

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, Revised January 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, Nadav Ben Zeev, 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, Iowa State University, and the International Methods Colloquium for helpful comments. We thank Masao Fukui, Emi Nakamura, and Jón Steinsson for sharing their code and data. 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 additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w27556.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Marco Stenborg Petterson, David G. Seim, and Jesse M. Shapiro. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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, Revised January 2020
JEL No. C22,C32,D41,Q11

ABSTRACT

We study the problem of learning about the effect of one market-level variable (e.g., price) on another (e.g., quantity) in the presence of shocks to unobservables (e.g., preferences). We show that economic intuitions about the plausible size of the shocks can be informative for the parameter of interest. We illustrate with applications to the grain market and the labor market.

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*

January 2021

Abstract

We study the problem of learning about the effect of one market-level variable (e.g., price) on another (e.g., quantity) in the presence of shocks to unobservables (e.g., preferences). We show that economic intuitions about the plausible size of the shocks can be informative for the parameter of interest. We illustrate with applications to the grain market and the labor market.

1 Introduction

Consider the problem of learning about the effect of one observed market-level variable p_t (e.g., log price) on another observed market-level variable q_t (e.g., log quantity demanded) from a finite time series $\{(p_t, q_t)\}_{t=1}^T$ with at least $T \geq 2$ periods. Economists often specify a linear model of the form

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

*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, Nadav Ben Zeev, 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, Iowa State University, and the International Methods Colloquium for helpful comments. We thank Masao Fukui, Emi Nakamura, and Jón Steinsson for sharing their code and data. 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.

where θ is an unknown slope (e.g., the price elasticity of demand) and ε_t is an unobserved factor (e.g., preferences). Models that can be cast into the form in equation (1) include Barro and Redlick’s (2011, equation 1) model of the effect of fiscal policy on economic growth, Fiorito and Zanella’s (2012, equation 3) model of the supply of labor, Roberts and Schlenker’s (2013a, equations 1 and 3) model of the supply and demand for food grains, and Autor et al.’s (2020, equation 2) model of the demand for skill, among many others.

Absent further restrictions, the data are uninformative about the slope θ . Economists often learn about θ by imposing restrictions on the evolution of ε_t , for example that it is unrelated to an observable instrument (e.g., Fiorito and Zanella 2012; Roberts and Schlenker 2013a), or that it is unrelated to p_t after accounting for time trends (e.g., Autor et al. 2020). These restrictions are typically motivated by economic intuitions about the determinants of ε_t .

In this paper we argue that economic intuitions about the size of fluctuations in ε_t can also be informative about θ and can therefore complement other approaches to learning about θ . Suppose, for example, that log prices p_t for a good vary considerably from year to year but log quantities q_t do not. Then a large price elasticity of demand θ implies larger fluctuations in ε_t than does a small price elasticity of demand. Large fluctuations in ε_t may be plausible if the good in question is a particular brand of scarf, preferences for which may change radically from year to year due to advertising campaigns, changes in fashion, etc. Large fluctuations in ε_t may be less plausible if the good in question is a standard agricultural commodity, preferences for which are likely more stable. In this latter case, economic intuitions about the size of fluctuations in ε_t may suggest a smaller value of θ .

Formally, we consider the implications of placing an upper bound $B \geq 0$ on a generalized power mean, with power greater than one, of the vector $(|\Delta\varepsilon_2|, \dots, |\Delta\varepsilon_T|)$ of absolute shocks to the unobserved factor, where $\Delta\varepsilon_t = \varepsilon_t - \varepsilon_{t-1}$ and Δ is the first difference operator. We show that any feasible bound B implies that θ lies in a closed, bounded interval. We provide a computationally tractable characterization of the endpoints of the interval. We further show that some bounds B can be inconsistent with the data, implying that, in some settings, we can place a lower bound on the size of the true shocks even with no knowledge of θ . We show that our approach extends readily to settings with a cross-sectional dimension (e.g., states or demographic groups), thus broadening the set of possible economic applications.

We illustrate the value of incorporating economic intuitions about the plausible size of shocks with two applications. The first is to the price elasticity of demand for staple grains, using the

data and model from Roberts and Schlenker (2013a). The second is to the crowding out of male employment by female employment, using the data and model from Fukui et al. (2020). In both of these applications, the authors impose linear models and approach estimation and inference using orthogonality restrictions with respect to excluded instruments. In both cases, we argue that, even without such restrictions, economically motivated bounds on the size of unobserved shocks are informative for the slope parameter θ . Reasonable bounds on the size of shocks imply bounds on θ that are consistent with, and in some cases even tighter than, the authors' own inferences.

We view the information about θ obtained by bounding the size of shocks as a complement to, rather than a replacement for, the approaches used by the authors. We show how to visualize in a single plot the bounds on θ implied by a range of bounds B on the size of shocks, so that researchers do not need to take a strong stand on the appropriate bound in order to use our approach.

We extend our approach in a few directions. In some cases, researchers may be interested in relaxing the assumption of a linear, separable model. We show how to obtain bounds on an average slope in the case where the model takes the nonlinear form $q_t = q(p_t) + \varepsilon_t$ or the nonseparable form $q_t = \tilde{q}(p_t, \varepsilon_t)$. Researchers may also be interested in learning about a function $\gamma(\theta)$ of the slope parameter θ . We show in an appendix that our approach extends readily to this situation.

The main contributions of this paper are to demonstrate that economic intuitions about the plausible size of shocks to unobservables are available and useful in important applications, and to propose a formal approach to exploiting these intuitions. We think that in cases such as the applications we consider, authors may be leaving useful information on the table by failing to exploit these intuitions.¹

Our formal setup is closely related to a large literature, mainly in electrical engineering and optimal control, that considers bounds on the size of unobservable noise in a system (see, e.g., Walter and Piet-Lahanier 1990; Milanese et al. 1996). The focus of much of this literature is on settings in which, unlike ours, exact computation of parameter bounds is impossible, and approximations are needed. In the paper, we highlight some specific connections between our characterizations and those in this and other related work.

Within economics, proposals to impose restrictions on the variability of unobservables go back at least to Marschak and Andrews (1944; see, e.g., equation 1.37), and are related to (though distinct from) approaches based on bounded support of the outcome variable (e.g., Manski 1990).

¹For example, Roberts and Schlenker (2013a, p. 2277) write, "Price fluctuations are proportionately much larger than quantity fluctuations.... This fact suggests that both demand and supply are inelastic." Roberts and Schlenker (2013a) do not formalize the logic behind this statement or develop its quantitative implications, as we do here.

More broadly, many canonical approaches to identification impose restrictions on the distribution of unobserved variables (see, e.g., Matzkin 2007; Tamer 2010), such as the assumption that the unobservables are uncorrelated with an observed instrument, have a correlation with the observed instrument that can be bounded or otherwise restricted (e.g., Conley et al. 2012; Nevo and Rosen 2012), or are independent of or uncorrelated with one another (e.g., Leamer 1981; Feenstra 1994; Feenstra and Weinstein 2017; MacKay and Miller 2019).² An appendix discusses some connections between these types of approaches and ours.

Our approach is also related to recent proposals to learn about parameters of interest by restricting the realization of unobservables rather than their distribution. In the structural vector autoregression setting, Ben Zeev (2018) considers restrictions on the time-series properties of an unobserved shock including the timing of its maximum value, 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), and 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 price elasticity of demand in a reference period under the assumption that demand is growing over time.³ Though related, none of these sets of restrictions coincides with those we consider here.

In ongoing work, Giacomini et al. (2020) study inferential issues that arise in the presence of restrictions on the realizations of unobservables. Our approach instead characterizes bounds on the parameter of interest that hold with certainty under the given bound on the size of the shocks, so, in common with much of the related engineering literature, issues of probabilistic inference do not arise in our setup.

The remainder of the paper is organized as follows. Section 2 presents our setup and results. Section 3 presents our applications. Section 4 presents an extension to more general functional relationships. Section 5 concludes. An appendix presents proofs of results stated in the text and discusses additional extensions and connections.

²See 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, and a related literature (e.g., Manski 1997) considers the implications of restrictions on functional form, including monotonicity.

³In 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 Setup and Characterization of Sets of Interest

For any D -dimensional vector v and any $k \geq 1$, write the generalized k -mean

$$M_k(v) = \left(\frac{1}{D} \sum_{d=1}^D v_d^k \right)^{1/k},$$

with $M_\infty(v) = \max_d \{v_d\}$ denoting the maximum value of the elements of v and $M_2(v)$ denoting their root mean squared value. Let $|v| = (|v_1|, \dots, |v_D|)$ denote the absolute value of the vector v .

Now for any $k > 1$, let $\hat{M}_k(\theta) = M_k(|\Delta\varepsilon(\theta)|)$ denote the k -mean of the absolute value of the vector $\Delta\varepsilon(\theta) = (\Delta\varepsilon_2(\theta), \dots, \Delta\varepsilon_T(\theta))$, where $\Delta\varepsilon_t(\theta) = \Delta q_t - \theta \Delta p_t$ is the value of the shock to the unobserved factor in period t implied by a given slope θ . Our main object of interest is the set of slopes

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

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

In some applications, we may wish to impose direct restrictions on the possible values of the slope θ , for example that $\theta \leq 0$ in the case of a demand function. To capture these direct restrictions we will suppose that $\theta \in \bar{\Theta} \subseteq \mathbb{R}$, where, for example, $\bar{\Theta} = \mathbb{R}_{\leq 0}$ in the case where we impose that $\theta \leq 0$, and $\bar{\Theta} = \mathbb{R}$ in the case where we impose no direct restrictions. A slope θ is compatible with the restriction that $\hat{M}_k(\theta) \leq B$ and with the direct restrictions if and only if it is contained in $\hat{\Theta}_k(B) \cap \bar{\Theta}$.

Given the model in equation (1), a bound B is compatible with the data, and with the direct restrictions on θ , if and only if $\hat{\Theta}_k(B) \cap \bar{\Theta}$ is nonempty. We let

$$\mathcal{B}(k, \bar{\Theta}) = \{B \in \mathbb{R} : \hat{\Theta}_k(B) \cap \bar{\Theta} \neq \emptyset\} \subseteq \mathbb{R}_{\geq 0}$$

denote the set of bounds B that are compatible with the data and with the direct restrictions on θ .

We assume throughout that $p_t \neq p_{t+1}$ for at least one $t < T$. This condition holds in our applications. If it fails, any bound that is compatible with the data is uninformative.⁴

We begin with the case of $k = \infty$, in which we bound the maximum absolute value of the shock. This case plays an important role in our applications, and yields a particularly simple form for the

⁴Specifically, if $\Delta p = 0$, then $\hat{M}_k(\theta) = M_k(|\Delta q|)$ for all $\theta \in \mathbb{R}$, so $\hat{\Theta}_k(B) = \mathbb{R}$ if $M_k(|\Delta q|) \leq B$ and $\hat{\Theta}_k(B) = \emptyset$ otherwise. Thus in this case $\mathcal{B}(k, \mathbb{R}) = [M_k(|\Delta q|), \infty)$.

sets of interest.

Proposition 1. *Let*

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

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

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

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

All proofs are given in Appendix A. The objects \underline{B}_∞ , $\underline{\theta}_\infty(B)$, and $\bar{\theta}_\infty(B)$ defined in Proposition 1 can be readily calculated on datasets of reasonable size. In the extreme case where the bounds on the shocks are achieved, the limit points $\underline{\theta}_\infty(B)$ and $\bar{\theta}_\infty(B)$ coincide, and $\hat{\Theta}_\infty(B)$ is a singleton.⁵

The objects characterized in Proposition 1 have antecedents in prior work. The interval $\hat{\Theta}_\infty(B)$ solves a special case of Milanese and Belforte's (1982) Problem B. The limit points $\underline{\theta}_\infty(B)$ and $\bar{\theta}_\infty(B)$ of the interval appear in the analysis of the linear regression model with uniformly distributed errors (Robbins and Zhang 1986). Walter and Piet-Lahanier (1996) study the computation of \underline{B}_∞ in a case with multiple unknown slope parameters.

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 unbounded and strictly decreasing on $(-\infty, \check{\theta}_k)$ and unbounded and strictly increasing on $(\check{\theta}_k, \infty)$ for $\check{\theta}_k = \arg \min_\theta \hat{M}_k(\theta)$.*

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

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

⁵More precisely, if $M_\infty(|\Delta \varepsilon(\theta)|) = B$, and in particular there are s, t such that $\Delta p_s, \Delta p_t \neq 0$, $\Delta \varepsilon_s = B \operatorname{sgn}(-\Delta p_s)$, and $\Delta \varepsilon_t = B \operatorname{sgn}(\Delta p_t)$, then $|\hat{\Theta}_\infty(B)| = 1$.

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

Proposition 2 shows that $\mathcal{B}(k, \mathbb{R})$ is a left-bounded interval whose limit point \underline{B}_k can be calculated by minimizing the function $\hat{M}_k(\theta)$. 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, bounded interval whose limit points can be calculated by solving the nonlinear equation $\hat{M}_k(\theta) = B$. By Lemma 1, on either side of $\check{\theta}_k$ and for $B > \underline{B}_k$ the equation is strictly monotone and has a unique solution, which simplifies computation. If we are in the extreme case where the bounds are achieved, i.e. $\hat{M}_k(\theta) = B$ at the true θ , then either $\underline{\theta}_k(B) = \theta$ or $\bar{\theta}_k(B) = \theta$, or both if $\hat{M}_k(\theta) = \underline{B}_k$. 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 $\underline{B}_2, \underline{\theta}_2(B), \bar{\theta}_2(B)$, and $\check{\theta}_2$ defined in Proposition 2 are available in closed form. For any D -dimensional vector $v \in \mathbb{R}^D$, let $\Delta v = (\Delta v_2, \dots, \Delta v_D) \in \mathbb{R}^{D-1}$. 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*

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

Observe that $\check{\theta}_2 = \underline{\theta}_2(\underline{B}_2) = \bar{\theta}_2(\underline{B}_2)$ 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 $\{(\Delta p_t, \Delta q_t)\}_{t=2}^T$.

Our approach extends readily to the case where we observe a finite time series $\{(p_{it}, q_{it})\}_{t=1}^T$ for each of a cross-section of units $i \in \{1, \dots, N\}$, such as countries or states. Let $\Delta \varepsilon_i(\theta) = (\Delta \varepsilon_{i2}(\theta), \dots, \Delta \varepsilon_{iT}(\theta))$, where $\Delta \varepsilon_{it} = \Delta q_{it} - \theta \Delta p_{it}$, and define $\hat{M}_{ik}(\theta) = M_k(|\Delta \varepsilon_i(\theta)|)$ correspondingly. Suppose we are prepared to impose a bound B_i on the size of the shocks in each unit i . If a different slope θ_i is thought to apply to each unit i , so that $q_{it} = \theta_i p_{it} + \varepsilon_{it}$, then we can simply

repeat the exercise described above, defining one set $\hat{\Theta}_{ik}(B_i) = \{\theta_i \in \mathbb{R} : \hat{M}_{ik}(\theta_i) \leq B_i\}$ for each unit i . If a common slope θ is thought to apply to each unit i , so that $q_{it} = \theta p_{it} + \varepsilon_{it}$, then we can form the set $\cap_{i=1}^N \hat{\Theta}_{ik}(B_i)$, which collects those slopes θ that are compatible with the bounds B_i on the size of the shocks in each unit i . Note that this treatment implicitly allows for the possibility that we impose the same bound for all units ($B_i = B$ for all i) or that we impose a bound only for a subset of units ($B_i = \infty$ for some i).

3 Applications

3.1 Price Elasticity of World Demand for Staple Food Grains

Roberts and Schlenker (2013a) estimate the price elasticity of world demand for staple food grains using annual data from 1960 through 2007. We use their code and data (Roberts and Schlenker 2013b), supplemented with data from the World Bank (2019a; 2019b) on annual world population and GDP. From these we construct a time series $\{(p_t^D, q_t^D)\}_{t=1}^T$, where p_t^D is 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 is the the log of the quantity of grains consumed in the world in year t , measured in calories per capita.⁶ We also construct a measure y_t of annual log world GDP per capita in 2010 US dollars.⁷

Roberts and Schlenker (2013a, equation 3) assume that the demand curve takes a log-linear form consistent with equation (1). Roberts and Schlenker (2013a) adopt an instrumental variables approach to estimating the price elasticity of demand θ^D , using the contemporaneous yield shock as an excluded instrument for price. Here we explore what we can learn about the price elasticity of demand by imposing bounds on the size of shocks to demand.

It is reasonable for economists to have intuitions about the size of shocks to world demand for staple grains. The major determinants of world demand for grain in the modern period are population and income (Johnson 1999; Valin et al. 2014). We measure demand on a per capita basis, leaving income as a major determinant. Engel’s Law holds that demand for food is income-inelastic. Forecasts summarized in Valin et al. (2014, Table 3) imply an income elasticity of world food crop demand ranging from 0.09 to 0.37.⁸ Taking the upper end of the range, the income-

⁶We 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.

⁷We deflate to 2010 US dollars using the consumer price index from Roberts and Schlenker (2013b).

⁸The 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

driven shock to log per-capita demand in year t has absolute value $|0.37\Delta y_t|$. The largest value of this shock over the sample period is $M_\infty(|0.37\Delta y|) \approx 0.05$. The root mean squared value is $M_2(|0.37\Delta y|) \approx 0.02$. Shocks substantially larger than these may seem implausible.

Figure 1 illustrates why these intuitions are informative for the price elasticity of demand θ^D . The figure plots the value of the shock $\Delta \varepsilon_t(\theta^D)$ in each year t implied by two benchmark values of θ^D : the point estimate $\hat{\theta}_{RS}^D = -0.066$ given in Roberts and Schlenker (2013a, Table 1, Column 2c), and the value $\theta^D = -1$ implying unit price elasticity. The shocks $\Delta \varepsilon_t(-1)$ to per capita world food grain demand implied by unit price elasticity are, to us, implausible, reaching values as high as 0.55, more than 10 times the largest income-driven shock, and implying that, at constant prices, the world changed its desired consumption of food grains by 55 percent on a per-capita basis in a single year! By contrast, the shocks $\Delta \varepsilon_t(-0.066)$ implied by Roberts and Schlenker's (2013a) point estimate appear much more reasonable.

Following the logic of Section 2, we can directly characterize the implications for the price elasticity θ^D of a given bound B^D on the size of the shocks. Figure 2 illustrates the construction of the bounds on θ^D implied by a bound of $B^D = 0.07$ on the maximum shock. This value of B^D is chosen to be about 1.4 times larger than the largest income-driven shock. The figure depicts a scatterplot of the first-differenced data $\{(\Delta p_t^D, \Delta q_t^D)\}_{t=2}^T$ with a dashed interval of radius $B^D = 0.07$ around each point. On this plot, a demand function is a line through the origin with nonpositive slope $\theta^D \in \bar{\Theta}^D = \mathbb{R}_{\leq 0}$. A bound of $B^D = 0.07$ on the maximum absolute value of the demand shock requires that the demand function pass through all of the dashed intervals. The set of demand functions that do so, i.e. those with $\theta^D \in \hat{\Theta}_\infty(0.07) \cap \bar{\Theta}^D$, are collected in the shaded region.

A bound of $B^D = 0.07$ on the maximum absolute value of the demand shock is informative for the price elasticity of demand θ^D . Such a bound implies that $\theta^D \in \hat{\Theta}_\infty(0.07) \cap \bar{\Theta}^D = [-0.122, 0]$. This interval contains Roberts and Schlenker's (2013a, Table 1, Column 2c) confidence interval of $[-0.107, -0.025]$ fairly tightly. Roberts and Schlenker (2013a) devote substantial attention to discussion of the exogeneity of their instrument and related sensitivity analysis. Our analysis suggests that arguing for a reasonable bound on the size of demand shocks in the grain market provides another way to defend Roberts and Schlenker's (2013a) conclusions.

Not all readers may accept the same bound B^D on the size of the shock. It is therefore appealing to display the implications for the price elasticity θ^D of many possible bounds B^D . Figure 3 shows

implied income elasticities therefore range from $\ln(1.08)/\ln(16000/6700) \approx 0.088$ to $\ln(1.38)/\ln(16000/6700) \approx 0.370$.

that this can be done in a single plot. Panel A depicts the interval $\hat{\Theta}_\infty(B^D) \cap \bar{\Theta}^D$ of elasticities compatible with each bound $B^D \in [0, 0.10]$ on the maximum absolute shock. Panel B depicts the interval $\hat{\Theta}_2(B^D) \cap \bar{\Theta}^D$ of elasticities compatible with each bound $B^D \in [0, 0.04]$ on the root mean squared shock. In each case, we choose the range of bounds so that the largest bound is around twice the size $M_k(|0.37\Delta y|)$ of the income-driven shocks, thus allowing for non-income-driven shocks to demand of about the same size as the income-driven shocks. For comparison, we also depict the point estimate and confidence interval from Roberts and Schlenker (2013a, Table 1, Column 2c). It is interesting that even the largest bounds B^D that we depict are informative, implying a price elasticity of demand smaller than 0.18 in absolute value.

Figure 3 also illustrates the interpretation of the set $\mathcal{B}(k, \bar{\Theta}^D)$, depicted as the solid portion of the x-axes. The data imply that the maximum absolute shock is at least 0.038 (Panel A) and the root mean squared shock is at least 0.017 (Panel B). These implications may be of direct economic interest, and rely only on equation (1) and the sign restriction that $\theta^D \leq 0$.

This application illustrates some aspects of our approach that are worth highlighting. One is that intuitions about the plausible size of shocks can be informed by data other than the time series in question. For example, estimates of the income elasticity of food demand can be informed by comparisons across countries at a point in time.⁹ Another is that the choice of reasonable bounds can be contextual (as well as subjective). For example, in earlier historical periods the income elasticity of food demand was likely larger (see, e.g., Logan 2006), which might suggest looser bounds on shocks to world demand.

Appendix B includes several extensions of our analysis of the grain market. Appendix B.1 develops bounds on the price elasticity of supply θ^S of staple grains based on bounds B^S on the size of shocks to supply, illustrated in Appendix Figure 1. Appendix B.2 considers the possibility that the researcher may be interested in a function of the elasticities θ^D and θ^S , illustrated in Appendix Figure 2 with an application to the “multiplier” parameter studied in Roberts and Schlenker (2013a). Lastly, Appendix Figure 3 illustrates the role of the choice of k by showing how the value of B^D needed to obtain a given lower bound on the price elasticity, and the value of $M_k(|0.37\Delta y|)$, vary with k .

⁹Muhammad 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.

3.2 Crowding Out of Male Employment by Female Employment

Fukui et al. (2020) estimate the crowding out θ^C of male employment by female employment using data on US states for 1970 and 2016. We use the code and data underlying Fukui et al.’s (2020) Table 3, provided to us by the authors. From these we obtain the cross-section $\{(\Delta f_i, \Delta m_i)\}_{i=1}^N$ consisting of the change Δf_i in the female employment-to-population ratio and the change Δm_i in the male employment-to-population ratio in each state i between 1970 and 2016.

Fukui et al. (2020, equation 5) specify a homogenous linear relationship between Δm_i and Δf_i of the form $\Delta m_i = \theta^C \Delta f_i + \Delta \varepsilon_i$.¹⁰ Fukui et al. (2020) take an instrumental variables approach to estimating the crowding out parameter θ^C , using various shifters of female employment as excluded instruments for Δf_i . Here we explore what we can learn about the crowding out parameter by imposing bounds on the size of shocks to male employment.

During the study period, female labor force participation expanded greatly. Across US states, the median change in the female employment-to-population ratio was 0.27, and the largest change was 0.44. The massive cultural and technological forces that contributed to this trend have been widely studied and documented (see, for example, the review in Greenwood et al. 2017). Although prime-age male labor force participation declined over this period (e.g., Binder and Bound 2019), the forces affecting male participation were arguably less dramatic than those affecting female participation.¹¹ Shocks to male employment on the same scale as those to female employment may therefore seem implausible.

Imposing that the absolute shock to male employment-to-population is less than or equal to some value B in all states means that $\theta^C \in \cap_{i=1}^N \hat{\Theta}_i(B)$, where the choice of k is now irrelevant as we only observe a single difference ($T = 2$) in each state i .¹² Imposing that crowding out is nonpositive means that $\theta^C \in \bar{\Theta} = \mathbb{R}_{\leq 0}$. Figure 4 depicts the interval $\cap_{i=1}^N \hat{\Theta}_i(B) \cap \bar{\Theta}$ for all $B \in [0, 0.23]$, or up to just over half of the largest change in Δf_i across all states. The figure shows that the bounds are informative. Suppose, for example, that we impose that no state’s male employment would have changed by more than $B = 0.14$ in the absence of changes in female employment. This bound is about half the median change in Δf_i and a bit under a third of the maximum change in Δf_i . Then the depicted set is $\cap_{i=1}^N \hat{\Theta}_i(0.14) \cap \bar{\Theta} = [-0.33, 0]$, which is contained within the confidence interval

¹⁰To cast this into the form in equation (1), suppose that male employment in each state obeys $m_{it} = \theta_i^C f_{it} + \varepsilon_{it}$, with $\theta_i^C = \theta^C$ for all i .

¹¹Juhn and Potter (2006, p. 32) write, “The biggest story in labor force participation rates in recent decades involves the labor force attachment of women.”

¹²That is, for any feasible bound B , we have that $\hat{\Theta}_{ik}(B) = \hat{\Theta}_i(B)$ for all $k > 1$.

of $[-0.35, 0.09]$ from Fukui et al.'s (2020, Table 3, Column 2) preferred specification, as is the set $\cap_{i=1}^N \hat{\Theta}_i(0.14) = [-0.33, 0.03]$. With bounds $B < 0.13$, the interval $\cap_{i=1}^N \hat{\Theta}_i(B) \cap \bar{\Theta}$ implies that there must be crowding out, i.e. that $\theta^C < 0$. The interval $\cap_{i=1}^N \hat{\Theta}_i(B) \cap \bar{\Theta}$ contains Fukui et al.'s (2020, Table 3, Column 2) preferred point estimate $\hat{\theta}_{FNS}^C$ unless B is less than 0.09.

It is also instructive to look at the shocks to male employment implied by a given value of θ^C . Suppose, for example, that $\theta^C = -0.5$, implying substantial crowding out. Then to rationalize the data, six states (Iowa, Wisconsin, Alaska, Nebraska, South Dakota, and Minnesota) must have experienced *positive* shocks to male employment of between 10 and 20 percentage points, and one (North Dakota) must have experienced a positive shock of over 20 percentage points. Recall that these shocks represent the implied change in male employment absent a change in female employment. Although there were some important positive influences on male employment over this period (such as the fracking boom, see, e.g., Bartik et al. 2019), such large, positive shocks to male employment across so many states seem difficult to square with the prevailing economic understanding of influences on male labor force participation over this period (e.g., Binder and Bound 2019).

Fukui et al. (2020, Section 4.3) devote significant attention to discussion and analysis of sources of possible correlation between their instrument and unobserved shocks to male employment. Our analysis shows that arguing that shocks to male employment were meaningfully smaller than shocks to female employment over the study period, or that very negative values of θ^C imply implausibly large shocks to male employment, provides another way to defend Fukui et al.'s (2020) conclusions.

4 Extension to More General Functional Relationships

4.1 Nonlinear Relationship

In the applications of Section 3, the authors assume a linear relationship between the observed variables of interest, as in equation (1). In some settings we may be interested in nonlinear relationships of the form

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

for $q(\cdot)$ an unknown function.

In this case a bound on the size of the shock can be used to derive a bound on the average slope

$\theta_{s,t}$ between any two periods $s < t$ with $p_s \neq p_t$. In particular, we can write

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

where

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

If $q(\cdot)$ is everywhere differentiable, then by the mean value theorem $\theta_{s,t} = q'(c)$ for some c strictly between p_s and p_t .

If we are prepared to impose an upper bound of B on the size of the shock between periods s and t , then we can obtain a bound on the average slope $\theta_{s,t}$ via the relation

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

The interval given in equation (4) has the same structure as the interval $\hat{\Theta}_k(B)$, for any k , in the linear case with $T = 2$.

The interval in equation (4) is informative in our application to the price elasticity of world demand for staple foods. Panel A of Figure 5 depicts the bounds on the average price elasticity $\theta_{t-1,t}^D$ if we assume that the shock in each year is no greater than $B^D = 0.07$, as in Panel A of Figure 3. In 80 percent of years t the analysis implies that demand is price-inelastic on average between years $t - 1$ and t in the sense that $\theta_{t-1,t}^D > -1$. Appendix Figure 5 depicts analogous bounds for the average crowding out parameter in the setting of Section 3.2.

Even more informative statements are possible if we are prepared to assume that $q(\cdot)$ is polynomial of known degree. Panel B of Figure 5 shows that even allowing for a six-degree polynomial, a substantial generalization of linearity, in 89 percent of years we can conclude that $\theta_{t-1,t}^D > -0.3$. Appendix Figure 4 depicts analogous bounds for the average price elasticity of supply $\theta_{t-1,t}^S$.

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

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

By the Minkowski inequality we have that

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

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

4.2 Nonlinear, Nonseparable Relationship

A further relaxation of the model in equation (1) can be written as

$$q_t = \tilde{q}(p_t, \varepsilon_t) \quad (5)$$

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 .¹³

It is again possible to bound the average slope $\tilde{\theta}_{s,t}$ between any two periods $s < t$ with $p_s \neq p_t$, where now

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

with

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

and

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

Here $\tilde{\theta}_{s,t}$ describes the average slope of $\tilde{q}(\cdot, \varepsilon_s)$ between p_s and p_t , fixing the unobserved factor at ε_s . The shock $\tilde{\varepsilon}_{t,t} - \tilde{\varepsilon}_{t,s}$ describes the effect on q_t of changing the unobserved factor from ε_s to ε_t , fixing the value of p_t .

If we are prepared to impose an upper bound of B on the size of $|\tilde{\varepsilon}_{t,t} - \tilde{\varepsilon}_{t,s}|$, then the resulting bounds on $\tilde{\theta}_{s,t}$ follow an analogous structure to the set in equation (4).¹⁴ In the context of our application to the price elasticity of world demand for staple food grains, this means that the intervals depicted in Panel A of Figure 5 can be interpreted as showing the bounds on $\tilde{\theta}_{s,t}$ implied by a bound of $B^D = 0.07$ on the change in quantity demanded at given prices p_t between periods $t-1$ and t .

¹³Fixing any such relationship $q_t = \tilde{q}_t(p_t, \zeta_t)$ for ζ_t an unobserved factor, let $\varepsilon_t = (\zeta_t, t)$ and define $\tilde{q}(\cdot, \cdot)$ so that $\tilde{q}(p_t, \varepsilon_t) = \tilde{q}_t(p_t, \zeta_t)$ for all ζ_t and t .

¹⁴Specifically,

$$\{\tilde{\theta}_{s,t} \in \mathbb{R} : |\tilde{\varepsilon}_{t,t} - \tilde{\varepsilon}_{t,s}| \leq B\} = \left[\frac{q_t - q_s}{p_t - p_s} - \frac{B}{|p_t - p_s|}, \frac{q_t - q_s}{p_t - p_s} + \frac{B}{|p_t - p_s|} \right].$$

5 Conclusions

Though often unobserved, shocks to economic variables have economic meaning, and economists will in some situations have intuitions about their size. We formalize an approach to using these intuitions to bound a slope parameter in an economic model. We illustrate the utility of the approach with applications to two important economic markets, where we argue that the approach can usefully complement existing approaches to inference. We extend the approach to the case of nonlinear or nonseparable models and show that it remains informative.

References

- Alexandratos, Nikos and Jelle Bruinsma. 2012. World agriculture towards 2030/2050: The 2012 revision. *ESA Working Paper No. 12-03*. Accessed in May 2020 at <http://www.fao.org/3/a-ap106e.pdf>.
- Antolín-Díaz, Juan and Juan F. Rubio-Ramírez. 2018. Narrative sign restrictions for SVARs. *American Economic Review* 108(10): 2802-29.
- Autor, David, Claudia Goldin and Lawrence F. Katz. 2020. Extending the race between education and technology. *AEA Papers and Proceedings* 110: 347-351.
- Barro, Robert J. and Charles J. Redlick. 2011. Macroeconomic effects from government purchases and taxes. *Quarterly Journal of Economics* 126(1): 51-102.
- Bartik, Alexander W., Janet Currie, Michael Greenstone, and Christopher R. Knittel. 2019. The local economic and welfare consequences of hydraulic fracturing. *American Economic Journal: Applied Economics* 11(4): 105-155.
- Ben Zeev, Nadav. 2018. What can we learn about news shocks from the late 1990s and early 2000s boom-bust period? *Journal of Economic Dynamics and Control* 87: 94-105.
- Binder, Ariel J. and John Bound. 2019. The declining labor market prospects of less-educated men. *Journal of Economic Perspectives* 33(2): 163-190.
- Conley, Timothy G., Christian B. Hansen, and Peter E. Rossi. 2012. Plausibly exogenous. *Review of Economics and Statistics* 94(1): 260-272.
- 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.
- Fiorito, Riccardo and Giulio Zanella. 2012. The anatomy of the aggregate labor supply elasticity. *Review of Economic Dynamics* 15(2): 171-187.
- Fukui, Masao, Emi Nakamura and Jón Steinsson. 2020. Women, wealth effects, and slow recoveries. *NBER Working Paper No. 25311*.
- Giacomini, Raffaella, Toru Kitagawa, and Matthew Read. 2020. Identification and inference under narrative restrictions . Slides from the Gary Chamberlain Online Seminar in Econometrics, presented October 30, 2020. Accessed on December 21, 2020 at <https://www.chamberlainseminar.org/past-seminars/autumn-2020#h.qxodbm9sauk5>.
- Greenwood, Jeremy, Nezih Guner, and Guillaume Vandenbroucke. 2017. Family economics writ large. *Journal of Economic Literature* 55(4): 1346-1434.
- 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.
- Juhn, Chinhui and Simon Potter. 2006. Changes in labor force participation in the United States. *Journal of Economic Perspectives* 20(3): 27-46.
- Leamer, Edward E. 1981. Is it a demand curve, or is it a supply curve? Partial identification through inequality constraints. *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*.
- Manski, Charles F.. 1990. Nonparametric bounds on treatment effects. *American Economic Review* 80(2): 319-323.
- Manski, Charles F.. 1997. Monotone treatment response. *Econometrica* 65(6): 1311-1334.
- Marschak, Jacob and William H. Andrews, Jr.. 1944. Random simultaneous equations and the theory of production. *Econometrica* 12(3/4): 143-205.
- 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. Ams-

terdam: Elsevier.

- Milanese, Mario and Gustavo Belforte. 1982. Estimation theory and uncertainty intervals evaluation in presence of unknown but bounded errors: Linear families of models and estimators. *IEEE Transactions on Automatic Control* 27(2): 408-414.
- Milanese, Mario, John Norton, H el ene Piet-Lahanier, and Eric Walter (eds.). 1996. *Bounding Approaches to System Identification*. Boston: Springer.
- 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 in May 2020 at <https://www.ers.usda.gov/webdocs/publications/47579/7637_tb1929.pdf?v=45>.
- 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*.
- Nevo, Aviv and Adam M. Rosen. 2012. Identification with imperfect instruments. *Review of Economics and Statistics* 94(3): 659-671.
- 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 in December 2019 at <<https://doi.org/10.3886/E112674V1>>.
- 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 Honor e, 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 Mensbrugge, Gerald C. Nelson, Helal Aham-


- mad, 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.
- Walter, Eric and H el ene Piet-Lahanier. 1990. Estimation of parameter bounds from bounded-error data: A survey. *Mathematics and Computers in Simulation* 32(5-6): 449-468.
- Walter, Eric and H el ene Piet-Lahanier. 1996. Chapter 12: Recursive robust minimax estimation. In Mario Milanese, John Norton, H el ene Piet-Lahanier, and Eric Walter (eds.), *Bounding Approaches to System Identification* 183-197. Boston: Springer.
- 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.
- 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>>.

Figure 1: Implied Shocks to World Demand for Food Grain Under Different Elasticities



Notes: The plot depicts the shocks to demand for grain implied by different values of the price elasticity of demand in the setting of Roberts and Schlenker (2013a) described in Section 3.1. Each series corresponds to the shocks $\Delta \varepsilon_t(\theta^D)$ to demand implied by a given value of the price elasticity of demand θ^D . We depict the shocks corresponding to the main point estimate of Roberts and Schlenker (2013a), denoted $\hat{\theta}_{RS}^D$, and corresponding to the case of unit elasticity, $\theta^D = -1$.


Figure 2: Constructing Bounds on an Elasticity from Bounds on Shocks




Notes: The plot illustrates the construction of bounds on the price elasticity of demand from bounds on the size of shocks to the demand for grain in the setting of Roberts and Schlenker (2013a) described in Section 3.1. The plot depicts a scatterplot of the data $\{(\Delta p_t^D, \Delta q_t^D)\}_{t=2}^T$ along with a shaded region showing the demand functions consistent with the upper bound $B^D = 0.07$ on the maximum shock. The dotted interval has radius B^D and the solid interval has its radius given by the maximum absolute value of the shock for the given period t consistent with B^D .

Figure 3: Implications of Bounds on Shocks to World Demand for Food Grain

Panel A: All Bounds $B^D \in [0, 0.10]$ on the Maximum Shock ($k = \infty$)




Panel B: All Bounds $B^D \in [0, 0.04]$ on the Root Mean Squared Shock ($k = 2$)



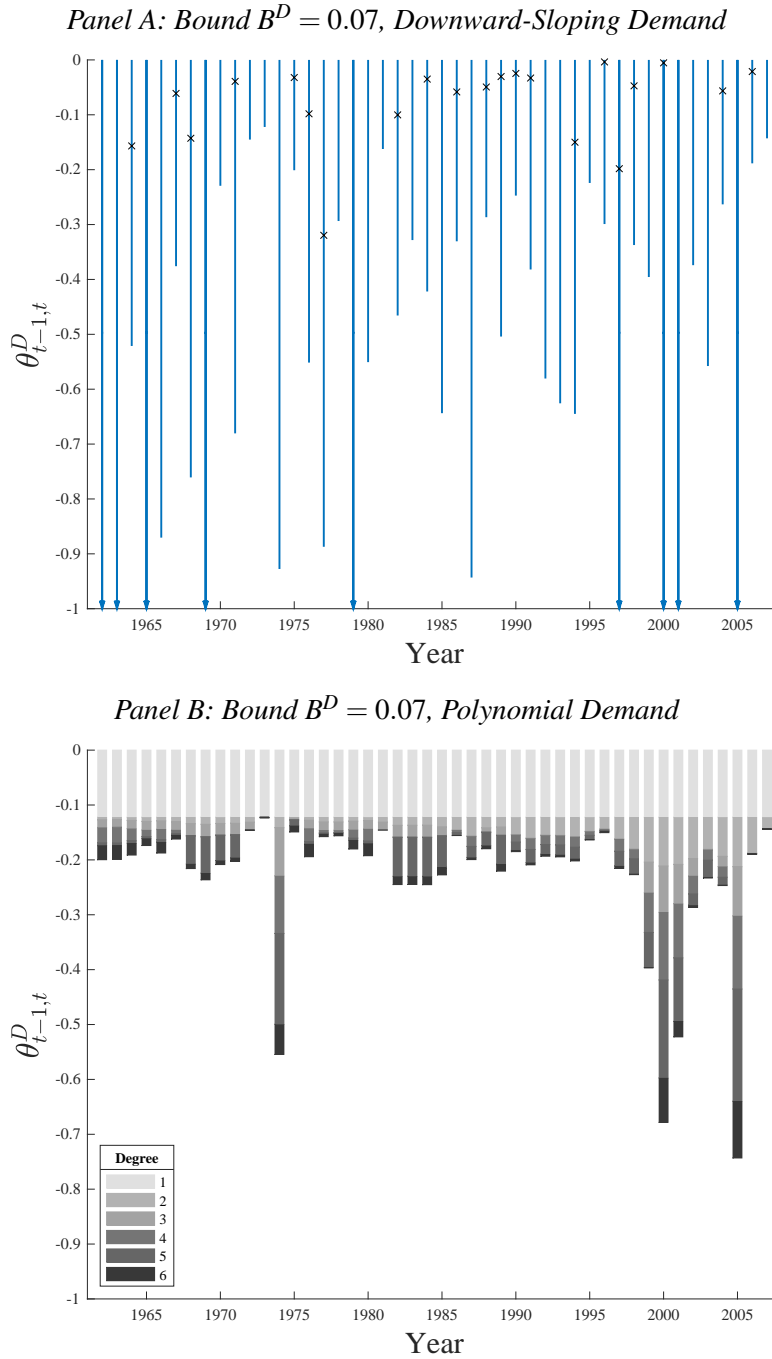
Notes: The plots illustrate implications of bounds on the size of shocks to the demand for grain in the setting of Roberts and Schlenker (2013a) described in Section 3.1. Panel A depicts the interval $\hat{\Theta}_\infty(B^D) \cap \bar{\Theta}^D$ implied by bounds $B^D \in [0, 0.10]$ on the maximum shock, where $\bar{\Theta}^D = \mathbb{R}_{\leq 0}$. The dashed vertical line is at twice the maximum absolute income-driven shock $M_\infty(|0.37\Delta y|)$. Panel B depicts the interval $\hat{\Theta}_2(B^D) \cap \bar{\Theta}^D$ implied by bounds $B^D \in [0, 0.04]$ on the root mean squared shock. The dashed vertical line is at twice the root mean squared income-driven shock $M_2(|0.37\Delta y|)$. In each plot, the horizontal line depicts the main estimate $\hat{\theta}_{RS}^D$ of the price elasticity of demand 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 $B^D \in \mathcal{B}(k, \bar{\Theta}^D)$ that are compatible with the data.

Figure 4: Implications of Bounds on Shocks to Male Employment



Notes: The plot illustrates implications of bounds on the size of shocks to male employment in the setting of Fukui et al. (2020) described in Section 3.2. The plot depicts the interval $\cap_{i=1}^N \hat{\Theta}_i(B) \cap \bar{\Theta}$ implied by bounds $B \in [0, 0.23]$ on the shock where $\bar{\Theta} = \mathbb{R}_{\leq 0}$. The dashed vertical line is at half the maximum absolute change in female employment $\max_i |\Delta f_i|$. The horizontal line depicts the main estimate $\hat{\theta}_{FNS}^C$ of the crowding out of male employment by female employment in Fukui et al. (2020), with the 95% confidence interval pictured as the shaded region. The solid portion of the x-axis corresponds to the bounds $B \in \mathcal{B}(k, \bar{\Theta})$ that are compatible with the data.

Figure 5: Relaxing Linearity of the Demand Function



Notes: Each plot depicts bounds on the average price elasticity of demand $\theta_{t-1,t}^D$ between each pair of adjacent years based on the assumption that the shock to world demand for staple food grains is no greater than $B^D = 0.07$. In Panel A, the depicted bounds are formed by intersecting the set in equation (4) with the sign restriction that the average price elasticity is nonpositive. Each line segment represents the interval of possible average price elasticities, with an arrow indicating that the interval contains price elasticities less than -1 , and a crosshatch indicating the value of $\Delta q_t / \Delta p_t$ when contained in the plotted range. In Panel B, we further impose that the function $q(\cdot)$ is a polynomial of known degree whose derivative is nonpositive everywhere on the closed interval from the lowest to the highest observed price. Each line segment represents the interval of possible average price elasticities under the given polynomial degree (from one to six).

A Proofs of Results Stated in the Text

Proof of Proposition 1

We have that

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

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

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

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

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

and therefore that $\underline{\theta}_\infty(0) \geq \bar{\theta}_\infty(0)$. Therefore there is a unique solution $\bar{B} \geq 0$ to $\underline{\theta}_\infty(\bar{B}) = \bar{\theta}_\infty(\bar{B})$. 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). $\lim_{\theta \rightarrow -\infty} \hat{M}_k(\theta) = \lim_{\theta \rightarrow \infty} \hat{M}_k(\theta) = \infty$.

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

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

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

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

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

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

We have that

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

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 $\check{\theta}_k \in \mathbb{R}$ such that $\check{\theta}_k = \arg \min_{\theta} \hat{M}_k(\theta)$.

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

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

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 $\check{\theta}_2$. Because $\hat{M}_2(\theta)$ is nonnegative and is differentiable in θ when $\hat{M}_2(\theta) > 0$, either $\hat{M}_2(\check{\theta}_2) = 0$ or $\hat{M}_2(\check{\theta}_2) > 0$ and $\frac{d}{d\theta} \hat{M}_2(\theta) |_{\theta=\check{\theta}_2} = 0$. In either case we have that

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

Because $\hat{s}_{pp} \neq 0$ we can also say that

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

It then follows that

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

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

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

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

The roots of this quadratic equation are given by

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

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

B Extensions of Analysis of World Market for Staple Food Grains

B.1 Price Elasticity of World Supply of Staple Food Grains

Here we explore the information about the price elasticity of supply $\theta^S \in \bar{\Theta}^S = \mathbb{R}_{\geq 0}$ that can be obtained from imposing a bound B^S on the size of shocks to supply. From the data described in Section 3.1 we construct the time series $\{(p_t^S, q_t^S)\}_{t=1}^T$, where p_t^S is 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 is the log of the quantity of grains produced in the world in year t , measured in calories per capita. We also obtain from Roberts and Schlenker (2013b) a measure of the shock Δg_t to agricultural yields in year t .¹⁵

A major source of shocks to the world supply of grain is variation in agricultural yields due to the weather (Roberts and Schlenker 2013a). The maximum absolute value of the yield shock over the sample period is 0.057, and the root mean squared 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 $[0, 0.08]$ for $k = 2$.

Appendix Figure 1 depicts the implications of the contemplated bounds for the price elasticity of supply θ^S . The structure parallels that of Figure 3. The contemplated bounds are again informative. All of the contemplated bounds imply that supply is price-inelastic, $\theta^S < 1$. Roberts and Schlenker (2013a, Table 1, Column 2c) estimate that the price elasticity of supply is $\hat{\theta}_{RS}^S = 0.097$ 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—more than twice the maximum yield shock—implies a price elasticity of at most

¹⁵We use the definition of the yield shock underlying Roberts and Schlenker's (2013a) Table 1, Column 2c.

0.130. The same bound on the price elasticity arises from a bound of $B^S = 0.043$ on the root mean squared shock, or more than 1.7 times the root mean squared yield shock.

B.2 Bounds on a Function of Elasticities

Roberts and Schlenker (2013a) devote attention to the “multiplier” $(|\theta^D| + \theta^S)^{-1}$, 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 $\gamma(\theta^D, \theta^S)$, such as $\gamma(\theta^D, \theta^S) = (|\theta^D| + \theta^S)^{-1}$,¹⁶ by forming the set


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

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.


¹⁶Another prominent example is the function $\gamma(\theta^D, \theta^S) = \theta^S (|\theta^D| + \theta^S)^{-1}$, which determines how the incidence of a tax is shared between consumers and producers (see, e.g., Weyl and Fabinger 2013).

Appendix Figure 1: Implications of Bounds on Shocks to World Supply of Food Grain

Panel A: All Bounds $B^S \in [0, 0.20]$ on the Maximum Shock ($k = \infty$)




Panel B: All Bounds $B^S \in [0, 0.08]$ on the Root Mean Squared Shock ($k = 2$)




Notes: The plots illustrate implications of bounds on the size of shocks to the supply of grain in the application of Roberts and Schlenker (2013a) described in Appendix B.1. Panel A depicts the interval $\hat{\Theta}_\infty(B^S) \cap \bar{\Theta}^S$ implied by bounds $B^S \in [0, 0.20]$ on the maximum shock, where $\bar{\Theta}^S = \mathbb{R}_{\geq 0}$. The dashed vertical line is at three times the maximum absolute yield shock $M_\infty(|\Delta g|)$. Panel B depicts the interval $\hat{\Theta}_2(B^S) \cap \bar{\Theta}^S$ implied by bounds $B^D \in [0, 0.08]$ on the root mean squared shock. The dashed vertical line is at three times the root mean squared yield shock $M_2(|\Delta g|)$. In each plot, the horizontal line depicts the main estimate $\hat{\theta}_{RS}^S$ of the price elasticity of supply 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 $B^D \in \mathcal{B}(k, \bar{\Theta}^S)$ that are compatible with the data.

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

Panel A: Bounds on the Maximum Shock ($k = \infty$)

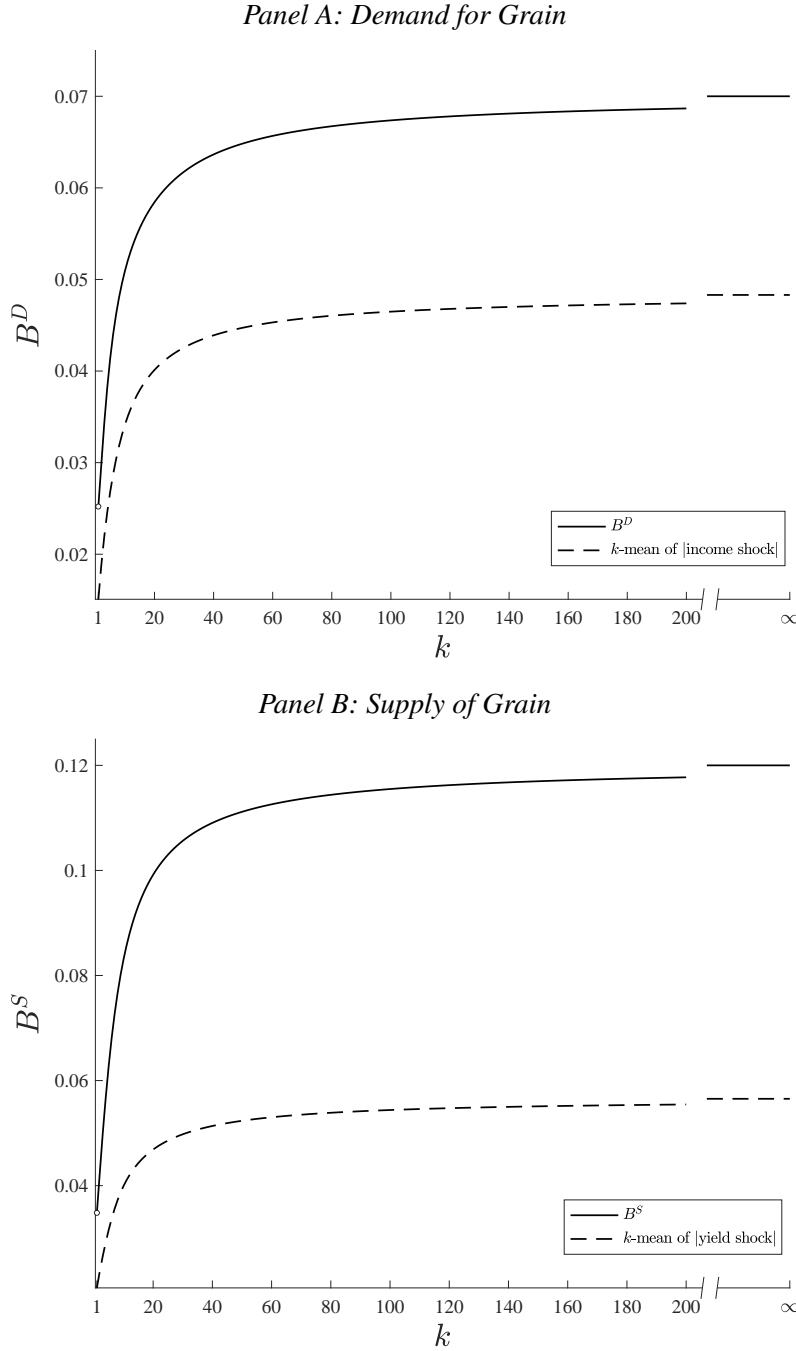


Panel B: Bounds on the Root Mean Squared Shock ($k = 2$)



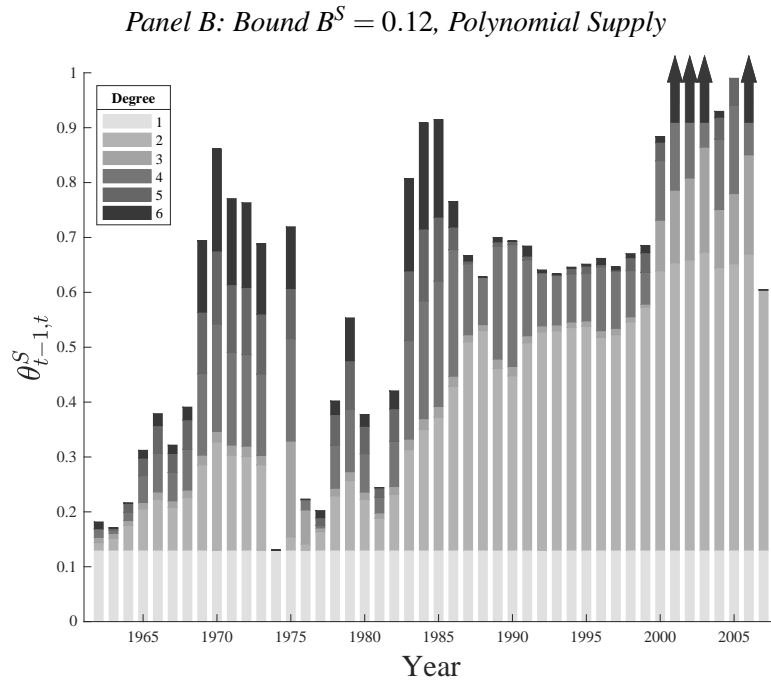
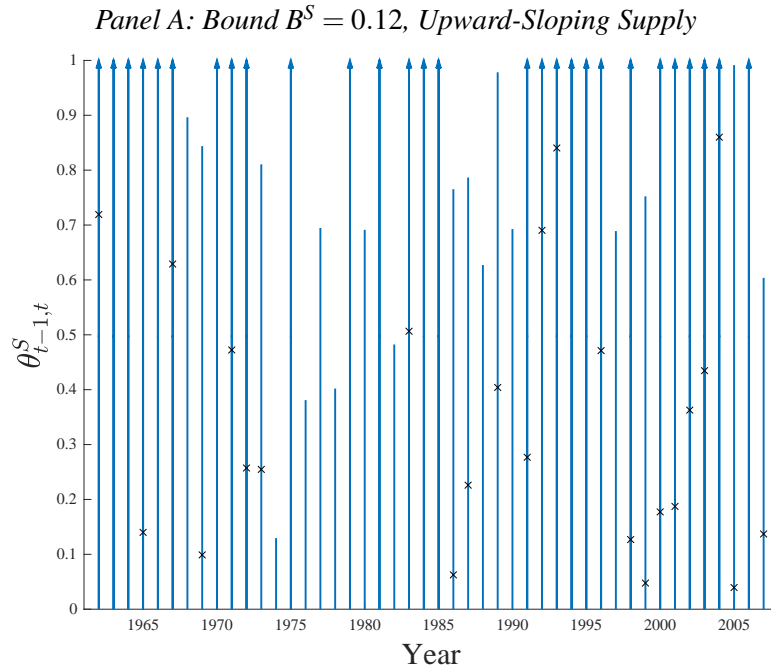
Notes: The plots illustrate implications of bounds on the size of shocks to the supply and demand of grain in the application of Roberts and Schlenker (2013a). Panel A considers bounds $B^D \in [0.035, 0.10]$, $B^S \in [0.095, 0.20]$ on the maximum value of the shock ($k = \infty$). Panel B considers bounds $B^D \in [0.015, 0.04]$, $B^S \in [0.040, 0.06]$ on the root mean squared shock ($k = 2$). In each plot, the black surface depicts the lowest value of the multiplier $\gamma(\theta^D, \theta^S) = (|\theta^D| + \theta^S)^{-1}$ that is compatible with elasticities $\theta^D \in \hat{\Theta}_k(B^D) \cap \bar{\Theta}^D$, $\theta^S \in \hat{\Theta}_k(B^S) \cap \bar{\Theta}^S$, i.e. the smallest element of the set $\hat{\Gamma}_k(B^D, B^S)$. The gray horizontal plane depicts the main estimate $\hat{\gamma}^{RS}$ of the multiplier in Roberts and Schlenker (2013a).

Appendix Figure 3: Bounds on Shocks to Demand and Supply of Grain, Varying k



Notes: 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 $M_k(|\Delta g|)$ of the absolute value of the yield shock. In both panels, values are plotted for $k \in (1, 200]$ and $k = \infty$.

Appendix Figure 4: Relaxing Linearity of the Supply Function

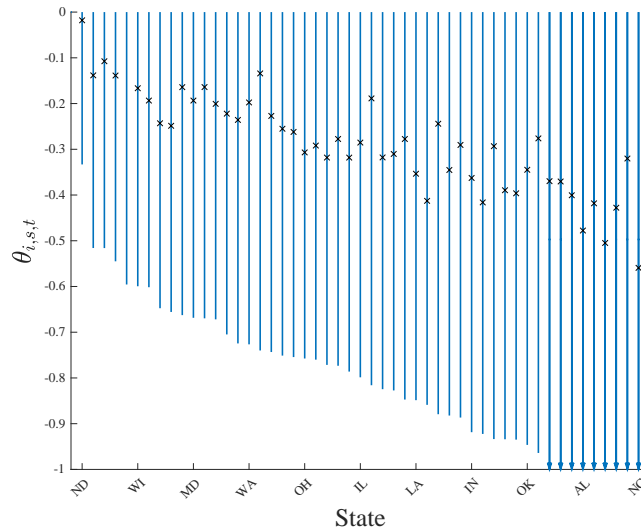


Notes: Each plot depicts bounds on the average price elasticity of supply $\theta_{t-1,t}^S$ between each pair of adjacent years based on the assumption that the shock to world supply of staple food grains is no greater than $B^S = 0.12$. In Panel A, the depicted bounds are formed by intersecting the set in equation (4) with the sign restriction that the average price elasticity is nonnegative. Each line segment represents the interval of possible average price elasticities, with an arrow indicating that the interval contains price elasticities greater than one, and a crosshatch indicating the value of $\Delta q_t / \Delta p_t$ when contained in the plotted range. In Panel B, we further impose that the function $q(\cdot)$ is a polynomial of known degree whose derivative is nonnegative everywhere on the closed interval from the lowest to the highest observed price. Each line segment represents the interval of possible average price elasticities under the given polynomial degree (from one to six), with an arrow indicating that the interval contains price elasticities greater than one.

C Extensions of Analysis of Crowding Out of Male Employment by Female Employment

Appendix Figure 5: Relaxing Linearity of the Crowding Out Function

Bound $B = 0.14$, Downward-Sloping Crowding Out Function



Notes: Each plot depicts bounds on the average crowding out parameter $\theta_{i,s,t}$ for each state i between years $s = 1970$ and $t = 2016$ based on the assumption that the shock to male employment was no greater than $B = 0.14$. The depicted bounds are formed by intersecting the set in equation (4) with the sign restriction that the average crowding out is nonpositive. Each line segment represents the interval of possible average crowding out parameters, with an arrow indicating that the interval contains a value less than -1 , and a crosshatch indicating the value of $\Delta m_i / \Delta f_i$ when contained in the plotted range. States are placed in ascending order according to the lowest value of the crowding out parameter included in the interval.

D Connections to Other Approaches

D.1 Orthogonality Restrictions

Let z_t be some observed variable transformed so that $M_1(\Delta z) = 0$ and $M_2(\Delta z) = 1$.¹⁷ Consider a restriction of the form

$$|M_1(\Delta \varepsilon(\theta) \circ \Delta z)| \leq C \quad (6)$$

where $C \geq 0$ is a scalar. An orthogonality restriction is such a restriction that takes $C = 0$.¹⁸

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, observe that, by the same argument as in the proof of Corollary 1, $\check{\theta}_2 = \arg \min_{\theta} \hat{M}_2(\theta)$ solves

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

For Δp_t standardized, equation (7) is equivalent to an orthogonality restriction with $\Delta z_t = \Delta p_t$.

D.2 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 $\Delta p_t^D = \Delta p_t^S = \Delta p_t$.¹⁹ Then

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

Intuitively, because any pair $(\theta^D, \theta^S) \in \hat{\Theta}_k(B^D) \times \hat{\Theta}_k(B^S)$ is consistent with the data, and by assumption the data are consistent with equilibrium, any pair $(\theta^D, \theta^S) \in \hat{\Theta}_k(B^D) \times \hat{\Theta}_k(B^S)$ must

¹⁷Beginning with a variable \tilde{z}_t we can take $z_t = M_2(\Delta \tilde{z} - M_1(\Delta \tilde{z}) J_{T-1,1})^{-1}(\tilde{z}_t - (t-1)M_1(\Delta \tilde{z}))$, for $J_{T-1,1}$ the $(T-1)$ -dimensional vector of ones.

¹⁸When $C = 0$, 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.

¹⁹In the world market for staple food grains, the quantity demanded and quantity supplied need not be equal at a point in time (and likewise for the demand price and the supply price) because grain can be stored and planting decisions are made in advance of consumption (Roberts and Schlenker 2013a).

also be consistent with equilibrium. In this sense, given a bound B^D on the size of the shocks $\Delta\varepsilon^D(\theta^D)$, there is no further information about θ^D to be obtained by placing a bound B^S on the size of the shocks $\Delta\varepsilon^S(\theta^S)$, and vice versa.

The situation is different if we are prepared to restrict the relationship between the shocks $\Delta\varepsilon_t^D(\theta^D)$ and the shocks $\Delta\varepsilon_t^S(\theta^S)$. For illustration, suppose that $M_1(\Delta q) = M_1(\Delta p) = 0$ and take the restriction that

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

If $R = 0$ then

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

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