NBER WORKING PAPER SERIES

BOUNDS ON A SLOPE FROM SIZE RESTRICTIONS ON ECONOMIC SHOCKS

Marco Stenborg Petterson
David G. Seim
Jesse M. Shapiro

Working Paper 27556
http://www.nber.org/papers/w27556

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2020

Kevin Murphy's lectures in price theory inspired this work. We acknowledge funding from the Population Studies and Training Center, the Eastman Professorship, and the JP Morgan Chase Research Assistant Program at Brown University, and the Jan Wallander and Tom Hedelius Foundation (Grant P2015-0095:1). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding sources. We thank Isaiah Andrews, Steve Berry, Ray Fair, Ed Glaeser, Phil Haile, Alex MacKay, Serena Ng, Emily Oster, Sharon Oster, Michael Roberts, Chris Snyder, Elie Tamer, and seminar participants at Brown University for helpful comments. We thank our dedicated research assistants for their contributions to this project. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w27556.ack

Bounds on a Slope from Size Restrictions on Economic Shocks
Marco Stenborg Petterson, David G. Seim, and Jesse M. Shapiro
NBER Working Paper No. 27556
July 2020
JEL No. C22,C32,D41,Q11

ABSTRACT

We study the problem of learning about the slope of a linear relationship between an outcome (e.g., quantity) and an input (e.g., price) when the outcome is subject to time-varying, unobserved economic shocks. We show that restrictions on the absolute magnitude of the economic shocks are informative for the value of the slope. We argue that such restrictions are reasonable in some economic situations. We illustrate with an application to the demand and supply of food grains. We consider extensions including to the case of a nonlinear relationship.

Marco Stenborg Petterson
Brown University
marco_stenborg_petterson@brown.edu

David G. Seim
Universitetsvägen 10 A, 106 91 Stockholm
Stockholm University
david.g.seim@gmail.com

Jesse M. Shapiro
Economics Department
Box B
Brown University
Providence, RI 02912
and NBER
jesse_shapiro_1@brown.edu

Bounds on a Slope from Size Restrictions on Economic Shocks

Marco Stenborg Petterson, *Brown University*[*]

David Seim, *Stockholm University, CEPR, and Uppsala University*

Jesse M. Shapiro, *Brown University and NBER*

July 2020

Abstract

We study the problem of learning about the slope of a linear relationship between an outcome (e.g., quantity) and an input (e.g., price) when the outcome is subject to time-varying, unobserved economic shocks. We show that restrictions on the absolute magnitude of the economic shocks are informative for the value of the slope. We argue that such restrictions are reasonable in some economic situations. We illustrate with an application to the demand and supply of food grains. We consider extensions including to the case of a nonlinear relationship.

1 Overview and Contribution

A canonical problem in empirical economics is to learn the slope of some outcome (e.g., quantity demanded) with respect to some input (e.g., price) in the presence of an unobserved factor (e.g., preferences) that may be related to the input. Concretely, say that we observe the outcome q_t and the input p_t in each of $T \geq 2$ consecutive periods t, so the data are $\{(p_t, q_t)\}_{t=1}^{T}$. We assume that

$$q_t = \theta p_t + \varepsilon_t \tag{1}$$

[*]Kevin Murphy's lectures in price theory inspired this work. We acknowledge funding from the Population Studies and Training Center, the Eastman Professorship, and the JP Morgan Chase Research Assistant Program at Brown University, and the Jan Wallander and Tom Hedelius Foundation (Grant P2015-0095:1). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding sources. We thank Isaiah Andrews, Steve Berry, Ray Fair, Ed Glaeser, Phil Haile, Alex MacKay, Serena Ng, Emily Oster, Sharon Oster, Michael Roberts, Chris Snyder, Elie Tamer, and seminar participants at Brown University for helpful comments. We thank our dedicated research assistants for their contributions to this project.

E-mail: marco_stenborg_petterson@brown.edu, david.seim@ne.su.se, jesse_shapiro_1@brown.edu.

for ε_t an unobserved factor, and that $p_t \neq p_{t+1}$ for at least one $t < T$. We wish to learn the value of the slope $\theta \in \overline{\Theta} \subseteq \mathbb{R}$, which quantifies the effect on the outcome q_t of changing the input p_t while holding constant the unobserved factor ε_t. The set $\overline{\Theta}$ encodes *a priori* economic restrictions (e.g., negativity) on the possible values of θ.

To fix ideas, suppose that q_t is the log quantity of a competitively supplied good, p_t is its log price, $\theta \in \overline{\Theta} = \mathbb{R}_{\leq 0}$ is its elasticity of demand, and ε_t reflects factors such as income, preferences, and prices of substitutes that determine the quantity demanded at a given price. Figure 1(A) depicts a hypothetical dataset with $T = 2$. As is well known (e.g., Leontief 1929, p. 26; see also Wright 1915 and Working 1927), any elasticity $\tilde{\theta} \in \overline{\Theta}$ can be rationalized with any pair of points $\{(p_1, q_1), (p_2, q_2)\}$, by taking $\varepsilon_1 = q_1 - \tilde{\theta} p_1$ and $\varepsilon_2 = q_2 - \tilde{\theta} p_2$. The figure illustrates this construction for three possible elasticities θ', θ'', and θ'''.

Any conjectured elasticity $\tilde{\theta} \in \overline{\Theta}$ thus implies a particular value of the demand shock $\varepsilon_2 - \varepsilon_1$, which is the change in the log quantity demanded, at a given price, between periods 1 and 2. Some values of the demand shock may be economically implausible. For example, in the world market for grain, the main influences on demand in recent decades have been population and income, and demand has been relatively income-inelastic (Johnson 1999; Valin et al. 2014). It may therefore be reasonable to conclude that, at a given price, per-capita demand has changed by no more than a few percent per year.

A bound on the size of the shock restricts the possible values of the elasticity. For Δ the first-difference operator, let $\Delta\varepsilon_t(\cdot)$ be the real-valued, data-dependent function that gives the value of the economic shock in period t implied by a given slope θ,

$$\Delta\varepsilon_t(\theta) = \Delta q_t - \theta \Delta p_t.$$

In our hypothetical dataset, the values of $\theta \in \mathbb{R}$ compatible with the restriction that $|\Delta\varepsilon_2(\theta)| \leq B$ for some bound $B \geq 0$ are given by the interval $\hat{\Theta}(B) = \frac{1}{\Delta p_2}[\Delta q_2 - B, \Delta q_2 + B]$. As depicted in Figure 1(B), each elasticity $\theta \in \hat{\Theta}(B)$ defines a line $\theta \Delta p_2$ through the origin that intersects the line segment connecting $(\Delta p_2, \Delta q_2 - B)$ to $(\Delta p_2, \Delta q_2 + B)$.

It is possible to extend this logic to the case where $T > 2$. For any $D-$dimensional vector v and any $k > 1$, write the generalized $k-$mean $M_k(v) = \left(\frac{1}{D}\sum_{d=1}^{D} v_d^k\right)^{1/k}$, with $M_\infty(v) = \max_d\{v_d\}$, and write the absolute value $|v| = (|v_1|, ..., |v_D|)$. Let $\hat{M}_k(\theta) = M_k(|\Delta\varepsilon(\theta)|)$ denote the $k-$mean of the absolute value of the vector of shocks $\Delta\varepsilon(\theta) = (\Delta\varepsilon_2(\theta), ..., \Delta\varepsilon_T(\theta))$. We can then define the set of slopes

$$\hat{\Theta}_k(B) = \left\{\theta \in \mathbb{R} : \hat{M}_k(\theta) \leq B\right\} \tag{2}$$

that are compatible with a given bound B on the value of $\hat{M}_k(\theta)$.

Section 2 shows that, when nonempty, the set $\hat{\Theta}_k(B)$ takes the form of a closed interval. For $\hat{\Theta}_\infty(B)$, where the bound is applied to the maximum absolute shock, and $\hat{\Theta}_2(B)$, where the bound is applied to the root mean squared shock, the limit points of the interval can be expressed in closed form. In the remaining cases, the limit points are the solutions to a tractable nonlinear equation. In all cases, under the given restriction on the magnitude of the shocks $\Delta\varepsilon(\theta)$, the true slope must lie in the interval $\hat{\Theta}_k(B)$.

Some bounds B are incompatible with the data in the sense that $\hat{\Theta}_k(B) \cap \overline{\Theta} = \emptyset$ or even $\hat{\Theta}_k(B) = \emptyset$. For example, if $T = 2$, $\overline{\Theta} = \mathbb{R}_{\leq 0}$, and $\Delta p_2, \Delta q_2 > 0$, then $\hat{\Theta}_k(B) \cap \overline{\Theta} = \emptyset$ for any $k > 1$ and any $B < \Delta q_2$, because $\hat{\Theta}_k(B) = \hat{\Theta}(B) = \frac{1}{\Delta p_2}[\Delta q_2 - B, \Delta q_2 + B]$. If $T = 3$ and the points $\{(p_t, q_t)\}_{t=1}^{T}$ do not lie on a single line, $\hat{\Theta}_k(B) = \emptyset$ for any $k > 1$ and any sufficiently small B, because $\hat{M}_k(\theta)$ is bounded away from zero. More generally, let $\mathscr{B}\left(k, \overline{\Theta}\right) = \left\{B : \hat{\Theta}_k(B) \cap \overline{\Theta} \neq \emptyset\right\}$ be the set of bounds B that are compatible with the data and with $\overline{\Theta}$. Section 2 shows that $\mathscr{B}(k, \mathbb{R})$ is a left-bounded interval and characterizes its limit point.

Section 3 applies the characterizations developed in Section 2 to the analysis of the demand and supply of food grains considered by Roberts and Schlenker (2013a). Roberts and Schlenker (2013a) estimate log-linear models of demand and supply under maintained exclusion restrictions and conclude that both are price-inelastic. We maintain log-linearity but do not impose exclusion restrictions. Instead we consider bounds on the size of demand and supply shocks motivated by prior evidence on the origins of these shocks during the sample period. These bounds are informative, implying that demand and supply are price-inelastic. The implied upper bounds on the absolute elasticities are close in magnitude to Roberts and Schlenker's point estimates (2013a). We also calculate the minimum size of demand and supply shocks necessary to rationalize the data. Finally, we show that the upper bounds we obtain on the absolute elasticities imply an informative lower bound for a function of the elasticities that Roberts and Schlenker (2013a) estimate.

Section 4 discusses extensions. In the cases where $q_t = q(p_t) + \varepsilon_t$ for some possibly nonlinear function $q(\cdot)$, or $q_t = q(p_t, \varepsilon_t)$ for some possibly nonseparable function $q(\cdot)$, we show how to bound the value of the average elasticity between any two periods. The resulting bounds are informative in the setting of Roberts and Schlenker (2013a).

The main contribution of this paper is to show that economically motivated bounds on the magnitude of unobserved shocks can be useful in learning about a structural slope parameter. Restrictions on the distribution of unobserved variables are central to many canonical approaches to identification (see, e.g., Matzkin 2007; Tamer 2010). A leading approach is to restrict the relationship between the unobserved shock and some other variable. An example (and the approach taken by Roberts and Schlenker 2013a) is the widely employed exclusion restriction that the unobserved shock is independent of or uncorrelated with an observed instrument. Another example is the assumption that demand and supply shocks are independent of or uncorrelated with one another, studied in Leamer (1981) and developed by Feenstra (1994), Feenstra and Weinstein (2017), MacKay and Miller (2019), and others.[1] Our approach does not directly restrict the relationship of the unobserved shocks to observed variables or to one another. Section 4 discusses the relationship of such restrictions to those we consider.

Although many approaches to identification employ restrictions on the distribution of unobservables, some employ restrictions on the realized values of unobservables, as we do here. Such restrictions can be especially appropriate in time-series settings where the researcher may have prior knowledge of the factors influencing the outcome during the sample period. In the structural vector autoregression setting, Antolín-Díaz and Rubio-Ramírez (2018) consider restrictions on the relative importance of a given shock in explaining the change in a given observed variable during a given time period (or periods). Ludvigson et al. (2020) consider inequality constraints on the absolute magnitude of shocks during a given period (or periods), as well as inequality constraints on the correlation between a shock and an observed variable. In the demand estimation setting, Mullin and Snyder (forthcoming) obtain bounds on the elasticity of demand in a reference period under the assumption that demand is growing over time.[2] None of these sets of restrictions coincides with those we consider here.

[1]See also Leontief (1929). Morgan (1990, Chapter 6) quotes a 1913 thesis by Lenoir which discusses how the relative variability of demand and supply shocks influences the correct interpretation of data on market quantities and prices. Leamer (1981) also imposes that the demand (supply) elasticity is negative (positive). A large literature (reviewed, for example, in Uhlig 2017) develops the implications of sign restrictions in a variety of settings.

[2]In our leading example of log-linear demand, this corresponds to the assumption that $\Delta\varepsilon_t > 0$ for all t. Mullin and Snyder (forthcoming) consider a variety of forms for demand in the reference period, including linear demand, demand known up to a scalar parameter, and concave demand.

2 Characterization of Sets of Interest

We begin with the case of $k = \infty$, in which we bound the maximum absolute value of the shock. In this case the sets of interest take a particularly simple form.

Proposition 1. *Let*

$$\underline{\theta}_\infty(B) = \max_{\{t:\Delta p_t \neq 0\}} \left\{ \frac{\Delta q_t}{\Delta p_t} - \frac{B}{|\Delta p_t|} \right\}$$
$$\overline{\theta}_\infty(B) = \min_{\{t:\Delta p_t \neq 0\}} \left\{ \frac{\Delta q_t}{\Delta p_t} + \frac{B}{|\Delta p_t|} \right\}$$

and let $\tilde{B} \geq 0$ *be the unique solution to* $\underline{\theta}_\infty(\tilde{B}) = \overline{\theta}_\infty(\tilde{B})$.

Then $\mathscr{B}(\infty, \mathbb{R}) = [\underline{B}_\infty, \infty)$ *for* $\underline{B}_\infty = \max\left\{\max_{\{t:\Delta p_t = 0\}}\{|\Delta q_t|\}, \tilde{B}\right\}$, *and for any* $B \in \mathscr{B}(\infty, \mathbb{R})$

$$\hat{\Theta}_\infty(B) = \left[\underline{\theta}_\infty(B), \overline{\theta}_\infty(B)\right].$$

All proofs are given in Appendix A. The objects $\underline{B}_\infty$, $\underline{\theta}_\infty(B)$, and $\overline{\theta}_\infty(B)$ defined in Proposition 1 can be readily calculated on datasets of reasonable size.[3]

We next consider the case of $k \in (1, \infty)$. Here we make use of the following properties of the function $\hat{M}_k(\theta)$:

Lemma 1. *For* $k \in (1, \infty)$, *the function* $\hat{M}_k(\theta)$ *is strictly decreasing on* $(-\infty, \breve{\theta}_k)$ *and strictly increasing on* $(\breve{\theta}_k, \infty)$ *for* $\breve{\theta}_k = \arg\min_\theta \hat{M}_k(\theta)$.

Lemma 1 implies that $\hat{M}_k(\theta)$ has a "bowl" shape, first decreasing to a unique global minimum and then increasing. The following characterization of $\hat{\Theta}_k(B)$ is then immediate.

Proposition 2. *For* $k \in (1, \infty)$, *the set* $\mathscr{B}(k, \mathbb{R})$ *is equal to* $[\underline{B}_k, \infty)$ *for* $\underline{B}_k = \min_\theta \hat{M}_k(\theta)$. *Moreover, for any* $B \in \mathscr{B}(k, \mathbb{R})$ *we have that*

$$\hat{\Theta}_k(B) = \left[\underline{\theta}_k(B), \overline{\theta}_k(B)\right]$$

where $\underline{\theta}_k(B), \overline{\theta}_k(B)$ *are the only solutions to* $\hat{M}_k(\theta) = B$, *with* $\breve{\theta}_k = \underline{\theta}_k(\underline{B}_k) = \overline{\theta}_k(\underline{B}_k)$.

Proposition 2 shows that $\mathscr{B}(k, \mathbb{R})$ is a left-bounded interval whose limit point $\underline{B}_k$ can be calculated by minimizing the function $\hat{M}_k(\theta)$. By Lemma 1, $\hat{M}_k(\theta)$ strictly increases as θ departs from

[3] The objects $\underline{\theta}_\infty(B)$ and $\overline{\theta}_\infty(B)$ also appear in the analysis of the linear regression model with uniformly distributed errors (Robbins and Zhang 1986).

$\breve{\theta}_k$, which simplifies computation. The limit point $\underline{B}_k$ has a direct economic interpretation as the minimum size of shocks necessary to rationalize the data.

Proposition 2 further shows that $\hat{\Theta}_k(B)$ is a closed interval whose limit points can be calculated by solving the nonlinear equation $\hat{M}_k(\theta) = B$. By Lemma 1, on either side of $\breve{\theta}_k$ and for $B > \underline{B}_k$ the equation is strictly monotone and has a unique solution, which simplifies computation. The sets characterized in Propositions 1 and 2 are related by the fact that $\hat{\Theta}_\infty(B) \subseteq \hat{\Theta}_k(B)$ for any $B \geq 0$ and $k \in (1, \infty)$.

In the special case of $k = 2$, in which we bound the root mean squared shock, the equation $\hat{M}_2(\theta) = B$ is quadratic, and so the objects $\underline{B}_2$, $\underline{\theta}_2(B)$, $\overline{\theta}_2(B)$, and $\breve{\theta}_2$ defined in Proposition 2 are available in closed form. For any $D-$dimensional vector $v \in \mathbb{R}^D$, let $\Delta v = (\Delta v_2, ..., \Delta v_D) \in \mathbb{R}^{D-1}$. For any $v, w \in \mathbb{R}^D$, let $\hat{s}_{vw} = M_1(\Delta v \circ \Delta w)$, where $\circ$ is the elementwise product.

Corollary 1. *For $k = 2$ we have that*

$$\begin{aligned}
\underline{\theta}_2(B) &= \frac{\hat{s}_{qp}}{\hat{s}_{pp}} - \sqrt{\left(\frac{\hat{s}_{qp}}{\hat{s}_{pp}}\right)^2 - \frac{1}{\hat{s}_{pp}}\left(\hat{s}_{qq} - B^2\right)} \\
\overline{\theta}_2(B) &= \frac{\hat{s}_{qp}}{\hat{s}_{pp}} + \sqrt{\left(\frac{\hat{s}_{qp}}{\hat{s}_{pp}}\right)^2 - \frac{1}{\hat{s}_{pp}}\left(\hat{s}_{qq} - B^2\right)} \\
\underline{B}_2 &= \sqrt{\hat{s}_{qq} - \left(\frac{\hat{s}_{qp}}{\hat{s}_{qp}}\right)^2 \hat{s}_{pp}} \\
\breve{\theta}_2 &= \frac{\hat{s}_{qp}}{\hat{s}_{pp}}.
\end{aligned}$$

Observe that $\breve{\theta}_2 = \underline{\theta}_2(\underline{B}_2) = \overline{\theta}_2(\underline{B}_2)$ corresponds to the slope of the ordinary least squares regression of Δq_t on Δp_t with no intercept, i.e., the line through the origin with best least-squares fit to the data $\{(\Delta p_t, \Delta q_t)\}_{t=2}^T$.

3 Application to Demand and Supply of Food Grains

Roberts and Schlenker (2013a) estimate the elasticities of demand and supply of staple grains using annual data from 1960 through 2007. Roberts and Schlenker (2013a, equations 1 and 3) assume that demand and supply curves take a log-linear form consistent with equation (1). Roberts and Schlenker (2013a) adopt an instrumental variables approach, using contemporaneous yield shocks as excluded instruments for price in the demand equation, and past yield shocks as excluded instruments for price in the supply equation. Here we explore the information about parameters of interest that can be obtained from bounds on the size of realized shocks, without imposing exclusion restrictions.

Our analysis relies on the code and data from Roberts and Schlenker (2013b), supplemented with data from the World Bank (2019a, 2019b) on annual world population and GDP. For the demand equation, the data $\left\{\left(p_t^D, q_t^D\right)\right\}_{t=1}^T$ consist of the log p_t^D of the average current-month futures price of grains delivered in year t, measured in 2010 US dollars per calorie, and the log quantity q_t^D of grains consumed in the world in year t, measured in calories per capita. For the supply equation, the data $\left\{\left(p_t^S, q_t^S\right)\right\}_{t=1}^T$ consist of the log p_t^S of the average one-year-ahead futures price of grains delivered in year t, measured in 2010 US dollars per calorie, and the log quantity q_t^S of grains produced in the world in year t, measured in calories per capita.[4] We compute world GDP per capita y_t in 2010 US dollars by deflating nominal GDP per capita by the CPI provided by Roberts and Schlenker (2013b).

We consider bounds on the size of demand shocks that are motivated by economic features of the demand for food grain during the sample period. The major determinants of world demand for grain in the modern period are population and income (Johnson 1999; Valin et al. 2014). Because we measure demand on a per capita basis, and because demand for grain is relatively income-inelastic (Johnson 1999; Valin et al. 2014), we expect relatively small annual shocks to the demand for grain. For example, forecasts summarized in Valin et al. (2014, Table 3) imply an income elasticity of world food crop demand ranging from 0.09 to 0.37.[5] Given the evolution of annual log world GDP per capita y_t over the study period, an income elasticity of demand of 0.37 implies that the income-driven shock to log per-capita demand has maximum absolute value $M_\infty(|0.37\Delta y|) \approx 0.05$ and root mean square value $M_2(|0.37\Delta y|) \approx 0.02$. Allowing for some non-income-driven shocks of similar magnitude, we consider bounds B^D on demand shocks in $[0, 0.10]$ for $k = \infty$ and in $[0, 0.04]$ for $k = 2$.

These bounds are context-specific. Larger bounds might be appropriate in historical periods with more income-elastic food demand (see, e.g., Logan 2006) or in periods that include major shocks such as world wars or global pandemics. The bounds are also *a priori* in the sense that they incorporate information that is not contained in the time-series data that we analyze. For example,

[4] We use the definitions of price and total calories from Roberts and Schlenker (2013a, Table 1, Column 2c), and divide total calories by world population to obtain calories per capita.

[5] The models summarized in Valin et al. (2014, Table 3) imply that an increase from \$6,700 to \$16,000 in world GDP over the period 2005-2050 will cause an increase in per capita food demand of between 8 and 38 percent. The implied income elasticities therefore range from $\ln(1.08)/\ln(16000/6700) \approx 0.088$ to $\ln(1.38)/\ln(16000/6700) \approx 0.370$.

estimates of the income elasticity of food demand can be informed by comparisons across countries at a point in time.[6] We consider smaller values of B^D in the case of $k = 2$ than in the case of $k = \infty$, consistent with the facts that $M_2(|0.37\Delta y|) < M_\infty(|0.37\Delta y|)$ and that $\hat{\Theta}_\infty(B) \subseteq \hat{\Theta}_2(B)$ for any $B \geq 0$.

Figure 2 depicts the implications of the contemplated bounds B^D for the elasticity of demand $\theta^D \in \overline{\Theta}_D = \mathbb{R}_{\leq 0}$. The first column of plots considers bounds B^D on the maximum value of the shock ($k = \infty$), and the second column of plots considers bounds B^D on the root mean squared shock ($k = 2$). In each column, the first row of plots depicts a scatterplot of the data $\left\{\left(\Delta p_t^D, \Delta q_t^D\right)\right\}_{t=2}^T$ along with the demand functions consistent with a particular value of B^D. The second row of plots depicts the interval $\hat{\Theta}_k\left(B^D\right) \cap \overline{\Theta}_D$ associated with each B^D in a given interval.

The first row of plots in Figure 2 shows that demand functions take the form of lines through the origin. A demand function is consistent with the restrictions $\hat{M}_k\left(\theta^D\right) \leq B^D$ and $\theta^D \leq 0$ if and only if its slope θ^D is in $\hat{\Theta}_k\left(B^D\right) \cap \overline{\Theta}_D$. A demand function with slope θ^D passes through the points $\left\{\left(\Delta p_t, \theta^D \Delta p_t\right)\right\}_{t=2}^T$. For any t, the set of points $\left\{\left(\Delta p_t^D, \theta^D \Delta p_t^D\right)\right\}_{\theta^D \in \hat{\Theta}_k(B^D)}$ defines a vertical line segment that is depicted (with a solid line) on the plot. For $k = \infty$, the line segment is contained in the segment $\left\{\left(\Delta p_t^D, \Delta q_t^D + b\right)\right\}_{b \in [-B^D, B^D]}$, which is also depicted (with a dotted line) for reference.

The second row of plots in Figure 2 shows that the bounds we contemplate are informative. Any of the contemplated bounds implies that demand is inelastic, $\left|\theta^D\right| < 1$. Imposing an exclusion restriction, Roberts and Schlenker (2013a, Table 1, Column 2c) estimate that the elasticity of demand is $\hat{\theta}_{RS}^D = -0.066$ with a confidence interval of $[-0.107, -0.025]$, also depicted in the plot. A bound of $B^D = 0.07$ on the maximum shock implies an elasticity not below -0.122. The same bound on the elasticity arises from a bound of $B^D = 0.030$ on the root mean squared shock. Appendix Figure 1 depicts the corresponding bound for other values of k, as well as (for reference) $M_k(|0.37\Delta y|)$.

The second row of plots in Figure 2 also illustrates the interpretation of the set $\mathscr{B}\left(k, \overline{\Theta}_D\right)$. The plots report that under *any* elasticity of demand the data require that the maximum absolute shock be at least 0.038 and the root mean square shock be at least 0.017. These conclusions rely only on equation (1) and the sign restriction that $\theta^D \leq 0$, and may be of direct economic interest.

We can also consider bounds on the size of supply shocks. A major source of shocks to the world supply of grain is variation in agricultural yields due to the weather (Roberts and Schlenker

[6] Muhammad et al. (2011) estimate a model of food demand using country-level data from 2005. Alexandratos and Bruinsma (2012, pp. 56-57) use cross-country variation to determine the relationship between calorie demand and per-capita expenditure in 2005/2007. Several of the models summarized in Valin et al. (2014, p. 56) use the studies by Muhammad et al. (2011) and Alexandratos and Bruinsma (2012) as source information on the income elasticity of demand for food.

2013a). Roberts and Schlenker (2013a) construct an annual measure of yield shocks.[7] The maximum absolute value of the yield shock over the sample period is 0.057, and the root mean square value of the yield shock is 0.024. Allowing for shocks that do not act through yield (e.g., changes in growing area), we consider bounds B^S on supply shocks in $[0, 0.20]$ for $k = \infty$ and in $[0, 0.06]$ for $k = 2$.

Figure 3 depicts the implications of the contemplated bounds B^S for the elasticity of supply $\theta^S \in \overline{\Theta}_S = \mathbb{R}_{\geq 0}$. The structure parallels that of Figure 2. The contemplated bounds are again informative. All of the contemplated bounds imply that supply is inelastic, $\theta^S < 1$. Roberts and Schlenker (2013a, Table 1, Column 2c) estimate that the elasticity of supply is $\hat{\theta}^S_{RS} = 0.097$ with a confidence interval of $[0.060, 0.134]$, also depicted in the plot. A bound of $B^S = 0.12$ on the maximum shock implies a supply elasticity of at most 0.130. The same bound on the elasticity arises from a bound of $B^S = 0.043$ on the root mean squared shock. A maximum absolute shock of at least 0.096 and a root mean square shock of at least 0.041 are necessary to rationalize the data.

Roberts and Schlenker (2013a) devote attention to the "multiplier" $\left(\left|\theta^D\right| + \theta^S\right)^{-1}$, which governs the effect on equilibrium prices of an exogenous change in quantity. Roberts and Schlenker (2013a) conclude that the estimated multiplier is economically substantial. We can determine the implications of bounds B^D, B^S for any known function $\gamma\left(\theta^D, \theta^S\right)$, such as $\gamma\left(\theta^D, \theta^S\right) = \left(\left|\theta^D\right| + \theta^S\right)^{-1}$,[8] by forming the set

$$\hat{\Gamma}_k\left(B^D, B^S\right) = \left\{\gamma\left(\theta^D, \theta^S\right) : \theta^D \in \hat{\Theta}_k\left(B^D\right) \cap \overline{\Theta}_D, \theta^S \in \hat{\Theta}_k\left(B^S\right) \cap \overline{\Theta}_S\right\}.$$

Appendix Figure 2 shows that the bounds we contemplate are informative in that they imply a large multiplier. Roberts and Schlenker (2013a, Table 1, Column 2c) estimate that the multiplier has a value of 6.31 with a confidence interval of $[4.6, 9.1]$. A bound of $B^D = 0.07$ on the maximum demand shock coupled with a bound of $B^S = 0.12$ on the maximum supply shock implies a lower bound on the multiplier of 3.97.

[7] We use the definition of the yield shock underlying Roberts and Schlenker's (2013a) Table 1, Column 2c.

[8] Another prominent example is the function $\gamma\left(\theta^D, \theta^S\right) = \theta^S\left(\left|\theta^D\right| + \theta^S\right)^{-1}$, which determines how the incidence of a tax is shared between consumers and producers (see, e.g., Weyl and Fabinger 2013).

4 Discussion and Extensions

Nonlinear relationship to the input. Our analysis focuses on a linear relationship between the output and the input, as assumed by Roberts and Schlenker (2013a) in our main application. In some settings we may be interested in nonlinear relationships of the form

$$q_t = q(p_t) + \varepsilon_t \tag{3}$$

for $q(\cdot)$ some unknown function. For any two periods $s < t$ we can write

$$q_t - q_s = \theta_{s,t}(p_t - p_s) + \varepsilon_t - \varepsilon_s$$

where

$$\theta_{s,t} = \frac{q(p_t) - q(p_s)}{p_t - p_s}$$

when $p_t \neq p_s$ and $\theta_{s,t}$ is defined arbitrarily otherwise. Thus if $p_t \neq p_s$ then $\theta_{s,t}$ describes the average slope of $q(\cdot)$ between p_s and p_t, and if $q(\cdot)$ is everywhere differentiable then by the mean value theorem $q'(c) = \theta_{s,t}$ for some c strictly between p_s and p_t.

It is immediate that

$$\{\theta_{s,t} \in \mathbb{R} : |\varepsilon_t - \varepsilon_s| \leq B\} = \left[\frac{q_t - q_s}{p_t - p_s} - \frac{B}{|p_t - p_s|}, \frac{q_t - q_s}{p_t - p_s} + \frac{B}{|p_t - p_s|}\right] \tag{4}$$

whenever $p_s \neq p_t$. The set given in equation (4) has the same structure as the interval $\hat{\Theta}(B)$ in the linear case with $T = 2$.

Appendix Figure 3 depicts the set given in equation (4) (intersected with the relevant sign restrictions) for pairs $(t-1, t)$ and for reference values of B in our application to Roberts and Schlenker (2013a). In 80 percent of years t the analysis implies that demand is inelastic between years $t-1$ and t in the sense that $\left|\theta_{t-1,t}\right| < 1$. In 39 percent of years t the analysis implies that supply is inelastic between years $t-1$ and t.

We may also be interested in a summary of the average slopes such as the mean $\overline{\theta} = M_1\left(\vec{\theta}\right)$ of the average slopes $\vec{\theta} = (\theta_{1,2}, ..., \theta_{T-1,T})$ between adjacent periods. We can write that

$$\Delta q_t = \overline{\theta}\Delta p_t + \left(\theta_{t-1,t} - \overline{\theta}\right)\Delta p_t + \Delta\varepsilon_t.$$

By the Minkowski inequality we have that

$$M_k\left(\left|\left(\vec{\theta}-\overline{\theta}\right)\circ\Delta p+\Delta\varepsilon\right|\right)\leq M_k\left(\left|\left(\vec{\theta}-\overline{\theta}\right)\circ\Delta p\right|\right)+M_k\left(\left|\Delta\varepsilon\right|\right).$$

Therefore if we are prepared to impose a bound $M_k\left(\left|\left(\vec{\theta}-\overline{\theta}\right)\circ\Delta p\right|\right)\leq V$ on the deviation of the average slopes from $\overline{\theta}$ and a bound $M_k\left(\left|\Delta\varepsilon\right|\right)\leq B$ on the size of the shocks, then we can say that the set of possible values of $\overline{\theta}$ is contained in $\hat{\Theta}_k\left(B+V\right)$.

Lastly, we may wish to restrict the form of $q\left(\cdot\right)$. Appendix Figure (4) depicts bounds on the average slopes $\theta_{t-1,t}$ between adjacent periods in our application to Roberts and Schlenker (2013a) under the assumption that $q\left(\cdot\right)$ is polynomial of known degree and under reference bounds on the maximum shock $\hat{M}_\infty\left(\left|\Delta\varepsilon\right|\right)$. Even allowing for a sixth-degree polynomial, in many periods the bounds depicted in Appendix Figure 4 are meaningfully tighter than those depicted in Appendix Figure 3.

Nonseparable relationship to a possibly nonscalar unobserved factor. A further relaxation of our model can be written as

$$q_t=\tilde{q}\left(p_t,\varepsilon_t\right)\tag{5}$$

where ε_t may now be non-scalar or even infinite-dimensional. The model in equation (5) can accommodate any functional relationship between q_t and p_t, including relationships that depend on the time period t.[9]

For any two periods $s<t$ we can now write

$$q_t-q_s=\tilde{\theta}_{s,t}\left(p_t-p_s\right)+\tilde{\varepsilon}_{t,t}-\tilde{\varepsilon}_{t,s}$$

where

$$\tilde{\theta}_{s,t}=\frac{\tilde{q}\left(p_t,\varepsilon_s\right)-\tilde{q}\left(p_s,\varepsilon_s\right)}{p_t-p_s}$$

when $p_t\neq p_s$ and $\tilde{\theta}_{s,t}$ is defined arbitrarily otherwise, and where

$$\tilde{\varepsilon}_{t,t}-\tilde{\varepsilon}_{t,s}=\tilde{q}\left(p_t,\varepsilon_t\right)-\tilde{q}\left(p_t,\varepsilon_s\right).$$

Thus if $p_t\neq p_s$ then $\tilde{\theta}_{s,t}$ describes the average slope of $\tilde{q}\left(\cdot,\varepsilon_s\right)$ between p_s and p_t, fixing the

[9]Fixing any such relationship $q_t=\tilde{q}_t\left(p_t,\zeta_t\right)$ for ζ_t an unobserved factor, let $\varepsilon_t=\left(\zeta_t,t\right)$ and define $\tilde{q}\left(\cdot,\cdot\right)$ so that $\tilde{q}\left(p_t,\varepsilon_t\right)=\tilde{q}_t\left(p_t,\zeta_t\right)$ for all ζ_t and t.

unobserved factor at ε_s. The difference $\tilde{\varepsilon}_{t,t} - \tilde{\varepsilon}_{t,s}$ describes the effect on the output of changing the unobserved factor from ε_s to ε_t, fixing the input at p_t.

If we are prepared to bound $|\tilde{\varepsilon}_{t,t} - \tilde{\varepsilon}_{t,s}|$ above by some amount B, then the resulting bounds on $\tilde{\theta}_{s,t}$ follow an analogous structure to the set in equation (4). In the context of our application to Roberts and Schlenker (2013a), this means that the intervals depicted in Appendix Figure 3 can be interpreted as showing the bounds on $\tilde{\theta}_{s,t}$ implied by reference bounds B on the changes in quantity demanded or quantity supplied at given prices p_t between periods $t-1$ and t.

Connection to orthogonality restrictions. Let z_t be some observed variable transformed so that $M_1(\Delta z) = 0$ and $M_2(\Delta z) = 1$.[10] Consider a restriction of the form

$$|M_1(\Delta\varepsilon(\theta) \circ \Delta z)| \leq C \tag{6}$$

where $C \geq 0$ is a scalar. An orthogonality restriction is such a restriction that takes $C = 0$.[11]

Restrictions of the form in (6) are related to those we consider in the sense that, from the Cauchy-Schwarz inequality and the fact that Δz is standardized,

$$(M_1(\Delta\varepsilon(\theta) \circ \Delta z))^2 \leq (M_2(\Delta\varepsilon(\theta)))^2.$$

Hence $M_2(\Delta\varepsilon(\theta)) = \hat{M}_2(\theta) \leq B$ implies that $|M_1(\Delta\varepsilon(\theta) \circ \Delta z)| \leq B$.

As a further connection, recall that $\breve{\theta}_k = \arg\min_\theta \hat{M}_k(\theta)$ for $k \in (1, \infty)$. By Lemma 1, $\breve{\theta}_k$ is the unique solution to

$$\frac{1}{T-1}\sum_{t=2}^{T} \Delta\varepsilon_t(\theta)\,|\Delta\varepsilon_t(\theta)|^{k-2}\,\Delta p_t = 0. \tag{7}$$

Observe that for $k = 2$ and Δp_t standardized, equation (7) is equivalent to an orthogonality restriction with $\Delta z_t = \Delta p_t$.

Connection to cross-equation restrictions. Let $\Delta\varepsilon_t^D(\theta^D) = \Delta q_t^D - \theta^D \Delta p_t^D$ and $\Delta\varepsilon_t^S(\theta^S) = \Delta q_t^S - \theta^S \Delta p_t^S$, and assume in the spirit of static competitive equilibrium that $\Delta q_t^D = \Delta q_t^S = \Delta q_t$ and

[10] Beginning with a variable $\tilde{z}_t$ we can take $z_t = M_2(\Delta\tilde{z} - M_1(\Delta\tilde{z})J)^{-1}(\tilde{z}_t - (t-1)M_1(\Delta\tilde{z}))$, for J a conformable vector of ones.

[11] When $C = 0$, the inequality in (6) implies that $\theta = M_1(\Delta q \circ \Delta z)/M_1(\Delta p \circ \Delta z)$ when this ratio—the linear instrumental-variables estimator—is well-defined. In practice restrictions of this form are more often applied to the distribution of unobserved shocks than to their realization.

$\Delta p_t^D = \Delta p_t^S = \Delta p_t$.[12] Then

$$\left\{\theta^D, \theta^S : M_k\left(\left|\Delta\varepsilon^D\left(\theta^D\right)\right|\right) \leq B^D, M_k\left(\left|\Delta\varepsilon^S\left(\theta^S\right)\right|\right) \leq B^S\right\} = \hat{\Theta}_k\left(B^D\right) \times \hat{\Theta}_k\left(B^S\right).$$

Intuitively, because any pair $\left(\theta^D, \theta^S\right) \in \hat{\Theta}_k\left(B^D\right) \times \hat{\Theta}_k\left(B^S\right)$ is consistent with the data, and the data are consistent with equilibrium, any pair $\left(\theta^D, \theta^S\right) \in \hat{\Theta}_k\left(B^D\right) \times \hat{\Theta}_k\left(B^S\right)$ must also be consistent with equilibrium. In this sense there is no further information about θ^D to be obtained by placing a bound B^S on the size of the shocks $\Delta\varepsilon^S\left(\theta^S\right)$, and vice versa.

The situation is different if we are prepared to restrict the relationship between the shocks $\Delta\varepsilon_t^D\left(\theta^D\right)$ and the shocks $\Delta\varepsilon_t^S\left(\theta^S\right)$. For example, in the context of the world market for food grain, we may be prepared to assume that the major source of demand shocks (say, income) is not strongly related to the major source of supply shocks (say, weather). For illustration, suppose that $M_1(\Delta q) = M_1(\Delta p) = 0$ and take the restriction that

$$\left|M_1\left(\Delta\varepsilon^D\left(\theta^D\right) \circ \Delta\varepsilon^S\left(\theta^S\right)\right)\right| \leq R. \tag{8}$$

If $R = 0$ then

$$\left(\theta^D - \breve{\theta}_2\right)\left(\theta^S - \breve{\theta}_2\right) = \left(\left(\frac{\hat{s}_{qp}}{\sqrt{\hat{s}_{pp}\hat{s}_{qq}}}\right)^2 - 1\right)\frac{\hat{s}_{qq}}{\hat{s}_{pp}}$$

which is analogous to Leamer (1981, equation 6). If $\theta^S \geq 0$ and $\theta^D \leq 0$, then, again following Leamer (1981), if $\breve{\theta}_2 < 0$, then $\theta^D \leq \breve{\theta}_2$, and if $\breve{\theta}_2 > 0$, then $\theta^S \geq \breve{\theta}_2$.

Interpretation of bounds on slope under tight bounds on shocks. Our approach assumes that the bound B holds, $\hat{M}_k(\theta) \leq B$, but not that it is tight, $\hat{M}_k(\theta) = B$. Tightness seems unlikely to arise in practice, but it can nevertheless be instructive to note how the sets we characterize behave under tightness.

Corollary 2. *If $M_\infty(|\Delta\varepsilon(\theta)|) = B$, and in particular there are s,t such that $\Delta p_s, \Delta p_t \neq 0$, $\Delta\varepsilon_s = B\,\mathrm{sgn}(-\Delta p_s)$, and $\Delta\varepsilon_t = B\,\mathrm{sgn}(\Delta p_t)$, then $\left|\hat{\Theta}_\infty(B)\right| = 1$.*

If $M_k(|\Delta\varepsilon(\theta)|) = B$ for some $k \in (1,\infty)$, then either $\underline{\theta}_k(B) = \theta$ or $\overline{\theta}_k(B) = \theta$, or both if $B = \underline{B}_k$.

These properties follow directly from the definitions in Propositions 1 and 2.

[12] In our application, the quantity supplied and quantity demanded need not be equal at a point in time (and likewise for the supply price and the demand price) because grain can be stored and planting decisions are made in advance (Roberts and Schlenker 2013a).

No variation in the input. We have assumed throughout that $p_t \neq p_{t+1}$ for some $t < T$. If this fails, any bound that is compatible with the data is uninformative. More precisely, if $\Delta p = 0$, then $\hat{\Theta}_k(B) = \mathbb{R}$ if $M_k(|\Delta q|) \leq B$ and $\hat{\Theta}_k(B) = \emptyset$ otherwise. Thus in this case $\mathscr{B}(k, \mathbb{R}) = [M_k(|\Delta q|), \infty)$.

5 Conclusions

We show that *a priori* bounds on the size of economic shocks can imply a restriction on a slope parameter in an economic model. In an application to the demand and supply of food grain, we show that economically motivated bounds on the size of shocks imply informative bounds on parameters of interest.

References

Alexandratos, Nikos and Jelle Bruinsma. 2012. World agriculture towards 2030/2050: The 2012 revision. *ESA Working Paper No. 12-03*. Accessed at <http://www.fao.org/3/a-ap106e.pdf> in May 2020.

Antolín-Díaz, Juan and Juan F. Rubio-Ramírez. 2018. Narrative sign restrictions for SVARs. *American Economic Review* 108(10): 2802-29.

Feenstra, Robert C. 1994. New product varieties and the measurement of international prices. *American Economic Review* 84(1): 157-177.

Feenstra, Robert C. and David E. Weinstein. 2017. Globalization, markups, and US welfare. *Journal of Political Economy* 125(4): 1040-1074.

Johnson, D. Gale. 1999. The growth of demand will limit output growth for food over the next quarter century. *Proceedings of the National Academy of Sciences* 96(11): 5915-5920.

Leamer, Edward E. 1981. Is it a demand curve, or is it a supply curve? Partial identification through inequality constraints. *The Review of Economics and Statistics* 63(3): 319-327.

Leontief, Wassily. 1929. Ein Versuch zur statistischen Analyse von Angebot und Nachfrage. *Weltwirtschaftliches Archiv* 30: 1-53.

Logan, Trevon D. 2006. Nutrition and well-being in the late nineteenth century. *Journal of Economic History* 66(2): 313-341.

Ludvigson, Sydney C., Sai Ma, and Serena Ng. 2020. Shock restricted structural vector-autoregressions. *National Bureau of Economic Research Working Paper* No. 23225.

MacKay, Alexander and Nathan H. Miller. 2019. Estimating models of supply and demand: Instruments and covariance restrictions. *Harvard Business School Working Paper* No. 19-051.

Matzkin, Rosa L. 2007. Chapter 73: Nonparametric identification. In James J. Heckman and Edward E. Leamer (eds.), *Handbook of Econometrics* Volume 6, Part B: 5307-5368. Amsterdam: Elsevier.

Morgan, Mary S. 1990. *The History of Econometric Ideas.* Cambridge: Cambridge University Press.

Muhammad, Andrew, James L. Seale, Jr., Birgit Meade, and Anita Regmi. 2011. International evidence on food consumption patterns: An update using 2005 International Comparison Program data. *Economic Research Technical Bulletin* No. 1929. Accessed at <https://www.ers.usda.gov/webdocs/publications/47579/7637_tb1929.pdf?v=45> in May 2020.

Mullin, Wallace P. and Christopher M. Snyder. Forthcoming. A simple method for bounding the elasticity of growing demand with applications to the analysis of historic antitrust cases. *American Economic Journal: Microeconomics*.

Roberts, Michael J. and Wolfram Schlenker. 2013a. Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate. *American Economic Review* 103(6): 2265-95.

Roberts, Michael J. and Wolfram Schlenker. 2013b. Replication data for: Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate. Accessed at <https://doi.org/10.3886/E112674V1> in December 2019.

Robbins, Herbert and Cun-Hui Zhang. 1986. Maximum likelihood estimation in regression with uniform errors. *Institute of Mathematical Statistics Lecture Notes-Monograph Series* 8: 365-85.

Tamer, Elie. 2010. Partial identification in econometrics. *Annual Review of Economics* 2: 167-195.

Uhlig, Harald. 2017. Shocks, sign restrictions, and identification. In Bo Honoré, Ariel Pakes, Monika Piazzesi, and Larry Samuelson (eds.), *Advances in Economics and Econometrics: Eleventh World Congress of the Econometric Society* Volume 2, Chapter 4: 95-127. Cambridge: Cambridge University Press.

Valin, Hugo, Ronald D. Sands, Dominique van der Mensbrugghe, Gerald C. Nelson, Helal Ahammad, Elodie Blanc, Benjamin Bodirsky, Shinichiro Fujimori, Tomoko Hasegawa, Petr Havlik, Edwina Heyhoe, Page Kyle, Daniel Mason-D'Croz, Sergey Paltsev, Susanne Rolinski, Andrzej Tabeau, Hans van Meijl, Martin von Lampe, and Dirk Willenbockel. 2014. The future of food demand: understanding differences in global economic models. *Agricultural Economics* 45(1): 51-67.

Weyl, E. Glen and Michal Fabinger. 2013. Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy* 121(3): 528-583.

Working, Elmer J. 1927. What do statistical "demand curves" show? *Quarterly Journal of Economics* 41(2): 212-235.

World Bank. 2019a. Population, total. *World Development Indicators. The World Bank Group*. Accessed on October 2, 2019 at <https://data.worldbank.org/indicator/SP.POP.TOTL>.

World Bank. 2019b. GDP (current US$). *World Development Indicators. The World Bank Group*. Accessed on October 24, 2019 at <https://data.worldbank.org/indicator/NY.GDP.MKTP.CD>.

Wright, Philip G. 1915. Reviews: Moore's economic cycles. *Quarterly Journal of Economics* 29(3): 631-641.

Figure 1: Illustration of a Two-Period Dataset

Panel A: Price and Quantity in Levels

Panel B: Quantity and Price in First Differences

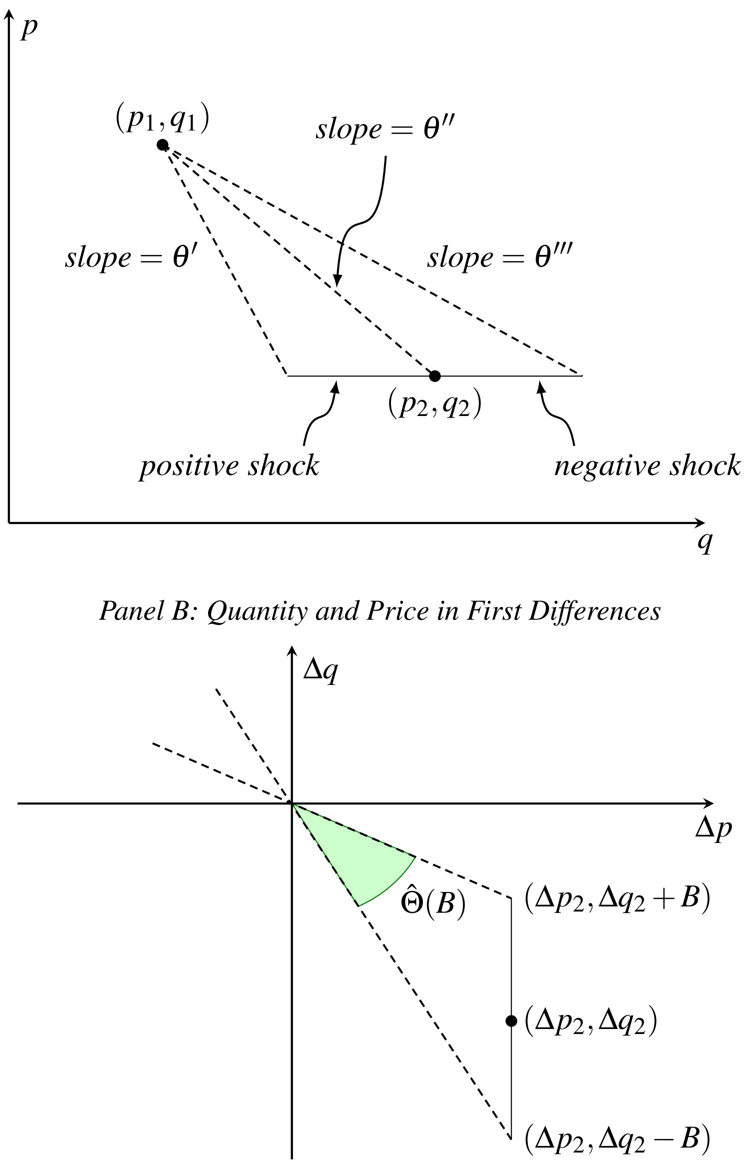

Notes: Panel A illustrates the interpretation of a two-period dataset $\{(p_1, q_1), (p_2, q_2)\}$ under the model in equation (1), with the price depicted on the vertical axis. Panel B illustrates the implications of a bound on the size of the shock of the form in equation (2) for the same two-period dataset depicted in Panel A, with the change in price depicted on the horizontal axis.

Figure 2: Bounds on Shocks to Demand for Grain

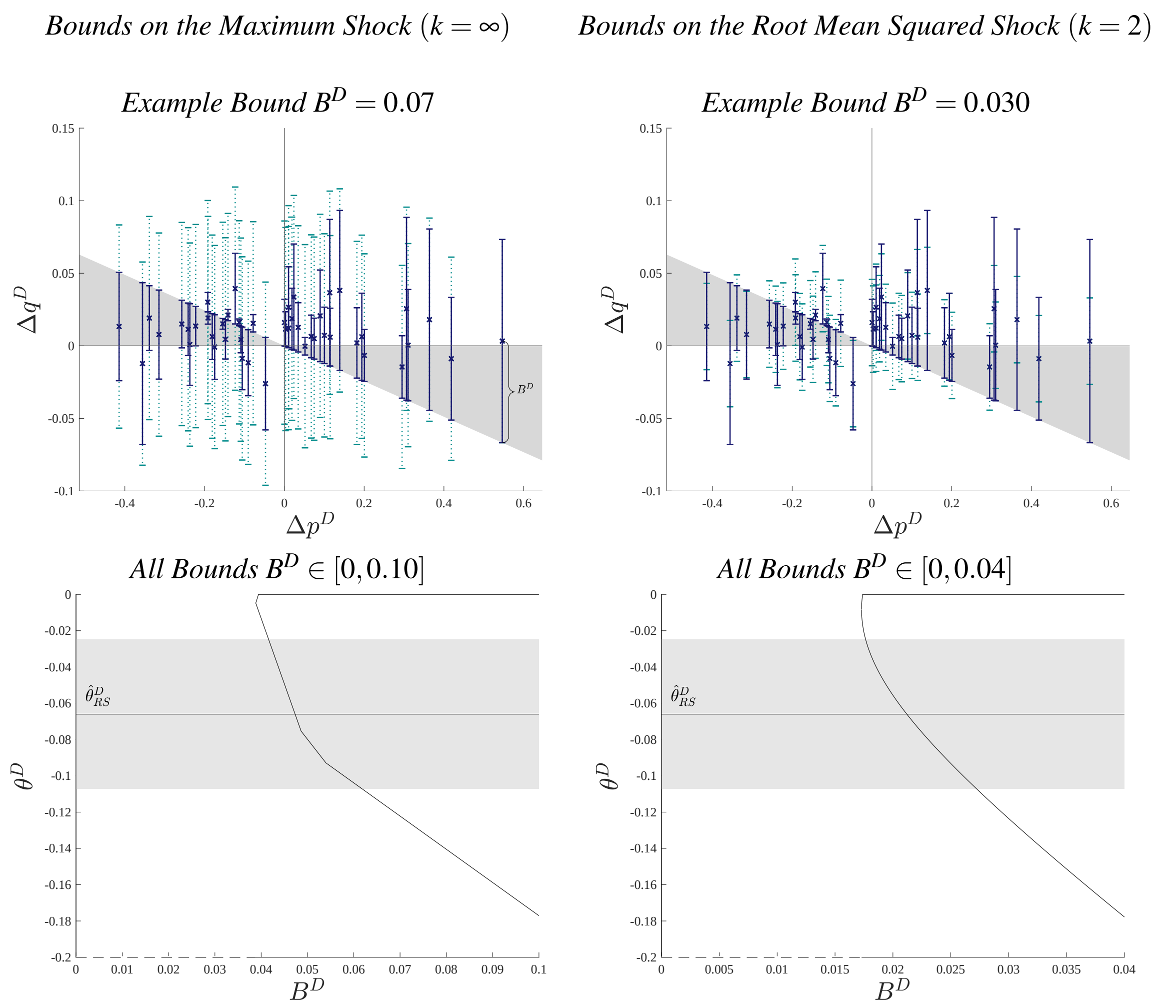

Note: The plots illustrate implications of bounds on the size of shocks to the demand for grain in the application of Roberts and Schlenker (2013a). The first column of plots considers bounds B^D on the maximum value of the shock ($k = \infty$), and the second column of plots considers bounds B^D on the root mean squared shock ($k = 2$). In each column, the first row of plots depicts a scatterplot of the data $\left\{ \left(\Delta p_t^D, \Delta q_t^D \right) \right\}_{t=2}^T$ along with a shaded region showing the demand functions consistent with a particular choice of B^D. The value $B^D = 0.030$ for $k = 2$ is chosen to imply the same minimum slope θ^D as the value $B^D = 0.07$ for $k = \infty$. The dotted interval has radius B^D and the solid interval has radius given by the maximum absolute value of the shock for the given period t consistent with B^D. The second row of plots depicts the interval $\hat{\Theta}_k(B) \cap \overline{\Theta}_D$. The horizontal line depicts the main estimate $\hat{\theta}^D_{RS}$ of the demand elasticity in Roberts and Schlenker (2013a), with the 95% confidence interval pictured as the shaded region. The solid portion of the x-axis corresponds to the bounds $B^D \in \mathscr{B}\left(k, \overline{\Theta}^D\right)$ that are compatible with the data. We define p_t^D as the log of the average current-month futures price of grains delivered in year t, measured in 2010 US dollars per calorie, and q_t^D as log quantity of grains consumed in the world in year t, measured in calories per capita.

Figure 3: Bounds on Shocks to Supply of Grain

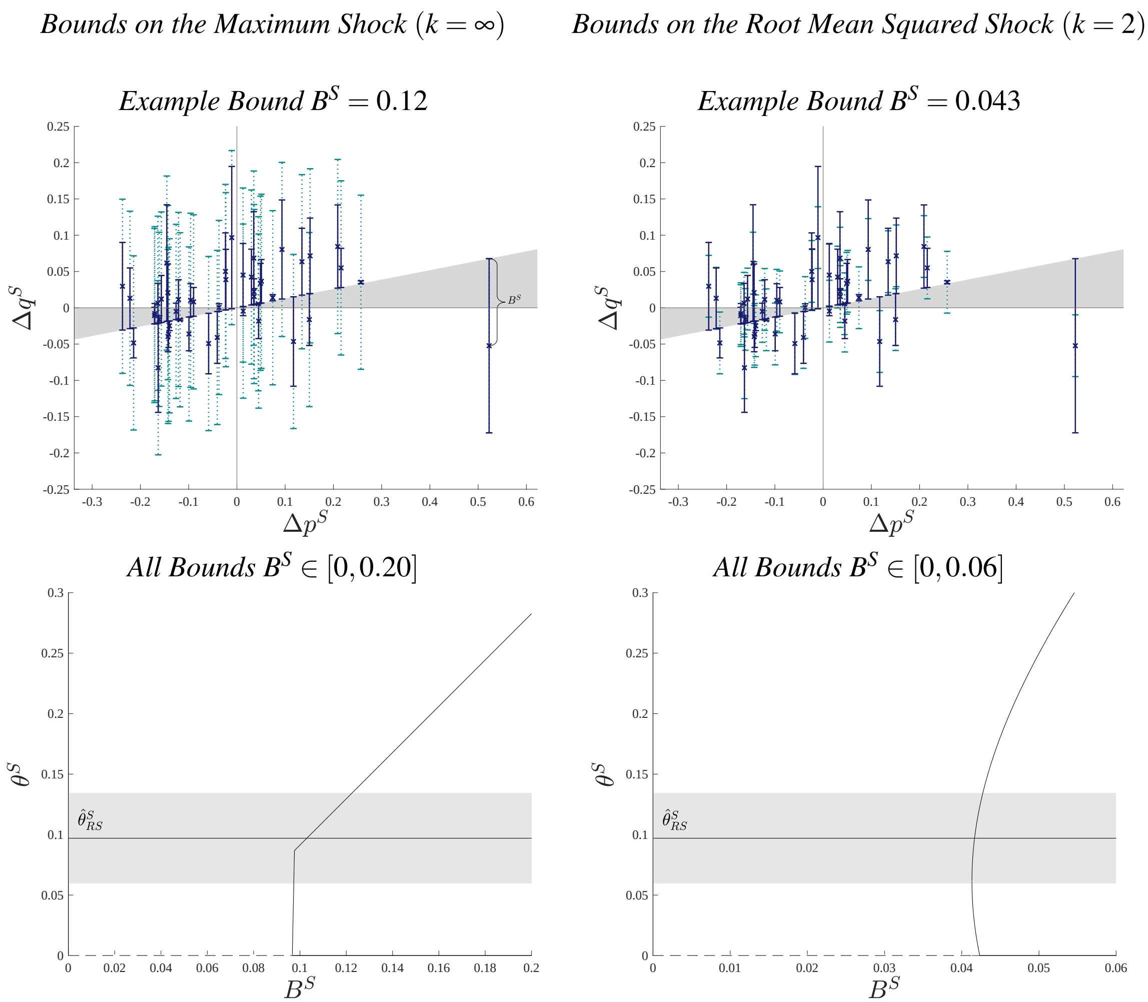

Note: The plots illustrate implications of bounds on the size of shocks to the supply of grain in the application of Roberts and Schlenker (2013a). The first column of plots considers bounds B^S on the maximum value of the shock ($k = \infty$), and the second column of plots considers bounds B^S on the root mean squared shock ($k = 2$). In each column, the first row of plots depicts a scatterplot of the data $\left\{\left(\Delta p_t^S, \Delta q_t^S\right)\right\}_{t=2}^{T}$ along with a shaded region showing the supply functions consistent with a particular choice of B^S. The value $B^S = 0.043$ for $k = 2$ is chosen to imply the same maximum slope θ^S as the value $B^S = 0.12$ for $k = \infty$. The dotted interval has radius B^S and the solid interval has radius given by the maximum absolute value of the shock for the given period t consistent with B^S. The second row of plots depicts the interval $\hat{\Theta}_k(B) \cap \overline{\Theta}_S$. The horizontal line depicts the main estimate $\hat{\theta}^S_{RS}$ of the supply elasticity in Roberts and Schlenker (2013a), with the 95% confidence interval pictured as the shaded region. The solid portion of the x-axis line corresponds to the bounds $B^S \in \mathscr{B}\left(k, \overline{\Theta}^S\right)$ that are compatible with the data. We define p_t^S as the log of the average one-year-ahead futures price of grains delivered in year t, measured in 2010 US dollars per calorie, and q_t^S as log quantity of grains produced in the world in year t, measured in calories per capita.

A Proofs of Results Stated in the Text

Proof of Proposition 1

We have that

$$\hat{M}_{\infty}(\theta) = \max_{t \in \{2,\ldots,T\}} \left(|\Delta q_t - \theta \Delta p_t|\right).$$

Therefore $\hat{M}_{\infty}(\theta) \leq B$ if and only if

$$-B \leq \Delta q_t - \theta \Delta p_t \leq B.$$

for all t. For a given t, if $\Delta p_t = 0$ this condition is equivalent to

$$\Delta q_t \in [-B, B],$$

whereas if $\Delta p_t \neq 0$ it is equivalent to

$$\theta \in \left[\frac{\Delta q_t}{\Delta p_t} - \frac{B}{|\Delta p_t|}, \frac{\Delta q_t}{\Delta p_t} + \frac{B}{|\Delta p_t|}\right].$$

Therefore if $B < |\Delta q_t|$ for some t with $\Delta p_t = 0$ then $\hat{\Theta}_{\infty}(B) = \emptyset$. So take $B \geq \max_{\{t:\Delta p_t = 0\}} |\Delta q_t|$. Let

$$\begin{aligned}
\underline{\theta}_{\infty}(B) &= \max_{\{t:\Delta p_t \neq 0\}} \left\{\frac{\Delta q_t}{\Delta p_t} - \frac{B}{|\Delta p_t|}\right\} \\
\overline{\theta}_{\infty}(B) &= \min_{\{t:\Delta p_t \neq 0\}} \left\{\frac{\Delta q_t}{\Delta p_t} + \frac{B}{|\Delta p_t|}\right\}.
\end{aligned}$$

If $\underline{\theta}_{\infty}(B) > \overline{\theta}_{\infty}(B)$ then $\hat{\Theta}_{\infty}(B) = \emptyset$; otherwise $\hat{\Theta}_{\infty}(B) = \left[\underline{\theta}_{\infty}(B), \overline{\theta}_{\infty}(B)\right]$. Notice that $\underline{\theta}_{\infty}(B)$ is continuous and strictly decreasing in B with $\lim_{B\to\infty} \underline{\theta}_{\infty}(B) = -\infty$ and that $\overline{\theta}_{\infty}(B)$ is continuous and strictly increasing in B with $\lim_{B\to\infty} \overline{\theta}_{\infty}(B) = \infty$. Notice further that

$$\begin{aligned}
\underline{\theta}_{\infty}(0) &= \max_{\{t:\Delta p_t \neq 0\}} \left\{\frac{\Delta q_t}{\Delta p_t}\right\} \\
\overline{\theta}_{\infty}(0) &= \min_{\{t:\Delta p_t \neq 0\}} \left\{\frac{\Delta q_t}{\Delta p_t}\right\}
\end{aligned}$$

and therefore that $\underline{\theta}_{\infty}(0) \geq \overline{\theta}_{\infty}(0)$. Therefore there is a unique solution $\tilde{B} \geq 0$ to $\underline{\theta}_{\infty}\left(\tilde{B}\right) = \overline{\theta}_{\infty}\left(\tilde{B}\right)$. The proposition then follows immediately.

Proof of Lemma 1

We proceed by establishing several elementary properties of the function $\hat{M}_k(\theta)$:

$$\hat{M}_k(\theta) = \left(\frac{1}{T-1}\sum_{t=2}^{T}|\Delta q_t - \theta \Delta p_t|^k\right)^{1/k}$$

for $k \in (1, \infty)$.

Property (i). $\hat{M}_k(\theta)$ is continuous in θ for all $\theta \in \mathbb{R}$.

This property follows because $\hat{M}_k(\theta)$ is a composite of continuous elementary operations.

Property (ii). $\lim_{\theta\to-\infty}\hat{M}_k(\theta) = \lim_{\theta\to\infty}\hat{M}_k(\theta) = \infty$.

Observe that for t' such that $\Delta p_{t'} \neq 0$,

$$\lim_{\theta\to-\infty}|\Delta q_{t'} - \theta\Delta p_{t'}|^k = \lim_{\theta\to\infty}|\Delta q_{t'} - \theta\Delta p_{t'}|^k = \infty$$

whereas for t'' such that $\Delta p_{t''} = 0$,

$$\lim_{\theta\to-\infty}|\Delta q_{t''} - \theta\Delta p_{t''}|^k = \lim_{\theta\to\infty}|\Delta q_{t''} - \theta\Delta p_{t''}|^k = |\Delta q_{t''}|^k.$$

The property then follows immediately because $\lim_{x\to\infty} x^{1/k} = \infty$ for $k > 0$, and by assumption $\Delta p_t \neq 0$ for some $t \in \{2, ..., T\}$.

Property (iii). $\left(\hat{M}_k(\theta)\right)^k$ is strictly convex in θ on $\mathbb{R}$.

We have that

$$\left(\hat{M}_k(\theta)\right)^k = \left(\frac{1}{T-1}\sum_{t=2}^{T}|\Delta q_t - \theta\Delta p_t|^k\right).$$

If $\Delta p_t = 0$ then the function $|\Delta q_t - \theta\Delta p_t|^k$ is trivially weakly convex in θ. Therefore it suffices to show that if $\Delta p_t \neq 0$ then the function $|\Delta q_t - \theta\Delta p_t|^k$ is strictly convex in θ. But this follows from the strict convexity of $|x|^k$ in x on $\mathbb{R}$ for $k > 1$, because if $f(x)$ is strictly convex in x then so is $f(ax+b)$ for $a \neq 0$.

Property (iv). There is $\breve{\theta}_k \in \mathbb{R}$ such that $\breve{\theta}_k = \arg\min_\theta \hat{M}_k(\theta)$.

Pick some $c' > \hat{M}_k(0)$. By properties (i) and (ii), there are at least two solutions to $c' = \hat{M}_k(\theta)$.

By property (iii), there are at most two solutions to $(c')^k = \left(\hat{M}_k(\theta)\right)^k$. Hence there are exactly two solutions to $c' = \hat{M}_k(\theta)$; denote these $\underline{\theta}(c'), \overline{\theta}(c')$, with $\underline{\theta}(c') < \overline{\theta}(c')$. Because the interval $\left[\underline{\theta}(c'), \overline{\theta}(c')\right]$ is compact, by properties (i) and (iii), $\left(\hat{M}_k(\theta)\right)^k$ has a minimum on $\left[\underline{\theta}(c'), \overline{\theta}(c')\right]$ at some unique $\breve{\theta}_k$ on the interior of $\left[\underline{\theta}(c'), \overline{\theta}(c')\right]$. But also by property (iii), $\left(\hat{M}_k(\theta)\right)^k > \left(\hat{M}_k\left(\breve{\theta}_k\right)\right)^k$ for any $\theta \notin \left[\underline{\theta}(c'), \overline{\theta}(c')\right]$, establishing that $\breve{\theta}_k = \arg\min_\theta \left(\hat{M}_k(\theta)\right)^k$ and hence $\breve{\theta}_k = \arg\min_\theta \left(\hat{M}_k(\theta)\right)$.

Property (v). $\hat{M}_k(\theta') > \hat{M}_k(\theta'')$ for any $\theta' < \theta'' < \breve{\theta}_k$ and $\hat{M}_k(\theta') < \hat{M}_k(\theta'')$ for any $\breve{\theta}_k < \theta' < \theta''$.

This is an immediate consequence of property (iii), applying the strict monotonicity of x^k on $\mathbb{R}_{\geq 0}$ for $k \in (1, \infty)$.

Proof of Proposition 2

This follows immediately from Lemma 1.

Proof of Corollary 1

We have that

$$\hat{M}_2(\theta) = \left(\frac{1}{T-1}\sum_{t=2}^{T}(\Delta q_t - \theta\Delta p_t)^2\right)^{1/2}.$$

By Lemma 1, $\hat{M}_2(\theta)$ has a unique global minimizer $\breve{\theta}_2$. Because $\hat{M}_2(\theta)$ is differentiable, the minimizer $\breve{\theta}_2$ must satisfy the first-order condition $\frac{d}{d\theta}\hat{M}_2(\theta)\mid_{\theta=\breve{\theta}_2} = 0$, which is equivalent to the equation

$$\hat{s}_{qp} - \breve{\theta}_2\hat{s}_{pp} = 0$$

and therefore

$$\breve{\theta}_2 = \frac{\hat{s}_{qp}}{\hat{s}_{pp}}$$

because $\hat{s}_{pp} \neq 0$. It also follows that

$$\begin{aligned}\underline{B}_2 = \hat{M}_2\left(\breve{\theta}_2\right) &= \hat{M}_2\left(\frac{\hat{s}_{qp}}{\hat{s}_{pp}}\right) \\ &= \sqrt{\hat{s}_{qq} - \left(\frac{\hat{s}_{qp}}{\hat{s}_{pp}}\right)^2 \hat{s}_{pp}}.\end{aligned}$$

Observe that, by the Cauchy-Schwarz inequality, this expression is real-valued.

Next, by Proposition 2, the bounds $\underline{\theta}_2(B), \overline{\theta}_2(B)$ solve $\hat{M}_2(\theta) = B$ which is equivalent to the quadratic equation

$$\left(\hat{s}_{qq} - B^2\right) - 2\theta\hat{s}_{qp} + \theta^2\hat{s}_{pp} = 0.$$

The roots of this quadratic equation are given by

$$\frac{\hat{s}_{qp}}{\hat{s}_{pp}} \pm \sqrt{\left(\frac{\hat{s}_{qp}}{\hat{s}_{pp}}\right)^2 - \frac{1}{\hat{s}_{pp}}\left(\hat{s}_{qq} - B^2\right)}.$$

Observe that these roots are real-valued whenever $B \geq \underline{B}_2$, thus completing the proof.

Proof of Corollary 2

For the first part, observe that by the definition of $\overline{\theta}_\infty(B)$ and $\underline{\theta}_\infty(B)$ in Proposition 1, it follows that

$$\overline{\theta}_\infty(B) - \underline{\theta}_\infty(B) = \min_{\{t:\Delta p_t \neq 0\}}\left\{\frac{\operatorname{sgn}(\Delta p_t)\,\Delta\varepsilon_t + B}{|\Delta p_t|}\right\} - \max_{\{t:\Delta p_t \neq 0\}}\left\{\frac{\operatorname{sgn}(\Delta p_t)\,\Delta\varepsilon_t - B}{|\Delta p_t|}\right\}.$$

If $\Delta\varepsilon_s = B\operatorname{sgn}(-\Delta p_s)$ for some s such that $\Delta p_s \neq 0$, then the first term is equal to zero. If $\Delta\varepsilon_t = B\operatorname{sgn}(\Delta p_t)$ for some t such that $\Delta p_t \neq 0$, then the second term is equal to zero. But then $\overline{\theta}_\infty(B) - \underline{\theta}_\infty(B) = 0$ and hence $\hat{\Theta}_\infty(B)$ is a singleton as desired.

For the second part, the desired conclusion follows from the fact that $\underline{\theta}_k(B), \overline{\theta}_k(B)$ are the only solutions to the equation $\hat{M}_k(\theta) = B$ and that these solutions coincide when $B = \underline{B}_k$.

Appendix Figure 1: Bounds on Shocks to Demand and Supply of Grain, Varying k

Panel A: Demand for Grain

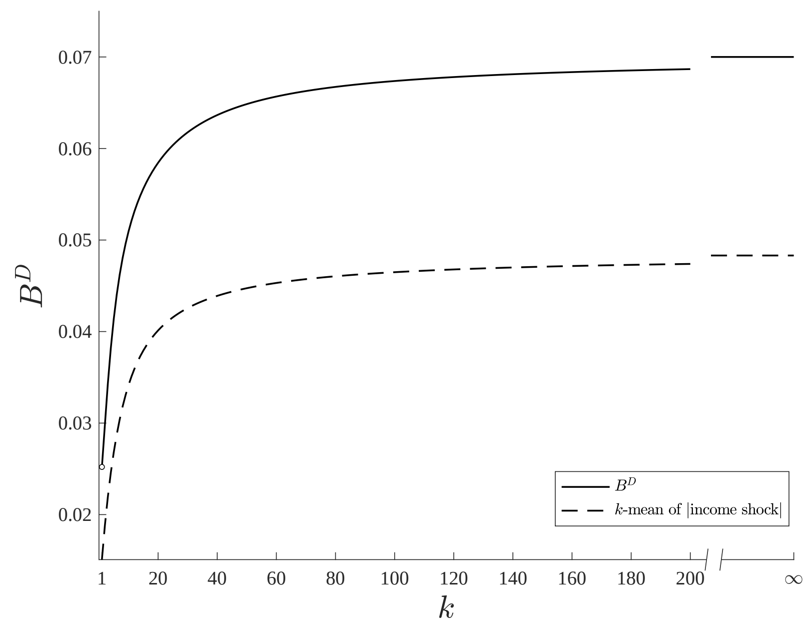

Panel B: Supply of Grain

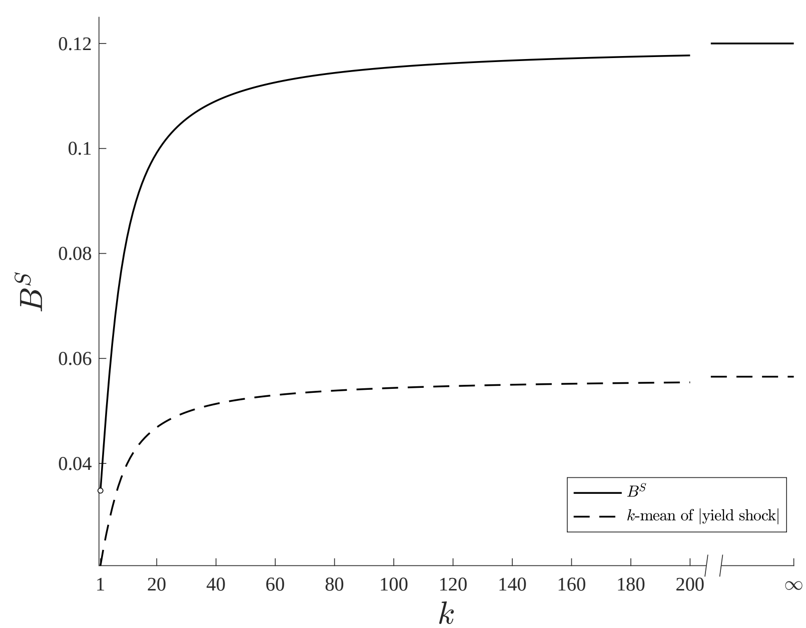

Note: The plots illustrate the bound B on the $k-$mean of the shock that implies a given bound on the slope θ in the application of Roberts and Schlenker (2013a). The solid line in Panel A depicts the bound B^D on the $k-$mean of the absolute value of the demand shock that implies the same lower bound on the demand elasticity θ^D as a bound B^D of 0.07 on the maximum absolute value of the shock. The dashed line in Panel A depicts the $k-$mean $M_k(|0.37\Delta y|)$ of the absolute value of the income shock. The solid line in Panel B depicts the bound B^S on the $k-$mean of the absolute value of the supply shock that implies the same upper bound on the supply elasticity θ^S as a bound B^S of 0.12 on the maximum absolute value of the shock. The dashed line in Panel B depicts the $k-$mean of the absolute value of the yield shock. In both panels, values are plotted for $k \in (1, 200]$ and $k = \infty$.

Appendix Figure 2: Implications of Bounds on Shocks for the Multiplier Parameter

Bounds on the Maximum Shock ($k = \infty$) *Bounds on the Root Mean Squared Shock* ($k = 2$)

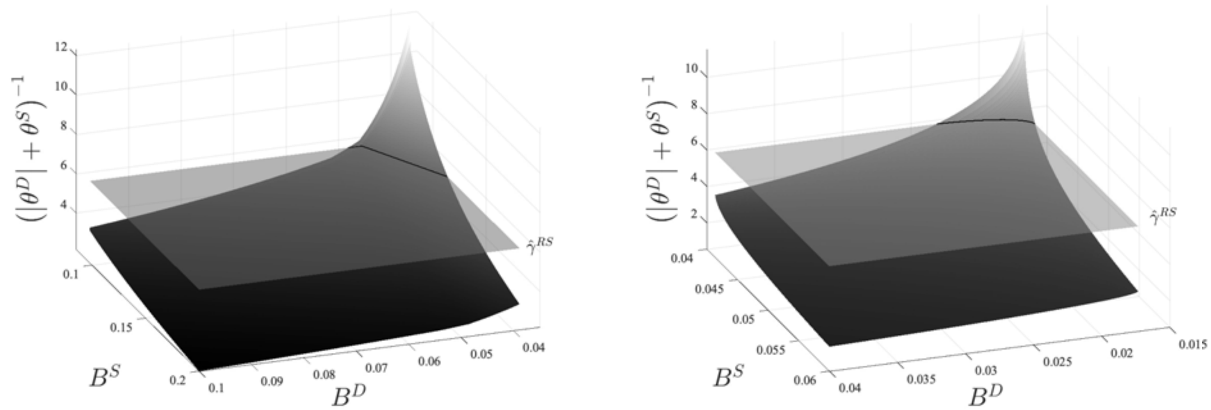

Note: The plots illustrate implications of bounds on the size of shocks to the supply and demand of grain in the application of Roberts and Schlenker (2013a). The left plot considers bounds $B^D \in [0.035, 0.10]$, $B^S \in [0.095, 0.20]$ on the maximum value of the shock ($k = \infty$), and the right plot considers bounds $B^D \in [0.015, 0.04]$, $B^S \in [0.040, 0.06]$ on the root mean squared shock ($k = 2$). Both plots depict the values of the multiplier $\left(\left|\theta^D\right| + \theta^S\right)^{-1}$ compatible with elasticities $\theta^D \in \hat{\Theta}_k\left(B^D\right) \cap \overline{\Theta}_D$, $\theta^S \in \hat{\Theta}_k\left(B^S\right) \cap \overline{\Theta}_S$, i.e. the set $\hat{\Gamma}_k\left(B^D, B^S\right)$ for $\gamma\left(\theta^D, \theta^S\right) = \left(\left|\theta^D\right| + \theta^S\right)^{-1}$. The horizontal plane depicts the main estimate $\hat{\gamma}^{RS}$ of the multiplier in Roberts and Schlenker (2013a).

Appendix Figure 3: Implications of Bounds on Shocks for the Average Elasticity Between Adjacent Years

Panel A: Demand for Grain ($B^D = 0.07$)

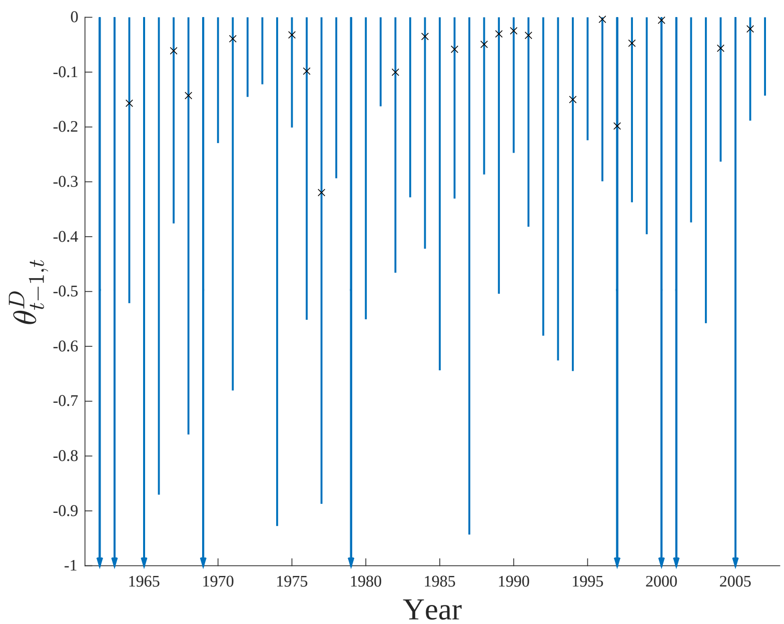

Panel B: Supply of Grain ($B^S = 0.12$)

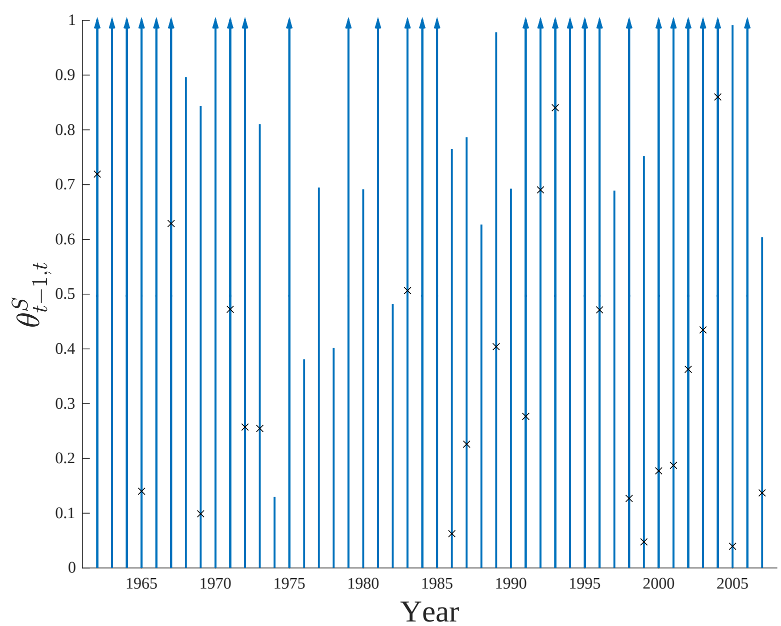

Notes: Panels A and B illustrate, respectively, the bounds on the average elasticity of demand and supply of grain between adjacent years in the application of Roberts and Schlenker (2013a). The depicted bounds intersect those in equation (4) with the relevant sign restriction for each elasticity. We use $B^D = 0.07$ and $B^S = 0.12$ as the bounds on the absolute value of shocks to demand and supply, respectively. Each line segment represents the interval of possible average elasticities between the given year and the preceding year, with an arrow indicating that the interval contains elasticities greater than one, and a crosshatch indicating the value of $\Delta q_t / \Delta p_t$ when contained in the plotted range.

Appendix Figure 4: Implications of Bounds on Shocks for the Average Elasticity Between Adjacent Years, Under Polynomial Restrictions

Panel A: Demand for Grain ($B^D = 0.07$)

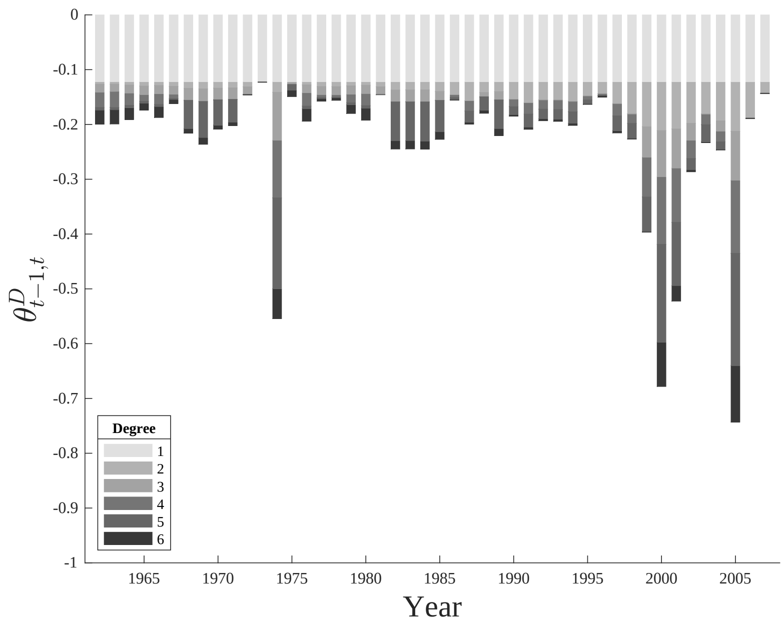

Panel B: Supply of Grain ($B^S = 0.12$)

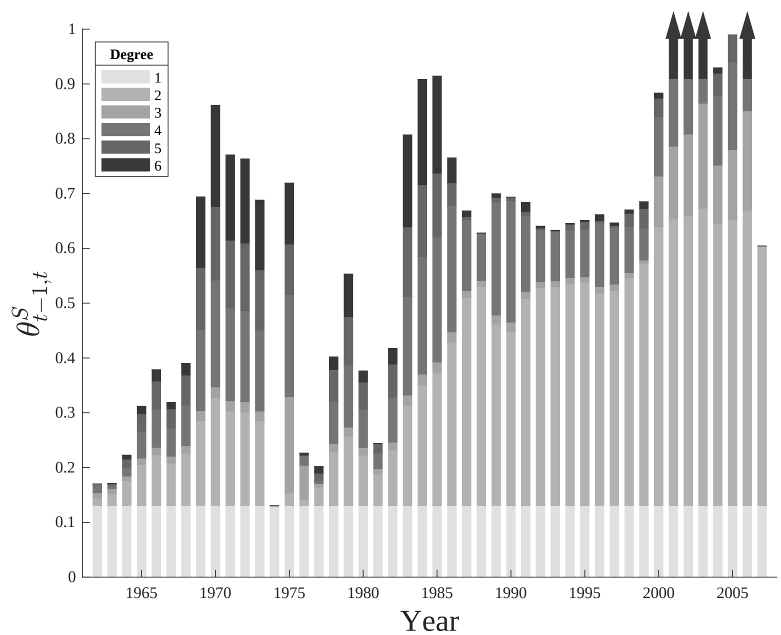

Notes: Panels A and B illustrate, respectively, the bounds on the average elasticity of demand and supply of grain between adjacent years in the application of Roberts and Schlenker (2013a). We use $B^D = 0.07$ and $B^S = 0.12$ as the bounds on the maximum absolute value of shocks to demand and supply, respectively. We assume that the function $q(\cdot)$ defined in equation (3) is a polynomial of known degree whose derivative is nonnegative (Panel A) or nonpositive (Panel B) everywhere on the closed interval from the lowest to the highest observed price. Each line segment represents the interval of possible average elasticities between the given year and the preceding year under the given polynomial degree (from one to six), with an arrow indicating that the interval includes absolute elasticities greater than one.