.
- 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, special issue Understanding International Commodity Price Fluctuations Edited by Rabah Arezki, Prakash Loungani, Rick van der Ploeg and Anthony J. Venables.
- FAO (2020). FAOSTAT, Food and Agriculture Organization of the United Nations statistical database. <http://www.fao.org/faostat/en/#home>, accessed June 24, 2020.
- Garcia, P., Irwin, S. H. and Smith, A. (2015). Futures market failure? *American Journal of Agricultural Economics*, 97(1), 40–64.
- Gardner, B. L. (1976). Futures prices in supply analysis. *American Journal of Agricultural Economics*, 58(1), 81–84.
- Ghanem, D. and Smith, A. (2022). Causality in structural vector autoregressions: Science or sorcery? *American Journal of Agricultural Economics*, 104(3), 881–904.
- Ghoshray, A. (2010). A reexamination of trends in primary commodity prices. *Journal of Development Economics*, 95(2), 242–251.
- Gouel, C. (2013a). Comparing numerical methods for solving the competitive storage model. *Computational Economics*, 41(2), 267–295.
- Gouel, C. (2013b). Optimal food price stabilisation policy. *European Economic Review*, 57, 118–134.
- Gouel, C. (2020). The value of public information in storable commodity markets: Application to the soybean market. *American Journal of Agricultural Economics*, 102(3), 846–865.
- Gouel, C. and Legrand, N. (2017). Estimating the competitive storage model with trending commodity prices. *Journal of Applied Econometrics*, 32(4), 744–763.
- Gourieroux, C., Monfort, A. and Renault, E. (1993). Indirect inference. *Journal of Applied Econometrics*, 8(S1), S85–S118.
- Guerra, E. A., Bobenrieth, E. S. A., Bobenrieth, 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). Implications of recent research on optimal storage rules. *Journal of Farm Economics*, 40(2), 290–300.
- Guvenen, F. and Smith, A. A. (2014). Inferring labor income risk and partial insurance from economic choices. *Econometrica*, 82(6), 2085–2129.
- Harrell, F. E. (2001). *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*. New York: Springer.
- Heiss, F. and Winschel, V. (2008). Likelihood approximation by numerical integration on sparse grids. *Journal of Econometrics*, 144(1), 62–80.
- Hendricks, N. P., Janzen, J. P. and Smith, A. (2015). Futures prices in supply analysis: Are instrumental variables necessary? *American Journal of Agricultural Economics*, 97(1), 22–39.
- IMF (2014). *World Economic Outlook: Recovery strengthens, Remains uneven*. Washington, DC: International Monetary Fund.
- Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *The American Economic Review*, 99(3), 1053–1069.
- Kilian, L. and Murphy, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29(3), 454–478.
- Knittel, C. R. and Pindyck, R. S. (2016). The simple economics of commodity price speculation. *American Economic Journal: Macroeconomics*, 8(2), 85–110.
- Lee, J., List, J. A. and Strazicich, M. C. (2006). Non-renewable resource prices: Deterministic or stochastic trends? *Journal of Environmental Economics and Management*, 51(3), 354–370.
- Lee, J. and Strazicich, M. C. (2003). Minimum Lagrange multiplier unit root test with two structural breaks. *The Review of Economics and Statistics*, 85(4), 1082–1089.
- Lee, J. and Strazicich, M. C. (2013). Minimum LM unit root test with one structural break. *Economics Bulletin*, 33(4), 2483–2492.
- Li, T. (2010). Indirect inference in structural econometric models. *Journal of Econometrics*, 157(1), 120–128.
- Lowry, M., Glauber, J. W., Miranda, M. J. and Helmberger, P. G. (1987). Pricing and storage of field crops: A quarterly model applied to soybeans. *American Journal of Agricultural Economics*, 69(4), 740–749.
- Maliar, L. and Maliar, S. (2014). Numerical methods for large-scale dynamic economic models. In K. Schmedders and K. L. Judd (eds.) *Handbook of Computational Economics*, volume 3, chapter 7, (pp. 325–477). Amsterdam: Elsevier.
- 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.
- Nordhaus, W. D. (1987). Forecasting efficiency: Concepts and applications. *The Review of Economics and Statistics*, 69(4), 667–674.
- Oglend, A. and Kleppe, T. S. (2017). On the behavior of commodity prices when speculative

- storage is bounded. *Journal of Economic Dynamics and Control*, 75, 52–69.
- Orcutt, G. H. and Winokur, H. S. (1969). First order autoregression: Inference, estimation, and prediction. *Econometrica*, 37(1), 1–14.
- Osborne, T. (2004). Market news in commodity price theory: Application to the Ethiopian grain market. *The Review of Economic Studies*, 71(1), 133–164.
- Perron, P. (1989). The great crash, the oil price shock, and the unit root hypothesis. *Econometrica*, 57(6), 1361–1401.
- Porteous, O. (2019). High trade costs and their consequences: An estimated dynamic model of African agricultural storage and trade. *American Economic Journal: Applied Economics*, 11(4), 327–366.
- Powell, M. J. D. (2009). *The BOBYQA algorithm for bound constrained optimization without derivatives*. Numerical analysis preprint DAMTP 2009/NA06, University of Cambridge.
- 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.
- Rotemberg, J. J. and Woodford, M. (1997). An optimization-based econometric framework for the evaluation of monetary policy. *NBER Macroeconomics Annual*, 12, 297–346.
- Routledge, B. R., Seppi, D. J. and Spatt, C. S. (2000). Equilibrium forward curves for commodities. *The Journal of Finance*, 55(3), 1297–1338.
- Ruge-Murcia, F. J. (2007). Methods to estimate dynamic stochastic general equilibrium models. *Journal of Economic Dynamics and Control*, 31(8), 2599–2636.
- Ruge-Murcia, F. J. (2020). Estimating nonlinear dynamic equilibrium models by matching impulse responses. *Economics Letters*, 197, 109624.
- Simonovska, I. and Waugh, M. E. (2014). The elasticity of trade: Estimates and evidence. *Journal of International Economics*, 92(1), 34–50.
- Smets, F. and Wouters, R. (2003). An estimated dynamic stochastic general equilibrium model of the euro area. *Journal of the European Economic Association*, 1(5), 1123–1175.
- Smets, F. and Wouters, R. (2007). Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *The American Economic Review*, 97(3), 586–606.
- Smith, A. A. (1993). Estimating nonlinear time-series models using simulated vector autoregressions. *Journal of Applied Econometrics*, 8(S1), S63–S84.
- Steinwender, C. (2018). Real effects of information frictions: When the States and the Kingdom became united. *The American Economic Review*, 108(3), 657–696.
- USDA (2020). United States Department of Agriculture, Production, Supply and Distribution database. <https://www.fas.usda.gov/data>, version of February, 2020.
- Williamson, L. and Williamson, P. (1942). What we eat. *Journal of Farm Economics*, 24(3), 698–703.

- 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.
- Wright, B. D. and Williams, J. C. (1984). The welfare effects of the introduction of storage. *The Quarterly Journal of Economics*, 99(1), 169–192.
- Zivot, E. and Andrews, D. W. K. (1992). Further evidence on the great crash, the oil-price shock, and the unit-root hypothesis. *Journal of Business & Economic Statistics*, 10(3), 251–270.

Appendix

A. Numerical methods

A. Algorithm

The proposed storage model includes three state variables, elastic supply, and isoelastic functions. These three features complicate its numerical resolution compared to most of the storage models in the literature. This model could be solved by a collocation method on a regular grid (see e.g., Gouel, 2013a); however, this would be too slow for being used in estimation methods involving simulations. We therefore develop a solution method that is specific to our model based on recent developments in the literature (Maliar and Maliar, 2014). Technically, it is based on linear interpolation on a sparse grid using Delaunay triangulation and a grid that is adapted to each set of parameters based on the ergodic distribution of the state variables. For each grid point, the equations are solved by derivative-free fixed-point iterations. The expectations operators are substituted by deterministic sums using sparse grid integration (Heiss and Winschel, 2008).

The interpolation grid is built using heuristics from the literature on numerical methods for large-scale dynamic models (Maliar and Maliar, 2014). However, it deviates from the existing methods to accommodate the specificity of the model in which only one state variable is endogenous. Two of the three state variables are exogenous shocks, so the grid points corresponding to these variables can be adjusted for each parameter change based on the new standard deviations. Only the grid points for the remaining state variable, net availability, are adjusted based on simulations from the ergodic distribution.

Taking account of these adaptations we can generate the grid in three steps. First, we construct a grid on the shocks $\{\varphi, \mu\}$ assuming $\sigma_\varphi = \sigma_\mu = 1$. This produces a Smolyak grid based on Heiss and Winschel's (2008) numerical integration programs. The grid can be scaled to different standard deviations. Note that we retain the integration weights for later use. Second, the model is simulated based on a previous solution (or guessed policy rules) which provides an availability series from which we calculate the mean \bar{s} and standard deviations σ_s . We generate a logarithmically-spaced availability vector between $\bar{s} - 4\sigma_s$ and $\bar{s} + 5\sigma_s$ which in our experience covers almost all simulated availabilities. A logarithmically-spaced vector will position more points in the low availability area where the cutoff of no stock is likely located than would a linearly-spaced vector. Assuming for availability a normal distribution with parameters \bar{s} and σ_s , we associate probability weights with each vector point based on the segments on which each vector point is centered. Third, we construct the full grid on the three state variables taking the tensor-product of the grid on shocks times the vector on availability. The same tensor-product is used to combine the probability weights. To trim the grid of low probability combinations, we use the weights and retain only the points with the highest probability weights.

The grid is a function of the policy rules so should be updated with policy rules until consistency.

However, since this is a costly step the grid is updated only once for each new set of parameters. Since the optimization algorithm used for the estimation involves smaller steps with convergence to the solution, this implies that close to the solution the grid converges to its configuration with full updating.

For conciseness, the following algorithm includes a few simplifications. Expectations operators are retained; in practice, they are replaced by simple weighted sums. We omit time subscripts: next period variables and shocks are indicated using the + exponent. We normalize the steady-state values to 1. The algorithm then runs as follows.

Step 1. Initialization step. Choose

- A convergence criterion $\varpi = 10^{-8}$ and a damping parameter $\lambda = 0.2$.
- A sparse grid on planting-time supply shocks and demand shocks $\{\varphi, \mu\}$, with associated probability weights.
- A sparse grid on shocks for numerical integration $\{\varphi^+, \epsilon^+, \mu^+\}$ with associated weights.
- Initial policy rules (guessed): $\mathcal{P}^1, \mathcal{X}^1, \mathcal{Q}^1$.

Step 2. New grid step. If $n = 1$, then update the interpolation grid.

Step 2.1. Use the policy rules and the transition equations to simulate the model over 50,000 periods (after excluding burn-in periods), keeping the same shocks each time the model is simulated to update the grid. Calculate the average availability \bar{s} and the standard deviation of availability σ_s .

Step 2.2. Generate a 12×1 logarithmically-spaced vector of availability between $\bar{s} - 4\sigma_s$ and $\bar{s} + 5\sigma_s$ and associate with each points a probability assuming a normal distribution with parameters \bar{s} and σ_s .

Step 2.3. Use a tensor product of the grid on shocks and the vector on availability to obtain a full grid and keep the 140 grid points with the highest probability weights. Divide availability by demand shock to obtain the grid points on net availability.

Step 2.4. Use the policy rules, \mathcal{P}^n and \mathcal{Q}^n , to adjust the response variables to the new grid $\{\tilde{s}, \varphi, \mu\}$:

$$p^{n-1} = \mathcal{P}^n(\tilde{s}, \varphi, \mu), \quad (\text{A1})$$

$$q^{e,m,n-1} = \mathcal{Q}^n(\tilde{s}, \varphi, \mu), \quad (\text{A2})$$

$$x^{m,n-1} = \max(0, (\tilde{s} - d(p^{n-1})) e^\mu). \quad (\text{A3})$$

Step 3. Solve for production and storage. Define $q^{e,1,n} = q^{e,m,n-1}$ and $x^{1,n} = x^{m,n-1}$. For each gridpoint, iterate on m according to the following steps:

Step 3.1. Calculate next-period price for a combination of interpolation grid points and integration nodes:

$$p^+ = \mathcal{P}^n \left(\left[(1 - \delta) x^{m-1,n} e^{-g_q} + q^{e,m-1,n} e^{\epsilon^+} \right] e^{-\mu^+}, \varphi^+, \mu^+ \right). \quad (\text{A4})$$

Step 3.2. Fixed-point iteration with damping:

$$q^{e,m,n} = (1 - \lambda) q^{e,m-1,n} + \lambda e^\varphi \left[E \left(p^+ e^{\epsilon^+} \right) \right]^{\alpha_S}, \quad (\text{A5})$$

$$x^{m,n} = (1 - \lambda) x^{m-1,n} + \lambda \max \left(0, \tilde{s} e^\mu - d \left(\beta (1 - \delta) e^{g_p} E \left(p^+ \right) - k \right) e^\mu \right). \quad (\text{A6})$$

If $\max (\|q^{e,m,n} - q^{e,m-1,n}\|_2, \|x^{m,n} - x^{m-1,n}\|_2) < \lambda \varpi$ or $m = m^{\max}$ then stop iterations and go to next step.

Step 4. Approximation step. Calculate prices as

$$p^n = d^{-1} \left(\tilde{s} - x^{m,n} e^{-\mu} \right), \quad (\text{A7})$$

from which we update the price function

$$\mathcal{P}^{n+1} (\tilde{s}, \varphi, \mu) = p^n. \quad (\text{A8})$$

We also update the production function

$$\mathcal{Q}^{n+1} (\tilde{s}, \varphi, \mu) = q^{e,m,n}. \quad (\text{A9})$$

Step 5. Terminal step.

If $n = 1$ or $\|p^n - p^{n-1}\|_2 \geq \varpi$ or $\max (\|q^{e,m,n} - q^{e,m-1,n}\|_2, \|x^{m,n} - x^{m-1,n}\|_2) \geq \lambda \varpi$ then increment n to $n + 1$ and go to step 3.

At the end of the algorithm, we use the most recent calculated values of $x^{m,n}$ and $E(p^+)$ to determine the storage rule, \mathcal{X} , and an approximation of the expected prices which are useful to simulate the model.

There are a few things to note about this algorithm. First, the stop criterion of the inner fixed point on production and storage implies that this fixed point may stop before convergence is achieved. We choose $m^{\max} = 5$ so that it occurs frequently. This is a useful procedure since production and storage levels do not need to be perfectly consistent with the price rule before the overall algorithm converges. It is better to stop after a few iterations when a reasonable guess can be made rather than solving for a perfect intermediary solution requiring many iterations. In addition, for poor price rules there may be no solution to this fixed point. However, to ensure that production and storage levels eventually converge to a level consistent with the price rule when the algorithm stops, this convergence is tested in step 5.

Second, due to the damping parameter the convergence criterion for step 3 needs to be stricter than the convergence criterion for the norm of $p^n - p^{n-1}$ in the final step. With the same convergence criterion, production and storage levels would not be sufficiently updated in the last steps of the algorithm and it would cycle infinitely between the inner and outer loops.

Third, the interpolation is made not on prices but on the logarithm of prices. This increases the precision in stockout situations where the price then becomes an isoelastic function of net availability. Therefore, a linear interpolation in logarithm will be exact in stockouts, while a linear interpolation in level would not. This detail is not included in the above algorithm.

B. Solution precision

Once a solution is obtained, its accuracy can be assessed by rewriting unit-free the equations (13) and (14) which give two measures of the Euler equation errors. Using a net availability and shocks series, $\{\tilde{s}_i, \varphi_i, \mu_i\}$, the storage and production equation errors can be assessed using (see Gouel, 2013a, for details of the derivation of these measures for the storage model)

$$EE_i^x = 1 - \frac{d (\max(d^{-1}(\tilde{s}_i), \beta(1-\delta)e^{g_p} E \mathcal{P}([(1-\delta)\mathcal{X}(\tilde{s}_i, \varphi_i, \mu_i)e^{-g_q} + \mathcal{Q}(\tilde{s}_i, \varphi_i, \mu_i)e^\epsilon]e^{-\mu}, \varphi, \mu) - k)) e^{\mu_i}}{\tilde{s}_i e^{\mu_i} - \mathcal{X}(\tilde{s}_i, \varphi_i, \mu_i)}, \quad (\text{A10})$$

$$EE_i^h = 1 - \frac{e^{\varphi_i} \{E[\mathcal{P}([(1-\delta)\mathcal{X}(\tilde{s}_i, \varphi_i, \mu_i)e^{-g_q} + \mathcal{Q}(\tilde{s}_i, \varphi_i, \mu_i)e^\epsilon]e^{-\mu}, \varphi, \mu)e^\epsilon]\}^{\alpha_s}}{\mathcal{Q}(\tilde{s}_i, \varphi_i, \mu_i)}. \quad (\text{A11})$$

To assess the precision of the algorithm, we simulate the model calibrated on our preferred estimation (4-knot spline in Table 8). We then sample 1,000 points from the ergodic distribution and use them to calculate the Euler equation errors defined above. Table A1 presents the average and the maximum errors expressed in base-10 logarithm. The accuracy in both equations is similar. At about -2 , maximum errors involve a \$1 error every \$100 consumption or production decisions. However, such high error rates are rare and are located close to cutoff situations of no storage. The average errors involve less than \$1 error every \$1,000 decisions, and are closer to \$1 error for every \$10,000 decisions. This is a satisfactory level of precision for this type of model, and as the Monte Carlo experiments show is sufficiently high for our estimation procedure to recover the true parameter values if the model is well specified.

Table A1 – Euler equations error ($\log_{10} |EE|$)

Equation	Average error	Max error
EE^x	-3.65	-1.77
EE^h	-3.73	-2.01

Notes: Calculated over 1,000 simulations from the model's ergodic distribution. The model parameters are from our preferred estimation (4-knot spline in Table 8).

B. Supplementary tables

Table A2 – Parameter bounds when minimizing the indirect inference objective

Parameter	Lower bound	Upper bound
ρ_μ	0	1
$\rho_{\eta,\omega}$	-1	1
σ_ω	0	1
σ_η	0	0.1
σ_ϵ	0	0.1
σ_ν	0	0.1
δ	0	1
k	0	$+\infty$
α_D	$-\infty$	0
α_S	0	$+\infty$

Table A3 – Additional Monte Carlo experiments with instrumental variables for $\sigma_\omega = 20\%$

	ρ_μ	$c_q - 1$	σ_ψ (%)	σ_ϑ (%)	σ_ν (%)	α_D	α_S
$T = 100$, {Supply: $E(F) = 19$, $E(p\text{-value}) = 0.07$ }, {Demand: $E(F) = 38$, $E(p\text{-value}) = 0.00$ }							
Mean	0.50	0.194	2.50	3.42	1.62	-0.071	0.084
St. dev.	0.12	0.093	0.17	0.37	0.21	0.015	0.036
RMSE (%)	23.84	51.322	6.90	11.30	13.48	21.325	45.259
SE	0.12	0.094	0.18		0.20	0.015	0.036
$T = 200$, {Supply: $E(F) = 39$, $E(p\text{-value}) = 0.01$ }, {Demand: $E(F) = 76$, $E(p\text{-value}) = 0.00$ }							
Mean	0.50	0.189	2.50	3.38	1.61	-0.070	0.081
St. dev.	0.07	0.062	0.12	0.24	0.14	0.010	0.023
RMSE (%)	14.24	34.025	4.86	7.08	8.95	13.716	29.057
SE	0.08	0.063	0.13		0.14	0.010	0.023
$T = 1000$, {Supply: $E(F) = 193$, $E(p\text{-value}) = 0.00$ }, {Demand: $E(F) = 384$, $E(p\text{-value}) = 0.00$ }							
Mean	0.50	0.181	2.50	3.35	1.60	-0.070	0.078
St. dev.	0.03	0.027	0.06	0.11	0.06	0.004	0.010
RMSE (%)	6.57	14.971	2.24	3.14	3.97	6.337	12.494
SE	0.03	0.027	0.06		0.06	0.004	0.010

Notes: See notes to table 1.

Table A4 – Additional Monte Carlo experiments with indirect inference and $\sigma_\omega = 20\%$ (auxiliary model based on OLS regressions)

	ρ_μ	$\rho_{\eta,\omega}$	σ_ω (%)	σ_η (%)	σ_ϵ (%)	σ_ν (%)	δ (%)	k (%)	α_D	α_S
$T = 100$	OID: 0.028									
Mean	0.50	-0.40	20.70	1.48	1.98	1.60	2.03	2.92	-0.071	0.081
St. dev.	0.08	0.19	3.61	0.28	0.23	0.17	1.36	1.81	0.011	0.018
RMSE (%)	16.43	47.90	18.37	18.55	11.62	10.56	68.13	60.46	15.951	22.608
ASE	0.07	0.21	3.52	0.29	0.25	0.19	13.44	12.48	0.014	0.018
$T = 200$	OID: 0.022									
Mean	0.50	-0.40	20.31	1.50	1.99	1.60	1.95	3.00	-0.070	0.080
St. dev.	0.06	0.14	2.50	0.18	0.16	0.13	1.26	1.45	0.008	0.013
RMSE (%)	11.07	35.19	12.62	12.03	7.82	7.92	63.25	48.32	11.047	16.566
ASE	0.05	0.14	2.31	0.20	0.17	0.13	10.21	9.30	0.010	0.012
$T = 1000$	OID: 0.012									
Mean	0.50	-0.40	20.05	1.50	1.99	1.60	1.86	3.03	-0.070	0.080
St. dev.	0.03	0.06	1.13	0.08	0.07	0.05	1.03	0.99	0.003	0.006
RMSE (%)	5.06	16.24	5.63	5.50	3.48	3.36	51.89	32.97	4.911	7.629
ASE	0.02	0.06	0.98	0.09	0.08	0.06	4.85	4.45	0.004	0.005

Notes: See notes to table 2.

Table A5 – Monte Carlo experiment with indirect inference approach (auxiliary model based on 2SLS regressions)

	ρ_μ	$\rho_{\eta,\omega}$	σ_ω (%)	σ_η (%)	σ_ϵ (%)	σ_ν (%)	δ (%)	k (%)	α_D	α_S
$T = 56$	$\sigma_\omega = 5\%$	OID: 0.021								
Mean	0.48	-0.43	4.89	1.48	1.98	1.59	1.94	2.92	-0.068	0.078
St. dev.	0.12	0.29	0.94	0.31	0.27	0.30	1.18	2.08	0.016	0.014
RMSE (%)	24.59	71.96	19.00	20.81	13.63	18.52	59.02	69.47	22.652	17.802
ASE	0.11	0.66	2.97	0.38	0.32	0.23	20.47	19.61	0.016	0.022
$T = 56$	$\sigma_\omega = 10\%$	OID: 0.019								
Mean	0.48	-0.41	10.31	1.48	1.97	1.58	1.95	2.91	-0.068	0.077
St. dev.	0.13	0.27	2.53	0.33	0.28	0.28	1.19	2.09	0.015	0.020
RMSE (%)	25.41	66.66	25.50	22.02	14.26	17.57	59.74	69.69	21.210	25.697
ASE	0.11	0.54	5.48	0.41	0.34	0.24	21.83	20.85	0.017	0.030
$T = 56$	$\sigma_\omega = 20\%$	OID: 0.015								
Mean	0.47	-0.39	21.98	1.48	1.96	1.58	2.04	2.83	-0.068	0.074
St. dev.	0.14	0.26	8.00	0.34	0.29	0.28	1.16	1.96	0.015	0.029
RMSE (%)	27.57	65.90	41.21	23.01	14.42	17.70	57.83	65.56	21.630	36.339
ASE	0.12	0.48	13.83	0.46	0.37	0.27	23.37	22.19	0.018	0.042
$T = 100$	$\sigma_\omega = 20\%$	OID: 0.004								
Mean	0.49	-0.38	21.98	1.49	1.96	1.59	1.91	2.98	-0.070	0.077
St. dev.	0.10	0.18	6.82	0.26	0.23	0.23	1.23	1.66	0.013	0.025
RMSE (%)	20.76	46.59	35.48	17.60	11.56	14.50	61.85	55.26	18.113	31.064
ASE	0.08	0.29	8.69	0.33	0.28	0.18	18.68	16.96	0.012	0.025
$T = 200$	$\sigma_\omega = 20\%$	OID: 0.004								
Mean	0.50	-0.39	20.73	1.50	1.99	1.61	1.96	2.96	-0.070	0.079
St. dev.	0.07	0.13	3.72	0.17	0.15	0.16	1.20	1.41	0.009	0.017
RMSE (%)	13.85	32.47	18.97	11.31	7.78	10.23	60.19	47.17	12.376	21.814
ASE	0.06	0.19	5.03	0.23	0.19	0.12	13.80	12.65	0.008	0.017
$T = 1000$	$\sigma_\omega = 20\%$	OID: 0.004								
Mean	0.50	-0.40	20.17	1.50	1.99	1.60	1.90	2.98	-0.070	0.080
St. dev.	0.03	0.06	1.44	0.08	0.07	0.07	0.98	0.98	0.004	0.008
RMSE (%)	6.51	14.61	7.23	5.17	3.41	4.18	49.33	32.64	5.251	9.707
ASE	0.02	0.08	2.09	0.10	0.09	0.05	6.93	6.24	0.003	0.007

Notes: See notes to table 2. For $T = 100$, 1 replication had to be dropped due to non-convergence.

Table A7 – Supply equation estimation, 1962–2007

	(1)	(2)	(3)
<i>Panel A. 2SLS</i>			
Supply elasticity b_q	0.102*** (0.032)	0.093** (0.035)	0.089*** (0.033)
Shock c_q	1.130*** (0.194)	1.177*** (0.204)	1.157*** (0.185)
<i>Panel B. First stage</i>			
Lagged shock b_{EP}	−3.908*** (1.232)	−3.579*** (1.062)	−3.619*** (1.089)
Shock c_{EP}	−2.908 (1.822)	−2.311 (1.483)	−2.382 (1.568)
<i>Panel C. OLS</i>			
Supply elasticity b_q	0.111*** (0.017)	0.086*** (0.018)	0.085*** (0.017)
Shock c_q	1.162*** (0.135)	1.157*** (0.141)	1.144*** (0.133)
$\sigma_{u_q^{2SLS}}$	0.018	0.016	0.015
$\sigma_{\vartheta^{2SLS}}$	0.032	0.032	0.031
$\sigma_{u_{EP}}$	0.159	0.140	0.143
$\sigma_{u_q^{OLS}}$	0.018	0.016	0.015
$\sigma_{\vartheta^{OLS}}$	0.033	0.032	0.031
First stage F -stat	10.065	11.367	11.037
p -value for Hausman test	0.805	0.846	0.880
p -value for Cumby-Huizinga test (panel A)	0.007	0.012	0.041
Observations	46	46	46
Spline knots	3	4	5

Notes: Standard errors robust to heteroskedasticity in parenthesis. ***, **, and * indicate significance at the 99%, 95%, and 90% levels, respectively. The knots are placed following Roberts and Schlenker (2013): 1963, 1984, and 2005 for 3 knots; 1962, 1976, 1992, and 2006 for 4 knots; and 1962, 1973, 1984, 1995, and 2006 for 5 knots.

Table A8 – Demand equation estimation, 1962–2007

	(1)	(2)	(3)
<i>Panel A. 2SLS</i>			
Demand elasticity b_c	-0.034 (0.033)	-0.045 (0.032)	-0.047 (0.032)
Lagged price c_c	0.022 (0.023)	0.012 (0.020)	0.011 (0.021)
Lagged demand d_c	0.977*** (0.228)	0.768*** (0.217)	0.633** (0.242)
<i>Panel B. First stage</i>			
Shock b_p	-3.926*** (0.939)	-3.743*** (1.056)	-3.819*** (1.041)
Lagged price c_p	0.599*** (0.146)	0.532*** (0.149)	0.547*** (0.152)
Lagged demand d_p	4.960** (2.034)	3.644* (2.007)	4.403* (2.446)
<i>Panel C. OLS</i>			
Demand elasticity b_c	0.000 (0.012)	-0.007 (0.011)	-0.005 (0.010)
Lagged price c_c	-0.003 (0.016)	-0.012 (0.015)	-0.016 (0.015)
Lagged demand d_c	0.743*** (0.131)	0.593*** (0.152)	0.433** (0.174)
<i>Panel D. 2SLS using Roberts and Schlenker's approach (eqs. (38) for 2nd stage and (42) for 1st)</i>			
Demand elasticity b_c	-0.033 (0.023)	-0.062* (0.031)	-0.059** (0.028)
$\sigma_{u_c}^{2SLS}$	0.015	0.015	0.015
σ_{u_p}	0.175	0.175	0.175
$\sigma_{u_c}^{OLS}$	0.014	0.013	0.012
$\sigma_{u_c}^{2SLS, RS}$	0.022	0.020	0.018
σ_{μ}^{2SLS}	0.072	0.023	0.019
First stage F -stat (panel A)	17.479	12.573	13.461
p -value for Hausman test (panel A)	0.274	0.221	0.117
p -value for Cumby-Huizinga test (panel A)	0.484	0.335	0.113
First stage F -stat (panel D)	20.780	17.059	17.156
p -value for Hausman test (panel D)	0.008	0.052	0.033
p -value for Cumby-Huizinga test (panel D)	0.000	0.035	0.067
Observations	46	46	46
Spline knots	3	4	5

Notes: Standard errors robust to heteroskedasticity in parenthesis, except for panel D where they are also robust to autocorrelation. The lagged demand estimates in panel A are bias adjusted (Orcutt and Winokur, 1969). **, *, and * indicate significance at the 99%, 95%, and 90% levels, respectively. The knots are placed following Roberts and Schlenker (2013): 1963, 1984, and 2005 for 3 knots; 1962, 1976, 1992, and 2006 for 4 knots; and 1962, 1973, 1984, 1995, and 2006 for 5 knots.

C. Sensitivity analyses

Table A9 – Estimation results for the indirect inference approach (auxiliary model based on 2SLS regressions)

	4-knot spline		5-knot spline	
	Estimate	Standard error	Estimate	Standard error
ρ_μ	0.477	(0.117)	0.443	(0.145)
$\rho_{\eta,\omega}$	-0.484	(0.420)	-0.434	(0.425)
σ_ω	0.285	(0.115)	0.278	(0.113)
σ_η	0.014	(0.005)	0.014	(0.005)
σ_ϵ	0.019	(0.004)	0.019	(0.004)
σ_ν	0.017	(0.003)	0.016	(0.003)
δ	0		0	
k	0.017	(0.013)	0.015	(0.013)
α_D	-0.044	(0.015)	-0.037	(0.015)
α_S	0.054	(0.020)	0.056	(0.020)
σ_φ	0.026	(0.005)	0.025	(0.005)
σ_ψ	0.024	(0.002)	0.024	(0.002)
σ_μ	0.019	(0.003)	0.018	(0.003)
σ_ϑ	0.032	(0.004)	0.032	(0.004)
OID p -value	0.658		0.774	

Notes: See notes to table 8.

Table A10 – Additional estimation results (auxiliary model based on OLS regressions and detrending with 4 knots)

	FAOSTAT data		FAOSTAT data without rice				USDA-PSD data			
	II – Only shrinkage ($k = 0$)		2SLS		II		2SLS		II	
	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE
g_p	-0.020				-0.020				-0.020	
g_q	0.025				0.027				0.025	
ρ_μ	0.714	(0.067)	0.530	(0.158)	0.674	(0.074)	0.533	(0.226)	0.738	(0.059)
$\rho_{\eta,\omega}$	-0.450	(0.315)			-0.396	(0.287)			-0.437	(0.341)
σ_ω	0.189	(0.032)			0.219	(0.042)			0.177	(0.026)
σ_η	0.015	(0.006)			0.021	(0.006)			0.013	(0.007)
σ_ϵ	0.020	(0.005)			0.024	(0.006)			0.020	(0.005)
σ_v	0.018	(0.003)	0.021	(0.005)	0.024	(0.004)	0.018	(0.004)	0.022	(0.004)
δ	0.038	(0.013)			0				0	
k	0				0.032	(0.012)			0.038	(0.015)
α_D	-0.064	(0.018)	-0.083	(0.035)	-0.087	(0.026)	-0.076	(0.033)	-0.089	(0.025)
α_S	0.086	(0.016)	0.088	(0.035)	0.096	(0.020)	0.086	(0.026)	0.079	(0.012)
σ_φ	0.027	(0.005)			0.037	(0.006)			0.024	(0.005)
σ_ψ	0.025	(0.002)	0.030	(0.003)	0.032	(0.003)	0.023	(0.002)	0.024	(0.002)
σ_μ	0.026	(0.005)	0.025	(0.006)	0.033	(0.007)	0.022	(0.006)	0.033	(0.006)
σ_θ	0.034	(0.004)	0.040		0.044	(0.005)	0.031		0.031	(0.003)
OID	0.060				0.068				0.169	

Notes: See notes to tables 6–8.

Table A11 – Estimation results by commodity (auxiliary model based on OLS regressions and detrending with 4 knots)

	Maize				Soybeans				Wheat			
	2SLS		II		2SLS		II		2SLS		II	
	Estimate	SE										
g_p			-0.020				-0.018				-0.021	
g_q			0.028				0.044				0.019	
ρ_μ	0.501	(0.192)	0.735	(0.062)	0.482	(0.195)	0.565	(0.102)	0.605	(0.197)	0.628	(0.085)
$\rho_{\eta,\omega}$			-0.960	(1.641)			0.201	(0.233)			-0.133	(0.245)
σ_ω			0.170	(0.028)			0.362	(0.098)			0.473	(0.160)
σ_η			0.013	(0.018)			0.025	(0.015)			0.035	(0.007)
σ_ϵ			0.039	(0.007)			0.040	(0.010)			0.024	(0.009)
σ_v	0.028	(0.005)	0.034	(0.005)	0.048	(0.013)	0.059	(0.017)	0.029	(0.010)	0.037	(0.010)
δ			0				0				0	
k			0.048	(0.016)			0.018	(0.018)			0.060	(0.028)
α_D	-0.110	(0.031)	-0.131	(0.033)	-0.090	(0.111)	-0.168	(0.118)	-0.096	(0.074)	-0.126	(0.048)
α_S	0.162	(0.057)	0.165	(0.032)	0.226	(0.181)	0.170	(0.054)	0.060	(0.052)	0.064	(0.024)
σ_φ			0.043	(0.010)			0.063	(0.011)			0.051	(0.008)
σ_ψ	0.041	(0.004)	0.042	(0.004)	0.047	(0.004)	0.047	(0.004)	0.040	(0.004)	0.042	(0.004)
σ_μ	0.033	(0.007)	0.051	(0.010)	0.055	(0.016)	0.071	(0.025)	0.037	(0.014)	0.047	(0.015)
σ_θ	0.057		0.058	(0.008)	0.073		0.074	(0.008)	0.047		0.056	(0.007)
OID			0.486				0.119				0.102	

Notes: See notes to tables 6–8.

Table A12 – Sensitivity of estimates to estimation moments (indirect inference approach with auxiliary model based on OLS regressions and detrending with 4 knots)

Coefficient	ρ_μ	$\rho_{\eta,\omega}$	σ_ω	σ_η	σ_ϵ	σ_ν	k	α_D	α_S
b_q	0.180	-0.066	-0.606	-0.037	0.034	0.071	0.198	0.012	0.545
c_q	0.044	6.265	-0.060	-0.116	0.042	-0.124	-0.690	-0.296	0.502
b_c	0.030	0.015	0.026	0.089	-0.012	0.002	0.326	0.175	-0.015
c_c	-0.017	0.026	0.009	-0.004	0.007	-0.013	-0.064	-0.048	-0.009
d_c	0.161	0.238	-0.034	0.095	-0.010	-0.134	-0.134	-0.070	0.081
b_p	-0.048	1.103	-0.134	-0.981	0.457	-0.159	0.430	-0.394	0.168
c_p	0.137	-0.107	0.218	-0.157	0.040	0.140	-0.306	0.371	-0.297
d_p	0.001	-0.004	-0.000	0.004	-0.001	-0.003	0.001	-0.004	-0.001
$b_{E,p}$	-0.009	-0.213	0.025	0.440	-0.249	0.014	0.157	0.062	-0.006
$c_{E,p}$	-0.041	0.754	0.012	-0.409	0.272	-0.201	-0.245	-0.419	0.021
σ_{u_q}	0.088	-0.554	0.408	-0.153	0.073	-0.093	-0.132	-0.226	0.602
σ_{u_c}	-0.216	0.885	0.031	-0.432	0.277	1.260	0.912	1.468	-0.035
σ_{u_p}	-0.236	-1.362	0.032	1.157	-0.581	0.171	2.205	-0.307	0.043
$\sigma_{u_{E,p}}$	0.490	0.052	0.599	-0.286	0.166	0.283	-0.054	0.273	-0.748
σ_{u_ψ}	-0.141	1.036	-0.086	0.604	1.057	-0.548	-0.852	-1.056	0.120

Notes: Measure of sensitivity of Andrews et al. (2017), normalized as elasticities of estimated parameters with respect to moments.

Table A13 – Sensitivity of estimates to estimation moments (indirect inference approach with auxiliary model based on OLS regressions and detrending with 4 knots, without per-unit storage costs)

Coefficient	ρ_μ	$\rho_{\eta,\omega}$	σ_ω	σ_η	σ_ϵ	σ_ν	δ	α_D	α_S
b_q	0.182	-0.057	-0.605	-0.039	0.034	0.072	0.165	0.027	0.542
c_q	0.052	6.242	-0.057	-0.187	0.089	-0.106	-0.647	-0.284	0.505
b_c	0.032	0.012	0.026	0.091	-0.016	-0.004	0.286	0.178	-0.015
c_c	-0.017	0.023	0.010	0.001	0.005	-0.012	-0.054	-0.048	-0.009
d_c	0.145	0.270	-0.037	0.081	0.001	-0.136	-0.121	-0.076	0.089
b_p	-0.047	1.128	-0.138	-0.979	0.463	-0.144	0.541	-0.389	0.177
c_p	0.126	-0.046	0.228	-0.178	0.061	0.129	-0.416	0.384	-0.308
d_p	0.001	-0.008	0.000	0.005	-0.003	-0.002	0.004	-0.002	-0.001
$b_{E,p}$	-0.010	-0.266	0.029	0.484	-0.279	0.004	0.118	0.045	-0.010
$c_{E,p}$	-0.039	0.747	0.011	-0.414	0.273	-0.168	-0.136	-0.378	0.020
σ_{u_q}	0.089	-0.557	0.402	-0.148	0.067	-0.076	-0.099	-0.205	0.607
σ_{u_c}	-0.214	0.774	0.031	-0.347	0.226	1.212	0.803	1.413	-0.031
σ_{u_p}	-0.195	-1.262	0.028	1.075	-0.524	0.064	2.015	-0.512	0.050
$\sigma_{u_{E,p}}$	0.473	0.040	0.604	-0.283	0.156	0.290	0.024	0.326	-0.754
σ_{u_ψ}	-0.141	1.120	-0.089	0.552	1.088	-0.463	-0.518	-0.966	0.125

Notes: Measure of sensitivity of Andrews et al. (2017), normalized as elasticities of estimated parameters with respect to moments.