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

Christophe Gouel Nicolas Legrand

A Model with structural trends

In this section, we develop a variant of our storage model in which the trends are not equilibrium trends but structural trends: trends that are demand and cost shifters. We show how the equilibrium trends used in the paper relate to structural trends. To simplify the exposition, we replace the various constant parameters in the cost and demand functions respectively with two parameters χ and π. The cost function is assumed to be $\chi H_t^{1+1/\alpha_S} \exp(g_c t/\alpha_S + \omega_t)/(1+1/\alpha_S)$, where g_c is the cost trend expressed as a shifter of the supply curve. Similarly, demand is given by $D_t = \pi P_t^{\alpha_D} \exp(g_d t + \mu_t)$, where g_d is the demand trend.

Under this cost structure, the producer's problem is

$$\max_{H_t \geq 0} \beta \, \mathrm{E}_t \left(P_{t+1} H_t \, \mathrm{e}^{\eta_t + \varepsilon_{t+1}} \right) - \chi H_t^{1+1/\alpha_S} \, \mathrm{e}^{g_c t/\alpha_S + \omega_t} / (1+1/\alpha_S), \tag{S.1}$$

which leads to the following first-order condition

$$\beta \, \mathrm{E}_t \left(P_{t+1} \, \mathrm{e}^{\eta_t + \varepsilon_{t+1}} \right) = \chi H_t^{1/\alpha_S} \, \mathrm{e}^{g_c t/\alpha_S + \omega_t} \tag{S.2}$$

and an expression of H_t:

$$H_t = \left[(\beta/\chi) \, \mathrm{E}_t \left(P_{t+1} \, \mathrm{e}^{\eta_t + \varepsilon_{t+1}} \right) \right]^{\alpha_S} \mathrm{e}^{-g_c t - \alpha_S \omega_t}. \tag{S.3}$$

To determine the trend, we study the deterministic growth path, over which all shocks are equal to their zero average, and no stocks are held because of the absence of uncertainty. In this situation,

market equilibrium implies $H_{t-1} = D_t$, in which we substitute the expression of H obtained above:

$$[(\beta/\chi)P_t]^{\alpha_S} \mathrm{e}^{-g_c(t-1)} = \pi P_t^{\alpha_D} \mathrm{e}^{g_d t}. \tag{S.4}$$

This leads to

$$P_t = \left[\pi(\chi/\beta)^{\alpha_S}\right]^{1/(\alpha_S-\alpha_D)} \mathrm{e}^{[g_d t + g_c(t-1)]/(\alpha_S-\alpha_D)}, \tag{S.5}$$

from which we can identify the price trend as the terms proportional to time in the exponential:

$$g_p = \frac{g_d + g_c}{\alpha_S - \alpha_D}. \tag{S.6}$$

Equation (S.6) shows that the equilibrium price trend emerges from any difference between the demand and cost trends (considering that $g_c < 0$). If demand increases faster than costs decrease, then there is a positive price trend, and conversely.

We replace the expression of P_t in the definition of final demand and obtain

$$D_t = \pi\left[\pi(\chi/\beta)^{\alpha_S}\right]^{\alpha_D/(\alpha_S-\alpha_D)} \mathrm{e}^{\alpha_D g_p t + g_d t - g_c \alpha_D/(\alpha_S-\alpha_D)}. \tag{S.7}$$

As before, we can identify the equilibrium quantity trend as the terms proportional to time in the exponential:

$$g_q = \frac{\alpha_S g_d + \alpha_D g_c}{\alpha_S - \alpha_D}. \tag{S.8}$$

Likewise, one can check that H_t follows the same trend. Obviously, a positive quantity trend arises from a positive demand trend or a negative cost trend.

We can also solve equations (S.6) and (S.8) for g_c and g_d, which gives

$$g_c = \alpha_S g_p - g_q \text{ and } g_d = g_q - \alpha_D g_p. \tag{S.9}$$

These equations imply that the equilibrium trends, g_q and g_p, that are used in the paper can be obtained by combining the structural trends and the elasticities. However, one benefit of working with equilibrium rather than structural trends is that the former can be estimated directly, contrary to the latter.

We obtain $g_c = -2.71\%$ and $g_d = 2.4\%$ based on our trend and preferred elasticity estimates.

B Numerical methods

B.1 Algorithm

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

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

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

highest probability weights.

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

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

Step 1. Initialization step. Choose

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

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

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

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

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

Step 2.4. Use the policy rules, $\mathscr{P}^n$ and $\mathscr{Q}^n$, to adjust the response variables to the new grid

$\{\tilde{s}, \varphi, \mu\}$:

$$p^{n-1} = \mathscr{P}^n(\tilde{s}, \varphi, \mu), \tag{S.10}$$

$$q^{e,m,n-1} = \mathscr{Q}^n(\tilde{s}, \varphi, \mu), \tag{S.11}$$

$$x^{m,n-1} = \max(0, (\tilde{s} - d(p^{n-1}))\,\mathrm{e}^{\mu}). \tag{S.12}$$

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

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

$$p^{+} = \mathscr{P}^n\left(\left(x^{m-1,n}\,\mathrm{e}^{-g_q} + q^{e,m-1,n}\,\mathrm{e}^{\varepsilon^{+}}\right)\mathrm{e}^{-\mu^{+}}, \varphi^{+}, \mu^{+}\right). \tag{S.13}$$

Step 3.2. Fixed-point iteration with damping:

$$q^{e,m,n} = (1-\lambda)\,q^{e,m-1,n} + \lambda\,\mathrm{e}^{\varphi}\left[\mathrm{E}\left(p^{+}\,\mathrm{e}^{\varepsilon^{+}}\right)\right]^{\alpha_S}, \tag{S.14}$$

$$x^{m,n} = (1-\lambda)\,x^{m-1,n} + \lambda\max\left(0, \tilde{s}\,\mathrm{e}^{\mu} - d\left(\beta\,\mathrm{e}^{g_p}\,\mathrm{E}\left(p^{+}\right) - k\right)\mathrm{e}^{\mu}\right). \tag{S.15}$$

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

Step 4. Approximation step. Calculate prices as

$$p^{n} = d^{-1}\left(\tilde{s} - x^{m,n}\,\mathrm{e}^{-\mu}\right), \tag{S.16}$$

from which we update the price function

$$\mathscr{P}^{n+1}(\tilde{s}, \varphi, \mu) = p^{n}. \tag{S.17}$$

We also update the production function

$$\mathscr{Q}^{n+1}(\tilde{s}, \varphi, \mu) = q^{e,m,n}. \tag{S.18}$$

Step 5. Terminal step.

If $n = 1$ or $\|p^n - p^{n-1}\|_2 \geq \varpi$ or $\max(\|q^{e,m,n} - q^{e,m-1,n}\|_2, \|x^{m,n} - x^{m-1,n}\|_2) \geq \lambda\varpi$ then

increment n to $n+1$ and go to step 3.

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

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

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

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

B.2 Solution precision

Once a solution is obtained, its accuracy can be assessed by rewriting unit-free the equations (14) and (15) which give two measures of the Euler equation errors. Using a net availability and shocks series, $\{\tilde{s}_i, \varphi_i, \mu_i\}$, the storage and production equation errors can be assessed using (see Gouel, 2013, for details of the derivation of these measures for the storage model)

$$EE_i^x = 1 - \frac{d\left(\max\left(d^{-1}\left(\tilde{s}_i\right), \beta\, \mathrm{e}^{g_p}\, \mathrm{E}\,\mathscr{P}\left(\left[\mathscr{X}\left(\tilde{s}_i, \varphi_i, \mu_i\right) \mathrm{e}^{-g_q} + \mathscr{Q}\left(\tilde{s}_i, \varphi_i, \mu_i\right) \mathrm{e}^{\varepsilon}\right] \mathrm{e}^{-\mu}, \varphi, \mu\right) - k\right)\right) \mathrm{e}^{\mu_i}}{\tilde{s}_i\, \mathrm{e}^{\mu_i} - \mathscr{X}\left(\tilde{s}_i, \varphi_i, \mu_i\right)}, \tag{S.19}$$

$$EE_i^h = 1 - \frac{\mathrm{e}^{\varphi_i} \left\{\mathrm{E}\left[\mathscr{P}\left(\left[\mathscr{X}\left(\tilde{s}_i, \varphi_i, \mu_i\right) \mathrm{e}^{-g_q} + \mathscr{Q}\left(\tilde{s}_i, \varphi_i, \mu_i\right) \mathrm{e}^{\varepsilon}\right] \mathrm{e}^{-\mu}, \varphi, \mu\right) \mathrm{e}^{\varepsilon}\right]\right\}^{\alpha_S}}{\mathscr{Q}\left(\tilde{s}_i, \varphi_i, \mu_i\right)}. \tag{S.20}$$

To assess the precision of the algorithm, we simulate the model calibrated on our preferred estimation (4-knot spline in Table 6). We then sample 1,000 points from the ergodic distribution

and use them to calculate the Euler equation errors defined above. Table S.1 presents the average and the maximum errors expressed in base-10 logarithm. The accuracy in both equations is similar. At about -2, maximum errors involve a \$1 error every \$100 consumption or production decisions. However, such high error rates are rare and are located close to cutoff situations of no storage. The average errors involve less than \$1 error every \$1,000 decisions, and are closer to \$1 error for every \$10,000 decisions. This is a satisfactory level of precision for this type of model, and as the Monte Carlo experiments show is sufficiently high for our estimation procedure to recover the true parameter values if the model is well specified.

Table S.1: Euler equations error ($\log_{10}|EE|$)

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

Notes: Calculated over 1,000 simulations from the model's ergodic distribution. The model parameters are from our preferred estimation (4-knot spline in Table 6).

C Monte Carlo experiments

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

All the experiments are based on 500 replications and most use the sample size $T = 56$ which corresponds to the actual length of the observed dataset used for inference, but longer samples are also used to analyze the asymptotic properties of our estimation approach. The model parameters chosen for the experiments are based on the estimates in section 5, except for σ_ω for which different values are chosen to illustrate some of the difficulties that can be encountered. The parameter values used are $\beta = 0.98$, $g_q = 2.5\%$, $g_p = -2\%$, $\rho_\mu = 0.5$, $\rho_{\eta,\omega} = -0.4$, $\sigma_\eta = 1.5\%$, $\sigma_\varepsilon = 2\%$, $\sigma_\upsilon = 1.6\%$, $k = 4\%$, $\alpha_D = -0.07$, and $\alpha_S = 0.08$. Since the cost shock, ω, is a crucial and unobserved determinant of the omitted variable bias in the supply equation, we run the Monte Carlo

experiments for three values of its standard deviation: $\sigma_\omega = \{0.05, 0.1, 0.2\}$, the latter being close to the value estimated in section 5. For these parameters, the proportion of stockouts is 17%, which indicates a setting with regular price spikes and important nonlinearities. For the indirect inference, the optimization for each replication starts from a different vector θ with values drawn randomly from continuous uniform distributions defined over the intervals extending 20% below and above the true values.

The results of the OLS and 2SLS approaches are reported in table S.2 panels A and B, and the results for the indirect inference approach are presented in table S.3. The results for the longer samples are contained in tables S.4–S.5. These tables show that, for the parameters common to both methods, the indirect inference approach is more precise than the OLS or 2SLS approaches, as evidenced by the lower root mean squared errors (RMSEs) obtained in either small or large samples. Note also that estimates of the demand and supply elasticities exploiting the indirect inference approach are not biased in either the small or the large samples which contrasts with the OLS estimates on which they are based.[1] This result confirms that in the context of indirect inference the elasticities are not just set equal to their OLS counterparts. More precisely, the approach relies on the information derived from b_c^{OLS} and c_q^{OLS} in combination with the other parameters, and delivers unbiased and consistent elasticity estimates.

Nevertheless, using the indirect inference approach, two parameters are difficult to estimate: the storage cost (k) and the correlation between the planting-time shocks ($\rho_{\eta,\omega}$). For the storage cost, the limited precision of the estimation of storage cost could stem from the fact that it is identified only indirectly, in part through its effect on the autocorrelation and volatility of prices, two moments which contribute also to the identification of other parameters. The full information approaches in Cafiero et al. (2015) and Gouel and Legrand (2017) provide lower RMSE for their storage cost parameter. However, in our context these approaches are not feasible given that they require observability of the planting-time shocks.

The parameter $\rho_{\eta,\omega}$ is estimated based on its effect on the auxiliary parameter c_q^{OLS} (see equation (37)). However, what matters for estimating $\rho_{\eta,\omega}$ is $c_q^{\text{OLS}} - 1 = \alpha_S \sigma_\eta (\sigma_\eta - \rho_{\eta,\omega}\sigma_\omega)/\sigma_\psi^2$ and this is not precisely estimated in the auxiliary model (table S.2). The estimates of $\rho_{\eta,\omega}$ will be affected not only by the uncertainty related to the estimates of $c_q^{\text{OLS}} - 1$, but also by the uncertainty related to the other parameter estimates which explains its high RMSE. However, the challenges related to estimating $\rho_{\eta,\omega}$ are of secondary importance. In section 2.4, this parameter is absent from the rational expectations problem expressed in compact form. The equilibrium price, expected price,

[1] In the case of the supply elasticity, the inconsistency caused by using the expected price to substitute for the true incentive price can be evaluated by using $\mathrm{E}_{t-1}(p_t \exp(\varepsilon_t))$ instead of $\mathrm{E}_{t-1}\, p_t$ in a large-sample setting. The bias caused by using the proxy variable is -0.7%.

Table S.2: Monte Carlo experiment with OLS and 2SLS estimations of the supply and demand equations

	ρ_μ	$c_q - 1$	σ_ψ (%)	σ_ϑ (%)	σ_υ (%)	α_D	α_S
Panel A. OLS							
$\sigma_\omega = 5\%$							
Mean	0.36	0.051	2.49	2.65	1.28	−0.020	0.068
St. dev.	0.13	0.023	0.23	0.25	0.13	0.012	0.005
RMSE (%)	37.88	41.783	9.40	9.62	21.76	74.103	16.873
SE	0.14	0.024	0.24		0.13	0.012	0.005
$\sigma_\omega = 10\%$							
Mean	0.37	0.068	2.49	2.76	1.29	−0.021	0.054
St. dev.	0.13	0.043	0.23	0.27	0.14	0.012	0.010
RMSE (%)	36.55	53.805	9.40	10.26	21.07	72.104	34.465
SE	0.14	0.045	0.24		0.13	0.011	0.009
$\sigma_\omega = 20\%$							
Mean	0.39	0.083	2.49	3.02	1.33	−0.025	0.020
St. dev.	0.12	0.078	0.23	0.31	0.14	0.011	0.016
RMSE (%)	33.71	69.192	9.40	13.53	19.18	66.852	77.165
SE	0.13	0.083	0.24		0.13	0.011	0.015
Panel B. 2SLS							
σ_ω =5%, {Supply: E(F) =15, E(p-value) =0.35}, {Demand: E(F) =22, E(p-value) =0.01}							
Mean	0.50	0.069	2.49	2.70	1.65	−0.073	0.080
St. dev.	0.18	0.031	0.23	0.26	0.36	0.026	0.014
RMSE (%)	36.21	45.798	9.40	9.75	22.91	37.747	17.786
SE	0.19	0.034	0.24		0.29	0.023	0.015
σ_ω =10%, {Supply: E(F) =14, E(p-value) =0.29}, {Demand: E(F) =21, E(p-value) =0.01}							
Mean	0.49	0.111	2.49	2.90	1.65	−0.073	0.082
St. dev.	0.18	0.061	0.23	0.31	0.36	0.026	0.027
RMSE (%)	35.37	57.722	9.40	10.84	22.70	37.272	34.118
SE	0.19	0.066	0.24		0.29	0.023	0.029
σ_ω =20%, {Supply: E(F) =12, E(p-value) =0.18}, {Demand: E(F) =20, E(p-value) =0.02}							
Mean	0.49	0.199	2.49	3.45	1.65	−0.073	0.088
St. dev.	0.17	0.135	0.23	0.60	0.36	0.026	0.059
RMSE (%)	33.65	74.389	9.40	18.14	22.74	37.201	74.947
SE	0.18	0.142	0.24		0.30	0.023	0.059

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

Table S.3: Monte Carlo experiment with indirect inference approach

	ρ_μ	$\rho_{\eta,\omega}$	σ_ω (%)	σ_η (%)	σ_ε (%)	σ_υ (%)	k (%)	α_D	α_S
$\sigma_\omega = 5\%$									
Mean	0.50	−0.44	5.06	1.47	1.98	1.61	4.12	−0.071	0.080
St. dev.	0.11	0.30	0.66	0.34	0.27	0.27	2.46	0.016	0.007
RMSE (%)	22.93	74.57	13.30	22.66	13.62	16.62	61.66	23.263	9.149
ASE	0.11	0.36	0.67	0.37	0.29	0.25	2.68	0.017	0.007
$\sigma_\omega = 10\%$									
Mean	0.50	−0.44	10.19	1.47	1.98	1.61	4.14	−0.071	0.081
St. dev.	0.11	0.28	1.62	0.35	0.28	0.26	2.49	0.016	0.014
RMSE (%)	22.75	70.42	16.31	23.61	14.23	16.14	62.26	23.114	17.192
ASE	0.11	0.32	1.55	0.38	0.30	0.24	2.71	0.017	0.014
$\sigma_\omega = 20\%$									
Mean	0.50	−0.43	21.09	1.48	1.97	1.62	4.18	−0.072	0.082
St. dev.	0.12	0.27	5.50	0.37	0.29	0.25	2.48	0.016	0.025
RMSE (%)	23.49	67.73	28.06	24.85	14.80	15.64	62.11	22.943	31.053
ASE	0.11	0.32	5.60	0.41	0.32	0.24	2.68	0.017	0.026

Notes: Monte Carlo experiment based on 500 replications, with a sample size $T = 56$. For $\sigma_\omega = 5\%$, 10%, and 20%, respectively 26, 26, and 25 replications had to be dropped due to converging to $k = 0$ for which it is not possible to calculate standard errors. True values: $\rho_\mu = 0.5$, $\rho_{\eta,\omega} = -0.4$, $\sigma_\eta = 1.5\%$, $\sigma_\varepsilon = 2\%$, $\sigma_\upsilon = 1.6\%$, $k = 4\%$, $\alpha_D = -0.07$, and $\alpha_S = 0.08$. The mean and standard deviations are respectively the average and standard deviations of the empirical parameter distribution. They are combined to calculate the RMSE expressed as a percentage of the true parameter value. ASE means asymptotic standard errors, based on equation (51), and represents the average standard errors calculated at the solutions.

demand, and production depend not on the specific value of the shock ω but rather on the aggregate shock φ. In a Monte Carlo experimental setting, it is possible to calculate the RMSE for σ_φ. At 18%—for $\sigma_\omega = 20\%$—this is similar to the RMSE for the other shocks. Overall, the empirical method seems to be appropriate for estimating the volatility of all the shocks, but the various errors will be compounded in $\rho_{\eta,\omega}$ which is difficult to estimate, although without consequences for the rest of the model.

Tables S.2–S.3 and S.4–S.5 show that both approaches have good asymptotic properties. The RMSE and their two components vanish "asymptotically"—i.e., as the sample length increases from 56, to 100, 200, and 1000—showing the consistency of both estimators (apart from a small bias in the supply elasticity discussed above).

The standard errors (rows SE in table S.2 and asymptotic standard errors (ASE) in table S.3) are similar to the standard deviations of the Monte Carlo estimates showing that for both methods the standard errors are consistent with the standard deviations in the population. The comparability of standard deviations and standard errors is an important result for two reasons. Reliable auxiliary model standard errors matter because in the indirect inference approach they directly enter the weighting matrix. Also, consistent indirect inference standard errors in the Monte Carlo analysis

suggests that the asymptotic formula we apply has a limited small-sample bias (table S.5 shows that with longer samples the biases are negligible).

Table S.4: Additional Monte Carlo experiments with instrumental variables for $\sigma_\omega = 20\%$

	ρ_μ	$c_q - 1$	σ_ψ (%)	σ_ϑ (%)	σ_υ (%)	α_D	α_S
T =100, {Supply: E(F) =20, E(p-value) =0.08}, {Demand: E(F) =34, E(p-value) =0.00}							
Mean	0.50	0.194	2.50	3.41	1.62	−0.071	0.084
St. dev.	0.12	0.090	0.17	0.36	0.22	0.016	0.035
RMSE (%)	23.94	49.689	6.90	10.80	13.83	22.616	43.931
SE	0.12	0.091	0.18		0.21	0.015	0.035
T =200, {Supply: E(F) =40, E(p-value) =0.02}, {Demand: E(F) =69, E(p-value) =0.00}							
Mean	0.50	0.188	2.50	3.38	1.62	−0.070	0.081
St. dev.	0.07	0.060	0.12	0.23	0.15	0.010	0.023
RMSE (%)	14.22	33.204	4.86	6.89	9.13	14.470	28.780
SE	0.08	0.062	0.13		0.14	0.010	0.023
T =1000, {Supply: E(F) =199, E(p-value) =0.00}, {Demand: E(F) =349, E(p-value) =0.00}							
Mean	0.50	0.181	2.50	3.35	1.60	−0.070	0.079
St. dev.	0.03	0.027	0.06	0.10	0.06	0.005	0.010
RMSE (%)	6.54	14.694	2.24	3.07	4.04	6.647	12.318
SE	0.04	0.026	0.06		0.06	0.004	0.009

Notes: See notes to table S.2.

We also ran Monte Carlo estimations for gradually increasing sizes of σ_ω to analyze its role in the parameter estimations. Table S.2 panels A and B show that an increase from 5 to 20% in σ_ω affects only the OLS and 2SLS performances for the supply-side parameters estimates. Varying σ_ω fleshes out the trade-off between consistency and precision in the supply elasticity estimates highlighted by Hendricks et al. (2015). What is gained in terms of reduced bias from using 2SLS is lost through higher volatility of the estimates, resulting in similar but lower RMSE for the OLS compared to the 2SLS. This is because a higher σ_ω implies a larger omitted variable bias but it also makes the lagged yield shocks a weaker instrument because their role in explaining price changes declines as the variance of cost shocks increases. For this choice of storage model parameters, deciding between estimating supply using OLS or 2SLS is difficult given that both approaches have some limitations. However, in the present context, as documented in table S.3, the indirect inference approach is much more robust to σ_ω with RMSEs that deteriorate less as this parameter increases.

In addition to the Monte Carlo results on the parameters, table S.2 displays some results about the 2SLS diagnostics statistics: the first-stage F-statistics and the p-value for the Hausman test of endogeneity. The F-statistics show that the instrument is much weaker for the supply equation than for the demand equation with a mean value closer to 10, and deteriorating with σ_ω. This is

Table S.5: Additional Monte Carlo experiments with indirect inference and $\sigma_\omega = 20\%$

	ρ_μ	$\rho_{\eta,\omega}$	σ_ω (%)	σ_η (%)	σ_ε (%)	σ_υ (%)	k (%)	α_D	α_S
$T = 100$									
Mean	0.50	−0.40	20.70	1.48	1.97	1.59	3.91	−0.070	0.080
St. dev.	0.08	0.19	3.50	0.28	0.23	0.17	1.79	0.011	0.017
RMSE (%)	16.76	46.43	17.87	19.02	11.81	10.49	44.78	16.365	21.505
ASE	0.08	0.21	3.88	0.27	0.22	0.18	2.06	0.012	0.019
$T = 200$									
Mean	0.50	−0.40	20.34	1.50	1.99	1.60	3.91	−0.070	0.080
St. dev.	0.06	0.14	2.48	0.19	0.16	0.13	1.38	0.008	0.013
RMSE (%)	11.20	35.11	12.52	12.44	7.96	8.03	34.58	11.785	15.740
ASE	0.06	0.13	2.53	0.18	0.15	0.12	1.53	0.008	0.013
$T = 1000$									
Mean	0.50	−0.40	20.09	1.50	1.99	1.59	3.84	−0.069	0.080
St. dev.	0.03	0.06	1.12	0.08	0.07	0.06	0.62	0.004	0.006
RMSE (%)	5.16	16.11	5.60	5.66	3.58	3.50	16.06	5.376	7.301
ASE	0.02	0.06	1.11	0.08	0.07	0.05	0.70	0.004	0.006

Notes: See notes to table S.3. For $T = 100$, 5 replications had to be dropped due to converging to $k = 0$ for which it is not possible to calculate standard errors.

consistent with the idea that a lagged yield, used to instrument the expected price in the supply equation, is a worse predictor of price than current yield is which is something that is also found on observations. The Hausman test shows that the null of exogenous expected prices is not rejected on average, in a context where we know that these prices are endogenous and should be instrumented. This result disappears in longer samples (table S.4) and is explained by the important standard errors that the supply elasticity displays in small samples.

D Additional results

D.1 Inspecting the auxiliary model

It is interesting to analyze the similarity between the estimates of the auxiliary model parameters based both on observations and simulations. Table S.6 reports the auxiliary parameters obtained respectively from the actual and the simulated data along with their standard errors estimated on the observations. Note that the standard errors column corresponds to the square root of the diagonal of $(1 + 1/\tau)\Omega^{-1}$. For each parameter we can calculate a t-statistic of equality of the coefficients and test for consistency of the auxiliary model (Gourieroux et al., 1993, Appendix 1). Apart from the parameters $b_{\mathrm{E}p}$ from equation (33) and d_p from equation (42), we cannot reject the null of equality

between the estimates based on observations and those based on simulations from the structurally estimated model.[2] Although some parameters differ a lot between the two columns (e.g., c_c or $c_{\mathrm{E}p}$), they are estimated imprecisely in the auxiliary model, and thus were given a small weight in the objective function which the indirect inference procedure minimizes.

Table S.6: Coefficients of the auxiliary model: estimation based on observations versus based on simulations

	Observations		Model
Coefficient	Estimate	Standard error	Estimate
b_q	0.058	0.013	0.048
c_q	1.103	0.087	1.147
σ_{u_q}	0.015	0.001	0.015
b_c	−0.021	0.013	−0.007
c_c	−0.005	0.011	0.011
d_c	0.547	0.119	0.534
σ_{u_c}	0.014	0.001	0.014
$b_{\mathrm{E}p}$	−3.783	0.864	−2.300
$c_{\mathrm{E}p}$	−2.382	0.948	−1.690
$\sigma_{u_{\mathrm{E}p}}$	0.165	0.022	0.160
b_p	−4.112	1.418	−4.441
c_p	0.486	0.103	0.456
d_p	−0.130	1.614	2.882
σ_{u_p}	0.180	0.028	0.180
σ_{u_ψ}	0.023	0.002	0.025
b_q^{2SLS}	0.075	0.026	0.086
b_c^{2SLS}	−0.065	0.026	−0.068

Notes: Standard errors equal to $\sqrt{(1+1/\tau)\operatorname{diag}\Omega^{-1}}$ for the upper panel. The lower panel presents the parameters estimated by 2SLS not present in the auxiliary model used for the estimation with standard errors robust to heteroskedasticity.

Although the auxiliary model used here involves only OLS estimations, it is useful to compare also the fit with the supply and demand elasticities estimated by 2SLS (see lower panel in table S.6). For both elasticities, t-tests would not reject the null hypothesis of no differences between estimates on observations and on simulations. However, we can note that the model tends to overestimate b_q^{2SLS} and to underestimate b_q estimated by OLS (albeit insignificantly for both). Since the difference between the supply elasticities estimated by OLS and 2SLS is supposed to increase with σ_ω, this

[2]The results would be similar if using for the standard errors the inverse of the square root of the diagonal of W, except that we would not reject the equality of $b_{\mathrm{E}p}$ between observations and simulations.

difference could confirm the possible overestimation of σ_ω highlighted in section 5.

These results suggest an overall good fit of the auxiliary model between observations and simulations, except for two parameters.

D.2 Sensitivity analyses

In this sensitivity analysis, we discuss the roles of data, the weighting matrix, and each moment in the estimation process.

D.2.1 Data

For the data, we have followed Roberts and Schlenker (2013) and considered the four most important crops for quantities but excluded rice from the price index because of the short sample of rice price futures. This could be a concern if the rice market behaves differently from the other markets. To verify this, we estimate the model on data from which the rice sector has been removed altogether. Similar results are obtained (table S.7) except that all shocks and elasticities are higher in absolute values (albeit not significantly). This could be explained by the fact that rice consumption and production are more stable than for the other crops, because of its almost exclusive use for food consumption and its large share of irrigated production which limits production shocks.

In addition to FAOSTAT, it is possible to obtain almost-global information about quantities from the USDA. The USDA Production, Supply and Distribution (PSD) database (USDA, 2020) provides information about a smaller sample of countries, which excludes some countries with minor contribution to the global food balance. Although it also allows using a longer sample, for comparability with FAOSTAT data, we maintain the same 1961–2017 sample. 2SLS and indirect inference results with USDA-PSD are available in table S.7. Results based on USDA data are extremely similar to those based on FAOSTAT data whatever the estimator. One noticeable difference is the more elastic demand. This difference could be due to differences in stock changes data, which would appear here as a difference in consumption, given that production is less susceptible to measurement errors. However, this different estimate does not help improve the model fit as studied in table 7.

We have carried out our estimations on a caloric aggregate following Roberts and Schlenker (2013), because it provided us with a 2SLS benchmark to compare our indirect inference estimates. However, aggregating commodities may create bias. In table S.8, we present the crop by crop results of 2SLS and indirect inference estimates for maize, soybeans, and wheat (the sample for rice is too short to obtain reliable results). At the commodity level, only the elasticities of maize are

Table S.7: Additional estimation results (detrending with 4 knots)

	FAOSTAT data without rice				USDA-PSD data			
	2SLS		II		2SLS		II	
	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE
g_p			−0.020				−0.020	
g_q			0.027				0.025	
ρ_μ	0.530	(0.158)	0.674	(0.087)	0.533	(0.226)	0.738	(0.071)
$\rho_{\eta,\omega}$			−0.396	(0.256)			−0.438	(0.344)
σ_ω			0.219	(0.042)			0.177	(0.027)
σ_η			0.021	(0.006)			0.013	(0.006)
σ_ε			0.024	(0.005)			0.020	(0.004)
σ_υ	0.021	(0.005)	0.024	(0.004)	0.018	(0.004)	0.022	(0.004)
k			0.032	(0.023)			0.038	(0.024)
α_D	−0.083	(0.035)	−0.086	(0.025)	−0.076	(0.033)	−0.089	(0.023)
α_S	0.088	(0.035)	0.096	(0.022)	0.086	(0.026)	0.079	(0.013)
σ_φ			0.037	(0.006)			0.024	(0.005)
σ_ψ	0.030	(0.003)	0.032	(0.003)	0.023	(0.002)	0.024	(0.002)
σ_μ	0.025	(0.006)	0.033	(0.008)	0.022	(0.006)	0.033	(0.008)
σ_ϑ	0.040		0.044	(0.005)	0.031		0.031	(0.003)

Notes: See notes to tables 4–6.

significant when estimated by 2SLS, but they are all precisely estimated with indirect inference, except for the demand elasticity of soybeans. Demand and supply are more elastic (except for the supply of wheat) at the commodity level, consistent with the idea that these crops are substitutes. The shocks tend also to be much larger since they are no longer smoothed by the aggregation. The conclusion that supply shocks are larger than demand shocks remain, but with a smaller difference between them. This could be explained by the fact that at the commodity level a supply shock for one commodity can become a demand shock for another: for example a bad wheat harvest could create increased demand for maize for feed.

D.2.2 Weighting matrix

Our indirect inference estimator uses for the weighting matrix the diagonal of the inverse of the covariance matrix of the auxiliary parameters. While this estimator is consistent for any positive definite matrix, its efficiency can be improved by using a weighting matrix closer to the optimal weighting matrix. With this in mind, we also implemented a two-step estimation with the weighting matrix in the second step being the model-restricted optimal covariance matrix defined in

Table S.8: Estimation results by commodity (detrending with 4 knots)

	Maize				Soybeans				Wheat			
	2SLS		II		2SLS		II		2SLS		II	
	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE
g_p			−0.020				−0.018				−0.021	
g_q			0.028				0.044				0.019	
ρ_μ	0.501	(0.192)	0.735	(0.075)	0.482	(0.195)	0.565	(0.115)	0.605	(0.197)	0.626	(0.089)
$\rho_{\eta,\omega}$			−0.961	(1.322)			0.201	(0.249)			−0.132	(0.212)
σ_ω			0.170	(0.030)			0.359	(0.125)			0.475	(0.163)
σ_η			0.013	(0.015)			0.025	(0.014)			0.035	(0.007)
σ_ε			0.039	(0.006)			0.040	(0.009)			0.024	(0.008)
σ_υ	0.028	(0.005)	0.034	(0.006)	0.048	(0.013)	0.058	(0.020)	0.029	(0.010)	0.036	(0.010)
k			0.048	(0.030)			0.017	(0.033)			0.059	(0.032)
α_D	−0.110	(0.031)	−0.131	(0.031)	−0.090	(0.111)	−0.165	(0.127)	−0.096	(0.074)	−0.124	(0.045)
α_S	0.162	(0.057)	0.165	(0.034)	0.226	(0.181)	0.171	(0.068)	0.060	(0.052)	0.064	(0.025)
σ_φ			0.043	(0.008)			0.063	(0.011)			0.051	(0.008)
σ_ψ	0.041	(0.004)	0.042	(0.004)	0.047	(0.004)	0.047	(0.005)	0.040	(0.004)	0.042	(0.004)
σ_μ	0.033	(0.007)	0.051	(0.011)	0.055	(0.016)	0.071	(0.029)	0.037	(0.014)	0.047	(0.015)
σ_ϑ	0.057		0.058	(0.007)	0.073		0.074	(0.009)	0.047		0.056	(0.007)

Notes: See notes to tables 4–6.

equation (52) and obtained from the first-step estimation based on the diagonal weighting matrix. The results are reported in table S.9. To ease comparison, the left columns reproduces the estimation results obtained using the preferred specification (i.e., table 6).

Table S.9: Sensitivity of the indirect inference estimation results to the weighting matrix (detrending with 4 knots)

	One-step		Two-step	
	Estimate	Standard error	Estimate	Standard error
ρ_μ	0.702	(0.081)	0.707	(0.080)
$\rho_{\eta,\omega}$	−0.443	(0.295)	−0.452	(0.280)
σ_ω	0.188	(0.031)	0.188	(0.031)
σ_η	0.014	(0.005)	0.015	(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)	0.035	(0.026)
α_D	−0.068	(0.019)	−0.064	(0.016)
α_S	0.086	(0.017)	0.088	(0.017)
σ_φ	0.027	(0.005)	0.028	(0.005)
σ_ψ	0.025	(0.002)	0.025	(0.003)
σ_μ	0.026	(0.006)	0.025	(0.005)
σ_ϑ	0.034	(0.004)	0.034	(0.004)

Notes: The columns "One–step" correspond to the benchmark results presented in table 6 based on a diagonal weighting matrix estimated on observations. The columns "Two-step" present indirect inference estimates using as weighting matrix the model-restricted weighting matrix defined in equation (52) and obtained from the simulations using the one-step estimation. See notes to table 6.

On the whole, the estimated parameters are very close, except for $\rho_{\eta,\omega}$, the covariance between the planting-time shocks. In addition, the parameters tend to be more precisely estimated, although the differences are very small. The lack of significant differences between the two sets of values confirms our choice of using the simpler and faster one-step approach.

D.2.3 Sensitivity to the estimating moments

Finally, to be transparent regarding the identification, we report Andrews et al.'s (2017) measure of sensitivity of the estimates to the estimating moments in table S.10. This measure is calculated as $\Lambda = -(J'WJ)^{-1}J'W$ and describes how estimated parameters change with the moments. To normalize this measure as elasticities of changes in parameters with respect to moments, we display $\text{diag}(\hat{\theta}^{-1})\Lambda\,\text{diag}(\hat{\zeta}_T)$, where diag transforms a vector in a diagonal matrix. Table S.10 presents this measure for our benchmark estimation. It confirms the intuitions about identification laid out in section 3.2, but it also shows that for most parameters, identification actually comes from a combination of moments. For example, the four moments to which the supply elasticity is the most sensitive are the three moments associated with the supply equation (b_q, c_q, and σ_{u_q}), with the expected signs, as well as $\sigma_{\mathrm{E}\,p}$ because more volatile expected prices would imply a lower supply elasticity.

Table S.10: Sensitivity of indirect inference estimates to estimation moments (detrending with 4 knots)

Coefficient	ρ_μ	$\rho_{\eta,\omega}$	σ_ω	σ_η	σ_ε	σ_υ	k	α_D	α_S
b_q	0.180	−0.065	−0.607	−0.038	0.035	0.070	0.199	0.011	0.546
c_q	0.044	6.282	−0.059	−0.113	0.040	−0.123	−0.688	−0.293	0.499
b_c	0.030	0.016	0.026	0.088	−0.012	0.002	0.324	0.175	−0.015
c_c	−0.017	0.026	0.009	−0.003	0.007	−0.013	−0.064	−0.048	−0.009
d_c	0.160	0.239	−0.034	0.096	−0.011	−0.135	−0.133	−0.071	0.082
b_p	−0.048	1.100	−0.134	−0.976	0.457	−0.160	0.426	−0.395	0.168
c_p	0.137	−0.107	0.217	−0.156	0.040	0.140	−0.302	0.371	−0.296
d_p	0.001	−0.004	−0.000	0.003	−0.001	−0.003	0.001	−0.004	−0.001
$b_{\mathrm{E}\,p}$	−0.041	0.753	0.012	−0.406	0.272	−0.201	−0.247	−0.418	0.020
$c_{\mathrm{E}\,p}$	−0.009	−0.210	0.024	0.439	−0.249	0.014	0.157	0.062	−0.006
σ_{u_q}	0.087	−0.555	0.409	−0.152	0.072	−0.094	−0.136	−0.227	0.601
σ_{u_c}	−0.216	0.892	0.030	−0.436	0.281	1.259	0.913	1.466	−0.034
σ_{u_p}	−0.237	−1.360	0.031	1.154	−0.583	0.173	2.205	−0.304	0.043
$\sigma_{u_{\mathrm{E}\,p}}$	0.490	0.056	0.597	−0.291	0.170	0.281	−0.050	0.270	−0.745
σ_{u_ψ}	−0.141	1.025	−0.084	0.613	1.053	−0.548	−0.857	−1.053	0.117

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

D.3 Historical decomposition

The model can be used to perform a historical decomposition, i.e., to extract the various shocks from the series.[3] This does not require indirect inference estimation per se. The linear regressions estimates would be sufficient as long as the residuals are given a structural interpretation as proposed in section 3.1. Figure S.1 depicts the shocks that are identified along with the log deviations of their price and quantity trends. Our preferred estimates are from section 5.1.2: the indirect inference estimates with 4-knot spline.[4]

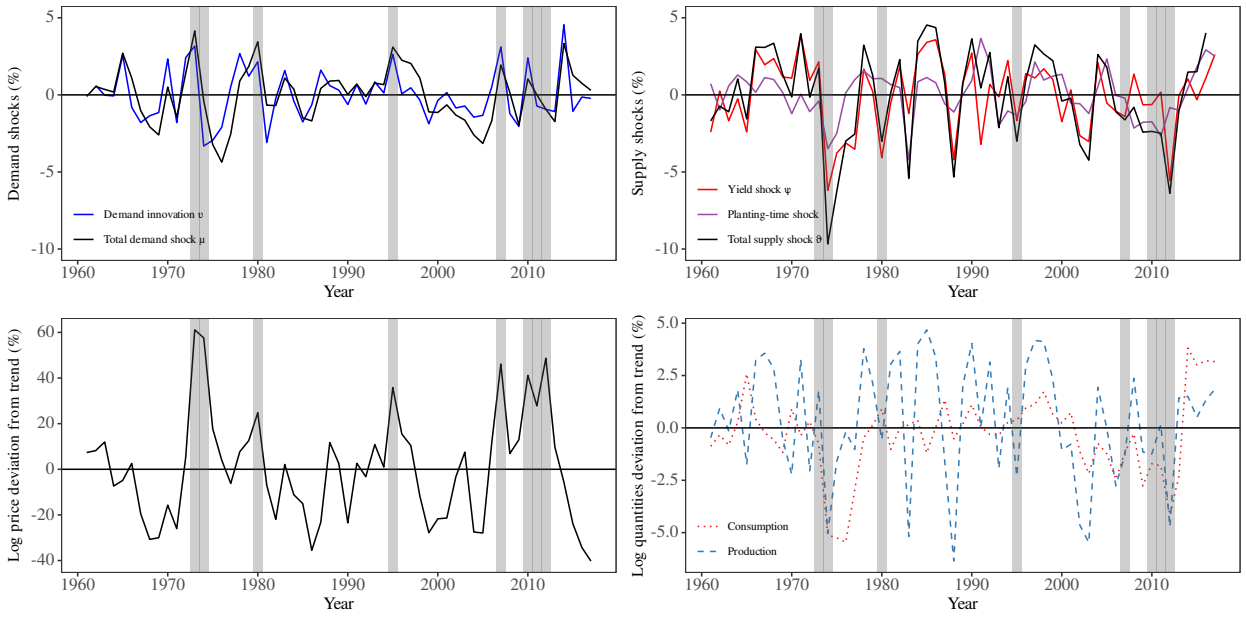

Figure S.1: Historical decomposition of the world price, production and consumption of grains into the various shocks. Gray areas denote price spike periods defined as log deviations from the trend greater than one standard deviation, 23.6%. The planting-time supply shock in purple corresponds to $\alpha_S(\eta_{t-1}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

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

First, in line with the estimated standard deviations of the shocks, supply disturbances are larger

[3] Bobenrieth et al. (2013) propose another kind of historical analysis by showing the consistency between observed and predicted stock levels. Although the results are encouraging, their approach requires an invertible model where it is possible to obtain all model variables from the observables. This is not the case here, where only a subset of variables can be recovered from the observables.

[4] Figure S.1 would nonetheless be very similar if created using instead the 2SLS estimates.

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

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

E Supplementary tables

Table S.11: Parameter bounds

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

TABLE S.12 Unit root tests results

Demand Price			Supply Price			Consumption			Production		
Panel A. LM unit root tests with quadratic trend and one or two trend breaks											
TB1	TB2	t-stat (lags)	TB1	TB2	t-stat (lags)	TB1	TB2	t-stat (lags)	TB1	TB2	t-stat (lags)
1979		$-5.17(1)^{**}$	1980		$-5.29(1)^{**}$	2006		$-3.97(1)$	2003		$-4.88(2)^{**}$
1971	2007	$-6.41(1)^{**}$	1972	2008	$-5.65(2)$	1984	2007	$-6.69(3)^{**}$	1982	2000	$-6.92(1)^{**}$
Panel B. Unit root tests on detrended data (t-stat)											
ADF	PP	KPSS	ADF	PP	KPSS	ADF	PP	KPSS	ADF	PP	KPSS
-2.87^{***}	-2.87^{***}	0.26^{***}	-2.59^{***}	-2.34	0.29^{***}	-2.32^{***}	-2.58	0.51^{***}	-2.6^{***}	-3.12^{***}	0.49^{***}
-3.15^{***}	-3.13^{***}	0.18^{***}	-2.87^{***}	-2.57	0.21^{***}	-1.47	-1.45	0.3^{***}	-2.3^{***}	-3.22^{***}	0.29^{***}
-3.91^{***}	-3.66^{***}	0.05	-3.58^{***}	-3.09^{***}	0.06	-3.3^{***}	-3.27^{***}	0.09	-5.38^{***}	-6.95^{***}	0.04
-4.13^{***}	-3.84^{***}	0.05	-3.68^{***}	-3.18^{***}	0.06	-4.3^{***}	-4.15^{***}	0.04	-5.82^{***}	-7.33^{***}	0.02

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

Table S.13: Supply equation estimation, 1962–2007

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

Notes: Standard errors robust to heteroskedasticity in parenthesis. The knots are placed following Roberts and Schlenker (2013): 1963, 1984, and 2005 for 3 knots; 1962, 1976, 1992, and 2006 for 4 knots; and 1962, 1973, 1984, 1995, and 2006 for 5 knots.

Table S.14: Demand equation estimation, 1962–2007

	(1)	(2)	(3)
Panel A. 2SLS			
Demand elasticity b_c	−0.034	−0.045	−0.047
	(0.033)	(0.032)	(0.032)
Lagged price c_c	0.022	0.012	0.011
	(0.023)	(0.020)	(0.021)
Lagged demand d_c	0.977	0.768	0.633
	(0.228)	(0.217)	(0.242)
Panel B. First stage			
Shock b_p	−3.926	−3.743	−3.819
	(0.939)	(1.056)	(1.041)
Lagged price c_p	0.599	0.532	0.547
	(0.146)	(0.149)	(0.152)
Lagged demand d_p	4.960	3.644	4.403
	(2.034)	(2.007)	(2.446)
Panel C. OLS			
Demand elasticity b_c	0.000	−0.007	−0.005
	(0.012)	(0.011)	(0.010)
Lagged price c_c	−0.003	−0.012	−0.016
	(0.016)	(0.015)	(0.015)
Lagged demand d_c	0.743	0.593	0.433
	(0.131)	(0.152)	(0.174)
Panel D. 2SLS using Roberts and Schlenker's approach (eqs. (39) for 2nd stage and (43) for 1st)			
Demand elasticity b_c	−0.033	−0.062	−0.059
	(0.023)	(0.031)	(0.028)
$\sigma_{u_c^{\text{2SLS}}}$	0.015	0.015	0.015
σ_{u_P}	0.175	0.175	0.175
$\sigma_{u_c^{\text{OLS}}}$	0.014	0.013	0.012
$\sigma_{u_c^{\text{2SLS, RS}}}$	0.022	0.020	0.018
$\sigma_{\mu^{\text{2SLS}}}$	0.072	0.023	0.019
First stage F-stat (panel A)	17.479	12.573	13.461
p-value for Hausman test (panel A)	0.274	0.221	0.117
p-value for Cumby–Huizinga test (panel A)	0.484	0.335	0.113
First stage F-stat (panel D)	20.780	17.059	17.156
p-value for Hausman test (panel D)	0.008	0.052	0.033
p-value for Cumby–Huizinga test (panel D)	0.000	0.035	0.067
Observations	46	46	46
Spline knots	3	4	5

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

References

Andrews, I., Gentzkow, M. and Shapiro, J. M. (2017). Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics*, 132(4), 1553–1592.

Bobenrieth, E. S. A., Wright, B. D. and Zeng, D. (2013). Stocks-to-use ratios and prices as indicators of vulnerability to spikes in global cereal markets. *Agricultural Economics*, 44(s1), 43–52.

Cafiero, C., Bobenrieth, E. S. A., Bobenrieth, J. R. A. and Wright, B. D. (2015). Maximum Likelihood estimation of the standard commodity storage model: Evidence from sugar prices. *American Journal of Agricultural Economics*, 97(1), 122–136.

Gouel, C. (2013). Comparing numerical methods for solving the competitive storage model. *Computational Economics*, 41(2), 267–295.

Gouel, C. and Legrand, N. (2017). Estimating the competitive storage model with trending commodity prices. *Journal of Applied Econometrics*, 32(4), 744–763.

Gourieroux, C., Monfort, A. and Renault, E. (1993). Indirect inference. *Journal of Applied Econometrics*, 8(S1), S85–S118.

Heiss, F. and Winschel, V. (2008). Likelihood approximation by numerical integration on sparse grids. *Journal of Econometrics*, 144(1), 62–80.

Hendricks, N. P., Janzen, J. P. and Smith, A. (2015). Futures prices in supply analysis: Are instrumental variables necessary? *American Journal of Agricultural Economics*, 97(1), 22–39.

Maliar, L. and Maliar, S. (2014). Numerical methods for large-scale dynamic economic models. In K. Schmedders and K. L. Judd (eds.) *Handbook of Computational Economics*, volume 3, chapter 7, (pp. 325–477). Amsterdam: Elsevier.

Michaelides, A. and Ng, S. (2000). Estimating the rational expectations model of speculative storage: A Monte Carlo comparison of three simulation estimators. *Journal of Econometrics*, 96(2), 231–266.

Orcutt, G. H. and Winokur, H. S. (1969). First order autoregression: Inference, estimation, and prediction. *Econometrica*, 37(1), 1–14.

Roberts, M. J. and Schlenker, W. (2013). Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate. *The American Economic Review*, 103(6), 2265–2295.

USDA (2020). United States Department of Agriculture, Production, Supply and Distribution database. https://www.fas.usda.gov/data, version of February, 2020.

Wright, B. D. (2014). Global biofuels: Key to the puzzle of grain market behavior. *Journal of*

Economic Perspectives, 28(1), 73–98.