The explication here draws more from the NBER technical working paper (#178) version of the paper.

■■■

MOTIVATION

How to *interpret* the IV estimates from a demand function. They focus on general situation where

1. Nonlinear demand function: potentially, slope of demand function (i.e., $\frac{\partial q}{\partial p}$ differs across p).
2. Slope of demand function also varies across time (or across cross-sectional units).

This is the same general situation as in the BLP paper. The BLP approach to estimating nonlinear demand functions is to write down a flexible functional form specification (the random coefficients logit model), and estimate it (and overcome computational hurdles associated with "nested" algorithm, simulation in the process).

The AGI approach, however, is strikingly different. Instead of writing down a model they instead focus on the questions:

What does the linear IV regression of q on p identify, even if the true (but unknown) demand function is not linear?

Focus in the paper is extending some of the insights from the "program evaluation" literature in labor economics, which attempts to uncover the "causal effects" of social programs on outcome variables. Here, think of price as a "treatment", and we investigate the "causal effect" of price on quantity demanded.

■■■

MODEL SETUP AND NOTATION

$q_t^d(p, z, x)$: demand function in period t
$q_t^s(p, z, x)$: supply function in period t

Data are a time series $t = 1, \ldots, T$ (can also interpret as a cross-sectional index). They observe q and p for a number of days in the Fulton Fish Market in NYC.

Consider 4 cases (in increasing generality):

1. Linear system with constant coefficients (across t):

$$\begin{aligned} q_t^d(p,z,x) &= \alpha_0 + \alpha_1 p + \alpha_2 z + \alpha_3 x + \epsilon_t \\ q_t^s(p,z,x) &= \beta_0 + \beta_1 p + \beta_2 z + \beta_3 x + \eta_t. \end{aligned}$$

2. Linear system with non-constant (time-varying) coefficients:

$$\begin{aligned} q_t^d(p,z,x) &= \alpha_{0t} + \alpha_{1t} p + \alpha_{2t} z + \alpha_{3t} x + \epsilon_t \\ q_t^s(p,z,x) &= \beta_{0t} + \beta_{1t} p + \beta_{2t} z + \beta_{3t} x + \eta_t. \end{aligned}$$

3. Nonlinear system with "constant shape" (across t, only intercept changes):

$$\begin{aligned} q_t^d(p,z,x) &= q^d(p,z,x) + \epsilon_t \\ q_t^s(p,z,x) &= q^s(p,z,x) + \eta_t. \end{aligned}$$

4. Nonlinear system with "time-varying shape":

$$\begin{aligned} q_t^d(p,z,x) &= q^d(p,z,x,\epsilon_t) \\ q_t^s(p,z,x) &= q^s(p,z,x,\eta_t). \end{aligned}$$

Note: the BLP random-coefficients logit demand functions would fall under case 4.

■■■

There are two dimensions of "heterogeneity" in effects of price p on demand q^d:

1. heterogeneity depending on value of p, fixing t (not an issue in the linear cases 1 and 2, but relevant in nonlinear cases 3 and 4)
2. heterogeneity across t, fixing p (only relevant in cases 2 and 4).

But problem is that: you do not know which case you are in. So question is: under worst-case scenario (case 4), can you still interpret the linear IV estimate?

■■■

BINARY INSTRUMENT CASE

It is useful to focus in some detail on the case where the instrument Z is binary-valued (0 or 1) for the demand function. We can draw some parallels between this case and the selection problem in the "program evaluation" literature, where usually both the endogenous x (the "treatment") as well as the instrument are binary-valued.

Assumption:

1. Regularity conditions on $q_t^d(p, z, x)$ and $q_t^s(p, z, x)$.
   (a) sequence $\left\{q_t^d(p, z, x), q_t^s(p, z, x), p_t, z_t, w_t\right\}$ has well-defined first and second moments (implies some sort of stationarity). If t is a cross-sectional index, usual assumption is that $\{\cdots\}$ are independent across t. (Serial correlation across t does not affect calculation of estimators, but affects variance-covariance matrix.)
   (b) Functions $q_t^d(p, z, x), q_t^s(p, z, x)$ are continuously differentiable in p
2. z_t is valid instrument in demand function
   (a) Independence: for all (p, z, x), z_t is jointly independent of $\left\{q_t^d(p, z, x), q_t^s(p, z, x)\right\}$ given x_t. That is:
   $$\begin{aligned} \forall \hat{p}, \hat{z}, \hat{x}: \; & f\left(q_t^d(\hat{p}, \hat{z}, \hat{x}), q_t^s(\hat{p}, \hat{z}, \hat{x}), z_t | x_t\right) = \\ & f\left(q_t^d(\hat{p}, \hat{z}, \hat{x}) | x_t\right) \cdot f\left(q_t^s(\hat{p}, \hat{z}, \hat{x}) | x_t\right) \cdot f\left(z_t | x_t\right). \end{aligned}$$
   Intuitively, fixing $\hat{p}, \hat{z}, \hat{x}$, randomness in $q_t^d(\hat{p}, \hat{z}, \hat{x})$ and $q_t^s(\hat{p}, \hat{z}, \hat{x})$ due to ϵ_t and η_t: so condition is that ϵ_t, η_t, and z_t are mutually independent, conditional on x_t.
   For standard linear model (case A), ϵ_t and η_t are usually assumed independent of z_t, unconditionally on x_t (stronger than here).
   (b) Exclusion: for all p, t,
   $$q_t^d(p, z = 1, x_t) = q_t^d(p, z = 0, x_t) \equiv q_t^d(p, x_t)$$
   that is, z is "excluded" from demand equation (it only has an effect on q_t^d indirectly, through the endogenous variable p).

In the standard linear model, this corresponds to the usual exclusion restriction that $\alpha_2 = 0$.

(c) For some periods t, $q_t^s(p_t, 1, x_t) \neq q_t^s(p_t, 0, x_t)$. That is, z is a "cost-shifter" which shifts the supply curve. Implies that movement in z cause movement in p (via the reduced form).

For the linear model (case A), assumptions 2a, 2c are like standard assumption on linear IV that the IV must be "uncorrelated with ϵ_t, but correlated with the endogenous p_t".

■■■

Given these assumption, next discuss identification. Focus on simplest case of

- binary instrument $z = \begin{cases} 0 & \text{calm weather at sea} \\ 1 & \text{stormy weather at sea} \end{cases}$

  Story is that weather is a "cost-shifter", which affects the price of fish only through the costs, but not correlated with demand-side unobservables. (For example, this would be invalid if people prefer to eat fish on cloudy days.)

- No x (focus just on estimation of demand function)

Hence, demand function just has one regressor, price p.

Linear IV estimator using z as instrument for regression of q on p is the "Wald" estimator

$$\begin{aligned} \hat{\alpha}_{1,0} &\equiv \frac{\frac{\sum_t z_t q_t}{\sum_t z_t} - \frac{\sum_t (1-z_t) q_t}{\sum_t (1-z_t)}}{\frac{\sum_t z_t p_t}{\sum_t z_t} - \frac{\sum_t (1-z_t) p_t}{\sum_t (1-z_t)}} \\ &\xrightarrow{p} \frac{E[q_t|z_t = 1] - E[q_t|z_t = 0]}{E[p_t|z_t = 1] - E[p_t|z_t = 0]} \equiv \alpha_{1,0} : \end{aligned}$$

the ratio of differences in means.

Now, with case 1, $\alpha_{1,0}$ is a consistent estimator of structural parameter α_1. Furthermore, *any* IV estimator is estimator of α_1. No ambiguity about interpretation.

What about more complicated models, (say) case 4?

- $\alpha_{1.0}$ is not an estimator of structural parameter.
- Furthermore, it identifies something which differs, depending on which instrument you use!

  More broadly, the results in this paper imply that ***it may be wise to*** **divorce** ***structural estimation from estimating "deep" population parameters (as advocated by the Lucas critique). The authors' point is that the linear IV estimator $\hat{\alpha}_{1,0}$ can identify some aspects of the*** **causal effects** ***of p on q, without identifying some deep structural parameter***.

■■■

In order to interpret $\hat{\alpha}_{1,0}$ (or its large-sample counterpart $\alpha_{1,0}$ for case 4, make additional assumptions:

The first set of assumption concerns **Equilibrium:**

1. observed price is market-clearing price:

$$q_t^d(p_t) - q_t^s(p_t, z_t) = 0, \ \forall t.$$

2. "potential prices": for each value of the instrument z, there is a unique market clearing price:

$$\forall z, t: \ \tilde{p}(z,t) \ s.t. \ q_t^d\left(\tilde{p}(z,t)\right) = q_t^s\left(\tilde{p}(z,t), z\right).$$

   $\tilde{p}(z,t)$ is "potential price" under any counterfactual event (z,t).

   Realized price $p_t = \tilde{p}(z_t, t)$.

   (This is like in the program evaluation literature, where only observe the realized treatment D.)

To motivate the second assumption, look at the denominator of the Wald estimator:

$$E\left[p_t | z_t = 1\right] = E\left[\tilde{p}(z_t, t) | z_t = 1\right]$$
$$E\left[p_t | z_t = 0\right] = E\left[\tilde{p}(z_t, t) | z_t = 0\right].$$

Wald estimator is well-defined only when denominator is nonzero, or when $E[p_t|z_t = 1] \neq E[p_t|z_t = 0]$.

This is ensured by the following **monotonicity** assumption:

$$\forall t: \ \tilde{p}(z, t) \text{ is nondecreasing (weakly increasing) in } z.$$

- This is a powerful assumption, which allows us to deduce an "average slope" interpretation of $\alpha_{1.0}$. A similar monotonicity condition plays an important role in the program evaluation context (Angrist and Imbens (1994)).
- However, this assumption is *untestable* because, for each t, you only observe either $\tilde{p}(0, t)$ or $\tilde{p}(1, t)$ (same as in the program evaluation context, where for each individual, you observe each individual only once, either with or without the treatment). Any story about how IV affects the endogenous variable is just a story, and untestable: but institutional details are useful to justify the story.
- In this sense, IV methods are "cross-sectional" methods, in the sense that you only observe each unit once. In a *panel* context, there is the possibility that for at least a subset of the units, you observe the outcome both with and without the treatment, allowing you to do a "difference-in-difference" analysis.

  In panel context, if you observe both IV and endogenous variables over time, for the same individual, then monotonicity assumption is potentially testable.

■■■

Lemma 1: Numerator of $\alpha_{1,0}$

$$E[q_t|z_t = 1] - E[q_t|z_t = 0] = E\left\{\int_{\tilde{p}(0,t)}^{\tilde{p}(1,t)} \frac{\partial q_t^d(s)}{\partial s} ds\right\}$$

The expectation on the RHS denotes averaging across t. The numerator is a *double average*: first, there is averaging over t; and second, there is averaging over the range of prices $[\tilde{p}(0, t), \tilde{p}(1, t)]$ induced by a shift in the instrument from 0 to 1.

For a given t, the term $\int_{\tilde{p}(0,t)}^{\tilde{p}(1,t)} \frac{\partial q_t^d(s)}{\partial s} ds$ is the average slope of the demand curve between $\tilde{p}(0, t)$ and $\tilde{p}(1, t)$. Note that this range is *different* for each t.

- Hence, $\alpha_{1,0}$ yields information about the demand curve only in the range of potential price variation induced by variation in the instrument $\Rightarrow$ don't know about the shape of the demand curve outside the range $[\tilde{p}(0,t), \tilde{p}(1,t)]$! (See Fig. 1 in the RES version of the paper.)

- Hence, for different instruments z, $\alpha_{1,0}$ is interpreted as different things. This differs from standard linear IV case with constant coefficients, where any valid IV estimator (regardless of which IV you use) yields a consistent estimator of the *same* parameter.

- Without monotonicity assumption, you do not know what the appropriate range of integration across prices is: for some t, you might integrate from $\tilde{p}(0,t)$ to $\tilde{p}(1,t)$, but for others, you might be integrating from $\tilde{p}(1,t)$ to $\tilde{p}(0,t)$.

- Furthermore, observations t for which $\tilde{p}(1,t) = \tilde{p}(0,t)$ (allowed under the monotonicity assumption) are not included in the averaging (because $\int_{\tilde{p}(0,t)}^{\tilde{p}(1,t)} \cdots = 0$). So weighted average is only over a particular sub-sample of the population. However, since you never observe both potential prices, you cannot pinpoint which sub-sample it is relevant for.

Proposition 1: Interpretation of $\alpha_{1,0}$:

$$\begin{aligned}
\alpha_{1,0} &= \frac{E\left\{\int_{\tilde{p}(0,t)}^{\tilde{p}(1,t)} \frac{\partial q_t^d(s)}{\partial s} ds\right\}}{E\tilde{p}(1,t) - E\tilde{p}(0,t)} \\
&= \frac{\int_0^\infty \left\{E\left[\frac{\partial q_t^d(s)}{\partial s}\mathbf{1}\left(s \in [\tilde{p}(0,t), \tilde{p}(1,t)]\right)\right]\right\} ds}{E\tilde{p}(1,t) - E\tilde{p}(0,t)} \\
&= \int_0^\infty \left\{E\left[\frac{\partial q_t^d(s)}{\partial s}\middle| s \in [\tilde{p}(0,t), \tilde{p}(1,t)]\right] \frac{F_{\tilde{p}(0,t)}(s) - F_{\tilde{p}(1,t)}(s)}{E\tilde{p}(1,t) - E\tilde{p}(0,t)}\right\} ds \\
&\equiv \int_0^\infty E\left[\frac{\partial q_t^d(s)}{\partial s}\middle| s \in [\tilde{p}(0,t), \tilde{p}(1,t)]\right] \omega(s) ds
\end{aligned}$$

where $F_{\tilde{p}(0,t)}(s)$ denote Prob $(\tilde{p}(0,t) \leq s)$, where the probability is taken over randomness in the t dimension.

Once again, there is double-average feature:

- given t, average slope of q_t^d between $\tilde{p}(0,t)$ and $\tilde{p}(1,t)$.
- given price $s \in [\tilde{p}(0,t), \tilde{p}(1,t)]$, average of $q_t^d(s)$ across t (randomness due to ϵ_t.

Weight function $\omega(s)$:

- Note that weight $\omega(s)$ is *not* a function of t.
- the weight on a price s is proportional to $F_{\tilde{p}(0,t)}(s) - F_{\tilde{p}(1,t)}(s)$: roughly speaking, weight is largest for prices which are "most likely" to fall between $\tilde{p}(0,t)$ and $\tilde{p}(1,t)$.

Examples:

- Case 1: $\alpha_{1,0} = \alpha_1$. No ambiguity.
- Case 2: $q_t^d(p) = \alpha_{0t} + \alpha_{1t} p + \epsilon_t$, so that

$$\alpha_{1,0} = \frac{E\left[\alpha_{1t}\left(\tilde{p}(1,t) - \tilde{p}(0,t)\right)\right]}{E\tilde{p}(1,t) - E\tilde{p}(0,t)} \neq E\alpha_{1t}!$$

  Intuitive outcome arises when (say) $(\tilde{p}(1,t) - \tilde{p}(0,t))$ is (mean) independent of α_{1t}.

■■■

MORE COMPLICATED CASES: continuous instrument case

Now assume that z takes on continuous values. $h_t(z) \equiv \tilde{p}(z,t)$ denotes potential price *function* in period r, for all possible values of z. (Note: even fixing z, $\tilde{p}(z,t)$ is still random, due to the random errors ϵ_t and η_t.)

Again, consider the univariate case (i.e. only one regressor, p, in demand equation).

Given continuous z, we consider the following Wald estimator, which is now defined for each value of z:

$$\alpha(z) = \lim_{\nu \to 0} \frac{E(q_t|z) - E(q_t|z-\nu)}{E(p_t|z) - E(p_t|z-\nu)}.$$

In principle, if you have lots and lots of data, the conditional expectation $E(q_t|z)$ can be approximated by the sample average of q_t over those observations where $z_t = z$:

$$E(q_t|z) \approx \frac{\sum_{t=1}^{T} \mathbf{1}\left(z_t = z\right) q_t}{\sum_{t=1}^{T} \mathbf{1}\left(z_t = z\right)}.$$

However, in practice, when you have a small dataset, you want to "smooth" the weights in the above expression $\left(\frac{\mathbf{1}(z_t=z)}{\sum_{t=1}^{T}\mathbf{1}(z_t=z)}\right)$. The most common way is "kernel smoothing":

$$E(q_t|z) \approx \hat{q}(z) \equiv \frac{1}{T}\sum_{t=1}^{T} v(z_t, z) q_t$$

which is a weighted average of the q's with the weights

$$v(z_t, z) \equiv \frac{\frac{1}{h}\mathcal{K}\left(\frac{z_t - z}{h}\right)}{\frac{1}{h}\sum_{t'=1}^{T}\mathcal{K}\left(\frac{z_{t'} - z}{h}\right)}$$

and the weights sum to 1, across observations t. In the above $\mathcal{K}(\cdots)$ denotes a "kernel function". $\hat{p}(z)$ defined analogously.

> **Technical aside:** a kernel function is a smoothing function, which is symmetric around zero, and integrates to 1 along its support. Common examples are:
>
> - Epanechnikov kernel: $\mathcal{K}(u) = \frac{3}{4}(1-u^2)\mathbf{1}(|u| \leq 1)$
> - Uniform kernel: $\mathcal{K}(u) = \frac{1}{2}\mathbf{1}(|u| \leq 1)$
> - Normal kernel: $\mathcal{K}(u) = (2\pi)^{\frac{1}{2}}\exp(-u^2/2)$

Then $\alpha(z)$ can be calculated in practice by

$$\alpha'(z) = \frac{\hat{q}'(z)}{\hat{p}'(z)}.$$

which can be approximated by the finite differences:

$$\alpha'(z) \approx \frac{\hat{q}(z+h) - \hat{q}(z))}{\hat{p}(z+h) - \hat{p}(z)}, \quad h \text{ small.}$$

Lemma 2: Interpretation of local IV estimator $\alpha(z)$:

$$\begin{aligned}\alpha(z) &= \frac{\partial E(q_t|z)/\partial z}{\partial E(p_t|z)/\partial z} \\ &= E\left\{ \left(\frac{\partial q_t^d(h_t(z))}{\partial p} \right) \cdot \frac{h_t'(z)}{E(h_t'(z))} \right\}.\end{aligned}$$

This is a weighted average slope of demand function. Averaging is only *across time*, keeping price "fixed" at the hypothetical level $h_t(z)$ (i.e. the potential price that would obtain if the value of the instrument were fixed at z). (Note: recall that fixing z, the potential price $h_t(z)$ is still random due to ϵ_t and η_t.)

■■■

The estimator considered in the previous section is "local", in the sense that it takes a different value depending on which value of the instrument z you focus on. We next consider ways to "average" this function across different values of z.

Specifically, consider a 2SLS-like procedure where instrument is an estimate of $E(p_t|z)$. (Note that this accommodates multiple instruments.)

Let $p_t^* \equiv E(p_t|z_t) - \mu_p$ denote instruments where μ_p is the across-time mean of p_t.

Consider the estimator

$$\hat{\alpha}_E = \sum_t q_t p_t^* \Big/ \sum_t p_t p_t^*.$$

Note that this estimator averages across different values of z, in a particular sort of way.

Proposition 3: Interpretation of α_E:

$$\alpha_E = \int_0^\infty \alpha(z)\omega(z)dz$$

where the weight function

$$\omega(z) = \frac{1}{VarE(p_t|z)} \cdot \{[E(p_t|z_t \geq z) - E(p_t|z_t \leq z)] \cdot Prob(z_t \geq z) \cdot [1 - Prob(z_t \geq z)] \cdot E(h_t'(z))\}.$$

■■■

When are these results useful?

- When you have a parametric model, you don't need these results, since there is no ambiguity in interpreting results (and you can define the appropriate nonlinear IV estimators for parameters). This is the approach taken in, eg., BLP.

- When you *don't* have a parametric model (or don't feel confident enough to write down one), these results show that linear IV estimators still have some interpretation as estimators of average slopes.

- This analysis is limited to a single product. what about multi-product markets? (One potential problem: monotonicity requirement more difficult to justify in oligopolistic, multi-product, multi-firm setting.)

References

Angrist, J., K. Graddy, and G. Imbens (2000): "The Interpretation of Instrumental Variables Estimators in Simulataneous Equations Models with an Application to the Demand for Fish," *Review of Economic Studies*, 67, 499–527.

Angrist, J., and G. Imbens (1994): "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 62, 467–476.