

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 March 2022

Kevin Murphy's lectures in price theory inspired this work. We acknowledge funding from the Jan Wallander and Tom Hedelius Foundation (Grant P2015-0095:1), the Population Studies and Training Center, the Eastman Professorship, and the JP Morgan Chase Research Assistant Program at Brown University, the Semester Undergraduate Program for Economics Research at Harvard University, and the National Science Foundation (Grant 1949047). 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, Amy Finkelstein, Peter Ganong, Ed Glaeser, Phil Haile, Toru Kitagawa, Pat Kline, Alex MacKay, Adam McCloskey, Emi Nakamura, Serena Ng, Matt Notowidigdo, Emily Oster, Sharon Oster, Mar Reguant, Michael Roberts, Jon Roth, Chris Snyder, Juan Carlos Suárez Serrato, Chad Syverson, Elie Tamer, seminar participants at Brown University, Iowa State University, the International Methods Colloquium, the NBER, the Interactive Online IO Seminar, the LACEA LAMES annual meeting, and especially NBER discussant Brigham Frandsen 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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 March 2022
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 about the parameter of interest. We illustrate with a main application to the grain 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
Department of Economics
Harvard University
Littauer Center
Cambridge, MA 02138
and NBER
jesse.m.shapiro@gmail.com

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, *Harvard University and NBER**

March 2022

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 about the parameter of interest. We illustrate with a main application to the grain market.

keywords: restrictions on shocks, demand, supply, identification

JEL codes: C20, J21, Q11

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

*Kevin Murphy's lectures in price theory inspired this work. We acknowledge funding from the Jan Wallander and Tom Hedelius Foundation (Grant P2015-0095:1), the Population Studies and Training Center, the Eastman Professorship, and the JP Morgan Chase Research Assistant Program at Brown University, the Semester Undergraduate Program for Economics Research at Harvard University, and the National Science Foundation (Grant 1949047). 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, Amy Finkelstein, Peter Ganong, Ed Glaeser, Phil Haile, Toru Kitagawa, Pat Kline, Alex MacKay, Adam McCloskey, Emi Nakamura, Serena Ng, Matt Notowidigdo, Emily Oster, Sharon Oster, Mar Reguant, Michael Roberts, Jon Roth, Chris Snyder, Juan Carlos Suárez Serrato, Chad Syverson, Elie Tamer, seminar participants at Brown University, Iowa State University, the International Methods Colloquium, the NBER, the Interactive Online IO Seminar, the LACEA LAMES annual meeting, and especially NBER discussant Brigham Frandsen 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@fas.harvard.edu.

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

$$q_t = \theta p_t + \varepsilon_t, \quad (1)$$

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, Goldin, and Katz’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, Goldin, and Katz 2020). These restrictions are typically motivated by economic intuitions about the determinants of ε_t .

In this paper we show that economic intuitions about the size of fluctuations in ε_t can also be informative 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. Because q_t is stable, fluctuations in θp_t must be offset by fluctuations in ε_t . It follows that a larger price elasticity of demand—a more negative θ —implies larger fluctuations in ε_t than does a smaller 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 price elasticity of demand—a less negative value of θ .

We formalize this logic by supposing we can place an upper bound $B \geq 0$ on a generalized power mean, with power at least 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 such 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 .

An economist interested in informing an audience (of, say, policymakers or other economists)

about θ can exhibit the size of shocks necessary to rationalize different values of θ or, alternatively, the values of θ consistent with each of a range of reasonable bounds B on the size of shocks. Such an exhibit can serve as a standalone method of learning about θ , or as a sensitivity analysis complementing another method.

We illustrate our approach with an application to the price elasticity of demand for staple grains following Roberts and Schlenker (2013a). Roberts and Schlenker (2013a) impose a linear model of the form in (1) and approach estimation and inference using orthogonality restrictions with respect to excluded instruments. Values of the demand elasticity θ much larger than Roberts and Schlenker's (2013a) point estimate imply shocks $\Delta\varepsilon_t$ that we consider implausibly large. Accordingly, a range of reasonable bounds on the size of shocks imply informative bounds on θ that are consistent with Roberts and Schlenker's (2013a) inferences. An online appendix includes a second application to the crowding out of male employment by female employment following Fukui, Nakamura, and Steinsson (2020), and an illustration of a data-driven approach to informing bounds based on Ellison and Ellison (2009a).

Untestable restrictions on unobservable variables seem inherently subjective, and we find it unlikely that all economists will agree on an exact bound B (Andrews, Gentzkow, and Shapiro 2020; Andrews and Shapiro 2021). It is for this reason that we advocate reporting the implications of a range of bounds B for the parameter θ , much as Conley, Hansen, and Rossi (2012) advocate reporting the implications of a range of violations of the exclusion restriction for the coefficient of interest in a linear instrumental variables model.

We also think it is unlikely that, in settings such as those we consider, economists will be unable to form useful intuitions about the plausible size of unobserved shocks to economic variables. Such intuitions may be informed by everyday experience (as in our scarf example), economic theory, or by prior evidence on the determinants of the outcome variable (as we illustrate in our application). If such intuitions exist, failing to apply them means that the economist is leaving potentially useful information on the table.

We extend our approach in a few directions. We show how to obtain bounds on an average slope in the case where the model takes the nonlinear form $q_t = q(p_t) + \varepsilon_t$. We discuss the implications of mismeasurement of economic variables. An online appendix discusses further extensions, including to incorporate covariates x_t , to allow for a nonseparable model $q_t = \tilde{q}(p_t, \varepsilon_t)$, and to obtain bounds on a function $\gamma(\cdot)$ of one or more slope parameters.

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 expect that our approach will be most useful to economists analyzing a time-series or panel of well-measured aggregate or market-level variables. Economists analyzing cross-sectional microdata, such as from a random sample survey of individuals, may find it difficult to motivate restrictions on the size of unmeasured economic variables analogous to those we consider here. Economists analyzing poorly measured variables may be able to adopt our formal approach, but, as we discuss in more detail in the paper, may require a statistical, in addition to economic, justification for restrictions on the size of shocks.

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, computation of exact 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 unobserved economic variables go back at least to Marschak and Andrews (1944; see, e.g., equation 1.37),¹ and are related to (though distinct from) approaches based on bounded support of the outcome variable (e.g., Manski 1990). 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, Hansen, and Rossi 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 2022).² The online 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

¹Wald (1940, section 7) considers related restrictions on the distribution of measurement errors.

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.

(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, Ma, and Ng (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 (2021) 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 the policy evaluation setting with a binary treatment, Manski and Pepper (2018) consider a set of restrictions, including a bound on the variation in a given unit’s counterfactual outcome between pairs of years that coincides with the restrictions we study in the two-period case.

Also in the structural vector autoregression setting, Giacomini, Kitagawa, and Read (2021) 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 a given bound on the size of the realized shocks $\Delta\varepsilon_t$. Therefore, in common with the closely related engineering literature that we reference above, issues of probabilistic inference do not arise in our main setup.

The remainder of the paper is organized as follows. Section 2 presents our setup and results. Section 3 presents our application. Section 4 presents extensions. Section 5 concludes. An appendix includes proofs of results stated in the text. An online appendix discusses additional extensions, applications, and connections.

2 Setup and Characterization of Sets of Interest

2.1 Setup

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

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

³In our leading example of log-linear demand, this corresponds to the assumption that $\Delta\varepsilon_t > 0$ for all t . Mullin and Snyder (2021) consider a variety of forms for demand in the reference period, including linear demand, demand known up to a scalar parameter, and concave demand.

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 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 $B \geq 0$ on the value of $\hat{M}_k(\theta)$. We focus on the case where the bound B holds with certainty, but note that our characterizations extend naturally to the case where the bound holds probabilistically.⁴

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 $B \geq 0$ 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

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

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 application. If it fails, any bound that is compatible with the data is uninformative.⁵

2.2 Bounds on the Maximum Absolute Value of the Shock

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

⁴By (2), $M_k(|\Delta\varepsilon|) \leq B$ implies $\theta \in \hat{\Theta}_k(B)$ and vice versa. Therefore $\Pr(\theta \in \hat{\Theta}_k(B)) = \Pr(M_k(|\Delta\varepsilon|) \leq B)$.

⁵Specifically, 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 $\mathcal{B}(k, \mathbb{R}) = [M_k(|\Delta q|), \infty)$.

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 the Appendix. 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.⁶

Remark. 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.

2.3 Bounds on Other Generalized Means of the Absolute Value of the Shock

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,*

⁶More precisely, if $\hat{M}_\infty(\theta) = B$ at the true θ , and in particular there are s, t such that $\Delta p_s, \Delta p_t \neq 0$, $\Delta \varepsilon_s = B \text{sgn}(-\Delta p_s)$, and $\Delta \varepsilon_t = B \text{sgn}(\Delta p_t)$, then $|\hat{\Theta}_\infty(B)| = 1$.

for any $B \in \mathcal{B}(k, \mathbb{R})$ we have that

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

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 $k \in (1, \infty)$. Online Appendix D.1 extends the analysis to the case of $k = 1$ and shows that $\hat{\Theta}_1(B)$ likewise takes the form of an interval.

Figure 1 illustrates the logic of Proposition 2 in a hypothetical example. The function $\hat{M}_k(\theta)$ reaches a minimum at \underline{B}_k , implying that any bound $B' < \underline{B}_k$ is incompatible with the data. A horizontal line at $B > \underline{B}_k$ intersects the function $\hat{M}_k(\theta)$ twice, defining the endpoints $\underline{\theta}_k(B), \bar{\theta}_k(B)$ of the interval $\hat{\Theta}_k(B)$.

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 $\check{\theta}_2$ described in Proposition 2 are available in closed form. Towards a characterization, 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. We then have the following.

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 $\{(\Delta p_t, \Delta q_t)\}_{t=2}^T$. Reporting $\hat{\Theta}_2(B)$ for $B \geq \underline{B}_2$ can therefore be seen as a form of sensitivity analysis with respect to an ordinary least squares estimate, relaxing the orthogonality of Δp_t and $\Delta \varepsilon_t$. Online Appendices E.1 and E.2 further discuss the connection between orthogonality restrictions and those we consider here.

3 Application to the Price Elasticity of World Demand for Staple Food Grains

3.1 Setting

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 log of the quantity of grains consumed in the world in year t , measured in calories per capita.⁷ We also construct a measure y_t of the log of the annual world GDP per capita in 2010 US dollars.⁸

⁷We use the definitions of price and total calories from Roberts and Schlenker (2013a, Table 1, Column 2c), and divide total calories by world population (World Bank 2019a) to obtain calories per capita.

⁸We deflate to 2010 US dollars using the consumer price index from Roberts and Schlenker (2013b).

Roberts and Schlenker (2013a, equation 3) assume that the demand curve takes a log-linear form consistent with equation (1). Their analysis treats the log-linear demand model as structural, using it, for example, to calculate the effect on equilibrium prices and consumer surplus of the US ethanol mandate (pp. 2278-9).

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. Their paper discusses the possibility that yields are endogenous to prices, for example because growers adjust crop densities in response to prices.⁹ Their paper includes extensive sensitivity analysis related to their choice of instrument (see, e.g., pp. 2274-5). Roberts and Schlenker’s (2013a) inclusion of extensive discussion and sensitivity analysis related to the identifying assumption suggests that, while reasonable, not all economists immediately accept it as true, and therefore that there may be room for alternative approaches to learning about the parameter θ^D .

3.2 Forming Intuitions about the Plausible Size of Shocks

Prior research can inform economic 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 (Engel 1857; Houthakker 1957) holds that the income elasticity of demand for food is less than one. Forecasts summarized in Valin et al. (2014, Table 3) imply an income elasticity of world food crop demand ranging from 0.09 to 0.37.¹⁰ Taking the upper end of the range, the income-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.

This discussion 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 data being analyzed. For example, estimates of the income elasticity of food demand can be informed by

⁹They write that “A potential shortcoming...is that yields themselves may be endogenous to price” (p. 2267) and that “...yields might themselves be endogenous, which would make yield deviations an invalid instrument” (Roberts and Schlenker 2013a, p. 2272).

¹⁰The 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 $\ln(1.38)/\ln(16000/6700) \approx 0.370$.

comparisons across countries at a point in time.¹¹ Another is that the choice of reasonable bounds can be contextual. For example, in earlier historical periods the income elasticity of food demand was likely larger (see, e.g., Logan 2006), so an economist studying data from such a period might wish to consider larger bounds B^D than an economist studying data from the modern period. A final aspect is that intuitions about the plausible size of shocks are subjective. Although we find it implausible that shocks to demand were much larger than those that can be explained by shocks to income alone, we do not think it is possible to defend a single numerical bound as the most reasonable one. We therefore explore the implications of a range of bounds.

3.3 Implications of Bounds on the Size of Shocks

Figure 2 illustrates why intuitions about the size of shocks to demand are informative about 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. An economist interested in informing an audience about the price elasticity of demand θ^D could present a plot similar to Figure 2, allowing the audience to evaluate the plausibility of the shocks implied by different values of θ^D .

Following the logic of Section 2, we can also directly characterize the implications for the price elasticity θ^D of a given bound B^D on the size of the shocks. Figure 3 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, $M_\infty(|0.37\Delta y|) \approx 0.05$. The figure depicts a scatterplot of the first-differenced data $\{(\Delta p_t^D, \Delta q_t^D)\}_{t=2}^T$. In first differences, a demand function is a line through the origin with nonpositive slope $\theta^D \in \overline{\Theta}^D = \mathbb{R}_{\leq 0}$. The figure also

¹¹Muhammad 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.

depicts a dotted interval with radius $B^D = 0.07$ around each point. A demand function consistent with a bound of $B^D = 0.07$ on the maximum absolute value of the demand shock is one that passes through all of the dotted intervals. The figure depicts a shaded region collecting all such demand functions, i.e., those with slope $\theta^D \in \hat{\Theta}_\infty(0.07) \cap \bar{\Theta}^D$.

A bound of $B^D = 0.07$ on the maximum absolute value of the demand shock is informative about 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.

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 4 does this. 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). An economist interested in informing an audience about the price elasticity of demand θ^D could present a plot similar to Figure 4, allowing the audience to evaluate the implications of different plausible bounds B^D on the size of the shocks, and to compare these implications to those of other approaches to learning about θ^D .

Figure 4 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 demand shock is at least 0.039 (Panel A) and the root mean squared demand 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 $\theta^D \leq 0$.

Online Appendix A includes several extensions of our analysis of the grain market. Online Appendix A.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 Online Appendix Figure 1. Online Appendix A.2 characterizes bounds on a function of the elasticities θ^D and θ^S , illustrated in Online Appendix Figure 2 with an application to the “multiplier” parameter studied in Roberts and Schlenker (2013a). Online Appendix A.3 discusses the possibility of orthogonalizing with respect to an observed covariate, illustrated in Online Appendix Figure 3 with an application to time trends considered in Roberts and Schlenker (2013a). Lastly, Online Appendix Figure 4 illustrates the role of k by showing how

$M_k(|0.37\Delta y|)$, and the value of B^D needed to obtain a given bound on the price elasticity, vary with k .

4 Extensions and Discussion

4.1 Nonlinear Model

In the setting of Section 3 and many others, the authors assume a linear relationship between the observed variables of interest, as in equation (1). In settings where the economic model instead implies a nonlinear relationship via a known strictly monotone link function, $q_t = f(\theta p_t + \varepsilon_t)$, we may proceed by inverting the link function, replacing q_t with $f^{-1}(q_t)$ in (1), as in Berry (1994).

In some settings we may instead be interested in nonlinear relationships of the form

$$q_t = q(p_t) + \varepsilon_t, \quad (3)$$

where $q(\cdot)$ is an unknown function.

In such settings, 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 4. 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 $\theta_{t-1,t}^D > -1$.

Even more informative statements are possible if we are prepared to assume that $q(\cdot)$ is a polynomial of known degree. Panel B of Figure 5 shows that even allowing for a polynomial of degree 6, a substantial generalization of linearity, in 89 percent of years we can conclude that $\theta_{t-1,t}^D > -0.3$.

Online Appendix D.2 shows that in the case of a nonseparable model, an analogue of the characterization in (4) can be obtained via a suitable reinterpretation of the economic quantities. Online Appendix D.3 further shows how to obtain a bound on the mean of the average slopes between adjacent periods by coupling a bound on the size of the shock with a bound on the variation in the slope of the function $q(\cdot)$.

4.2 Mismeasured Variables

Suppose that the economist observes $\hat{q}_t = q_t + \mu_t$ for μ_t an unobserved measurement error.¹² The economist can proceed as in Section 2, now treating B as a bound on the size $M_k(|\Delta\varepsilon_t + \Delta\mu_t|)$ of the absolute value $|\Delta\varepsilon_t + \Delta\mu_t|$ of the shock to the composite unobservable comprised of both the unobserved economic factor ε_t and the unobserved measurement error μ_t . The presence of measurement error may necessitate using different values of B and k than would be appropriate in its absence. For example, it may be that the economist is prepared to impose a bound B on $M_\infty(|\Delta\varepsilon_t|)$ but not on $M_\infty(|\Delta\varepsilon_t + \Delta\mu_t|)$, say because occasional extreme economic shocks are not plausible but occasional severe mismeasurement is plausible. In such a setting, the economist may prefer to impose a bound on, say, $M_2(|\Delta\varepsilon_t + \Delta\mu_t|)$, guided by intuitions about the plausible size of measurement error in a typical period.

Suppose that measurement error is present but the economist fails to account for it. If some bound B applies to the size $M_k(|\Delta\varepsilon_t|)$ of the economic shock, but only a looser bound $B' \geq B$ is appropriate for the size $M_k(|\Delta\varepsilon_t + \Delta\mu_t|)$ of the composite shock, then the economist using the bound B that is too tight will obtain an interval for θ that is too tight, because $\hat{\Theta}_k(B) \subseteq \hat{\Theta}_k(B')$.

Suppose next that it is p_t , rather than q_t , that is potentially mismeasured. In general, this situation does not fit in the framework of Section 2, and therefore requires a different approach or

¹²Starting from any true value q_t and measured value \hat{q}_t , we can always define $\mu_t \equiv \hat{q}_t - q_t$.

characterization. A partial exception is the case where the measurement error in p_t takes a known statistical form, such as when it comes from sampling variation. Online Appendix Section B.3 discusses such a situation in the context of a specific application.

5 Conclusions

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 a linear economic model that nests many models used in empirical research. We illustrate the utility of the approach with an application, where we argue that the approach can usefully complement existing approaches to learning about the parameter of interest. We extend the approach to the case of nonlinear models and show that it remains informative.

References


- Alexandratos, Nikos and Jelle Bruinsma. 2012. World agriculture towards 2030/2050: The 2012 revision. *FAO ESA Working Paper No. 12-03*. <http://www.fao.org/3/a-ap106e.pdf>. Accessed in May 2020.
- Andrews, Isaiah, Matthew Gentzkow, and Jesse Shapiro. 2020. Transparency in structural research. *Journal of Business & Economic Statistics* 38(4): 711-22.
- Andrews, Isaiah and Jesse Shapiro. 2021. A model of scientific communication. *Econometrica* 89: 2117-42
- 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: 347-51.
- Barro, Robert J. and Charles J. Redlick. 2011. Macroeconomic effects from government purchases and taxes. *Quarterly Journal of Economics* 126(1): 51-102.
- 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.
- Berry, Steven T. 1994. Estimating discrete-choice models of product differentiation. *RAND Journal of Economics* 25(2): 242-62.

- Conley, Timothy G., Christian B. Hansen, and Peter E. Rossi. 2012. Plausibly exogenous. *Review of Economics and Statistics* 94(1): 260-72.
- Ellison, Glenn and Sara Fisher Ellison. 2009a. Search, obfuscation, and price elasticities on the internet. *Econometrica* 77(2): 427-52
- Engel, Ernst. 1857. Die Productions- und Consumtionsverhältnisse des Königreichs Sachsen. *Zeitschrift des statistischen Bureaus des Königlich Sächsischen Ministerium des Inneren* 8-9: 1-54.
- Feenstra, Robert C. 1994. New product varieties and the measurement of international prices. *American Economic Review* 84(1): 157-77.
- Feenstra, Robert C. and David E. Weinstein. 2017. Globalization, markups, and US welfare. *Journal of Political Economy* 125(4): 1040-74.
- Fiorito, Riccardo and Giulio Zanella. 2012. The anatomy of the aggregate labor supply elasticity. *Review of Economic Dynamics* 15(2): 171-87.
- Fukui, Masao, Emi Nakamura, and Jón Steinsson. 2020. Women, wealth effects, and slow recoveries. *National Bureau of Economic Research Working Paper No. 25311*.
- Giacomini, Raffaella, Toru Kitagawa, and Matthew Read. 2021. Identification and inference under narrative restrictions. <https://arxiv.org/abs/2102.06456>. Accessed on April 21 2021.
- Houthakker, Hendrik S. 1957. An international comparison of household expenditure patterns, commemorating the centenary of Engel's law. *Econometrica* 25(4): 532-51.
- 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): 5915-20.
- 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): 319-27.
- 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): 313-41.
- 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. 2022. 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): 319-23.
- Manski, Charles F. 1997. Monotone treatment response. *Econometrica* 65(6): 1311-34.
- Manski, Charles F. and John V. Pepper. 2018. How do right-to-carry laws affect crime rates? Coping with ambiguity using bounded-variation assumptions. *Review of Economics and Statistics* 100(2): 232-44.
- 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: 5307-68. Amsterdam: 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): 408-14.
- 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. *USDA Economic Research Service Technical Bulletin* No. 1929. https://www.ers.usda.gov/webdocs/publications/47579/7637_tb1929.pdf?v=45. Accessed in May 2020.
- Mullin, Wallace P. and Christopher M. Snyder. 2021. A simple method for bounding the elasticity of growing demand with applications to the analysis of historic antitrust cases. *American Economic Journal: Microeconomics* 13(4): 172-217.
- Nevo, Aviv and Adam M. Rosen. 2012. Identification with imperfect instruments. *Review of Economics and Statistics* 94(3): 659-71.
- 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. *American Economic Association* [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E112674V1>.
- Robbins, Herbert and Cun-Hui Zhang. 1986. Maximum likelihood estimation in regression with uniform errors. In John Van Ryzin (ed.), *Institute of Mathematical Statistics Lecture Notes* 8: 365-85. Hayward, CA: Institute of Mathematical Statistics.
- Tamer, Elie. 2010. Partial identification in econometrics. *Annual Review of Economics* 2: 167-95.
- 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 Mensbrugge, 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.
- Wald, Abraham. 1940. The fitting of straight lines if both variables are subject to error. *Annals of Mathematical Statistics* 11(3): 284-300.
- 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): 449-68.
- Walter, Eric and H el ene Piet-Lahanier. 1996. Recursive robust minimax estimation. In Mario Milanese, John Norton, H el ene Piet-Lahanier, and Eric Walter (eds.), *Bounding Approaches to System Identification* 183-97. Boston: Springer.
- World Bank. 2019a. Population, total. *World Development Indicators*. *The World Bank Group*. <https://data.worldbank.org/indicator/SP.POP.TOTL>. Accessed on October 2, 2019.
- World Bank. 2019b. GDP (current US\$). *World Development Indicators*. *The World Bank Group*. <https://data.worldbank.org/indicator/NY.GDP.MKTP.CD>. Accessed on October 24, 2019.

Figure 1: Illustration of Proposition 2




Notes: The plot illustrates the logic of Proposition 2 for a hypothetical example. For some $k \in (1, \infty)$, the plot shows the k -mean of the absolute value of the shocks, $\hat{M}_k(\theta)$, as a function of the unknown slope, θ . Given an upper bound B on $\hat{M}_k(\theta)$, we can infer that the slope θ must lie in the shaded interval $\hat{\Theta}_k(B) = [\underline{\theta}_k(B), \bar{\theta}_k(B)]$. Moreover, any bound B that is below $\underline{B}_k = \hat{M}_k(\check{\theta}_k)$ is incompatible with the data because it lies below $\hat{M}_k(\theta)$ for all θ .

Figure 2: 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. Each series corresponds to the shocks $\Delta \epsilon_t(\theta^D)$ to demand implied by a given value of the price elasticity of demand θ^D . We depict the shocks implied by the point estimate of Roberts and Schlenker (2013a, Table 1, Column 2c), denoted $\hat{\theta}_{RS}^D$, and the shocks implied by unit-elastic demand, $\theta^D = -1$.


Figure 3: 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. The cross-hatches depict a scatterplot of the data $\{(\Delta p_t^D, \Delta q_t^D)\}_{t=2}^T$. The dotted interval around each cross-hatch has radius $B^D = 0.07$. The shaded region depicts all demand functions consistent with an upper bound of $B^D = 0.07$ on the maximum absolute value of the demand shock. These are the downward-sloping lines that pass through the origin and through all of the dotted intervals, i.e., the lines through the origin with slope $\theta^D \in \hat{\Theta}_\infty(0.07) \cap \bar{\Theta}^D$ for $\bar{\Theta}^D = \mathbb{R}_{\leq 0}$.

Figure 4: 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 ($k = \infty$)



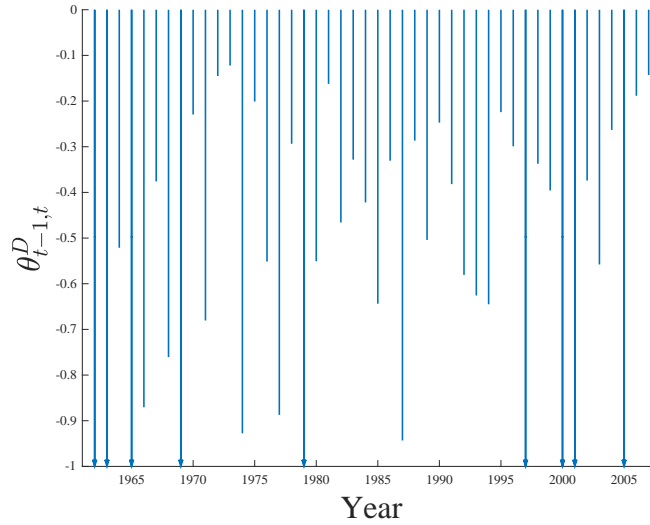
Panel B: All Bounds $B^D \in [0, 0.04]$ on the Root Mean Squared Shock ($k = 2$)




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. 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 point estimate $\hat{\theta}_{RS}^D$ of the price elasticity of demand in Roberts and Schlenker (2013a, Table 1, Column 2c), and the shaded region depicts the associated 95% confidence interval. 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 5: Relaxing Linearity of the Demand Function

Panel A: Bound $B^D = 0.07$, Downward-Sloping Demand



Panel B: Bound $B^D = 0.07$, Polynomial Demand



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 absolute 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 . 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).

Appendix: 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 $t \geq 2$. For a given $t \geq 2$, 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 $t \geq 2$ 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 $\underline{\theta}_\infty(B)$ and $\bar{\theta}_\infty(B)$ be as defined in the statement of the proposition. 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

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

Therefore there is a unique solution $\tilde{B} \geq 0$ to $\underline{\theta}_\infty(\tilde{B}) = \bar{\theta}_\infty(\tilde{B})$. 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.

Online Appendix for

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, *Harvard University and NBER*

March 2022

A Extensions of Analysis of World Market for Staple Food Grains

A.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 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 .¹

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 $k = 2$.

Online Appendix Figure 1 depicts the implications of the contemplated bounds for the price elasticity of supply θ^S . The structure parallels that of Figure 4. 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

*E-mail: marco_stenborg_petterson@brown.edu, david.seim@ne.su.se, jesse_shapiro@fas.harvard.edu

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

elasticity of at most 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.

A.2 Bounds on a Function of Two 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 $\gamma(\theta^D, \theta^S) = (|\theta^D| + \theta^S)^{-1}$,² by forming the set

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

Online 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.

A.3 Orthogonalization with Respect to Covariates

Let $\{x_t\}_{t=1}^T$ be an observed sequence of values of a (possibly vector-valued) covariate. For any θ , let $\Delta\varepsilon_{\perp}(\theta)$ be the component of $\Delta\varepsilon(\theta)$ orthogonal to $\Delta x = (\Delta x'_2, \dots, \Delta x'_T)$.³ If we are prepared to impose an upper bound of $B_{\perp} \geq 0$ on the k -mean of $|\Delta\varepsilon_{\perp}(\theta)|$, then we may form the set $\{\theta \in \mathbb{R} : M_k(|\Delta\varepsilon_{\perp}(\theta)|) \leq B_{\perp}\}$ of parameters θ that are consistent with this bound.

We may loosely think of B_{\perp} as a bound on the portion of the shocks that cannot be “explained” (statistically) by the covariates. The economic interpretation of a bound $B_{\perp} \geq 0$ on the size of the orthogonalized shocks $\Delta\varepsilon_{\perp}(\theta)$ is different from that of a bound $B \geq 0$ on the size of the overall shocks $\Delta\varepsilon(\theta)$. Which type of bound will be of interest in a given application will therefore depend on whether it is easier to form economic intuitions about the size of $\Delta\varepsilon_{\perp}(\theta)$ or about the size of $\Delta\varepsilon(\theta)$.


²Another 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).

³That is, $\Delta\varepsilon_{\perp}(\theta) = \Delta\varepsilon(\theta) - \Delta x (\Delta x' \Delta x)^{-1} \Delta x' \Delta\varepsilon(\theta)$.


In their model of world food demand, Roberts and Schlenker (2013a, Table 1, Column 2c) include as a control a restricted cubic spline. Panel A of Online Appendix Figure 3 depicts the implications of imposing a bound B_{\perp} on the maximum absolute value of the component of the demand shock that is orthogonal to the components of this spline. Panel B of Online Appendix Figure 3 depicts the implications of imposing a bound B_{\perp} on the maximum absolute value of the component of the supply shock that is orthogonal to the control variables included in Roberts and Schlenker's (2013a, Table 1, Column 2c) model of supply.

Online 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 Online Appendix A.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 estimate $\hat{\theta}_{RS}^S$ of the price elasticity of supply in Roberts and Schlenker (2013a, Table 1, Column 2c), and the shaded region depicts the associated 95% confidence interval. 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.

Online 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) described in Online Appendix A.2. Panel A considers bounds $B^D \in [0.035, 0.10]$, $B^S \in [0.085, 0.20]$ on the maximum value of the shock ($k = \infty$). Panel B considers bounds $B^D \in [0.015, 0.04]$, $B^S \in [0.040, 0.08]$ 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 point estimate $\hat{\gamma}^{RS}$ of the multiplier in Roberts and Schlenker (2013a, Table 1, Column 2c).

Online Appendix Figure 3: Implications of Bounds on Orthogonalized Shocks to World Demand and Supply for Food Grain

Panel A: All Bounds $B_{\perp}^D \in [0, 0.10]$ on the Maximum Orthogonalized Demand Shock ($k = \infty$)

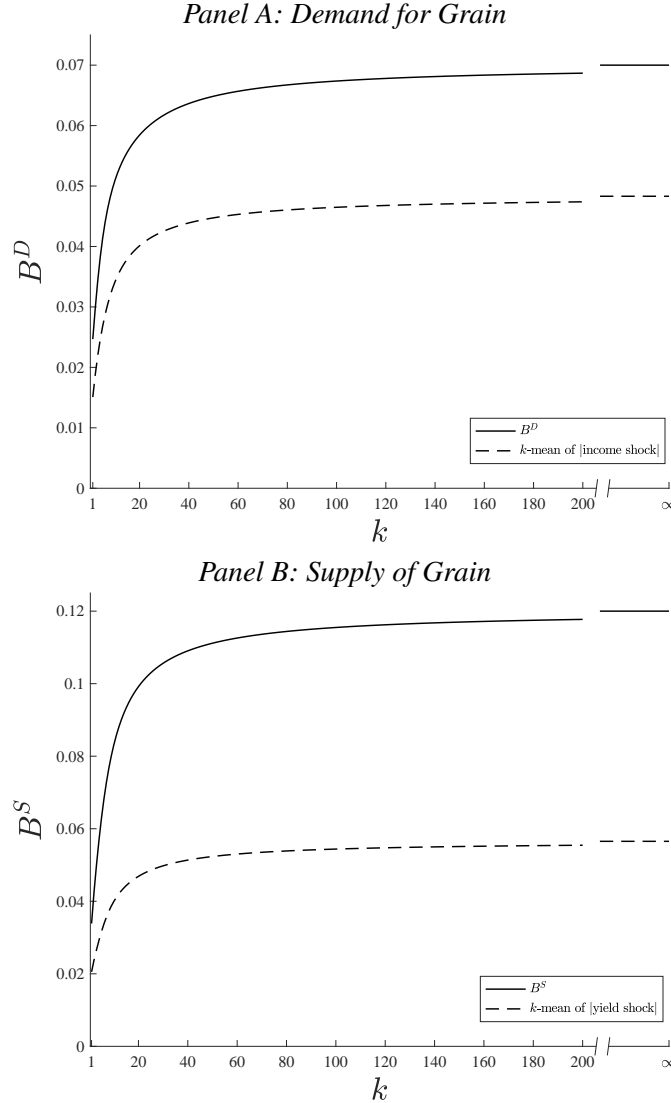


Panel B: All Bounds $B_{\perp}^S \in [0, 0.20]$ on the Maximum Orthogonalized Supply Shock ($k = \infty$)



Notes: The plot illustrates implications of bounds on the size of orthogonalized shocks to the demand and supply for grain in the setting of Roberts and Schlenker (2013a), following the approach described in Online Appendix A.3. Panel A depicts the interval $\{\theta \in \bar{\Theta}^D : M_{\infty}(|\Delta \varepsilon_{\perp}^D(\theta)|) \leq B_{\perp}^D\}$ implied by bounds $B_{\perp}^D \in [0, 0.10]$ on the maximum absolute orthogonalized shock to demand, where $\bar{\Theta}^D = \mathbb{R}_{\leq 0}$. Panel B depicts the interval $\{\theta \in \bar{\Theta}^S : M_{\infty}(|\Delta \varepsilon_{\perp}^S(\theta)|) \leq B_{\perp}^S\}$ implied by bounds $B_{\perp}^S \in [0, 0.20]$ on the maximum absolute orthogonalized shock to supply, where $\bar{\Theta}^S = \mathbb{R}_{\geq 0}$. In each plot, we orthogonalize with respect to the first difference of the covariates x_t specified in Roberts and Schlenker (2013a, Table 1, Column 2c). In Panel A, x_t consists of the components of a five-knot restricted cubic spline. In Panel B, x_t additionally includes the yield shock g_t . In each plot, the horizontal line depicts the point estimate $\hat{\theta}_{RS}^D$ or $\hat{\theta}_{RS}^S$ of the price elasticity of demand or supply, respectively, in Roberts and Schlenker (2013a, Table 1, Column 2c), and the shaded region depicts the associated 95% confidence interval. The solid portion of the x-axis corresponds to the bounds B_{\perp}^D or B_{\perp}^S that are compatible with the data.

Online Appendix Figure 4: 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 $k = \infty$.

B Extension to Panel Data

B.1 Setup

Our approach extends readily to the case where we observe a finite time series $\{(p_{it}, q_{it})\}_{t=1}^{T_i}$ for each of a cross-section of units $i \in \{1, \dots, N\}$, such as countries or states. Let $\Delta \varepsilon_i(\theta) = (\Delta \varepsilon_{i2}(\theta), \dots, \Delta \varepsilon_{iT_i}(\theta))$, 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 repeat the exercise in Section 2, 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 allows for imposing the same bound for all units ($B_i = B$ for all i), different bounds for different units ($B_i \neq B_j$ for some $i \neq j$), or no bound for some units ($B_i = \infty$ for some i). Note also that, because we treat all variables in first differences, the analysis is unchanged if we envision that $q_{it} = \alpha_i + \theta_i p_{it} + \varepsilon_{it}$ for some unit-specific intercept α_i .

Our approach also extends readily to the case where the economist wishes to impose different bounds on the size of shocks in different time periods. To see this, note that if we partition the set $\{1, \dots, T\}$ of periods into cells $i \in \{1, \dots, N\}$, each containing a contiguous set of periods $\{\underline{t}_i, \dots, \bar{t}_i\}$, then we can proceed as in the case of panel data, with the cells i of the partition now playing the role of the cross-sectional units.

B.2 Application to Crowding Out of Male Employment by Female Employment

Fukui, Nakamura, and Steinsson (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, Nakamura, and Steinsson'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, Nakamura, and Steinsson (2020, equation 5) specify a homogenous linear relationship

between Δm_i and Δf_i of the form $\Delta m_i = \theta^C \Delta f_i + \Delta \varepsilon_i$.⁴ Fukui, Nakamura, and Steinsson (2020) adopt 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 Δf_i in the female employment-to-population ratio was 0.27, and the largest change was 0.44. The major cultural and technological forces that contributed to this trend have been widely studied and documented (see, for example, the review in Greenwood, Guner, and Vandenbroucke 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.⁵ 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 .⁶ Imposing that crowding out is nonpositive means that $\theta^C \in \bar{\Theta} = \mathbb{R}_{\leq 0}$. Online Appendix Figure 5 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-to-population would have changed by more than $B = 0.14$ in the absence of changes in female employment-to-population. 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 of $[-0.35, 0.09]$ from Fukui, Nakamura, and Steinsson'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, Nakamura, and Steinsson's (2020, Table 3, Column 2) preferred point estimate $\hat{\theta}_{FNS}^C = -0.13$ unless B is less than 0.09.⁷

⁴To 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 .

⁵Juhn and Potter (2006, p. 32) write, "The biggest story in labor force participation rates in recent decades involves the labor force attachment of women."

⁶That is, for any feasible bound B , we have that $\hat{\Theta}_{ik}(B) = \hat{\Theta}_i(B)$ for all $k \geq 1$.

⁷Fukui, Nakamura, and Steinsson (forthcoming, Table 3, Column 2) report a revised point estimate of $\hat{\theta}_{FNS}^C = -0.18$ with a confidence interval of $[-0.34, -0.02]$.


It is also instructive to examine the shocks to male employment-to-population 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-to-population 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-to-population absent a change in female employment-to-population. 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, Nakamura, and Steinsson (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 inform conclusions about θ^C .

B.3 Accounting for Sampled Data

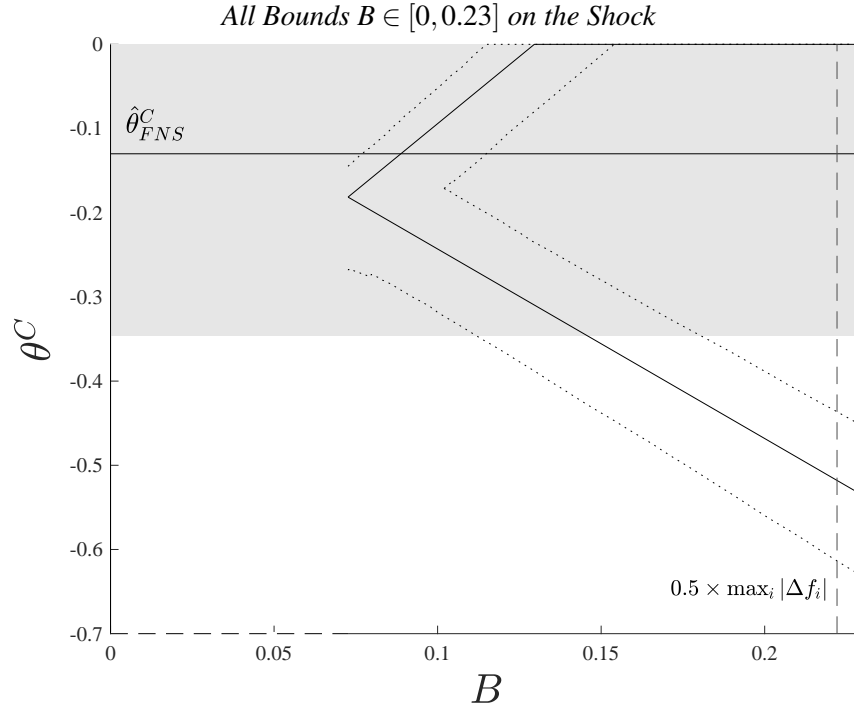
Fukui, Nakamura, and Steinsson (2020, Section 2) measure the variables Δf_i and Δm_i using survey microdata. Because the survey microdata come from a random sample we can approximate the sampling variation in the measured variables. Online Appendix Figure 6 depicts a bootstrap estimate of the variation in the computed bounds on the crowding out parameter θ^C induced by sampling variation in the measures of Δf_i and Δm_i . Because the survey sample is fairly large, in this application we estimate that the influence of sampling variation is modest compared to the information contained in the bounds.

Online Appendix Figure 5: 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, Nakmura, and Steinsson (2020) described in Online Appendix B.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-to-population $\max_i |\Delta f_i|$. The horizontal line depicts the point estimate $\hat{\theta}_{FNS}^C$ of the crowding out of male employment by female employment in Fukui, Nakmura, and Steinsson (2020, Table 3, Column 2), and the shaded region depicts the associated 95% confidence interval. The solid portion of the x-axis corresponds to the bounds $B \in \mathcal{B}(k, \bar{\Theta})$ that are compatible with the data.

Online Appendix Figure 6: Implications of Bounds on Shocks to Male Employment, Accounting for Sampled Data



Notes: The plot illustrates the implications of sampling uncertainty for the bounds on the size of shocks to male employment in the setting of Fukui, Nakmura, and Steinsson (2020) described in Online Appendix B.2. Following Online Appendix Figure 5, the solid lines depict 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 dotted lines around the bounds depict, respectively, the 2.5th and 97.5th percentiles of the upper and lower bounds in the sampling distribution of the variables Δf_i and Δm_i . We obtain these percentiles from a nonparametric bootstrap with 1000 replicates. In each replicate, we draw individuals with replacement from the survey microdata from which Δf_i and Δm_i are calculated, and recompute the variables on the resampled data. The dashed vertical line is at half the maximum absolute change in female employment $\max_i |\Delta f_i|$. The horizontal line depicts the point estimate $\hat{\theta}_{FNS}^C$ of the crowding out of male employment by female employment in Fukui, Nakmura, and Steinsson (2020, Table 3, Column 2), and the shaded region depicts the associated 95% confidence interval. The solid portion of the x-axis corresponds to the bounds $B \in \mathcal{B}(k, \bar{\Theta})$ that are compatible with the data in the full sample. We depict the interval $\cap_{i=1}^N \hat{\Theta}_i(B) \cap \bar{\Theta}$ only for $B \in \cap_{i=1}^N \hat{\Theta}_i(B) \cap \bar{\Theta}$. We compute percentiles only among those bootstrap replicates in which the respective bound is well-defined.

C Data-driven Bounds with Infrequent Changes

C.1 Setup

In cases where p_t changes infrequently it may be possible to inform the bound B using the data. Divide the periods $\{2, \dots, T\}$ into two mutually exclusive and exhaustive groups, with $\mathcal{S} = \{t \in 2, \dots, T : \Delta p_t = 0\}$ collecting periods in which there has been no change in p_t and $\mathcal{T} = \{t \in 2, \dots, T : \Delta p_t \neq 0\}$ collecting the rest. We have already assumed that \mathcal{T} is nonempty; for the purpose of this section we further assume that \mathcal{S} is also nonempty.

Now let

$$\hat{M}_k^{\mathcal{S}} = \left(\frac{1}{|\mathcal{S}|} \sum_{t \in \mathcal{S}} |\Delta q_t|^k \right)^{1/k} = \left(\frac{1}{|\mathcal{S}|} \sum_{t \in \mathcal{S}} |\Delta \varepsilon_t|^k \right)^{1/k},$$

where the second equality uses the property of (1) that if $\Delta p_t = 0$ for some t , then $\Delta q_t = \Delta \varepsilon_t$ regardless of θ . Further let

$$\hat{M}_k^{\mathcal{T}}(\theta) = \left(\frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} |\Delta \varepsilon_t(\theta)|^k \right)^{1/k}.$$

Then it may be reasonable to use the value of $\hat{M}_k^{\mathcal{S}}$ to inform a choice of bound on $\hat{M}_k^{\mathcal{T}}(\theta)$, for example by supposing that $\hat{M}_k^{\mathcal{T}}(\theta) \leq \lambda \hat{M}_k^{\mathcal{S}}$ for some scalar $\lambda \geq 1$. We caution that if p_t depends on ε_t , for example due to optimization or market equilibrium, direct restrictions such as $\lambda = 1$ need not be economically appealing.⁸

C.2 Application to Online Sales of Memory Modules

Ellison and Ellison (2009a) study the elasticity of demand for computer memory modules sold by an internet retailer using daily data for dates in the period from May 2000 through May 2001. We focus on the demand for low-quality memory modules from a single website owned by the retailer. From Ellison and Ellison's (2009b) code and data, we construct a time series $\{(p_t, q_t)\}_{t=1}^T$, where p_t is the log of the average transaction price of low-quality 128MB PC100 memory modules sold by the website on day t , and q_t is the log of the daily quantity sold, undefined for the 7 out of $T = 343$ days on which no modules in this category were sold.

⁸Ottonello and Song (2022) discuss approaches to identification based on changes in the variability of unobserved shocks.


Ellison and Ellison (2009a, p. 440) assume that, when demand is positive, the demand curve takes a log-linear form consistent with equation (1), though with ε_t including a function of the price's rank on a price-search website, which in turn depends on the price p_t , as well as functions of the prices of other types of memory modules sold by the retailer. Ellison and Ellison (2009a) approach identification by assuming that an unobserved multiplicative structural error is mean-independent of a vector of covariates including either prices or excluded instruments. Ellison and Ellison (2009a, p. 441) discuss the interpretation of these identifying assumptions in a market with infrequent price changes. Here we explore using the size of shocks on dates without price changes to inform beliefs about the size of shocks on dates with price changes.

On the $|\mathcal{S}| = 171$ days in which the price of the modules is unchanged from the preceding day we find that the size of shocks is $\hat{M}_\infty^{\mathcal{S}} = 2.89$ and $\hat{M}_2^{\mathcal{S}} = 0.72$. Online Appendix Figure 7 uses these as a point of reference for the construction of bounds on the shocks $\hat{M}_k^{\mathcal{S}}(\theta)$ during periods with price changes. The vertical axis exhibits the bounds on θ and the horizontal axis exhibits the multiple λ that we use in constructing the bounds. In the case of both $k = \infty$ and $k = 2$, the data imply that $\hat{M}_k^{\mathcal{S}}(\theta) > \hat{M}_k^{\mathcal{S}}$, meaning that the bound $\hat{M}_k^{\mathcal{S}}(\theta) \leq \lambda \hat{M}_k^{\mathcal{S}}$ for $\lambda = 1$ is inconsistent with the data. We find, however, that allowing λ in a neighborhood of one yields informative bounds on θ ; these bounds exclude Ellison and Ellison's (2009a, Table III) point estimate up to $\lambda = 1.58$.


Unlike in the application in Section 3, the setting here is one in which bounds on the size of shocks can imply a lower bound on the absolute value $|\theta|$ of the price elasticity. Indeed, for $k = 2$ and $\lambda \leq 1.38$, we find that demand must be elastic.

Online Appendix Figure 7: Implications of Bounds on Shocks to Memory Module Demand

Panel A: All Bounds $\lambda \in [1, 1.5]$ on the Relative Maximum Shock ($k = \infty$)



Panel B: All Bounds $\lambda \in [1, 1.5]$ on the Relative Root Mean Squared Shock ($k = 2$)



Notes: The plot illustrates implications of bounds on the relative size of shocks to memory module demand in the setting of Ellison and Ellison (2009a) described in Online Appendix C.2. The plot depicts the nonpositive values of θ consistent with bounds $\hat{M}_k^{\mathcal{P}}(\theta) \leq \lambda \hat{M}_k^{\mathcal{S}}$ on the shock where $\hat{M}_k^{\mathcal{S}}$ is the k -mean of the shock in periods with no price change, $\hat{M}_k^{\mathcal{P}}(\theta)$ is the k -mean of the shock in periods with a price change, and we consider values $\lambda \in [1, 1.5]$. The horizontal line depicts the point estimate $\hat{\theta}_{EE}$ of the elasticity in Ellison and Ellison (2009a, Table III). The solid portion of the x-axis corresponds to the relative bounds λ that are compatible with the data.

D Extensions of Formal Approach

D.1 Bounds on the Mean Absolute Deviation ($k = 1$)

Proposition 3. For $k = 1$, the set $\mathcal{B}(1, \mathbb{R})$ is equal to $[\underline{B}_1, \infty)$ for $\underline{B}_1 = \min_{\theta} \hat{M}_1(\theta)$. Moreover, for any $B \in \mathcal{B}(\underline{B}_1, \mathbb{R})$ we have that

$$\hat{\Theta}_1(B) = [\underline{\theta}_1(B), \bar{\theta}_1(B)]$$

where $\underline{\theta}_1(B), \bar{\theta}_1(B)$ are finite.

The proof is similar to that of Lemma 1 and Proposition 2, but accounts for the fact that the function $\hat{M}_1(\theta)$ need not have a unique minimum.

Proof of Proposition 3

We begin by establishing several elementary properties of the function $\hat{M}_1(\theta)$:

$$\hat{M}_1(\theta) = \left(\frac{1}{T-1} \sum_{t=2}^T |\Delta q_t - \theta \Delta p_t| \right).$$

Property (i). $\hat{M}_1(\theta)$ is continuous in θ for all $\theta \in \mathbb{R}$.

This property follows because $\hat{M}_1(\theta)$ is a composite of continuous elementary operations.

Property (ii). $\lim_{\theta \rightarrow -\infty} \hat{M}_1(\theta) = \lim_{\theta \rightarrow \infty} \hat{M}_1(\theta) = \infty$.

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

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

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

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

The property then follows immediately because by assumption $\Delta p_t \neq 0$ for some $t \in \{2, \dots, T\}$.

Property (iii). $\hat{M}_1(\theta)$ is convex in θ on \mathbb{R} .

This follows from the convexity of $|x|$ in x on \mathbb{R} , because if $f(x)$ is convex in x then so is $f(ax + b)$.

Now pick $c > \hat{M}_1(0)$. The set $\{\theta \in \mathbb{R} : \hat{M}_1(\theta) \leq c\}$ is closed by (i) and convex by (iii), and so $\min_{\{\theta \in \mathbb{R} : \hat{M}_1(\theta) \leq c\}} \hat{M}_1(\theta)$ must exist. But by (i), (ii), and (iii), $\hat{M}_1(\theta) \geq \min_{\{\theta \in \mathbb{R} : \hat{M}_1(\theta) \leq c\}} \hat{M}_1(\theta)$, so that $\min_{\theta} \hat{M}_1(\theta)$ must exist. Therefore let $\underline{B}_1 = \min_{\theta} \hat{M}_1(\theta)$ and note that $\mathcal{B}(1, \mathbb{R}) = [\underline{B}_1, \infty)$.

Next pick $B \in \mathcal{B}(\underline{B}_1, \mathbb{R})$. The set $\hat{\Theta}_1(B) = \{\theta \in \mathbb{R} : \hat{M}_1(\theta) \leq B\}$ is closed by (i), so define $\underline{\theta}_1(B), \bar{\theta}_1(B)$ as its extreme points. The set $\hat{\Theta}_1(B)$ is convex by (iii), so $\hat{\Theta}_1(B) = [\underline{\theta}_1(B), \bar{\theta}_1(B)]$.

D.2 Nonseparable Model

Relative to Section 4.1, 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 .⁹

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).¹⁰ In the context of our application to the price elasticity of world demand for staple food grains, this means that the intervals

⁹Fixing 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 .

¹⁰Specifically,

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

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 .

D.3 Bounds on an Average Slope

In the setting of Section 4.1, a bound on the size of the shock, coupled with a bound on the variation in the slope of the function $q(\cdot)$, can be used to bound 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. Specifically, 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 $M_k(|\Delta \varepsilon|) \leq B$ on the size of the shocks and a bound $M_k \left(\left| (\vec{\theta} - \bar{\theta}) \circ \Delta p \right| \right) \leq V$ on the scaled deviation of the average slopes from $\bar{\theta}$, then we can say that $\bar{\theta} \in \hat{\Theta}_k(B + V)$.¹¹

¹¹If $q(\cdot)$ is linear, as in equation (1), then $V = 0$.

E Connections to Other Approaches

E.1 Orthogonality Restrictions

Let z_t be some observed variable transformed so that $M_1(\Delta z) = 0$ and $M_2(\Delta z) = 1$.¹² 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 $C = 0$.¹³

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 $\Delta z_t = \Delta p_t$.

E.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 $\Delta p_t^D = \Delta p_t^S = \Delta p_t$.¹⁴ 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 such pair must also be consistent with

¹²Beginning 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.

¹³When $C = 0$, 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.

¹⁴In the world market for staple food grains, the quantity demanded and quantity supplied need not be equal at a given 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).

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 $\theta^S \geq \check{\theta}_2$.

References (Not Appearing in Main Paper)

- 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): 105-55.
- Binder, Ariel J. and John Bound. 2019. The declining labor market prospects of less-educated men. *Journal of Economic Perspectives* 33(2): 163-90.
- Ellison, Glenn and Sara Fisher Ellison. 2009b. Supplement to “Search, obfuscation, and price elasticities on the internet.” *Econometric Society* [publisher], Econometric Society [distributor]. <https://www.econometricsociety.org/content/supplement-search-obfuscation-and-price-elasticities-internet-0>.
- Fukui, Masao, Emi Nakamura, and Jón Steinsson. Forthcoming. Women, wealth effects, and slow recoveries. *American Economic Journal: Macroeconomics*. <https://www.aeaweb.org/articles?id=10.1257/mac.20200312>. Accessed on February 17 2022.
- Greenwood, Jeremy, Nezih Guner, and Guillaume Vandenbroucke. 2017. Family economics writ large. *Journal of Economic Literature* 55(4): 1346-434.
- Juhn, Chinhui and Simon Potter. 2006. Changes in labor force participation in the United States. *Journal of Economic Perspectives* 20(3): 27-46.

- Ottonello, Paolo and Wenting Song. 2022. Financial intermediaries and the macroeconomy: Evidence from a high-frequency identification. *National Bureau of Economic Research Working Paper* No. 29638.
- 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): 528-83.