Robust Measures for Welfare Analysis

Zi Yang Kang — *Harvard*

Shoshana Vasserman — *Stanford GSB*

Motivation

"Economists have made remarkable progress over the last several decades in developing empirical techniques that provide compelling **evidence of causal effects**—the so-called **'credibility revolution"** in empirical work…

But while it is interesting and important to know what the effects of a policy are, we are often also interested in a **normative question** as well: Is the policy a **good** idea or a **bad** idea?

…What is the **welfare impact of the policy**?"

—Finkelstein and Hendren (2020)

Motivation

- Measuring welfare requires taking a stance on behavior at unobserved points.

Motivation

- Measuring welfare requires taking a stance on behavior at unobserved points.

- Many papers impose "standard" functional form assumptions.

  → Linear interpolation: $D_{\text{linear}}(p) = A - \beta p.$

    - Harberger (1964); Hackmann et al. (2015); Amiti et al. (2019); Hahn and Metcalfe (2021).

  → Isoelastic interpolation: $D_{\text{isoelastic}}(p) = Ap^{-\varepsilon}.$

    - Hausman (1981); Hausman et al. (1997); Brynjolfsson et al. (2003); Fajgelbaum et al. (2020).

Motivation

- Measuring welfare requires taking a stance on behavior at unobserved points.
- Many papers impose "standard" functional form assumptions.
  - → Linear interpolation: $D_{\text{linear}}(p) = A - \beta p.$
    - Harberger (1964); Hackmann et al. (2015); Amiti et al. (2019); Hahn and Metcalfe (2021).
  - → Isoelastic interpolation: $D_{\text{isoelastic}}(p) = Ap^{-\varepsilon}.$
    - Hausman (1981); Hausman et al. (1997); Brynjolfsson et al. (2003); Fajgelbaum et al. (2020).

How robust are welfare estimates to the choice of functional form assumption?

This Paper

- We establish measures of **robustness** for quantitative welfare conclusions.
  - How much **variability** in the demand curve can there be before the conclusion flips?
- We parametrize variability through conditions on **gradients** and **curvature**.
  - In each case, we obtain a single-dimensional statistic of relative robustness.

This Paper

- We establish measures of **robustness** for quantitative welfare conclusions.
  - How much **variability** in the demand curve can there be before the conclusion flips?
- We parametrize variability through conditions on **gradients** and **curvature**.
  - In each case, we obtain a single-dimensional statistic of relative robustness.
- To guarantee robustness, we establish **welfare bounds**.
  - These bounds are **robust**: they give the *best-case* and *worst-case* welfare estimates that are consistent with any demand curve within a class of variability.
  - These bounds are also **simple**: we can compute them in closed form.

Framework

Potential Outcomes for Demand: An Experimental Ideal

- Suppose we randomly assign prices for a good to two groups:
  - Group $t = 0$ gets price p_0.
  - Group $t = 1$ gets price p_1.
  - We observe individual i buying y_{it} units at her assigned price p_t.

Potential Outcomes for Demand: An Experimental Ideal

- Suppose we randomly assign prices for a good to two groups:
  - Group $t = 0$ gets price p_0.
  - Group $t = 1$ gets price p_1.
  - We observe individual i buying y_{it} units at her assigned price p_t.
- Consider the *potential outcomes*:

$$y_i = \begin{cases} y_{i1} & \text{if } t = 1, \\ y_{i0} & \text{if } t = 0. \end{cases}$$

Potential Outcomes for Demand: An Experimental Ideal

- Suppose we randomly assign prices for a good to two groups.

- Consider the *potential outcomes*:

$$y_i = \begin{cases} y_{i1} & \text{if } t = 1, \\ y_{i0} & \text{if } t = 0. \end{cases}$$

- Define aggregate demand:

$$D(p_t) = \mathbf{E}[y_{it}] \qquad \text{for } t = 0, 1.$$

- With sample estimator:

$$\hat{D}(p_t) = \frac{1}{n_t} \sum_{i=1}^{n_t} y_{it} \qquad \text{for } t = 0, 1.$$

Potential Outcomes for Demand: An Experimental Ideal

- Our goal is to estimate the difference in consumer surplus between the two groups.

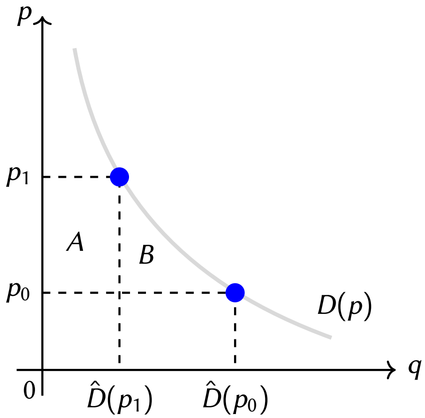

- With $D(p)$, the difference in CS is equal to:

$$\underbrace{\text{area } A}_{=(p_1-p_0)\hat{D}(p_1)} + \text{area } B = \int_{p_0}^{p_1} D(p)\, \mathrm{d}p.$$

- **Main challenge:**

  $D(p)$ isn't identified between p_0 and p_1.

Common Approach: Linear Interpolation

- Our goal is to estimate the difference in consumer surplus between the two groups.

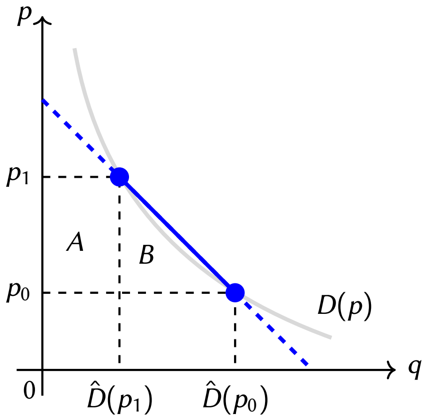

- Estimate regression:

$$y_{it} = \theta_1 - \theta_2 p_t + \epsilon_{it}.$$

- Integrate under $\hat{D}(p) = \hat{\theta}_1 - \hat{\theta}_2 p$ (w.r.t. p):

$$\widehat{\Delta CS}_{\text{linear}} = \frac{1}{2}(p_1 - p_0)\left[\hat{D}(p_1) + \hat{D}(p_0)\right].$$

Common Approach: Isoelastic Interpolation

- Our goal is to estimate the difference in consumer surplus between the two groups.

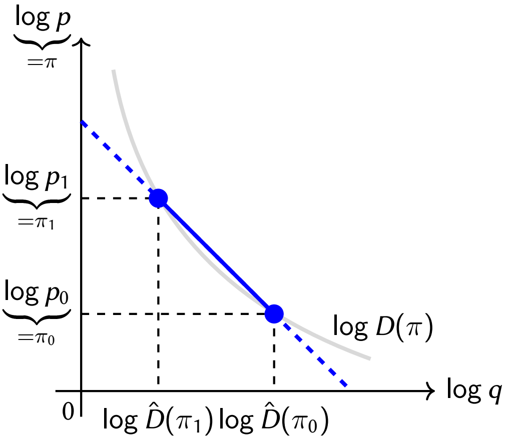

- Estimate regression:

$$\log(y_{it}) = \theta_1 - \theta_2 \log(p_t) + \epsilon_{it}.$$

- Integrate under $\hat{D}(\log p) = \hat{\theta}_1 p^{-\hat{\theta}_2}$ (w.r.t. p):

$$\widehat{\Delta \text{CS}}_{\text{isoelastic}} = \frac{(p_1 \hat{q}_1 - p_0 \hat{q}_0)\log(p_1/p_0)}{\log(\hat{q}_1/\hat{q}_0) + \log(p_1/p_0)},$$

$$\text{where} \quad \hat{q}_t = \hat{D}(\log p_t).$$

How different are these functional forms?

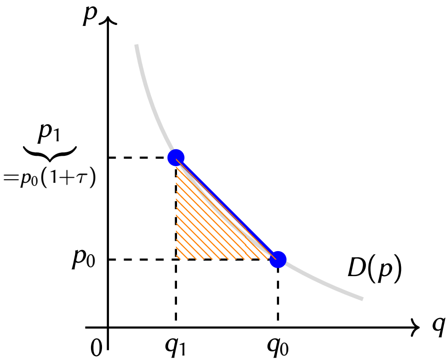

- Example from Trump tariffs: Amiti et al. (2019).
- **Setting:** 2018 trade war involved tariffs as high as 30–50%.
- **Question:** What was the DWL due to tariffs?
- **Approach:** Compare monthly prices and quantities by item in 2017 vs. 2018.
- **Method:** Approximate $D(p)$ with a linear curve; integrate under the curve.

DWL estimates based on different functional forms

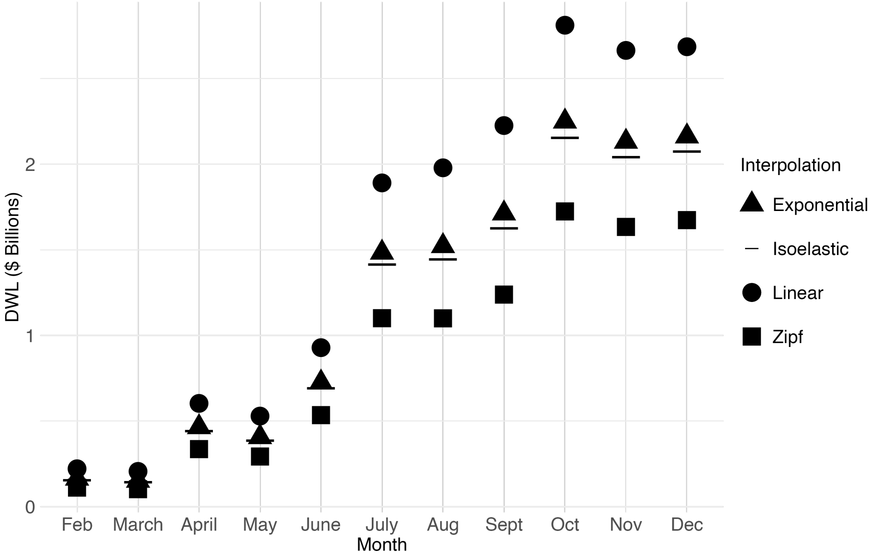

Parametrizing variability in demand curves

- Two commonly used functional form assumptions are linear and isoelastic demand.
  - Linear demand: constant gradient, zero curvature. $\rightsquigarrow$ of demand w.r.t. price
  - Isoelastic demand: constant gradient, zero curvature. $\rightsquigarrow$ of log-demand w.r.t. log-price

Parametrizing variability in demand curves

- Two commonly used functional form assumptions are linear and isoelastic demand.
  - Linear demand: constant gradient, zero curvature. $\rightsquigarrow$ of demand w.r.t. price
  - Isoelastic demand: constant gradient, zero curvature. $\rightsquigarrow$ of log-demand w.r.t. log-price

Generalization: $A(q)$ is affine in $B(p)$, where A, B are continuous and increasing.

$\rightsquigarrow$ **E.g.,** $A = B = \text{id}$ (linear); $A = B = \log$ (isoelastic); $A = \log$, $B = \text{id}$ (exponential)…

$\rightsquigarrow$ Would welfare conclusions derived under these functional forms continue to hold if:

- $A(q)$ had **non-constant gradient** in $B(p)$?
- $A(q)$ had **non-zero curvature** in $B(p)$?

Range of gradients along the demand curve

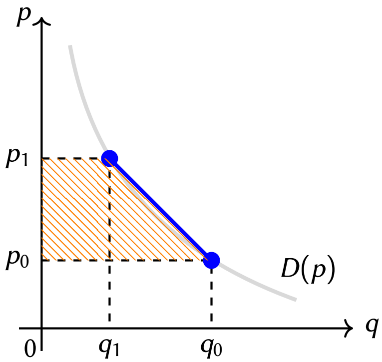

Under the assumption of linear demand, suppose

$$\Delta \mathsf{CS}_{\mathsf{linear}} - W < 0.$$

This assumes $D'(p) = \text{constant} = -\beta_{\text{avg}}$ for all p.

Range of gradients along the demand curve

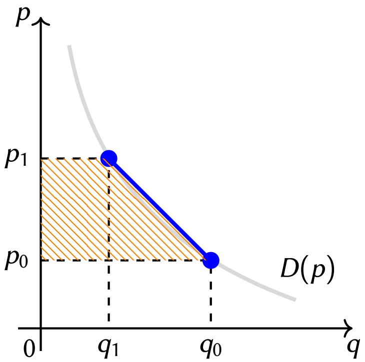

Under the assumption of linear demand, suppose

$$\Delta \mathrm{CS}_{\text{linear}} - W < 0.$$

This assumes $D'(p) = \text{constant} = -\beta_{\text{avg}}$ for all p.

What is the smallest r s.t.

$$D'(p) \in [-\beta_{\text{avg}}/(1-r), -\beta_{\text{avg}}(1-r)], \quad r \in [0, 1],$$

but the curve $D(p)$ flips the conclusion:

$$\Delta \mathrm{CS} - W \geq 0?$$

Range of gradients along the demand curve

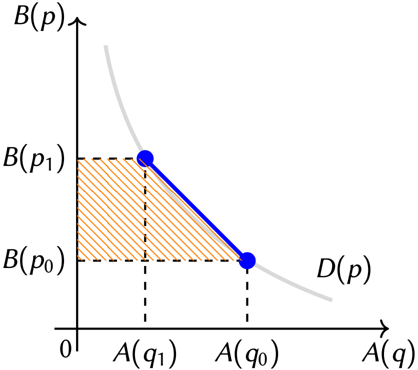

Under the assumption that $A(q)$ is affine in $B(p)$, suppose

$$\Delta \text{CS} - W < 0.$$

This assumes that the gradient of A vs. B is constant.

What is the smallest r s.t. the gradient of A vs. B is in

$$[-\beta_{\text{avg}}/\,(1-r)\,, -\beta_{\text{avg}}\,(1-r)], \quad r \in [0, 1],$$

but the curve $D(p)$ flips the conclusion:

$$\Delta \text{CS} - W \geq 0?$$

Robustness in Gradients

Welfare bounds for robustness in gradients

Suppose that the graph of A v.s. B has a gradient bounded between $\underline{\beta}$ and $\overline{\beta}$, *i.e.*,

$$\frac{A'(D(p))D'(p)}{B'(p)} \in [\underline{\beta}, \overline{\beta}] \quad \text{for } p \in [p_0, p_1].$$

What does this imply about the largest and smallest possible values of ΔCS?

Welfare bounds for robustness in gradients

Suppose that the graph of A v.s. B has a gradient bounded between $\underline{\beta}$ and $\overline{\beta}$, *i.e.*,

$$\frac{A'(D(p))D'(p)}{B'(p)} \in [\underline{\beta}, \overline{\beta}] \quad \text{for } p \in [p_0, p_1].$$

What does this imply about the largest and smallest possible values of ΔCS?

Theorem (welfare bounds for gradients).

Under the above assumption, the largest and smallest possible values of the change in consumer surplus ΔCS are attained by **2-piece A-B-linear interpolations**.

Defining 1-piece and 2-piece interpolations

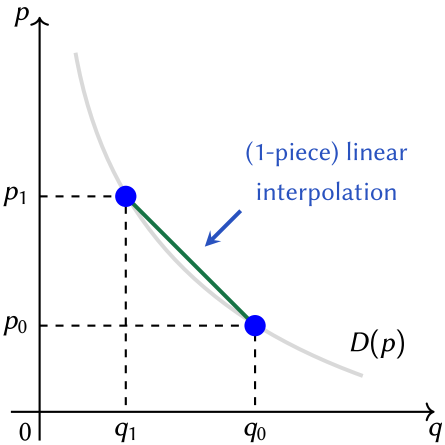

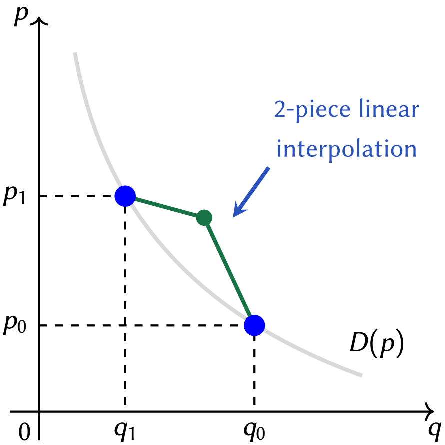

Defining 1-piece and 2-piece interpolations

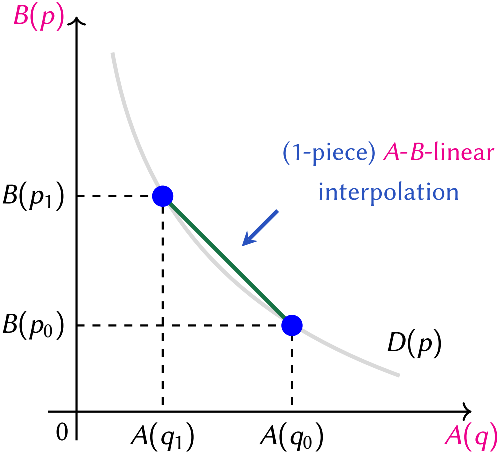

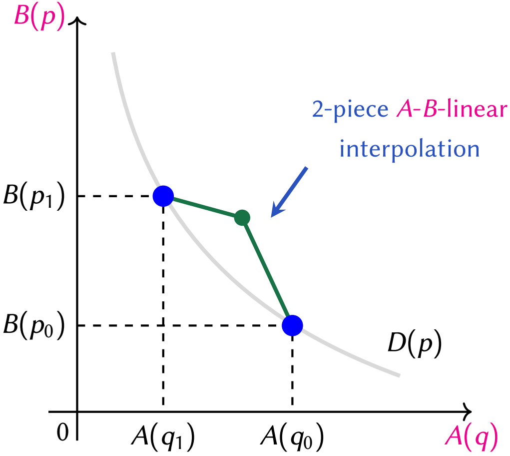

Theorem (welfare bounds for gradients).

Under the above assumption, the largest and smallest possible values of the change in consumer surplus ΔCS are attained by **2-piece *A*-*B*-linear interpolations**.

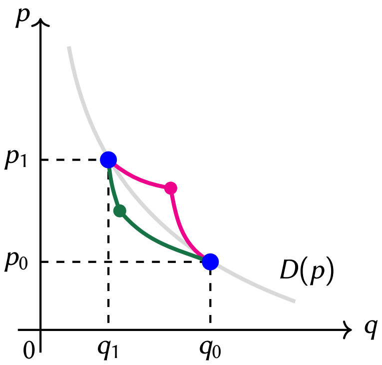

Theorem (welfare bounds for gradients).

Under the above assumption, the largest and smallest possible values of the change in consumer surplus ΔCS are attained by **2-piece *A*-*B*-linear interpolations**.

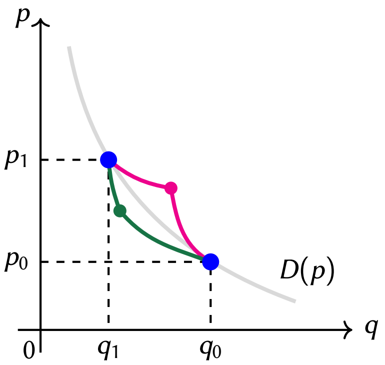

Geometric derivation of welfare bounds ◂ Back

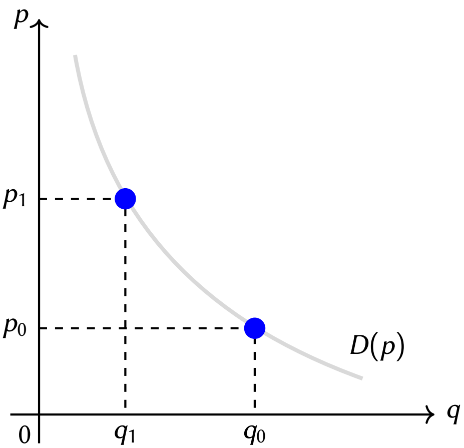

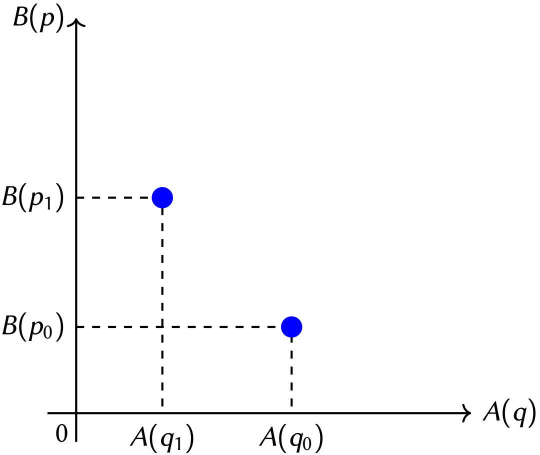

Geometric derivation of welfare bounds

Back

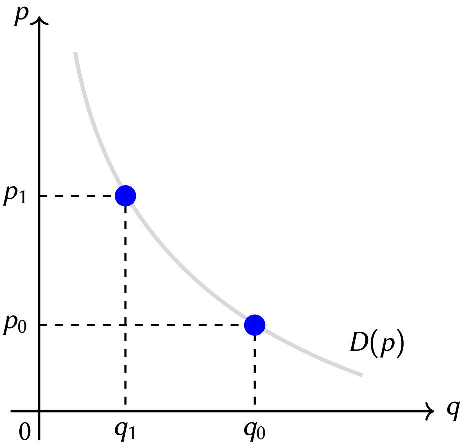

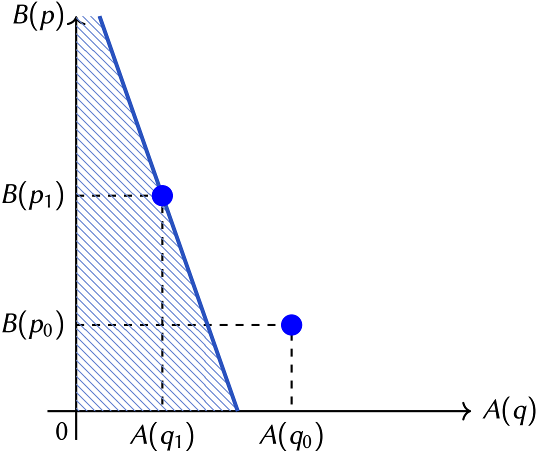

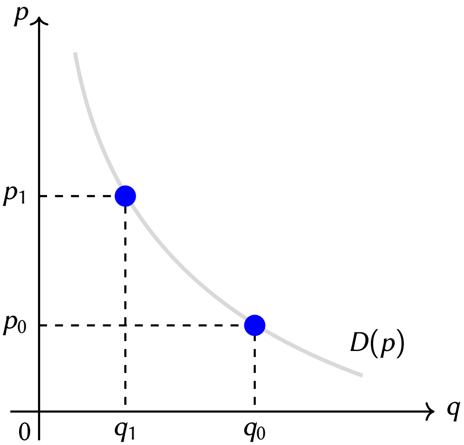

p
p_1
p_0
$D(p)$
q
0
q_1
q_0

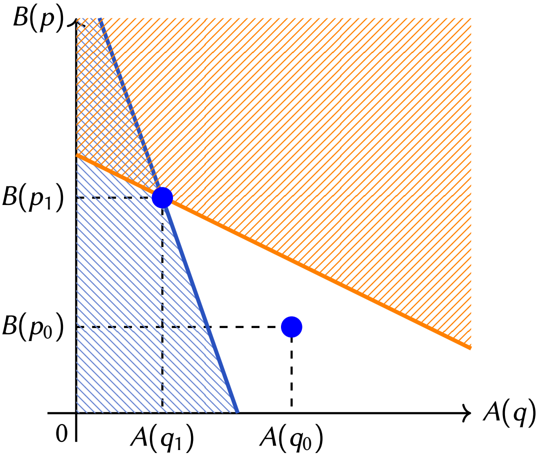

$B(p)$
$B(p_1)$
$B(p_0)$
$A(q)$
0
$A(q_1)$
$A(q_0)$

Geometric derivation of welfare bounds

Back

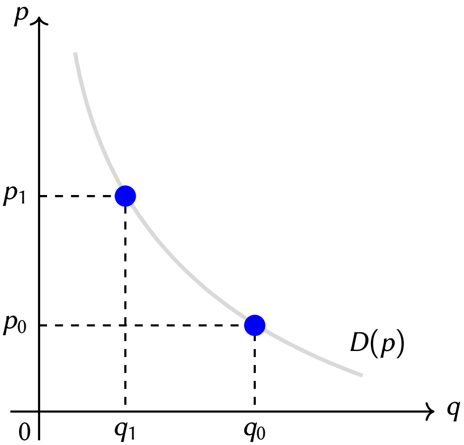

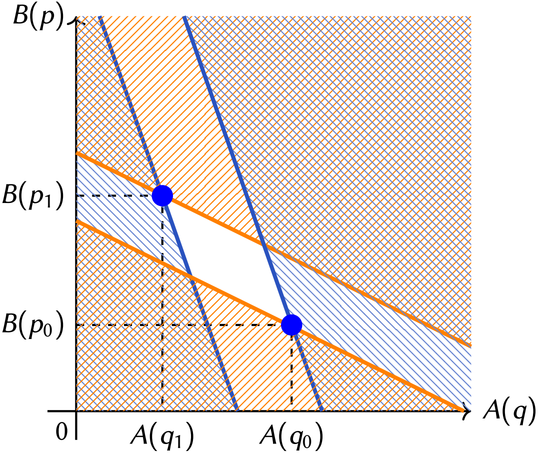

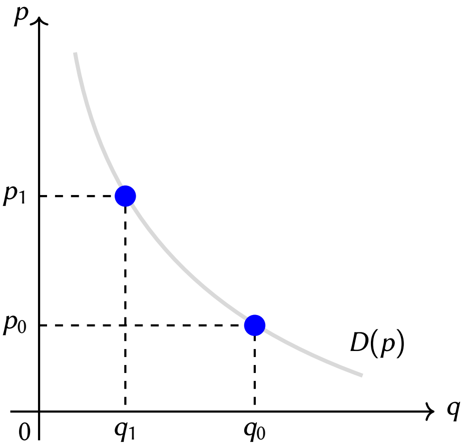

p
p_1
p_0
$D(p)$
q
0
q_1
q_0

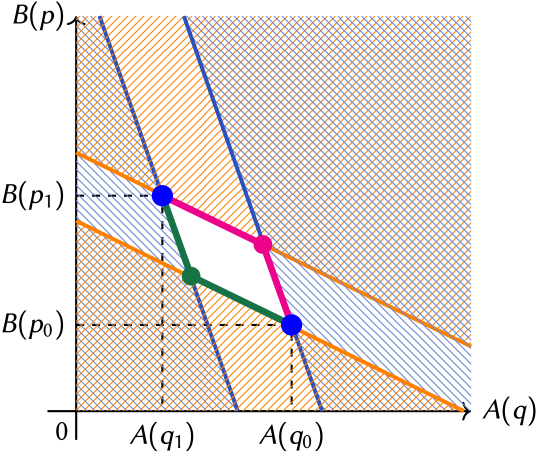

$B(p)$
$B(p_1)$
$B(p_0)$
$A(q)$
0
$A(q_1)$
$A(q_0)$

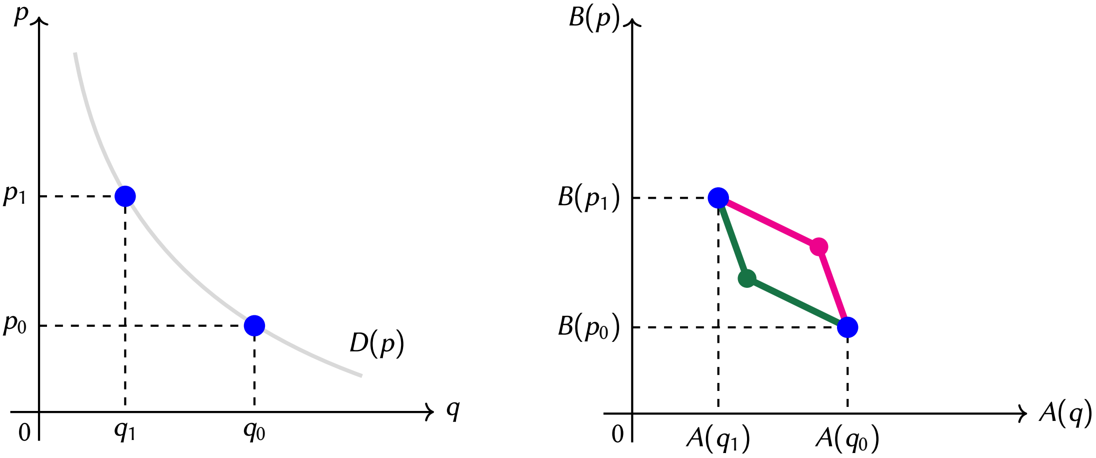
p
p1
p0
0
q1
q0
q
D(p)
B(p)
B(p1)
B(p0)
0
A(q1)
A(q0)
A(q)

Geometric derivation of welfare bounds

Back

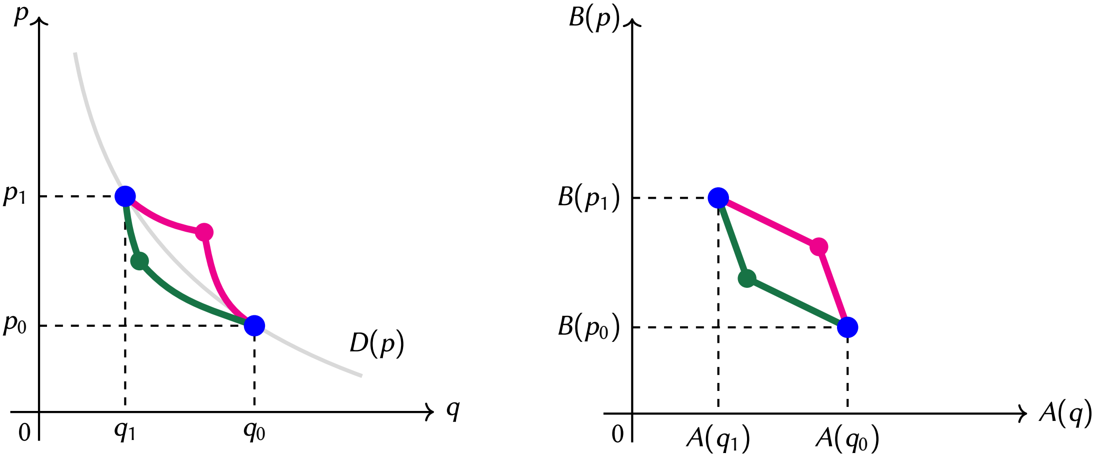

What if we have more price points?

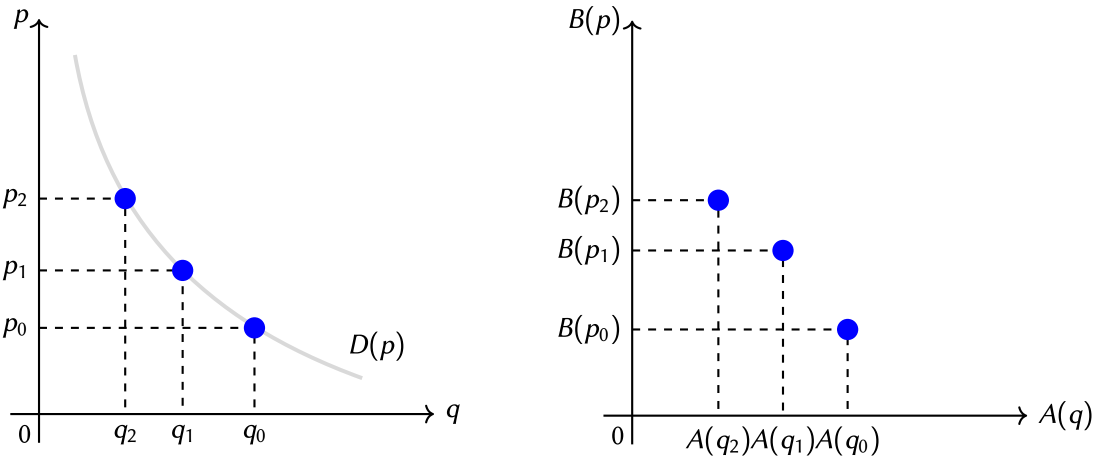

What if we have more price points?

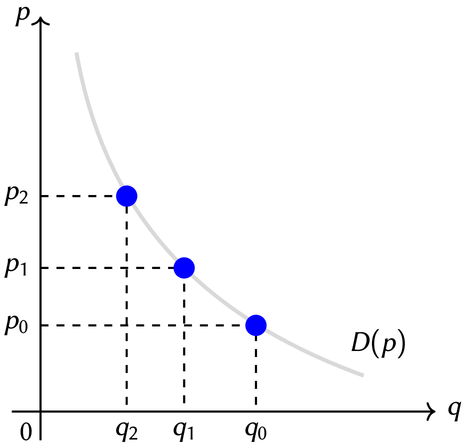

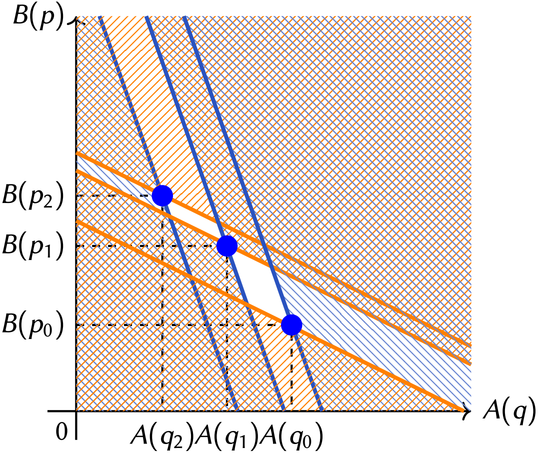

What if we have more price points?

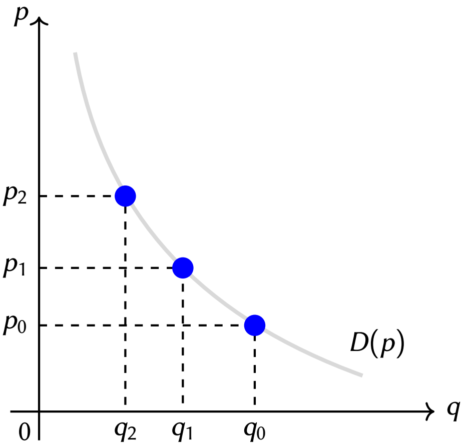

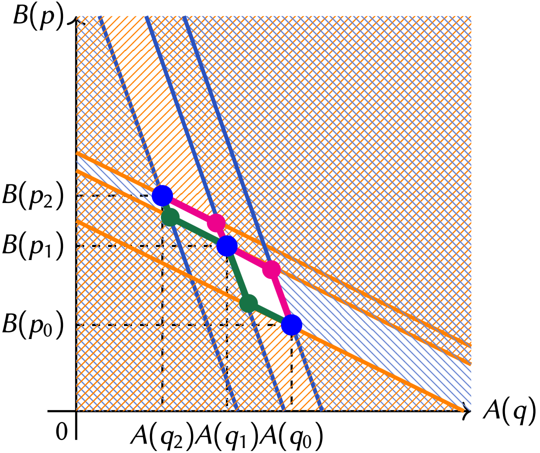

What if we have more price points?

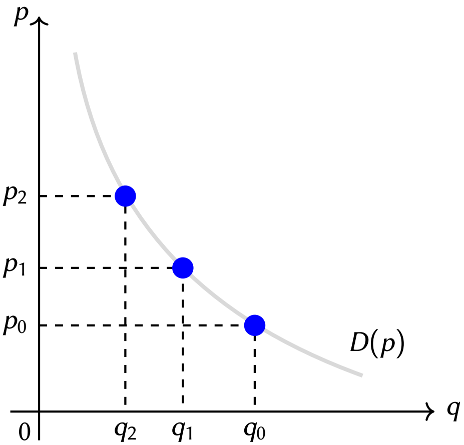

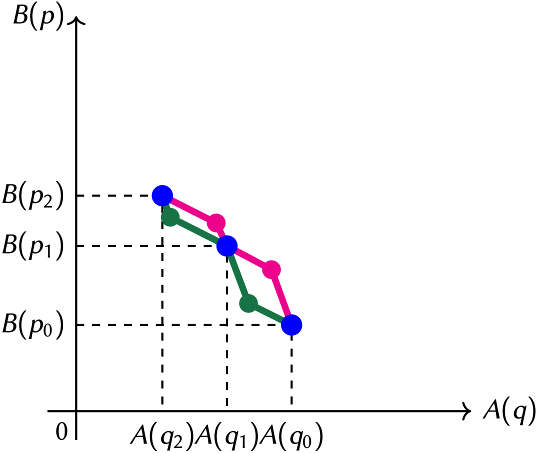

What if we have more price points?

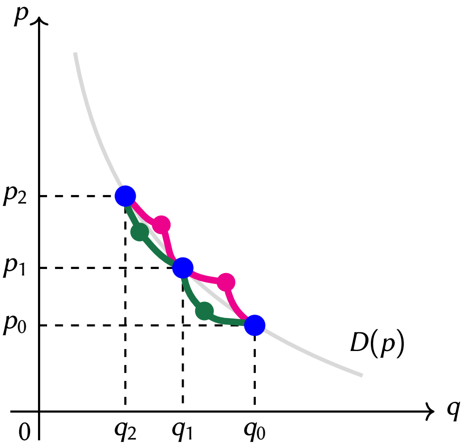

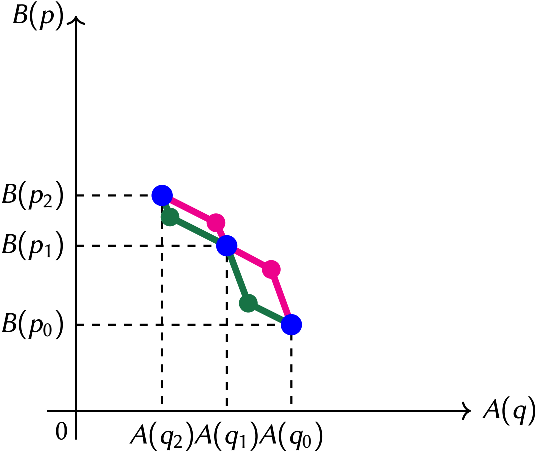

How does this work in practice?

- Suppose we observe price + quantity data for a good in a few markets at $t = \{0, 1\}$
- For now: suppose there was an exogenous price shock at $t = 1$
  - *e.g.* an import tariff (w/ pass through 1)
  - *e.g.* a local subsidy/discount in an experiment or promotion

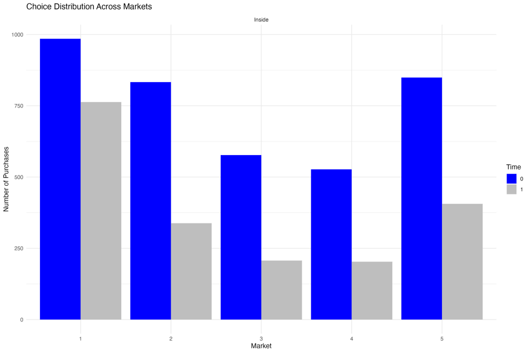

How does this work in practice?

- We observe different pre/post price points in each market...
- But the markets are also different...

- We observe different pre/post price points in each market...
- But the markets are also different...

⇝ In this example: RCL logit with market FEs

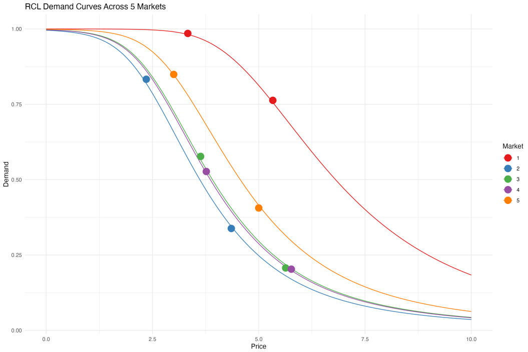

How does this work in practice?

- We don’t have really have enough data for BLP

$\Rightarrow$ What do we do?

Common Approach: Linear Interpolation

- Our goal is to estimate the difference in consumer surplus between the two groups.

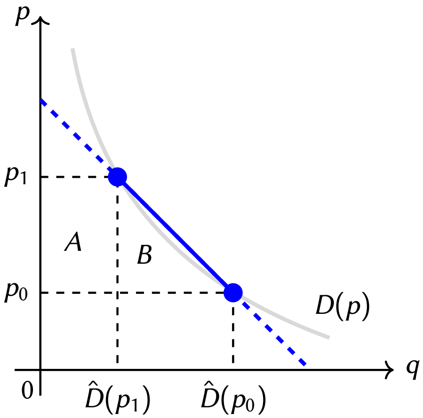

- Estimate regression:

$$y_{it} = \theta_1 - \theta_2 p_t + \epsilon_{it}.$$

- Integrate under $\hat{D}(p) = \hat{\theta}_1 - \hat{\theta}_2 p$ (w.r.t. p):

$$\widehat{\Delta CS}_{\text{linear}} = \frac{1}{2}(p_1 - p_0)\left[\hat{D}(p_1) + \hat{D}(p_0)\right].$$

How does this work in practice?

- A common approach: (diff in diff) linear regression:

$$q_{mt} = \alpha p_{mt} + \text{FE}_m + \nu_{mt} \tag{1}$$

$\Rightarrow$ interpretation: α is the average treatment effect of Δp

$\Rightarrow$ interpretation: α is the average *gradient* of the demand curve(s)

How does this work in practice?

- A common approach: (diff in diff) linear regression:

$$q_{mt} = \alpha p_{mt} + \mathsf{FE}_m + \nu_{mt} \tag{1}$$

$\Rightarrow$ interpretation: α is the average treatment effect of Δp

$\Rightarrow$ interpretation: α is the average *gradient* of the demand curve(s)

$\leadsto$ we can assume demand is linear/isoelastic/etc., and extrapolate

How does this work in practice?

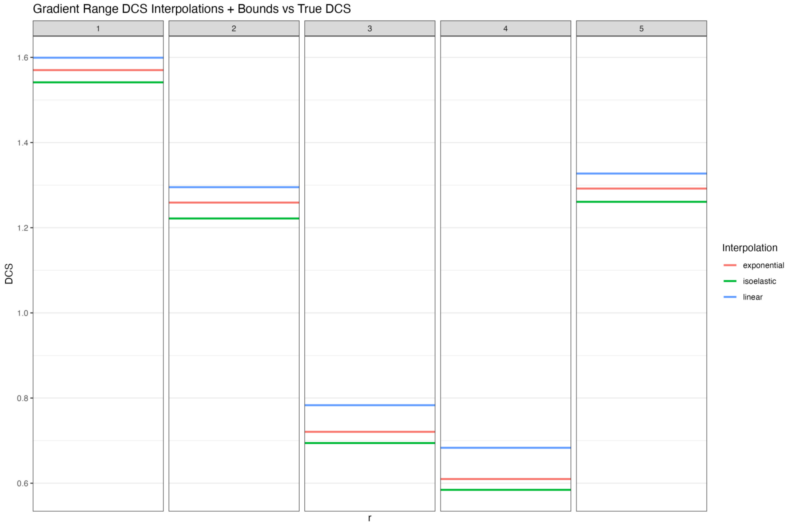

How does this work in practice?

- A common approach: (diff in diff) linear regression:

$$q_{mt} = \alpha p_{mt} + \mathrm{FE}_m + \nu_{mt} \tag{2}$$

$\Rightarrow$ interpretation: α is the average *gradient* of the demand curve(s)

$\Rightarrow$ we can assume demand is linear/isoelastic/etc., and extrapolate

$\rightsquigarrow$ Is this a good approximation?

How does this work in practice?

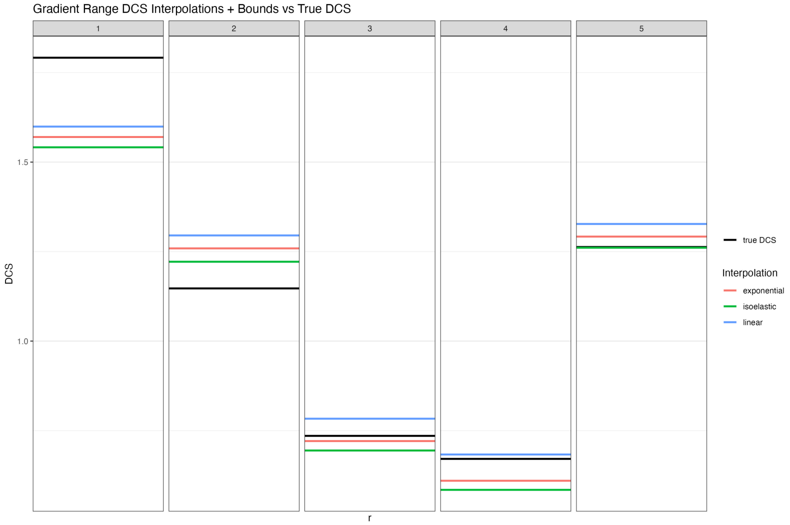

How does this work in practice?

- A common approach: (diff in diff) linear regression:

$$q_{mt} = \alpha p_{mt} + \mathsf{FE}_m + \nu_{mt} \tag{3}$$

⇒ interpretation: α is the average *gradient* of the demand curve(s)

⇒ we can assume demand is linear/isoelastic/etc., and extrapolate

⇒ Is this a good approximation?

- ⇝ In practice, we can't know the truth
- ⇝ But we can construct bounds to see how far off we might be

How does this work in practice?

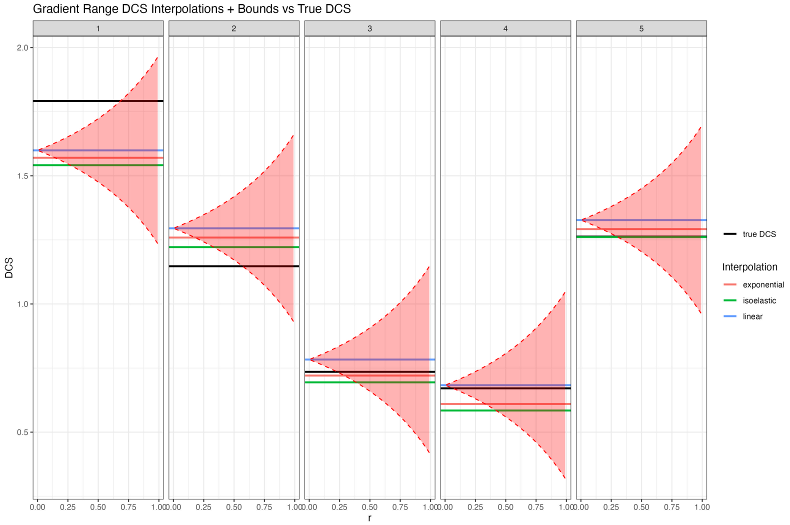

How does this work in practice?

For each market...

- Take p_0, p_1, q_0 and impute $q_1 = q_0 + \hat{\alpha}\Delta p$
- For $r \in [0, 1]$, compute bounds on ΔCS w/ Theorem 1

How does this work in practice?

For each market…

- Take p_0, p_1, q_0 and impute $q_1 = q_0 + \hat{\alpha}\Delta p$
- For $r \in [0, 1]$, compute bounds on ΔCS w/ Theorem 1

Suppose the price shock had a positive welfare externality $\overline{W}$

- **Policy question:** Is the externality benefit $\overline{W}$ bigger than the cost ΔCS?

How does this work in practice?

For each market...

- Take p_0, p_1, q_0 and impute $q_1 = q_0 + \hat{\alpha}\Delta p$
- For $r \in [0, 1]$, compute bounds on ΔCS w/ Theorem 1

Suppose the price shock had a positive welfare externality $\overline{W}$

- **Policy question:** Is the externality benefit $\overline{W}$ bigger than the cost ΔCS?
- **Robustness question:**
  What is the minimum gradient range s.t. ΔCS is guaranteed to be below $\overline{W}$?

How does this work in practice?

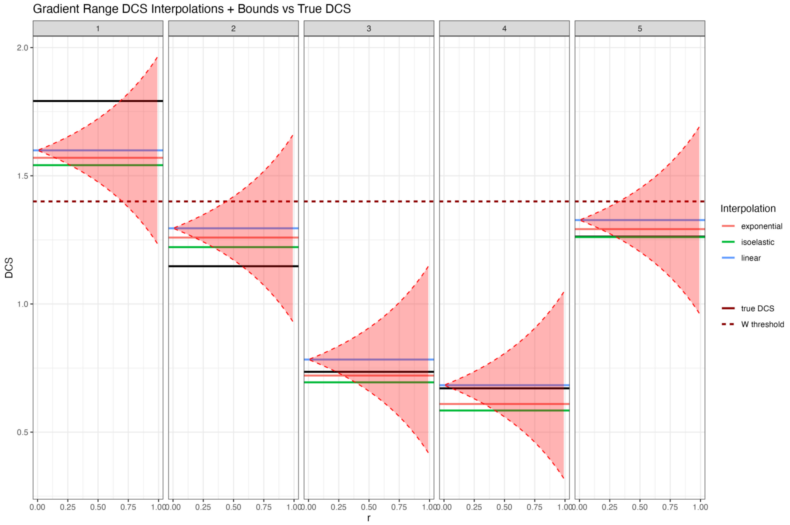

How does this work in practice?

For each market...

- Take p_0, p_1, q_0 and impute $q_1 = q_0 + \hat{\alpha}\Delta p$
- For $r \in [0, 1]$, compute bounds on ΔCS w/ Theorem 1

Suppose the price shock had a positive welfare externality $\overline{W}$

- Policy question: Is the externality benefit $\overline{W}$ bigger than the cost ΔCS?
- Robustness question:
  What is the minimum gradient range s.t. ΔCS is guaranteed to be below $\overline{W}$?

Note: Only the upper bound on the magnitude of ΔCS matters for this question

How does this work in practice?

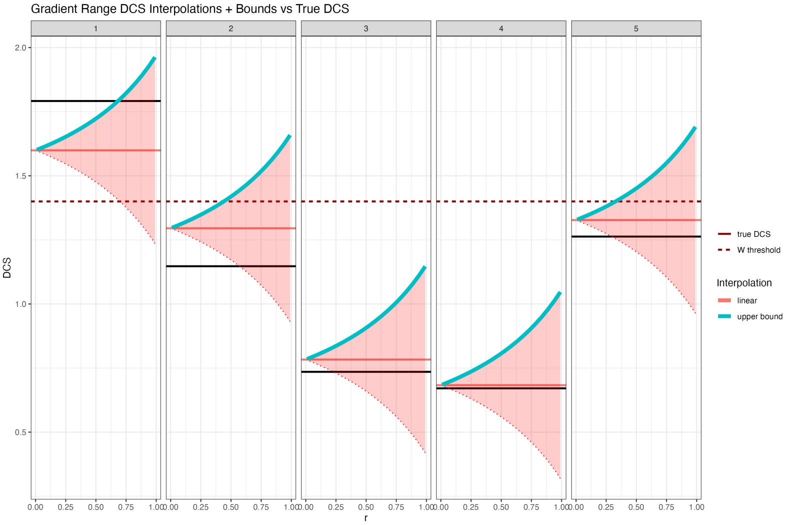

How does this work in practice?

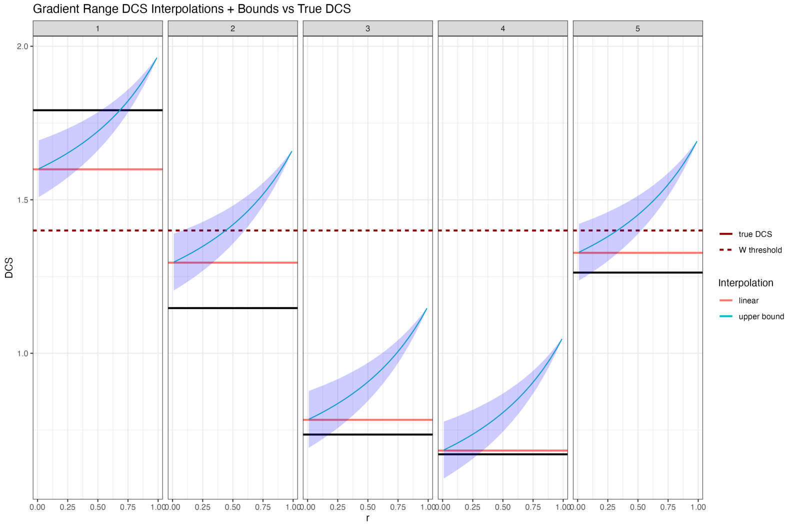

Where did that confidence band come from?

- The projection of q_1 has uncertainty

$$\text{SE}(\hat{q}_1) = \text{SE}(\hat{\alpha}) \times |\Delta p|$$

- $\Delta\text{CS}(\hat{q}_1, r)$ is continuous function of $\hat{q}_1$

⇝ Delta Method → standard errors on $\Delta\text{CS}(\hat{q}_1, r)$

$$\text{SE}(\Delta\text{CS}(\hat{q}_1, r)) = \left| \frac{\partial\, \Delta\text{CS}(\hat{q}_1, r)}{\partial q_1} \right| \times \text{SE}(\hat{q}_1)$$

⇝ Or (Bayesian) bootstrap the whole thing

- $\Rightarrow$ What if I want to use log units in the regression?
  - $\rightsquigarrow$ Elasticity range bounds (on the log-log ATE)

How does this work in practice?

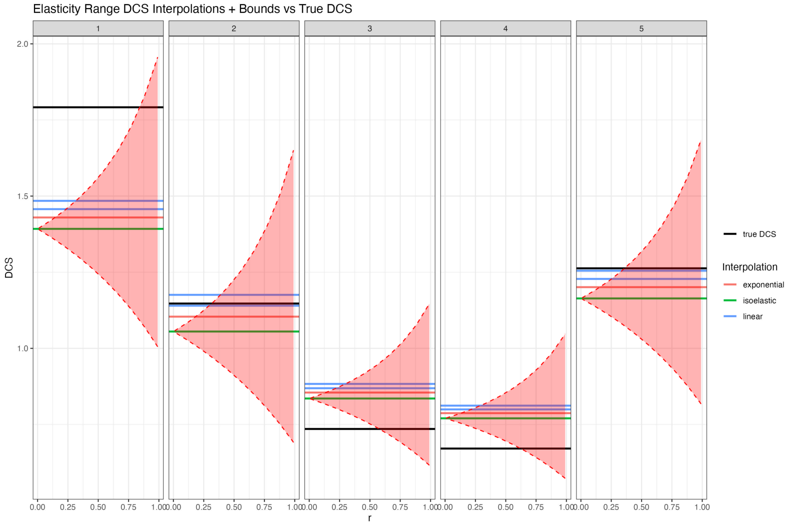

How does this work in practice?

- What if I want to use log units in the regression?
  - Elasticity range bounds (on the log-log ATE)

$\Rightarrow$ What if I don't have an exogenous price shock?

How does this work in practice?

- A common approach: IV regression

$$\mathbb{1}(\text{purchase})_{imt} = \alpha p_{imt} + \text{FE}_m + \nu_{imt} \tag{4}$$

$$p_{imt} = p_{m0} + Z_{imt}\Delta p + \epsilon_{imt} \tag{5}$$

- interpretation: α is the local average treatment effect of Δp (under IV monotonicity)
- interpretation: α is the average *gradient* of the demand curve(s)

How does this work in practice?

- A common approach: IV regression

$$\mathbb{1}(\text{purchase})_{imt} = \alpha p_{imt} + \text{FE}_m + \nu_{imt} \tag{4}$$

$$p_{imt} = p_{m0} + Z_{imt}\Delta p + \epsilon_{imt} \tag{5}$$

- interpretation: α is the local average treatment effect of Δp (under IV monotonicity)
- interpretation: α is the average *gradient* of the demand curve(s)

$\Rightarrow$ The rest goes the same as before

How does this work in practice?

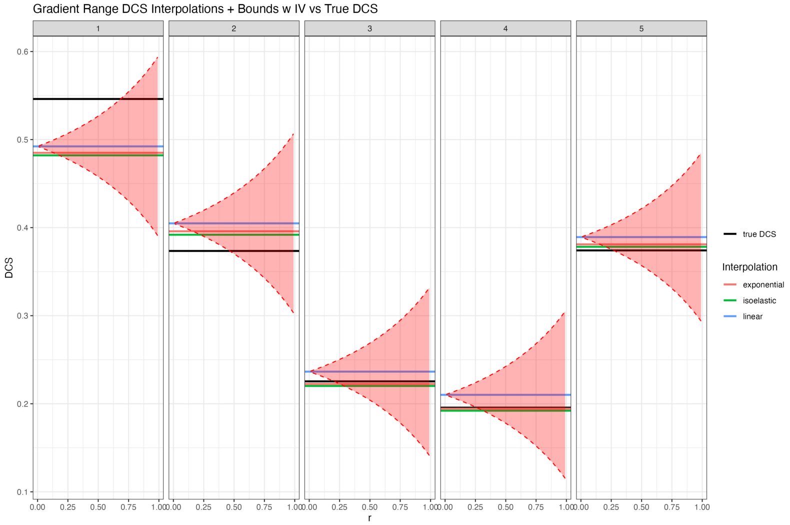

How does this work in practice?

- What if I want to use log units in the regression?
  - Elasticity range bounds (on the log-log ATE)
- What if I don't have an exogenous price shock?

$\Rightarrow$ What about second derivatives?

Robustness in Curvature

Welfare bounds for robustness in curvature

Suppose that the graph of A v.s. B has a second derivative bounded between $\underline{\gamma}$ and $\overline{\gamma}$:

$$\frac{1}{B'(p)} \frac{\mathrm{d}}{\mathrm{d}p} \left[\frac{A'(D(p))D'(p)}{B'(p)} \right] \in [\underline{\gamma}, \overline{\gamma}] \quad \text{for } p \in [p_0, p_1].$$

$$\text{where } -\infty < \underline{\gamma} \leq 0 \leq \overline{\gamma} < +\infty.$$

What does this imply about the largest and smallest possible values of $\Delta \mathrm{CS}$?

Welfare bounds for robustness in curvature

Suppose that the graph of A v.s. B has a second derivative bounded between $\underline{\gamma}$ and $\overline{\gamma}$:

$$\frac{1}{B'(p)} \frac{\mathrm{d}}{\mathrm{d}p} \left[\frac{A'(D(p))D'(p)}{B'(p)} \right] \in [\underline{\gamma}, \overline{\gamma}] \quad \text{for } p \in [p_0, p_1].$$

$$\text{where } -\infty < \underline{\gamma} \leq 0 \leq \overline{\gamma} < +\infty.$$

What does this imply about the largest and smallest possible values of ΔCS?

Theorem (welfare bounds for curvature).

Under the above assumption, the largest and smallest possible values of the change in consumer surplus ΔCS are attained by demand curves whose **gradients**, in units of $A(q)/B(p)$, are either **1-piece** or **2-piece linear interpolations**.

Explicit characterization of welfare bounds for curvature

Define the gradients (in units of $A(q)/B(p)$), $h^*, h_* : [B(p_0), B(p_1)] \to \mathbb{R}$, as follows:

$$h^*(s) = \begin{cases} -\dfrac{A(q_0)-A(q_1)}{B(p_1)-B(p_0)} - \dfrac{\underline{\gamma}}{2}\left[B(p_0)+B(p_1)\right] & \text{if } \underline{\gamma} \geq -\dfrac{2\left[A(q_0)-A(q_1)\right]}{\left[B(p_1)-B(p_0)\right]^2}, \\ \begin{cases} -\underline{\gamma}\left[B(p_1) - \sqrt{\frac{2[A(q_1)-A(q_0)]}{\underline{\gamma}}}\right] & \text{if } s > B(p_1) - \sqrt{\frac{2[A(q_1)-A(q_0)]}{\underline{\gamma}}}, \\ -\underline{\gamma}s & \text{if } s \leq B(p_1) - \sqrt{\frac{2[A(q_1)-A(q_0)]}{\underline{\gamma}}}, \end{cases} & \text{if } \underline{\gamma} < -\dfrac{2\left[A(q_0)-A(q_1)\right]}{\left[B(p_1)-B(p_0)\right]^2}; \end{cases}$$

$$h_*(s) = \begin{cases} \begin{cases} -\overline{\gamma}s & \text{if } s > B(p_0) + \sqrt{\frac{2[A(q_0)-A(q_1)]}{\overline{\gamma}}}, \\ -\overline{\gamma}\left[B(p_0) + \sqrt{\frac{2[A(q_0)-A(q_1)]}{\underline{\gamma}}}\right] & \text{if } s \leq B(p_0) + \sqrt{\frac{2[A(q_0)-A(q_1)]}{\overline{\gamma}}}, \end{cases} & \text{if } \overline{\gamma} \geq \dfrac{2\left[A(q_0)-A(q_1)\right]}{\left[B(p_1)-B(p_0)\right]^2}, \\ -\dfrac{A(q_0)-A(q_1)}{B(p_1)-B(p_0)} - \dfrac{\overline{\gamma}}{2}\left[B(p_0)+B(p_1)\right] & \text{if } \overline{\gamma} < \dfrac{2\left[A(q_0)-A(q_1)\right]}{\left[B(p_1)-B(p_0)\right]^2}. \end{cases}$$

Then:

$$\begin{cases} \overline{\Delta \mathrm{CS}} = \displaystyle\int_{p_0}^{p_1} A^{-1}\left(A(q_0) + \int_{B(p_0)}^{B(p)} \left[h^*(s) + \underline{\gamma}s\right]\, \mathrm{d}s\right)\, \mathrm{d}p, \\ \underline{\Delta \mathrm{CS}} = \displaystyle\int_{p_0}^{p_1} A^{-1}\left(A(q_0) + \int_{B(p_0)}^{B(p)} \left[h_*(s) + \overline{\gamma}s\right]\, \mathrm{d}s\right)\, \mathrm{d}p. \end{cases}$$

Curvature bounds in our simulated example

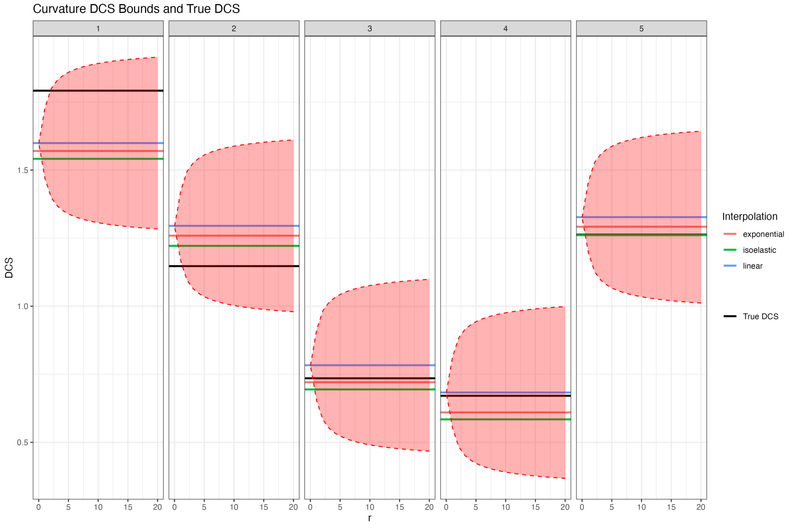

This talk: a special case

- **Special case:** parameterize curvature by **ρ-concavity** and **ρ-convexity**.
  - Equivalent to setting $A(q) = q^\rho/\rho$ and $B(p) = p$ in our framework:
$$D(p) \text{ is } \rho\text{-concave/convex} \iff \frac{q^\rho}{\rho} \text{ is concave/convex in } p.$$
  - Introduced in the economics literature by Caplin and Nalebuff (1991a,b).
  - $\rho \in \mathbb{R}$ parametrizes how "concave" or "convex" a function is.
  - Examples: $\rho = 0$ (log-concavity/convexity); $\rho = 1$ (concavity/convexity).

How robust are welfare conclusions to curvature?

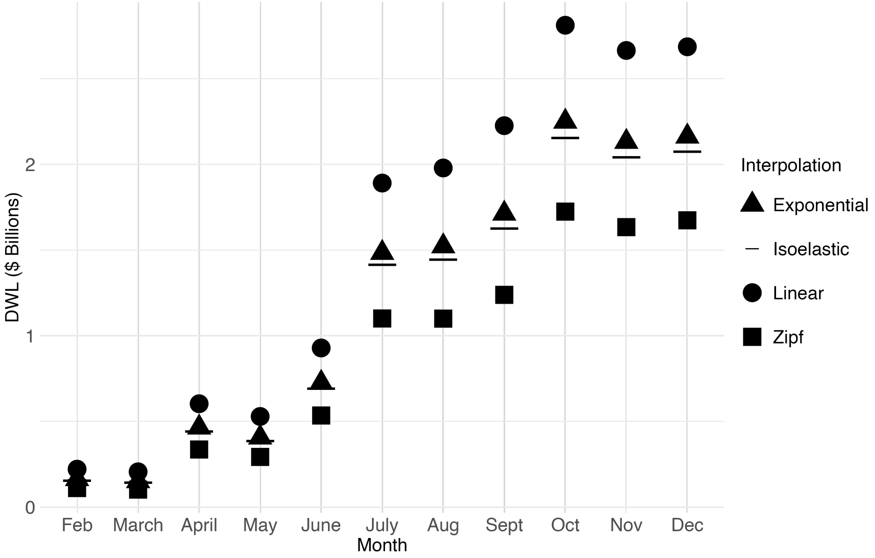

How robust are welfare conclusions to curvature?

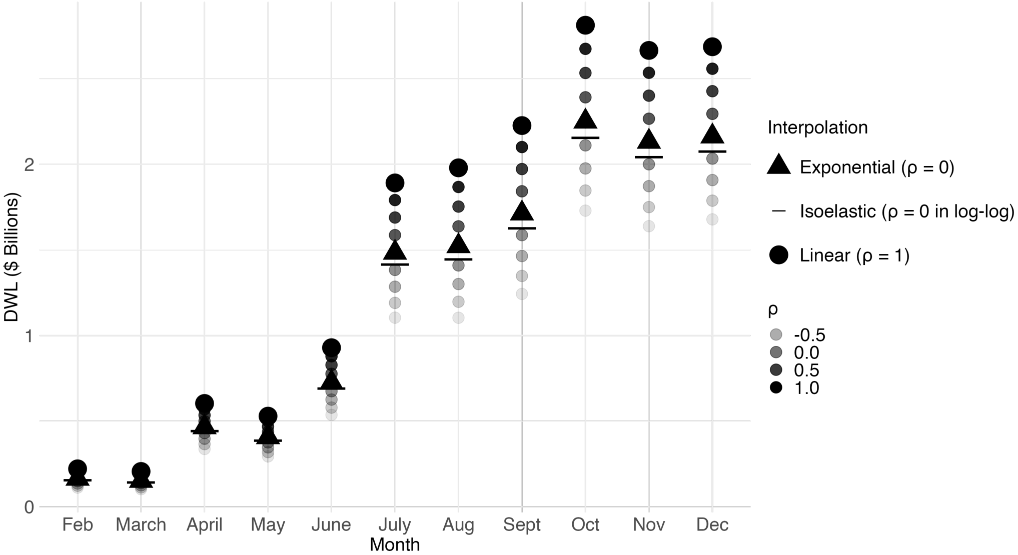

How robust are welfare conclusions to curvature?

- **Parameterize curvature w/ ρ-concavity/convexity** (Caplin and Nalebuff, 1991b)
  - The more *convex* $D(p)$ is, the *smaller* ΔCS is
  - The more *concave* $D(p)$ is, the *larger* ΔCS is
  - We parametrize "more" with ρ

How robust are welfare conclusions to curvature?

- **Parameterize curvature w/ ρ-concavity/convexity** (Caplin and Nalebuff, 1991b)
  - The more *convex* $D(p)$ is, the *smaller* ΔCS is
  - The more *concave* $D(p)$ is, the *larger* ΔCS is
  - We parametrize "more" with ρ
- How concave can $D(p)$ be to flip the conclusion $\Delta CS_{linear} - W > 0$?
  - Given ρ, characterize the lower bound on ΔCS

  $\Rightarrow$ The lower bound is attained by a ρ-linear curve

  $\Rightarrow$ Find smallest ρ s.t. $\Delta CS_{\rho} - W \leq 0$

Welfare bounds implied by ρ-curvature of demand in price

Theorem (welfare bounds for ρ-convex demand).

If demand is ρ-convex in price, the lower bound is given by a 2-piece ρ-linear interpolation and the upper bound is given by a 1-piece ρ-linear interpolation.

Theorem (welfare bounds for ρ-concave demand).

If demand is ρ-concave in price, the lower bound is given by a 1-piece ρ-linear interpolation and the upper bound is given by a 2-piece ρ-linear interpolation.

Special cases:

- $\rho = 0$: **exponential** interpolation is extremal for log-convex and log-concave demand.
- $\rho = 1$: **linear** interpolation is extremal for convex and concave demand.

Bounding welfare with demand curvature: ρ-concave demand

Theorem (welfare bounds for ρ-concave demand).

If demand is ρ-concave in price, the lower bound is given by a 1-piece ρ-linear interpolation and the upper bound is given by a 2-piece ρ-linear interpolation.

Recall:

- $D(p)$ is **ρ-concave** if $D'(p)\,[D(p)]^{\rho-1}$ is decreasing in p.
- $D(p)$ is **ρ-linear** if $D(p) = \left[q_0^{\rho} - \beta\,(p - p_0)\right]^{1/\rho}$ for some $\beta \geq 0$.

▸ Skip Proof

Step #1: change of variables

Variable change:

$$h(p) = -D'(p)\,[D(p)]^{\rho-1} \implies [D(p)]^{\rho} = q_0^{\rho} - \rho \int_{p_0}^{p} h(s)\,\mathrm{d}s.$$

Step #1: change of variables

Variable change:

$$h(p) = -D'(p)\,[D(p)]^{\rho-1} \implies [D(p)]^{\rho} = q_0^{\rho} - \rho \int_{p_0}^{p} h(s)\,\mathrm{d}s.$$

Constraint (on the mean of h)**:**

$$\mathcal{H} = \left\{ h \text{ is increasing s.t. } \int_{p_0}^{p_1} h(s)\,\mathrm{d}s = \frac{q_0^{\rho} - q_1^{\rho}}{\rho} \right\}.$$

Step #1: change of variables

Variable change:

$$h(p) = -D'(p)\,[D(p)]^{\rho-1} \implies [D(p)]^{\rho} = q_0^{\rho} - \rho \int_{p_0}^{p} h(s)\,\mathrm{d}s.$$

Constraint (on the mean of h)**:**

$$\mathcal{H} = \left\{ h \text{ is increasing s.t. } \int_{p_0}^{p_1} h(s)\,\mathrm{d}s = \frac{q_0^{\rho} - q_1^{\rho}}{\rho} \right\}.$$

Welfare:

$$\begin{cases} \overline{\Delta\mathrm{CS}} = \max\limits_{h \in \mathcal{H}} \displaystyle\int_{p_0}^{p_1} \left[q_0^{\rho} - \rho \int_{p_0}^{p} h(s)\,\mathrm{d}s \right]^{1/\rho} \mathrm{d}p, \\ \underline{\Delta\mathrm{CS}} = \min\limits_{h \in \mathcal{H}} \displaystyle\int_{p_0}^{p_1} \left[q_0^{\rho} - \rho \int_{p_0}^{p} h(s)\,\mathrm{d}s \right]^{1/\rho} \mathrm{d}p. \end{cases}$$

Step #2: establishing a partial order

Definition: $h_1 \succeq h_2$ if h_1 is a mean-preserving spread of h_2, *i.e.*,

$$h_1 \succeq h_2 \iff \int_{p_0}^{p} h_1(s)\,\mathrm{d}s \geq \int_{p_0}^{p} h_2(s)\,\mathrm{d}s \qquad \forall\, p \in [p_0, p_1].$$

- This defines a *partial order* on $\mathcal{H}$.
  - $\Rightarrow$ Can think of this as second-order stochastic dominance.
  - $\Rightarrow$ Because h is increasing, can think of h as a CDF (appropriately shifted and scaled).

Step #2: connecting to welfare

Lemma: The welfare objective is decreasing in the partial order $\succeq$, *i.e.*,

$$h_1 \succeq h_2 \implies \int_{p_0}^{p_1} \left[q_0^\rho - \rho \int_{p_0}^{p} h_1(s)\, \mathrm{d}s \right]^{1/\rho} \mathrm{d}p \leq \int_{p_0}^{p_1} \left[q_0^\rho - \rho \int_{p_0}^{p} h_2(s)\, \mathrm{d}s \right]^{1/\rho} \mathrm{d}p.$$

Proof: Pointwise comparison of the integrands.

Step #2: connecting to welfare

Lemma: The welfare objective is decreasing in the partial order $\succeq$, *i.e.*,

$$h_1 \succeq h_2 \implies \int_{p_0}^{p_1} \left[q_0^\rho - \rho \int_{p_0}^{p} h_1(s)\, \mathrm{d}s \right]^{1/\rho} \mathrm{d}p \leq \int_{p_0}^{p_1} \left[q_0^\rho - \rho \int_{p_0}^{p} h_2(s)\, \mathrm{d}s \right]^{1/\rho} \mathrm{d}p.$$

Proof: Pointwise comparison of the integrands.

Corollary. The lower (*resp.*, upper) bound is attained by iteratively applying mean-preserving spreads (*resp.*, mean-preserving contractions) to $h(p)$.

Step #3: deriving the *lower* bound

Consider the density that generates $h(p)$, where $h(p)$ is viewed as a CDF:

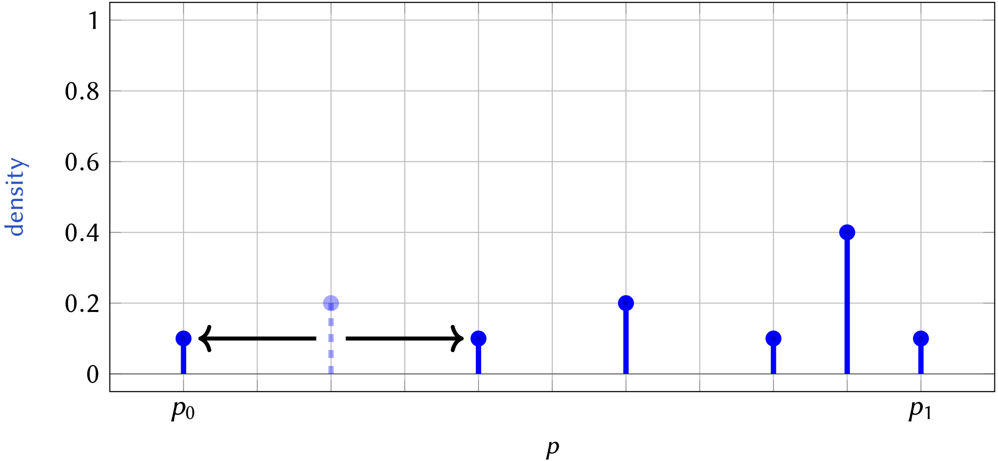

Step #3: deriving the *lower* bound

Consider the density that generates $h(p)$, where $h(p)$ is viewed as a CDF:

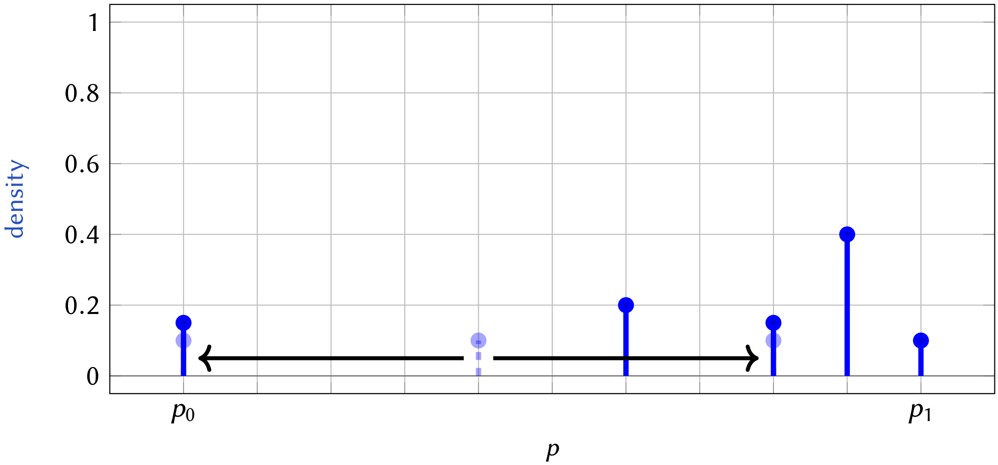

Step #3: deriving the *lower* bound

Consider the density that generates $h(p)$, where $h(p)$ is viewed as a CDF:

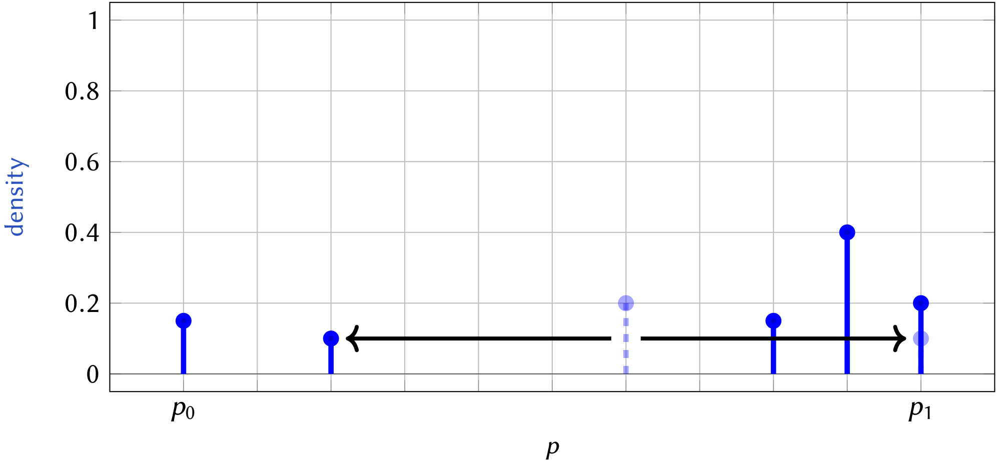

Step #3: deriving the *lower* bound

Consider the density that generates $h(p)$, where $h(p)$ is viewed as a CDF:

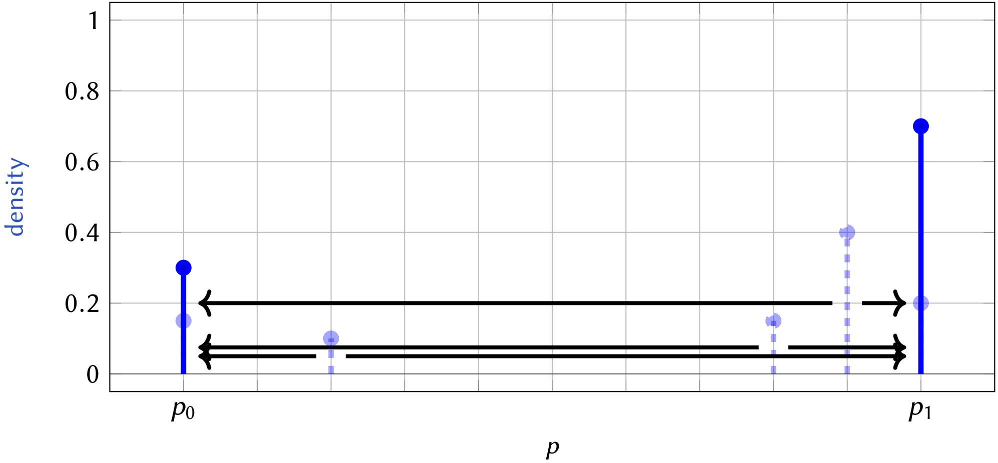

Step #3: deriving the *lower* bound

Consider the density that generates $h(p)$, where $h(p)$ is viewed as a CDF:

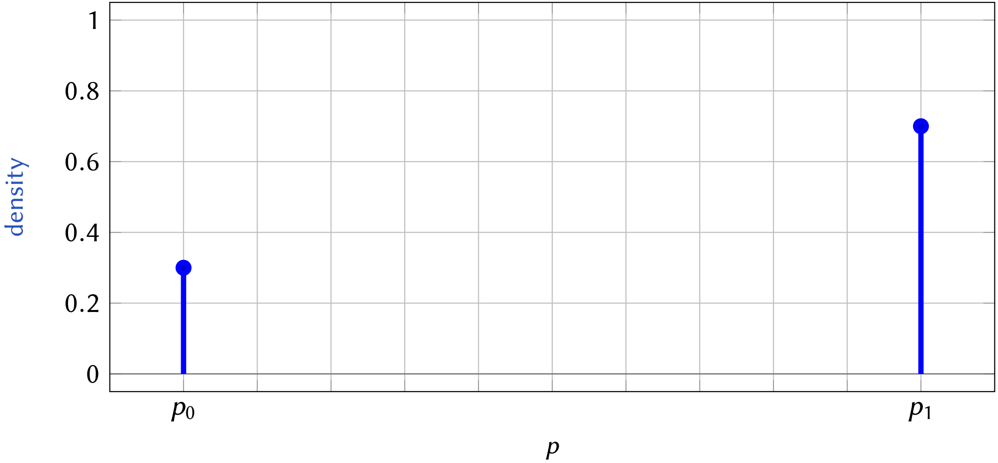

So the $h(p)$ that attains the **lower bound on welfare** is **constant** between p_0 and p_1:

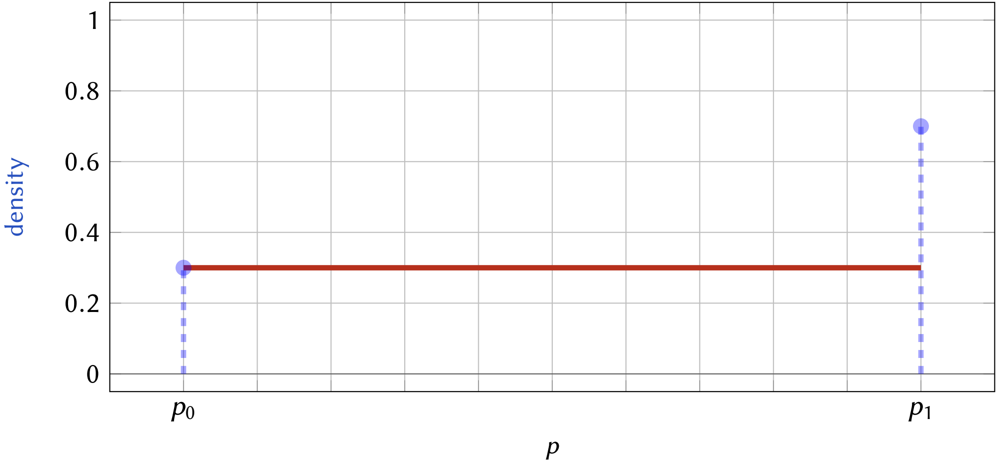

Step #3: deriving the *upper* bound

Similarly, the $h(p)$ that attains the **upper bound on welfare** is a **step function**.

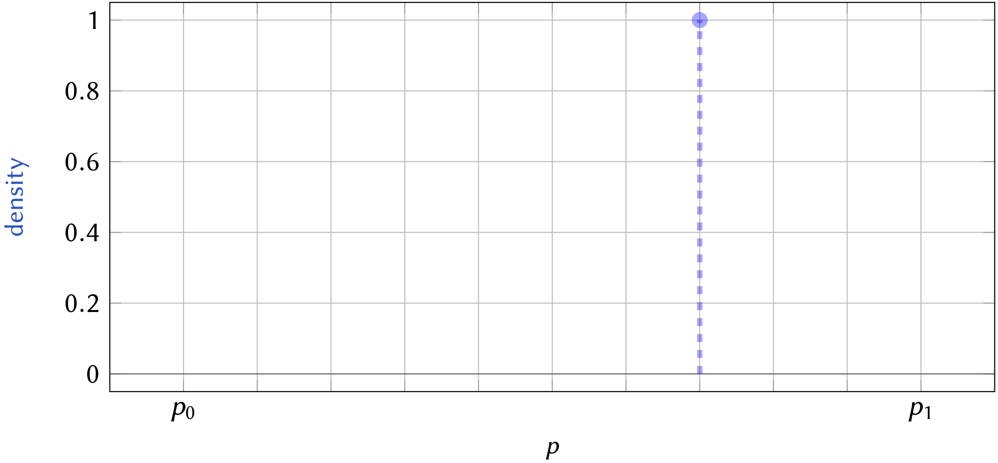

Step #3: deriving the *upper* bound

Similarly, the $h(p)$ that attains the **upper bound on welfare** is a **step function**.

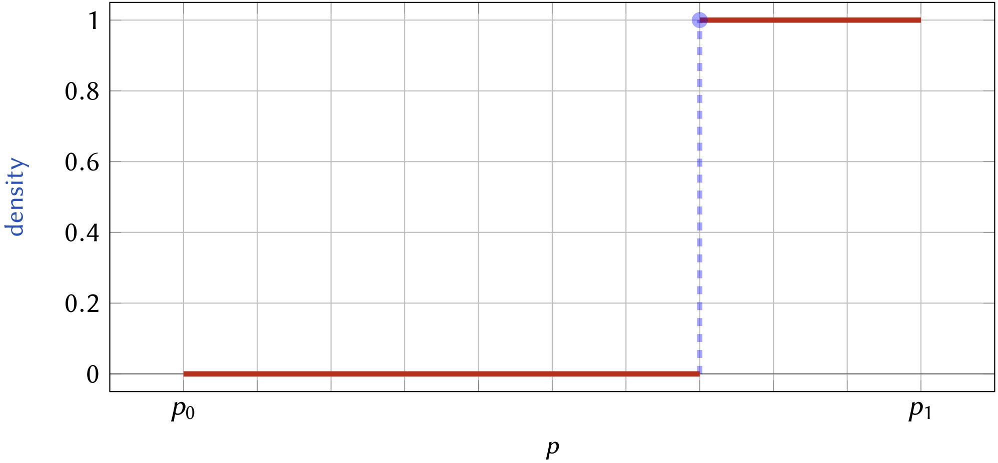

Step #3: deriving welfare bounds

- Mapping back from $h(p)$ into demand curves $D(p)$:

$$h(p) \text{ is constant in } p \iff -D'(p)\,[D(p)]^{\rho-1} \text{ is constant in } p$$
$$\iff D(p) = \left[q_0^\rho - \beta\,(p - p_0)\right]^{1/\rho}.$$

Note:

$$q_1^\rho = q_0^\rho - \beta\,(p_1 - p_0) \implies \beta = \frac{q_0^\rho - q_1^\rho}{p_1 - p_0}.$$

Step #3: deriving welfare bounds

- Mapping back from $h(p)$ into demand curves $D(p)$:

$$h(p) \text{ is constant in } p \iff -D'(p)\,[D(p)]^{\rho-1} \text{ is constant in } p$$
$$\iff D(p) = [q_0^\rho - \beta\,(p - p_0)]^{1/\rho}.$$

Note:

$$q_1^\rho = q_0^\rho - \beta\,(p_1 - p_0) \implies \beta = \frac{q_0^\rho - q_1^\rho}{p_1 - p_0}.$$

- This proves the bounds for ρ-concave demand:
  - The **lower bound** is attained by a 1-piece ρ-linear interpolation.
  - The **upper bound** is attained by a 2-piece ρ-linear interpolation.
- The same proof strategy works for ρ-convex demand.

Example: evaluating the deadweight loss of the Trump tariffs

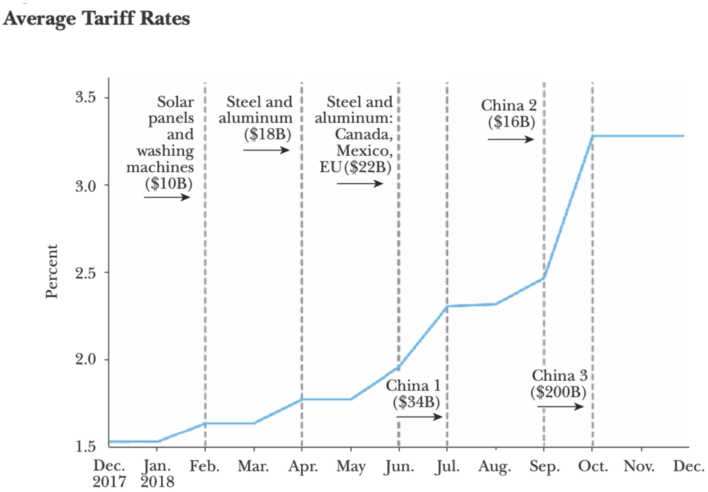

Source: Amiti, Redding and Weinstein (2019)

Example: evaluating the deadweight loss of the Trump tariffs

How Many Tariff Studies Are Enough?

The trade war hits consumers and exports, two more papers say.

By The Editorial Board
Jan. 20, 2020 4:39 pm ET

PRINT TEXT 169

Source: WSJ Editorial Board

Interpreting the tariff DWL Estimates

- **Contextualizing numbers.** The tariff revenue gained over 2018 is \$15.6 billion.
  - An isoelastic interpolation yields a DWL estimate of \$12.6 billion
  - A linear interpolation yields a DWL estimate of \$16.8 billion.
- **Positive Welfare Criterion.** Could added domestic manufacturing wages make up for the DWL?
  - Suppose the trade war recouped the 35,400 manufacturing jobs lost over the 2010s
  - $\rightsquigarrow$ \$1.86 billion/year assuming a \$52,500 average wage
  - $\rightsquigarrow$ Could this exceed the DWL?

Could the tariffs be worth it?

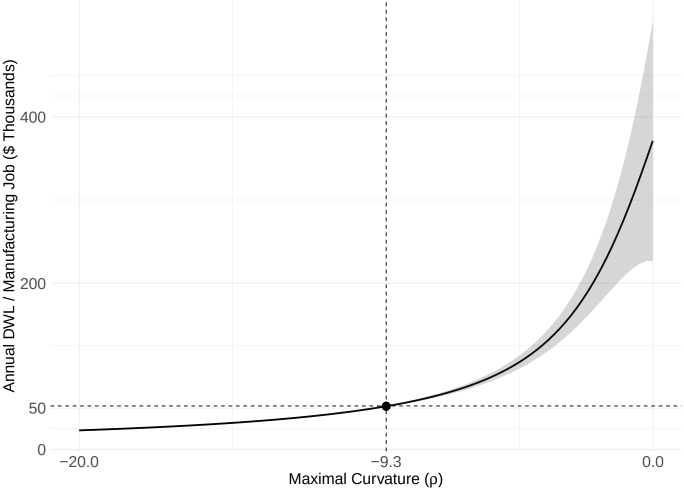

Welfare bounds implied by ρ-curvature of log-demand in log-price

Theorem (welfare bounds for ρ-convex demand).

If **log-demand** is ρ-convex in **log-price**, the lower bound is given by a 2-piece ρ-isoelastic interpolation and the upper bound is given by a 1-piece ρ-isoelastic interpolation.

Theorem (welfare bounds for ρ-concave demand).

If **log-demand** is ρ-concave in **log-price**, the lower bound is given by a 1-piece ρ-isoelastic interpolation and the upper bound is given by a 2-piece ρ-isoelastic interpolation.

Special case:

- $\rho = 1$: **isoelastic** interpolation is extremal for demand with decreasing elasticity (Marshall's second law) and demand with increasing elasticity.

Common interpolations as assumptions on demand curvature

Theorem (Bounding functions for concave-like curvatures).

The **lower** bound for the change in consumer surplus are attained by:

- **concave demand:** a *linear* interpolation; $D(p) = \theta_1 - \theta_2 p$
- **log-concave demand:** an *exponential* interpolation; $D(p) = \theta_1 e^{-\theta_2 p}$
- **decreasing MR:** a *constant MR (zipf)* interpolation; $D(p) = \theta_1 (p - \theta_2)^{-1}$
- **decreasing elasticity:** a *isoelastic* interpolation; $D(p) = \theta_1 p^{-\theta_2}$

Relationships between curvature assumptions

Concave-like assumptions

(A1) **Decreasing elasticity**

(A2) **Decreasing MR**

(A3) **Log-concave demand**

(A4) **Concave demand**

$$(\text{A4}) \Longrightarrow (\text{A3}) \Longrightarrow (\text{A1})$$
$$(\text{A3}) \Longrightarrow (\text{A2})$$

Convex-like assumptions

(A6) **Convex demand**

(A7) **Log-convex demand**

$$(\text{A7}) \Longrightarrow (\text{A6}).$$

Bounding the tariff DWL across countries and products

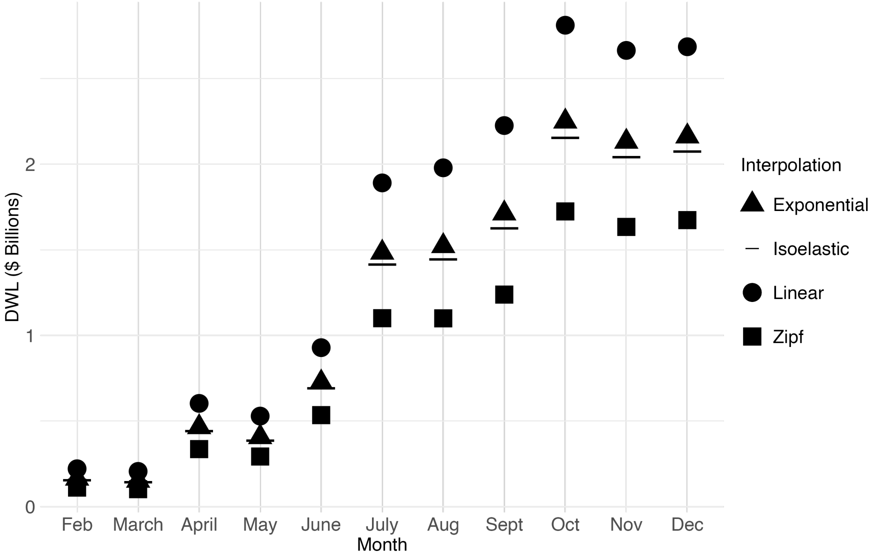

Further extensions: welfare beyond ΔCS

#1. Producer surplus works just as well as CS.

#2. Can handle heterogeneity + distributional questions.

#3. Can handle alternative welfare measures like EV and CV.

#4. Can handle multiple objectives at once.

- $\rightsquigarrow$ E.g., Pareto-weighted consumer surplus + DWL.

#5. Can handle multi-product markets.

- $\rightsquigarrow$ At least under constraints on cross-price and own-price elasticities.

Summing up

- **This paper.** Develops a framework to bound welfare based on economic reasoning.
- **Building on previous work.** Hope to make the case that everyone should use this.
- **Use cases.** Draw/assess conclusions from empirical objects commonly estimated.
- **Future work.** We're excited about this.
  - Robustness for structural IO-style problems (e.g., inference with endogenous pricing, merger screens, welfare in horizontally differentiated good markets).
  - Robustness for new goods and price indices (e.g., the CPI).
  - Robustness for larger macro models (e.g., extending ACR, ACDR).

Amiti, Mary, Stephen J. Redding, and David E. Weinstein, "The Impact of the 2018 Tariffs on Prices and Welfare," *Journal of Economic Perspectives*, 2019, *33* (4), 187–210.

Brynjolfsson, Erik, Yu Hu, and Michael D. Smith, "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers," *Management Science*, 2003, *49* (11), 1580–1596.

Caplin, Andrew and Barry Nalebuff, "Aggregation and Imperfect Competition: On the Existence of Equilibrium," *Econometrica*, 1991, *59* (1), 25–59.

__ **and** __ , "Aggregation and Social Choice: A Mean Voter Theorem," *Econometrica*, 1991, *59* (1), 1–23.

Fajgelbaum, Pablo D., Pinelopi K. Goldberg, Patrick J. Kennedy, and Amit K. Khandelwal, "The Return to Protectionism," *Quarterly Journal of Economics*, 2020, *135* (1), 1–55.

Finkelstein, Amy and Nathaniel Hendren, "Welfare Analysis Meets Causal Inference," *Journal of Economic Perspectives*, 2020, *34* (4), 146–67.

Hackmann, Martin B., Jonathan T. Kolstad, and Amanda E. Kowalski, "Adverse Selection and an Individual Mandate: When Theory Meets Practice," *American Economic Review*, 2015, *105* (3), 1030–66.

Hahn, Robert W. and Robert D. Metcalfe, "Efficiency and Equity Impacts of Energy Subsidies," *American Economic Review*, 2021, *111* (5), 1658–88.

Harberger, Arnold C., "The Measurement of Waste," *American Economic Review: Papers & Proceedings*, 1964, *54* (3), 58–76.

Hausman, Jerry A., "Exact Consumer's Surplus and Deadweight Loss," *American Economic Review*, 1981, *71* (4), 662–676.

— , **Ariel Pakes, and Gregory L. Rosston**, "Valuing the Effect of Regulation on New Services in Telecommunications," *Brookings Papers on Economic Activity, Microeconomics*, 1997, *1997*, 1–54.

Willig, Robert D., "Consumer's Surplus Without Appology," *American Economic Review*, 1976, *66* (4), 589–597.

Mapping CS to EV/CV when income effects are small

Consumer surplus provides bounds for equivalent and compensating variations.

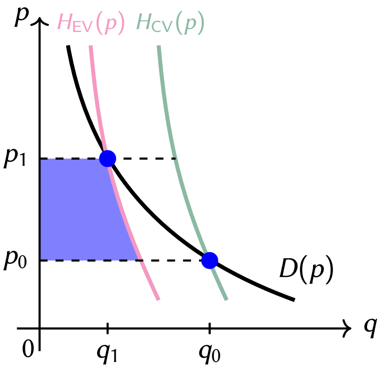

- **Generally:** EV $\leq$ CS $\leq$ CV.

Consumer surplus provides bounds for equivalent and compensating variations.

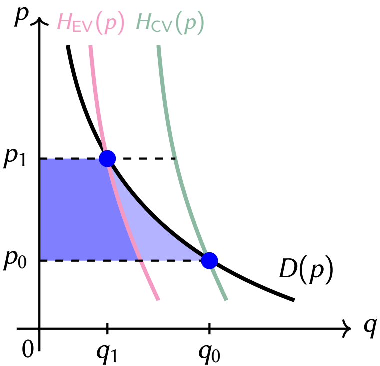

- **Generally:** EV $\leq$ CS $\leq$ CV.

Consumer surplus provides bounds for equivalent and compensating variations.

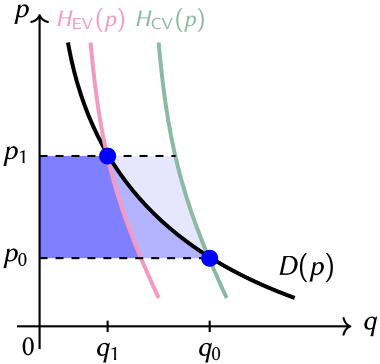

- **Generally:** EV $\leq$ CS $\leq$ CV.

Mapping CS to EV/CV when income effects are small

Consumer surplus provides bounds for equivalent and compensating variations.

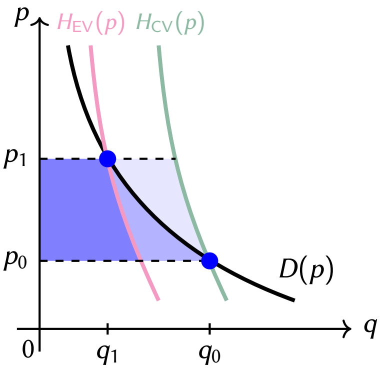

- **Generally:** EV $\leq$ CS $\leq$ CV.
- When income effects are 0 (e.g., with quasilinearity): EV = CS = CV.
- When income effects are ≈ 0: EV $\approx$ CS $\approx$ CV (Willig, 1976) (also if demand is pretty inelastic).

We can compute EV/CV bounds under assumptions about the *Hicksian* demand curve.

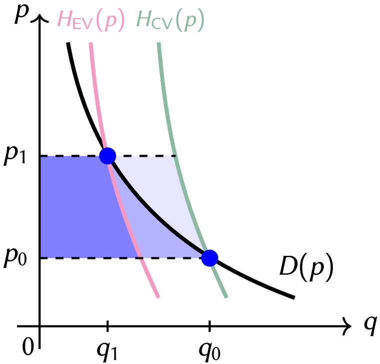

- **But!** we don't observe counterfactual expenditures.
- Need to bound $e(p_1, u_0)$ for CV.
- Need to bound $e(p_0, u_1)$ for EV.
- This maps to our "1-point" extension.

◂ Basic Model ▸ Skip to End