

FUTURES PRICES IN SUPPLY ANALYSIS: ARE INSTRUMENTAL VARIABLES NECESSARY?

NATHAN P. HENDRICKS, JOSEPH P. JANZEN, AND AARON SMITH

Crop yield shocks are partially predictable—high planting-time futures prices have tended to indicate that yield would be below trend. As a result, regressions of total caloric production on futures prices produce estimates of the supply elasticity that are biased downwards by up to 75%. Regressions of the world's growing area on futures prices have a much smaller bias of about 20% because although yield shocks are partially predictable, this predictability has a relatively small effect on land allocation. We argue that the preferred method for estimating the crop supply elasticity is to use regressions of growing area on futures prices and to include the realized yield shock as a control variable. An alternative method for bias reduction is to use instrumental variables (IVs). We show that the marginal contribution of an IV to bias reduction is small—IVs are not necessary for futures prices in supply analysis.

Key words: Endogeneity, futures prices, supply response.

JEL codes: Q11.

Since the pioneering work of [Wright \(1928\)](#), weather has been recognized as an exogenous supply shifter that can be used to identify the price elasticity of agricultural demand. Using demand shifters to identify the agricultural supply elasticity has been less common—perhaps because demand shifters can explain little of the variation in agricultural prices between years. Instead, agricultural supply analysis has typically used a planting-time measure of the price that farmers expect to receive and has assumed that these expected prices are exogenous to supply. In his seminal work, [Nerlove \(1958\)](#) suggests including a lagged price and a lagged dependent variable as explanatory variables based on a supply model with adaptive expectations. Since [Gardner \(1976\)](#), it has become standard practice to use planting-time (or pre-planting) futures prices

of contracts for post-harvest delivery in econometric models of supply based on the argument that the futures price equals the farmer's rational expectation of post-harvest prices.

The plausible exogeneity of the futures price received little scrutiny until recent work by [Roberts and Schlenker \(2013\)](#), who argue that futures markets incorporate production expectations and thereby make the futures price endogenous to supply. These authors propose using weather shocks in the previous year as an instrumental variable (IV) on the grounds that past weather shocks affect the futures price through an inventory channel. Moreover, they approximate weather shocks using deviations of yield from trend and apply their model to the world's supply of calories incurred from maize, soybeans, wheat, and rice from 1961–2007; their results indicate substantial endogeneity bias. Roberts and Schlenker estimate supply elasticities in the range of 0.020–0.051 using ordinary least squares (OLS), but obtain estimates 2–4 times larger using two-stage least squares (2SLS).

This article's main contribution is to show that, as applied to world caloric supply, using IVs is unnecessary. Our result is in contrast to [Roberts and Schlenker \(2013\)](#), who argue that IVs reduce substantial endogeneity bias

Nathan P. Hendricks is an assistant professor in the Department of Agricultural Economics at Kansas State University. Joseph P. Janzen is an assistant professor in the Department of Agricultural Economics and Economics at Montana State University. Aaron Smith is a professor in the Department of Agricultural and Resource Economics at the University of California, Davis. The authors thank the editor, anonymous reviewers, and seminar participants at the AAEA Annual Meeting for helpful comments. The authors are also thankful for excellent research assistance provided by Kayode Ajewole and Krishna Pokharel. Correspondence may be sent to: nph@ksu.edu.

Amer. J. Agr. Econ. 97(1): 22–39; doi: 10.1093/ajae/aau062

Published online July 23, 2014

© The Author (2014). Published by Oxford University Press on behalf of the Agricultural and Applied Economics Association. All rights reserved. For permissions, please e-mail: journals.permissions@oup.com

in supply models. Roberts and Schlenker compare an OLS regression of total production on the futures price to a 2SLS regression that includes the current-year realized yield shock as a control variable and the prior-year yield shock as an instrument. We replicate and extend these authors' analysis and show that the control variable does all the work. When we estimate their model using OLS and include the current-year yield shock as a control variable, we obtain almost identical results to those from 2SLS. Using IVs is mostly innocuous in this context, although it results in a much wider confidence interval about the estimated supply elasticity than the OLS estimate.

To explain this finding, we decompose the supply elasticity into parts due to changes in (a) total growing area, (b) the composition of growing area across crops and countries, and (c) deviations of yield from trend (i.e., yield shocks). Applying this decomposition to the data reveals that the futures price is endogenous to total production due to a predictable component in yield shocks that is negatively correlated with the futures price.

When included as a control variable in a model of total production, the current-year realized yield shock serves as a proxy for the omitted (and unobservable) expected yield shock. But if the realized yield shock is included on the right-hand-side of the regression, it cancels out the component of production that comes from yield shocks. After removing yield shocks from production, we are left with growing area and trend yield. Thus, by construction the only source of supply response to the futures price is through changes in the magnitude and composition of growing area. Therefore, the supply elasticity may be best estimated using growing area rather than production as the dependent variable. Using data from Roberts and Schlenker (2013), we find that regressions of growing area on the futures price give a slightly larger supply elasticity (about 20% larger) when the yield shock is included as a control and the difference in the estimates is statistically significant at the 10% level. These results indicate a small endogeneity bias in regressions of growing area on futures prices, which may be reduced by including the realized yield shock as a control.

Our results resolve an apparent contrast between Roberts and Schlenker, who model production response, and Choi and

Helmberger (1993), who raise concerns about the endogeneity of the futures price in a model of acreage response. Choi and Helmberger estimate a system of equations for consumption demand, inventory demand, the futures price, acreage, and demand for seed, which they apply to the United States soybean market. In contrast to Roberts and Schlenker, Choi and Helmberger (1993) find little difference in OLS and three-stage least squares estimates of acreage response to price.

We suggest that the preferred method for estimating agricultural supply elasticities is to add the realized yield shock to regressions of growing area on futures prices to reduce any potential endogeneity bias. Some caveats are worth noting. First, the realized yield shock is a noisy proxy for the *predictable* component of the yield shock, which is the omitted variable in the regression. Thus, the control does not eliminate the bias (Wooldridge 2002) and, because it reduces the amount of noise in the model, it could magnify the importance of the predictable component of the yield shock in the model error, thereby increasing bias. Second, our results assume that yield deviations from trend are not affected by futures prices. Even if farmers do not adjust their input use during the growing season in response to prices, there are several ways that yield shocks are affected by prices. For example, farmers may expand growing area to land with lower productivity or may increase the growing area of a crop by changing from rotating crops to continuous cropping, which generally decreases yields. In such cases, it may be preferable to include weather or pest variables that are observed prior to planting as control variables.

Our decomposition of the supply elasticity is also useful for understanding methods to estimate the aggregate supply elasticity. Even if yield shocks are not affected by prices, then estimating an aggregate regression of growing area on futures prices does not represent the total supply elasticity because the growing area response may be spatially correlated with trend yields. We show that the total supply elasticity can be correctly estimated from a regression that uses aggregate data and in which the dependent variable is the total growing area times the area weighted average trend yield across regions. It can also be correctly estimated as the production weighted average of disaggregate estimates of area response to price.

Econometric Modeling of Global Agricultural Supply

We base our analysis on the model specification in Roberts and Schlenker (2013), who study the world production of calories from maize, soybeans, wheat, and rice using a simple regression equation relating global agricultural supply to prices. These authors regress the log of total caloric production on the log of a calorie-weighted index of United States futures prices while controlling for deterministic trends. This specification and weighting scheme is appropriate for some purposes but not for others. An alternative strategy is to weight by a value measure other than calories to reflect the fact that wheat and rice are more readily transformed into desirable human food than maize or soybeans. Other studies may analyze individual commodities or the supply in individual countries (e.g., Hausman 2012; Hendricks, Smith, and Sumner 2014). Our framework is pertinent to all of these cases, each of which requires aggregating heterogeneous quantities using an implicit or explicit weighting scheme. Even studies of the supply of a single grain commodity in a single country require aggregation over heterogeneous growing conditions and heterogeneous output quality.

A myriad of complexity underlies the Roberts and Schlenker supply equation. In response to a change in the futures price index, farmers throughout the world change the number of acres planted to these crops and the mix of crops within those acres. When making these decisions, farmers account for the heterogeneous productivity of the land they operate and predictions about growing season temperature and rainfall. Moreover, planting and harvest times vary across countries, so farmers may have some information about the likely size of the crop in other parts of the world when making these decisions.

To understand how these factors affect global agricultural supply response, we decompose world caloric production in year t into three components: (a) total growing area (A_t), (b) average trend caloric production per unit of land (trend yield; Y_t), and (c) the average proportional deviation from trend yield (Ψ_t); Y_t and Ψ_t are weighted averages of their country-crop counterparts. For crop c in country i in year t , we write the growing area as A_{cit} , the trend yield as Y_{cit} , and the proportional deviation from trend yield

as Ψ_{cit} . Caloric production from crop c in country i is $Q_{cit} = A_{cit}\kappa_c Y_{cit}\Psi_{cit}$, where κ_c denotes the number of calories in one unit of crop c , and world caloric production is $Q_t = \sum_i \sum_c A_{cit}\kappa_c Y_{cit}\Psi_{cit}$. We decompose world caloric production as $Q_t = A_t Y_t \Psi_t$, where

$$(1) \quad A_t = \sum_i \sum_c A_{cit}$$

$$(2) \quad Y_t = \frac{\sum_i \sum_c A_{cit}\kappa_c Y_{cit}}{\sum_i \sum_c A_{cit}}$$

$$(3) \quad \Psi_t = \frac{\sum_i \sum_c A_{cit}\kappa_c Y_{cit}\Psi_{cit}}{\sum_i \sum_c A_{cit}\kappa_c Y_{cit}}.$$

Using the notation that lower case objects represent the logarithm of upper case objects, in the remainder of the article we work with the decomposition $q_t = a_t + y_t + \psi_t$, where $q_t \equiv \ln(Q_t)$, $a_t \equiv \ln(A_t)$, $y_t \equiv \ln(Y_t)$, and $\psi_t \equiv \ln(\Psi_t)$.

Farmers' crop acreage allocation decisions affect q_t through changes in both a_t and y_t , the first two components of our decomposition. We define the country-level trend yield for each crop (Y_{cit}) as a deterministic function of time (as in Roberts and Schlenker 2013). Even though trend yield for a particular crop in a particular country does not depend on price, log average trend yield (y_t) may be affected by price if the spatial variation in acreage response to price is correlated with the spatial variation in yields. For example, if countries with more productive climates and soils have a larger growing area response to price, then average trend yield increases when price increases. Similarly, if countries that specialize in higher yielding crops have a larger growing area response to price, then average trend yield increases when price increases.

The log average yield shock (ψ_t), the third component in our decomposition of q_t , is almost identical to the yield shock variable (ω_t) in the supply equation of Roberts and Schlenker. These authors construct ω_t as the average across countries of log yield deviations from trend, whereas we define ψ_t as the log of average yield deviations from trend. The difference between ω_t and ψ_t does not matter much empirically. The correlation between the two variables in our dataset is 0.997, and regression estimates of the supply elasticity are essentially identical using either variable. We use ψ_t because it allows

for the decomposition of global caloric production in equations (1)–(3) to analyze the channels through which price affects caloric production.¹

Log yield shocks (ψ_{cit}) are determined mostly by weather. An exceptionally hot and/or dry growing season causes yield to be far below trend. Roberts and Schlenker assume the yield shock is independent of growing area (a_{cit}) and trend yields (y_{cit}), so the average deviation from trend yield (ψ_t) is not affected by price. This assumption would fail if farmers respond to output price shocks by changing inputs such as fertilizer and labor in ways that affect yield. The assumption would also fail if shocks to output prices cause expansion onto cropland of different-than-average quality, thereby causing country-level yield to deviate from trend. Roberts and Schlenker provide some evidence to suggest that any yield response to price is likely negligible compared to yield variation due to weather.

With the exogeneity of ψ_{cit} , aggregate production responds to price only through changes in the amount and composition of growing area. Farmers allocate land to crops based on expectations of prices and yields at harvest, so these two expectations affect total growing area directly, and they affect trend yield through their effects on the composition of growing area. We write linear regression equations for a_t and y_t as follows:

$$(4) \quad a_t = \alpha^a + \beta^a p_{\tau t} + \gamma^a \psi_{\tau t} + f^a(t) + u_t^a$$

$$(5) \quad y_t = \alpha^y + \beta^y p_{\tau t} + \gamma^y \psi_{\tau t} + f^y(t) + u_t^y$$

where $p_{\tau t} \equiv \ln(E_{\tau}[p_t])$, $\psi_{\tau t} \equiv \ln(E_{\tau}[\psi_t])$, E_{τ} denotes expectations conditional on information available at planting time, and $f(t)$ is a flexible time trend. Thus, $p_{\tau t}$ is the log of the expected harvest price and $\psi_{\tau t}$ is the log of the expected yield shock at time τ . Following Roberts and Schlenker (2013), we assume that the futures market provides an unbiased expectation of harvest prices, thus $p_{\tau t}$ denotes the log of a calorie-weighted average of harvest-time futures contracts trading at time τ . If yield shocks are not foreseeable, then $\psi_{\tau t} = 0$ and the associated term drops out of equations (4) and (5).

Total log production equals the sum of the logs of the three components in equations

(1)–(3), that is, $q_t = a_t + y_t + \psi_t$. Exogeneity of the yield shock combined with equations (4) and (5) implies that a model for world caloric supply is:

$$(6) \quad q_t = \alpha + \beta p_{\tau t} + \gamma \psi_{\tau t} + \psi_t + f(t) + u_t$$

where $\beta = \beta^a + \beta^y$ and similarly for other parameters of the supply equation.

The supply elasticity parameter (β) in equation (6) represents global agricultural supply response to exogenous price shocks. This parameter can be estimated consistently by OLS if two conditions hold: (a) the supply shocks embedded in u_t are unanticipated by the futures market ($E_{\tau}[u_t] = 0$), and (b) predictable yield shocks do not affect planting decisions ($\psi_{\tau t} = 0$ or $\gamma = 0$), or a suitable proxy exists for the predictable component of yield shocks. Neither of these conditions are likely to hold exactly. In Roberts and Schlenker, $p_{\tau t}$ is the futures price one year prior to the realization of harvested production in year t . Southern Hemisphere planted acreage is known at time τ , as is the Northern Hemisphere planted acreage of winter wheat. Therefore, the supply shocks embedded in u_t are partially predictable. We show in this article that yield shocks are also partially predictable ($\psi_{\tau t} \neq 0$), although they have only small effects on growing area (i.e., γ is small).

Bias of OLS and IV Estimators

Instrumental variables may overcome the OLS bias implied by the predictability of u_t and ψ_t . Roberts and Schlenker (2013) argue that the previous year's yield shock affects the futures price through interannual storage but should not directly affect production decisions. These authors suggest using the lagged yield shock as an instrument for the futures price. An alternative, or complementary, approach would be to use the observed yield shock ψ_t to proxy for the unobservable $\psi_{\tau t}$.

We thus have three possible regression equations with which to estimate β in equation (6):

$$(7) \quad q_t = \alpha_1 + \beta_1 p_{\tau t} + f_1(t) + u_{1t}$$

$$(8) \quad q_t = \alpha_2 + \beta_2 p_{\tau t} + \gamma_2 \psi_t + f_2(t) + u_{2t}$$

$$(9) \quad q_t = \alpha_3 + \beta_3 p_{\tau t} + \gamma_3 \psi_t + \delta_3 \varepsilon_t + f_3(t) + u_{3t}$$

¹ Because the average of a logarithm does not equal the logarithm of an average, $q_t \neq a_t + y_t + \omega_t$.

where $\varepsilon_{\tau t}$ denotes errors from the first stage regression

$$(10) \quad p_{\tau t} = \theta_0 + \theta_1 \psi_{t-1} + \theta_2 \psi_t + g(t) + \varepsilon_{\tau t}.$$

Equations (7)–(9) are defined such that the error terms (u_{it}) are uncorrelated with the respective right-hand side variables, so the parameters represent the probability limit of the OLS estimator. Equation (9) is the second stage of a 2SLS estimation procedure, for which equation (10) is the first stage, that is, it is an IVs estimator that uses ψ_{t-1} as an instrument for price.² We assume that the exclusion restriction holds to make ψ_{t-1} a valid instrument for $p_{\tau t}$, that is, $E[\psi_{t-1} u_t | \psi_t] = 0$.

We derive the conditions under which β_1 , β_2 , and β_3 equal β . To allow supply shocks embedded in u_t to be predictable, we define v_t such that $E[\varepsilon_{\tau t} v_t] = 0$ in the equation $u_t = \delta \varepsilon_{\tau t} + v_t$, where $\delta < 0$. This condition implies that we can rewrite the equation of interest (6) as

$$(11) \quad q_t = \alpha + \beta p_{\tau t} + \gamma \psi_{\tau t} + \psi_t + \delta \varepsilon_{\tau t} + f(t) + v_t$$

where v_t is uncorrelated with the right-hand side variables. We can now treat equations (7)–(9) as regressions with omitted variables relative to equation (11) and apply standard omitted variable bias formulas. Specifically, the bias in β_1 , β_2 , and β_3 equals the dot product of coefficients on any omitted variables in equation (11) with the coefficients on $p_{\tau t}$ in auxiliary regressions of each omitted variable on the included variables.

Understanding the bias requires expressions for the parameters in these auxiliary regressions. To this end, we define η_t as the surprise in the yield shock, such that $\Psi_t = E_{\tau}[\Psi_t] \eta_t$, where $E_{\tau}[\eta_t] = 1$. It follows that

$$(12) \quad \psi_t = \psi_{\tau t} + \ln(\eta_t)$$

where $cov[\psi_{\tau t}, \ln(\eta_t)] = 0$ as long as η_t is independent of information available at τ . In addition, because it is an important parameter for characterizing the bias, we define the parameter π to represent the predictability

of yield shocks. Specifically, π is the coefficient on price in the following hypothetical regression

$$(13) \quad \psi_t = \mu + \pi p_{\tau t} + h(t) + e_t.$$

Replacing ψ_t with $\psi_{\tau t}$ in this regression would produce the same price coefficient π because η_t is independent of information available at τ . We expect $\pi < 0$ because predictable increases in yield cause the futures price to decrease.

Using these definitions, omitted variables bias formulas (e.g., Wooldridge 2002) imply that

$$(14) \quad \beta_1 = \beta + (1 + \gamma)\phi_{11} + \delta\phi_{12}$$

$$(15) \quad \beta_2 = \beta + \gamma\phi_{21} + \delta\phi_{22}$$

$$(16) \quad \beta_3 = \beta + \gamma\phi_{31}$$

where

$$\phi_{11} = \pi \quad \phi_{12} = \frac{\sigma_{\varepsilon}^2}{\sigma_{p_{\tau}}^2}$$

$$\phi_{21} = \frac{\pi \sigma_{\eta}^2}{\sigma_{\psi}^2 - \sigma_{p_{\tau}}^2 \pi^2}$$

$$\phi_{22} = \frac{\sigma_{\varepsilon}^2 \sigma_{\psi}^2}{\sigma_{p_{\tau}}^2 \sigma_{\psi}^2 - \sigma_{p_{\tau}}^4 \pi^2}$$

$$\phi_{31} = [\pi \sigma_{\eta}^2 - \pi \rho_1 \sigma_{\psi}^2 - \sigma_{\eta}^2 (\pi - \pi_1 \rho_1) (1 - \rho_1)^{-1}] / [\sigma_{\psi}^2 - \sigma_{p_{\tau}}^2 \pi^2 - \sigma_{p_{\tau}}^{-2} \sigma_{\psi}^2 \sigma_{\varepsilon}^2]$$

where $\sigma_{\varepsilon}^2 \equiv var[\varepsilon_{\tau t}]$, $\sigma_{\eta}^2 \equiv var[\ln(\eta_t)]$, $\sigma_{\psi}^2 \equiv var[\psi_t]$, $\sigma_{p_{\tau}}^2 \equiv var[\tilde{p}_{\tau t}]$ is the variance of detrended prices, $\pi_1 \equiv E[\tilde{p}_{\tau t} \psi_{t-1}] \sigma_{\tau p}^{-2}$ is the coefficient from a regression of ψ_{t-1} on $\tilde{p}_{\tau t}$, and $\rho_1 \equiv E[\psi_{t-1} \psi_t] \sigma_{\psi}^{-2} = E[\psi_{t-1} \psi_{\tau t}] \sigma_{\psi}^{-2}$ is the first-order autocorrelation coefficient for ψ_t (derivations of these formulas are available in the supplementary online appendix).

Equations (14)–(16) reveal that all three estimators underestimate the true supply elasticity in general, but are consistent in some special cases. The parameters π , γ , and δ are critical for determining the magnitudes of the bias. Table 1 summarizes the implications for the relevant combinations of these parameters. First, π equals zero if yield shocks are unanticipated by the futures market, and is negative if such shocks are partially predictable. If $\pi = 0$, then yield

² To see this, recall the textbook method for implementing a Hausman test (e.g., Wooldridge 2002, p. 118–122).

shocks must have zero autocorrelation $\rho_1 = 0$. In this case, the 2SLS estimator in (9) is a consistent estimator and the two OLS estimators in equations (7) and (8) are biased only if supply shocks embedded in u_t are predictable (i.e., $\delta < 0$; see rows (i) and (ii) of table 1). Thus, in this case, the 2SLS framework introduced by Roberts and Schlenker produces consistent estimates of supply parameters. Moreover, this case implies that controlling for current yield shocks (ψ_t) has no effect on the OLS estimator, that is, $\beta_1 = \beta_2$ because $\phi_{12} = \phi_{22}$.

If $\pi < 0$ and $\gamma = 0$, then yield shocks are partially predictable but they have no effect on land allocation. In this case, 2SLS is again consistent but there is a significant difference between the two OLS estimators (see rows (iii) and (iv) of table 1). When the current yield shock is omitted from the model, the coefficient on price (β_1) has a greater bias because the omitted variable is correlated with price. The difference between β_2 and β_3 in this setting depends on the value of δ . If δ is small, then β_2 and β_3 will be similar. A small δ means that the endogeneity bias comes mostly from correlation between expected yield shocks and prices. Controlling for yield shocks removes this component of bias, leaving little bias for 2SLS to correct.

If $\delta = 0$, but $\gamma > 0$ and $\pi < 0$, then 2SLS would have a smaller bias than the OLS estimator that controls for the current yield shock if ϕ_{31} were less negative than ϕ_{21} . On the other hand, if ϕ_{31} were more negative than ϕ_{21} , then 2SLS would have a greater bias than OLS (see row (v) of table 1). Comparing, ϕ_{21} and ϕ_{31} , we see that ϕ_{31} has a less negative numerator but also a less positive denominator, so the relative bias is in general unsigned.³ We also note at this point that the

bias expressions we derive are for the probability limit of the estimator. Thus, for 2SLS it essentially assumes that the instrument is infinitely strong. The finite sample bias of the 2SLS estimator is likely worse than the asymptotic bias.

Because we assume that yield shocks are exogenous, we interpret $\pi < 0$ as indicating that predicted yield shocks cause changes in prices. Suppose this assumption fails, and the causation runs in the opposite direction, that is, a negative π reflects the response of yield shocks to price through substitution between lands of different qualities.⁴ The supply elasticity in this case is $\beta + (1 + \gamma)\pi$, which is the sum of the price effects on growing area and on yield deviations from trend. If growing area shocks are exogenous to price ($\delta = 0$), then equation (14) shows that the simple regression in equation (7) produces a consistent estimate of the supply elasticity. Adding the yield shock control as in equation (8) or instrumenting using the lagged yield shock as in equation (9) cause the elasticity to be biased upwards by partialling out the component of the price effect that works through yield. These latter estimators would overestimate the supply elasticity by omitting the tendency of yield to decrease when acreage increases. We explore this possibility empirically, along with the other cases discussed above.

Decomposing the Supply Elasticity

Equation (6) essentially includes ψ_t on both the left and right-hand sides. On the left-hand side, total caloric production is $q_t = a_t + y_t + \psi_t$, but we control for ψ_t because average yield shocks may be predictable. Given that this variable is assumed to be exogenous to prices, there seems little reason to include it on the left-hand side of the model. Modeling the acreage allocation directly using equation (1) and (2), or the sum of these two equations, would be cleaner than adding the noise induced by ψ_t only to take it back out again. Even if yield shocks respond to price because farmers adjust inputs in ways that affect deviations of yield

³ The 2SLS bias would be worse than OLS if, after we control for ε_t , the partial correlation between price and the predictable component of yield becomes stronger. When $\delta = 0$, the 2SLS regression (9) effectively uses ε_t to proxy for ψ_t , whereas the OLS regression in equation (8) leaves ψ_t as an omitted variable. If this proxy absorbs price variation that is uncorrelated with the omitted variable and thereby amplifies the partial correlation between price and the omitted variable, then bias could increase. Put another way, ψ_{t-1} is not a valid instrument for price because it likely has a small correlation with production through the predictable component of yield shocks, so including it as an instrument could worsen bias. A similar phenomenon arises when comparing β_1 and β_2 for the case with $\delta < 0$ and $\pi = 0$. Adding the current yield shock as a control makes the bias worse ($\phi_{22} > \phi_{12}$) because controlling for ψ_t increases the partial correlation between price and ε_t . See Wooldridge (2002, p. 64) for more on the conditions under which imperfect proxies can worsen bias.

⁴ Price could increase yield deviations from trend (i.e., yield shocks) by changing the use of inputs such as fertilizer, or it could decrease yield shocks by encouraging expansion of acreage onto marginal land, thereby causing average yield to deviate from trend. We focus here on the latter case ($\pi < 0$) because it matches our empirical finding of a negative correlation between yield and price.

Table 1. Characterizing the Bias in OLS and IV Estimators

	Sign of Bias			Difference	
	$\beta_1 - \beta$	$\beta_2 - \beta$	$\beta_3 - \beta$	$\beta_2 - \beta_1$	$\beta_3 - \beta_2$
(i) $\pi = 0$, any γ , $\delta = 0$	0	0	0	0	0
(ii) $\pi = 0$, any γ , $\delta < 0$	neg	neg	0	0	pos
(iii) $\pi < 0$, $\gamma = 0$, $\delta = 0$	neg	0	0	pos	0
(iv) $\pi < 0$, $\gamma = 0$, $\delta < 0$	neg	neg	0	pos	pos
(v) $\pi < 0$, $\gamma > 0$, $\delta = 0$	neg	neg	neg	pos	undef
(vi) $\pi < 0$, $\gamma > 0$, $\delta < 0$	neg	neg	neg	pos	undef

Note: The abbreviation “neg” denotes a negative sign and “pos” denotes a positive sign; β_1 is the price coefficient in the OLS regression omitting the yield shock (equation (7)), β_2 is the price coefficient in the OLS regression including the yield shock (equation (8)), and β_3 is the price coefficient in the two-stage least squares regression (equation (9)); π measures the predictability of yield shocks, γ measures the effect of predictable yield shocks on land allocation, and δ measures the effect of predictable supply shocks other than to yield.

from trend, this effect is eliminated from the model once we control for ψ_t .

In the models shown in equations (8) and (9), which control for ψ_t , the effect of price on production works partially through its affect on total growing area as in equation (1) and partially through its effect on trend yield through changing the composition of the growing area as in equation (2). To estimate the relative importance of these two channels, we estimate the following analogs of equations (8) and (9):

$$\begin{aligned}
 (17) \quad a_t &= \alpha_2^a + \beta_2^a p_{t\tau} + \gamma_2^a \psi_t + f_2^a(t) + u_{2t}^a \\
 a_t &= \alpha_3^a + \beta_3^a p_{t\tau} + \gamma_3^a \psi_t + \delta_3^a \varepsilon_{t\tau} \\
 (18) \quad &+ f_3^a(t) + u_{3t}^a \\
 (19) \quad y_t &= \alpha_2^y + \beta_2^y p_{t\tau} + \gamma_2^y \psi_t + f_2^y(t) + u_{2t}^y \\
 y_t &= \alpha_3^y + \beta_3^y p_{t\tau} + \gamma_3^y \psi_t + \delta_3^y \varepsilon_{t\tau} \\
 (20) \quad &+ f_3^y(t) + u_{3t}^y.
 \end{aligned}$$

By construction, the sums of the parameters across dependent variables in these equations equal their analog in equations (8) and (9), for example, $\beta_2^a + \beta_2^y = \beta_2$. Thus, these models provide a simple decomposition of the estimated price responses. Analogous bias expressions to those in equations (15) and (16) can be obtained by replacing β and γ with their analogs from equations (4) and (5).

The model in equation (7) does not control for ψ_t , so the implied effect of price on production from this model could work through all three components, a_t , y_t , and ψ_t . To estimate the relative importance of these three channels, we estimate the following analogs of equation (7) for a_t and y_t :

$$(21) \quad a_t = \alpha_1^a + \beta_1^a p_{t\tau} + f_1^a(t) + u_{1t}^a$$

$$(22) \quad y_t = \alpha_1^y + \beta_1^y p_{t\tau} + f_1^y(t) + u_{1t}^y$$

along with equation (13), which is the analogous equation for ψ_t . The relevant bias expressions are those in equation (14) except with $(1 + \gamma)$ replaced by γ .

It is common for agricultural economists to estimate acreage response equations similar to equation (21). In such equations, the futures price variable is endogenous if the futures market incorporates some available information about forthcoming yield shocks and those predicted yield shocks affect growing area. Roberts and Schlenker (2013) motivate this source of endogeneity using the example of soybean rust in the United States. The discovery of soybean rust in the United States before the 2004 planting season may have induced farmers to reduce soybean acres and thus may have affected futures prices. This point applies to the literature that uses panel data where the dependent variable is the share of a region planted to crops (e.g., Wu and Segerson 1995; Hardie and Parks 1997; Holt 1999; Miller and Plantinga 1999), or where the dependent variable is a discrete variable indicating the crop planted (e.g., Wu et al. 2004; Lubowski, Plantinga, and Stavins 2008; Hendricks, Smith, and Sumner 2014).

The estimate of π in equation (13) provides direct evidence on the predictability of the yield shock. Even if the true yield response to price is zero, OLS estimation of π will likely produce a negative coefficient on the futures price if yield shocks are predictable. To some extent, growing-season weather may be predictable before planting. For example, some droughts persist over several years. In addition to predictions of weather, yield forecasts may also reflect other factors such as pest pressure. The discovery of soybean

rust in the United States provides one such example.

Aggregate versus Disaggregate Production Data

Next we consider the implications of the scale and the scope of the data used in supply analysis for the magnitude of the endogeneity bias. The scale of the data refers to the aggregation level of the individual units of observation (e.g., field level or country level). The scope of the data refers to the geographic extent of the dataset (e.g., regional or global).

The scale of the data does not impact the magnitude of endogeneity bias in the case of a static model. Consider three different estimates of the global supply response: (a) an estimate with aggregate time-series data, (b) a pooled estimate using data for all individual units and years, imposing common coefficients across all individual units, and (c) the supply response averaged across separate regressions for each individual unit. The three estimators converge to the same parameter when the aggregate time-series model is the aggregate of static and linear submodels, therefore the endogeneity bias is the same whatever the scale of the data used in the analysis (see the online appendix for a proof).⁵ This result may not hold when aggregating dynamic models.

The scope of the data, however, does impact the magnitude of endogeneity bias. The correlation between expected yield shocks and the futures price likely differs across regions. Endogeneity bias is smaller in regions where yield shocks are uncorrelated with global yield shocks, and the bias is larger for regions where yield shocks are highly correlated with global yield shocks. Thus, the magnitude of endogeneity bias depends on the regions contained in the data analysis.

Which Futures Price to Use?

Choosing when to observe the futures price variable is an important aspect for specifying the supply equation. Several studies use the

price of a harvest-time contract traded at planting-time (e.g., Gardner 1976; Choi and Helmberger 1993; Goodwin, Vandever, and Deal 2004; Hausman 2012), while others use the price of a harvest-time contract traded prior to planting (e.g., Orazem and Miranowski 1994; Holt 1999; Wu et al. 2004; Hendricks, Smith, and Sumner 2014). Gardner (1976) viewed this choice primarily as an issue of determining when production decisions were made. A more relevant consideration for the econometrician is to choose a price that is subject to fewer endogeneity concerns. Futures markets can incorporate additional information about exogenous factors affecting planted area—such as excess rainfall preventing planting—at the time of planting rather than in the months preceding planting. In the United States, each year the USDA National Agricultural Statistics Service (NASS) releases a prospective plantings report during the last week of March. Based on this consideration it seems preferable to use a futures price traded prior to March to reduce endogeneity concerns in the United States. This point was also made by Orazem and Miranowski (1994).

In their empirical application, Roberts and Schlenker (2013) estimate the world supply of calories using the price of a December or November contract (depending on the commodity) traded one year prior to delivery. For example, the futures price for maize in 2007 is the average December 2006 price of the December 2007 Chicago Board of Trade contract. Roberts and Schlenker (2013) use production data from the Food and Agricultural Organization (FAO), which reports crop production by country according to the calendar year the crop was harvested. For example, Brazilian soybeans planted in October 2006 and harvested in April 2007 are recorded by the FAO as 2007 production. While the Roberts and Schlenker (2013) futures price reflects expectations prior to planting in the Northern Hemisphere, it also reflects realized crop acreage in the Southern Hemisphere and winter-wheat acreage in the Northern Hemisphere.

Estimation and Results

We replicate the OLS and IV supply estimates of Roberts and Schlenker to better understand the potential endogeneity of futures prices in the estimation of supply

⁵ The log of total production is not equivalent to the sum of log production across all fields, so aggregation is not exact when estimating the model in logs, but this likely has minimal effect on results. Our proof in the online appendix uses a model in levels so that aggregation is exact. If all of the right-hand side variables only vary across time, then all three estimators provide identical coefficient estimates.

response. We do not replicate their demand analysis since we are only concerned with supply analysis in this article. We only replicate results using production data from the Food and Agriculture Organization (FAO) of the United Nations. We also estimate similar supply equations with area, trend yield, and the yield shock as dependent variables to decompose bias in the estimated supply elasticity when total production is the dependent variable in the supply equation (equations (17)–(22) and (13)). We further decompose supply elasticity estimates using country-by-country regressions.

Comparison of OLS and IV Estimates

Following Roberts and Schlenker, total caloric production is the sum of the production of maize, rice, soybeans, and wheat using the caloric conversion factors from Williamson and Williamson (1942). We have data from 1961–2007, but we lose one year because the lagged yield shock is used as an instrument. In all regressions, we use the same 46 years, 1962–2007. We construct the Roberts and Schlenker (2013) yield shock ω_t as the weighted average of country and crop-specific log yield shocks as described in their article. We construct our yield shock ψ_t as the log of the weighted average of country and crop-specific yield shocks (i.e., the log of equation 3).

The futures price is the caloric weighted average of the price of maize, soybeans, and wheat.⁶ We use futures contracts with a delivery month of December for maize and wheat and a delivery month of November for soybeans. The futures price is the average price in December one-year prior to delivery. Where price data are missing, we use the monthly average futures price nearest to December of the year prior to the harvest year. Following Roberts and Schlenker (2013) we use cubic splines to model the flexible trends; we report results for splines with 3, 4, and 5 knots. We also used the data posted by Roberts and Schlenker to compare estimation procedures and obtained very similar results (results reported in the online appendix).

We compare alternative estimation procedures in table 2. Panel A reports OLS

estimates of equation (7), which omits the yield shock. Panel B reports OLS estimates of equation (8), which includes the yield shock as a control. Panel C reports 2SLS (IV) estimates of equation (9). Our estimates in columns (1a)–(1c) of panels A and C are remarkably similar to the supply estimates in columns (1a)–(1c) of table 1 and table 4 of Roberts and Schlenker (2013). Results in columns (1a)–(1c) in table 2 use the Roberts and Schlenker (2013) yield-shock measure ω_t , whereas results in columns (2a)–(2c) use ψ_t as the yield shock. Estimates of the supply elasticity differ little whether we use ω_t or ψ_t as the yield shock. For the remainder of our results, we use ψ_t as the yield shock because it allows for the decomposition in (1)–(3).

The OLS estimates of the supply elasticity in panel A of table 2 are much smaller than estimates in panels B or C, indicating substantial predictability in yield shocks. If such predictability were not present ($\pi = 0$), then our bias expressions imply that panels A and B would produce similar results. We report the p-value for an omitted variable bias test between the models in panels A and B.⁷ We reject the null hypothesis of no omitted variable bias at the 1% level for all the specifications of spline knots.

From the bias expressions in equations (14)–(16), we expect a substantial difference between the estimates in panels B and C if growing area shocks are predictable (i.e., if $\delta < 0$). A natural source of such predictability comes from the fact that Southern Hemisphere planted acreage is known at time τ , as is the Northern Hemisphere planted acreage of winter wheat. However, the OLS estimates in panel B are very similar to 2SLS estimates, implying that the effect of such predictability on growing area is small (i.e., δ is close to zero). In fact, OLS estimates are slightly larger than 2SLS estimates in columns (2a) and (2c). We also report the p-value from a Hausman test of endogeneity that tests between the models in panels B and C. We cannot reject the null hypothesis of exogeneity for any of the specifications of spline knots.

In summary, the futures price is endogenous to total production due to a predictable component in observed deviations from trend

⁶ Roberts and Schlenker do not use the futures price of rice when constructing an average price because rice futures did not trade before 1986.

⁷ We perform this test by treating equations (7) and (8) as a pair of seemingly unrelated regressions. We stack the regressions, estimate by OLS, and compute a t-statistic for equality of the β coefficients. We cluster the standard errors by year.

Table 2. Estimates of World Caloric Supply with Alternative Models

	ω_t as Yield Shock			ψ_t as Yield Shock		
	(1a)	(1b)	(1c)	(2a)	(2b)	(2c)
Panel A. OLS Omitting Yield Shock						
Supply Elast.	0.049** (0.023)	0.023 (0.026)	0.022 (0.026)	0.049** (0.023)	0.023 (0.026)	0.022 (0.026)
Panel B. OLS Including Yield Shock						
Supply Elast.	0.110*** (0.013)	0.089*** (0.014)	0.088*** (0.012)	0.112*** (0.013)	0.090*** (0.014)	0.089*** (0.012)
Shock	1.216*** (0.111)	1.194*** (0.106)	1.223*** (0.087)	1.334*** (0.121)	1.313*** (0.114)	1.338*** (0.096)
Panel C. Two-stage Least Squares						
Supply Elast.	0.107*** (0.024)	0.104*** (0.025)	0.088*** (0.019)	0.108*** (0.024)	0.103*** (0.025)	0.088*** (0.020)
Shock	1.202*** (0.131)	1.240*** (0.117)	1.222*** (0.094)	1.315*** (0.145)	1.356*** (0.128)	1.335*** (0.105)
p-value for test of omitted variable bias ($H_0 = \text{no bias}$)	0.007	0.007	0.009	0.006	0.006	0.008
p-value for Hausman test ($H_0 = \text{exogeneity}$)	0.873	0.490	0.995	0.881	0.530	0.995
Observations	46	46	46	46	46	46
Spline Knots	3	4	5	3	4	5

Note: Standard errors were computed under the assumption of homoscedasticity and no serial correlation. Newey-West standard errors with two lags were very similar. Asterisks *, **, and *** denote significance at the 1%, 5%, and 10% levels, respectively.

yield ($\pi < 0$). Other predictable shocks to growing area have little influence ($\delta \approx 0$), so there is little need to apply IVs estimation once current yield shocks are controlled for. Moreover, because the 2SLS estimates of standard errors are roughly 75% larger than standard errors of OLS, these results argue against using 2SLS. This conclusion holds under the Roberts and Schlenker assumption that yield shocks do not respond to price. In the next section, we address this assumption and decompose the econometric biases in supply elasticity estimation due to acreage, trend yield, and yield shocks.

Decomposing the Bias

Table 3 reports the decomposition of the sources of endogeneity found by estimating equations (13) and (17)–(22). The supply elasticity estimates in panel A are OLS estimates that omit the yield shock as a control, panel B are OLS estimates with the yield shock as a control, and panel C are 2SLS estimates with the yield shock as a control, and the lagged yield shock as an instrument. The sum of the estimates in columns (1), (2), and (3) of table 3 equal the corresponding aggregate estimates in column (2) of table 2.

Table 3 provides some evidence that the futures price is partially endogenous to growing area ($\gamma^a > 0$ and $\gamma^y > 0$). The elasticity estimates of the growing area response to price are about 0.012 smaller (a 13–19% reduction) if the yield shock is omitted as a control (columns 1a–1c of panels A and B) and elasticity estimates of the average trend yield are about 0.005 smaller (columns 2a–2c of panels A and B; a 13–19% reduction). These sources of endogeneity may arise if producers adjust their planted acreage in anticipation of a yield shock and futures prices respond to the anticipated change in planted acreage. We conduct a test for omitted variable bias for each growing area model and obtain p-values in the range of 0.055–0.095. For the average trend yield models, the corresponding p-values range from 0.102–0.138.

Table 3 indicates roughly 75% of the difference between the supply elasticity estimates with and without the yield shock control (panels A and B) is due to the predictability of yield shocks. Regressing the yield shock on the futures price provides an elasticity estimate of roughly -0.050 (columns 3a–3c of panel A). This estimate is an order of magnitude too large to be interpreted as a yield

Table 3. Decomposition of Supply Estimates

	a_t as LHS Variable			y_t as LHS Variable			ψ_t as LHS Variable		
	(1a)	(1b)	(1c)	(2a)	(2b)	(2c)	(3a)	(3b)	(3c)
2-4 5-7 8-10									
Panel A. OLS Omitting Yield Shock Supply Elast.	0.070*** (0.010)	0.053*** (0.011)	0.051*** (0.010)	0.026*** (0.005)	0.022*** (0.006)	0.021*** (0.006)	-0.048*** (0.015)	-0.051*** (0.018)	-0.050*** (0.017)
Panel B. OLS Including Yield Shock Supply Elast.	0.081*** (0.011)	0.064*** (0.012)	0.063*** (0.010)	0.031*** (0.006)	0.026*** (0.007)	0.026*** (0.006)	0	0	0
Shock	0.242** (0.102)	0.227** (0.095)	0.246*** (0.084)	0.092* (0.052)	0.086 (0.053)	0.092* (0.051)	1	1	1
Panel C. Two-stage Least Squares Supply Elast.	0.086*** (0.021)	0.082*** (0.021)	0.072*** (0.018)	0.021** (0.011)	0.021* (0.011)	0.017 (0.011)	0	0	0
Shock	0.263** (0.122)	0.288*** (0.109)	0.274*** (0.093)	0.052 (0.064)	0.068 (0.06)	0.061 (0.057)	1	1	1
Panel A - Panel B	-0.011	-0.012	-0.012	-0.004	-0.004	-0.005	-0.048	-0.051	-0.050
Panel B - Panel C	-0.005	-0.018	-0.008	0.010	0.005	0.009	0	0	0
p-value for test of omitted variable bias ($H_0 =$ no bias)	0.086	0.095	0.055	0.106	0.138	0.102	0.003	0.003	0.006
p-value for Hausman test ($H_0 =$ exogeneity)	0.778	0.315	0.588	0.326	0.625	0.342	N/A	N/A	N/A
Observations	46	46	46	46	46	46	46	46	46
Spline Knots	3	4	5	3	4	5	3	4	5

Note: Standard errors were computed under the assumption of homoscedasticity and no serial correlation. Newey-West standard errors with two lags were very similar. Asterisks *, **, and *** denote significance at the 1%, 5%, and 10% levels, respectively.

response to price. Such an interpretation would require that new land brought into production in response to a price rise has close to zero yield. The estimated growing area responses in columns (1a–1c) are 0.051–0.070; it is only possible to increase acreage and decrease yield by the same percentage if the marginal land has zero yield. This interpretation becomes even less plausible when we consider supply response by country since most of the supply response in these estimates comes from the United States.

It seems clear that the reason for a negative estimated yield elasticity is because futures markets incorporate some forecast of the expected yield shock, so the futures price tends to be higher in years with a negative yield shock. Assuming that the true yield elasticity is zero, a statistical test for endogeneity is simply to test whether the coefficient on log futures price is equal to zero. In table 3 we report p-values for this test in the range of 0.003–0.006.

Differences between parameter estimates in panels B and C are small. The OLS estimates of growing area response to price that control for the yield shock are slightly smaller than 2SLS estimates, but p-values for a test of endogeneity range from 0.315–0.778. The comparison of results in panels B and C reinforces the notion that predictable supply shocks embedded in u_t are minimal ($\delta \approx 0$). The OLS estimates of average trend yield response to price are actually larger than 2SLS estimates.

Comparing columns (1a–1c) to columns (2a–2c) in panel B, we see that 70% of the estimated supply response is due to changes in total growing area ($0.063/(0.063 + 0.026) = 0.7$), and the remainder is due to changes in average trend yield. The effect on average trend yield reveals changes in the composition of growing area. The composition of the growing area may change if area response to price is heterogeneous and correlated with trend yields. Two-stage least squares estimates in panel C indicate that roughly 82% of the supply response is due to changes in total growing area.

To summarize, these results suggest that regressions of world growing area on futures prices may have a bias of about 20% that can be mitigated by controlling for realized yield shocks. Regressions of total production on futures prices are subject to a much greater bias because, although yield shocks are

partially predictable, this predictability has a relatively small effect on land allocation.

Disaggregate Growing Area Results

By construction, the supply response estimated by Roberts and Schlenker (2013) works entirely through growing area response to price. The world supply response to price is not, however, equal to the world growing area response to price because price changes alter the composition of output. Our decomposition picks up this output-composition effect through the log-average-trend-yield term (y_t). Alternatively, we could estimate the world supply response to price using a production-weighted average of disaggregate estimates of area response to price. Zellner (1969) showed that regression estimates with aggregate time series data and the average of regression estimates with disaggregate data are both consistent estimators of the aggregate coefficient in the case of a linear, static model with heterogeneous coefficients.⁸

We construct a panel dataset of caloric production using FAO production data. To maintain consistent production regions throughout the sample period, we aggregate production from countries that were formerly part of the Union of Soviet Socialist Republics (USSR), or formerly part of the Socialist Federal Republic of Yugoslavia (Yugoslav SFR). Countries that produced less than 0.5% of world caloric production, on average, were aggregated into one of two regions—“Rest of North” or “Rest of South”—depending on hemisphere. We assign each country to the Northern or Southern Hemisphere based on the average planting and harvest dates for maize and rice from Sacks et al. (2010). Countries where the planting date is earlier than harvest date within a calendar year are assigned to the Northern Hemisphere.⁹ Our dataset is a balanced panel of 31 countries or regions from 1961–2007.

⁸ Our aggregate regressions are not exactly the sum of country-specific regressions because the sum of logs is not equivalent to the log of a sum, but this is not likely to be a major concern in practice.

⁹ In other words, countries are grouped into the Northern Hemisphere if planting maize and rice occurs early in the calendar year and harvest occurs late in the calendar year. We use the dates for the crop that represents the largest portion of total production in the few cases where the crop calendar for maize and rice provide conflicting results for the hemisphere.

Table 4. Estimates of World Caloric Supply as the Production Weighted Average of Disaggregate Estimates of Growing Area Response to Price

	a_{it} as LHS Variable (1)	$a_{it} + y_{it}$ as LHS Variable (2)	a_{cit} as LHS Variable (3)
Panel A. OLS Omitting Yield Shock			
Supply Elast.	0.073*** (0.017)	0.072*** (0.016)	0.088*** (0.016)
Panel B. OLS Including Yield Shock			
Supply Elast.	0.088*** (0.013)	0.089*** (0.013)	0.102*** (0.014)
Shock	0.306*** (0.095)	0.346*** (0.097)	0.281** (0.112)
Panel C. Two-stage Least Squares			
Supply Elast.	0.087*** (0.026)	0.088* (0.052)	0.101*** (0.030)
Shock	0.300** (0.141)	0.340 (0.247)	0.278* (0.160)
Observations	1,426	1,426	5,166
Spline Knots	5	5	5

Note: Standard errors computed using a bootstrap method described in text. Asterisks *, **, and *** denote significance at the 1%, 5%, and 10% levels, respectively.

Table 4 shows estimates of the aggregate supply elasticity using three different dependent variables for the disaggregate area response to price. In column (1), we estimate the growing area response to price for each country or region in our panel. The aggregate supply elasticity is the average of these estimates weighted by caloric production of each country.¹⁰ Estimates of the supply response in column (1) account for changes in aggregate caloric supply due to area response to price being correlated with average caloric yields across countries. For example, countries with a larger area response to price may also have larger caloric yields than average.

The supply responses in column (1) of table 4 do not, however, account for changes in aggregate caloric supply due to changes in the composition of crop area within countries. For example, the proportion of growing area planted to a particular crop within a country may tend to increase as the average caloric price increases and the average caloric yield of that crop differs from the other crops. The estimates in columns (2) and (3) of table 4 account for the possibility of changes

of crop composition within each country. In column (2), we estimate growing area and trend caloric yield response to price for each country—analogue to columns (1c) plus (2c) in table 3. In column (3), we estimate growing area response to price for each crop within each country.¹¹

The right-hand side variables in table 4 are exactly the same as before. We only report estimates using a trend with five spline knots. Standard errors for the aggregate supply elasticities are generated using a pairs bootstrap, clustered by year, with 1,000 replications. The bootstrap procedure randomly draws, with replacement, data across all countries by year. For each set of resampled data, we estimate the aggregate supply elasticity. The standard error is estimated as the standard deviation of supply elasticities across bootstrap replications. Resampling the data by year (i.e., clustering by year) allows for spatial correlation of the errors across all countries and is thus less restrictive than methods that specify a spatial weights matrix. Bester, Conley, and Hansen (2011) propose clustering as a method of accounting for spatially correlated errors. Our bootstrap

¹⁰ That is, our dependent variable is $a_{it} \equiv \ln(\sum_c A_{cit})$ and our weights are $\frac{\sum_i \sum_c A_{cit}^{k_c} Y_{cit} \psi_{cit}}{\sum_i \sum_c \sum_c A_{cit}^{k_c} Y_{cit} \psi_{cit}}$. Alternatively, an estimate of the world area response to price (analogous to table 3 columns 1a–1c) can be calculated as a weighted average of country-specific estimates of area response to price using growing area of each country as weights instead of production.

¹¹ For column (3), we estimate a weighted average of area response to price across all crops in all countries, that is, our dependent variable is $a_{cit} \equiv \ln(A_{cit})$ and the weights are $\frac{\sum_i A_{cit}^{k_c} Y_{cit} \psi_{cit}}{\sum_i \sum_c \sum_c A_{cit}^{k_c} Y_{cit} \psi_{cit}}$.

method assumes independence of the errors between years.¹²

The estimates in columns (1) and (2) of table 4 differ only by a small amount. This small difference implies that changes in the crop composition within countries play a small role in determining the aggregate supply response to price. Thus, the composition effects in (2a)–(2c) of table 3 come from changes in the composition of growing area across countries. Columns (2) and (3) would be asymptotically identical if our dependent variables were measured in levels, rather than logs. Thus, the small differences we observe are due to functional form differences—the sum of logs is not equal to the log of a sum. In general, the differences between columns (1), (2), and (3) are small. Next, we compare the estimates using the disaggregate area response to price with the previous estimates that used data on aggregate caloric production.

Results in panel A of table 4 that omit the yield shock as a control are much larger than estimates in panel A of table 2. As in table 3, most of the endogeneity of the futures price with respect to caloric production is because yield shocks are predictable and affect the futures price—estimating area response to price mitigates this endogeneity concern. Estimates of the aggregate supply elasticity in panels B and C of table 4 are similar to estimates in panels B and C of table 2, which empirically confirms that the estimated aggregate supply response is the production weighted average of disaggregate area response to price.

The OLS estimates in panel A of table 4 are smaller than the corresponding estimates in panel B. As in table 3, this finding indicates a bias of up to 20% due to expected yield shocks affecting anticipated growing area. The fact that growing area in the Southern Hemisphere is known at the time the futures price is determined could also induce bias ($\delta < 0$). Two-stage least squares estimates would reduce the asymptotic bias from this endogeneity. We find that, as with our earlier estimates, the results differ little between panels B and C in table 4, which indicates $\delta \approx 0$. These results suggest that the large body of literature that estimates area response to price may not suffer from

endogeneity concerns to the same degree as do supply estimates where the dependent variable is total production.

We can use the disaggregate estimates underlying table 4 to better understand the source of world supply response. Since most of the aggregate supply response is accounted for as the weighted average of country-specific estimates of growing area response to price (column 1 of table 4), we investigate these country-specific estimates in more detail. Figure 1 shows estimates of the supply elasticity for each country or region with 95% confidence intervals where the flexible trend is specified with 5 spline knots. For some countries 2SLS provides larger estimates of the area response to price, but for others 2SLS provides smaller estimates. Generally, estimates do not differ substantially if we include the yield shock or omit the yield shock from OLS regressions.

Estimates for those countries that produce most of the world's calories are particularly relevant. The top-five countries or regions produce, on average, roughly 65% of the world's calories from these crops, as follows: USA (23%), China (20%), India (9%), the former USSR (7%), and the Rest of the North (7%). Estimates of supply response in China, India, and the Rest of the North are minimal. Supply response in the former USSR is larger, but the largest estimated supply response is in the United States. This finding may in part reflect the fact that we estimate the elasticity with respect to United States prices.

Table 5 shows results for growing area response to price in the United States for the different estimators and different specifications of spline knots. The world supply elasticity due to area response to price in the United States is the area response multiplied by the proportion of caloric production in the United States. For 5 spline knots, this amount is $0.295 \times 0.23 = 0.068$, so approximately 77% ($0.068/0.088 = 0.77$) of the estimated world supply elasticity is due to area response to price in the United States.

The 2SLS estimates in panel C of our table 5 are very similar to the results reported by Roberts and Schlenker (2013) in panel B of their table 3.¹³ The OLS estimates that include the yield shock as a control (panel B

¹² The estimated standard errors in tables 2 and 3 changed only slightly when we used the Newey-West estimator to allow for autocorrelation, so this assumption is likely to be a reasonable approximation.

¹³ One difference in our specification is that we use ψ_t as the yield shock instead of ω_t , but this makes a negligible difference in the results.

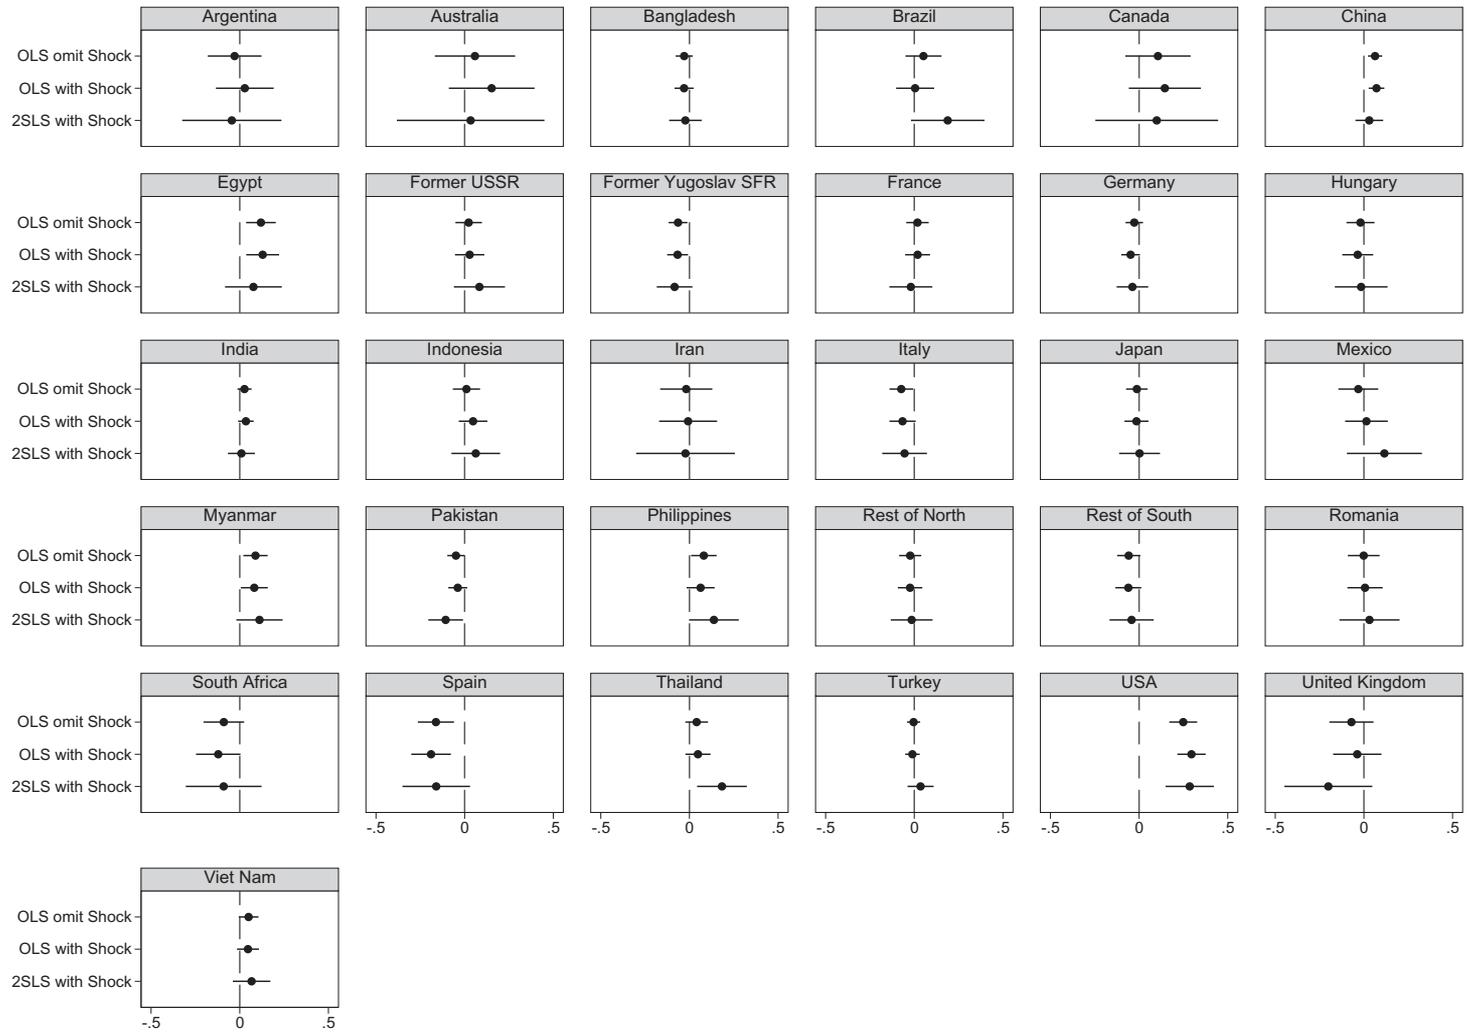


Figure 1. Country-specific estimates of growing area response to price

Table 5. Estimates of Growing Area Response to Price in the United States

	(1)	(2)	(3)
Panel A. OLS Omitting Yield Shock			
Supply Elast.	0.275*** (0.034)	0.245*** (0.040)	0.249*** (0.040)
Panel B. OLS Including Yield Shock			
Supply Elast.	0.320*** (0.036)	0.291*** (0.041)	0.295*** (0.041)
Shock	0.938*** (0.334)	0.909*** (0.331)	0.927*** (0.336)
Panel C. Two-stage Least Squares			
Supply Elast.	0.300*** (0.067)	0.293*** (0.071)	0.285*** (0.069)
Shock	0.853** (0.398)	0.914** (0.368)	0.892** (0.368)
p-value for test of omitted variable bias ($H_0 = \text{no bias}$)	0.038	0.047	0.055
p-value for Hausman test ($H_0 = \text{exogeneity}$)	0.744	0.980	0.877
Observations	46	46	46
Spline Knots	3	4	5

Note: Standard errors were computed under the assumption of homoscedasticity and no serial correlation. Newey-West standard errors with two lags were very similar. Asterisks *, **, and *** denote significance at the 1%, 5%, and 10% levels, respectively.

of table 5) are slightly larger than 2SLS estimates when we use 3 or 5 spline knots in the trend. Hausman tests do not reject the null hypothesis of the exogeneity of the futures price. The OLS estimates that omit the yield shock as a control (panel A of table 5) provide estimates that are about 15% smaller. We reject the null hypothesis of no omitted variable bias at the 5% level for trend specifications with 3 or 4 spline knots. Since the dependent variable is growing area in these regressions, all of the omitted variable bias is due to the endogeneity of the futures price with growing area.

We also considered a specification in which we estimated country-specific area response to price using different futures prices for countries in the Northern and Southern Hemispheres (results not reported). The futures price used for each hemisphere is the price of a contract for delivery after harvest trading prior to planting for the respective planting, and harvest periods for the Northern and Southern Hemisphere. This futures price should reduce endogeneity concerns for the Southern Hemisphere because the futures price is observed before the growing area is known. Results using different futures prices for each hemisphere differed little with our results in table 4. One explanation for the small difference between these results is that most of the world's caloric production (87%) occurs in the Northern Hemisphere.

Conclusions

The futures price is endogenous in agricultural supply analysis if supply shocks that are predictable prior to planting are omitted from the regression. Including the current-year realized yield shock as a control reduces this bias, but does not eliminate it because the realized yield shock is a noisy proxy for an unobservable omitted variable, that is, the expected yield shock. Instrumental variables estimates that use the lagged yield shock as an instrument are also biased if yields are autocorrelated because the lagged yield shock is correlated with the noise of the proxy. Although both of these estimators are biased, they may be preferred to estimators that use demand shifters as instruments since demand shifters may have little explanatory power of the futures price. Even if the instruments are uncorrelated with the omitted variables, two-stage least squares estimates are biased in finite samples and the bias can be large when the instruments are weak (Bound, Jaeger, and Baker 1995).

We replicate the OLS and two-stage least squares results of Roberts and Schlenker for world supply of maize, soybeans, wheat, and rice from 1961–2007 and consider an intermediate case of OLS with the realized yield shock as a control. There is little difference in OLS estimates that control for the realized yield shock and two-stage least squares

estimates. When total global production is the dependent variable, OLS estimates of the supply elasticity that omit the realized yield shock are biased downward substantially. Decomposing total global production reveals that this bias is due to a negative and significant relationship between price and realized deviations from trend yield. Interpreting this relationship as an expansion of low productivity cropland does not seem plausible given the magnitude of the estimated relationship. An alternative explanation is that futures prices reflect some expected or predictable component of realized yield shocks. We find that bias in estimates of the supply elasticity are reduced when growing area rather than total production is the dependent variable. The remaining small bias is mitigated by adding the realized yield shock as a control.

In summary, we argue that the aggregate crop supply elasticity can be well estimated by OLS regression of growing area on a futures price with controls for a flexible trend and realized deviations from trend yield. This recommendation should not be applied without qualification. Any regression estimate of a supply elasticity is specific to the type of quantity and price variation in the sample, the location in space of the price measured, and the time period covered. By removing a flexible deterministic trend from the data to reduce the risk of spurious correlation, the model in this article necessarily removes the most persistent price shocks. Thus, the estimates reflect quantity responses to mostly short-run price shocks. The elasticity with respect to persistent price shocks may be larger.

References

- Bester, C.A., T.G. Conley, and C.B. Hansen. 2011. Inference with Dependent Data using Cluster Covariance Estimators. *Journal of Econometrics* 165: 137–151.
- Bound, J., D.A. Jaeger, and R.M. Baker. 1995. Problems with Instrumental Variables Estimation when the Correlation between the Instruments and the Endogenous Explanatory Variable is Weak. *Journal of the American Statistical Association* 90: 443–450.
- Choi, J.S., and P.G. Helmberger. 1993. Acreage Response, Expected Price Functions, and Endogenous Price Expectations. *Journal of Agricultural and Resource Economics* 18: 37–46.
- Gardner, B.L. 1976. Futures Prices in Supply Analysis. *American Journal of Agricultural Economics* 58: 81–84.
- Goodwin, B.K., M.L. Vandever, and J.L. Deal. 2004. An Empirical Analysis of Acreage Effects of Participation in the Federal Crop Insurance Program. *American Journal of Agricultural Economics* 86: 1058–1077.
- Hardie, I.W., and P.J. Parks. 1997. Land Use with Heterogeneous Land Quality: An Application of an Area Base Model. *American Journal of Agricultural Economics* 79: 299–310.
- Hausman, C. 2012. Biofuels and Land Use Change: Sugarcane and Soybean Acreage Response in Brazil. *Environmental and Resource Economics* 51: 163–187, doi: 10.1007/s10640-011-9493-7.
- Hendricks, N.P., A. Smith, and D.A. Sumner. 2014. Crop Supply Dynamics and the Illusion of Partial Adjustment. *American Journal of Agricultural Economics* doi: 10.1093/ajae/aau024.
- Holt, M.T. 1999. A Linear Approximate Acreage Allocation Model. *Journal of Agricultural and Resource Economics* 24: 383–397.
- Lubowski, R.N., A.J. Plantinga, and R.N. Stavins. 2008. What Drives Land-use Change in the United States? A National Analysis of Landowner Decisions. *Land Economics* 84: 529–550.
- Miller, D.J., and A.J. Plantinga. 1999. Modeling Land Use Decisions with Aggregate Data. *American Journal of Agricultural Economics* 81: 180–194.
- Nerlove, M. 1958. *The Dynamics of Supply: Estimation of Farmers' Response to Price*. Baltimore, MD: The Johns Hopkins Press.
- Orazem, P.F., and J.A. Miranowski. 1994. A Dynamic Model of Acreage Allocation with General and Crop-specific Soil Capital. *American Journal of Agricultural Economics* 76: 385–395.
- Roberts, M.J., and W. Schlenker. 2013. Identifying Supply and Demand Elasticities of Agricultural Commodities: Implications for the US Ethanol Mandate. *American Economic Review* 103: 2265–2295.
- Sacks, W.J., D. Deryng, J.A. Foley, and N. Ramankutty. 2010. Crop Planting Dates: An Analysis of Global Patterns.

- Global Ecology and Biogeography* 19: 607–620.
- Williamson, L., and P. Williamson. 1942. What We Eat. *Journal of Farm Economics* 24: 698–703.
- Wooldridge, J.M. 2002. *Econometric Analysis of Cross Section and Panel Data, 1st ed.* Cambridge, MA: MIT Press.
- Wright, P.G. 1928. *The Tariff on Animal and Vegetable Oils.* New York: Macmillan.
- Wu, J., R.M. Adams, C.L. Kling, and K. Tanaka. 2004. From Microlevel Decisions to Landscape Changes: An Assessment of Agricultural Conservation Policies. *American Journal of Agricultural Economics* 86: 26–41.
- Wu, J., and K. Segerson. 1995. The Impact of Policies and Land Characteristics on Potential Groundwater Pollution in Wisconsin. *American Journal of Agricultural Economics* 77: 1033–1047.
- Zellner, A. 1969. On the Aggregation Problem: A New Approach to a Troublesome Problem. In *Economic Models, Estimation and Risk Programming: Essays in Honor of Gerhard Tintner*, ed. K. Fox, J. Sengupta, and G. Narasimham, 365–374. Berlin: Springer-Verlag.