Submitted to *Quantitative Economics*

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

CHRISTOPHE GOUEL
INRAE, AgroParisTech, Paris-Saclay Applied Economics, Université Paris-Saclay and CEPII

NICOLAS LEGRAND
Department of Agricultural and Applied Economics, Virginia Tech and INRAE, SMART, Institut Agro

We develop an indirect inference approach relying on a linear supply and demand model serving as an auxiliary model to provide the first full empirical test of the rational expectations commodity storage model. We build a rich storage model that incorporates a supply response and four structural shocks and show that exploiting information on both prices and quantities is critical for relaxing previous restrictive identifying assumptions and assessing the empirical consistency of the model's features. Finally, we carry out a structural estimation on the aggregate index of the world's most important staple food products. Our estimations show that supply shocks are the main drivers of food market dynamics and that our storage model is consistent with most of the moments in the data, including the high price persistence so far the subject of a long-standing puzzle.

KEYWORDS. Commodity price dynamics, indirect inference, Monte Carlo analysis, storage.

JEL CLASSIFICATION. C51, C52, Q11.

1. INTRODUCTION

Speculative storage by allowing the transfer of commodities from one period to another and by allowing prices to react immediately to news about future market conditions is a crucial determinant of commodity price dynamics. While this insight is well recognized empirically (see, e.g., Kilian and Murphy, 2014, Letta et al., 2022), the theory underpinning this behavior is far from being empirically validated. Despite being widely used in many applied and policy works (Gouel, 2013, Porteous, 2019, Steinwender, 2018), the framework provided by the rational expectations storage model was rejected by the first estimations of Deaton and Laroque (1992, 1996). Deaton and Laroque found that a simple storage model while able to account qualitatively for many of the stylized facts

Christophe Gouel: christophe.gouel@inrae.fr
Nicolas Legrand: nlegrand90@vt.edu
We are grateful for helpful comments and feedback from three anonymous referees, as well as Alain Carpentier, Dan Cao, Joe Janzen, Delphine Lautier, and John Rust, and participants at various seminars (Georgetown University, Penn State, SMART, University of California Davis, and University of Illinois Urbana-Champaign) and conferences (Commodity & Energy Markets Association meeting, NCCC-134 2021 conference, and RCEA Time Series Workshop). This work was supported by the INRAE ECOSOCIO department.

of commodity price dynamics is not able to match the level of price autocorrelation observed in the data. Subsequent work offered some solutions to raise the persistence induced by the model and better match this central feature of the data (Cafiero et al., 2011, 2015, Gouel and Legrand, 2017, Bobenrieth et al., 2021), but all these studies build on Deaton and Laroque's approach where the model is estimated only on prices, which requires restrictive identifying assumptions, such as assuming only one source of shocks when using likelihood-based techniques, and prevents estimation of all the model parameters. This seriously limits storage models' usefulness in studying how price fluctuations are driven by the underlying shifts in supply and demand, in assessing the respective importance of these supply and demand shocks, and in using this structural model to run counterfactual and welfare analyses for policy purposes.

In this work, we build and estimate a rich rational expectations storage model with the aim of assessing its empirical validity beyond its ability to fit price dynamics. To achieve this, we depart from the standard model setup estimated so far. Specifically, we extend the simple storage model to include: (i) a supply response, (ii) long-run trends in prices and quantities, (iii) a persistent demand shock, and (iv) three supply shocks with different timings. Next, we demonstrate how to leverage the information contained in the joint dynamics of quantities and prices to identify all the structural parameters of the model. Finally, we take our enhanced storage model to five time series representing the global grains market, which we capture through an aggregate index of the world's most important staple food products.[1] We find that our model successfully reproduces the observed high price autocorrelation, with the transfer of inventories over time playing a crucial role in explaining this phenomenon. However, this is only achieved when combined with the other features of the model. We also show that, when fully specified, the storage model can effectively match the key moments of the global food market. Importantly, with our econometric strategy that exploits the joint dynamics of price and quantity, we can empirically assess the overall consistency of the model's combined extensions while identifying more formally which ones help to match the moments in the data.

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

Despite the richness of our model compared to most models in the storage literature, it remains quite stylized, and particularly compared to the number of observables. More

[1] These are maize, rice, soybeans, and wheat.

precisely, with only four shocks driving the fluctuations of five observables, the model presents a stochastic singularity. This is an obstacle to a likelihood-based estimation since, by construction, the model could not be expected to account for the richness of the data. A classical solution to this issue would consist in adding measurement errors to the observables. However, this paper being the first structural estimation of a rich storage model, we prefer to analyze its empirical performance more transparently. Instead, to deal with the stochastic singularity, we adopt an estimation approach that can be applied despite it, which allows us to choose the dimensions of the data to match, and which remains fully transparent with respect to the factors driving the estimation. This approach is the indirect inference proposed in Gourieroux et al. (1993) and Smith (1993). It is a simulated moment-based method in which the model is estimated by targeting parameter estimates from an auxiliary model. Put simply, indirect inference is based on the use of an auxiliary model as a statistical model which provides a rich description of the features in the data. This auxiliary model, which here is the supply and demand model of Roberts and Schlenker (2013), is estimated on both the true data and on simulated data from the structural model, and the structural model parameters are adjusted to minimize the distance between both sets of estimates from the auxiliary model. This approach allows us to exploit an econometric literature where intuitions about which moment is driving a parameter estimation are more explicit than full-information techniques. Also interesting with this approach is that it can be applied in the absence of information about stocks which are generally not available or too noisy to be of use.

We apply this indirect inference approach on the data used by Roberts and Schlenker (2013), which includes five observed variables: price, expected price, demand, production, and yield shock. This allows us to estimate all the parameters of the model. Using these estimates, we present two sets of results. First, we evaluate the ability of our extended storage model to capture the empirical time series properties of both price and quantity data. We assess the performance of the estimated storage model by comparing the covariances based on model simulations and those based on observations. Generally, the covariances are similar for simulations and observations, suggesting that the model is able to mimic the main moments in the data. Interestingly, our results raise a new puzzle: the model proves unable to match the correlations between price and quantities, consumption as well as production, which are much lower in the data than in the model.

Second, a credible solution to the price autocorrelation puzzle can be found by accounting for some key features of the international grains market. Based on our estimations, we can rank the different factors according to their relative contributions to the observed one-year autocorrelation in prices, with storage being the largest contributor, followed by the long-run trend in prices, autocorrelated demand shocks, and supply news shocks.

Our work relates to three strands of research. The first strand studies the theoretical and empirical properties of storage models. Our model builds on earlier studies that introduce similar features separately. For example, Wright and Williams's (1982) competitive storage model includes an elastic supply. Williams and Wright (1991), Chambers and Bailey (1996), Deaton and Laroque (1996), and Routledge et al. (2000) introduce autocorrelated shocks. Several papers (e.g., Lowry et al., 1987, Osborne, 2004, Gouel, 2020) use production shocks with different timings. Dvir and Rogoff (2014) develop a storage model with trending quantities, and Bobenrieth et al. (2021) introduce a supply trend that generates quantity and price trends. Relative to this literature, our use of information on both price and quantities enables us to disentangle the effects of the core storage theory from the set of auxiliary assumptions needed for inference. Indeed, along three dimensions—the persistence of the demand shock, the supply elasticity, and the size and cross-correlation of the supply shocks—the dynamics of quantities play a critical role because price data alone cannot identify any of them.

The second strand is a literature that uses structural vector autoregressions (SVAR) to study commodity markets. This approach, one of the most popular for the empirical analysis of commodity markets, is used, for example, to study the role of supply and demand shocks in commodity markets (Kilian, 2009, Carter et al., 2017, Baumeister and Hamilton, 2019), the role of news shocks (Känzig, 2021), and the role of speculative storage (Kilian and Murphy, 2014, Cross et al., 2022). Compared to this SVAR literature, our paper provides one of the first fully structural approach in the commodity price literature allowing to identify the various shocks in a theoretically consistent way (another paper doing it with a structural model, but without storage and for the oil market, is Bornstein et al., 2023) and to analyze the role of speculative storage.

Last, our approach bridges two literature: the literature on the estimation of storage models and the literature on the estimation of dynamic stochastic general equilibrium (DSGE) models, which conceptually and numerically are close to storage models. The estimation of storage models has been so far restricted to small models too stylized to capture the richness of these markets. This was also the case for DSGE models up to the contributions of Smets and Wouters (2003, 2007), who show how to build and estimate DSGE model with rich stochastic structures. We follow Smets and Wouters by adding a rich set of structural shocks to a storage model. Our estimation approach also borrows from the DSGE literature where indirect inference is commonly applied.[2] In this literature, the auxiliary model is often a SVAR and the estimations depend on targeting the impulse responses (e.g., Rotemberg and Woodford, 1997, Christiano et al., 2005, Ruge-Murcia, 2020).[3] In our case, we show that a system of linear equations based on the instrumental variable model in Roberts and Schlenker (2013) is enough to capture the dynamic relationships of interest (as in Guvenen and Smith, 2014), including the strong nonlinearities.[4] However, a SVAR should also work since Carter et al. (2017) use this framework to approximate a storage model and Ghanem and Smith (2022) adapted

[2]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. However, as Michaelides and Ng (2000) followed Deaton and Laroque by estimating their model only on prices, the various auxiliary models they consider are all based on univariate time-series models.

[3]Indirect inference is also used in macroeconomics for a purpose other than estimating models: testing them (e.g., see Le et al., 2011). In this case, indirect inference consists in using a Wald statistic to test an estimated DSGE model against a reference statistical model, typically a VAR.

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

in a SVAR Roberts and Schlenker's IV model, which provides the basis for our auxiliary 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. Section 4 describes the data and gives descriptive statistics. Section 5 discusses the estimation results, assesses the model fit on moments not included in the estimation, and analyzes the role of storage in price dynamics. Section 6 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 been 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 4.2). We allow equilibrium quantity and price to have trends to account for these long-run dynamics. We work with trends in equilibrium variables because these trends can be directly estimated from the data, contrary to structural trends such as demand and cost trends, which are not observable.[5] Since, for simulation purposes, the storage model must be a stationary model, 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*

2.1.0.1 *Producers* A representative producer makes its production decision and pays for inputs one period before bringing its output to the market. The production choice represented by the acreage is made in period t and denoted H_t. The producer decision is affected by two shocks: η_t, a planting-time yield shock, and ω_t, a cost shock. The planting-time yield shock represents the component of yield shock that is observable by the producer when planting, for example, related to the field conditions during planting, the groundwater level, and the seasonal weather forecasts. Roberts and Schlenker (2013) also take the example of the soybean rust, which is observable from the previous growing season. The cost shock is also observable by the producer and aggregates a variety of 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 a correlation between η_t and ϵ_{t+1} in the model. If producers are efficient forecasters (in the sense

[5]See Appendix A for the same model developed with structural trends; this Appendix links equilibrium and structural trends.

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 also associated 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

$$\max_{H_t \geq 0} \beta \, \mathrm{E}_t \left(P_{t+1} H_t \, \mathrm{e}^{\eta_t + \epsilon_{t+1}} \right) - \Gamma_t \left(H_t \right) \mathrm{e}^{\omega_t + g_p t}, \tag{1}$$

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. The solution to this problem is given by the following first-order condition

$$\beta \, \mathrm{e}^{\eta_t} \, \mathrm{E}_t \left(P_{t+1} \, \mathrm{e}^{\epsilon_{t+1}} \right) = {\Gamma_t}' \left(H_t \right) \mathrm{e}^{\omega_t + g_p t} . \tag{2}$$

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

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

2.1.0.2 *Storers* For the storage sector, we assume free entry, competitive behavior, and risk-neutrality. To store an amount $X_t \geq 0$ from period t until $t+1$ competitive storers incur several costs. They incur an opportunity cost because they have to buy one period before being able to sell. Following most of the storage literature (Gustafson, 1958, Steinwender, 2018, Wright and Williams, 1982, 1984), we assume that storers incur a physical cost of storage proportional to the stored quantity, $k\bar{P}_t X_t$, where $\bar{P}_t$ is the price on the deterministic growth path (i.e., in the absence of shocks) and $k \geq 0$ is the per-unit physical storage cost expressed as a percentage of this price. To be compatible with a model that ultimately could be expressed in terms of stationary variables, the per-unit storage cost must be assumed either to be null (the assumption adopted in Bobenrieth et al., 2021) or as adopted here to follow the same trend as the price. We assume no deterioration of stored grains, and the working paper version of this article shows that it is impossible to estimate separately a per-unit storage cost and a rate of deterioration, because the same moments identify the two parameters (Gouel and Legrand, 2022).[6]

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

$$\max_{X_t \geq 0} \mathrm{E}_t \left[\left(\beta P_{t+1} - P_t - k\bar{P}_t \right) X_t \right]. \tag{3}$$

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

$$\beta \, \mathrm{E}_t \, P_{t+1} - P_t - k\bar{P}_t \leq 0, \ \ = 0 \text{ if } X_t > 0. \tag{4}$$

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

Total marginal storage costs equal $k\bar{P}_t - (\beta - 1) \, \mathrm{E}_t \, P_{t+1}$, which shows that a key difference between per-unit storage cost and the opportunity cost lies in the fact that the opportunity cost is a storage cost that rise with the (expected) price level.

2.1.0.3 *Final demand* Non-speculative demand for commodities can be affected by a variety of shocks: income, policy (e.g., public support for biofuels), and preference shocks (see e.g. Carter et al., 2011, Chen et al., 2010, Gilbert, 1989). For parsimony, we gather these different shocks in one demand shock μ_t, and since such shocks are likely to be persistent, we assume μ_t to be autocorrelated. 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 $\upsilon \sim \mathcal{N}\left(0, \sigma_\upsilon^2\right)$:

$$\mu_{t+1} = \rho_\mu \mu_t + \upsilon_{t+1}. \tag{5}$$

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

$$X_{t-1} + H_{t-1} \, \mathrm{e}^{\eta_{t-1} + \epsilon_t} = D_t \left(P_t \right) \mathrm{e}^{\mu_t} + X_t. \tag{6}$$

[6]We do not consider the possibility of an upper bound on storage capacities (Oglend and Kleppe, 2017) because, contrary to oil and gas, grains can be stored outside dedicated facilities. In addition, we follow the tradition of Wright and Williams and Deaton and Laroque by assuming away any negative (nonlinear) storage cost related to the concept of "convenience yield". The latter refers to the value of having stocks close at hand despite a seeming loss i.e., at a spread between expected and current prices below the full carrying costs (Kaldor, 1939, Working, 1949, Brennan, 1958). Assuming away convenience yield has the benefit of parsimony as only one parameter k can represent storage costs. The limit is that expected prices can only fall with respect to current prices in the absence of inventories (see Williams, 1986, for a full analysis of the concept of convenience yield).

2.2 *Stationary model*

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

$$P_{t+1} = p_{t+1}\,\mathrm{e}^{g_p t}, \tag{7}$$

$$D_t\left(P_t\right) = \mathrm{e}^{g_q t}\, d\left(p_t\right), \tag{8}$$

$$\Gamma_t{}'\left(H_t\right) = \gamma'\left(h_t\right), \tag{9}$$

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

For reasons of market equilibrium, all equilibrium quantities—final consumption, production, and stocks—must share the same multiplicative trend g_q for equation (6) to be made stationary. Note that we are working here with equilibrium trends, trends in equilibrium quantities and prices. The fact that in equilibrium, all quantities share the same trend does not preclude that demand and production costs may be exposed to different structural trends (see Appendix A). The equilibrium quantity trend results from the combined effects of demand and production cost trends that are not modeled here (see equation (S.8)). Similarly, the price trend emerges when demand and cost trends do not perfectly offset each other (see equation (S.6)).

Defining detrended stocks and acreage using

$$X_t = x_t \exp\left(g_q t\right) \text{ and } H_{t-1} = h_{t-1} \exp\left(g_q t\right), \tag{10}$$

we can replace the trending quantities by their detrended counterparts in the above market clearing equation (6):

$$x_{t-1}\,\mathrm{e}^{-g_q} + h_{t-1}\,\mathrm{e}^{\eta_{t-1}+\epsilon_t} = d\left(p_t\right)\mathrm{e}^{\mu_t} + x_t. \tag{11}$$

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

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

$$\beta\,\mathrm{e}^{g_p}\,\mathrm{E}_t\, p_{t+1} - p_t - k\bar{p} \le 0, \;\; = 0 \text{ if } x_t > 0. \tag{12}$$

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 4—raises the opportunity cost because, since prices tend to decrease over time, a higher expected price is required to maintain the same level of stocks. Associated with the price trend, the condition $g_p < -\log\beta$ ensures that inventories are costly and is a necessary condition for the existence of a stationary rational expectations equilibrium,[7] which is always satisfied for decreasing trends.[8]

In equation (11), 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

$$s_t \equiv x_{t-1}\,\mathrm{e}^{-g_q} + h_{t-1}\,\mathrm{e}^{\eta_{t-1}+\epsilon_t}. \tag{13}$$

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

$$h_t : \beta\,\mathrm{e}^{\eta_t-\omega_t}\,\mathrm{E}_t\left(p_{t+1}\,\mathrm{e}^{\epsilon_{t+1}}\right) = \gamma'\left(h_t\right), \tag{14}$$

$$x_t : \beta\,\mathrm{e}^{g_p}\,\mathrm{E}_t\,p_{t+1} - p_t - k\bar{p} \leq 0, \;\; = 0 \text{ if } x_t > 0, \tag{15}$$

$$p_t : s_t = x_t + d\left(p_t\right)\mathrm{e}^{\mu_t}. \tag{16}$$

It can be seen that, in the stationary model, while the price trend is equivalent to a change in the opportunity cost of storage, the quantity trend does not directly affect the incentives. However, it affects them indirectly through its scaling of past stocks in equation (13). 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*

For consistency with Roberts and Schlenker's framework and for simplicity, we assume that the stationary demand function takes an isoelastic form such that

$$d\left(p_t\right) = \bar{d}\left(\frac{p_t}{\bar{p}}\right)^{\alpha_D}, \tag{17}$$

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

$$\gamma'\left(h_t\right) = \beta\bar{p}\left(\frac{h_t}{\bar{d}}\right)^{1/\alpha_S}, \tag{18}$$

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

[7] It corresponds to the assumption 2 of Deaton and Laroque (1992).
[8] See Bobenrieth et al. (2021) for a thorough treatment of the behavior of a storage model with price trend.

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

$$\frac{s_t}{\bar{d}} = \frac{x_{t-1}}{\bar{d}} \mathrm{e}^{-g_q} + \frac{h_{t-1}}{\bar{d}} \mathrm{e}^{\eta_{t-1}+\epsilon_t}, \tag{19}$$

$$\frac{h_t}{\bar{d}} = \left[\mathrm{e}^{\eta_t-\omega_t} \mathrm{E}_t \left(\frac{p_{t+1}}{\bar{p}} \mathrm{e}^{\epsilon_{t+1}}\right)\right]^{\alpha_S}, \tag{20}$$

$$\beta \mathrm{e}^{g_p} \mathrm{E}_t \left(\frac{p_{t+1}}{\bar{p}}\right) - \frac{p_t}{\bar{p}} - k \leq 0, \quad = 0 \text{ if } \frac{x_t}{\bar{d}} > 0, \tag{21}$$

$$\frac{s_t}{\bar{d}} = \frac{x_t}{\bar{d}} + \left(\frac{p_t}{\bar{p}}\right)^{\alpha_D} \mathrm{e}^{\mu_t}. \tag{22}$$

From these equations, we see that the only effect of the deterministic steady-state quantity ($\bar{d}$) and price ($\bar{p}$) is that they scale the value of the variables.

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

2.4 *Model solution*

Equations (5) and (19)–(22) represent a nonlinear rational expectations system based on the exogenous state variable μ_t, the endogenous state variable s_t, and the response variables h_t, x_t, and p_t driven by the innovations $\{\eta_t, \omega_t, \epsilon_t, \upsilon_t\}$. 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 space can be employed. Given that for the numerical solution we use a projection method, it is important for speed and precision to reduce if possible 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 $q^e_{t+1} = \mathrm{E}_t\, q_{t+1} \exp(-\sigma^2_\epsilon/2) = h_t \exp(\eta_t)$, which is the expected production corrected for the mean harvest-time shock and which is given by

$$q^e_{t+1} = \bar{d}\, \mathrm{e}^{\eta_t} \left[\mathrm{e}^{\eta_t-\omega_t} \mathrm{E}_t \left(\frac{p_{t+1}}{\bar{p}} \mathrm{e}^{\epsilon_{t+1}}\right)\right]^{\alpha_S}. \tag{23}$$

In this case, the transition equation is defined as

$$s_{t+1} = x_t \mathrm{e}^{-g_q} + q^e_{t+1} \mathrm{e}^{\epsilon_{t+1}}. \tag{24}$$

We combine the two planting-time shocks that appear in equation (23) to form the aggregate planting-time shock $\varphi_t \equiv (1+\alpha_S)\eta_t - \alpha_S \omega_t$. φ_t summarizes the effective planting-time shocks, is observable by the producer, and allows for a further simplification of the supply equation:

$$q^e_{t+1} = \bar{d}\, \mathrm{e}^{\varphi_t} \left[\mathrm{E}_t \left(\frac{p_{t+1}}{\bar{p}} \mathrm{e}^{\epsilon_{t+1}}\right)\right]^{\alpha_S}. \tag{25}$$

We can see also that in the absence of demand for stock, the market clearing equation (16) collapses to $s_t = d(p_t)\exp(\mu_t)$. 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, $\tilde{s}_t \equiv s_t \exp(-\mu_t)$, i.e., availability in the market corrected for the demand shock.

From the above, we see that it is possible to reduce the number of state variables to 3 by replacing η_t and ω_t by φ_t. We also substitute the availability by the net availability, therefore we define the policy functions on the set of state variables $\{\tilde{s}_t, \varphi_t, \mu_t\}$:

$$q^e_{t+1}/\bar{d} = \mathcal{Q}\left(\tilde{s}_t, \varphi_t, \mu_t\right), \tag{26}$$

$$x_t/\bar{d} = \mathcal{X}\left(\tilde{s}_t, \varphi_t, \mu_t\right), \tag{27}$$

$$p_t/\bar{p} = \mathcal{P}\left(\tilde{s}_t, \varphi_t, \mu_t\right). \tag{28}$$

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:

$$\begin{aligned}\mathcal{P}\left(\tilde{s}_t, \varphi_t, \mu_t\right) = \max\Bigg\{ & \left(\tilde{s}_t/\bar{d}\right)^{1/\alpha_D}, \\ \beta\, \mathrm{e}^{g_p}\, \mathrm{E}_t\left[\mathcal{P}\left(\left[\mathcal{X}\left(\tilde{s}_t, \varphi_t, \mu_t\right)\mathrm{e}^{-g_q} + \mathcal{Q}\left(\tilde{s}_t, \varphi_t, \mu_t\right)\mathrm{e}^{\epsilon_{t+1}}\right]\mathrm{e}^{-\mu_{t+1}}, \varphi_{t+1}, \mu_{t+1}\right)\right] - k, & \Bigg\},\end{aligned} \tag{29}$$

$$\mathrm{e}^{\varphi_t}\left\{\mathrm{E}_t\left[\mathcal{P}\left(\left[\mathcal{X}\left(\tilde{s}_t, \varphi_t, \mu_t\right)\mathrm{e}^{-g_q} + \mathcal{Q}\left(\tilde{s}_t, \varphi_t, \mu_t\right)\mathrm{e}^{\epsilon_{t+1}}\right]\mathrm{e}^{-\mu_{t+1}}, \varphi_{t+1}, \mu_{t+1}\right)\mathrm{e}^{\epsilon_{t+1}}\right]\right\}^{\alpha_S} = \mathcal{Q}\left(\tilde{s}_t, \varphi_t, \mu_t\right). \tag{30}$$

Equation (29) 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. In other words, unlike in a model where there is a single i.i.d. shock driving all the commodity price fluctuations, the threshold price above which there is no storage is no longer constant and depends on the demand and planting-time supply shocks.

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 B, we propose a new solution method to the storage model based on recent developments in the literature (Maliar and Maliar, 2014) which satisfies this tradeoff.

3. Econometric procedure

Not all of the storage model variables are observable. For example, stock levels are available from the United States Department of Agriculture (USDA) statistics but for many countries they are based on USDA estimates in the absence of official statistics, and so are likely to be affected by measurement errors.[9] In this paper, we use the five observable variables proposed by Roberts and Schlenker (2013): price, expected price, consumption, production, and yield shock: $[p_t, \mathrm{E}_t\, p_{t+1}, c_t, q_t, \psi_t]$. The consumption variable will be built using information about stock variations. While stock variations can be affected by measurement errors, those are less important than for stock levels in which errors come from the estimate of initial stock levels plus the accumulation of errors in past stock variations.

Our storage model includes fourteen parameters, nine of which are estimated in combination and gathered in the row-vector $\theta \in \Theta$. 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-Martín (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 4.2).

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_υ) 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. The small- and large-sample properties of these two estimation strategies are studied with a Monte Carlo analysis in Appendix C.

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.

[9]The 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.

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 (20) is

$$\log q_t = \log\left(h_{t-1}\,\mathrm{e}^{\psi_t}\right) = \log\left(\bar{d}/\bar{p}^{\alpha_S}\right) + \alpha_S\left(\eta_{t-1} - \omega_{t-1}\right) + \alpha_S \log\left(\mathrm{E}_{t-1}\left(p_t\,\mathrm{e}^{\epsilon_t}\right)\right) + \psi_t. \quad (31)$$

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

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

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}. Under the model assumptions, lagged yield shocks are correlated with the expected prices because storage implies that past yield shocks have contemporaneous effects on prices through the availability in the market, and they are not correlated with the planting-time shocks and thus with the residuals. The first-stage equation is

$$\log\left(\mathrm{E}_{t-1}\, p_t\right) = a_{\mathrm{E}\,p} + b_{\mathrm{E}\,p}\psi_{t-1} + c_{\mathrm{E}\,p}\psi_t + u_{\mathrm{E}\,p,t}. \quad (33)$$

This supply-side estimation strategy deserves a few comments. First, substituting the expected price $\mathrm{E}_{t-1}\, p_t$ for the producer incentive price $\mathrm{E}_{t-1}[p_t \exp(\epsilon_t)]$ creates 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. The bias is expected to be negative.[10] The size of this bias can only be assessed using the policy function of the storage model once all parameter values are known. For the parameter values estimated in this paper, the bias is an order of magnitude smaller than the standard errors.[11] The Monte Carlo analysis in the Appendix 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 $b_q \log(\mathrm{E}_{t-1}\, p_t) = \alpha_S \log(\mathrm{E}_{t-1}[p_t \exp(\epsilon_t)])$, we can write

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

[10] $\log(\mathrm{E}_{t-1}[p_t \exp(\epsilon_t)])$ can be decomposed in $\log(\mathrm{E}_{t-1}\, p_t) + \log(1 + \mathrm{cov}_{t-1}(p_t, \exp(\epsilon_t))/\, \mathrm{E}_{t-1}\, p_t)$ and since a positive harvest-time yield shock should be associated with lower prices and conversely, we can expect $\mathrm{cov}_{t-1}(p_t, \exp(\epsilon_t)) < 0$ and $\log(\mathrm{E}_{t-1}[p_t \exp(\epsilon_t)]) < \log(\mathrm{E}_{t-1}\, p_t)$, which leads b_q to be a downward-biased estimate of supply elasticity. See Gouel (2020, Appendix) for a more detailed analysis of this bias.

[11] One cannot exclude larger biases for other combinations of parameters.

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

$$c_q = \frac{\operatorname{cov}\left(\log q_t - \alpha_S \log\left(\mathrm{E}_{t-1} p_t\right), \psi_t\right)}{\operatorname{var} \psi_t} \tag{35}$$

$$= \frac{\operatorname{cov}\left(\alpha_S \left(\eta_{t-1} - \omega_{t-1}\right) + \psi_t, \psi_t\right)}{\operatorname{var} \psi_t} \tag{36}$$

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

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

Similarly, the residuals can be used to obtain a measure of the total supply shock, which we denote ϑ. As for c_q, we can reorganize equation (32) to get

$$\log q_t - \log\left(\bar{d}/\bar{p}^{\alpha_S}\right) - b_q \log\left(\mathrm{E}_{t-1} p_t\right) = c_q \psi_t + u_{q,t} = \alpha_S \left(\eta_{t-1} - \omega_{t-1}\right) + \psi_t = \vartheta_t. \tag{38}$$

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, the occurrence of stockouts could raise concerns about the instrument's strength. In the complete absence of stocks, prices are not intertemporally linked, and past yield shocks do not influence current prices, so they are not a relevant instrument. If stockouts are occasional, the model alternates between the presence of stocks in which lagged yields are correlated with expected prices and the absence of stocks in which there is possibly no correlation. Thus, the more frequent the stockouts, the weaker the instrument. Whether lagged yields are an instrument strong enough in this setup is an empirical issue discussed in section 5 and the Monte Carlo analysis in the Appendix. Hendricks et al. (2015) raise a related issue by noting 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 useful for estimating supply elasticity.

3.1.2 *Consumption* From equation (17), logged consumption (denoted c_t), is given by

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

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

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

The fact that υ_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

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

with the associated first stage

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

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

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

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

Since equation (41) includes a lagged dependent variable, a condition for d_c to be consistently estimated is the absence of serial correlation in the residuals, which will be tested using the test proposed by Cumby and Huizinga (1992) which is valid for models that have endogenous regressors. Even in the absence of serial correlation in the residuals, standard estimators of autoregressive models are biased in finite sample. We correct for the finite sample bias using Orcutt and Winokur's (1969) formula: $\hat{\rho}_\mu = (1 + T\hat{d}_c)/(T-3)$, where T is the sample length.

3.2 *Indirect inference approach*

Indirect inference requires the 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 estimated by OLS:[12]

$$\log q_t = a_q + b_q \log\left(\mathrm{E}_{t-1}\, p_t\right) + c_q \psi_t + u_{q,t}, \tag{44}$$

$$\log\left(\mathrm{E}_{t-1}\, p_t\right) = a_{\mathrm{E}\,p} + b_{\mathrm{E}\,p}\psi_{t-1} + c_{\mathrm{E}\,p}\psi_t + u_{\mathrm{E}\,p,t}, \tag{45}$$

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

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

$$\psi_t = a_\psi + u_{\psi,t}. \tag{48}$$

A fundamental assumption of the indirect inference method is that the parameters from the auxiliary model bring information about all structural parameters.[13] This assumption

[12]Note that instead of this auxiliary model, we could have used one in which the supply and demand equations are estimated by 2SLS, as laid out in the previous section. This approach leads to similar results but with larger standard errors caused by the loss of precision related to the instrumentation. See the working paper version of this article for details (Gouel and Legrand, 2022).

[13]Formally, this assumption requires that the derivative of the binding function which links the structural to the auxiliary parameters ($\partial\zeta(\theta)/\partial\theta$) is full-column rank (Assumption A4 in Gourieroux et al., 1993).

explains why, in addition to the supply and demand equations, we have included the first-stage equations and equation (48): by having one equation associated with each observable, we ensure sufficient moments in the auxiliary model.

Using the selected auxiliary model, we can define the objective using a subset of the model parameters, which excludes the intercepts since these are only informative about the steady-state values which we normalize to unity: $\zeta = [b_q, c_q, \sigma_{u_q}, b_{\mathrm{E}\,p}, c_{\mathrm{E}\,p}, \sigma_{u_{\mathrm{E}\,p}}, b_c, c_c, d_c, \sigma_{u_c}, b_p, c_p, d_p, \sigma_{u_p}, \sigma_{u_\psi}]$. Note that this auxiliary model is almost equivalent to a subset of the parameters of a VAR(1) model.[14]

This auxiliary model has two important benefits. First, since it involves only linear regressions, it is trivial to estimate and avoids burdening the indirect inference procedure with a computationally costly auxiliary model.[15] Second, it is transparent regarding the relationships between the auxiliary model and the storage model parameters. For instance, b_q is asymptotically equal to the supply elasticity plus the omitted variable bias, which is a function of the size of the cost shock σ_ω. From equation (37), c_q and similarly σ_{u_q} are both nonlinear combinations of α_S, σ_ϵ, σ_η, σ_ω, and $\rho_{\epsilon,\omega}$. From Hendricks et al. (2015), $c_{\mathrm{E}\,p}$ relates 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_υ, themselves informed by d_c and σ_{u_c}. In equation (47), c_p connects to the first-order autocorrelation of $\log p$, which depends directly on the storage cost k, conditional on the other parameters. Intuitively, lower storage costs imply more storage and, in turn, a higher price autocorrelation and vice versa (Gouel and Legrand, 2017, Figure 2). In equation (48), we get σ_{u_ψ}, which equals $\sqrt{\sigma_\epsilon^2 + \sigma_\eta^2}$ and ensures that the model aims to fit the standard deviation of yields, an aggregate shock we observe. Finally, the inclusion of $\sigma_{u_{\mathrm{E}\,p}}$ and σ_{u_p} is almost equivalent to including the standard deviations of the price and the expected price in the objective. It ensures that the estimated model will also aim to fit these targets.

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

$$\hat{\theta} = \operatorname*{arg\,min}_{\theta \in \Theta} \left[\hat{\zeta}_T - \frac{1}{\tau} \sum_{i=1}^{\tau} \hat{\zeta}_T^i(\theta) \right]' W \left[\hat{\zeta}_T - \frac{1}{\tau} \sum_{i=1}^{\tau} \hat{\zeta}_T^i(\theta) \right], \tag{49}$$

where W is a 15×15 symmetric nonnegative definite weighting matrix. This estimator minimizes the weighted distance between the auxiliary model parameters estimated

[14]We thank the editor for this insight.

[15]See Li (2010) and Guvenen and Smith (2014) for other papers that rely on linear equations estimated by OLS as the auxiliary model in an indirect inference setting. In addition, based on a Monte Carlo analysis of various estimators of the storage model, Michaelides and Ng (2000, Section 6) advise against using nonlinear auxiliary models because of the risk of non-convergence and the increased computation time.

using actual data and those estimated using data simulated from our structural storage model.

For the weighting matrix used in the objective W, we use a diagonal matrix with elements that are the inverse of the variance of the estimate of ζ_T. We calculate this using the formulas for standard errors robust to heteroskedasticity for the standard regression parameters ($b_q, c_q, b_{\mathrm{E}p}, c_{\mathrm{E}p}, b_c, c_c, d_c, b_p, c_p, d_p$), and using the following formulas for the standard deviations ($\sigma_{u_q}, \sigma_{u_{\mathrm{E}p}}, \sigma_{u_c}, \sigma_{u_p}, \sigma_{u_\psi}$):

$$\mathrm{var}\left(\sigma\right) = \frac{\sigma^2}{2\left(T-l\right)}, \tag{50}$$

where $T-l$ is the degree of freedom of the corresponding regression. Classical minimum distance estimators, of which the indirect inference is part, suffer from small-sample bias when the optimal weighting matrix calculated on observations is used (Altonji and Segal, 1996). As a result, a common simplification in the indirect inference literature is to use a diagonal of the optimal weighting matrix calculated on observations (see, e.g., Christiano et al., 2005, Ruge-Murcia, 2020). While this choice of a diagonal matrix limits the small-sample bias, it also leads to neglecting interactions between parameters. Still, this issue remains limited in our setting where the number of target parameters does not exceed by much the number of structural parameters. We further explore the sensitivity to the choice of the weighting matrix in Appendix D.2.2.

As for the optimization routine, at every step of the minimization a new set of parameters θ is proposed. For this new θ, a numerical solution of the storage model is computed using the algorithm proposed in Appendix B.1. The resulting policy functions are then used to simulate the model starting from the deterministic steady state and using random shocks drawn at the beginning of the estimation procedure and kept fixed throughout.

In line with Gourieroux et al. (1993), the variance-covariance matrix for the parameter estimates converges asymptotically to

$$\left(1+\frac{1}{\tau}\right)\left(J'WJ\right)^{-1}J'W\Omega^{-1}WJ\left(J'WJ\right)^{-1}, \tag{51}$$

where $J = \left(1/\tau\right)\sum_{i=1}^{\tau}\mathrm{E}[\partial\hat{\zeta}_T^i\left(\theta\right)/\partial\theta]$ is a 15×9 full rank matrix, evaluated by central difference at $\theta=\hat{\theta}$, and Ω is the optimal weighting matrix. We estimate the optimal weighting matrix at the solution by Monte Carlo simulations using the following formula (Le et al., 2011)

$$\Omega^{-1} = \frac{1}{K}\sum_{i=1}^{K}\left[\hat{\zeta}_T^i\left(\hat{\theta}\right) - \frac{1}{K}\sum_{j=1}^{K}\hat{\zeta}_T^j\left(\hat{\theta}\right)\right]\left[\hat{\zeta}_T^i\left(\hat{\theta}\right) - \frac{1}{K}\sum_{j=1}^{K}\hat{\zeta}_T^j\left(\hat{\theta}\right)\right]', \tag{52}$$

where $K=1,000$ is the number of samples of size $T+1+t^{\mathrm{burn}}$ we simulate to calculate this matrix. Calculating the optimal weighting matrix this way has two benefits. First, with K large enough, it removes the short-sample bias inherent to the estimation on observations.

Second, as shown by Le et al. (2011, 2016), in presence of misspecifications, this model-restricted variance-covariance matrix has more power to reject the null hypothesis that the model is correct than the data-based variance-covariance matrix.[16]

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 exploring parameter values outside their domain of definition and those that would make solving the model difficult (see table S.11). Furthermore, to limit the risk of finding only a local optimum, the optimization algorithm starts from 500 different initial values of θ, except for the Monte Carlo experiments in the Appendix, which uses a unique starting point. Finally, although it is costly to solve for the rational expectations equilibrium of the model, it is less costly to simulate from it. Therefore, we choose $\tau = 200$ to minimize the simulation-related uncertainty in the estimates.

4. 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.

4.1 *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):

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

[16]We could also have used this optimal weighting matrix to test the overidentification restrictions by building a Wald statistic as proposed in Gourieroux et al. (1993, Section 4). However, our Monte Carlo experiments point to important size distortions for the associated chi-square test, making it unreliable (this is a frequent problem in this type of structural approach, see, e.g., Michaelides and Ng, 2000, Ruge-Murcia, 2007, 2012).

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

$$\rho_{lit} = \frac{A_{lit}\kappa_l \hat{Y}_{lit}}{\sum_{l'}\sum_{i'} A_{l'i't}\kappa_{l'}\hat{Y}_{l'i't}} \tag{54}$$

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, which can induce spurious correlations due to mixing different marketing seasons (Guerra et al., 2015). The model includes two prices: the current price P_t, which is the price received by the farmers at harvest time and paid by consumers, and the expected price $\mathrm{E}_{t-1} P_t$, which corresponds to the farmers' rational expectations at planting time about the price P_t they will receive at harvest time. Since Gardner (1976), it is common to use futures prices in place of the unobservable expected price. This is a valid approach if futures prices are unbiased predictor of spot prices, which is not true for all commodities but is true for the commodity prices studied here according to Chinn and Coibion (2014).[17] Given the annual time-frame of the model, we take futures contracts with a one-year horizon. For consistency, P_t is the corresponding futures contract at delivery. Following Roberts and Schlenker (2013), we use prices from the Chicago Board of Trade futures for the main month following each crop harvest (i.e., December for maize and wheat, November for rice and soybeans).[18] Monthly prices are obtained by averaging the daily prices observed during each month. Futures prices for rice started trading in

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

[18]At 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.

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):

$$P_t = \frac{\sum_{l\neq \text{rice}} \left(\sum_i \rho_{lit}\right) P_{lt|t}/\kappa_l}{\sum_{l\neq \text{rice}} \sum_i \rho_{lit}} \text{ and } \mathrm{E}_{t-1} P_t = \frac{\sum_{l\neq \text{rice}} \left(\sum_i \rho_{lit}\right) P_{lt|t-1}/\kappa_l}{\sum_{l\neq \text{rice}} \sum_i \rho_{lit}}, \tag{55}$$

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

4.2 *Nonstationarity*

The storage model in sections 2 and 3 assumed that the logs of price and quantity are characterized by linear time trends, and provided theoretical predictions for deviations of the variables from this linear trend. As we will show in this section, the actual trends in the data are more complicated. Our empirical strategy is to compare the model's predictions for deviations from a linear time trend with the observed deviations from a cubic spline. This section discusses our reasons for doing this and the possible ramifications.

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 an extensive 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 takeaway 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.[19] We thus 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). These LM tests allow for one or two structural breaks with or without a linear or quadratic deterministic trend under both the null and alternative hypotheses.

Although our econometric models call for variables in logarithms, the literature widely acknowledges that unit root tests are highly sensitive to data transformation, which is also likely to transform the underlying trends (Corradi and Swanson, 2006). For example, while quantities in levels exhibit an almost linear trend up to the mid-2000s, this is less evident in logarithms. We, therefore, apply the variable tests in levels (results are reported in Appendix table S.12, panel A). The null hypothesis of difference stationarity is rejected for all the variables, either with one or two breaks, using the bootstrap critical values given by Lee et al. (2006).[20] More precisely, for production and consumption, the unit

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

[20] Based on 5,000 replications of sample sizes $T = 100$.

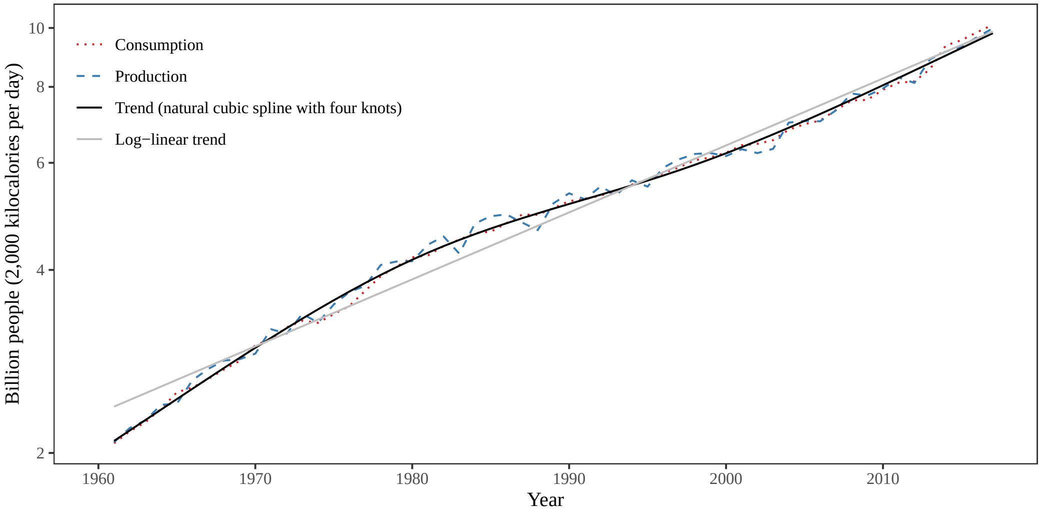

FIGURE 1. World caloric production and consumption, and their trend for 1961–2017. The y-axis is the number of people that hypothetically could be fed 2,000 kilocalories per day diet based on consumption of only the four commodities.

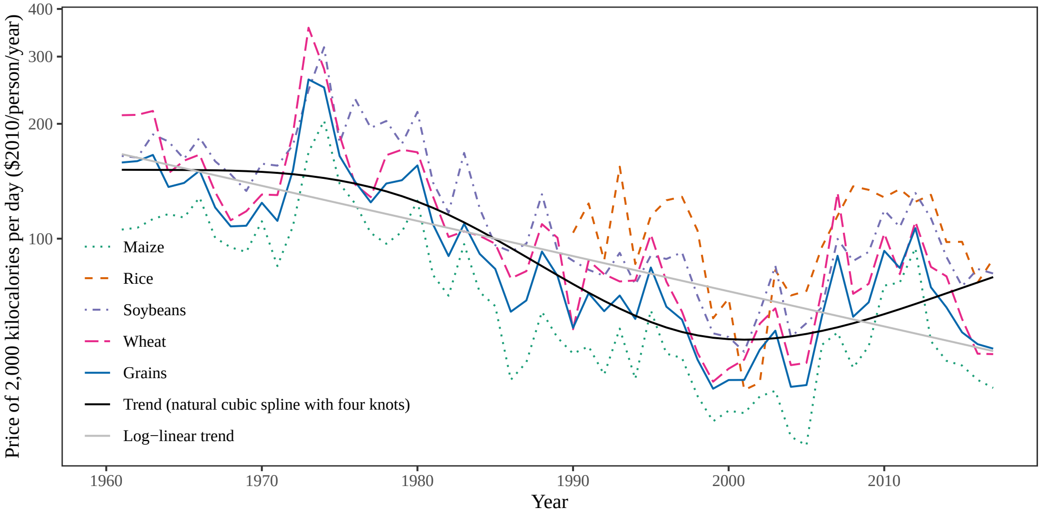

FIGURE 2. Real caloric prices at delivery. The y-axis is the annual cost of 2,000 kilocalories per day.

root assumption is rejected at the 5% significance level 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 hypothesis at the 5% level with a single estimated break occurring in 1979 and 1980.[21]

[21]It is interesting to note 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 also applies to consumption in relation to a

Overall, these tests support our choice of deterministic trends modeling. However, there is a mismatch between the log-linear trends assumed in the model and the trend flexibility needed to make the data stationary. Indeed, observed variables detrended using a log-linear trend cannot be considered stationary (see table S.12, first row of panel B), which prevents the direct mapping of the model variables to the observations. Put another way, the observed variables need to be detrended using a more flexible trend specification than what is accounted for in the theory, and we explain below the origins and consequences of this discrepancy.

Using log-linear trends in the model is motivated by the fact that an infinite horizon model with such trends can be expressed as a set of stationary equations, as proposed in section 2.2. This is no longer true with more flexible deterministic trends, such as splines or polynomials. Our storage model is the first to include long-run demand and cost trends, leading to rich long-run dynamics. Moreover, we are unaware of any existing approaches to build a storage model with more flexible trends that would remain compatible with an infinite horizon rational expectations framework.[22] Such a discrepancy between the flexibility of theoretical and empirical trends is common in macroeconomic models, where a trend that is theoretically consistent with a growth path may not be flexible enough to render the data stationary (Fernández-Villaverde et al., 2016, Section 8.4).

Since DSGE and storage models are primarily built to explain short-run fluctuations, the fact that they do not match long-run variations might be of secondary concern. However, this is not always the case not only because long-run fluctuations can impact short-run incentives, as shown by Bobenrieth et al. (2021) in the context of the storage model, but also because removing long-run variations from the data is not always innocuous (Canova, 2014). In our case, this discrepancy may have consequences, but, in the absence of a more general model, we cannot quantify them. One likely example of such discrepancy is the period starting in 2007, when biofuels mandates in Europe and United States led to higher demand for agricultural commodities and higher prices (Wright, 2014). In the absence of a regime-switching trend, this new regime is not accounted for in the model and part of the higher prices is absorbed by the spline trend. We leave further investigation of this issue to future work.

In practice, we handle the trends as follows. 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, in the esti-

regime change in 2007, which followed the implementation of the biofuels mandates in Europe and the United States (Wright, 2014).

[22]A couple of recent contributions to the econometrics of trending storage models have not faced the same problem (Bobenrieth et al., 2021, 2022). More precisely, Bobenrieth et al. (2021, 2022) develop estimation methods applied to a trending storage model and estimate the trend alongside some structural parameters, with the second paper introducing a one-step method that allows for the calculation of standard errors. However, these approaches are applicable provided the trends in the structural model are consistent with the observed trends, which requires commodities with log-linear price trends—an assumption rejected in our sample.

mations.[23] We further confirm 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. The results are reported in table S.12, panel B, with an increasing degree of flexibility going from the top to the bottom of each column. Except for the variables detrended with a linear trend or with three knots, we find all transformed series to be stationary at the 1% significance level. In other words, a natural cubic spline with three knots—i.e., flexibility equivalent to a quadratic trend—is not flexible enough to transform price and quantity data into stationary series. Since the four-knot spline involves the minimum flexibility needed to obtain stationary data, this is our preferred trend specification. As a robustness check, we also test for more and less flexible trends.

For the instrumental variables approach, following Roberts and Schlenker, we augment all the first- and second-stage equations with trend variables generated by natural cubic splines with three to five knots. For the indirect inference approach, we proceed in two steps. First, we detrend all variables by regressing their logarithms on spline variables and taking the exponential of the residuals as the detrended values. Second, we estimate the auxiliary model on these detrended variables.[24]

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

4.3 *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 1 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 crop prices are correlated more strongly to the grain index than to the prices of any of the other crops, except for the correlation between rice and soybeans. 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.

[23]Unless 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.

[24]Note that, in our setting, the combination of an auxiliary model based on linear regressions and a detrending using splines implies, by the Frisch–Waugh–Lovell theorem, that it is equivalent to jointly estimate trend and auxiliary parameters on observations or to do it in two steps as done here. This property ensures the consistency of the trend's treatment between the instrumental variable and indirect inference approaches.

TABLE 1. Correlation coefficients of detrended real prices at delivery, 1961–2017 (except rice, 1986–2017)

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

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

Table 2 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. The storage model's inability to match these high serial correlation levels in prices for a range of storable commodities initially led Deaton and Laroque (1992, 1996) to reject this model.[25] 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.[26] Production and yield shocks have small and insignificant autocorrelation in line with our model assumption of supply shocks without serial correlation.

TABLE 2. Autocorrelation and standard deviation of log detrended caloric data, 1961–2017

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

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

[25] Deaton and Laroque ruled out the presence of trend in commodity prices and so rejected the storage model on its failure to match much larger autocorrelations.

[26] The high consumption persistence is robust to detrending in level or logarithm and to the number of spline knots.

Table 3 displays the correlation coefficients of all the detrended variables in logarithm. The correlations with obvious counterparts in the model have the expected signs. Current and expected prices are strongly correlated, consistent with equation (4) in the presence of inventories frequently held. 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 3. Correlation coefficients of log detrended caloric data, 1961–2017

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

5. Estimation

5.1 *Structural parameters*

Before analyzing the results obtained by indirect inference in section 5.1.2, 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_υ, ρ_μ, and σ_ϑ), and indirect information about the others.

5.1.1 *Instrumental variable estimations* Tables 4 and 5 present the supply and demand estimates. To enable comparison with Roberts and Schlenker (2013), we replicate these estimates in Appendix (Tables S.13 and S.14) for a shorter sample (1962–2007) which corresponds to the sample length they used. The Appendix tables have some minor differences with the Table 1 in Roberts and Schlenker. These are due to two deviations from their approach: a slightly different procedure to construct the yield shock (in line with Hendricks et al., 2015), and the detrending of yields using a 4-knot spline rather than a 3-knot spline which is more consistent with our longer sample.

Table 4 reports the estimations of the supply equation. For the 2SLS estimates, the Cumby–Huizinga test rejects the hypothesis of residuals without serial correlation. We nevertheless report standard errors and diagnostic tests that are robust only to heteroskedasticity. Not only is this the most conservative choice in this particular setting but it also allows us to use the same type of standard errors for the supply and demand equations as well as for the weighting matrix of the indirect inference approach. 2SLS estimates of the supply elasticity are around 0.08, slightly lower than the values obtained by Roberts and Schlenker (2013). However, comparison with table S.13 shows that the

TABLE 4. Supply equation estimation

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

Notes: Standard errors robust to heteroskedasticity in parenthesis.

difference is entirely explained by our longer sample. The c_q estimates are always above $1 + \alpha_S$ (although not significantly). According to the discussion in section 3.1.1, this indicates a negative correlation between the two planting-time shocks (η and ω). The estimations using four and five knots are similar but present small differences with the estimations using three knots which is in line with the previous stationarity test results. Consistent with Hendricks et al.'s (2015) insights, the OLS and 2SLS supply elasticity estimates show only small and insignificant differences indicating that using the yield shock as a control variable helps to mitigate the omitted variable bias. This is further confirmed by the Hausman test which fails to reject the null of exogenous expected prices. However, the Monte Carlo analysis in the Appendix shows that in such short samples, the Hausman test tends to fail to reject the exogeneity despite prices being endogenous by construction. Therefore, we do not follow the Hausman test and for the comparisons that will follow our benchmark estimate is the 2SLS with four knots. For this specification, total supply shocks have a standard deviation σ_ϑ equal to 0.031, a value slightly above the standard deviation of production in table 2.

Table 5 presents the estimation results of the demand equation. The demand elasticity estimates are higher in absolute values than in Roberts and Schlenker (2013), which again seems to result from using a longer sample (see table S.14). We use equation (41) to

TABLE 5. Demand equation estimation

	(1)	(2)	(3)
Panel A. 2SLS			
Demand elasticity b_c	-0.051	-0.065	-0.060
	(0.028)	(0.026)	(0.027)
Lagged price c_c	0.041	0.019	0.014
	(0.016)	(0.014)	(0.014)
Lagged demand d_c	1.054	0.535	0.442
	(0.070)	(0.159)	(0.203)
Panel B. First stage			
Shock b_p	-4.287	-4.112	-4.014
	(0.882)	(0.937)	(1.056)
Lagged price c_p	0.569	0.486	0.498
	(0.087)	(0.105)	(0.111)
Lagged demand d_p	1.446	-0.130	0.523
	(0.745)	(1.690)	(2.012)
Panel C. OLS			
Demand elasticity b_c	-0.012	-0.021	-0.018
	(0.010)	(0.010)	(0.010)
Lagged price c_c	0.015	-0.005	-0.010
	(0.010)	(0.011)	(0.011)
Lagged demand d_c	0.949	0.547	0.413
	(0.044)	(0.118)	(0.162)
Panel D. 2SLS using Roberts and Schlenker's approach (eqs. (39) for 2nd stage and (43) for 1st)			
Demand elasticity b_c	-0.069	-0.079	-0.066
	(0.049)	(0.023)	(0.023)
$\sigma_{u_c^{\text{2SLS}}}$	0.018	0.016	0.016
σ_{u_P}	0.180	0.180	0.180
$\sigma_{u_c^{\text{OLS}}}$	0.016	0.014	0.013
$\sigma_{u_c^{\text{2SLS, RS}}}$	0.049	0.020	0.017
$\sigma_{\mu^{\text{2SLS}}}$		0.019	0.018
First stage F-stat (panel A)	23.627	19.252	14.443
p-value for Hausman test (panel A)	0.137	0.043	0.054
p-value for Cumby–Huizinga test (panel A)	0.851	0.199	0.057
First stage F-stat (panel D)	16.668	27.501	22.935
p-value for Hausman test (panel D)	0.000	0.029	0.052
p-value for Cumby–Huizinga test (panel D)	0.000	0.014	0.045
Observations	56	56	56
Spline knots	3	4	5

Notes: Standard errors robust to heteroskedasticity in parenthesis, except for panel D where they are also robust to autocorrelation. The lagged demand estimates in panel A are bias adjusted (Orcutt and Winokur, 1969).

estimate both the demand elasticity and autocorrelation of the demand shock. This contrasts with Roberts and Schlenker (2013) who use 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.[27] The Cumby–Huizinga test cannot reject the hypothesis of absence of serial correlation in the residuals for equation (41) (but not for equation (39)), which is a necessary condition for the consistent estimation of autoregressive terms. Estimates of the autocorrelation of the demand shocks differ depending on the number of knots. ρ_μ estimated along with a 3-knot spline is not statistically different from 1, indicating a nonstationary demand. This value confirms the results in section 4.2, which shows that a 3-knot spline is not sufficiently flexible to obtain stationary series.[28] A higher number of knots reduces ρ_μ by reducing the autocorrelation in the data, but at 0.53 (0.16) and 0.44 (0.20) for four and five knots the estimates are similar. The last parameter which can be identified from the demand estimation is the standard error of the demand shock. Using 4- and 5-knot splines, σ_υ (estimated by $\sigma_{u_c^{2\text{SLS}}}$) is about 0.016, which is slightly lower than the volatility of consumption observed in the raw data reported in table 2.

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

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

5.1.2 *Indirect inference estimations* We followed Roberts and Schlenker by presenting the instrumental variable results for natural cubic spline trends with three to five knots. However, both the unit-root tests and the estimates from table 5 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 are presented in table 6.

Most parameters are estimated precisely for both trend specifications despite the rather short sample size. The exceptions are the correlation between the planting-time

[27] Monte Carlo simulations (not reported in Appendix) show that using equation (41) instead of equation (39) leads to slightly smaller RMSE, consistent with the fact that more spherical residuals should make the estimator more efficient.

[28] Since ρ_μ is estimated above 1 for a 3-knot spline, it is not possible to calculate $\sigma_\mu^{2\text{SLS}}$ in table 5.

TABLE 6. Estimation results for the indirect inference approach

	4-knot spline		5-knot spline	
	Estimate	Standard error	Estimate	Standard error
ρ_μ	0.702	(0.081)	0.682	(0.086)
$\rho_{\eta,\omega}$	−0.443	(0.295)	−0.370	(0.289)
σ_ω	0.188	(0.031)	0.185	(0.030)
σ_η	0.014	(0.005)	0.014	(0.005)
σ_ϵ	0.020	(0.004)	0.020	(0.004)
σ_υ	0.019	(0.003)	0.018	(0.003)
k	0.037	(0.026)	0.034	(0.026)
α_D	−0.068	(0.019)	−0.059	(0.018)
α_S	0.086	(0.017)	0.086	(0.016)
σ_φ	0.027	(0.005)	0.026	(0.004)
σ_ψ	0.025	(0.002)	0.025	(0.002)
σ_μ	0.026	(0.006)	0.024	(0.006)
σ_ϑ	0.034	(0.004)	0.033	(0.004)

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

shocks ($\rho_{\eta,\omega}$) and the per-unit storage cost (k). This limited precision for these two parameters is consistent with the large standard errors obtained in table S.3 for the Monte Carlo analysis. For $\rho_{\eta,\omega}$, this is also consistent with the lack of precision in table 4 of the estimates of $c_q - 1$ from which $\rho_{\eta,\omega}$ is derived.

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

The volatility of the cost shock σ_ω is about 19% which is an order of magnitude larger than the estimates of the other shocks. However, the cost shock has no direct effect on quantities. Making it comparable to the other shocks requires its multiplication by α_S which produces 1.6% with four and five knots that is a contribution similar to the planting-time yield shock ($(1+\alpha_S)\sigma_\eta$). In the Monte Carlo analysis, such a large cost shock would make the 2SLS estimation of the supply equation very imprecise because the lagged yield shock would be a weak instrument, and could also create a wide gap between the OLS and the 2SLS estimates. This is not fully consistent with the results in table 4 where the OLS and 2SLS estimates are similar, indicating possible overestimation of σ_ω. The planting-time shocks η and ω can be aggregated in the shock φ. The standard deviation of φ exceeds the standard deviation of harvest-time yield shock σ_ϵ, which indicates that the majority of supply shocks is known before deciding to produce. Finally, these three supply shocks can be aggregated together. The last row in table 6 shows that the standard deviation of the resulting total supply shock ϑ is about 30% larger than the standard deviation of the demand shock, μ. The historical role of demand and supply shocks in past food price crises is discussed in Appendix D.3 using a historical decomposition

The per-unit storage cost (k) is estimated at 3.7% of the steady-state price with four knots. By combining the opportunity costs related to the interest rate and the price trend, we obtain an estimated total annual storage cost of around 7.6% at the steady state ($k + 1 - \beta \exp g_p$). 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%, exactly the opportunity cost implied by the downward price trend. The cost created by the positive quantity trend also contributes to higher storage costs but cannot be characterized analytically and is thus ignored in this discussion. Similar to the price trend, assuming a constant interest rate different from the 2% assumed here would lead to the same total annual storage costs but with a different physical cost, k. Simulations based on a model calibrated on the estimated parameters predict an occurrence of stockouts of 11%.[29]

Overall, these results suggest that our indirect inference approach returns fairly precise parameter estimates that are reasonably consistent with the 2SLS estimates. Appendix D.1 examines how close the parameters of the auxiliary model are when estimated on observations and on simulations. Appendix D.2 addresses some of the potential concerns about these estimations, including the roles of data, the weighting matrix, and each moment in the estimation process. All sensitivity analyses confirm the picture presented here. Since the differences across trend specifications are small, all subsequent analyses are based on the estimation using the 4-knot spline, our preferred trend specification.

5.2 *Inspecting the model fit on other moments*

We next assess the performance of the estimated storage model by comparing the variances and covariances based on model simulations and those based on observations (as typically done following the estimation of DSGE models, e.g., Smets and Wouters, 2003). Recall that so far the empirical performance of estimated storage models was judged based only on their ability to replicate price-based moments given that only prices were used for the estimations. By focusing on second-order moments calculated up to one lag for each of our 5 observables, our empirical setting now allows evaluation of the model fit over 40 moments (which could be mapped to the 40 parameters of a VAR(1) model). The results of this exercise are presented in table 7 which includes all the moments calculated on the detrended observations, their standard deviation calculated by bootstrap, and the corresponding moments from the simulated model. 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.[30] The majority of the moments are similar for observations

[29] Since the solution method involves linear interpolation over a sparse grid, it cannot precisely identify stockouts. Instead of zero stocks, very small values will be predicted. Thus, a stockout is defined here as a stock level below 1E-4, which corresponds to 0.1% of the average demand.

[30] A more formal test would have followed Le et al. (2011, 2016) by using indirect inference to test the storage model against a VAR(1) model based on the five observables. Since our informal moments comparison leads to a clear rejection of the estimated model, there is no need for a more sophisticated approach at this stage.

and simulations, indicating that our extended storage model is generally able to replicate the main dynamics in the data. This applies in particular to the first-order autocorrelation of price, the subject of long-standing debates since Deaton and Laroque (1992).[31]

TABLE 7. Comparison of actual and model-based second-order moments

Moment	Observed	St. dev.	Simulated	Moment	Observed	St. dev.	Simulated
$\sigma_{\ln p}$	0.236	0.023	0.262	$\phi_{\ln p,\ln c}(1)$	−0.469	0.125	0.191
$\sigma_{\ln c}$	0.019	0.002	0.018	$\phi_{\ln p,\ln q}(1)$	0.104	0.156	−0.014
$\sigma_{\ln q}$	0.028	0.002	0.031	$\phi_{\ln p,\ln \mathrm{E}\, p}(1)$	0.643	0.069	0.628
$\sigma_{\ln \mathrm{E}\, p}$	0.193	0.018	0.180	$\phi_{\ln p,\psi}(1)$	−0.274	0.142	−0.183
σ_{ψ}	0.024	0.002	0.025	$\phi_{\ln c,\ln p}(1)$	−0.326	0.109	0.205
$\phi_{\ln p}(1)$	0.576	0.110	0.560	$\phi_{\ln c,\ln q}(1)$	0.184	0.110	0.299
$\phi_{\ln c}(1)$	0.642	0.146	0.568	$\phi_{\ln c,\ln \mathrm{E}\, p}(1)$	−0.300	0.118	0.182
$\phi_{\ln q}(1)$	0.042	0.140	−0.011	$\phi_{\ln c,\psi}(1)$	0.304	0.127	0.187
$\phi_{\ln \mathrm{E}\, p}(1)$	0.652	0.116	0.607	$\phi_{\ln q,\ln p}(1)$	−0.257	0.110	0.216
$\phi_{\psi}(1)$	0.146	0.142	0.001	$\phi_{\ln q,\ln c}(1)$	0.323	0.110	0.353
$\phi_{\ln p,\ln c}(0)$	−0.488	0.102	0.083	$\phi_{\ln q,\ln \mathrm{E}\, p}(1)$	−0.212	0.116	0.093
$\phi_{\ln p,\ln q}(0)$	−0.406	0.103	−0.182	$\phi_{\ln q,\psi}(1)$	0.067	0.134	−0.142
$\phi_{\ln p,\ln \mathrm{E}\, p}(0)$	0.939	0.017	0.871	$\phi_{\ln \mathrm{E}\, p,\ln p}(1)$	0.566	0.094	0.535
$\phi_{\ln p,\psi}(0)$	−0.534	0.118	−0.454	$\phi_{\ln \mathrm{E}\, p,\ln c}(1)$	−0.508	0.116	0.293
$\phi_{\ln c,\ln q}(0)$	0.395	0.109	0.591	$\phi_{\ln \mathrm{E}\, p,\ln q}(1)$	0.070	0.147	0.044
$\phi_{\ln c,\ln \mathrm{E}\, p}(0)$	−0.452	0.106	0.284	$\phi_{\ln \mathrm{E}\, p,\psi}(1)$	−0.358	0.129	−0.137
$\phi_{\ln c,\ln \psi}(0)$	0.529	0.116	0.463	$\phi_{\ln \psi,\ln p}(1)$	−0.162	0.108	−0.120
$\phi_{\ln q,\ln \mathrm{E}\, p}(0)$	−0.271	0.115	−0.024	$\phi_{\ln \psi,\ln c}(1)$	0.334	0.127	0.124
$\phi_{\ln q,\psi}(0)$	0.775	0.050	0.831	$\phi_{\ln \psi,\ln q}(1)$	−0.115	0.122	0.002
$\phi_{\ln \mathrm{E}\, p,\psi}(0)$	−0.500	0.118	−0.291	$\phi_{\ln \psi,\ln \mathrm{E}\, p}(1)$	−0.203	0.115	−0.226

Notes: Moments calculated over 100,000 sample observations from the asymptotic distribution simulated with a storage model calibrated with the indirect inference estimates with a 4-knot spline from table 6. $\phi(1)$ denotes first-order serial correlation and $\phi_{i,j}(l) = \mathrm{cor}(i_{t-l}, j_t)$ denotes l^{th}-order correlation between variable i and j. Statistics involving $\mathrm{E}\, p$ refer to $\mathrm{E}_t\, p_{t+1}$, e.g., $\phi_{\ln p,\ln \mathrm{E}\, p}(0) = \mathrm{cor}(\ln p_t, \ln \mathrm{E}_t\, p_{t+1})$. Standard deviation calculated by bootstrapping the dataset of detrended variables using 5,000 bootstrap replications.

However, it can be seen that the storage model fails to match some moments (14 lie outside the 10% bootstrap confidence interval including 11 outside the 5% confidence interval). These moments mostly relate to two aspects. Six moments are related to consumption and its (lagged) covariance with current and expected prices. In particular, the model fit related to the negative correlation between consumption and spot prices is problematic: $\mathrm{cor}(\ln p_t, \ln c_t) = -0.49$ on observations but 0.08 on simulations. Logically, given the strong autocorrelation of both prices and consumption combined with the strong correlation between current and expected prices, this issue persists with a lag as well as if we consider expected instead of current prices.[32] A similar problem arises

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

[32] It is worth noting that consumption is actually a reconstructed variable based on the difference between production and stock variations. In other words, part of this mismatch might simply be due to an artifact of the data construction and measurement errors in the global stock variations.

for four moments related to production and its (lagged) covariance with current and expected prices.

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

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

5.3 *The role of storage in market dynamics*

The introduction of many new features in our storage model calls for investigation of their respective contributions to the price and quantity dynamics generated by the model. In this section, we explore the role of storage in the movement of prices based on the alternative exclusion of the various model features. To save space, we restrict the discussion to six moments of interest: price autocorrelation which since Deaton and Laroque (1992) is the benchmark metric used to assess the performance of the storage model, price, consumption as well as production volatilities, and the correlation between price and consumption, and price and production. $\phi_{\ln p,\ln c}(0)$ and $\phi_{\ln p,\ln q}(0)$ are of particular interest because in the previous section we showed that the model struggles to match these moments; thus, it is helpful to examine which model characteristics is driving their behavior. Table 8 reports the results of this exercise as well as the same moments calculated on the raw and detrended data for comparison.[33]

Switching off the model features one at a time allows us to analyze their respective contributions 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 (i.e., difference between the rows “Trending data” and “Detrended data”). Regarding the remaining serial correlation explained by the benchmark model, the various model features contribute jointly and nonlinearly to the persistence. This prevents us from uniquely decomposing their contributions. Nonetheless, we can informally rank them in descending order of importance

[33]Note that the row “Detrended data” represents the same information as the column “Observed” in table 7.

TABLE 8. Role of model assumptions in price and quantity dynamics

Data or model	$\phi_{\ln p}(1)$	$\sigma_{\ln p}$	$\sigma_{\ln c}$	$\sigma_{\ln q}$	$\phi_{\ln p,\ln c}(0)$	$\phi_{\ln p,\ln q}(0)$
Trending data	0.87	0.46	–	–	–	–
Detrended data	0.58	0.24	0.019	0.028	-0.49	-0.41
1. Benchmark	0.56	0.26	0.018	0.031	0.08	-0.18
2. $\rho_\mu = 0$	0.38	0.21	0.017	0.029	-0.33	-0.50
3. $\rho_\mu = 0, \sigma_\upsilon = \sigma_\mu$	0.38	0.23	0.022	0.030	-0.04	-0.38
4. $\alpha_S = 0$	0.65	0.30	0.014	0.024	0.10	-0.17
5. $g_q = 0$	0.57	0.26	0.018	0.031	0.08	-0.17
6. $g_p = 0$	0.60	0.24	0.018	0.032	0.19	-0.14
7. $k = 0.018$	0.60	0.24	0.018	0.032	0.19	-0.14
8. $\sigma_\eta = 0$	0.53	0.25	0.017	0.027	0.19	-0.12
9. $\sigma_\omega = 0$	0.54	0.25	0.016	0.026	0.20	-0.11
10. $\sigma_\eta = 0, \sigma_\epsilon = \sigma_\psi$	0.53	0.26	0.017	0.030	0.09	-0.20
11. $\sigma_\omega = \sigma_\eta = 0, \sigma_\epsilon = \sigma_\psi$	0.51	0.26	0.017	0.028	0.15	-0.16
12. $\rho_\mu = 0, \sigma_\upsilon = \sigma_\mu, \alpha_S = 0$	0.25	0.20	0.022	0.026	0.06	-0.36
13. $\rho_\mu = 0, \sigma_\upsilon = \sigma_\mu, \alpha_S = 0, \sigma_\eta = 0, \sigma_\epsilon = \sigma_\psi, g_q = 0$	0.26	0.21	0.022	0.028	0.00	-0.40
14. $\rho_\mu = 0.535, \sigma_\upsilon = 0.016$	0.46	0.23	0.016	0.030	-0.26	-0.39
15. $k = \infty$	0.16	0.45	0.025	0.025	-0.59	-0.59

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 6, except for the parameter values indicated in the first column.

as follows: the smoothing effect of storage, the autocorrelation coefficient of demand innovations, and the presence of planting-time shocks.

Comparing model 1 to model 2 shows the contribution of demand persistence, which increases the price autocorrelation from 0.38 to 0.56. This result contrasts with Deaton and Laroque's (1996) estimation results for a model with autocorrelated supply shocks. Indeed, Deaton and Laroque found that almost all of the serial correlation in prices was attributable to the persistence of shocks, not speculative storage. The key difference with our findings lies in our use of quantities as observables. This ensures that any shock autocorrelation must be compatible with the quantity dynamics, which is not necessarily the case if we rely solely on price data. The contribution of planting-time shocks is much more modest and can be observed from the difference between models 10 and 1, with an increase from 0.53 to 0.56. Planting-time shocks contribute to price persistence by linking periods: shocks at planting time affect production and, therefore, prices in the following period. However, since these shocks are immediately observed, they also affect current prices due to the intertemporal link created by storage. Storage is thus key to inducing price persistence. In model 15, where storage is shut down, the autocorrelation decreases to 0.16. This figure should be compared to the price autocorrelation in model 13, in which all model features, except for storage, are turned off, making it the closest to the simplest model estimated in Deaton and Laroque (1996).[34] In this setup—where only storage explains price persistence—the price autocorrelation is less than half of its benchmark

[34]The only differences are the use of an isoelastic demand function and multiplicative shocks, while Deaton and Laroque (1996) use a linear demand function and additive shocks.

value. This demonstrates that storage alone does not explain price persistence; rather, it does so only in combination with other features. This also suggests that to match the true persistence of prices, estimation of the simpler version of the storage model considered in the literature so far would require lower storage costs. Finally, the inclusion of a supply response has an ambiguous effect on price autocorrelation. Comparing the benchmark model to model 4 shows 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 with 12—a supply response increases price persistence.

The simulations of the estimated model raise a new puzzle about the inability of the model to match the price-consumption correlation. This moment is explained by the respective roles of the demand and supply shocks in driving price movements, combined with the demand elasticity. At the extreme without demand shocks, the correlation would be -1. Therefore, removing planting-time shocks (models 8–11) or the supply response (model 4) would only decrease the role of supply shocks and exacerbate the problems related to this moment. Some improvement can be achieved by removing the persistence of the demand shock (models 2–3) or increasing the storage cost (model 15), but both lead to a lower fit of the price autocorrelation. The indirect inference approach overestimates ρ_μ by 0.168 and σ_υ by 0.003 compared with the 2SLS approach. Comparing models 2 and 3 with the benchmark shows that overestimation of ρ_μ would contribute only a little to solving this puzzle. However, setting the size of the demand shock equal to its 2SLS estimate level, in addition to ρ_μ (model 14), would bring the simulated moment closer to the observed moment, inside the 99% bootstrap confidence interval but outside the 95% interval. In other words, the covariance mismatch between consumption and price might be due in part to the overstatement of both the persistence and variance of the demand shocks.

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

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

Overall, the effects of the various model features on consumption and production volatility have the expected signs. We next discuss the effects of the model variants not considered so far. In model 5, the positive trend on quantities g_q is removed. As discussed in section 2.2 this boils down to decreasing storage costs which slightly increases price persistence. In model 6, the negative trend on price g_p is removed. Because the price trend directly affects the storers' incentives, for a value similar to g_q it has a stronger impact on the autocorrelation in prices. In addition, a comparison of models 6 with 7 shows that the impact of the price trend g_p is very similar to the effect of a corresponding decrease in the per-unit storage cost k (i.e., decreasing it by $\beta[1-\exp(g_p)]$).

6. Conclusions

This paper proposes a new empirical strategy to estimate a rational expectations storage model. It requires five observables (current price, expected price, production, consumption, and supply shock) and reliance on a simple linear supply and demand model as the auxiliary model in an indirect inference approach. Including quantities as well as prices within the set of observables is crucial because it allows estimation of all the model parameters which is important to empirically validate the model and run counterfactual simulations for policy applications. Although the key role of storage for mediating the dynamics of commodity prices has long been acknowledged and has been exploited widely in finance and economics, so far a full empirical validation of a rational expectations storage model has not been carried out. To apply our approach, we chose the empirical setting of the global grains market following Roberts and Schlenker (2013), who use an instrumental variable strategy motivated by storage theory. While they estimate only a subset of the structural parameters, their strategy provides a good benchmark for comparing our indirect inference estimates. We also used their estimating equations to choose our auxiliary model.

Our results show that the long-standing price autocorrelation puzzle highlighted by Deaton and Laroque (1992, 1996) can be solved convincingly by accounting for sufficient features of the market for grains, such as (in descending order of importance): storage, a long-run price trend, autocorrelated demand shocks, and producers' incentive shocks associated with an elastic supply.

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

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

References

Altonji, Joseph G. and Lewis M. Segal (1996), “Small-sample bias in GMM estimation of covariance structures.” *Journal of Business & Economic Statistics*, 14 (3), 353–366, URL https://www.jstor.org/stable/1392447. [17]

Barro, Robert J. and Xavier Sala-i-Martín (1990), “World real interest rates.” In *NBER Macroeconomics Annual 1990* (Olivier Jean Blanchard and Stanley Fischer, eds.), volume 5, chapter 1, 15–74, MIT Press, Cambridge, MA. [12]

Baumeister, Christiane and James D. Hamilton (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, URL https://www.aeaweb.org/articles?id=10.1257/aer.20151569. [4, 36]

Bobenrieth, Eugenio S. A., Juan R. A. Bobenrieth, Ernesto A. Guerra, Brian D. Wright, and Di Zeng (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. [2, 4, 6, 9, 22, 34]

Bobenrieth, Eugenio S. A., Juan R. A. Bobenrieth, and Brian D. Wright (2002), “A commodity price process with a unique continuous invariant distribution having infinite mean.” *Econometrica*, 70 (3), 1213–1219, URL http://www.jstor.org/stable/2692312. [10]

Bobenrieth, Juan R. A., Eugenio S. A. Bobenrieth, Andrés F. Villegas, and Brian D. Wright (2022), “Estimation of endogenous volatility models with exponential trends.” *Mathematics*, 10 (15), 2647. [22]

Bornstein, Gideon, Per Krusell, and Sergio Rebelo (2023), “A world equilibrium model of the oil market.” *The Review of Economic Studies*, 90 (1), 132–164. [4]

Brennan, Michael J. (1958), “The supply of storage.” *The American Economic Review*, 48 (1), 50–72, URL http://www.jstor.org/stable/1812340. [7]

Cafiero, Carlo, Eugenio S. A. Bobenrieth, Juan R. A. Bobenrieth, and Brian D. Wright (2011), “The empirical relevance of the competitive storage model.” *Journal of Econometrics*, 162 (1), 44–54, URL http://ideas.repec.org/a/eee/econom/v162y2011i1p44-54.html. [2, 11]

Cafiero, Carlo, Eugenio S. A. Bobenrieth, Juan R. A. Bobenrieth, and Brian D. Wright (2015), “Maximum Likelihood estimation of the standard commodity storage model: Evidence from sugar prices.” *American Journal of Agricultural Economics*, 97 (1), 122–136, URL http://ajae.oxfordjournals.org/content/early/2014/08/13/ajae.aau068.abstract. [2]

Canova, Fabio (2014), "Bridging DSGE models and the raw data." *Journal of Monetary Economics*, 67, 1–15, URL http://www.sciencedirect.com/science/article/pii/S030439321400083X. [22, 35]

Carter, Colin A., Gordon C. Rausser, and Aaron Smith (2011), "Commodity booms and busts." *Annual Review of Resource Economics*, 3 (1), 87–118, URL http://ideas.repec.org/a/anr/reseco/v3y2011p87-118.html. [7]

Carter, Colin A., Gordon C. Rausser, and Aaron Smith (2017), "Commodity storage and the market effects of biofuel policies." *American Journal of Agricultural Economics*, 99 (4), 1027–1055, URL http://dx.doi.org/10.1093/ajae/aaw010. [4]

Chambers, Marcus J. and Roy E. Bailey (1996), "A theory of commodity price fluctuations." *Journal of Political Economy*, 104 (5), 924–957, URL http://www.jstor.org/stable/2138947. [3]

Chen, Yu-Chin, Kenneth S. Rogoff, and Barbara Rossi (2010), "Can exchange rates forecast commodity prices?" *The Quarterly Journal of Economics*, 125 (3), 1145–1194, URL http://qje.oxfordjournals.org/content/125/3/1145.abstract. [7]

Chinn, Menzie D. and Olivier Coibion (2014), "The predictive content of commodity futures." *Journal of Futures Markets*, 34 (7), 607–636. [19]

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2005), "Nominal rigidities and the dynamic effects of a shock to monetary policy." *Journal of Political Economy*, 113 (1), 1–45, URL http://www.journals.uchicago.edu/doi/abs/10.1086/426038. [4, 17]

Corradi, Valentina and Norman R. Swanson (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, URL https://www.sciencedirect.com/science/article/pii/S0304407605000461. Common Features. [20]

Cross, Jamie L., Bao H. Nguyen, and Trung Duc Tran (2022), "The role of precautionary and speculative demand in the global market for crude oil." *Journal of Applied Econometrics*, 37 (5), 882–895, URL https://onlinelibrary.wiley.com/doi/abs/10.1002/jae.2905. [4]

Cumby, Robert E. and John Huizinga (1992), "Testing the autocorrelation structure of disturbances in ordinary least squares and instrumental variables regressions." *Econometrica*, 60 (1), 185. [15]

Deaton, Angus and Guy Laroque (1992), "On the behaviour of commodity prices." *The Review of Economic Studies*, 59 (1), 1–23, URL http://www.jstor.org/stable/2297923. [1, 2, 4, 9, 19, 24, 31, 32, 35]

Deaton, Angus and Guy Laroque (1996), "Competitive storage and commodity price dynamics." *Journal of Political Economy*, 104 (5), 896–923, URL http://ideas.repec.org/a/ucp/jpolec/v104y1996i5p896-923.html. [1, 3, 7, 24, 33, 35]

DeJong, David N., John C. Nankervis, N. E. Savin, and Charles H. Whiteman (1992), "Integration versus trend stationary in time series." *Econometrica*, 60 (2), 423–433, URL http://www.jstor.org/stable/2951602. [20]

Dvir, Eyal and Kenneth Rogoff (2014), "Demand effects and speculation in oil markets: Theory and evidence." *Journal of International Money and Finance*, 42, 113–128, URL http://www.sciencedirect.com/science/article/pii/S0261560613001095. [4]

FAO (2020), "FAOSTAT, Food and Agriculture Organization of the United Nations statistical database." http://www.fao.org/faostat/en/#home. Accessed June 24, 2020. [18]

Fernández-Villaverde, Jesús, Juan F. Rubio-Ramírez, and Frank Schorfheide (2016), "Solution and estimation methods for DSGE models." In *Handbook of Macroeconomics* (John Taylor and Harald Uhlig, eds.), volume 2, chapter 9, 527–724, Elsevier. [22]

Garcia, Philip, Scott H. Irwin, and Aaron Smith (2015), "Futures market failure?" *American Journal of Agricultural Economics*, 97 (1), 40–64, URL http://ajae.oxfordjournals.org/content/early/2014/08/14/ajae.aau067.abstract. [19]

Gardner, Bruce L. (1976), "Futures prices in supply analysis." *American Journal of Agricultural Economics*, 58 (1), 81–84, URL http://www.jstor.org/stable/1238581. [19]

Ghanem, D. and Aaron Smith (2022), "Causality in structural vector autoregressions: Science or sorcery?" *American Journal of Agricultural Economics*, 104 (3), 881–904, URL https://asmith.ucdavis.edu/research/causality-structural-vector-autoregressions-science-or-sorcery. [4]

Ghoshray, Atanu (2010), "A reexamination of trends in primary commodity prices." *Journal of Development Economics*, 95 (2), 242–251, URL http://ideas.repec.org/a/eee/deveco/v95y2011i2p242-251.html. [20]

Gilbert, Christopher L (1989), "The impact of exchange rates and developing country debt on commodity prices." *The Economic Journal*, 99 (397), 773–784, URL http://ideas.repec.org/a/ecj/econjl/v99y1989i397p773-84.html. [7]

Gouel, Christophe (2013), "Optimal food price stabilisation policy." *European Economic Review*, 57, 118–134, URL http://www.sciencedirect.com/science/article/pii/S0014292112001419. [1]

Gouel, Christophe (2020), "The value of public information in storable commodity markets: Application to the soybean market." *American Journal of Agricultural Economics*, 102 (3), 846–865. [3, 13]

Gouel, Christophe and Nicolas Legrand (2017), "Estimating the competitive storage model with trending commodity prices." *Journal of Applied Econometrics*, 32 (4), 744–763. [2, 16]

Gouel, Christophe and Nicolas Legrand (2022), "The role of storage in commodity markets: Indirect inference based on grains data." Working Papers 2022-04, CEPII, URL https://www.cepii.fr/CEPII/fr/publications/wp/abstract.asp?NoDoc=13445. [7, 15]

Gourieroux, Christian, Alain Monfort, and Eric Renault (1993), "Indirect inference." *Journal of Applied Econometrics*, 8 (S1), S85–S118, URL https://onlinelibrary.wiley.com/doi/abs/10.1002/jae.3950080507. [3, 15, 17, 18]

Guerra, Ernesto Alex, E. S. A. Bobenrieth, J. R. A. Bobenrieth, and Carlo Cafiero (2015), "Empirical commodity storage model: the challenge of matching data and theory." *European Review of Agricultural Economics*, 42 (4), 607–623, URL http://erae.oxfordjournals.org/content/42/4/607.abstract. [19]

Gustafson, Robert L. (1958), "Implications of recent research on optimal storage rules." *Journal of Farm Economics*, 40 (2), 290–300, URL http://www.jstor.org/stable/1234919. [6]

Guvenen, Fatih and Anthony A. Smith (2014), "Inferring labor income risk and partial insurance from economic choices." *Econometrica*, 82 (6), 2085–2129. [4, 16]

Harrell, Frank E (2001), *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*. Springer, New York. [23]

Hendricks, Nathan P., Joseph P. Janzen, and Aaron Smith (2015), "Futures prices in supply analysis: Are instrumental variables necessary?" *American Journal of Agricultural Economics*, 97 (1), 22–39, URL http://ajae.oxfordjournals.org/content/97/1/22.abstract. [14, 16, 18, 19, 25, 26]

IMF (2014), *World Economic Outlook: Recovery strengthens, Remains uneven*. International Monetary Fund, Washington, DC. [12]

Kaldor, Nicholas (1939), "Speculation and economic stability." *The Review of Economic Studies*, 7 (1), 1–27, URL http://www.jstor.org/stable/2967593. [7]

Kilian, Lutz (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, URL http://www.jstor.org/stable/25592494. [4, 36]

Kilian, Lutz and Daniel P. Murphy (2014), "The role of inventories and speculative trading in the global market for crude oil." *Journal of Applied Econometrics*, 29 (3), 454–478. [1, 4, 36]

Knittel, Christopher R. and Robert S. Pindyck (2016), "The simple economics of commodity price speculation." *American Economic Journal: Macroeconomics*, 8 (2), 85–110. [35]

Känzig, Diego R. (2021), "The macroeconomic effects of oil supply news: Evidence from OPEC announcements." *The American Economic Review*, 111 (4), 1092–1125, URL https://www.aeaweb.org/articles?id=10.1257/aer.20190964. [4]

Le, Vo Phuong Mai, David Meenagh, Patrick Minford, and Michael Wickens (2011), "How much nominal rigidity is there in the US economy? Testing a new keynesian DSGE model using indirect inference." *Journal of Economic Dynamics and Control*, 35 (12), 2078–2104. [4, 17, 18, 30]

Le, Vo Phuong Mai, David Meenagh, Patrick Minford, Michael Wickens, and Yongdeng Xu (2016), "Testing Macro Models by Indirect Inference: A Survey for Users." *Open Economies Review*, 27 (1), 1–38. [18, 30]

Lee, Junsoo, John A. List, and Mark C. Strazicich (2006), "Non-renewable resource prices: Deterministic or stochastic trends?" *Journal of Environmental Economics and Management*, 51 (3), 354–370, URL http://www.sciencedirect.com/science/article/pii/S0095069605000914. [20]

Lee, Junsoo and Mark C. Strazicich (2003), "Minimum Lagrange multiplier unit root test with two structural breaks." *The Review of Economics and Statistics*, 85 (4), 1082–1089, URL http://www.jstor.org/stable/3211829. [20]

Lee, Junsoo and Mark C. Strazicich (2013), "Minimum LM unit root test with one structural break." *Economics Bulletin*, 33 (4), 2483–2492, URL https://ideas.repec.org/a/ebl/ecbull/eb-13-00296.html. [20]

Letta, Marco, Pierluigi Montalbano, and Guillaume Pierre (2022), "Weather shocks, traders' expectations, and food prices." *American Journal of Agricultural Economics*, 104 (3), 1100–1119, URL https://onlinelibrary.wiley.com/doi/abs/10.1111/ajae.12258. [1]

Li, Tong (2010), "Indirect inference in structural econometric models." *Journal of Econometrics*, 157 (1), 120–128. [16]

Lowry, Mark, Joseph W. Glauber, Mario J. Miranda, and Peter G. Helmberger (1987), "Pricing and storage of field crops: A quarterly model applied to soybeans." *American Journal of Agricultural Economics*, 69 (4), 740–749, URL http://www.jstor.org/stable/1242183. [3]

Maliar, Lilia and Serguei Maliar (2014), "Numerical methods for large-scale dynamic economic models." In *Handbook of Computational Economics* (Karl Schmedders and Kenneth L. Judd, eds.), volume 3, chapter 7, 325–477, Elsevier, Amsterdam. [11]

Michaelides, Alexander and Serena Ng (2000), "Estimating the rational expectations model of speculative storage: A Monte Carlo comparison of three simulation estimators." *Journal of Econometrics*, 96 (2), 231–266, URL http://ideas.repec.org/a/eee/econom/v96y2000i2p231-266.html. [4, 16, 18]

Nordhaus, William D. (1987), "Forecasting efficiency: Concepts and applications." *The Review of Economics and Statistics*, 69 (4), 667–674. [6]

Oglend, Atle and Tore Selland Kleppe (2017), "On the behavior of commodity prices when speculative storage is bounded." *Journal of Economic Dynamics and Control*, 75, 52–69. [7]

Orcutt, Guy H. and Herbert S. Winokur (1969), "First order autoregression: Inference, estimation, and prediction." *Econometrica*, 37 (1), 1–14. [15, 27]

Osborne, Theresa (2004), "Market news in commodity price theory: Application to the Ethiopian grain market." *The Review of Economic Studies*, 71 (1), 133–164, URL http://www.jstor.org/stable/3700714. [3]

Perron, Pierre (1989), "The great crash, the oil price shock, and the unit root hypothesis." *Econometrica*, 57 (6), 1361–1401, URL http://www.jstor.org/stable/1913712. [20]

Porteous, Obie (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. [1]

Powell, Michael J. D. (2009), "The BOBYQA algorithm for bound constrained optimization without derivatives." Numerical analysis preprint DAMTP 2009/NA06, University of Cambridge. [18]

Roberts, Michael J. and Wolfram Schlenker (2013), "Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate." *The American Economic Review*, 103 (6), 2265–2295. [2, 3, 4, 5, 9, 12, 13, 15, 18, 19, 20, 22, 23, 25, 26, 27, 28, 35]

Rotemberg, Julio J. and Michael Woodford (1997), "An optimization-based econometric framework for the evaluation of monetary policy." *NBER Macroeconomics Annual*, 12, 297–346. [4]

Routledge, Bryan R., Duane J. Seppi, and Chester S. Spatt (2000), "Equilibrium forward curves for commodities." *The Journal of Finance*, 55 (3), 1297–1338, URL http://dx.doi.org/10.1111/0022-1082.00248. [3]

Ruge-Murcia, Francisco J. (2007), "Methods to estimate dynamic stochastic general equilibrium models." *Journal of Economic Dynamics and Control*, 31 (8), 2599–2636. [18]

Ruge-Murcia, Francisco J. (2012), "Estimating nonlinear DSGE models by the simulated method of moments: With an application to business cycles." *Journal of Economic Dynamics and Control*, 36 (6), 914–938, URL http://www.sciencedirect.com/science/article/pii/S0165188912000231. [18]

Ruge-Murcia, Francisco J. (2020), "Estimating nonlinear dynamic equilibrium models by matching impulse responses." *Economics Letters*, 197, 109624. [4, 17]

Smets, Frank and Raf Wouters (2003), "An estimated dynamic stochastic general equilibrium model of the euro area." *Journal of the European Economic Association*, 1 (5), 1123–1175, URL http://www.jstor.org/stable/40004854. [4, 30]

Smets, Frank and Rafael Wouters (2007), "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *The American Economic Review*, 97 (3), 586–606, URL http://www.jstor.org/stable/30035013. [4]

Smith, Anthony A., Jr. (1993), "Estimating nonlinear time-series models using simulated vector autoregressions." *Journal of Applied Econometrics*, 8 (S1), S63–S84. [3]

Steinwender, Claudia (2018), "Real effects of information frictions: When the States and the Kingdom became united." *The American Economic Review*, 108 (3), 657–696. [1, 6]

Williams, Jeffrey C. (1986), *The Economic Function of Futures Markets*. Cambridge University Press. [7]

Williams, Jeffrey C. and Brian D. Wright (1991), *Storage and commodity markets*. Cambridge University Press. [3]

Williamson, Lucille and Paul Williamson (1942), "What we eat." *Journal of Farm Economics*, 24 (3), 698–703. [18]

Working, Holbrook (1949), "The theory of price of storage." *The American Economic Review*, 39 (6), 1254–1262, URL http://www.jstor.org/stable/1816601. [7]

Wright, Brian D. (2014), "Global biofuels: Key to the puzzle of grain market behavior." *Journal of Economic Perspectives*, 28 (1), 73–98, URL http://www.aeaweb.org/articles.php?doi=10.1257/jep.28.1.73. [22]

Wright, Brian D. and Jeffrey C. Williams (1982), "The economic role of commodity storage." *The Economic Journal*, 92 (367), 596–614, URL http://www.jstor.org/stable/2232552. [3, 6, 7]

Wright, Brian D. and Jeffrey C. Williams (1984), "The welfare effects of the introduction of storage." *The Quarterly Journal of Economics*, 99 (1), 169–192, URL http://www.jstor.org/stable/1885726. [6]

Zivot, Eric and Donald W K Andrews (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. [20]