



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

NCCC-134

APPLIED COMMODITY PRICE ANALYSIS, FORECASTING AND MARKET RISK MANAGEMENT

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

by

Christophe Gouel and Nicolas Legrand

Suggested citation format:

Gouel, C. and N. Legrand. 2021. "The Role of Storage in Commodity Markets: Indirect Inference Based on Grains Data." Proceedings of the NCCC-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management. [<http://www.farmdoc.illinois.edu/nccc134>].

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

Christophe Gouel and Nicolas Legrand*

Paper prepared for the NCCC-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management, 2021.

Copyright 2021 by Christophe Gouel and Nicolas Legrand. All rights reserved. Readers may make verbatim copies of this document for non-commercial purposes by any means, provided that this copyright notice appears on all such copies.

*Christophe Gouel (christophe.gouel@inrae.fr) is a Senior Research Fellow at INRAE, Économie Publique. Nicolas Legrand (nicolas.legrand@inrae.fr) is a Research Fellow at INRAE, SMART. This work was supported by the INRAE ECOSOCIO department.

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

This paper develops and estimates a rational expectations storage model with rich dynamic features. The model incorporates elastic supply, long-run demand and cost trends, and four structural shocks. It is estimated by indirect inference using a linear supply and demand model as an auxiliary model. This approach deviates from the common practice of estimating storage models only on prices and allows all parameters to be estimated. The estimation is carried out on a market representing the caloric aggregate of four basic staples: maize, rice, soybean, 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 which is the subject of longstanding debate. However, it fails to replicate the strength of the negative correlation between consumption and price.

Keywords: commodity price dynamics, indirect inference, Monte Carlo analysis, storage.

1 Introduction

In addition to being one of the first rational expectations models (Gustafson, 1958), the speculative storage model is the workhorse model used to analyze issues related to commodity prices dynamics. It is used for empirical work on commodity prices (Steinwender, 2018) and for policy analysis (Gouel, 2013b; Porteous, 2019). Having a sound model of storable commodity prices matters in many contexts: from governments in countries heavily reliant on exporting a narrow range of commodities to those who want to assess the likely impact of policies that affect commodity markets (e.g., biofuels). However, despite the importance of such a model in the economics and finance literature on commodity prices, empirical validations of it remain elusive. The first estimations by Deaton and Laroque (1992, 1996) cast doubt on its relevance because of its inability to generate the observed high levels of autocorrelation in prices. Subsequent work challenged these initial results and offered solutions to the autocorrelation puzzle (Cafiero et al., 2011, 2015; Gouel and Legrand, 2017). However, these studies build on Deaton and Laroque’s approach where the model is estimated only on prices, which requires restrictive identifying assumptions and does not allow estimation of all the model parameters. Using only prices as observables reduces the storage model to a pricing model, not considering that it is also a supply and demand model extended to a dynamic setting. Thus, the existing estimations say nothing about the ability of the model to match moments other than those related to prices.

This paper attempts to analyze the empirical validity of the storage model beyond its ability to fit price dynamics. This requires the inclusion of other observables than just prices, and developing a new estimation method to account for the additional information. Our chosen method is an indirect inference approach which estimates the storage model by matching the estimated parameters of an auxiliary model inspired from Roberts and Schlenker (2013). We apply this approach to the set of observed variables in Roberts and Schlenker (2013)—i.e., price, expected price, demand, production, and yield shock—corresponding to a caloric aggregate of maize, rice, soybeans, and wheat, the leading four sources of energy in human diets worldwide. A corollary of this estimation strategy is the need for a richer storage model than the one estimated so far. Therefore, we develop a new storage model with elastic supply, long-run demand and cost trends, and four shocks with heterogeneous timing.

The new model and estimation strategy allow a more credible solution to the price autocorrelation puzzle than proposed in previous work, since the model features explaining the high price persistence have to be consistent also with the non-price moments. After decomposing the serial correlation into the respective contributions of the various model's components, our benchmark estimations show that approximately 33% of the observed 0.87 one-year autocorrelation can be explained by a trend in prices, 10% by autocorrelated demand shocks, 6% by the presence of planting-time shocks, and 51% by speculative storage. Thus, storage while being crucial for inducing price persistence needs to be paired with other model features which can be estimated only on datasets that include information contained in the joint price and quantity dynamics. The literature so far emphasizes this moment, but after accounting for other observables there are many other moments to match. We show that the model is able to fit most second-order moments (with covariance lagged up to one year). However, this paper raises a new puzzle: the model proves unable to match the correlation between price and consumption, which is much lower in the data than in the model.

Our storage model differs from the large existing literature triggered by the seminal work by [Deaton and Laroque \(1992\)](#). [Deaton and Laroque](#) estimate the storage model only on price, which imposes a simple theoretical structure: an inelastic supply affected by a single shock at harvest time. If one considers additional observables, this simple theoretical structure leads to a trivial estimation.¹ For example, because demand is not affected by any disturbance, a regression of demand on price would directly recover the demand elasticity. In contrast, in our approach the inference of the storage model results in the estimation of a dynamic supply and demand equilibrium model, where correct identification requires accounting for unobservable shifts in each curve. A recent innovation in this literature is [Roberts and Schlenker's \(2013\)](#) application of storage theory to identify an appropriate instrument to estimate supply elasticities in storable commodity markets. Therefore, in our model we introduce a range of shocks, guided by both the theoretical structure implicit behind [Roberts and Schlenker's](#) estimation approach and by the moments in the data. Our model includes four shocks: an autocorrelated demand shock, two yield shocks, the first occurring during planting while the other during harvesting, and a cost shock. The model also includes a supply response along with trends in the demand and marginal cost functions. Although this model is a radical departure from the simple versions of the storage model estimated previously, it builds on studies that introduce similar features separately. For example, [Wright and Williams's](#) competitive storage model (1982) includes a responsive 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. One key difference between this paper and the existing literature is our ability to assess empirically the overall consistency of the model's combined features while identifying which ones help to match the moments in the data.

Despite the richness of our model compared to most models in the storage literature, it remains small in size, and particularly in terms of the number of observables. With only four shocks for five observables, the model is misspecified. Thus, in a likelihood-based approach, it would present a stochastic singularity, and by construction, could not be expected to account for the richness of

¹If stochastic singularity issues are neglected.

the data. To deal with this misspecification, we adopt an estimation approach that does not rely on the calculation of a likelihood, 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 \(1993\)](#), and is a simulated moment-based method in which the model is estimated by targeting parameter estimates from an auxiliary model. In simple terms, 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 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. In our setting, the auxiliary model is a supply and demand model, relying on linear equations, inspired by the instrumental variable model in [Roberts and Schlenker \(2013\)](#). This allows us to exploit an econometric literature where intuitions about which moment is driving a parameter estimation are more explicit than full-information techniques. Some model parameters such as the elasticities, and the autocorrelation and standard deviation of the demand shock can be estimated using two-stage least squares (2SLS) following [Roberts and Schlenker](#), an approach we employ also to provide a benchmark for comparing our indirect inference approach. Although a 2SLS estimation leaves unidentified many of the storage model parameters, they are present in non-linear combinations of the first- and second-stage parameters and can be recovered by indirect inference.

This indirect inference approach allows us to estimate all the model's parameters.² The elasticities obtained using indirect inference, i.e., once accounting for other restrictions in the data, are close to those obtained by 2SLS, which supports [Roberts and Schlenker's](#) low supply and demand elasticity estimates for this bundle of staple grains. In the case of the yield shock, the planting-time shock explains only 17% of its standard deviation of 2.6%, with the remaining 83% due to disturbance at harvest time. Combining planting-time yield with cost shocks leads to a total planting-time shock similar in size to the production shocks occurring later in the season. A large size for planting-time shocks matters because earlier shocks leave more time for production and stock levels to be adjusted to the future market scarcity than if the uncertainty is resolved at a much later moment.

This 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. Still, 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. However, this method is commonly applied in the estimation of dynamic stochastic general equilibrium (DSGE) models which conceptually and numerically are close to the storage model. In this literature, the auxiliary model is often a vector autoregression (VAR) model 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, a system of linear equations is enough to capture all the moments of interest (as in [Güvenen and Smith, 2014](#)), but a VAR model would likely work also since [Ghanem and Smith \(2020\)](#) adapted in a VAR framework [Roberts and Schlenker's](#) model, which provides the basis for our auxiliary model.

²[Lowry et al. \(1987\)](#) recognized that most of the storage model parameters could be estimated using 2SLS and apply this insight to the model calibration. However, they do not estimate all the parameters nor do they consider the issue of model validation. Similarly, [Steinwender \(2018\)](#) combines price and quantity data to estimate most of the parameters of a storage-trade model.

The rest of the paper is as follows. Section 2 describes the storage model. Section 3 presents 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. Roberts and Schlenker (2013) take the example of the soybean rust which is observable from the previous growing season. Hendricks et al. (2015) note that in such an annual model this can also be related to the crops already planted, for example in the Southern Hemisphere or winter wheat in the Northern Hemisphere. The cost shock aggregates a variety of unrepresented 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

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

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

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

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 $\psi_{t+1} = \eta_t + \varepsilon_{t+1}$ 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 some incentive shocks (e.g., because of changes to agricultural and trade policies or because of price changes in competing crops).

Storers Competitive storers are risk-neutral. To store an amount $X_t \geq 0$ from period t until $t + 1$ competitive storers incur an opportunity cost and 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 balance growth path (i.e., in the absence of shocks) and k is the per-unit physical storage cost expressed as a percentage of this price. Assuming rational expectations and taking account of the non-negativity constraint on storage yields the following arbitrage condition

$$\beta E_t P_{t+1} - P_t - k\bar{P}_t \leq 0, = 0 \text{ if } X_t > 0. \quad (3)$$

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.

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 $v \sim \mathcal{N}(0, \sigma_v^2)$:

$$\mu_{t+1} = \rho_\mu \mu_t + v_{t+1}. \quad (4)$$

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

$$X_{t-1} + H_{t-1} e^{\eta_{t-1} + \varepsilon_t} = D_t(P_t) e^{\mu_t} + X_t. \quad (5)$$

2.2 Stationary model

Trends and steady-state growth Detrended variables and functions are denoted in lower case and relate to their trending counterparts based on the following relations

$$P_{t+1} = p_{t+1} e^{g_p t}, \quad (6)$$

$$D_t(P_t) = e^{g_q t} d(p_t), \quad (7)$$

$$\Gamma_t'(H_t) = \gamma'(h_t), \quad (8)$$

where g_q is the assumed rate of growth of quantities. In equation (6), 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. 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 (5):

$$x_{t-1} e^{-g_q} + h_{t-1} e^{\eta_{t-1} + \varepsilon_t} = d(p_t) e^{\mu_t} + x_t. \quad (9)$$

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, 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 (3) can be expressed with detrended variables as

$$\beta e^{g_p} E_t p_{t+1} - p_t - k\bar{p} \leq 0, \quad = 0 \text{ if } x_t > 0. \quad (10)$$

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 to the price trend, $g_p < -\log \beta$ becomes a necessary condition for the existence of a stationary rational expectations equilibrium, which is always satisfied for decreasing trends.

Stationary equilibrium In equation (9), 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

$$s_t \equiv x_{t-1} e^{-g_q} + h_{t-1} e^{\eta_{t-1} + \varepsilon_t}. \quad (11)$$

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

$$h_t : \beta e^{\eta_t - \omega_t} E_t (p_{t+1} e^{\varepsilon_{t+1}}) = \gamma' (h_t), \quad (12)$$

$$x_t : \beta e^{g_p} E_t p_{t+1} - p_t - k\bar{p} \leq 0, = 0 \text{ if } x_t > 0, \quad (13)$$

$$p_t : s_t = x_t + d(p_t) e^{\mu_t}. \quad (14)$$

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. 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

$$d(p_t) = \bar{d} \left(\frac{p_t}{\bar{p}} \right)^{\alpha_D}, \quad (15)$$

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:

$$\gamma' (h_t) = \beta \bar{p} \left(\frac{h_t}{\bar{d}} \right)^{1/\alpha_S}, \quad (16)$$

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

$$\frac{s_t}{\bar{d}} = \frac{x_{t-1}}{\bar{d}} e^{-g_q} + \frac{h_{t-1}}{\bar{d}} e^{\eta_{t-1} + \varepsilon_t}, \quad (17)$$

$$\frac{h_t}{\bar{d}} = \left[e^{\eta_t - \omega_t} E_t \left(\frac{p_{t+1}}{\bar{p}} e^{\varepsilon_{t+1}} \right) \right]^{\alpha_S}, \quad (18)$$

$$\beta e^{g_p} E_t \left(\frac{p_{t+1}}{\bar{p}} \right) - \frac{p_t}{\bar{p}} - k \leq 0, = 0 \text{ if } \frac{x_t}{\bar{d}} > 0, \quad (19)$$

$$\frac{s_t}{\bar{d}} = \frac{x_t}{\bar{d}} + \left(\frac{p_t}{\bar{p}} \right)^{\alpha_D} e^{\mu_t}. \quad (20)$$

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 (4) and (17)–(20) represent a non-linear rational expectations system based on the variables μ_t, s_t, h_t, x_t , and p_t driven by the innovations $\{\eta_t, \omega_t, \varepsilon_t, \nu_t\}$. 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 $q_{t+1}^e = E_t q_{t+1} \exp(-\sigma_\varepsilon^2/2) = h_t \exp(\eta_t)$, which is the expected production corrected for the mean harvest-time shock given by

$$q_{t+1}^e = \bar{d} e^{\eta_t} \left[e^{\eta_t - \omega_t} E_t \left(\frac{p_{t+1}}{\bar{p}} e^{\varepsilon_{t+1}} \right) \right]^{\alpha_S}. \quad (21)$$

In this case, the transition equation is defined as

$$s_{t+1} = x_t e^{-gq} + q_{t+1}^e e^{\varepsilon_{t+1}}. \quad (22)$$

We combine the two planting-time shocks that appear in equation (21) to form the aggregate planting-time shock $\varphi_t = (1 + \alpha_S)\eta_t - \alpha_S\omega_t$. This allows for a further simplification of the supply equation:

$$q_{t+1}^e = \bar{d} e^{\varphi_t} \left[E_t \left(\frac{p_{t+1}}{\bar{p}} e^{\varepsilon_{t+1}} \right) \right]^{\alpha_S}. \quad (23)$$

We can see also that in the absence of demand for stock, the market clearing equation (14) 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 $\{\tilde{s}_t, \eta_t, \mu_t\}$:

$$q_{t+1}^e / \bar{d} = \mathcal{Q}(\tilde{s}_t, \varphi_t, \mu_t), \quad (24)$$

$$x_t / \bar{d} = \mathcal{X}(\tilde{s}_t, \varphi_t, \mu_t), \quad (25)$$

$$p_t / \bar{p} = \mathcal{P}(\tilde{s}_t, \varphi_t, \mu_t). \quad (26)$$

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:

$$\mathcal{P}(\tilde{s}_t, \varphi_t, \mu_t) = \max \left\{ \beta e^{s_p} E_t \left[\mathcal{P} \left((\mathcal{X}(\tilde{s}_t, \varphi_t, \mu_t) e^{-s_q} + \mathcal{Q}(\tilde{s}_t, \varphi_t, \mu_t) e^{\varepsilon_{t+1}}) e^{-\mu_{t+1}}, \varphi_{t+1}, \mu_{t+1} \right) \right] - k, (\tilde{s}_t / \bar{d})^{1/\alpha_D} \right\}, \quad (27)$$

$$e^{\varphi_t} \left\{ E_t \left[\mathcal{P} \left((\mathcal{X}(\tilde{s}_t, \varphi_t, \mu_t) e^{-s_q} + \mathcal{Q}(\tilde{s}_t, \varphi_t, \mu_t) e^{\varepsilon_{t+1}}) e^{-\mu_{t+1}}, \varphi_{t+1}, \mu_{t+1} \right) e^{\varepsilon_{t+1}} \right] \right\}^{\alpha_S} = \mathcal{Q}(\tilde{s}_t, \varphi_t, \mu_t). \quad (28)$$

Equation (27) 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.

Let us now consider the cutoff situation of no storage but where the storage arbitrage equation still holds. Then equation (27) simplifies to

$$\beta e^{s_p} E_t \left[\mathcal{P} \left(\mathcal{Q}(\tilde{s}_t, \varphi_t, \mu_t) e^{\varepsilon_{t+1} - \mu_{t+1}}, \varphi_{t+1}, \mu_{t+1} \right) \right] - k = (\tilde{s}_t / \bar{d})^{1/\alpha_D}. \quad (29)$$

This equation defines the cutoff net availability $\tilde{s}^*(\varphi, \mu)$ as a function of the contemporaneous planting-time supply and demand shocks. For a given pair of shocks $\{\varphi, \mu\}$, any net availability below $\tilde{s}^*(\varphi, \mu)$ implies the absence of a demand for stocks because prices are too high to make speculation profitable. Similarly, we can define a cutoff price policy function $p^*(\varphi, \mu) = \bar{p}(\tilde{s}^*(\varphi, \mu))^{1/\alpha_D}$ above which there is no storage. Finally, in the absence of stocks, the expected production is a function only of the planting-time supply and demand shocks, $q^{e,*}(\varphi, \mu)$, and is the solution to

$$\bar{d} e^{\varphi_t} \left\{ E_t \left[\mathcal{P} \left(q^{e,*}(\varphi_t, \mu_t) e^{\varepsilon_{t+1} - \mu_{t+1}}, \varphi_{t+1}, \mu_{t+1} \right) e^{\varepsilon_{t+1}} \right] \right\}^{\alpha_S} = q^{e,*}(\varphi_t, \mu_t). \quad (30)$$

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 USDA statistics but for many countries they are based on USDA estimates in the absence of official

statistics.³ 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 vector $\theta \in \Theta$. Our storage model includes fourteen parameters, nine 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)**.

3.1.1 Production

Expressed in logarithm, the supply equation (18) is

$$\log q_t = \log (h_{t-1} e^{\psi_t}) = \log (\bar{d}/\bar{p}^{\alpha_S}) + \alpha_S (\eta_{t-1} - \omega_{t-1}) + \alpha_S \log (E_{t-1} (p_t e^{\varepsilon_t})) + \psi_t. \quad (31)$$

In this equation, $\eta_{t-1} - \omega_{t-1}$ and $E_{t-1} (p_t e^{\varepsilon_t})$ 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

³The 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.

estimation equation

$$\log q_t = a_q + b_q \log(E_{t-1} p_t) + c_q \psi_t + u_{q,t}. \quad (32)$$

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} . 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 under the model assumptions, they are not correlated with the planting-time shocks and thus with the residuals. The first-stage equation is

$$\log(E_{t-1} p_t) = a_{E_p} + b_{E_p} \psi_{t-1} + c_{E_p} \psi_t + u_{E_p,t}. \quad (33)$$

This supply-side estimation strategy deserves a few comments. First, substituting the expected price $E_{t-1} p_t$ for the producer incentive price $E_{t-1}(p_t e^{\varepsilon_t})$ 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 $b_q \log(E_{t-1} p_t) = \alpha_S \log(E_{t-1} p_t e^{\varepsilon_t})$, we can write

$$\log q_t - b_q \log(E_{t-1} p_t) = a_q + c_q \psi_t + u_{q,t} = \log(\bar{d}/\bar{p}^{\alpha_S}) + \alpha_S(\eta_{t-1} - \omega_{t-1}) + \psi_t. \quad (34)$$

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

$$c_q = \frac{\text{cov}(\log q_t - \alpha_S \log(E_{t-1} p_t), \psi_t)}{\text{var} \psi_t} \quad (35)$$

$$= \frac{\text{cov}(\alpha_S(\eta_{t-1} - \omega_{t-1}) + \psi_t, \psi_t)}{\text{var} \psi_t} \quad (36)$$

$$= 1 + \alpha_S \sigma_\eta \frac{\sigma_\eta - \rho_{\eta,\omega} \sigma_\omega}{\sigma_\psi^2}. \quad (37)$$

This 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 (32) to get

$$\log q_t - \log(\bar{d}/\bar{p}^{\alpha_S}) - b_q \log(E_{t-1} p_t) = c_q \psi_t + u_{q,t} = \alpha_S(\eta_{t-1} - \omega_{t-1}) + \psi_t = \vartheta_t. \quad (38)$$

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.

3.1.2 Consumption

From (15), logged consumption (denoted c_t), is given by

$$\log c_t = \log(d(p_t) e^{\mu_t}) = \log(\bar{d}/\bar{p}^{\alpha_D}) + \alpha_D \log p_t + \mu_t. \quad (39)$$

By calculating $\log c_t - \rho_\mu \log c_{t-1}$ and using equation (4), we can recover the innovation v_t in the demand equation:

$$\log c_t = (1 - \rho_\mu) \log(\bar{d}/\bar{p}^{\alpha_D}) + \alpha_D \log p_t - \alpha_D \rho_\mu \log p_{t-1} + \rho_\mu \log c_{t-1} + v_t. \quad (40)$$

The fact that v_t is unobservable but correlated with p_t implies that an OLS estimation of equation (40) would again lead to an omitted variable bias. We solve this by instrumenting prices with the yield shocks. Thus, the estimation equation is

$$\log c_t = a_c + b_c \log p_t + c_c \log p_{t-1} + d_c \log c_{t-1} + u_{c,t}, \quad (41)$$

with the associated first stage

$$\log p_t = a_p + b_p \psi_t + c_p \log p_{t-1} + d_p \log c_{t-1} + u_{p,t}. \quad (42)$$

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 (39) is estimated directly using

$$\log p_t = a_p + b_p \psi_t + u_{p,t} \quad (43)$$

as first stage, since [Roberts and Schlenker](#)'s focus is on the demand elasticity and not the other parameters. These methods are mostly equivalent in terms of estimating the demand elasticity.

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:

$$\log q_t = a_q + b_q \log(E_{t-1} p_t) + c_q \psi_t + u_{q,t}, \quad (44)$$

$$\log(E_{t-1} p_t) = a_{E_p} + b_{E_p} \psi_{t-1} + c_{E_p} \psi_t + u_{E_p,t}, \quad (45)$$

$$\log c_t = a_c + b_c \log p_t + c_c \log p_{t-1} + d_c \log c_{t-1} + u_{c,t}, \quad (46)$$

$$\log p_t = a_p + b_p \psi_t + c_p \log p_{t-1} + d_p \log c_{t-1} + u_{p,t}, \quad (47)$$

$$\psi_t = a_\psi + u_{\psi,t}. \quad (48)$$

The model includes both the first and second-stage equations presented previously, and equation (48) 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 (44) and (46) 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 accurate, 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 (44)–(48) estimated by OLS. However, we retain the first-stage equations in the system because they contain information not provided in the other equations, and because 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.⁴ 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 advantages. 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 (37), c_q and similarly σ_{u_q} are both nonlinear combinations of α_S , σ_ε , σ_η , σ_ω , and $\rho_{\varepsilon,\omega}$. In [Hendricks et al. \(2015\)](#), c_{E_p} is related to the predictability of the yield shocks, and thus to σ_η . In equation (46), 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 (47), c_p is linked to the first-order autocorrelation of $\log p$, which conditional

⁴See also [Simonovska and Waugh \(2014\)](#) for an estimation approach in which a biased auxiliary model is used to obtain an unbiased simulated estimator.

on the other parameters, depends directly on the storage cost k . More precisely, a lower storage cost implies more storage and hence a higher price autocorrelation (Gouel and Legrand, 2017, Figure 2), and vice versa. In equation (48), $\sigma_{u_\psi}^2 = \sigma_\varepsilon^2 + \sigma_\eta^2$. Finally, the inclusion of $\sigma_{u_{Ep}}$ 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 $\zeta_{\tau T}(\theta)$ denotes the counterpart of ζ_T estimated on artificial data generated by the storage model for a given set of parameters θ . The size of the simulated sample then is $\tau T + 1 + t^{\text{burn}}$ with the integers $\tau \geq 1$ and $t^{\text{burn}} \geq 1$. The first t^{burn} simulations are dropped as burn-in periods to remove the influence of the initial state. The final $\tau T + 1$ 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

$$\hat{\theta} = \arg \min_{\theta \in \Theta} \left[\hat{\zeta}_T - \hat{\zeta}_{\tau T}(\theta) \right]' W \left[\hat{\zeta}_T - \hat{\zeta}_{\tau T}(\theta) \right], \quad (49)$$

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 section A.1. 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 Gouriéroux et al. (1993), and assuming that W is the optimal weighting matrix, the variance-covariance matrix for the parameter estimates converges asymptotically to

$$\left(1 + \frac{1}{\tau} \right) (J' W J)^{-1}, \quad (50)$$

where $J = E[\partial \zeta_{\tau T}(\theta) / \partial \theta]$ is a 15×9 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 and autocorrelation for the standard regression parameters ($b_q, c_q, b_{Ep}, c_{Ep}, b_c, c_c, d_c, b_p, c_p, d_p$), and using the following formulas for the standard deviations ($\sigma_{u_q}, \sigma_{u_{Ep}}, \sigma_{u_c}, \sigma_{u_p}, \sigma_{u_\psi}$):

$$\text{var}(\sigma^{\text{OLS}}) = \frac{(\sigma^{\text{OLS}})^2}{2(T-l)} \quad \text{and} \quad \text{var}(\sigma^{2\text{SLS}}) = \frac{(\sigma^{2\text{SLS}})^2}{2(T-l)R_p^2}, \quad (51)$$

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).

The 15 parameters included in the auxiliary model and the 9 to be estimated in the storage model means that there are overidentification restrictions, which will be tested using the statistics

$$\frac{T\tau}{1+\tau} \min_{\theta \in \Theta} \left[\hat{\xi}_T - \hat{\xi}_{\tau T}(\theta) \right]' W \left[\hat{\xi}_T - \hat{\xi}_{\tau T}(\theta) \right], \quad (52)$$

which follows asymptotically a chi-square distribution with 6 degrees of freedom (Gourieroux et al., 1993).

Since it is costly to evaluate the objective in equation (49), 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 which would make it difficult to define or solve the model (see table A2). Furthermore, to limit the risk of finding only a local optimum, with the exception of the Monte Carlo experiments in the next section which uses a unique starting point, the optimization algorithm starts from 500 different initial values of θ . 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 our preferred estimates in section 6 for ρ_μ , σ_v , α_D , α_S , and σ_ψ , the trend estimates for g_q and g_p , and an arbitrary choice for the other parameters. The parameter values used are $\beta = 0.98$, $g_q = 2.54\%$, $g_p = -2.01\%$, $\rho_\mu = 0.5$, $\rho_{\eta,\omega} = -0.3$, $\sigma_\eta = 1.2\%$, $\sigma_\varepsilon = 2\%$, $\sigma_v = 1.6\%$, $k = 5\%$, $\alpha_D = -0.065$, and $\alpha_S = 0.058$. 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: $\sigma_\omega = \{0.05, 0.1, 0.2\}$. 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.⁵ 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} , c_q^{OLS} in combination with the other parameters, and delivers unbiased and consistent elasticity estimates.⁶

Nevertheless, using the indirect inference approach, two parameters are difficult to estimate: the storage cost (k), and the correlation between the planting-time shocks ($\rho_{\eta,\omega}$) with more than 55% and 104% RMSE, explained mostly by very volatile estimates given the very small biases. The difficulty involved in estimating k could stem from the fact that it is identified only indirectly, in part through its 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 this parameter. However, in our context they 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 (37)). 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 21%—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 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 the supply elasticity discussed above).

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. This 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

⁵In 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 (31) where $E_{t-1}(p_t \exp(\varepsilon_t))$ is replaced by $E_{t-1} p_t$. At -1.5% , this bias is small under these parameters.

⁶This 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	σ_ψ (%)	σ_v (%)	σ_ϑ (%)	$c_q - 1$	α_D	α_S
<i>Panel A. OLS</i>							
$\sigma_\omega = 5\%$							
Mean	0.34	2.33	1.24	2.41	0.027	-0.018	0.052
St. dev.	0.13	0.23	0.14	0.24	0.017	0.011	0.003
RMSE (%)	41.31	9.89	24.10	9.94	54.370	75.140	12.435
SE	0.12	0.22	0.12		0.017	0.009	0.003
$\sigma_\omega = 10\%$							
Mean	0.35	2.33	1.25	2.48	0.037	-0.018	0.044
St. dev.	0.13	0.23	0.14	0.25	0.033	0.011	0.006
RMSE (%)	40.36	9.89	23.73	10.03	68.651	74.161	25.947
SE	0.12	0.22	0.12		0.033	0.009	0.006
$\sigma_\omega = 20\%$							
Mean	0.36	2.33	1.26	2.66	0.046	-0.020	0.024
St. dev.	0.13	0.23	0.14	0.27	0.062	0.011	0.011
RMSE (%)	37.81	9.89	22.69	10.87	83.784	71.326	62.092
SE	0.12	0.22	0.12		0.062	0.009	0.010
<i>Panel B. 2SLS</i>							
$\sigma_\omega = 5\%$							
Mean	0.45	2.33	1.63	2.43	0.035	-0.068	0.058
St. dev.	0.18	0.23	0.33	0.24	0.020	0.022	0.009
RMSE (%)	37.27	9.89	20.74	9.99	59.239	33.927	16.112
SE	0.17	0.22	0.27		0.021	0.018	0.009
$\sigma_\omega = 10\%$							
Mean	0.45	2.33	1.63	2.54	0.055	-0.068	0.060
St. dev.	0.18	0.23	0.33	0.26	0.040	0.022	0.018
RMSE (%)	36.83	9.89	20.76	10.18	74.455	33.972	31.303
SE	0.17	0.22	0.28		0.042	0.018	0.017
$\sigma_\omega = 20\%$							
Mean	0.45	2.33	1.63	2.86	0.097	-0.068	0.064
St. dev.	0.17	0.23	0.34	0.36	0.083	0.022	0.039
RMSE (%)	35.91	9.89	21.15	12.92	90.071	34.666	68.062
SE	0.16	0.22	0.29		0.086	0.019	0.036

Notes: Monte Carlo experiment based on 500 replications, with a sample size $T = 56$. True values: $\rho_\mu = 0.5$, $\sigma_\psi = 2.33\%$, $\sigma_v = 1.6\%$, $\alpha_D = -0.065$, and $\alpha_S = 0.058$. The values of σ_ϑ and c_q vary with σ_ω as follows $\sigma_\vartheta = \{2.47, 2.61, 3.01\}$ and $c_q = \{1.035, 1.054, 1.092\}$ 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.

small-sample bias (tables A4 and A5 show that any small-sample biases is even less apparent in

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_v (%)	k (%)	α_D	α_S
$\sigma_\omega = 5\%$	OID: 0.043								
Mean	0.46	-0.32	5.05	1.17	1.94	1.59	4.70	-0.063	0.058
St. dev.	0.12	0.33	0.65	0.35	0.25	0.27	2.76	0.015	0.004
RMSE (%)	24.71	109.20	13.00	29.31	12.93	17.04	55.43	23.605	7.483
ASE	0.08	0.65	0.66	0.44	0.25	0.22	1.76	0.013	0.004
$\sigma_\omega = 10\%$	OID: 0.049								
Mean	0.46	-0.34	10.14	1.16	1.95	1.59	4.65	-0.063	0.058
St. dev.	0.12	0.31	1.54	0.37	0.27	0.27	2.77	0.015	0.008
RMSE (%)	25.12	105.79	15.44	30.76	13.49	16.70	55.78	23.336	14.473
ASE	0.08	0.68	1.44	0.44	0.26	0.22	1.76	0.013	0.007
$\sigma_\omega = 20\%$	OID: 0.038								
Mean	0.45	-0.34	20.79	1.15	1.95	1.59	4.63	-0.063	0.059
St. dev.	0.12	0.31	5.20	0.39	0.28	0.26	2.83	0.015	0.016
RMSE (%)	25.68	104.23	26.31	32.83	14.14	16.07	57.10	22.982	27.492
ASE	0.08	0.60	4.58	0.47	0.27	0.22	1.83	0.013	0.014

Notes: Monte Carlo experiment based on 500 replications, with a sample size $T = 56$. For $\sigma_\omega = 5\%$, 10% , and 20% , respectively 33, 27, and 26 replications had to be dropped due to non-convergence. True values: $\rho_\mu = 0.5$, $\rho_{\eta,\omega} = -0.3$, $\sigma_\eta = 1.2\%$, $\sigma_\varepsilon = 2\%$, $\sigma_v = 1.6\%$, $k = 5\%$, $\alpha_D = -0.065$, and $\alpha_S = 0.058$. 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 (50), and represents the average standard errors calculated at the solutions. OID is the empirical size of the chi-square test of overidentifying restrictions.

longer samples), although the downward bias on the standard error of k might be a concern. The empirical size of the OID statistic is 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).⁷

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 makes the lagged yield shocks a weaker instrument because their role in explaining price changes declines. 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 more robust to σ_ω , and has RMSEs which deteriorate less as this parameter increases.

⁷This size-distortion issue related to specification tests is acknowledged in the literature (see, e.g., Ruge-Murcia, 2007; Michaelides and Ng, 2000).

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 2SLS regression based model for the robustness checks.

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. \(2015\)](#):

$$\Psi_t = \frac{\sum_l \sum_i A_{lit} \kappa_l Y_{lit}}{\sum_l \sum_i A_{lit} \kappa_l \hat{Y}_{lit}} = \sum_l \sum_i \rho_{lit} \Psi_{lit}, \quad (53)$$

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

$$\rho_{lit} = \frac{A_{lit} \kappa_l \hat{Y}_{lit}}{\sum_{l'} \sum_{i'} A_{l'i'} \kappa_{l'} \hat{Y}_{l'i'}} \quad (54)$$

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; however, this 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. If we ignore the risk premium and the basis risk, the expected price can be approximated by futures prices, and 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 retrieved 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).⁸ 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 (54):

$$P_t = \frac{\sum_{l \neq \text{rice}} (\sum_i \rho_{lit}) P_{lt|t} / \kappa_l}{\sum_{l \neq \text{rice}} \sum_i \rho_{lit}} \text{ and } E_{t-1} P_t = \frac{\sum_{l \neq \text{rice}} (\sum_i \rho_{lit}) P_{lt|t-1} / \kappa_l}{\sum_{l \neq \text{rice}} \sum_i \rho_{lit}}, \quad (55)$$

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.⁹ 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

⁸At 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.

⁹It 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](#)).

structural breaks with or without a linear or quadratic deterministic trend under both the null and alternative hypotheses.




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.

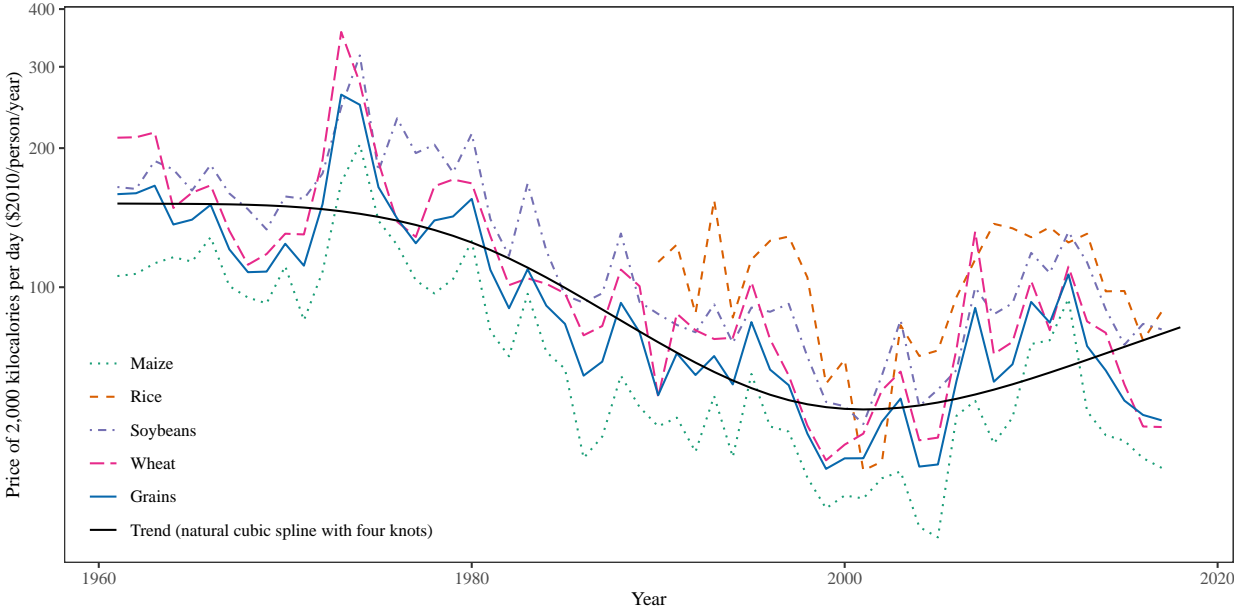


Figure 2: Real caloric prices at delivery. The y-axis is the annual cost of 2,000 kilocalories per day.

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).¹⁰ 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 1980.¹¹

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 balanced 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.¹² 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.¹³ 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 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

¹⁰Based on 5,000 replications of sample sizes $T = 100$.

¹¹It 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).

¹²Unless 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.

¹³Recall that in the KPSS test the null is a trend-stationary series.

(consumption and production) on a common linear trend and similarly with the logged prices to estimate $g_p = -2.01\%$.

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, 1990–2017)

Commodity	Maize	Rice	Soybeans	Wheat
Maize				
Rice	0.661			
Soybeans	0.858	0.766		
Wheat	0.790	0.598	0.775	
Grains	0.923	0.680	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.

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

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	0.168	0.236
Supply price ($\log(E_t p_{t+1})$)	0.652	0.235	0.193
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	0.051	0.024

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 appear to have the expected signs. Current and expected prices are strongly correlated, consistent with equation (3) 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 .

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 ($\log(E_t p_{t+1})$)	0.939			
Consumption ($\log(c_t)$)	−0.488	−0.452		
Production ($\log(q_t)$)	−0.406	−0.271	0.395	
Yield shock (ψ_t)	−0.534	−0.499	0.528	0.775

6 Estimation

6.1 Structural parameters

Before analyzing the results for 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 , and ρ_μ), and indirect information about the others.

6.1.1 Instrumental variable estimations

Supply and demand equations can be estimated on raw trending data. Following [Roberts and Schlenker](#), we augment each equation including the first stages, by adding the trend variables

generated by the 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 three deviations from their approach: a different procedure to construct the yield shock (in line with Hendricks et al., 2015), detrending of yields using a 4-knot spline rather than a 3-knot spline consistent with our longer sample, and use of R instead of Stata which implies a different treatment of the residuals autocorrelation for the robust standard errors.

Table 6: Supply equation estimation

	(1)	(2)	(3)
<i>Panel A. 2SLS</i>			
Supply elasticity b_q	0.087* (0.048)	0.074*** (0.021)	0.081*** (0.020)
Shock c_q	1.145*** (0.092)	1.145*** (0.097)	1.130*** (0.106)
<i>Panel B. First stage</i>			
Lagged shock b_{EP}	-4.057*** (0.623)	-3.789*** (0.569)	-3.829*** (0.468)
Shock c_{EP}	-2.461 (2.025)	-2.370 (1.564)	-2.336 (1.585)
<i>Panel C. OLS</i>			
Supply elasticity b_q	0.135*** (0.015)	0.058*** (0.015)	0.060*** (0.015)
Shock c_q	1.290*** (0.093)	1.097*** (0.078)	1.072*** (0.085)
$\sigma_{u_q^{2SLS}}$	0.028	0.015	0.015
$\sigma_{u_{EP}}$	0.229	0.165	0.167
$\sigma_{u_q^{OLS}}$	0.026	0.015	0.015
$\sigma_{\delta^{OLS}}$	0.039	0.030	0.029
First stage F -stat	9.230	15.468	15.347
p -value for Hausman test	0.166	0.417	0.306
Observations	56	56	56
Spline knots	3	4	5

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

Table 6 reports the estimations of the supply equation. 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 (line Shock in panels A and C) 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). Consistent with Hendricks et al.'s insights (2015), the OLS and

Table 7: Demand equation estimation

	(1)	(2)	(3)
<i>Panel A. 2SLS</i>			
Demand elasticity b_c	-0.052** (0.024)	-0.065** (0.026)	-0.061** (0.027)
Lagged price c_c	0.041*** (0.015)	0.019 (0.013)	0.015 (0.012)
Lagged demand d_c	0.980*** (0.061)	0.488*** (0.136)	0.400** (0.150)
<i>Panel B. First stage</i>			
Shock b_p	-4.265*** (0.973)	-4.093*** (0.963)	-3.994*** (1.092)
Lagged price c_p	0.568*** (0.078)	0.484*** (0.069)	0.497*** (0.074)
Lagged demand d_p	1.444* (0.818)	-0.144 (1.749)	0.501 (2.098)
<i>Panel C. OLS</i>			
Demand elasticity b_c	-0.012 (0.009)	-0.021** (0.008)	-0.018** (0.008)
Lagged price c_c	0.015 (0.009)	-0.005 (0.011)	-0.010 (0.011)
Lagged demand d_c	0.949*** (0.041)	0.547*** (0.109)	0.413*** (0.133)
<i>Panel D. 2SLS using Roberts and Schlenker's approach (eqs. (39) for 2nd stage and (43) for 1st)</i>			
Demand elasticity b_c	-0.069* (0.039)	-0.079*** (0.029)	-0.066** (0.028)
$\sigma_{u_c^{2SLS}}$	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
$\sigma_{\mu^{2SLS}}$	0.090	0.019	0.017
First stage F -stat (panel A)	15.363	13.910	13.106
p -value for Hausman test (panel A)	0.047	0.014	0.017
First stage F -stat (panel D)	13.735	20.252	17.765
p -value for Hausman test (panel D)	0.001	0.008	0.018
Observations	56	56	56
Spline knots	3	4	5

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

2SLS supply elasticity estimates show only small differences indicating that using the yield shock as a control variable helps to mitigate the omitted variable bias. This is confirmed by the Hausman

test which does not reject the null of exogenous expected prices. The estimations using four and five knots are similar and different from the estimations using three knots which is in line with the previous stationarity test results. Therefore, in what follows our benchmark estimate is the OLS with four knots. For this specification, total supply shocks have a standard deviation σ_{ϑ} equal to 0.03, 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 appears to be explained by our longer sample (table A8). We use equation (41) to estimate both the demand elasticity and autocorrelation of the demand shock. This contrasts with Roberts and Schlenker (2013) who uses equation (39) 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 two approaches. Estimates of the autocorrelation of the demand shocks differ depending on the number of knots. ρ_{μ} using a 3-knot spline is not statistically different from 1 indicating non-stationary demand, which is in line with 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 0.49 (0.14) and 0.4 (0.15) for four and five knots the estimates are similar. The final parameter which can be identified from the demand estimation is the standard error of the demand shock. Using 4- and 5-knot splines, it seems that σ_v (estimated by $\sigma_{u_c^{2SLS}}$) is about 0.016, which is slightly lower than the volatility 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 insignificant, 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 (for both equations instrumented but the correlations do not change if the supply equation is estimated using OLS). This low correlation means that the 2SLS and 3SLS results are very similar and thus are not reported here.¹⁴ 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

¹⁴This equivalence differs from Roberts and Schlenker (2013) who finds different results between the 3SLS and 2SLS estimates. This gap is explained mostly by the fact that Stata imposes the same set of instruments on all equations in 3SLS, implying that in Roberts and Schlenker the 3SLS does not correspond exactly to a feasible generalized least squares version of the 2SLS.

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.

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	0.572	(0.087)	0.472	(0.112)
$\rho_{\eta,\omega}$	-0.628	(0.520)	-0.526	(0.530)
σ_ω	0.217	(0.042)	0.224	(0.044)
σ_η	0.011	(0.007)	0.010	(0.007)
σ_ε	0.023	(0.004)	0.024	(0.004)
σ_v	0.015	(0.002)	0.014	(0.002)
k	0.032	(0.013)	0.029	(0.013)
α_D	-0.044	(0.012)	-0.034	(0.010)
α_S	0.075	(0.016)	0.070	(0.015)
σ_φ	0.025	(0.007)	0.023	(0.006)
σ_ψ	0.026	(0.005)	0.026	(0.004)
σ_μ	0.019	(0.003)	0.016	(0.002)
σ_ϑ	0.035	(0.006)	0.033	(0.005)
OID p -value	0.034		0.033	

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

Most parameters are estimated precisely for both trend specifications despite the rather short sample size.¹⁵ The exceptions are the magnitude of the planting-time yield shock σ_η and its correlation ($\rho_{\eta,\omega}$) with the other planting-time shock (ω). The fact that $\rho_{\eta,\omega}$ is not precisely estimated is not surprising, given the large RMSE value 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. Although σ_η is less precisely identified than expected given the Monte Carlo experiment results, this has no consequences for the aggregated planting-time shock σ_φ which is accurately estimated (see lower panel in the table).

The parameters estimated using both methods (i.e., ρ_μ , σ_v , α_D , and α_S), 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

¹⁵Recall from the Monte Carlo analysis that the standard error of k tends to be underestimated meaning that its apparent precision must be interpreted with caution.

estimates of ρ_μ are sufficiently higher and the estimates of α_D sufficiently lower to be a potential concern and could indicate some model misspecification on the demand side. This is confirmed later by the poor fit of some demand-related moments.

The volatility of the cost shock σ_ω is about 22% which is an order of magnitude larger than the estimates of the other shocks. However, the cost shock has no direct effect on quantities. Making it comparable to the other shocks requires its multiplication by α_S which produces 1.6% with four and five knots or a more than 35% larger contribution than 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 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 φ shock. The standard deviation of φ is equivalent to the standard deviation of harvest-time yield shock σ_ε . 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 almost twice as large as the standard deviation of the demand shock.

The physical per-unit storage cost (k) is estimated a 3% of the steady-state price. 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% at the steady state $(k + 1 - \beta e^{g_p})$. 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 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.

We can compare the results obtained using the OLS estimates for all equations as the auxiliary model with the results obtained from estimating supply and demand equations using 2SLS. The latter estimates are available in table A9 in the Appendix. Comparing tables 8 and A9, none of the parameter estimates can be considered significantly different. However, 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 less precise estimates.

Overall, these results suggest that our indirect inference approach returns fairly precise parameter estimates which, with some caveats, 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

It is interesting also to check the similarity between the estimates of the auxiliary model parameters based on observations and based on simulations. The overidentification test rejects the model specification (in table 8 but not in table A9), although the test is biased against the rejection (according to the Monte Carlo experiment).

To more clearly identify the parameters causing this rejection, table 9 reports the respective auxiliary parameters obtained 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 demand elasticity parameter, b_c , 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 widely between the two columns (e.g., b_{Ep} or d_p), they are estimated imprecisely in the auxiliary model, and thus were given 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	0.015	0.034
c_q	1.097	0.078	1.139
σ_{u_q}	0.015	0.001	0.015
b_c	-0.021	0.008	-0.007*
c_c	-0.005	0.011	0.006
d_c	0.547	0.109	0.516
σ_{u_c}	0.014	0.001	0.013
b_{Ep}	-2.370	1.564	-1.223
c_{Ep}	-3.789	0.569	-2.963
$\sigma_{u_{Ep}}$	0.165	0.017	0.159
b_p	-4.093	0.963	-4.837
c_p	0.484	0.069	0.467
d_p	-0.144	1.749	2.525
σ_{u_p}	0.180	0.018	0.184
σ_{u_ψ}	0.024	0.002	0.026
b_q^{2SLS}	0.075	0.021	0.075
b_c^{2SLS}	-0.065	0.026	-0.044

Notes: Standard errors robust to heteroskedasticity and autocorrelation for the parameters and based on equation (51) 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 also compare the fit with the supply and demand elasticities estimated using 2SLS (see lower panel in table 9). The fit for the supply elasticity is perfect. Given that the model tends to underestimate b_q estimated by OLS (albeit insignificantly) but not 2SLS, this difference between OLS and 2SLS confirms the possible overestimation of σ_ω highlighted previously. Due to its large standard error, the demand elasticity estimated by 2SLS is not significantly different between observations and simulations, but is biased downward with a similar magnitude to the OLS estimates.

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

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 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 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 rich 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 \(1992\)](#).¹⁶

However, it can be seen that the storage model fails to match some moments. As suggested by the discrepancies between the OLS and the indirect inference estimates, these moments are related mostly 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 only -0.07 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 and if we consider expected instead of current prices. There are some problems with other moments: the model-based variance in consumption is too small compared to observations, the model underestimates the very high positive correlation between current and future prices, and overstates the positive correlation between production and yield shocks. Both correlation issues are potentially of less concern because in those cases the covariances on simulations remain an order of magnitude closer to those in the data although they fall outside the confidence intervals.

The correlation between consumption and price is governed in the model by the demand elasticity and the size of the demand shocks relative to the supply shocks. In the absence of demand shocks, the correlation would be -1 . The higher the variance of demand shocks, the higher the correlation which turns positive for demand shocks with sufficiently large variance. Similarly, for a given volatility of demand shocks, the more inelastic consumption demand, the stronger its correlation with price. We have shown that the indirect inference estimations lead to higher demand shock autocorrelation and more inelastic demand compared to those obtained using 2SLS. These differences between 2SLS and indirect inference contribute to explaining the difficulty related to

¹⁶However, 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	0.268
$\sigma_{\ln c}$	0.019	0.002	0.015**
$\sigma_{\ln q}$	0.028	0.002	0.032
$\sigma_{\ln E p}$	0.193	0.018	0.182
σ_{ψ}	0.024	0.002	0.026
$\phi_{\ln p}(1)$	0.576	0.110	0.504
$\phi_{\ln c}(1)$	0.642	0.146	0.481
$\phi_{\ln q}(1)$	0.042	0.140	-0.092
$\phi_{\ln E p}(1)$	0.652	0.116	0.575
$\phi_{\psi}(1)$	0.146	0.142	0.002
$\phi_{\ln p, \ln c}(0)$	-0.488	0.102	-0.059***
$\phi_{\ln p, \ln q}(0)$	-0.406	0.103	-0.318
$\phi_{\ln p, \ln E p}(0)$	0.939	0.017	0.853***
$\phi_{\ln p, \psi}(0)$	-0.534	0.118	-0.527
$\phi_{\ln c, \ln q}(0)$	0.395	0.109	0.439
$\phi_{\ln c, \ln E p}(0)$	-0.452	0.106	0.120***
$\phi_{\ln c, \ln \psi}(0)$	0.529	0.116	0.410
$\phi_{\ln q, \ln E p}(0)$	-0.271	0.115	-0.221
$\phi_{\ln q, \psi}(0)$	0.775	0.050	0.870**
$\phi_{\ln E p, \psi}(0)$	-0.500	0.118	-0.427
$\phi_{\ln p, \ln c}(1)$	-0.469	0.125	0.083***
$\phi_{\ln p, \ln q}(1)$	0.104	0.156	-0.132
$\phi_{\ln p, \ln E p}(1)$	0.643	0.069	0.575
$\phi_{\ln p, \psi}(1)$	-0.274	0.142	-0.249
$\phi_{\ln c, \ln p}(1)$	-0.326	0.109	0.020***
$\phi_{\ln c, \ln q}(1)$	0.184	0.110	0.212
$\phi_{\ln c, \ln E p}(1)$	-0.300	0.118	-0.001**
$\phi_{\ln c, \psi}(1)$	0.304	0.127	0.194
$\phi_{\ln q, \ln p}(1)$	-0.257	0.110	0.137***
$\phi_{\ln q, \ln c}(1)$	0.323	0.110	0.222
$\phi_{\ln q, \ln E p}(1)$	-0.212	0.116	0.035**
$\phi_{\ln q, \psi}(1)$	0.067	0.134	-0.178**
$\phi_{\ln E p, \ln p}(1)$	0.566	0.094	0.479
$\phi_{\ln E p, \ln c}(1)$	-0.508	0.116	0.154***
$\phi_{\ln E p, \ln q}(1)$	0.070	0.147	-0.103
$\phi_{\ln E p, \psi}(1)$	-0.358	0.129	-0.230
$\phi_{\ln \psi, \ln p}(1)$	-0.162	0.108	-0.104
$\phi_{\ln \psi, \ln c}(1)$	0.334	0.127	0.079*
$\phi_{\ln \psi, \ln q}(1)$	-0.115	0.122	0.002
$\phi_{\ln \psi, \ln E p}(1)$	-0.203	0.115	-0.178

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.

fitting the consumption-price correlation and confirm likely model misspecification affecting the demand side in particular.

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 about 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 the 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 five 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. $\sigma_{\ln c}$ and $\phi_{\ln p, \ln c}(0)$ are of particular interest because in the previous section we showed that the model struggles to match them; thus, it is helpful to examine which model characteristics is driving their behavior. [Table 11](#) reports the results of this exercise and the results for 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 for this moment are the autocorrelation coefficient of the demand innovations (model 3), the presence of planting-time shocks (model 11), and the smoothing effect of storage (model 16). 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 10% of the price autocorrelation, planting-time shocks account for 6%, and storage accounts for the remaining 51%. So, storage is key to induce price persistence but only explains half of the actual serial correlation, which means that the other model features matter too. This result contrasts with [Deaton and Laroque's](#) estimation results ([1996](#)) 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 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

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

Data or model	$\phi_{\ln p}(1)$	$\sigma_{\ln p}$	$\sigma_{\ln c}$	$\sigma_{\ln q}$	$\phi_{\ln p, \ln c}(0)$
Trending data	0.87	0.46	–	–	–
Detrended data	0.58	0.24	0.019	0.028	–0.49
1. Benchmark	0.50	0.27	0.015	0.032	–0.07
2. $\rho_\mu = 0$	0.41	0.24	0.015	0.031	–0.35
3. $\rho_\mu = 0, \sigma_v = \sigma_\mu$	0.41	0.25	0.017	0.031	–0.20
4. $\alpha_S = 0$	0.54	0.28	0.013	0.026	0.07
5. $g_q = 0$	0.51	0.27	0.015	0.032	–0.07
6. $g_p = 0$	0.54	0.24	0.015	0.033	0.03
7. $k = 0.012$	0.54	0.24	0.015	0.033	0.03
8. $\sigma_\eta = 0$	0.47	0.26	0.014	0.029	0.02
9. $\sigma_\omega = 0$	0.46	0.26	0.014	0.027	0.05
10. $\sigma_\eta = 0, \sigma_\varepsilon = \sigma_\psi$	0.46	0.26	0.015	0.031	–0.03
11. $\sigma_\omega = \sigma_\eta = 0, \sigma_\varepsilon = \sigma_\psi$	0.45	0.26	0.014	0.028	0.04
12. $\rho_\mu = 0, \sigma_v = \sigma_\mu, \alpha_S = 0$	0.38	0.25	0.018	0.028	–0.24
13. $\rho_\mu = 0, \sigma_v = \sigma_\mu, \alpha_S = 0, \sigma_\eta = 0, \sigma_\varepsilon = \sigma_\psi, g_q = 0$	0.32	0.23	0.017	0.028	–0.17
14. $\alpha_D = -0.065$	0.51	0.23	0.015	0.032	–0.26
15. $\alpha_D = -0.065, \rho_\mu = 0.488, \sigma_v = 0.016$	0.48	0.23	0.016	0.032	–0.27
16. $k = \infty$	0.07	0.68	0.026	0.026	–0.79

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.

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 elastic supply (model 4) would only decrease the role of supply shocks and exacerbate the problems related to this moment. Some improvement can be achieved by reducing the autocorrelation of the demand shock (models 2–3) or increasing the storage cost (model 16), but both are detrimental to price autocorrelation fit. The indirect inference approach overestimates ρ_μ by 0.082 and α_D by 0.021 compared with the 2SLS approach. Comparing models 2 and 3 with the benchmark shows that overestimation of ρ_μ would contribute little to solving this puzzle. However, setting the demand elasticity equal to its 2SLS estimate level (model 14) would be enough to bring the simulated moment closer to the observed moment. In other words, the covariance mismatch between consumption and price might be due in part to understatement of the demand elasticity. Imposing that all demand parameters plus the demand elasticity are fixed at their 2SLS estimates (model 15) does not help more to resolve the puzzle.

Price volatility is 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 in our model would be 152% higher (model 16). 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 equates with decreasing storage costs and 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 closest to the model in [Deaton and Laroque \(1992\)](#): it includes neither persistent shocks nor planting-time disturbances, and includes an inelastic supply. In this specification, serial price correlation decreases significantly from 0.5 to 0.32 which suggests that to match true persistence of prices, estimation of a simpler 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. The linear regressions estimates are 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 6.1.1: OLS estimates for supply and 2SLS for demand with 4-knot spline.¹⁷

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 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 to large demand shocks. This applies also to the most recent price spikes (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 ($< -3.5\%$): 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.

¹⁷Figure 3 would nonetheless be very similar if created using instead 2SLS estimates for supply.

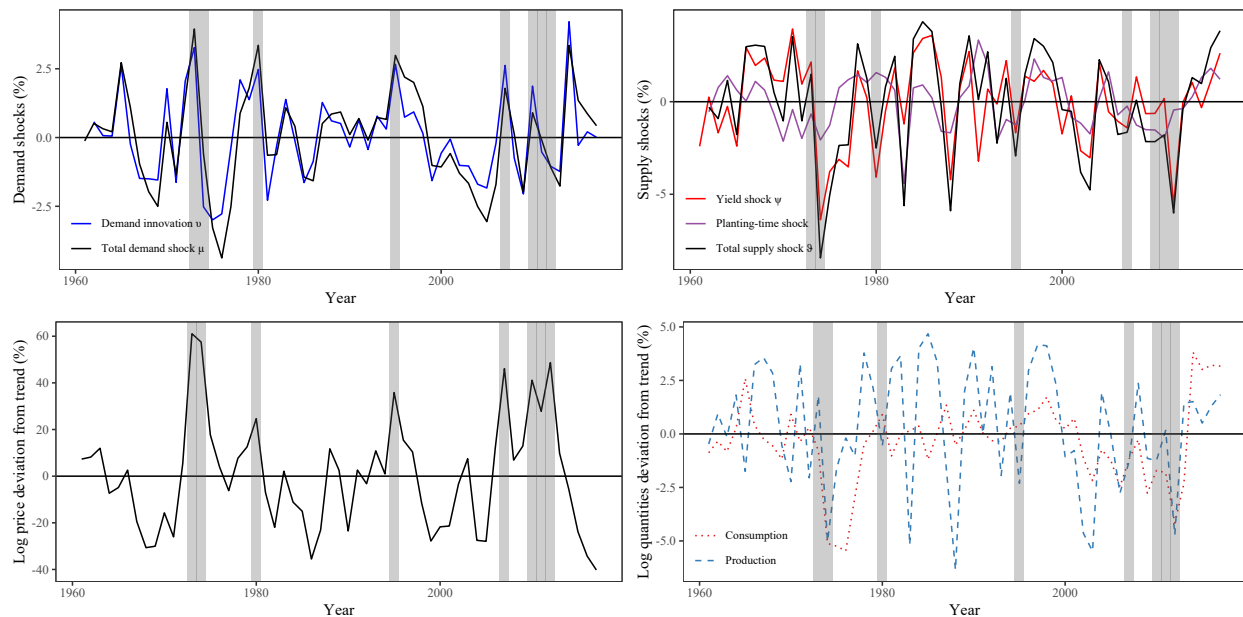


Figure 3: Historical decomposition of the world price, production and consumption of grains into the various shocks. Gray areas denote price spike periods defined as log deviations from the trend greater than one standard deviation, 23.6%. The planting-time supply shock in purple corresponds to $\alpha_S(\eta_{t-1}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

This demonstrates the importance of storage to buffer against supply shortages. In the absence of inventories, a -3.5% supply shock would lead to a 73% price increase because inelastic consumption would have to respond one-to-one to the supply shortfall.

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.¹⁸ 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

¹⁸In 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 \(2020\)](#).

integration constant, as the sum of the following three surpluses:

$$w_t = \underbrace{-\bar{d} \frac{p_t^{1+\alpha_D}}{1+\alpha_D} \bar{p}^{-\alpha_D} e^{\mu_t}}_{\text{Consumer surplus}} + \underbrace{p_t x_{t-1} - (p_t + k\bar{p}) x_t}_{\text{Storer profit}} + \underbrace{p_t h_{t-1} e^{\eta_{t-1} + \varepsilon_t} - \gamma(h_t) e^{\omega_t}}_{\text{Producer profit}}. \quad (56)$$

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 (11) for simplification, we can derive a unit-free expression of instantaneous welfare:

$$\frac{w_t}{\bar{p}\bar{d}} = \underbrace{-\frac{(p_t/\bar{p})^{1+\alpha_D}}{1+\alpha_D} e^{\mu_t} + \frac{p_t s_t - x_t}{\bar{p} \bar{d}}}_{\text{Consumer efficiency gains}} \underbrace{\overbrace{-k \frac{x_t}{\bar{d}}}^{\text{Storage costs}}}_{\text{Production costs}} - \beta \frac{(h_t/\bar{d})^{1+1/\alpha_S}}{1+1/\alpha_S} e^{\omega_t}. \quad (57)$$

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

$$W_t = (1 - \beta) w_t / (\bar{p}\bar{d}) + \beta E_t W_{t+1}. \quad (58)$$

Equation (57) can be evaluated over any state variables using the policy functions defined in section 2.4. Equation (58) 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 0.71% in annual steady-state consumption, the overall welfare effects are modest. However, the distributional effects are large and show a 9.26% increase in consumer surplus and a corresponding decrease in producer profit. This change is related mostly to the change in the mean price: absence of storage reduces the mean price by 8.28% 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 display only the efficiency changes and shows that the gains are shared equally between consumer efficiency gains and production costs.

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.

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
0.71	9.26	-8.55	0.40	-0.14	0.45

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 $k = \infty$.

8 Conclusions

This paper proposes a new empirical strategy to estimate the rational expectations storage model. It features 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, full empirical validation of the storage model has not been carried out. For this first full estimation of the storage model, we chose the empirical setting of the 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 last 60 years. We found that total supply shocks are twice as large as demand shocks, but that all price spikes have been associated to large positive demand shocks.

Our estimations demonstrate also that while our proposed storage model is able to rationalize many of the observed moments, it fails to reproduce the observed negative correlation between price and consumption. 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 two sources of misspecification in our model. One is the deterministic arbitrage relationship assumed for storage which creates stochastic singularity. 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 \(2016\)](#). Another 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](#)).

The present paper follows [Roberts and Schlenker \(2013\)](#) and focuses on the market for grains, but our empirical methodology in principle could be applied as well 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 help link the rational expectations storage literature to the estimation of VAR models 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.

Appendix

A Numerical methods

A.1 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 grid between $\bar{s} - 4\sigma_s$ and $\bar{s} + 5\sigma_s$ which in our experience covers almost all simulated availabilities. A logarithmically spaced grid will position more points in the low availability area where the cutoff of no stock is likely located than would a regular grid. Assuming for availability a normal distribution with parameters \bar{s} and σ_s , we associate probability weights to each grid point based on the segments on which each grid 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 grid 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 $\{\varphi^+, \varepsilon^+, \mu^+\}$ 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 logarithmically spaced grid of availability between $\bar{s} - 4\sigma_s$ and $\bar{s} + 5\sigma_s$ and associate to each grid 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 grid 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 $\{\tilde{s}, \varphi, \mu\}$:

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

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

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

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:

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

Step 3.2. Fixed-point iteration with damping:

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

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

If $\max \left(\|q^{e,m,n} - q^{e,m-1,n}\|_2, \|x^{m,n} - x^{m-1,n}\|_2 \right) < \lambda \varpi$ 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 $\max \left(\|q^{e,m,n} - q^{e,m-1,n}\|_2, \|x^{m,n} - x^{m-1,n}\|_2 \right) \geq \lambda \varpi$ 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.

A.2 Solution precision

Once a solution is obtained, its accuracy can be assessed by rewriting unit-free the equations (12) and (13) 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)

$$EE_i^x = 1 - \frac{d \left(\max \left(d^{-1}(\tilde{s}_i), \beta e^{s_p} \mathbb{E} \mathcal{P} \left([\mathcal{X}(\tilde{s}_i, \varphi_i, \mu_i) e^{-s_q} + \mathcal{Z}(\tilde{s}_i, \varphi_i, \mu_i) e^\varepsilon] e^{-\mu}, \varphi, \mu \right) - k \right) \right) e^{\mu_i}}{\tilde{s}_i e^{\mu_i} - \mathcal{X}(\tilde{s}_i, \varphi_i, \mu_i)}, \quad (\text{A10})$$

$$EE_i^h = 1 - \frac{e^{\varphi_i} \{ \mathbb{E} [\mathcal{P} \left([\mathcal{X}(\tilde{s}_i, \varphi_i, \mu_i) e^{-s_q} + \mathcal{Z}(\tilde{s}_i, \varphi_i, \mu_i) e^\varepsilon] e^{-\mu}, \varphi, \mu \right) e^\varepsilon] \}^{\alpha_s}}{\mathcal{Z}(\tilde{s}_i, \varphi_i, \mu_i)}. \quad (\text{A11})$$

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	-3.83	-1.89
EE^h	-3.71	-2.08

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
k	0	$+\infty$
α_D	$-\infty$	0
α_S	0	$+\infty$

Table A3: Additional MC experiments with instrumental variables for $\sigma_\omega = 20\%$

	ρ_μ	σ_ψ (%)	σ_ν (%)	σ_ϑ (%)	$c_q - 1$	α_D	α_S
<i>T</i> = 100							
Mean	0.47	2.33	1.61	2.82	0.094	-0.066	0.060
St. dev.	0.12	0.18	0.23	0.24	0.056	0.015	0.022
RMSE (%)	25.32	7.54	14.13	8.42	60.356	22.401	38.850
SE	0.12	0.17	0.20		0.060	0.013	0.023
<i>T</i> = 200							
Mean	0.49	2.32	1.61	2.80	0.095	-0.066	0.059
St. dev.	0.09	0.12	0.15	0.16	0.041	0.009	0.015
RMSE (%)	17.41	5.21	9.21	5.82	44.379	14.393	25.177
SE	0.08	0.12	0.14		0.042	0.009	0.015
<i>T</i> = 1000							
Mean	0.50	2.33	1.61	2.80	0.091	-0.065	0.057
St. dev.	0.04	0.05	0.06	0.07	0.017	0.004	0.006
RMSE (%)	7.20	2.22	3.96	2.51	18.570	5.885	11.061
SE	0.04	0.05	0.06		0.018	0.004	0.006

Notes: See notes to table 1.

Table A4: Additional MC 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 (%)	k (%)	α_D	α_S
$T = 100$	OID: 0.028								
Mean	0.48	-0.32	20.54	1.14	1.98	1.57	4.57	-0.063	0.058
St. dev.	0.09	0.23	3.35	0.31	0.21	0.18	2.06	0.011	0.011
RMSE (%)	17.84	77.32	16.97	26.31	10.59	11.11	42.00	16.876	19.245
ASE	0.06	0.40	3.13	0.33	0.20	0.16	1.33	0.010	0.010
$T = 200$	OID: 0.022								
Mean	0.49	-0.31	20.26	1.18	1.99	1.59	4.67	-0.064	0.058
St. dev.	0.06	0.17	2.32	0.20	0.14	0.13	1.52	0.008	0.008
RMSE (%)	11.45	57.47	11.67	16.90	7.14	8.40	31.03	11.914	13.785
ASE	0.04	0.18	2.11	0.21	0.15	0.11	0.94	0.007	0.007
$T = 1000$	OID: 0.012								
Mean	0.50	-0.30	20.09	1.20	1.99	1.59	4.68	-0.064	0.058
St. dev.	0.03	0.08	1.07	0.09	0.06	0.06	0.69	0.004	0.004
RMSE (%)	5.19	25.51	5.39	7.20	3.18	3.86	15.15	5.645	6.468
ASE	0.02	0.07	0.92	0.09	0.07	0.05	0.43	0.003	0.003

Notes: See notes to table 2. For $T = 100$ and $T = 200$, respectively 3 and 1 replications had to be dropped due to non-convergence.

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 (%)	k (%)	α_D	α_S
$T = 56$	$\sigma_\omega = 5\%$	OID: 0.021							
Mean	0.45	-0.31	5.32	1.16	1.96	1.65	4.73	-0.065	0.058
St. dev.	0.12	0.30	1.04	0.35	0.27	0.34	2.74	0.016	0.010
RMSE (%)	27.02	99.64	21.85	29.51	13.53	21.75	55.00	24.478	17.840
ASE	0.09	0.91	1.79	0.46	0.29	0.23	2.29	0.011	0.009
$T = 56$	$\sigma_\omega = 10\%$	OID: 0.019							
Mean	0.45	-0.31	11.14	1.16	1.97	1.64	4.74	-0.065	0.058
St. dev.	0.13	0.29	3.85	0.36	0.27	0.33	2.76	0.016	0.017
RMSE (%)	27.65	95.85	40.10	30.03	13.83	21.06	55.54	24.725	29.299
ASE	0.09	0.80	4.42	0.50	0.29	0.23	2.29	0.011	0.014
$T = 56$	$\sigma_\omega = 20\%$	OID: 0.015							
Mean	0.44	-0.31	23.36	1.15	1.96	1.63	4.90	-0.065	0.059
St. dev.	0.13	0.29	8.61	0.37	0.27	0.33	2.89	0.016	0.026
RMSE (%)	29.61	98.32	46.23	30.93	13.52	20.88	57.74	24.528	44.445
ASE	0.10	0.69	11.56	0.53	0.31	0.24	2.59	0.013	0.022
$T = 100$	$\sigma_\omega = 20\%$	OID: 0.004							
Mean	0.47	-0.30	22.57	1.15	1.98	1.62	4.69	-0.066	0.058
St. dev.	0.10	0.20	7.01	0.30	0.21	0.26	2.00	0.013	0.019
RMSE (%)	21.07	68.07	37.31	25.19	10.76	16.14	40.54	19.361	32.732
ASE	0.07	0.46	8.03	0.36	0.23	0.17	1.74	0.009	0.016
$T = 200$	$\sigma_\omega = 20\%$	OID: 0.004							
Mean	0.49	-0.31	20.96	1.18	2.00	1.62	4.76	-0.065	0.058
St. dev.	0.07	0.15	3.78	0.18	0.14	0.17	1.43	0.008	0.012
RMSE (%)	13.90	51.40	19.48	15.27	7.14	10.83	29.08	12.610	21.389
ASE	0.05	0.23	4.82	0.23	0.16	0.12	1.25	0.006	0.011
$T = 1000$	$\sigma_\omega = 20\%$	OID: 0.004							
Mean	0.50	-0.31	20.36	1.19	2.00	1.60	4.74	-0.065	0.057
St. dev.	0.03	0.07	1.55	0.08	0.06	0.07	0.65	0.003	0.006
RMSE (%)	6.41	23.37	7.98	6.63	3.17	4.44	14.10	5.389	9.629
ASE	0.02	0.08	2.00	0.10	0.07	0.05	0.57	0.003	0.005

Notes: See notes to table 2. For $T = 56$, 24, 17, and 21 replications had to be dropped due to non-convergence, respectively in the cases of $\sigma_\omega = 5\%$, 10%, and 20%. For $T = 100$, 3 replications had to be dropped due to non-convergence.

Table A6: Unit root tests results

Demand Price		Supply Price			Consumption			Production			
<i>Panel A. LM unit root test with quadratic trend and one or two trend breaks</i>											
TB1	TB2	t-stat (lags)	TB1	TB2	t-stat (lags)	TB1	TB2	t-stat (lags)	TB1	TB2	t-stat (lags)
1979		-5.17(1)**	1980		-5.3(1)**	2006		-3.97(1)	2003		-4.88(2)**
1971	2007	-6.41(1)**	1972	2008	-5.68(2)	1984	2007	-6.69(3)**	1982	2000	-6.93(1)**
<i>Panel B. Unit root tests on detrended data (t-stat)</i>											
ADF	PP	KPSS	ADF	PP	KPSS	ADF	PP	KPSS	ADF	PP	KPSS
-3.15***	-3.13***	0.18***	-2.86***	-2.57	0.21***	-1.47	-1.45	0.3***	-2.31***	-3.22***	0.29***
-3.91***	-3.66***	0.05	-3.57***	-3.09***	0.06	-3.3***	-3.27***	0.09	-5.38***	-6.96***	0.04
-4.13***	-3.84***	0.05	-3.67***	-3.18***	0.06	-4.3***	-4.15***	0.04	-5.82***	-7.33***	0.02

Notes: In panel B, each line corresponds to one of the trend specifications considered for detrending the data in logarithm, namely the natural cubic spline with, from top to bottom, 3, 4 and 5 knots. ***, **, and * indicate significance at the 99%, 95%, and 90% levels, respectively. For each test, the lag length has been selected using the common general-to-specific strategy.

Table A7: Supply equation estimation, 1962–2007

	(1)	(2)	(3)
<i>Panel A. 2SLS</i>			
Supply elasticity b_q	0.101*** (0.024)	0.092*** (0.022)	0.088*** (0.022)
Shock c_q	1.120*** (0.144)	1.167*** (0.145)	1.148*** (0.122)
<i>Panel B. First stage</i>			
Lagged shock b_{EP}	-3.922*** (0.392)	-3.588*** (0.598)	-3.625*** (0.511)
Shock c_{EP}	-2.896 (1.840)	-2.294 (1.519)	-2.366 (1.629)
<i>Panel C. OLS</i>			
Supply elasticity b_q	0.110*** (0.016)	0.086*** (0.018)	0.084*** (0.015)
Shock c_q	1.153*** (0.111)	1.150*** (0.116)	1.136*** (0.099)
$\sigma_{u_q}^{2SLS}$	0.018	0.016	0.015
$\sigma_{u_{EP}}$	0.161	0.141	0.144
$\sigma_{u_q}^{OLS}$	0.018	0.016	0.015
σ_{ϑ}^{OLS}	0.033	0.032	0.031
First stage F -stat	15.390	16.495	15.998
p -value for Hausman test	0.703	0.801	0.853
Observations	46	46	46
Spline knots	3	4	5

Notes: Standard errors robust to heteroskedasticity and autocorrelation 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 (0.029)	−0.046 (0.030)	−0.048 (0.030)
Lagged price c_c	0.022 (0.020)	0.012 (0.018)	0.012 (0.019)
Lagged demand d_c	0.895*** (0.204)	0.698*** (0.181)	0.572*** (0.188)
<i>Panel B. First stage</i>			
Shock b_p	−3.900*** (0.884)	−3.720*** (0.992)	−3.791*** (0.958)
Lagged price c_p	0.598*** (0.111)	0.530*** (0.114)	0.545*** (0.109)
Lagged demand d_p	4.943** (2.084)	3.628* (1.981)	4.370** (2.083)
<i>Panel C. OLS</i>			
Demand elasticity b_c	0.000 (0.011)	−0.007 (0.009)	−0.005 (0.007)
Lagged price c_c	−0.003 (0.013)	−0.012 (0.013)	−0.016 (0.013)
Lagged demand d_c	0.743*** (0.122)	0.593*** (0.128)	0.433*** (0.143)
<i>Panel D. 2SLS using Roberts and Schlenker's approach (eqs. (39) for 2nd stage and (43) for 1st)</i>			
Demand elasticity b_c	−0.033* (0.017)	−0.062 (0.041)	−0.059 (0.036)
$\sigma_{u_c}^{2SLS}$	0.015	0.015	0.015
σ_{u_p}	0.174	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}	0.034	0.020	0.018
First stage F -stat (panel A)	11.436	10.106	10.467
p -value for Hausman test (panel A)	0.090	0.051	0.018
First stage F -stat (panel D)	14.405	12.193	12.408
p -value for Hausman test (panel D)	0.011	0.010	0.005
Observations	46	46	46
Spline knots	3	4	5

Notes: See notes to table A7.

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	0.434	(0.109)	0.386	(0.124)
$\rho_{\eta,\omega}$	-0.587	(0.621)	-0.621	(0.770)
σ_ω	0.260	(0.083)	0.255	(0.078)
σ_η	0.010	(0.007)	0.009	(0.008)
σ_ε	0.022	(0.004)	0.022	(0.004)
σ_v	0.016	(0.003)	0.015	(0.003)
k	0.019	(0.013)	0.017	(0.013)
α_D	-0.036	(0.014)	-0.029	(0.013)
α_S	0.060	(0.016)	0.062	(0.016)
σ_φ	0.024	(0.007)	0.023	(0.008)
σ_ψ	0.024	(0.005)	0.024	(0.005)
σ_μ	0.018	(0.004)	0.017	(0.004)
σ_ϑ	0.032	(0.006)	0.032	(0.006)
OID p -value	0.569		0.628	

Notes: See notes to table 8.

References

- 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.
- 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.
- Canova, F. (2014). Bridging DSGE models and the raw data. *Journal of Monetary Economics*, 67, 1–15.
- Chambers, M. J. and Bailey, R. E. (1996). A theory of commodity price fluctuations. *Journal of Political Economy*, 104(5), 924–957.
- 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.
- 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.

- Ghanem, D. and Smith, A. (2020). *Causality in Structural Vector Autoregressions: Science or Sorcery?* Working paper.
- 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.
- 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.
- 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.

- 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.