American Economic Journal: Microeconomics 2023, 15(3): 552–572
https://doi.org/10.1257/mic.20210365

Bounds on a Slope from Size Restrictions on Economic Shocks†

By Marco Stenborg Petterson, David Seim, and Jesse M. Shapiro*

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. (*JEL* D83, E23, G13, Q11)

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

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

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

* Petterson: CSEF and University of Naples Federico II (email: marcostenborg.petterson@unina.it); Seim: Stockholm University, CEPR, and Uppsala University (email: david.seim@su.se); Shapiro: Harvard University and NBER (email: jesse_shapiro@fas.harvard.edu). John Asker was coeditor for this article. 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.
† Go to https://doi.org/10.1257/mic.20210365 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

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 $\left(|\Delta\varepsilon_2|, \ldots, |\Delta\varepsilon_T|\right)$ 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

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

of or uncorrelated with one another (e.g., Leamer 1981; Feenstra 1994; Feenstra and Weinstein 2017; MacKay and Miller 2023).[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 (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 I presents our setup and results. Section II presents our application. Section III presents extensions. Section IV concludes. An Appendix includes proofs of results stated in the text. An online Appendix discusses additional extensions, applications, and connections. An accompanying python package, PyBounds, facilitates adoption of our approach (Petterson et al. 2022).

[2] See also Leontief (1929). Morgan (1990, chapter 6) quotes a 1913 thesis by Lenoir that 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.

[3] In 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.

I. Setup and Characterization of Sets of Interest

A. *Setup*

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

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

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

Now let $\hat{M}_k(\theta) = M_k(|\boldsymbol{\Delta}\boldsymbol{\varepsilon}(\theta)|)$ denote the k-mean of the absolute value of the vector $\boldsymbol{\Delta}\boldsymbol{\varepsilon}(\theta) = (\Delta\varepsilon_2(\theta), \ldots, \Delta\varepsilon_T(\theta))$, 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

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

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

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

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

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

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]

B. *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.

PROPOSITION 1: *Let*

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

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

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

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

All proofs are given in the Appendix. The objects $\underline{B}_\infty$, $\underline{\theta}_\infty(B)$, 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 1: 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 $\underline{B}_\infty$ in a case with multiple unknown slope parameters.

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

[6] More 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 $|\hat{\Theta}_\infty(B)| = 1$.

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

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

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

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

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

Proposition 2 shows that $\mathcal{B}(k, \mathbb{R})$ is a left-bounded interval the limit point $\underline{B}_k$ of which 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 the limit points of which can be calculated by solving the nonlinear equation $\hat{M}_k(\theta) = B$. By Lemma 1, on either side of $\breve{\theta}_k$ and for $B > \underline{B}_k$ the equation is strictly monotone and has a unique solution, which simplifies computation. 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 $\overline{\theta}_k(B) = \theta$ 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 $\underline{\theta}_k(B), \overline{\theta}_k(B)$ 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 $\underline{B}_2$, $\underline{\theta}_2(B)$, $\overline{\theta}_2(B)$, and $\breve{\theta}_2$ described in Proposition 2 are available in closed form. Toward a characterization, for any D-dimensional vector $\mathbf{v} \in \mathbb{R}^D$, let $\boldsymbol{\Delta}\mathbf{v} = (\Delta v_2, \ldots, \Delta v_D) \in \mathbb{R}^{D-1}$. For any $\mathbf{v}, \mathbf{w} \in \mathbb{R}^D$, let $\hat{s}_{vw} = M_1(\boldsymbol{\Delta}\mathbf{v} \circ \boldsymbol{\Delta}\mathbf{w})$, where $\circ$ is the elementwise product. We then have the following.

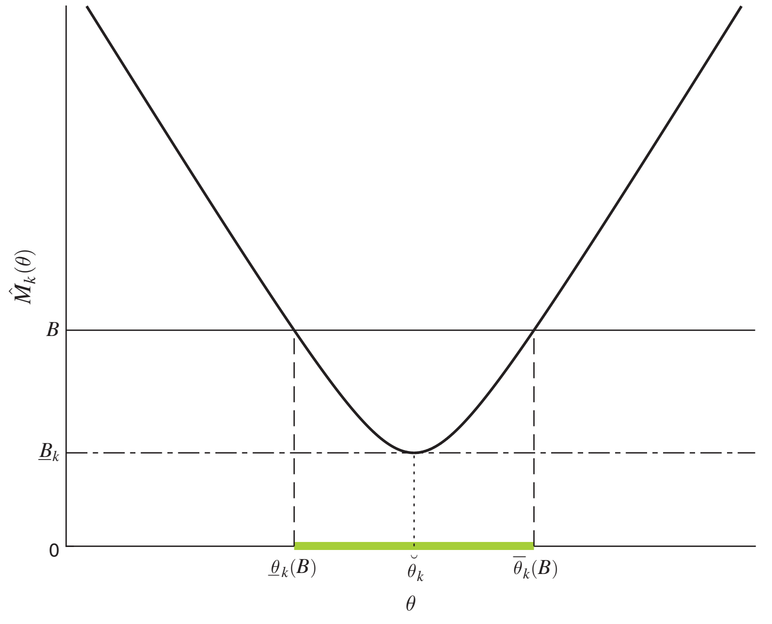

FIGURE 1. ILLUSTRATION OF PROPOSITION 2

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 $\hat{\Theta}_k(B) = [\underline{\theta}_k(B), \overline{\theta}_k(B)]$. Moreover, any bound B that is below $\underline{B}_k = \hat{M}_k(\breve{\theta}_k)$ is incompatible with the data because it lies below $\hat{M}_k(\theta)$ for all θ.

COROLLARY 1: *For $k = 2$ we have that*

$$\underline{\theta}_2(B) = \frac{\hat{s}_{qp}}{\hat{s}_{pp}} - \sqrt{\left(\frac{\hat{s}_{qp}}{\hat{s}_{pp}}\right)^2 - \frac{1}{\hat{s}_{pp}}\left(\hat{s}_{qq} - B^2\right)},$$

$$\overline{\theta}_2(B) = \frac{\hat{s}_{qp}}{\hat{s}_{pp}} + \sqrt{\left(\frac{\hat{s}_{qp}}{\hat{s}_{pp}}\right)^2 - \frac{1}{\hat{s}_{pp}}\left(\hat{s}_{qq} - B^2\right)},$$

$$\underline{B}_2 = \sqrt{\hat{s}_{qq} - \left(\frac{\hat{s}_{qp}}{\hat{s}_{pp}}\right)^2 \hat{s}_{pp}},$$

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

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

Appendixes E.1 and E.2 further discuss the connection between orthogonality restrictions and those we consider here.

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

A. *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, b) on annual world population and GDP. From these we construct a time series $\left\{\left(p_t^D, q_t^D\right)\right\}_{t=1}^{T}$, 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]

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 (Roberts and Schlenker 2013a, 2278–79).

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., Roberts and Schlenker 2013a, 2274–75). 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.

B. *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

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

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

[9] They 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, 2272).

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 $M_\infty(|0.37\mathbf{\Delta y}|) \approx 0.05$. The root mean squared value is $M_2(|0.37\mathbf{\Delta y}|) \approx 0.02$. 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 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.

C. *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 $\hat{\theta}^D_{RS} = -0.066$ given in Roberts and Schlenker (2013a, column 2c of table 1), 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 I, 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.

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

[11] Muhammad et al. (2011) estimate a model of food demand using country-level data from 2005. Alexandratos and Bruinsma (2012, 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, 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.

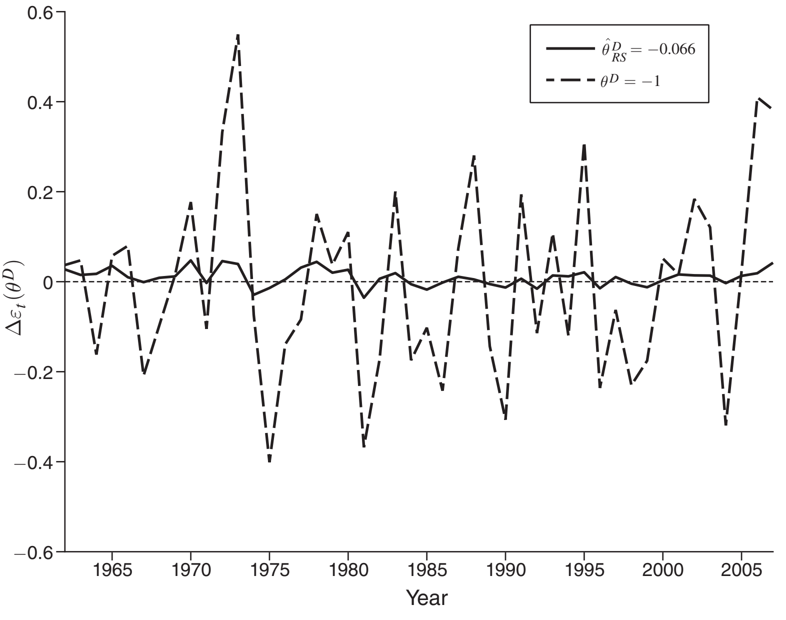

Figure 2. Implied Shocks to World Demand for Food Grain under Different Elasticities

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

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, $M_\infty(|0.37\boldsymbol{\Delta}\mathbf{y}|) \approx 0.05$. The figure depicts a scatterplot of the first-differenced data $\left\{\left(\Delta p_t^D, \Delta q_t^D\right)\right\}_{t=2}^T$. In first differences, a demand function is a line through the origin with nonpositive slope $\theta^D \in \bar{\Theta}^D = \mathbb{R}_{\leq 0}$. The figure also 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, column 2c of table 1) 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.

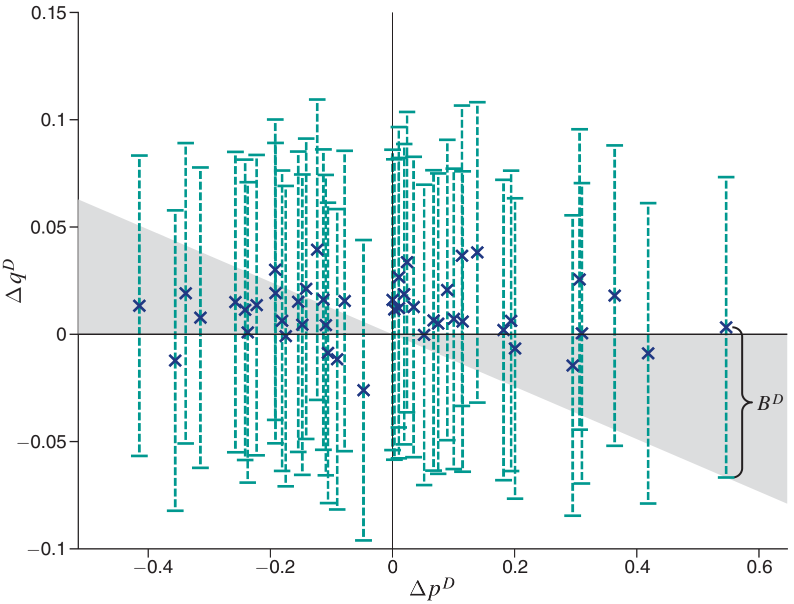

FIGURE 3. CONSTRUCTING BOUNDS ON AN ELASTICITY FROM BOUNDS ON SHOCKS

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 II. The crosshatches depict a scatterplot of the data $\{(\Delta p_t^D, \Delta q_t^D)\}_{t=2}^T$. The dotted interval around each crosshatch 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 $\bar{\Theta}^D = \mathbb{R}_{\leq 0}$.

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 $M_k(|0.37\boldsymbol{\Delta}\mathbf{y}|)$ 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, column 2c of table 1). 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

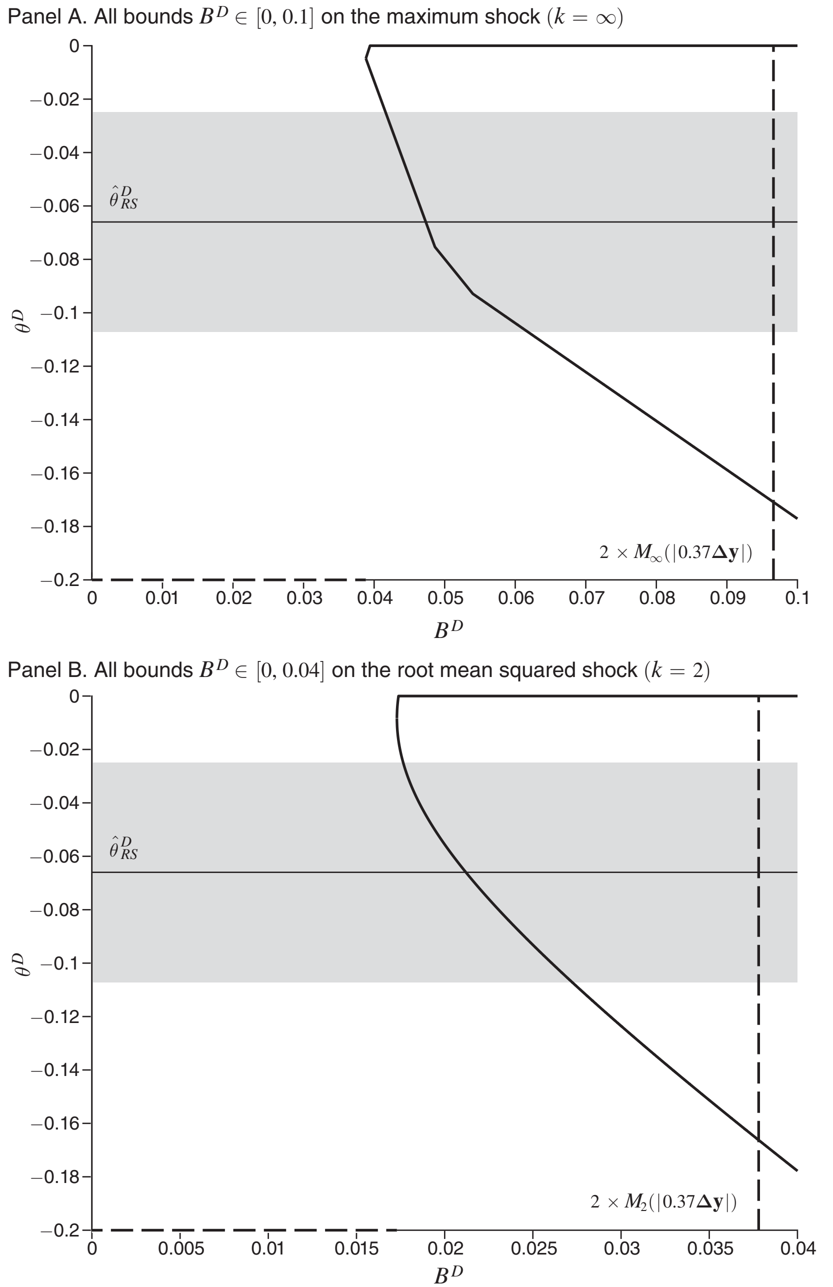

Figure 4. Implications of Bounds on Shocks to World Demand for Food Grain

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 II. Panel A depicts the interval $\hat{\Theta}_\infty(B^D) \cap \bar{\Theta}^D$ implied by bounds $B^D \in [0, 0.1]$ 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 $M_\infty(|0.37\mathbf{\Delta y}|)$. 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 $M_2(|0.37\mathbf{\Delta y}|)$. In each plot, the horizontal line depicts the point estimate $\hat{\theta}^D_{RS}$ of the price elasticity of demand in Roberts and Schlenker (2013a, column 2c of table 1), and the shaded region depicts the associated 95 percent 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.

of staple grains based on bounds B^S on the size of shocks to supply, illustrated in online Appendix Figure A1. Online Appendix A.2 characterizes bounds on a function of the elasticities θ^D and θ^S, illustrated in online Appendix Figure A2 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 A3 with an application to time trends considered in Roberts and Schlenker (2013a). Lastly, online Appendix Figure A4 illustrates the role of k by showing how $M_k(|0.37\boldsymbol{\Delta}\mathbf{y}|)$ and the value of B^D needed to obtain a given bound on the price elasticity vary with k.

III. Extensions and Discussion

A. *Nonlinear Model*

In the setting of Section II 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

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

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

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

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 $\theta^D_{t-1,t}$ 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 $\theta^D_{t-1,t} > -1$.

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 $\theta^D_{t-1,t} > -0.3$.

Online Appendix D.2 shows that in the case of a nonseparable model, an analog 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 $q(\cdot)$.

B. *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 I, 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 I and therefore requires a different approach or 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.

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

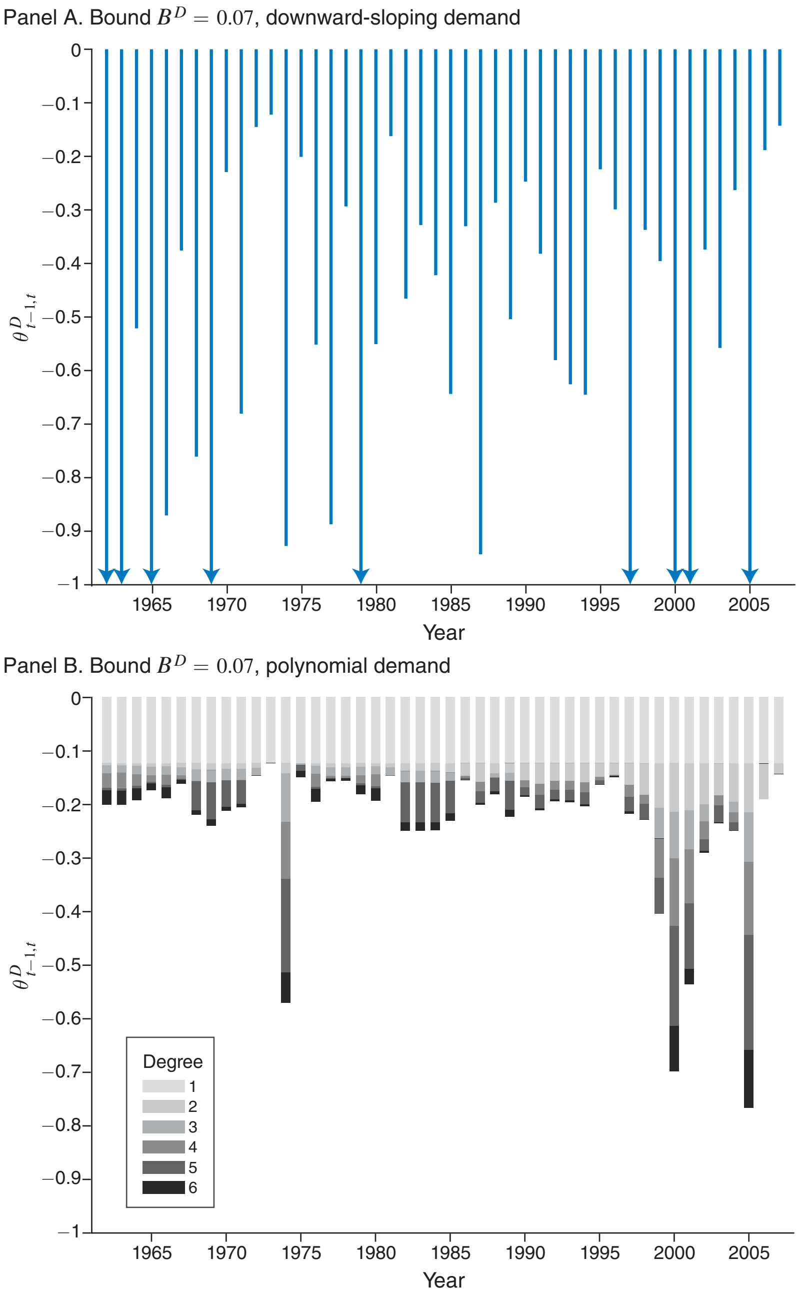

Figure 5. Relaxing Linearity of the Demand Function

Notes: Each plot depicts bounds on the average price elasticity of demand $\theta^D_{t-1,t}$ 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, the derivative of which 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 1 to 6).

IV. 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.

Appendix: Proofs of Results Stated in Text

A. *Proof of Proposition 1*

We have that

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

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

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

for all $t \geq 2$. 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 $\hat{\Theta}_\infty(B) = \varnothing$. 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 $\hat{\Theta}_\infty(B) = \varnothing$; otherwise, $\hat{\Theta}_\infty(B) = \left[\underline{\theta}_\infty(B), \bar{\theta}_\infty(B)\right]$. Notice that $\underline{\theta}_\infty(B)$ is continuous and strictly decreasing in B with $\lim_{B\to\infty} \underline{\theta}_\infty(B) = -\infty$ and that $\bar{\theta}_\infty(B)$ is continuous and strictly increasing in B with $\lim_{B\to\infty} \bar{\theta}_\infty(B) = \infty$. 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. ∎

B. *Proof of Lemma 1*

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

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

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

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

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

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

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

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

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

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

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

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

We have that

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

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

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

Pick some $c' > \hat{M}_k(0)$. By Properties (i) and (ii), there are at least two solutions to $c' = \hat{M}_k(\theta)$. By Property (iii), there are at most two solutions to $(c')^k = \left(\hat{M}_k(\theta)\right)^k$. Hence, there are exactly two solutions to $c' = \hat{M}_k(\theta)$; denote these $\underline{\theta}(c'), \overline{\theta}(c')$, with $\underline{\theta}(c') < \overline{\theta}(c')$. Because the interval $\left[\underline{\theta}(c'), \overline{\theta}(c')\right]$ is compact, by Properties (i) and (iii), $\left(\hat{M}_k(\theta)\right)^k$ has a minimum on $\left[\underline{\theta}(c'), \overline{\theta}(c')\right]$ at some unique $\breve{\theta}_k$ on the interior of $\left[\underline{\theta}(c'), \overline{\theta}(c')\right]$. But also by Property (iii), $\left(\hat{M}_k(\theta)\right)^k > \left(\hat{M}_k(\breve{\theta}_k)\right)^k$ for any $\theta \notin \left[\underline{\theta}(c'), \overline{\theta}(c')\right]$, and hence, $\breve{\theta}_k = \arg\min_\theta\left(\hat{M}_k(\theta)\right)$.

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

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

C. *Proof of Proposition 2*

This follows immediately from Lemma 1. ∎

D. *Proof of Corollary 1*

We have that

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

By Lemma 1, $\hat{M}_2(\theta)$ has a unique global minimizer $\breve{\theta}_2$. Because $\hat{M}_2(\theta)$ is nonnegative and is differentiable in θ when $\hat{M}_2(\theta) > 0$, either $\hat{M}_2(\breve{\theta}_2) = 0$ or $\hat{M}_2(\breve{\theta}_2) > 0$ and $\frac{d}{d\theta}\hat{M}_2(\theta)|_{\theta=\breve{\theta}_2} = 0$. In either case, we have that

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

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

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

It then follows that

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

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

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

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

The roots of this quadratic equation are given by

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

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

REFERENCES

Alexandratos, Nikos, and Jelle Bruinsma. 2012. "World Agriculture towards 2030/2050: The 2012 Revision." FAO ESA Working Paper 12-03.

Andrews, Isaiah, and Jesse M. Shapiro. 2021. "A Model of Scientific Communication." *Econometrica* 89 (5): 2117–42.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro. 2020. "Transparency in Structural Research." *Journal of Business and Economic Statistics* 38 (4): 711–22.

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." In *Zeitschrift Des Statistischen Bureaus Des Königlich Sächsischen Ministerium Des Inneren*, Vol. 8–9, edited by Ernst Engel, 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." NBER Working Paper 25311.

Giacomini, Raffaella, Toru Kitagawa, and Matthew Read. 2021. "Identification and Inference under Narrative Restrictions." *arXiv*: 2102.06456.

Houthakker, H.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." NBER Working Paper 23225.

MacKay, Alexander, and Nathan H. Miller. 2023. "Estimating Models of Supply and Demand: Instruments and Covariance Restrictions." Harvard Business School Working Paper 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 *Handbook of Econometrics*, Vol. 6, Part B, edited by James J. Heckman and Edward E. Leamer, 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, Hélène Piet-Lahanier, and Eric Walter. 1996. *Bounding Approaches to System Identification*. Boston, MA: Springer.

Morgan, Mary S. 1990. *The History of Econometric Ideas*. Cambridge, UK: 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*. Washington, DC: United States Department of Agriculture.

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.

Petterson, Marco Stenborg, David Seim, and Jesse M. Shapiro. 2023. "Replication data for: Bounds on a Slope from Size Restrictions on Economic Shocks." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E168561V1.

Petterson, Marco Stenborg, David Seim, Jesse M. Shapiro, and Nathan Sun. 2022. "PyBounds." JMSLab. https://github.com/JMSLab/PyBounds (accessed in April 2022).

Robbins, Herbert, and Cun-Hui Zhang. 1986. "Maximum Likelihood Estimation in Regression with Uniform Errors." In *Institute of Mathematical Statistics Lecture Notes*, Vol. 8, edited by John Van Ryzin, 365-85. Hayward, CA: Institute of Mathematical Statistics.

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. "Identifying Supply and Demand Elasticities of Agricultural Commodities: Implications for the US Ethanol Mandate: Dataset." *American Economic Review*. https://www.aeaweb.org/articles?id=10.1257/aer.103.6.2265.

Tamer, Elie. 2010. "Partial Identification in Econometrics." *Annual Review of Economics* 2 (1): 167–95.

Uhlig, Harald. 2017. "Shocks, Sign Restrictions, and Identification." In *Advances in Economics and Econometrics: Eleventh World Congress of the Econometric Society*, Vol. 2, edited by Bo Honoré, Ariel Pakes, Monika Piazzesi, and Larry Samuelson, 95–127. Cambridge, UK: Cambridge University Press.

Valin, Hugo, Ronald D. Sands, Dominique van der Mensbrugghe, Gerald C. Nelson, Helal Ahammad, Elodie Blanc, Benjamin Bodirsky, et al. 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 Hélène 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 Hélène Piet-Lahanier. 1996. "Recursive Robust Minimax Estimation." In *Bounding Approaches to System Identification*, edited by Mario Milanese, John Norton, Hélène Piet-Lahanier, and Eric Walter, 183–97. Boston, MA: Springer.

World Bank. 2019a. "Population, Total." World Bank Group. https://data.worldbank.org/indicator/SP.POP.TOTL (accessed on October 2, 2019).

World Bank. 2019b. "GDP (Current US$)." World Bank Group. https://data.worldbank.org/indicator/NY.GDP.MKTP.CD (accessed on October 24, 2019).