

Online Appendix for
Bounds on a Slope from Size Restrictions on Economic Shocks

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

A. EXTENSIONS OF ANALYSIS OF WORLD MARKET FOR STAPLE FOOD GRAINS

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

Here we explore the information about the price elasticity of supply $\theta^S \in \overline{\Theta}^S = \mathbb{R}_{\geq 0}$ that can be obtained from imposing a bound B^S on the size of shocks to supply. From the data described in Section II 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 A1 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 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.

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.

*Petterson: Brown University (email: marco.stenborg_petterson@brown.edu); Seim: Stockholm University, CEPR, and Uppsala University (email: david.seim@su.se); Jesse M. Shapiro (email: jesse_shapiro@fas.harvard.edu)

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

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 A2 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.

3. Orthogonalization with Respect to Covariates

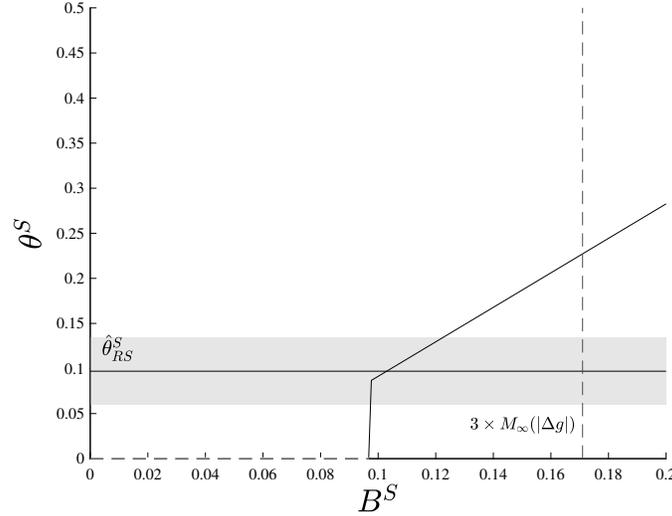
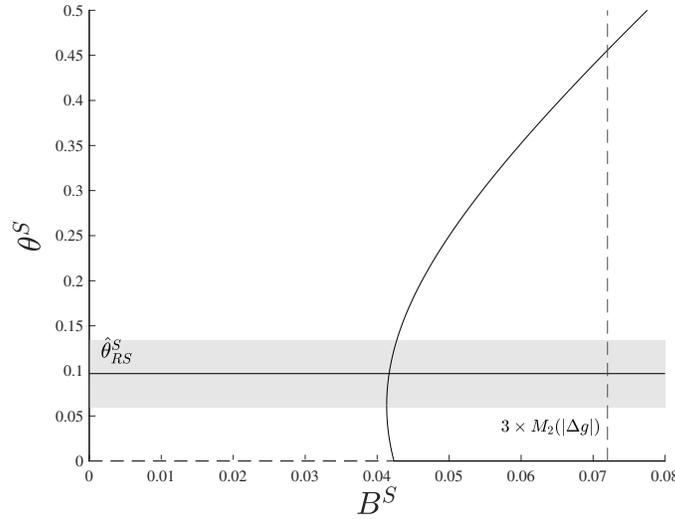
Let $\{x_t\}_{t=1}^T$ be an observed sequence of values of a (possibly vector-valued) covariate. For any θ , let $\Delta\varepsilon_{\perp}(\theta)$ be the component of $\Delta\varepsilon(\theta)$ orthogonal to $\Delta\mathbf{x} = (\Delta x'_2, \dots, \Delta x'_T)$.³ If we are prepared to impose an upper 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 $\Delta\varepsilon(\theta)$.

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 A3 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 A3 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.

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).

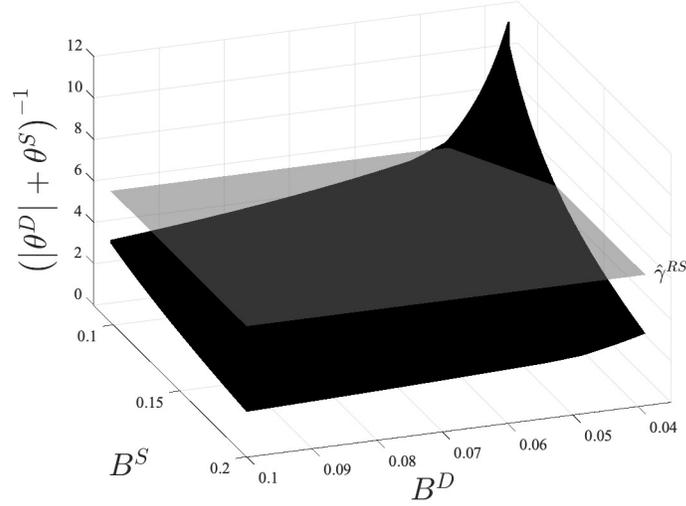
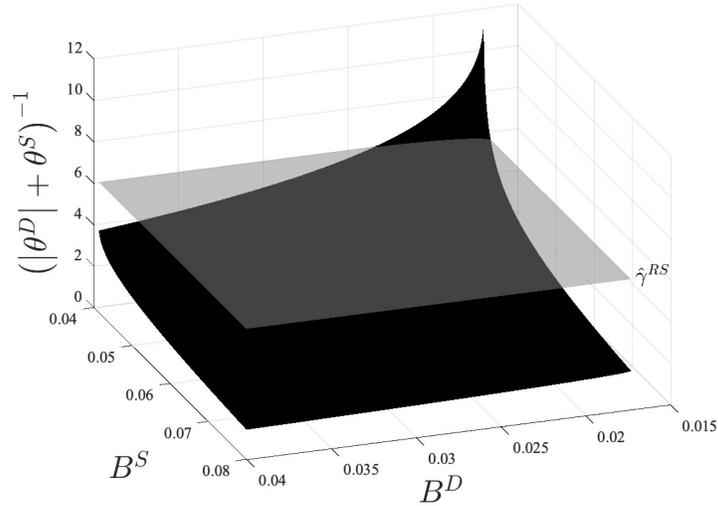
³That is, $\Delta\varepsilon_{\perp}(\theta) = \Delta\varepsilon(\theta) - \Delta\mathbf{x}(\Delta\mathbf{x}'\Delta\mathbf{x})^{-1}\Delta\mathbf{x}'\Delta\varepsilon(\theta)$.

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 ($k = 2$)

ONLINE APPENDIX FIGURE A1. IMPLICATIONS OF BOUNDS ON SHOCKS TO WORLD SUPPLY OF FOOD GRAIN

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 $\bar{\Theta}^S = \mathbb{R}_{>0}$.

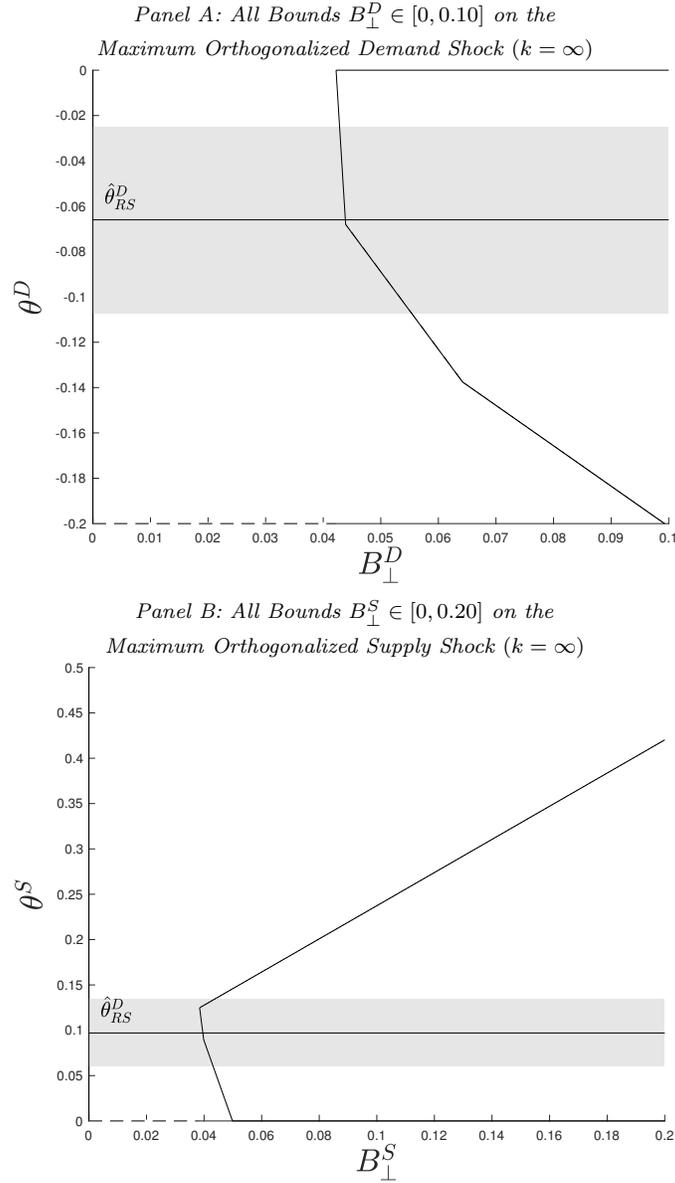
The dashed vertical line is at three times the maximum absolute yield shock $M_\infty(|\Delta g|)$. Panel B depicts the interval $\hat{\Theta}_2(B^S) \cap \bar{\Theta}^S$ implied by bounds $B^D \in [0, 0.08]$ on the root mean squared shock. The dashed vertical line is at three times the root mean squared 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.

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

ONLINE APPENDIX FIGURE A2. IMPLICATIONS OF BOUNDS ON SHOCKS FOR THE MULTIPLIER PARAMETER

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 ($k = \infty$).

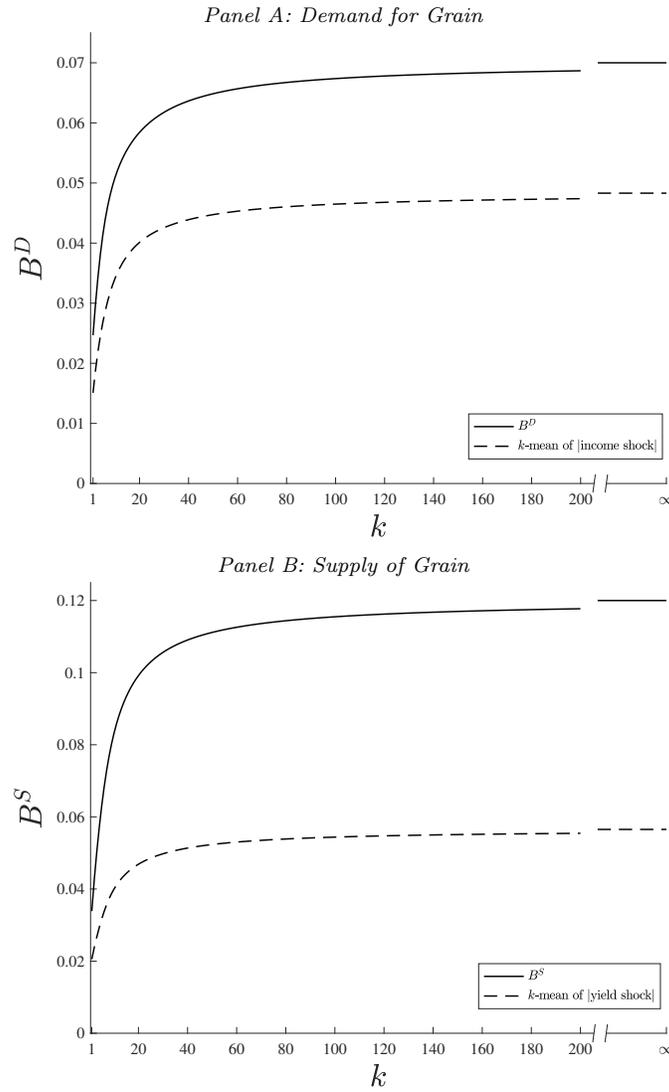
Panel B considers bounds $B^D \in [0.015, 0.04]$, $B^S \in [0.040, 0.08]$ on the root mean squared shock ($k = 2$). In each plot, the black surface depicts the lowest value of the multiplier $\gamma(\theta^D, \theta^S) = (|\theta^D| + \theta^S)^{-1}$ that is compatible with elasticities $\theta^D \in \hat{\Theta}_k(B^D) \cap \bar{\Theta}^D$, $\theta^S \in \hat{\Theta}_k(B^S) \cap \bar{\Theta}^S$, 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 A3. IMPLICATIONS OF BOUNDS ON ORTHOGONALIZED SHOCKS TO WORLD DEMAND AND SUPPLY FOR FOOD GRAIN

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 $\{\theta \in \bar{\Theta}^D : M_{\infty}(|\Delta \epsilon_{\perp}^D(\theta)|) \leq B_{\perp}^D\}$ 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 $\{\theta \in \bar{\Theta}^S : M_{\infty}(|\Delta \epsilon_{\perp}^S(\theta)|) \leq B_{\perp}^S\}$ 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 A4. 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

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 $\Delta \boldsymbol{\varepsilon}_i(\theta) = (\Delta \varepsilon_{i2}(\theta), \dots, \Delta \varepsilon_{iT_i}(\theta))$, where $\Delta \varepsilon_{it} = \Delta q_{it} - \theta \Delta p_{it}$, and define $\hat{M}_{ik}(\theta) = M_k(|\Delta \boldsymbol{\varepsilon}_i(\theta)|)$ 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 I, 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.

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 (Fukui, Nakamura, and Steinsson 2021). 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

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 .

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 B1 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

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.

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 $k \geq 1$.

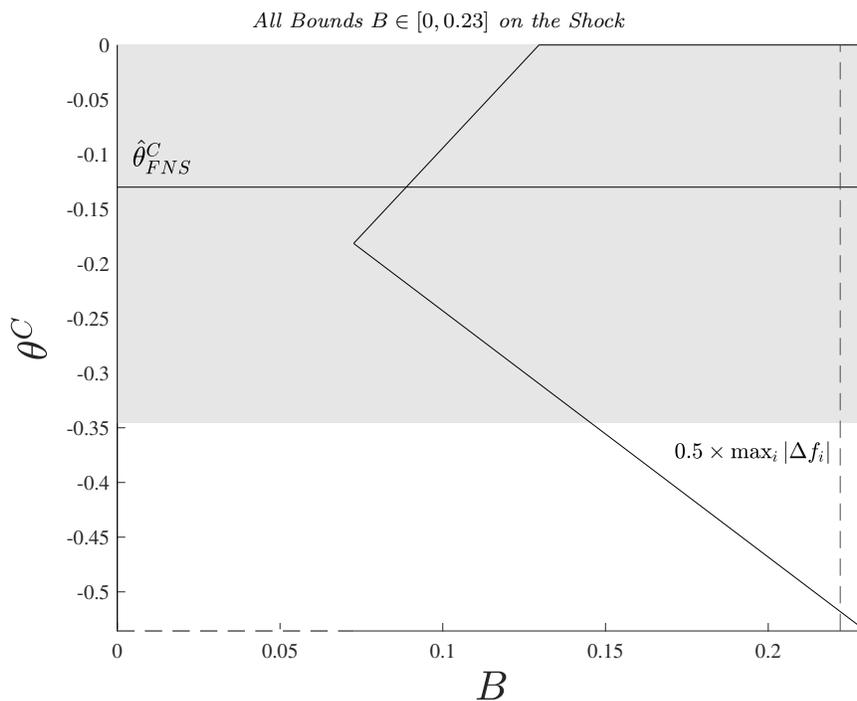
⁷Fukui, Nakamura, and Steinsson (forthcoming, Table 3, Column 2) report a revised point estimate of $\hat{\theta}_{FNS}^C = -0.18$ with a confidence interval of $[-0.34, -0.02]$.

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 .

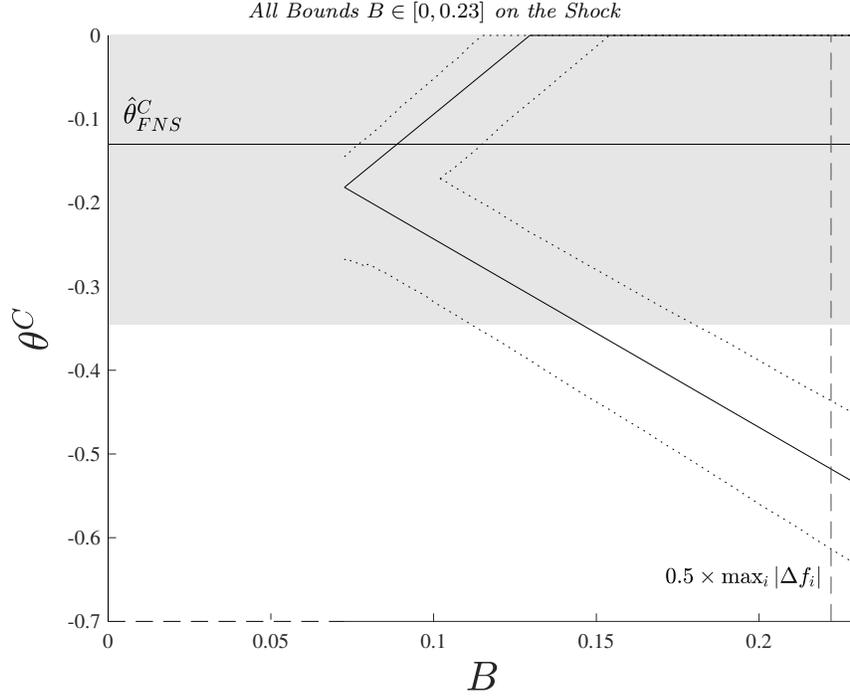
3. Accounting for Sampled Data

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 B2 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 B1. 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, Nakamura, 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, Nakamura, 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 B2. 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, Nakamura, and Steinsson (2020) described in Online Appendix B.2. Following Online Appendix Figure B1, 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, Nakamura, 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

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

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.

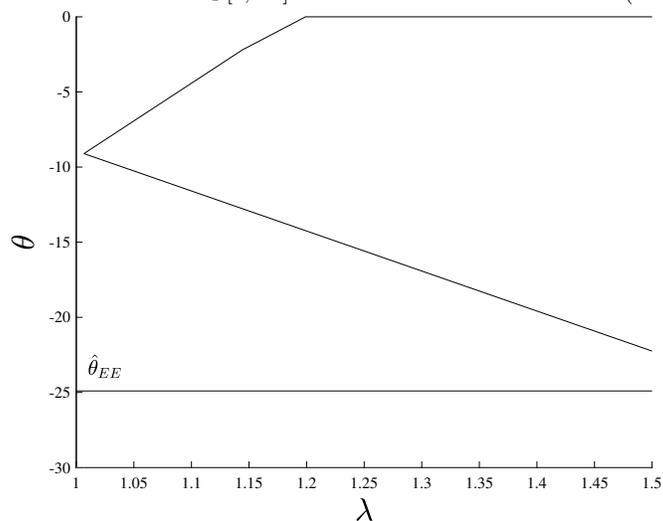
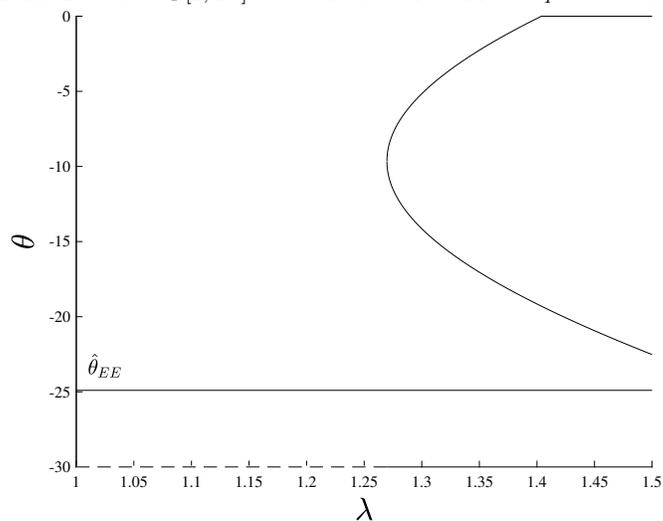
Ellison and Ellison (2009a, p. 440) assume that, when demand is positive, the demand curve takes a log-linear form consistent with equation (1), though

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

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 C1 uses these as a point of reference for the construction of bounds on the shocks $\hat{M}_k^{\mathcal{T}}(\theta)$ 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 $\hat{M}_k^{\mathcal{T}}(\theta) > \hat{M}_k^{\mathcal{S}}$, meaning that the bound $\hat{M}_k^{\mathcal{T}}(\theta) \leq \lambda \hat{M}_k^{\mathcal{S}}$ 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 II, 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.

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 ($k = 2$)

ONLINE APPENDIX FIGURE C1. IMPLICATIONS OF BOUNDS ON SHOCKS TO MEMORY MODULE DEMAND

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 $\hat{M}_k^T(\theta) \leq \lambda \hat{M}_k^S$ on the shock where \hat{M}_k^S is the k -mean of the shock in periods with no price change, $\hat{M}_k^T(\theta)$ 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

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

2. Nonseparable Model

Relative to Section III.A, a further relaxation of the model in equation (1) can be written as

$$(D1) \quad q_t = \tilde{q}(p_t, \varepsilon_t)$$

where ε_t may now be non-scalar or even infinite-dimensional. The model in equation (D1) 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 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 .

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,

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

3. Bounds on an Average Slope

In the setting of Section III.A, 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

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

Therefore if we are prepared to impose a bound $M_k(|\Delta \boldsymbol{\varepsilon}|) \leq B$ on the size of the shocks and a bound $M_k \left(\left| (\vec{\theta} - \bar{\theta}) \circ \Delta \mathbf{p} \right| \right) \leq V$ 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

1. Orthogonality Restrictions

Let z_t be some observed variable transformed so that $M_1(\Delta \mathbf{z}) = 0$ and $M_2(\Delta \mathbf{z}) = 1$.¹² Consider a restriction of the form

$$(E1) \quad |M_1(\Delta \boldsymbol{\varepsilon}(\theta) \circ \Delta \mathbf{z})| \leq C$$

where $C \geq 0$ is a scalar. An orthogonality restriction is such a restriction that takes $C = 0$.¹³

Restrictions of the form in (E1) are related to those we consider in the sense that, from the Cauchy-Schwarz inequality and the fact that $\Delta \mathbf{z}$ is standardized,

$$(M_1(\Delta \boldsymbol{\varepsilon}(\theta) \circ \Delta \mathbf{z}))^2 \leq (M_2(\Delta \boldsymbol{\varepsilon}(\theta)))^2.$$

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

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

$$(E2) \quad \frac{1}{T-1} \sum_{t=2}^T \Delta \varepsilon_t(\theta) \Delta p_t = 0.$$

For Δp_t standardized, equation (E2) is equivalent to an orthogonality restriction with $\Delta z_t = \Delta p_t$.

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

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

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

¹²Beginning with a variable \tilde{z}_t we can take $z_t = M_2(\Delta \tilde{\mathbf{z}} - M_1(\Delta \tilde{\mathbf{z}}) \mathbf{J}_{T-1,1})^{-1}(\tilde{z}_t - (t-1)M_1(\Delta \tilde{\mathbf{z}}))$, for $\mathbf{J}_{T-1,1}$ the $(T-1)$ -dimensional vector of ones.

¹³When $C = 0$, the inequality in (E1) implies that $\theta = M_1(\Delta \mathbf{q} \circ \Delta \mathbf{z}) / M_1(\Delta \mathbf{p} \circ \Delta \mathbf{z})$ 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).

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

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

$$(E3) \quad |M_1(\Delta \varepsilon^D(\theta^D) \circ \Delta \varepsilon^S(\theta^S))| \leq R.$$

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.** 2021. Data for "Women, wealth effects, and slow recoveries." Received from the authors on October 13, 2020 and April 21, 2021.
- 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.
- Ottoneo, 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.