

Estimating Models of Supply and Demand: Instruments and Covariance Restrictions[†]

By ALEXANDER MACKAY AND NATHAN H. MILLER*

We consider the identification of empirical models of supply and demand with imperfect competition. We show that a restriction on the covariance between unobserved demand and cost shocks can resolve endogeneity and identify the price parameter. We demonstrate how to employ this approach in estimation, and we compare it to the method of instrumental variables. Our formal results also indicate that weaker covariance restrictions can bound the price parameter. We illustrate the covariance restriction approach with applications to ready-to-eat cereal, cement, and airlines. (JEL C51, D12, L13, L61, L66, L93)

A fundamental challenge in identifying models of supply and demand is that firms can adjust markups in response to demand shocks. Even if marginal costs are constant, this source of price endogeneity generates upward-sloping supply in settings with imperfect competition. Thus, the empirical relationship between prices and quantities represents not a demand curve but rather a mixture of demand and supply. Researchers typically address this challenge by using supply-side instruments to estimate demand and then using the supply model to recover marginal costs and simulate counterfactuals (e.g., Berry, Levinsohn, and Pakes 1995; Nevo 2001).

In this paper, we develop an alternative identification strategy that exploits covariance restrictions between demand-side and supply-side structural error terms. We first establish a function linking the price parameter and the covariance of unobservable cost and demand shocks for a broad class of oligopoly models. We then show how this relationship can identify the price parameter and pin down the slope of demand in estimation. A key distinction between our approach and the use of instrumental variables is that we interpret endogenous variation in prices and quantities through the lens of the model, rather than relying on an additional (observed) variable to isolate exogenous variation in price. The core intuition is that the supply side

*MacKay: Department of Economics, University of Virginia (email: mackay@virginia.edu); Miller: McDonough School of Business, Georgetown University (email: nathan.miller@georgetown.edu). Leeat Yariv was coeditor for this article. We thank Chris Adams, Steven Berry, Renata Gaineddenova, Charlie Murry, Matt Osborne, Chuck Romeo, Gloria Sheu, Karl Schurter, Jesse Shapiro, Michael Sullivan, Jeff Thurk, Andrew Sweeting, Matthew Weinberg, and Nathan Wilson for helpful comments. We also thank seminar and conference participants at Harvard University, MIT, the University of Maryland, the Barcelona GSE Summer Forum, and the NBER Summer Institute. Previous versions of this paper were circulated with the title “Instrument-Free Demand Estimation.” Replication files are available from the AEA Data and Code Repository (MacKay and Miller 2025).

[†]Go to <https://doi.org/10.1257/mic.20230024> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

of the model dictates how prices respond to demand shocks, shaping the relative variation of quantities and prices in the data. We explore the promise and limitations of the covariance restriction approach to estimation, both theoretically and in the context of three empirical applications that we draw from the literature.

In Section I, we outline the data-generating process for our baseline model and provide formal identification results. The model can accommodate standard empirical demand systems, such as logit and random coefficients logit, among others. The supply-side assumptions nest different models of conduct for oligopolists with constant marginal costs, including differentiated-products Bertrand competition and Cournot competition. In this setting, prices are endogenous because they respond to a demand shock (the demand-side “structural error term”) that is unobserved to the econometrician.

We prove that the price parameter solves a quadratic equation in which the coefficients are functions of observables and the covariance between demand and cost shocks. With a restriction on the covariance term, the price parameter is identified up to (at most) two points. Under reasonable conditions, the price parameter is the more negative solution, and point identification is obtained. The price parameter can be computed directly from an analytical solution, or the covariance restriction can be recast as an orthogonality condition and estimation can proceed with the method of moments. We show how the empirical variation in (transformed) quantities and prices is informative about the price parameter. All else equal, more elastic demand yields greater variation in quantities relative to variation in prices.¹

Point identification can thus be achieved with a fully specified model and a covariance restriction on unobserved shocks. Whether the covariance restriction is credible depends on the details of the economic environment. For example, in the presence of capacity constraints, a positive demand shock can increase marginal costs, yielding an overall positive correlation. Using weaker assumptions, we show how our results can be used to achieve partial identification and bound the price parameter. If the econometrician can sign the correlation between unobserved demand and cost shocks, then one-sided bounds can be placed on the price parameter. Furthermore, our results show when it is possible to rule out some values of the price parameter without any assumption about this correlation.

In Section II, we compare the assumptions under which the covariance restriction approach obtains point identification to the corresponding assumptions of the instrumental variable approach. Both approaches employ restrictions on the correlation structure of unobservables. As typically defined, valid instruments satisfy an exclusion condition and a relevance condition. The exclusion restriction for instrumental variables is similar to the covariance restriction when the econometrician uses cost variation as an instrument. However, the instrumental variable approach requires the econometrician to observe exogenous variation in order to identify the model (e.g., Wooldridge 2010; Berry and Haile 2014). The covariance restriction approach avoids this requirement by imposing the stronger assumption that the residual unobserved variation is uncorrelated. This assumption also precludes the need for an

¹This is true even if the empirical relationship between prices and quantities is upward sloping. Our method will still recover the correct downward-sloping demand curve.

empirical relevance condition; all of the observed variation in prices and quantities is used in estimation. By contrast, the instrumental variable approach can suffer from bias if the instrument and price are weakly correlated.

After formalizing these distinctions, we compare our identification strategy to an approach that uses “residual instruments” recovered under a covariance restriction (Hausman and Taylor 1983) and to the oft-used instruments of Hausman (1996). Finally, we use Monte Carlo simulations to show that, when the necessary assumptions are met, a covariance restriction outperforms instruments in finite samples, especially when instruments become weak.

In Section III, we provide extensions to the baseline model. First, we consider a more general class of covariance restrictions that might be employed in practice. For example, for the Hausman instruments—prices in related markets—to be valid, three economic assumptions about the correlation structure between unobserved shocks must be satisfied. We discuss how one could employ these assumptions directly in a method of moments estimator rather than rely on an instrumental variables implementation using observed prices.

We also discuss models in which costs are not constant in quantity. In such cases, the response of prices to demand shocks is mediated by the slope of the cost curve. One approach to resolve this issue is to explicitly model the (nonconstant) marginal cost function. A covariance restriction may then be credible, though identification requires additional moments for any parameters that enter the nonconstant portion of marginal costs. Alternatively, one could forgo the estimation of the marginal cost function and, instead, invoke the bounds approach.

We apply these methods in a series of empirical applications (Section IV). The three settings that we have selected—ready-to-eat (RTE) cereals, cement, and airlines—differ in a variety of ways that influence our implementation. With RTE cereals, marginal costs can plausibly be modeled as constant, so we proceed with estimation under a covariance restriction, using fixed effects to absorb potentially confounding variation. With cement, capacity constraints imply that marginal costs may increase with quantities. We follow an approach developed in the literature and model this effect explicitly, after which we view a covariance restriction as credible. Finally, with airlines, we apply a bounds approach that uses weaker assumptions about the relationship between demand and cost shocks, as the stronger assumptions used in the two preceding applications may not be credible for that industry. In each case, we show how covariance restrictions support inferences about the price parameter.

Together, these results show how covariance restrictions can help overcome a significant obstacle for empirical research: the need to find valid instruments for price. In our three applications, we obtain parameter values that are consistent with instrumental variable estimates. Outside of this paper, Döpfer et al. (forthcoming) show that covariance restrictions deliver estimates comparable to those in the industrial organization literature (for cereal, yogurt, and beer) as well as in the international trade literature (coffee). De los Santos, O’Brien, and Wildenbeest (2024) use covariance restrictions in the context of e-books. Like other electronic goods, e-books can have substantial fluctuations in demand over time that are unrelated to changes in marginal costs. These examples indicate that covariance restrictions can deliver reasonable estimates in different settings.

Empirical models of imperfect competition typically have other key parameters that characterize heterogeneity of consumer preferences. To identify these parameters, researchers have used micro-moments constructed from the observed behavior of individual consumers (e.g., Backus, Conlon, and Sinkinson 2021; Döpfer et al. forthcoming) and “second-choice” data from surveys about consumers’ ordered preferences over products (e.g., Grieco, Murry, and Yurukoglu 2023). These strategies identify the consumer heterogeneity parameters but do not resolve price endogeneity (Berry and Haile 2024). Thus, the covariance restriction approach that we examine is a useful complement to the use of detailed consumer data.²

To put our results in context, covariance restrictions were analyzed in early research on the identification of linear systems of equations, including supply and demand models of perfect competition (e.g., Koopmans 1949; Koopmans, Rubin, and Leipnik 1950).³ With perfect competition, the supply curve is upward sloping due to increasing costs of production. With upward-sloping supply and downward-sloping demand, two separate restrictions are required for identification (Hausman and Taylor 1983). If, instead, price endogeneity arises due to the markup adjustments that occur in models of imperfect competition, then (as we show) a single restriction is sufficient for identification.

The strategy of using supply-side restrictions to reduce identification requirements has parallels in a handful of other articles. Leamer (1981) examines a linear model of perfect competition and provides conditions under which the price parameters can be bounded using only the endogenous variation in prices and quantities. Feenstra (1994) considers the case of monopolistic competition with constant markups, and a number of applications in the trade literature extend this constant-markup approach (e.g., Broda and Weinstein 2006, 2010; Soderbery 2015).⁴ Zoutman, Gavrilova, and Hopland (2018) return to perfect competition and show that under a standard assumption in models of taxation, both supply and demand can be estimated with exogenous variation in a single tax rate. At a high level, our approach to estimation with covariance restrictions relates to Petterson, Seim, and Shapiro (2023), who show how to bound structural parameters based on beliefs about the magnitudes of unobserved shocks. Our research complements these articles by developing results for imperfect competition with adjustable markups.⁵

²Alternatively, if instruments are constructed from the characteristics of competing products (e.g., Berry, Levinsohn, and Pakes 1995; Gandhi and Houde 2023), then the covariance restriction could be incorporated using the generalized method of moments (GMM) as an additional identifying restriction.

³Many articles advanced this research agenda, which began at the Cowles Foundation, in subsequent decades (e.g., Fisher 1963, 1965; Wegge 1965; Rothenberg 1971; Hausman and Taylor 1983; Hausman, Newey, and Taylor 1987). More recently, Matzkin (2016) examines covariance restrictions in semiparametric models.

⁴There are interesting historical antecedents to this trade literature. Leamer (1981) attributes an early version of his results to Schultz (1928). The identification argument of Feenstra (1994) is also proposed in Leontief (1929). Frisch (1933) provides an important econometric critique.

⁵Some applications in industrial organization identify demand-side parameters with the assistance of supply-side assumptions (e.g., Thomadsen 2005; Cho et al. 2018; Li et al. 2022). Among these, Thomadsen (2005) assumes that there are not unobserved demand shocks, and Cho et al. (2018) assume that there are no unobserved cost shocks; both are special cases of the covariance restriction approach.

I. Model and Identification

A. Data-Generating Process

The model examines supply and demand across markets. Markets can be conceptualized as (for example) separate locations, time periods, or both. In each market t , there is a set $\mathcal{J}_t = \{0, 1, \dots, J_t\}$ of products available for purchase. The market $t = 1, \dots, T$ is defined by (\mathcal{J}_t, χ_t) , where

$$\chi_t = \{\mathbf{x}_t, \mathbf{w}_t, \boldsymbol{\xi}_t, \boldsymbol{\eta}_t\}$$

contains product and market characteristics. Among these, $\mathbf{x}_t = [x_{1t}; \dots; x_{J_t t}]$ is a $J_t \times K$ matrix of (nonprice, exogenous) product-market characteristics that are observable to the econometrician, \mathbf{w}_t contains observable variables that arise in some specifications of the model, and $\boldsymbol{\xi}_t = (\xi_{1t}, \dots, \xi_{J_t t})$ and $\boldsymbol{\eta}_t = (\eta_{1t}, \dots, \eta_{J_t t})$ are mean-zero $J_t \times 1$ vectors of unobservable product-level or market-level characteristics. We sometimes refer to the unobservable characteristics as “structural error terms.” Let each $\xi_{jt}, \eta_{jt} \in \mathbb{R}$ and each $x_{jt} \in \mathbb{R}^K$ be a $K \times 1$ (row) vector. We assume that the first element in each x_{jt} equals one—i.e., that the characteristics in \mathbf{x}_t include a constant. The dimension of \mathbf{w}_t depends on the modeling specification. Without loss of generality, let $\mathcal{J}_t = \mathcal{J} = \{0, 1, \dots, J\}$.

Prices and quantities are determined endogenously by market participants. Let $\mathbf{p}_t = (p_{1t}, \dots, p_{J_t t})$ be a vector of prices and $\mathbf{q}_t = (q_{1t}, \dots, q_{J_t t})$ be a vector of quantities, with $p_{jt}, q_{jt} \in \mathbb{R}$. Both prices and quantities are observable to the econometrician. The parameters of the model are in the set $\boldsymbol{\theta} = \{\theta_1, \theta_2\}$. Following Nevo (2001), we let θ_1 include parameters that affect demand and supply linearly in a manner that we specify below, whereas we use θ_2 for additional parameters that enter with some specifications of the model. An example of the latter is the nesting parameter that enters if demand is nested logit (e.g., Berry 1994). Our main identification results are for θ_1 . Covariance restrictions also can help pin down θ_2 in some settings.

On the demand side, we assume that the quantity of each product is determined by $q_{jt} = \sigma^{(j,t)}(\mathbf{p}_t, \chi_t; \boldsymbol{\theta})$, where each $\sigma^{(j,t)}$ is a demand function. We also assume that, for every (j, t) , there exists a known function $h^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2)$ that is increasing in the quantity of the product (q_{jt}). The function can be interpreted as providing transformed quantities as a function of \mathbf{w}_t and θ_2 . We provide examples of $h^{(j,t)}$ later, as its form depends on the demand system. The substantive restriction we place on demand is that $h^{(j,t)}$ is constructed such that the following equality is satisfied everywhere:

$$(1) \quad h^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2) = \alpha p_{jt} + x_{jt} \beta + \xi_{jt},$$

where α and β are parameters contained in θ_1 and $\alpha < 0$ (i.e., demand slopes down). Thus, we assume that a known function can map quantities to an index that

is linear in prices. Models that satisfy these conditions are used regularly in the empirical industrial organization literature.⁶

Equation (1) embeds two important assumptions. First is that the unknowns θ_1 , ξ_t , and η_t do not enter $h^{(j,t)}$ directly (but can enter indirectly through \mathbf{q}_t). Thus, $h^{(j,t)}$ can be constructed given observables and knowledge of θ_2 . Second is that the right-hand side of equation (1) is linear. This restriction allows us to use linear regression results to construct an analytic expression for α . Though separability of ξ_{jt} and p_{jt} is important, it is not critical that the expression be linear in prices, as we show in Appendix A.

On the supply side, we decompose prices into markups and marginal costs:

$$(2) \quad p_{jt} = \mu_{jt} + mc_{jt}.$$

Consistent with equilibrium behavior in a broad class of oligopoly models, we assume that for each (j, t) , there exists a known function $\lambda^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2)$ such that the following equation holds:

$$(3) \quad \mu_{jt} = -\frac{1}{\alpha} \lambda^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2).$$

The substantive restrictions imposed are the multiplicative separability of $1/\alpha$ and that the unknowns θ_1 , ξ_t , and η_t do not enter $\lambda^{(j,t)}$ directly (but, as with demand, they can enter indirectly through \mathbf{q}_t). Therefore, $\lambda^{(j,t)}$ can be constructed given observables and knowledge of θ_2 . In our applications and in Appendix A, we show how to construct $\lambda^{(j,t)}(\cdot)$ in a variety of specific contexts.⁷

To fix ideas, consider the canonical model of logit demand with oligopoly price competition. Quantities are given by $q_{jt} = s_{jt}M_t$, where s_{jt} is the market share of the product and M_t (contained in \mathbf{w}_t) is the size of market t . The left-hand side of equation (1) is constructed as $h^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2) \equiv \ln(s_{jt}) - \ln(s_{0t})$, where $s_0 = 1 - \sum_{k \in \mathcal{J}} s_{kt}$ is the market share of the “outside good.” With logit demand, there are no parameters in θ_2 , and market size is the only variable in \mathbf{w}_t . The $h^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2)$ function provides the utility that the average consumer would obtain from the product. Likewise, on the supply side, $\lambda^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2)$ provides the markups *measured in utils*, while dividing by $-\alpha$ obtains markups measured in units of currency. With single-product firms, the logit markup is given by $\mu_{jt} = -\frac{1}{\alpha} \frac{1}{1 - s_{jt}}$, and thus $\lambda^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2) = 1/(1 - s_{jt})$.

We initially maintain that marginal costs are constant and linear in the product characteristics:

$$(4) \quad mc_{jt} = x_{jt} \gamma + \eta_{jt},$$

⁶See the discussion of logit demand in the paragraph following equation (3). Nested logit, random coefficients logit, linear demand, and constant elasticity demand are nested within the general model or accommodated with straightforward generalizations (Appendix A).

⁷For example, with single-product Bertrand pricing and differentiable demand, we have the general expression $\mu_{jt} = -\frac{1}{\alpha} \frac{dq_{jt}/dq_{jt}}{dq_{jt}/dq_{jt}} q_{jt} = -\frac{1}{\alpha} \frac{dh^{(j,t)}}{dq_{jt}} q_{jt}$, yielding $\lambda^{(j,t)} = \frac{dh^{(j,t)}}{dq_{jt}} q_{jt}$. Appendix A provides the relevant forms for specific demand systems, addresses the construction of $\lambda^{(j,t)}$ with multiproduct firms, and covers a generalized model of oligopoly that nests both Bertrand competition in prices and Cournot competition in quantities.

where γ is contained in θ_1 . We later allow marginal costs to depend on quantities (Section IIIB).

Combining equations (2)–(4), the supply side of the model implies that the following equation is satisfied for each product j and market t :

$$(5) \quad \lambda^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2) = -\alpha p_{jt} + \alpha x_{jt} \gamma + \alpha \eta_{jt}$$

This *supply relationship* characterizes how prices and quantities respond to shifts in demand (holding fixed α and marginal costs) given the behavior of firms. It can, for example, capture optimal price-setting behavior when individual firms have market power, and unlike the supply curve for perfectly competitive firms, the supply relationship can lie above the marginal cost curve when plotted in price-quantity space.⁸ Because the supply relationship expresses (transformed) quantities as a linear function of prices and characteristics, it is the analog to the demand relationship of equation (1).

Together, equations (1) and (5) provide the conditions that jointly determine prices and quantities. The supply-side behavior captured by equation (5) does not necessarily have to correspond to equilibrium behavior, but, when it does, these equations yield equilibrium outcomes. The framework covers many of the empirical models of industrial organization (Appendix A). Nonetheless, some models are excluded. For example, in models with constant elasticity demand, one cannot construct $h^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2)$ that satisfies equation (1). It is possible, however, to find a function that satisfies a related restriction: that the right-hand side is linear in log prices, and log prices are additively separable from the demand-side structural error term. We discuss how to extend our results to this and other cases in Appendix A.

Equations (1) and (5) also illuminate potential strategies to identify θ_1 . If a characteristic x^k shifts marginal costs ($\gamma^k \neq 0$) but is excluded from demand ($\beta^k = 0$), then it is a valid instrument and can be used to estimate equation (1). A variable not in x that is correlated with η but not ξ could also be a valid instrument. Conversely, a variable that shifts demand but is excluded from marginal costs can be a valid instrument in the estimation of equation (5). Thus, both cost shifters and demand shifters can provide exogenous variation that identifies θ_1 . Similarly, “markup shifters” that create variation in $\lambda^{(j,t)}$ or $h^{(j,t)}$ can be valid instruments for either the supply side or the demand side.⁹ Another possibility, which we develop below, is to place a covariance restriction between the structural error terms from both equations.

B. Identification with Covariance Restrictions

We now consider the identification of the model using a covariance restriction, focusing on the linear parameters $(\alpha, \beta, \gamma) = \theta_1$. We assume that the econometrician

⁸The difference between marginal costs and the supply relationship is the (perceived) inframarginal loss in revenue for selling an additional unit of quantity. Bresnahan (1982) refers to the inverse of this equation, with price on the left-hand side, as the “supply relation” and notes that it generalizes to different models of firm conduct. See Appendix A6 for a figure and additional discussion.

⁹These might include functions of other products’ characteristics (Berry, Levinsohn, and Pakes 1995; Gandhi and Nevo 2021) or competitive events such as mergers, entry, or exit (e.g., Miller and Weinberg 2017).

knows θ_2 . The identification of θ_2 has been considered in other research (e.g., Berry, Levinsohn, and Pakes 1995; Berry and Haile 2024; Gandhi and Houde 2023), and we return to the prospect that covariance restrictions may identify θ_2 in the context of the empirical applications.

Stacking objects across markets, the econometrician observes vectors of prices and quantities (\mathbf{P} and \mathbf{Q}), a matrix of nonprice characteristics (\mathbf{X}), and (possibly) other observables (\mathbf{W}). Using observables, the model, and θ_2 , the econometrician can evaluate the demand and supply transformations $h^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2)$ and $\lambda^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2)$ for every (j, t) pair. We denote these values h_{jt} and λ_{jt} and treat them as observed. We use \mathbf{H} to denote the stacked JT vector for h_{jt} .

The structural error terms can be decomposed as follows:

$$(6) \quad \xi_{jt} = \xi_j + \xi_t + \Delta\xi_{jt},$$

$$(7) \quad \eta_{jt} = \eta_j + \eta_t + \Delta\eta_{jt},$$

which incorporates product-specific persistent components (e.g., higher quality or higher cost), market-specific components (greater demand in a year and/or region), and an orthogonal mean-zero residual term. Define an augmented characteristics matrix, $\tilde{\mathbf{x}}_t$, as including the K observed covariates and a full set of dummy variables for products ($J - 1$) and markets ($T - 1$). Stacking across markets, we obtain $\tilde{\mathbf{X}}$ with dimension $JT \times (K + J + T - 2)$.

We assume that the augmented characteristics are exogenous, in the sense that $E[\Delta\xi_{jt}|\tilde{\mathbf{x}}_t] = E[\Delta\eta_{jt}|\tilde{\mathbf{x}}_t] = 0$ for all $j = 1, \dots, J$, as is commonly maintained in the literature. We also assume that $E[\mathbf{v}_t'\mathbf{v}_t]$ has full rank, where $\mathbf{v}_t = [\mathbf{p}_t, \tilde{\mathbf{x}}_t]$ is a $J \times (K + J + T - 1)$ matrix that combines prices with the augmented characteristics. Two immediate implications of these assumptions are that $E[\tilde{\mathbf{x}}_t'\tilde{\mathbf{x}}_t]$ has full rank and that β and γ are trivially identified given knowledge of α , following the standard arguments for linear regression.

We now focus on the identification of the price parameter, α . We leverage a single moment (a covariance restriction), which allows us to express our results using variance and covariance terms among random variables. This notation is familiar from univariate regressions. We obtain univariate regression analogs by taking a single variable (p_{jt}), projecting it on the other characteristics ($\tilde{\mathbf{x}}_{jt}$), and then considering a regression with the residualized values as the single regressor. By the Frisch-Waugh-Lovell theorem (and exogeneity of $\tilde{\mathbf{x}}_t$), the resulting coefficient estimate is identical to one obtained in the full multivariate regression. The residuals from a regression of p_{jt} on $\tilde{\mathbf{x}}_{jt}$ are given by

$$(8) \quad \mathbf{P}^* = \mathbf{P} - \tilde{\mathbf{X}} [\tilde{\mathbf{X}}' \tilde{\mathbf{X}}]^{-1} \tilde{\mathbf{X}}' \mathbf{P},$$

and they provide the component of price that is orthogonal to characteristics and fixed effects. Later in this paper, we residualize other variables in the same fashion. Throughout, we will use the superscript $*$ to denote the residuals obtained from a regression of a variable on $\tilde{\mathbf{x}}_{jt}$.

An implication of the rank condition is that the augmented characteristics do not fully explain prices. Therefore, the unconditional variance of p_{jt}^* is positive—i.e., $\text{var}[p^*] > 0$.¹⁰ Here, and throughout the remainder of the paper, we omit jt subscripts when the variable occurs in a covariance or variance expression or in the body of the text—e.g., p refers to p_{jt} and x refers to x_{jt} . Note that p^* is distinguished from the length JT vector \mathbf{P}^* by the use of lowercase and a lack of bold font.

We now formalize our first identification result, which links the OLS estimate to the price parameter. Following the discussion above, the probability limit ($T \rightarrow \infty$) of the OLS estimate of α obtained from a regression of h on p and \tilde{x} is

$$(9) \quad \alpha^{OLS} \equiv \frac{\text{cov}[p^*, h]}{\text{var}[p^*]} = \alpha + \frac{\text{cov}[p^*, \Delta\xi]}{\text{var}[p^*]}.$$

The corresponding OLS residuals are given by $\Delta\xi^{OLS} = \mathbf{H} - \mathbf{V}[\mathbf{V}'\mathbf{V}]^{-1}\mathbf{V}'\mathbf{H}$.

We now construct a function that maps the price coefficient to a specific value for the covariance of the residual structural error terms:

PROPOSITION 1: *The probability limit of the OLS estimate can be written as a function of α , the residuals from an OLS regression, prices and quantities, and a covariance term:*

$$(10) \quad \alpha^{OLS} = \alpha - \frac{1}{\alpha + \frac{\text{cov}[p^*, \lambda]}{\text{var}[p^*]}} \frac{\text{cov}[\Delta\xi^{OLS}, \lambda]}{\text{var}[p^*]} + \alpha \frac{1}{\alpha + \frac{\text{cov}[p^*, \lambda]}{\text{var}[p^*]}} \frac{\text{cov}[\Delta\xi, \Delta\eta]}{\text{var}[p^*]}.$$

Therefore, α solves the following quadratic equation:

$$(11) \quad 0 = \alpha^2 + \left(\frac{\text{cov}[p^*, \lambda]}{\text{var}[p^*]} - \alpha^{OLS} + \frac{\text{cov}[\Delta\xi, \Delta\eta]}{\text{var}[p^*]} \right) \alpha + \left(-\alpha^{OLS} \frac{\text{cov}[p^*, \lambda]}{\text{var}[p^*]} - \frac{\text{cov}[\Delta\xi^{OLS}, \lambda]}{\text{var}[p^*]} \right).$$

All proofs are in Supplemental Appendix D. The terms in equation (11) are well defined under our rank condition and, aside from α and $\text{cov}[\Delta\xi, \Delta\eta]$, they have straightforward empirical analogs.

¹⁰This expression refers to the unconditional variance of p^* (over products and markets). The unconditional variance is defined as $\text{var}[p^*] \equiv E[(p^*)^2] - E[p^*]^2$, where the expectations are taken over markets and products. The second component is zero because we assume that x includes a constant. The empirical analog is $\frac{1}{JT} \sum_{t=1}^T \sum_{j=1}^J (p_{jt}^*)^2$.

There are at most two solutions for α for any given value of $\text{cov}[\Delta\xi, \Delta\eta]$. Further, in most empirical models, α is likely to be the lower root. The following result provides formal conditions under which this is guaranteed:

PROPOSITION 2: *The parameter α is the lower root of equation (11) if and only if*

$$(12) \quad -\frac{1}{\alpha} \text{cov}[\Delta\xi, \Delta\eta] \leq \text{cov}\left[p^*, \Delta\eta - \frac{1}{\alpha}\Delta\xi\right],$$

and, furthermore, α is the lower root of equation (11) if

$$(13) \quad 0 \leq \alpha^{OLS} \text{cov}[p^*, \lambda] + \text{cov}[\Delta\xi^{OLS}, \lambda].$$

In the first condition, it is helpful to think of $-\Delta\xi/\alpha$ as the residual demand-side structural error term, scaled so that units are equivalent to those of marginal costs (and price). If $\text{cov}[\Delta\xi, \Delta\eta] = 0$, the condition holds as long as prices tend to increase with demand and marginal costs, as occurs in most empirical models. For example, the condition holds when demand is linear. Thus, α is likely the lower root of equation (11) in most applications.

The second condition is derived using properties of the quadratic formula. Because the terms in equation (13) are constructed from data, the sufficient condition can be estimated and used to test (and possibly reject) the null hypothesis that multiple negative roots exist. Henceforth, we assume that α is the lower root of equation (11).

The implication of this result—a one-to-one function mapping α to $\text{cov}[\Delta\xi, \Delta\eta]$ —is that the price coefficient can be recovered with information about the correlation between residual demand and cost shocks in models with imperfect competition. Conversely, moments that pin down the price parameter also pin down the value of $\text{cov}[\Delta\xi, \Delta\eta]$.

C. Estimation with Covariance Restrictions

Estimation can proceed with the method of moments (a general approach) by recasting the information about the covariance term as an orthogonality condition. One possibility is that demand-side and supply-side structural error terms are uncorrelated: $\text{cov}[\Delta\xi, \Delta\eta] = 0$. Equivalently, this can be expressed as $E[\Delta\xi_{jt} \Delta\eta_{jt}] = 0$.

Under this assumption, the method of moments estimator uses the empirical analog of this condition and attempts to minimize its contribution to the objective function. For a case with one moment and one parameter, the method of moments estimate of α is given by

$$(14) \quad \hat{\alpha}^{CR} = \arg \min_{\tilde{\alpha} < 0} \left[\frac{1}{T} \frac{1}{J} \sum_t \sum_{j \in \mathcal{J}} \Delta\xi_{jt}(\tilde{\alpha}) \Delta\eta_{jt}(\tilde{\alpha}) \right]^2,$$

where $\Delta\xi_{jt}(\tilde{\alpha})$ and $\Delta\eta_{jt}(\tilde{\alpha})$ can be recovered from residualized (transformed) quantities and prices, given the candidate parameter under consideration. Some care must

be taken to ensure convergence to the lower root. The generalized method of moments (GMM) may also be used with additional moments or when estimating multiple parameters jointly, in which case the sample moment may be weighted against other components of the GMM objective function using the standard approach.

An implication of Proposition 1 is that this estimate is consistent for the price parameter, i.e., $\hat{\alpha}^{CR} \rightarrow \alpha$. This is notable because, in general, the inclusion of a moment in a method of moments approach does not imply the consistent identification of an additional parameter.¹¹ By contrast, a restriction on the covariance between the structural error terms provides identification of the price parameter a priori. Pushing the result further, one could use the analytical expression in equation (11) to directly compute the coefficient estimate.¹²

This method of moments approach is employed in our first application to estimate models with random-coefficients logit demand and Bertrand pricing. In these models, β and γ can be estimated with OLS regression once $\hat{\alpha}^{CR}$ is obtained. Döpfer et al. (forthcoming) also employ the method of moments approach and illustrate how the additional parameters in θ_2 can be estimated with GMM using the nested fixed-point approach of Berry, Levinsohn, and Pakes (1995). For each candidate $\tilde{\theta}_2$, the covariance restriction estimator is applied to obtain $\hat{\alpha}^{CR}(\tilde{\theta}_2)$. In the outer loop, an estimate of θ_2 is pinned down by micro-moments in the GMM objective function. More generally, covariance restrictions can be applied in conjunction with instruments, and additional moments allow for efficiency improvements and specification tests (e.g., Hausman 1978; Hansen 1982).

The empirical variation that identifies α is the relative variation of (transformed) quantities and prices. When $\text{cov}[\Delta\xi, \Delta\eta] = 0$, we obtain the following formal result:

PROPOSITION 3: *If $\text{cov}[\Delta\xi, \Delta\eta] = 0$, then a first-order approximation to the probability limit of the method of moments estimator is*

$$(16) \quad \alpha^{CR} \approx -\sqrt{\frac{\text{var}[h^*]}{\text{var}[p^*]}}$$

Intuition can be gleaned from the simultaneous equations representation of the model, using equations (1) and (5). Rearranging these to obtain inverse demand and inverse supply relationships, we have

$$(17a) \quad p_{jt} = \frac{1}{\alpha} h^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2) - \frac{1}{\alpha} x_{jt} \beta - \frac{1}{\alpha} \xi_{jt} \quad (\text{Demand})$$

$$(17b) \quad p_{jt} = -\frac{1}{\alpha} \lambda^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2) + x_{jt} \gamma + \eta_{jt} \quad (\text{Supply}).$$

¹¹ To highlight this, consider that Berry and Haile (2024) identify a class of moment conditions (micro-moments) that can pin down consumer heterogeneity but provides no identifying information about the price parameter.

¹² The probability limit of the coefficient estimate is given by:

$$(15) \quad \alpha^{CR} = \frac{1}{2} \left(\alpha^{OLS} - \frac{\text{cov}[p^*, \lambda]}{\text{var}[p^*]} - \sqrt{\left(\alpha^{OLS} + \frac{\text{cov}[p^*, \lambda]}{\text{var}[p^*]} \right)^2 + 4 \frac{\text{cov}[\Delta\xi^{OLS}, \lambda]}{\text{var}[p^*]}} \right),$$

which obtains from an application of the quadratic formula.

By inspection, α determines the slope of both equations. A large α corresponds to a flatter inverse demand schedule (i.e., price-sensitive consumers) and a flatter inverse supply relationship (i.e., less market power). Uncorrelated shifts in such schedules tend to generate more variation in quantity than in price. By contrast, a small α corresponds to steeper inverse demand and supply relationships, such that uncorrelated shifts generate more variation in price than quantity. Formally connecting these observations generates an approximation of the lower root based on the ratio of variances. We illustrate this argument using a numerical example in Appendix B.

D. Partial Identification: Bounds

We now show how our formal identification results can be used to construct bounds on the price parameter, which may be useful for inference when a covariance restriction along the lines of $\text{cov}[\Delta\xi, \Delta\eta] = 0$ is not plausible. We first consider bounds that utilize prior knowledge of the sign of the correlation between $\Delta\xi$ and $\Delta\eta$. Next, we show how the model and the data together may bound the price coefficient without any additional information.

In the first case, we assume that the econometrician can sign the correlation between $\Delta\xi$ and $\Delta\eta$. This situation might arise, for example, if factor prices are influenced by macroeconomic conditions, such that there is a link between the unobserved demand-side and supply-side error terms that is difficult to model explicitly. With a prior of the sign of $\text{cov}[\Delta\xi, \Delta\eta]$, bounds can be placed on α . This is because there is a one-to-one mapping between the value of $\text{cov}[\Delta\xi, \Delta\eta]$ and the lower root of equation (11):

LEMMA 1 (Monotonicity): *Under Assumptions 1 and 2, a valid lower root of equation (11) (i.e., one that is negative) is decreasing in $\text{cov}[\Delta\xi, \Delta\eta]$. The range of the function is $(0, -\infty)$.*

Thus, if idiosyncratic demand and costs are correlated, such as through capacity constraints ($\text{cov}[\Delta\xi, \Delta\eta] \geq 0$), then one-sided bounds can be placed on α . More generally, let $r(m)$ be the lower root of the quadratic in equation (11), evaluated at $\text{cov}[\Delta\xi, \Delta\eta] = m$. Then $\text{cov}[\Delta\xi, \Delta\eta] \geq m$ produces $\alpha \in (-\infty, r(m)]$, and $\text{cov}[\Delta\xi, \Delta\eta] \leq m$ produces $\alpha \in [r(m), 0)$. The lower root, $r(m)$, can be estimated with the method of moments.¹³

In the second case, it can be that some values of the price parameter are unable to rationalize the data for *any* amount of correlation between $\Delta\xi$ and $\Delta\eta$. These values can be ruled out. Thus, the demand and supply assumptions alone may be informative about the plausible range of α . Formally, this occurs when the quadratic from equation (11) does not have a lower root and thus has no valid solution for α . To see why, represent the quadratic from equation (11) as $az^2 + bz + c$. By assumption, one root is $\alpha < 0$. As $a = 1$, the quadratic is U-shaped. If $c < 0$, then the existence of a negative root is guaranteed. However, if $c > 0$, then b must be

¹³Nevo and Rosen (2012) develop conceptually similar bounds for estimation with imperfect instruments, defined as instruments that are less correlated with the structural error term than with the endogenous regressor.

positive and sufficiently large for a negative root to exist. This places restrictions on $\text{cov}[\Delta\xi, \Delta\eta]$, which is a component of b . From the monotonicity result (Lemma 1), we can use the excluded values of $\text{cov}[\Delta\xi, \Delta\eta]$ to rule out values of α .

We now state the result formally:

PROPOSITION 4 (Model-Based Bound): *The model and data alone may bound $\text{cov}[\Delta\xi, \Delta\eta]$ from below. The bound is given by*

$$\text{cov}[\Delta\xi, \Delta\eta] > \text{var}[p^*] \alpha^{OLS} - \text{cov}[p^*, \lambda] + 2\text{var}[p^*] \sqrt{\left(-\alpha^{OLS} \frac{\text{cov}[p^*, \lambda]}{\text{var}[p^*]} - \frac{\text{cov}[\Delta\xi^{OLS}, \lambda]}{\text{var}[p^*]} \right)}.$$

The bound exists if and only if the term inside the radical is nonnegative. Further, through equation (11), this lower bound on $\text{cov}[\Delta\xi, \Delta\eta]$ provides an upper bound on α .

A model-based upper bound for α is obtained by evaluating the lower root of equation (11) at the model-based bound of $\text{cov}[\Delta\xi, \Delta\eta]$. In practice, priors over the the covariance of unobserved shocks may be combined with model-based bounds to further restrict the identified set.

E. Discussion

The covariance restriction approach to estimation requires the econometrician to assess whether a covariance restriction between the structural error terms of the model is reasonable given the institutional details of the setting. It also requires the econometrician to specify the supply side in order to estimate demand parameters. We discuss these two requirements in this section.

There are settings in which the econometrician may have reason to think that the primary components of costs are uncorrelated with the drivers of demand. For example, consider demand for coffee beans in the United States. The cost of production primarily depends on weather conditions and other agricultural concerns in Brazil, Colombia, Vietnam, and Guatemala.¹⁴ Demand in the United States is largely unrelated to such factors, so it may be reasonable to assume that demand shocks are orthogonal to marginal cost shocks.

On the other hand, there are environments where an assumption that demand and cost shocks are uncorrelated would be problematic. First, when products vary in quality, we would typically expect that products with higher (unobserved) quality are produced at higher costs. If the econometrician only has the use of cross-sectional data, then fixed effects (ξ_j, η_j) cannot be used to absorb the confounding variation, and the estimates would be biased. A second setting that

¹⁴Potential supply shocks include severe frosts, high temperatures, below-average rainfall, excessive rainfall, plant diseases, pests, and fertilizer costs. See <https://apps.fas.usda.gov/psdonline/circulars/coffee.pdf>.

would generate correlation in demand and supply shocks is one in which marginal costs vary with quantity produced. A third situation is simply when, even after addressing the above issues, the econometrician anticipates that the residual shocks will be correlated.

In the above circumstances, an understanding that the structural error terms are likely correlated may also correspond to a belief that the correlation has a particular sign or, potentially, falls within a specific range of values. In such cases, the econometrician can pair this understanding with our analytical results to place bounds on the structural parameters, as discussed above. We explore this approach with an application in Section IVC.

Further, in other cases, the correlation between demand and supply shocks can be accounted for with fixed effects or by explicitly modeling features that drive the correlation. In Section IIA, we discuss an example with online retail where fixed effects control for the correlation between demand and supply conditions. Similar arguments may be made for electronic goods or consumer products sold in brick-and-mortar stores (Döpfer et al. forthcoming; De los Santos, O'Brien, and Wildenbeest 2024). Later in this paper, we provide an extension to show how a known marginal cost function can be incorporated in the model (Section IIIB). In an application to the cement industry, we find that, after upward-sloping marginal costs are specified in the supply model, the covariance restriction approach yields an estimate similar to that obtained when employing the instruments used in the literature (Section IVB).

We now discuss the requirement that the econometrician specify the supply side of the model. In practice, empirical research in industrial organization often employs supply models to calculate markups or conduct counterfactuals. For these purposes, imposing the relevant supply model when estimating demand is not a substantive additional assumption about the economic environment. However, there are cases in which demand-side estimates are of interest independently, and the researcher may be hesitant to impose a supply model in these cases. To explore misspecification bias, we perform Monte Carlo exercises with a logit demand system and a misspecified supply model, which we detail in Appendix C. If we assume Bertrand competition when the true model features joint profit maximization (i.e., perfect collusion), then the average bias in the price parameter is less than 4 percent. The potential for bias appears to be mitigated by the fact that the covariance restriction approach also uses the demand side of the model. These simulations indicate that, in some cases, the bias from supply-side misspecification may not be large.

Finally, it is worthwhile to consider these two requirements in comparison to the requirements for instrumental variables. When estimating demand with instruments (e.g., cost shifters), the econometrician does not need to fully specify a supply model. This reduces the chance that supply-side assumptions materially affect the estimates. On the other hand, the instrumental variable approach also requires an assessment about the covariance structure of unobservables, similar to the covariance restriction approach. In the case of instruments, this assessment manifests as the exclusion restriction. We explore the connection between covariance restrictions and instruments in the next section.

II. Relationship to Instrumental Variables

The covariance restriction approach to estimation interprets observed endogenous variation in quantities and prices through the lens of an economic model. This differs from the instrumental variable approach, which seeks to isolate exogenous variation in prices. It also employs different assumptions than those used to obtain “residual instruments” or Hausman instruments. We provide formal distinctions in this section. We also use Monte Carlo simulations to illustrate that a covariance restriction can outperform instruments in finite samples, especially when the instrument become weak.

A. Motivating Example

We begin with an economic model in which a traditional instrumental variable approach and the covariance restriction approach can yield consistent estimates. Consider a simplified setting in which a profit-maximizing monopolist faces linear demand in each of $t = 1, \dots, T$ markets. The demand schedule takes the form

$$(18) \quad q_{jt} = \alpha p_{jt} + \xi_j + \xi_t + \Delta\xi_{jt},$$

and marginal costs take the form

$$(19) \quad mc_{jt} = \eta_j + \eta_t + \Delta\eta_{jt},$$

where quantity demanded (q_{jt}) and price (p_{jt}) are observed (we retain the subscript j for notational consistency). Thus, the model includes product- and market-specific shocks but no observable characteristics other than price.

The key identification challenge for estimating α is that prices reflect unobserved demand shocks. In our example, equilibrium prices are given by $p_{jt} = \frac{1}{2}mc_{jt} + \frac{1}{2|\alpha|}(\xi_j + \xi_t + \Delta\xi_{jt})$. Thus, prices are higher for higher-quality products ($\text{cov}[p_{jt}, \xi_j] > 0$) and in high-demand markets ($\text{cov}[p_{jt}, \xi_t] > 0$). In settings with sufficient observations, these correlations can be controlled for with fixed effects for products and markets. However, the primary concern about price endogeneity remains if $\text{cov}[p_{jt}, \Delta\xi_{jt}] \neq 0$.

Either instrumental variables or a covariance restriction could be used to resolve price endogeneity, after accounting for fixed effects. A standard instrumental variable approach is to obtain auxiliary data about a component of costs that is orthogonal to $\Delta\xi_{jt}$ and to use it as an instrument. For illustrative purposes, suppose that one could measure $\Delta\eta_{jt}$ directly. Then, $\Delta\eta_{jt}$ would be a valid instrument for p_{jt} in the demand equation as long as $\text{cov}[\Delta\xi_{jt}, \Delta\eta_{jt}] = 0$. Alternatively, our results show that the covariance restriction $\text{cov}[\Delta\xi_{jt}, \Delta\eta_{jt}] = 0$ can be employed directly in estimation, without observing $\Delta\eta_{jt}$. Both approaches can be motivated by the same restriction on the covariance structure of unobservables.

To fix ideas, consider an online retailer that sells coffee tables made from two different materials—e.g., wicker and solid wood. Consumers may prefer one product

to another (ξ_j), and overall demand for the retailer's products may vary across markets (ξ_i). On the supply side, products vary in procurement and distribution costs (η_j), and marginal costs vary across markets due to differences in distribution networks and fuel costs (η_i). The online retailer sells the products for different prices in each market in response to these features and local demand and supply conditions, $\Delta\xi_{jt}$ and $\Delta\eta_{jt}$. For the online retailer, residual product-market variation in costs is due to the interaction of product characteristics with features of the local distribution networks. Depending on the material, similarly sized coffee tables can differ significantly in weight, which affects the amount of fuel needed for shipping. Thus, the product-market cost variation ($\Delta\eta_{jt}$) can be approximated by $weight_j \times (fuel\ cost)_t$.

In this case, local market demand has no obvious link to idiosyncratic fluctuations in distribution costs. Based on this, one could estimate demand by first obtaining data on product-level characteristics (weight) and market-level features (fuel costs), and then using the interaction of the two to generate cost-shifter instruments in a standard instrumental variable approach. When controlling for product quality and market-level demand, it would be necessary to construct a measure with idiosyncratic across-market variation by product; otherwise, the instrument would be fully absorbed by the fixed effects.¹⁵

Instead of collecting these auxiliary data to use as an instrument, one could impose the covariance restriction $cov[\Delta\xi_{jt}, \Delta\eta_{jt}] = 0$, directly following Proposition 1. The identification assumption relies on the same logic as the instrument—that idiosyncratic product-market differences in costs are orthogonal to idiosyncratic product-market differences in preferences. Similar reasoning is used by Döpfer et al. (forthcoming) when estimating demand for consumer products. After using fixed effects to account for obvious linkages between demand and costs, the residual supply-side structural error features product-specific changes in input costs and distribution costs, both of which have been exploited as instruments in recent research (e.g., Miller and Weinberg 2017; Backus, Conlon, and Sinkinson 2021). Thus, these examples—along with the applications in Section IV—demonstrate how a similar justification for the validity of instrumental variables may be used to motivate the covariance restriction approach.

B. Excluded Instruments

We now provide formal distinctions between the sets of assumptions that underlie each approach. An instrument is an observable variable that satisfies both an

¹⁵Thus, it is often not sufficient to use a measure of η_j or η_i as an instrument. In practice, the econometrician often observes only a portion of marginal costs, in which case the instrument z_{jt} can be expressed as a component of the full structural error, $\Delta\eta_{jt} = z_{jt} + \tilde{\Delta}\eta_{jt}$. The unobserved component, $\tilde{\Delta}\eta_{jt}$, may be interpreted as measurement error—for example, if the interaction of fuel costs with weight is only a first-order approximation of actual shipping costs. Isolating a component of costs is an advantage when $cov[\Delta\xi_{jt}, z_{jt}] = 0$, but $cov[\Delta\xi_{jt}, \tilde{\Delta}\eta_{jt}] \neq 0$ as the instrument still yields a consistent estimate, while the covariance restriction approach may be biased. This bias becomes small when the orthogonal component z explains a greater share of idiosyncratic cost shocks. Thus, it can be helpful to understand the components that contribute the most to the marginal cost residual, even if they are unobserved in the data.

exclusion condition and a relevance condition (e.g., Wooldridge 2010). Using the model of equations (18) and (19), these two conditions can be expressed as

$$(20a) \quad E[\Delta\xi_{jt}z_{jt}] = 0,$$

$$(20b) \quad E[p_{jt}^*z_{jt}] \neq 0,$$

where, again, p^* denotes the residuals from the linear projection of p on \tilde{x} . We focus on the case of a single instrument, z . Without loss of generality, we express z_{jt} as a component of the supply-side structural error, $\Delta\eta = z + \tilde{\Delta}\eta$, where $\tilde{\Delta}\eta$ is the remaining unobserved component.¹⁶ Condition (20a) is, in general, distinct from the covariance restriction $E[\Delta\xi_{jt}\Delta\eta_{jt}] = 0$. There are settings in which one holds while the other does not.

Under these assumptions, the probability limit of the two-stage least squares (2SLS) estimate of α , given the instrument z , is

$$(21) \quad \alpha^{IV} = \frac{\text{cov}[h^*, z]}{\text{cov}[p^*, z]} = \frac{\text{cov}[h^*, \hat{p}^*]}{\text{var}[\hat{p}^*]},$$

where \hat{p}^* is the residual from the linear projection of \hat{p} on \tilde{x} , and \hat{p} is defined as the predicted values from the linear projection of p on z and \tilde{x} .¹⁷ The first expression shows that the 2SLS estimate equals the ratio of the coefficient obtained in the reduced-form regression of h^* on z to the coefficient obtained in a first-stage regression of p^* on z . The second expression provides a reformulation that is useful for our purposes: the 2SLS estimate equals the coefficient from a regression of h^* on \hat{p}^* , the (residualized) first-stage predicted values.

The instrumental variable approach has two empirical requirements beyond what is needed for the covariance restriction approach. First, the excluded instrument must be observed in the data, as the corresponding estimators are constructed as a function of z . Second, for α^{IV} to be well defined, the relevance condition (20b) must hold so that $\text{cov}[p^*, z] \neq 0$. That is, z must add explanatory power for prices above and beyond the variables in \tilde{x} .¹⁸

Inspection of equation (21) provides a way to compare the relevance condition to the assumptions of the covariance restriction approach. In the equation, the denominator of the second expression has the term $\text{var}[\hat{p}^*]$, which is similar to the term $\text{var}[p^*]$ in the denominator of equation (11). Each must be greater than zero for

¹⁶The instrument must be linearly independent from \tilde{x} ; otherwise, equation (20b), which can alternatively be expressed as $E[p_{jt}^*z_{jt}]$, would be violated. For the model in equations (18) and (19), \tilde{x} contains a constant and dummy variables that capture market and product fixed effects. When the model incorporates product-market varying characteristics (x), a variable x^k is a candidate supply-side instrument when $\gamma_k \neq 0$ and $\beta_k = 0$. In this case, we can simply redefine \tilde{x} such that it does not include x^k . This maps z to the above interpretation (as a component of $\Delta\eta$) and avoids the need for more cumbersome notation in this section. Note that omitting x^k from \tilde{x} and including it in $\Delta\eta$ does not affect the validity of a covariance restriction when $z = x^k$ satisfies condition (20a).

¹⁷In terms of data, $\hat{P}^* = \hat{P} - \tilde{X}[\tilde{X}'\tilde{X}]^{-1}\tilde{X}'\hat{P}$ and $\hat{P} = \tilde{Z}[\tilde{Z}'\tilde{Z}]^{-1}\tilde{Z}'\hat{P}$ for $\tilde{Z} = [\tilde{X}\tilde{Z}]$.

¹⁸An implication is that choosing the instrument $z_{jt} = \Delta\eta_{jt}$ will not work when there is no residual variation in costs, i.e., $E[p_{jt}^*\Delta\eta_{jt}] = 0$. However, in this case, the covariance restriction approach can still obtain point identification.

the corresponding estimator to be well defined.¹⁹ The condition that $\text{var}[\hat{p}^*] > 0$ is stronger than the condition that $\text{var}[p^*] > 0$. In fact, when the former holds, the latter is strictly implied.²⁰ From the Frisch-Waugh-Lovell theorem, we have $\text{var}[p^*] \geq \text{var}[\hat{p}^*]$. Thus, the covariance restriction approach can proceed under weaker assumptions about the conditional variance of prices.

Moreover, even if condition (20b) is satisfied in the limit, the instrumental variable estimator can exhibit asymptotic bias in finite samples (e.g., Keane and Neal 2024). This is an important consideration in practice, and many papers have been devoted to addressing the “weak instrument” problem when this condition is tenuously satisfied (e.g., Bound, Jaeger, and Baker 1995; Staiger and Stock 1997; Stock and Yogo 2005). The covariance restriction approach can side-step this issue because all of the residual variation in p is used to construct the estimate. In applications, $\text{var}[\hat{p}^*]$ may be small relative to $\text{var}[p^*]$ if z is constructed from one of many components of costs.

Thus far, we have framed a valid instrument as satisfying $E[\Delta\xi_{jt}z_{jt}] = 0$. However, in the context of the model, an alternative is to use an instrument that satisfies $E[\Delta\eta_{jt}z_{jt}] = 0$. Such an instrument could be taken from the demand side of the model or constructed based on markup shifters, as in Berry, Levinsohn, and Pakes (1995). This approach uses the supply side of the model—specifically, the first-order conditions of equation (5)—to estimate the price parameter. We highlight this possibility because it shows how either demand-side variation or supply-side variation can be used to pin down the price parameter, so long as the appropriate exclusion restriction can be applied. The covariance restriction approach exploits both supply-side variation and demand-side variation implicitly, using a single restriction.

C. Residual Instruments

We now compare our approach to the “residual instruments” approach, which uses a covariance restriction to achieve identification of simultaneous equations. Wooldridge (2010, 258) focuses on the case of two linear equations.²¹ To connect to this analysis, we rewrite demand and supply from our motivating example to obtain the following:

$$(22a) \quad q_{jt} = \alpha_1 p_{jt} + \xi_j + \xi_t + \Delta\xi_{jt} \quad (\text{Demand}),$$

$$(22b) \quad q_{jt} = -\alpha_2 p_{jt} + \alpha_2(\eta_j + \eta_t + \Delta\eta_{jt}) \quad (\text{Supply}),$$

where, again, the supply relationship comes directly from the monopolist’s first-order condition for profit maximization. Because we first consider the general

¹⁹As we discuss in Section IB, $\text{var}[p^*] > 0$ is an implication of $E[\mathbf{v}'_i \mathbf{v}_i]$ having full rank.

²⁰To see this, note that, by way of counterexample, if \bar{x} perfectly predicts p such that $\text{var}[p^*] = 0$, there is no residual variation for an instrument to explain.

²¹The discussion in Wooldridge (2010) builds on a substantial literature on covariance restrictions in linear simultaneous equation models (e.g., Hausman and Taylor 1983; Hausman, Newey, and Taylor 1987). More recent research generalizes these results (Matzkin 2016).

simultaneous equations approach, the price coefficients are allowed to vary across the equations.

The residual instruments approach to identification relies on the following moments for some observable z :

$$(23a) \quad E[z_{jt} \Delta \eta_{jt}] = 0,$$

$$(23b) \quad E[z_{jt} p_{jt}^*] \neq 0,$$

$$(23c) \quad E[\Delta \xi_{jt} \Delta \eta_{jt}] = 0.$$

The first two moments are analogous to instrumental variables conditions (20a) and (20b), though in this case z is a demand-side instrument and excluded from the supply equation. Here we assume that z is an observed component of the demand-side structural error, $\Delta \xi = z + \tilde{\Delta} \xi$. The third moment corresponds to the covariance restriction that $\text{cov}[\Delta \xi, \Delta \eta] = 0$. These three moments identify the model.

Estimation is typically described as proceeding in two steps. In the first step, a standard instrumental variables regression with z as the instrument is used to identify α_2 and the supply equation. In the second step, the demand equation (22a) is estimated using the residuals from the first step, $\widehat{\Delta \eta}$, as an instrument. The residuals meet the necessary exclusion restriction in the second step due to the covariance restriction. Thus, in this framework, covariance restrictions have been interpreted as providing excluded instruments.²² As the above demonstrates, in addition to the covariance restriction, the residual instruments approach requires an additional two conditions about the existence of a valid instrument. In practice, it must also be the case that z and $\widehat{\Delta \eta}$ explain p to a substantial degree; otherwise, each step could suffer from the weak instruments problem.

By contrast, our approach to estimation with covariance restrictions recognizes a theoretical connection between the slopes of demand and supply that is implied by the economic model: $\alpha_1 = \alpha_2 = \alpha$. In this case, α is point identified with only one restriction: $\text{cov}[\Delta \xi, \Delta \eta] = 0$. With this approach, there is no need for an excluded instrument (z).²³ This provides a path for identification under a different set of assumptions while also avoiding the finite-sample challenges of weak instruments.

D. Hausman Instruments

A number of articles in industrial organization have relied on prices in related markets as instruments in demand estimation (Gandhi and Nevo 2021). Typically, “related markets” refer to distinct geographic areas. In that setting, the price of a

²²This interpretation has been influential. For example, McFadden states in lecture notes (dated 1999) that, “Even covariance matrix restrictions can be used in constructing instruments. For example, if you know that the disturbance in an equation you are trying to estimate is uncorrelated with the disturbance in another equation, then you can use a consistently estimated residual from the second equation as an instrument.” See https://eml.berkeley.edu/~mcfadden/e240b_f01/ch6.pdf.

²³However, as discussed earlier, either a demand-side or supply-side instrument would also identify the model.

product in some market s can be a valid instrument for the price of the same product in market t if marginal costs are correlated across markets (e.g., due to shared production facilities) but demand is not. Such instruments are often referred to as “Hausman instruments” due to their use in Hausman (1996).

Here, we assume that the econometrician employs product and market fixed effects.²⁴ For Hausman instruments to be valid, there must exist pairs of markets $\{(t, s)\}$ such that

$$(24a) \quad E[\Delta\xi_{jt} \Delta\xi_{js}] = 0,$$

$$(24b) \quad E[\Delta\xi_{jt} \Delta\eta_{js}] = 0,$$

$$(24c) \quad E[\Delta\eta_{jt} \Delta\eta_{js}] \neq 0.$$

Condition (24a) states that the demand-side error terms are uncorrelated across markets, condition (24b) states that the demand-side error term of one market is uncorrelated with the supply-side error term in another market, and condition (24c) states that supply-side error terms are correlated across markets.

If these conditions are satisfied, then p_{js} is a valid (excluded) instrument for p_{jt} in the demand equation. Thus, analogous to the residual instruments approach in the previous section, this approach leverages assumptions about the correlation structure of unobservables to generate excluded instruments. This approach can suffer from the weak instruments problem, as we discuss in Section IIB.

Similar to our approach, the Hausman instruments are justified by assumptions about the correlational structure of demand and cost shocks. The set of assumptions are distinct: the three above conditions could be met when $E[\Delta\xi_{jt} \Delta\eta_{jt}] \neq 0$, or conditions (24a)–(24c) may not be satisfied while $E[\Delta\xi_{jt} \Delta\eta_{jt}] = 0$. One advantage of the covariance restriction approach is that only a single restriction is required. A second advantage is that it avoids the potential weak instrument problem that Hausman instruments may be subject to in finite samples.

Note that the use of p_{js} as an excluded instrument does not necessarily exploit all of the variation implied by conditions (24a)–(24c). As we discuss in Section IC, it is possible to use such orthogonality conditions directly with the method of moments rather than using them to justify an observed variable (p_{js}) as an instrument. We explore generalizations of our approach with these and other covariance restrictions in Section IIIA.

E. Finite Sample Comparison

We use Monte Carlo simulations to illustrate the finite sample performance of covariance restrictions relative to excluded instruments. For excluded instruments, we consider both traditional supply-side instruments and the demand-side instruments discussed in Section IIB. In both cases, we assume that the available instrument is

²⁴In practice, researchers sometimes use Hausman instruments that reflect variation across all markets, in which case it is necessary to assume that $\xi_j = 0$.

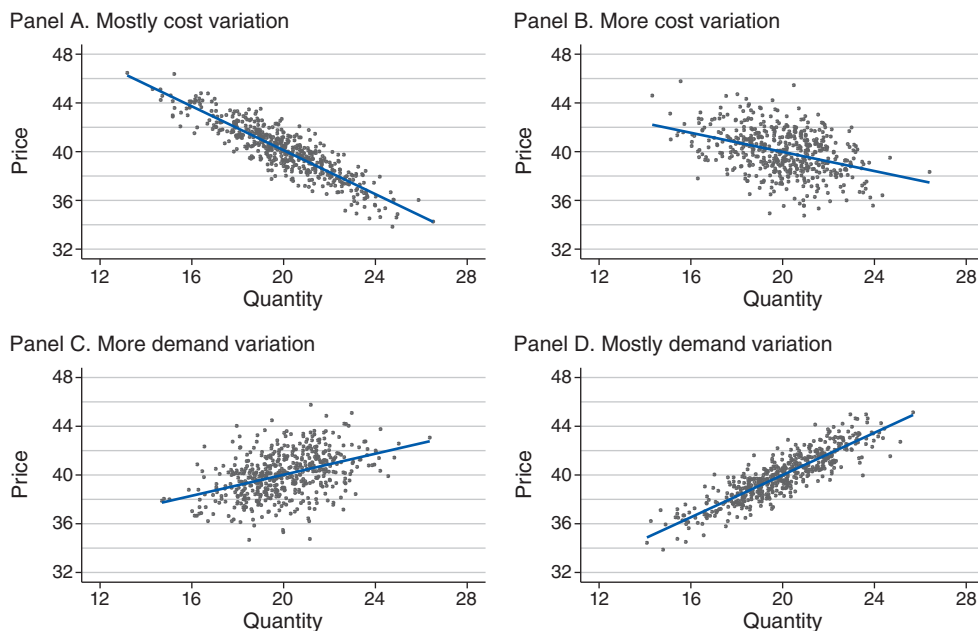


FIGURE 1. PRICES AND QUANTITIES IN THE MONOPOLY MODEL

Notes: This figure displays equilibrium prices and quantities under four different specifications for the distribution of unobserved shocks to demand and marginal costs. The line in each figure indicates the slope obtained by OLS regression.

fully efficient, in that it captures all of the relevant exogenous variation—i.e., $z = \Delta\eta$ for a supply-side instrument and $z = \Delta\xi$ for the demand-side instrument. That is, we allow the econometrician to fully observe cost shocks when estimating demand or demand shocks when estimating the supply relationship.

For our simulations, we use the monopoly model of equations (18) and (19) for a single product. We normalize the time fixed effects (ξ_t, η_t) to zero, and we set $\alpha = 1$, $\xi_j = 60$, and $\eta_j = 20$. We assume that $\Delta\xi_{jt}$ and $\Delta\eta_{jt}$ are mean-zero independent normal distributions with standard deviations σ_ξ and σ_η . We consider four specifications: (i) $\sigma_\xi = 1$ and $\sigma_\eta = 4$, (ii) $\sigma_\xi = 2$ and $\sigma_\eta = 3$, (iii) $\sigma_\xi = 3$ and $\sigma_\eta = 2$, and (iv) $\sigma_\xi = 4$ and $\sigma_\eta = 1$. Moving from (i) to (iv), demand-side variation increases and supply-side variation decreases.

As is well known, if both supply and demand variation are present, then equilibrium outcomes provide a “cloud” of data points that need not correspond to the demand curve. To illustrate, we present one simulation of 500 observations from each specification in Figure 1, along with the fit of an OLS regression of quantity on price. The expected values for the OLS estimator in each scenario are -0.882 , -0.385 , 0.385 , and 0.882 . With greater demand-side variation, the endogeneity bias is larger.²⁵ Using the condition that $\text{cov}[\Delta\xi, \Delta\eta] = 0$, we can correct for endogeneity and

²⁵Inspection of Figure 1 further suggests there may be a connection between OLS bias and goodness of fit. Indeed, starting with equation (16), a few lines of additional algebra obtain $\alpha \approx -|\alpha^{OLS}|/\sqrt{R^2}$, where R^2 is from the residual OLS regression of h^* on p^* . The approximation is exact with linear demand. This reformulation fails

TABLE 1—SMALL-SAMPLE PROPERTIES: RELATIVE VARIATION IN DEMAND AND SUPPLY SHOCKS

Observations	$\text{var}[\eta] \gg \text{var}[\xi]$		$\text{var}[\eta] > \text{var}[\xi]$		$\text{var}[\eta] < \text{var}[\xi]$		$\text{var}[\eta] \ll \text{var}[\xi]$	
	(i)		(ii)		(iii)		(iv)	
<i>Panel A. Covariance restrictions</i>								
25	-1.006	(0.100)	-1.019	(0.198)	-1.017	(0.199)	-1.004	(0.102)
50	-1.003	(0.069)	-1.010	(0.134)	-1.008	(0.136)	-1.002	(0.069)
100	-1.002	(0.047)	-1.005	(0.094)	-1.006	(0.095)	-1.001	(0.049)
500	-1.000	(0.021)	-1.001	(0.041)	-1.001	(0.041)	-1.001	(0.021)
<i>Panel B. Supply shifters (IV-1)</i>								
25	-1.007	(0.107)	-1.044	(0.314)	-1.273	(3.399)	-0.820	(13.379)
50	-1.003	(0.074)	-1.021	(0.202)	-1.112	(0.623)	-1.369	(10.661)
100	-1.002	(0.050)	-1.010	(0.137)	-1.057	(0.345)	-1.509	(6.676)
500	-1.000	(0.022)	-1.003	(0.060)	-1.009	(0.138)	-1.080	(0.444)
<i>Panel C. Demand shifters (IV-2)</i>								
25	-0.835	(12.357)	-1.303	(3.667)	-1.040	(0.315)	-1.005	(0.109)
50	-1.299	(11.845)	-1.116	(0.561)	-1.018	(0.203)	-1.003	(0.073)
100	-1.557	(6.517)	-1.052	(0.343)	-1.012	(0.139)	-1.001	(0.052)
500	-1.071	(0.420)	-1.011	(0.137)	-1.002	(0.060)	-1.001	(0.023)

Notes: Results are based on 10,000 simulations of data for each specification and number of observations. The demand curve is $q_{jt} = \alpha p_{jt} + \xi_{jt}$ with $\alpha = -1$ and $\xi_{jt} = \xi_j + \Delta\xi_{jt}$. Marginal costs are $mc_{jt} = \eta_{jt}$, where $\eta_{jt} = \eta_j + \Delta\eta_{jt}$. We consider a single product ($j = 1$) and vary the number of markets/observations from 25 to 500. IV-1 estimates are calculated using 2SLS with cost shocks ($\Delta\eta$) as an instrument in the demand equation. Analogously, IV-2 estimates are calculated using 2SLS with demand shocks ($\Delta\xi$) as an instrument in the supply relationship. We specify $\Delta\xi$ and $\Delta\eta$ as mean-zero independent normal distributions with standard deviations σ_ξ and σ_η . We consider four specifications: (i) $\sigma_\xi = 1$ and $\sigma_\eta = 4$, (ii) $\sigma_\xi = 2$ and $\sigma_\eta = 3$, (iii) $\sigma_\xi = 3$ and $\sigma_\eta = 2$, and (iv) $\sigma_\xi = 4$ and $\sigma_\eta = 1$. Moving from (i) to (iv), demand-side variation increases and supply-side variation decreases.

construct a covariance restriction estimator for α as $-\sqrt{\text{var}[q]/\text{var}[p]}$, because in this context the approximation of equation (16) is exact. The corresponding estimates for each simulation in Figure 1 are -0.98 , -1.00 , -0.98 , and -1.01 , close to the true parameter values.

We consider sample sizes of 25, 50, 100, and 500 observations. For each specification and sample size, we randomly draw 10,000 datasets, and for each we estimate the model with a covariance restriction, with a supply-side instrument, and with a demand-side instrument. For the covariance restriction, we estimate α using $-\sqrt{\text{var}[q]/\text{var}[p]}$, as above. For a supply-side instrument, we estimate demand with 2SLS, using the cost shock $\Delta\eta$ as the instrument. For the demand-side instrument, we estimate the supply relationship with 2SLS, using $\Delta\xi$ as the instrument. All three approaches rely on the same orthogonality condition: $E[\Delta\xi_{jt}\Delta\eta_{jt}] = 0$.

Table 1 provides the mean and (empirical) standard error of the point estimates for each specification and approach.²⁶ Panel A shows that the covariance restriction approach to estimation yields estimates that are consistently close to the true value. Panel B shows that with supply-side instruments, small-sample bias becomes substantial with smaller datasets and less variance in the cost shock. This is due to

if $R^2 = 0$, but numerical results indicate robustness for values of R^2 that are approximately zero. We thank Peter Hull for suggesting this connection.

²⁶To avoid outliers arising from the weak instrument problem, we bound the estimates of α on the range $[-100, 100]$. For specifications that suffer from weak instruments, this biases the standard errors toward zero. This affects specifications where the estimated standard error is greater than one—i.e., in 8 of 48 specifications.

a weak instrument—for example, the mean first-stage F -statistics in specification (iv) are 2.8, 4.2, 7.4, and 32.1 for markets with 25, 50, 100, and 500 observations, respectively. Panel C shows that with demand-side instruments, small-sample bias becomes substantial with smaller datasets and less variance in the demand shock, which is also due to a weak instruments problem.

Thus, in settings where instruments perform poorly, a covariance restriction may still provide a precise estimate when the assumptions about the environment are correct. In our simulations, the covariance restriction has smaller standard deviations than either instrumental variables strategy. Intuitively, even though exogenous variation is not observed, the covariance restriction approach exploits how variation from demand and supply is reflected in equilibrium prices and quantities.

III. Extensions

In this section, we provide two extensions. First, we discuss a broader class of potential covariance restrictions. Second, we consider the case of marginal cost functions. We use both extensions in the applications of Section IV.

A. Generalized Covariance Restrictions

Thus far, our analysis has focused on covariance restrictions between own demand and cost shocks. Our results demonstrate that this moment is expected to generate a consistent estimate of the price parameter. We now consider different covariance restrictions that generalize the approach. Though these other covariance restrictions do not provide a similar guarantee of point identification, they may work well in certain settings. Additionally, they may pin down other parameters of interest (e.g., those in θ_2) in addition to the price parameter.

Consider the assumptions (24a)–(24c) that are required for the Hausman instruments. Rather than using these assumptions to motivate the use of an instrument, the assumptions could be employed directly in a method of moments estimator, where, as in equation (14), the estimated residuals are generated from an econometric model for a candidate parameter. This approach has an advantage over the Hausman instruments approach in that the estimator would utilize all of the variation implied by the identifying moments.

Alternatively, it may be reasonable to assume that the variance of the demand shock does not depend on the level of the cost shock, and vice versa, which generates the moments $E_{jt}[\Delta\xi_{jt}^2 \Delta\eta_{jt}]$ and $E_{jt}[\Delta\xi_{jt} \Delta\eta_{jt}^2]$. Or it may be reasonable to assume that average shocks are uncorrelated across groups of products, i.e., $E_{gt}[\overline{\Delta\xi_{gt}} \overline{\Delta\eta_{gt}}] = 0$, where $\overline{\Delta\xi_{gt}}$ and $\overline{\Delta\eta_{gt}}$ are the mean demand and cost shocks for products in group g .

Finally, it may be useful to consider cross-product covariance restrictions, i.e.,

$$(25) \quad E_t[\Delta\xi_{jt} \Delta\eta_{kt}] = 0 \quad \forall j \neq k.$$

These restrictions state that the demand shock for product j is uncorrelated with the cost shock for product k . The expectation in equation (25) can be taken over t and

k to obtain J restrictions, or over markets (as written above) to obtain $J \times (J - 1)$ restrictions.²⁷

We also note that our analytical results do not require a covariance value of zero. Proposition 1 can be used to construct consistent estimates for any ς for which $E[\Delta\xi_{jt}\Delta\eta_{jt}] = \varsigma$. In certain situations, it may be possible to employ an estimate of the correlation in demand and cost shocks to identify the price parameter. For example, Berry, Levinsohn, and Pakes (1995) report that this correlation is 0.17. For a similar empirical setting, it may be reasonable to invoke Proposition 1 to obtain an estimate of α conditional on this value.

B. Marginal Cost Functions

If marginal cost depends on quantity, then an assumption that $\text{cov}[\Delta\xi, \Delta\eta] = 0$ may not be credible in the baseline model of Section I. Here, we demonstrate how it is possible impose additional structure to control for the relationship for quantity and marginal costs and obtain identification. Consider the case in which marginal costs can be expressed as the following function:

$$(26) \quad mc_{jt} = x_{jt}\gamma + g(q_{jt}, \mathbf{w}_t; \tau) + \eta_{jt},$$

where $g(\cdot)$ is a function that depends on quantity, the data in \mathbf{w}_t , and the vector of parameters τ that is contained in θ_2 . The supply relationship becomes

$$(27) \quad \lambda^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2) = -\alpha p_{jt} + \alpha x_{jt}\gamma + \alpha g(q_{jt}, \mathbf{w}_t; \tau) + \alpha \eta_{jt}.$$

In this augmented model, both markup adjustments and varying marginal costs contribute to price endogeneity. If the econometrician omits $g(\cdot)$ from the model, then the residual cost shock is $\widetilde{\Delta\eta}_{jt} \equiv g(q_{jt}, \mathbf{w}_t; \tau) + \Delta\eta_{jt}$. Then, $\text{cov}[\Delta\xi, \Delta\eta] = 0$ does not imply that $\text{cov}[\Delta\xi, \widetilde{\Delta\eta}] = 0$, and the approach of Section I may not produce consistent estimates.

However, tracing through the steps developed in Section IB, we can show that α is identified by the covariance restriction $\text{cov}[\Delta\xi, \Delta\eta] = 0$ for any value of τ , given knowledge of $g(\cdot)$. In that scenario, $g(\cdot)$ can be calculated from the data, given τ . Let g_{jt} (or simply g) denote the values of $g(\cdot)$ for each j and t , given τ . The OLS regression of h on p and \tilde{x} yields a price coefficient with the following probability limit:

$$(28) \quad \alpha^{OLS} = \alpha - \frac{1}{\alpha} \frac{\text{cov}[\Delta\xi, \lambda]}{\text{var}[p^*]} + \frac{\text{cov}[\Delta\xi, g]}{\text{var}[p^*]} + \frac{\text{cov}[\Delta\xi, \Delta\eta]}{\text{var}[p^*]}.$$

This equation can be reformulated such that the demand-side error term, $\Delta\xi$, is replaced with the probability limit of OLS residuals, $\Delta\xi^{OLS}$, creating an analog to

²⁷ Similarly, the own-product restrictions may be assumed to hold separately by product, providing J restrictions.

equation (10). Rearranging terms and assuming $\text{cov}[\Delta\xi, \Delta\eta] = 0$ then yields an analog to equation (11):

COROLLARY 1: *If marginal costs take the semilinear form of equation (26) and $\text{cov}[\Delta\xi, \Delta\eta] = 0$, then α solves the following quadratic equation:*

$$0 = \left(1 - \frac{\text{cov}[p^*, g]}{\text{var}[p^*]}\right) \alpha^2 + \left(\frac{\text{cov}[p^*, \lambda]}{\text{var}[p^*]} - \alpha^{OLS} + \alpha^{OLS} \frac{\text{cov}[p^*, g]}{\text{var}[p^*]} + \frac{\text{cov}[\Delta\xi^{OLS}, g]}{\text{var}[p^*]}\right) \alpha + \left(-\alpha^{OLS} \frac{\text{cov}[p^*, \lambda]}{\text{var}[p^*]} - \frac{\text{cov}[\Delta\xi^{OLS}, \lambda]}{\text{var}[p^*]}\right).$$

Given θ_2 and, thus, h and g , there are at most two solutions for α , and the lower root can be targeted in most applications. The method of moments can be used to jointly estimate α and τ , using the covariance restriction to identify α and auxiliary moments to identify τ . The auxiliary moments can consist of excluded instruments or the generalized covariance restrictions discussed above. We explore such a cost function approach to estimation in an application to cement (Section IVB).

IV. Empirical Applications

We provide three empirical applications to demonstrate how covariance restrictions can inform inference. The three settings—RTE cereals, cement, and airlines—differ in a variety of ways that influence our implementation. With RTE cereals, we proceed with estimation under $\text{cov}[\Delta\xi, \Delta\eta] = 0$, assuming constant marginal costs and using fixed effects to absorb potentially confounding variation, as discussed in Section II. With cement, capacity constraints imply that marginal costs can increase with quantities. We follow an approach developed in the literature and model this effect explicitly, after which $\text{cov}[\Delta\xi, \Delta\eta] = 0$ becomes credible (as in Section IIIB). Finally, with airlines, the relationship between demand shocks and prices can be complicated; instead of modeling it directly, we apply a bounds approach (as in Section ID).

A. Ready-to-Eat Cereals

We choose RTE cereals for our first application because, with panel data and appropriate fixed effects, a covariance assumption appears credible, for reasons that we explain below. Furthermore, it allows us to develop the covariance restriction approach to estimation in the context of the random coefficients logit demand model (Berry, Levinsohn, and Pakes 1995). We use the pseudo-real cereals data of Nevo (2000) and compare estimates obtained with a covariance restriction to

those obtained with the provided instruments. There are 24 products, 47 cities, and 2 quarters.²⁸

Let the indirect utility that consumer i receives from product j in market t (a combination of a quarter and a city) be given by

$$(29) \quad u_{ijt} = \delta_{jt}(x_{jt}, p_{jt}, \xi_{jt}; \theta_1) + \phi_{ijt}(x_{jt}, p_{jt}, \nu_i, D_i; \theta_2) + \epsilon_{ijt},$$

where δ_{jt} denotes a common component and ϕ_{ijt} provides consumer-specific utility as a function of data and θ_2 . These components are specified as

$$\delta_{jt} = \alpha p_{jt} + x_{jt} \beta + \xi_j + \Delta \xi_{jt},$$

$$\phi_{ijt} = [p_{jt} x_{jt}] (\Pi D_i + \Sigma \nu_i),$$

such that consumer-specific utility is linear in the parameters $[\alpha_i; \beta_i] = [\alpha; \beta] + \Pi D_i + \Sigma \nu_i$. Consumers can pick any one of the inside goods ($j = 1, \dots, 24$) or an outside good ($j = 0$) that provides indirect utility of $u_{i0t} = \epsilon_{i0t}$. Mapping to the notation of Section I, α and β are contained in θ_1 , Π and Σ are in θ_2 , and the data \mathbf{x}_t , \mathbf{p}_t , $\{D_i\}$, and $\{\nu_i\}$ are included in \mathbf{w}_t .

Demand is expressed in terms of market shares.²⁹ Let N_t denote the number of consumers in a market, which we assume to be large. Under the assumption that unobserved shock ϵ_{ijt} is distributed iid type 1 extreme value, the market share for product j in market t ($j \neq 0$) can be written as

$$(30) \quad s_{jt} = \varsigma_{jt}(\boldsymbol{\delta}_t, \mathbf{w}_t; \theta_2) \equiv \frac{1}{N_t} \sum_{i=1}^{N_t} \frac{\exp\{\delta_{jt} + \phi_{ijt}(\mathbf{w}_t; \theta_2)\}}{1 + \sum_{k=1}^J \exp\{\delta_{kt} + \phi_{ikt}(\mathbf{w}_t; \theta_2)\}}.$$

Stacking across products, we obtain the vector-valued equation $\mathbf{s}_t = \boldsymbol{\varsigma}_t(\boldsymbol{\delta}_t, \mathbf{w}_t; \theta_2)$. Because each $\varsigma_{jt}(\cdot)$ is strictly increasing in δ_{jt} , this equation can be inverted to obtain $\boldsymbol{\delta}_t(\mathbf{s}_t, \mathbf{w}_t; \theta_2)$. Each element δ_{jt} corresponds to $h_{jt} = h^{(j,t)}(\mathbf{s}_t, \mathbf{w}_t; \theta_2)$ in the notation of Section I. Thus, in implementation, the contraction mapping of Berry, Levinsohn, and Pakes (1995) can obtain the $J \times 1$ vector \mathbf{h}_t , given \mathbf{s}_t , \mathbf{w}_t , and θ_2 .

On the supply side of the model, marginal costs are given by

$$(31) \quad mc_{jt} = \eta_j + \Delta \eta_{jt}.$$

Prices are set by multiproduct firms in Bertrand competition. Following the general results for multiproduct firms in Appendix A5, equilibrium markups take the form specified in equation (3), such that $\lambda^{(j,t)}$ can be expressed as a function of \mathbf{s}_t , \mathbf{w}_t , and

²⁸ See also Dubé, Fox, and Su (2012); Knittel and Metaxoglou (2014); and Conlon and Gortmaker (2020). We focus on the “restricted” specification of Conlon and Gortmaker (2020), which addresses a multicollinearity problem by imposing that the parameter on $Price \times Income^2$ takes a value of zero. We refer readers to Nevo (2000) for details on the data.

²⁹ Following the discussion in Section IIA, market shares can be converted into quantities using market size: $q_{jt} = s_{jt} M_t$. Thus, if the data have quantities rather than market shares, then M_t must be in \mathbf{w}_t .

TABLE 2—POINT ESTIMATES FOR RTE CEREAL

Variable	Means	Standard deviations	Interactions with demographics		
			Income	Age	Child
<i>Panel A. Available instruments</i>					
Price	−32.019 (2.304)	1.803 (0.920)	4.187 (4.638)	−	11.755 (5.198)
Constant	−	0.375 (0.120)	3.101 (1.054)	1.198 (1.048)	−
Sugar	−	0.004 (0.012)	−0.190 (0.035)	0.028 (0.032)	−
Mushy	−	0.086 (0.193)	1.495 (0.648)	−1.539 (1.107)	−
<i>Panel B. Covariance restrictions</i>					
Price	−36.230 (1.216)	1.098 (1.069)	14.345 (1.677)	−	26.905 (1.384)
Constant	−	0.051 (0.230)	−0.156 (0.286)	1.072 (0.239)	−
Sugar	−	0.003 (0.014)	−0.084 (0.018)	−0.004 (0.010)	−
Mushy	−	0.130 (0.161)	0.301 (0.196)	−0.085 (0.103)	−

Notes: This table reports point estimates for the random-coefficients logit demand system estimated using the Nevo (2000) dataset. Panel A employs the available instruments. Panel B employs covariance restrictions.

θ_2 . For additional details, see Appendix A2. Note that we follow Nevo (2000) and exclude market fixed effects from demand and supply.

We use the covariance restriction $\text{cov}[\Delta\xi_{jt}, \Delta\eta_{jt}] = 0$ in estimation. The supply-side structural error term incorporates some of the cost-shifter instruments that have been used in the recent literature, including time-varying, product-specific shipping costs (Miller and Weinberg 2017) and the time-varying prices of product-specific ingredients (Backus, Conlon, and Sinkinson 2021). Given the fixed effects, these cost shifters can be conceptualized as providing the variation that is exploited in estimation. Furthermore, it may be reasonable to think that the marginal costs of consumer products are roughly constant, as often maintained in the literature (Villas-Boas 2007; Chevalier, Kashyap, and Rossi 2003; Hendel and Nevo 2013; Miller and Weinberg 2017; Backus, Conlon, and Sinkinson 2021).

The parameters for estimation include α and β (contained in θ_1) and also Π and Σ (contained in θ_2). Identification of θ_1 is obtained under the exogeneity of \tilde{x} as conditions (1) and (5) are satisfied. Additional identifying assumptions are needed for θ_2 . Some recent applications use micro-moments constructed from the observed behavior of individual consumers (e.g., Backus, Conlon, and Sinkinson 2021; Döppler et al. forthcoming) or survey data on the products that consumers view to be their “second choice” (e.g., Grieco, Murry, and Yurukoglu 2023). Both of these strategies identify Π and Σ but not the price parameter (Berry and Haile 2024). This separability allows for a two-step approach to estimation in which the price parameter is estimated after the other parameters. An alternative strategy is to use instruments constructed from competitor characteristics

(e.g., Berry, Levinsohn, and Pakes 1995; Gandhi and Houde 2023) to identify the additional parameters. As none of these options are available to us given the data and specification, we pursue an alternative approach based on a generalization of the covariance restriction assumption.

Specifically, we extend to all cross-product pairs the assumption that residual demand and cost shocks are uncorrelated, such that $\text{cov}[\Delta\xi_{jt}, \Delta\eta_{kt}] = 0$ for all j, k . The joint restrictions are valid if the demand shock of each product is orthogonal to its own marginal cost shock and to those of all other products. As there are 24 products in each market, the full covariance matrix of demand and cost shocks provides a sufficient number of moments to estimate the 12 nonlinear parameters in the specification.

Table 2 summarizes the results of estimation based on instruments (panel A) and covariance restrictions (panel B). In the application, D_i consists of measures of log income and age and an indicator for whether the individual is a child; ν_i is drawn as a standard normal to capture unobserved demographics; and x contains a constant, sugar content, and an indicator for whether the cereal gets mushy with milk. Both identification strategies yield similar mean own-price demand elasticities: -3.70 with instruments and -3.61 with covariance restrictions. Overall, the different approaches produce similar patterns for the coefficients. Most of the point estimates under covariance restrictions fall in the 95 percent confidence intervals implied by the specification with instruments, including that of the mean price parameter. The standard errors are noticeably smaller with covariance restrictions, which likely reflects that the covariance restriction approach to estimation more fully exploits the variation in the data. We conclude that in this setting—where a covariance restriction appears credible—estimation with covariance restrictions and with instruments produce similar results.

B. *The Cement Industry*

Our second empirical application considers a setting in which marginal costs increase with output. We build on the marginal cost specification from Section IIIB, in which the upward-sloping part of the cost function can be modeled explicitly. To illustrate, we consider the setting and data of Fowlie, Reguant, and Ryan (2016)—henceforth, FRR—which examines market power in the cement industry.

The data are a balanced panel of 520 region-year observations for 20 regions over 1984–2009, with the regions corresponding to selected urban areas in the United States. There are an average of 4.65 cement firms located in each region-year.³⁰

The model features Cournot competition among cement plants facing capacity constraints. Let the market demand curve in region r and year t have a logit form:

$$(32) \quad h^{(rt)}(Q_{rt}, \mathbf{w}_{rt}; \theta_2) \equiv \ln(Q_{rt}) - \ln(M_r - Q_{rt}) = x_{rt}\beta + \alpha p_{rt} + \xi_r + \Delta\xi_{rt},$$

³⁰See FRR for details on the data, which are available for download as part of the replication package.

where $Q_{rt} = \sum_{j \in \mathcal{J}} q_{jrt}$ is total quantity produced in the region-year, M_r is the “market size” of the region (and is contained in θ_2), and the only characteristic in x_{rt} is a constant.³¹ Further, we allow marginal costs to vary with quantity according to

$$(33) \quad mc_{jrt} = x_{rt}\gamma + \alpha p_{rt} + g_{jrt}(q_{jrt}, \mathbf{w}_{rt}; \tau) + \Delta\eta_{jrt}.$$

We follow FRR in the specification of the cost function and included fixed effects. In particular, we assume that g is a “hockey stick” function, $g_{jrt}(q_{jrt}, \mathbf{w}_{rt}; \tau) \equiv 2\tau \mathbf{1}\{q_{jrt}/k_{jr} > 0.9\}(q_{jrt}/k_{jr} - 0.9)$, where k_{jr} and q_{jrt}/k_{jr} are capacity and utilization, respectively. Marginal costs are constant if utilization is less than 90 percent. Above this threshold, marginal costs increase linearly in quantities at a rate determined by $\tau \geq 0$. Mapping to the notation of Section I, τ is a scalar element of θ_2 , and \mathbf{w}_{rt} includes k_{jr} for each firm (in addition to M_r).

As in our baseline model, correlation between price and the demand-side structural error term can arise due to both markup adjustments and the effect of demand on marginal costs. However, due to the presence of $g_{jrt}(\cdot)$ in the cost function, the latter channel exists even under the covariance restriction $\text{cov}[\Delta\xi_{rt}, \overline{\Delta\eta_{rt}}] = 0$, where $\overline{\Delta\eta_{rt}} = \frac{1}{J}\Delta\eta_{jrt}$. If $g_{jrt}(\cdot)$ is known or can be identified with additional moments, then the covariance restriction is sufficient to resolve price endogeneity, as the model informs the markup adjustments. In estimation, we maintain the covariance restriction at the market level.

Our demand and supply framework of equations (1) and (5) admits Cournot competition. As only market-level price and cost measures are observed, one must use the mean firm-level quantity $\bar{q}_{rt} = \frac{1}{J}Q_{rt}$ to obtain an expression for mean market-level markups and λ . In particular, when firms compete in quantities, we obtain $\lambda_{rt} = \frac{1}{J} \frac{dh}{dq} Q_{rt}$. Section IIIB establishes the necessary results to incorporate increasing marginal costs into our framework. In our implementation, we assume that $\psi = 800$, such that our $g_{jrt}(\cdot)$ function is close to what is used in FRR, and then use a method of moments estimator.

In the context of the cement industry, whether the covariance restriction is reasonable may depend primarily on the relationship between construction activity (a shifter of unobserved demand) and the prices of coal and electricity (determinants of unobserved marginal cost). There is a theoretical basis to assume orthogonality: for example, if coal suppliers have limited market power and roughly constant marginal costs, then coal prices should not respond much to demand for cement. Indeed, this is the identification argument of FRR, as coal and electricity prices are included in the set of excluded instruments. Consistent with this, the time series of coal prices between 1980 and 2010 is not obviously correlated with macroeconomic conditions (e.g., Miller, Osborne, and Sheu 2017).

We find that the covariance restriction approach yields a demand elasticity of -1.15 , with a standard error of 0.16 .³² This is close to the 2SLS estimate of -1.07

³¹ We use logit demand rather than the constant elasticity demand of FRR to allow for adjustable markups. The 2SLS results are unaffected by this choice. In our implementation, we assume $M_r = 2 \times \max_r\{Q_{rt}\}$.

³² We obtain bootstrapped standard errors based on 200 random samples constructed by drawing from the data with replacement.

(standard error 0.18) that we obtain using the FRR instruments: coal prices, natural gas prices, electricity prices, and wage rates. That the two approaches generate similar estimates may reflect that the identifying assumptions themselves are similar, with the main difference being whether the cost shifters are treated as observed (2SLS) or unobserved (covariance restrictions). By contrast, we obtain a demand elasticity of -0.47 and a standard error of 0.14 using OLS. If we use the covariance restriction approach without accounting for the presence of $g_{jmt}(\cdot)$, we obtain a demand elasticity of -0.90 and a standard error of 0.13 , which is in between the OLS and 2SLS estimates and demonstrates how accounting for marginal cost functions can matter for estimation results.

C. The Airline Industry

In our third empirical application, we examine demand for airline travel using the setting and data of Aguirregabiria and Ho (2012)—henceforth, AH. The economics of the industry suggest that the covariance restriction $\text{cov}[\Delta\xi, \Delta\eta] = 0$ would not be credible. The reason is that airlines bear an opportunity cost when they sell a seat, because that seat can no longer be sold at a higher price to another passenger (Williams 2022). Thus, all else equal, greater demand generates higher marginal costs, inclusive of the opportunity cost. Absent a model of these opportunity costs, it would be difficult to achieve point identification using the covariance restriction. Instead, we illustrate how to proceed in such cases by constructing bounds on the parameters of interest.

The data feature 2,950 city pairs in the United States observed over the four quarters of 2004. A market is a city-pair-quarter combination. Products are classified into the following groups: nonstop flights, one-stop flights, and the outside good. On average, there are 7.98 products per market (not including the outside good), including 2.04 nonstop products.³³

The nested logit demand system can be expressed as

$$(34) \quad h^{(j,mt)}(\mathbf{q}_{mt}, \mathbf{w}_{mt}; \theta_2) \equiv \ln s_{jmt} - \ln s_{0mt} - \sigma \ln \bar{s}_{jmt|g} \\ = \alpha p_{jmt} + x_{jmt} \beta + \xi_{a(j)} + \xi_{mt} + \Delta \xi_{jmt},$$

where s_{jmt} is the market share of product j in market m in period t . The conditional market share, $\bar{s}_{j|g} = s_j / \sum_{k \in g} s_k$, is the the choice probability of product j given that its “group” of products, g , is selected. The outside good is indexed as $j = 0$. Consumer preferences vary by airline ($\xi_{a(j)}$) and by route quarter (ξ_{mt}), which we account for with fixed effects. Higher values of σ increase within-group consumer substitution relative to across-group substitution.³⁴

³³We thank Victor Aguirregabiria for providing the data, which derive from the DB1B database maintained by the Department of Transportation. Replication is not exact, because the sample differs somewhat from what is used in the AH publication and because we employ a different set of fixed effects in estimation.

³⁴The covariates include an indicator for nonstop itineraries, the distance between the origin and destination cities, and a measure of the airline’s “hub sizes” at the origin and destination cities. In estimation, we include airline

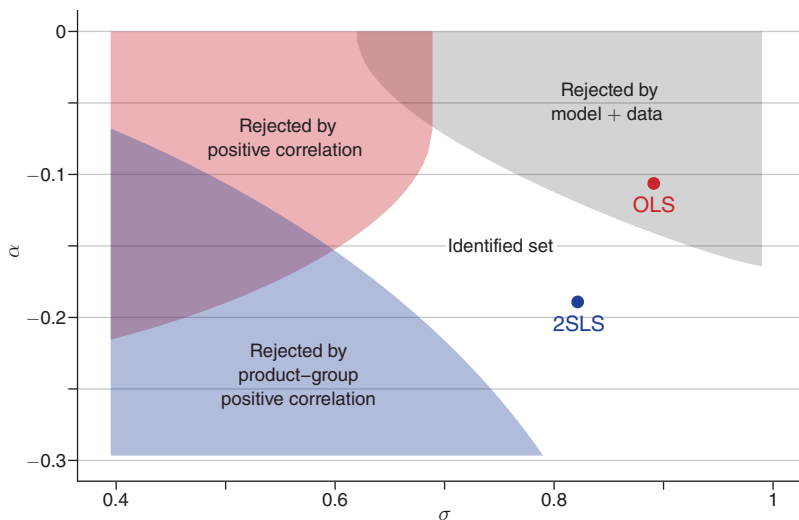


FIGURE 2. ANALYSIS OF BOUNDS IN THE AIRLINES INDUSTRY

Notes: This figure displays candidate parameter values for (σ, α) . The gray region indicates the set of parameters that cannot generate the observed data from the assumptions of the model. The red region indicates the set of parameters that generate $\text{cov}[\Delta\xi, \Delta\eta] < 0$, and the blue region indicates parameters that generate $\text{cov}[\overline{\Delta\xi}, \overline{\Delta\eta}] < 0$. The identified set is obtained by rejecting values in the above regions under the assumption of (weakly) positive correlation. For context, the OLS and the 2SLS estimates are plotted. The parameter σ can only take values on $[0, 1)$.

We impose three sets of bounds. First, we assume that product-level demand and cost shocks are weakly positive, i.e., $\text{cov}[\Delta\xi_{jmt}, \Delta\eta_{jmt}] \geq 0$, based on the role of opportunity costs in the industry. Second, if the correlation in product-level shocks is weakly positive, it is also reasonable to assume that the correlation in group-level shocks is also weakly positive. That is, overall demand for nonstop flights in a market may drive up the opportunity costs for nonstop flights. Thus, building on Section IIIA, we apply the group-level inequality

$$(35) \quad E_{gmt} [\overline{\Delta\xi}_{gmt} \overline{\Delta\eta}_{gmt}] \geq 0,$$

where $\overline{\Delta\xi}_{gmt} = \frac{1}{|g|} \sum_{j \in g} \Delta\xi_{jmt}$ and $\overline{\Delta\eta}_{gmt} = \frac{1}{|g|} \sum_{j \in g} \Delta\eta_{jmt}$ are the mean demand and cost shocks within a group-market-period. Finally, we combine these bounds with the model-based bounds developed in Section ID. We then construct an identified set by rejecting values of the parameters (α, σ) that fail to generate the data or that deliver negative correlations between costs and demand.

Figure 2 displays the rejected regions based on both the model and our assumptions on unobserved shocks. The gray region corresponds to the parameter values rejected by the model-based bounds; the model itself rejects some values of α if

fixed effects and route \times quarter fixed effects. Market size, which determines the market share of the outside good, is equal to the total population in the origin and destination cities combined.

$\sigma \geq 0.62$. As σ becomes larger, a more negative α is required to rationalize the data. The red region corresponds to parameter values that generate negative correlation between demand and supply shocks; this region is rejected under the prior that $\text{cov}\left[\frac{\Delta\xi_{jmt}}{\Delta\xi_{gmt}}, \frac{\Delta\eta_{jmt}}{\Delta\eta_{gmt}}\right] \geq 0$. The blue region provides the corresponding set for the prior $\text{cov}\left[\frac{\Delta\xi_{gmt}}{\Delta\xi_{jmt}}, \frac{\Delta\eta_{gmt}}{\Delta\eta_{jmt}}\right] \geq 0$ and is similarly rejected.

The three regions overlap, but no region is a subset of another. The nonrejected values provide the identified set. We rule out values of σ less than 0.599 for any value of α , as these lower values cannot generate positive correlation in both product-level and product-group-level shocks. Thus, the bounds serve to reject the logit model ($\sigma = 0$) in favor of nested logit, even with relatively limited information about the covariance of shocks.

Similarly, we obtain an upper bound on α of -0.067 across all values of σ . Combined, these bounds indicate that the mean own-price elasticity is less than -0.537 . For context, we plot the OLS and the 2SLS estimates in Figure 2. The OLS estimate falls in a rejected region and can be ruled out by the model alone. The 2SLS estimate falls within the identified set. This result is not mechanical, as these point estimates are generated with non-nested assumptions.

V. Conclusion

We have shown that restrictions on the covariance between unobserved demand and cost shocks can resolve price endogeneity and allow for consistent estimation in models of imperfect competition. Unlike the method of instrumental variables, the covariance restriction approach does not require the econometrician to construct an observed variable that is correlated with an exogenous portion of price. Instead, the endogenous variation in quantity and price is interpreted through the lens of the model to recover the structural parameters. Our three empirical applications demonstrate how the covariance restriction approach can be used in different environments.

More broadly, our analysis shows how imposing a supply-side model provides feasible paths to identification. Our formal results illustrate how demand-side instruments, in addition to covariance restrictions, can be sufficient to resolve price endogeneity. We also establish model-free bounds, in which the model and the data jointly can reject certain values of the price parameter, without the need for additional identifying assumptions. Conditional on meeting these bounds, there is typically a unique mapping between the price coefficient and the covariance of demand and costs shocks. In such settings, the covariance term can act as a free parameter to rationalize different values of the price coefficient.

The appeal of the covariance restriction approach relative to alternatives depends on data availability and the institutional details of the industry under study. In cases where cost shifters are observed, instrumental variables can recover demand parameters with only an informal understanding of supply. By contrast, the covariance restriction approach leverages all of the observed variation in prices and quantities, but it requires a formal supply-side model. Our results provide paths to identification that may facilitate research in areas for which strong supply-side instruments are unavailable.

APPENDIX A. DEMAND SYSTEM APPLICATIONS

The demand system of equation (1) is sufficiently flexible to nest monopolistic competition with linear demand (e.g., as in the motivating example) and the discrete choice demand models that support much of the empirical research in industrial organization. The demand assumption can also be modified to allow for semilinearity in a transformation of prices, $f(p_{jt})$:

$$(A1) \quad h^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2) = \alpha f(p_{jt}) + x_{jt}\beta + \xi_{jt}.$$

Under this modified assumption, it is possible to employ a method of moments approach to estimate the structural parameters. In certain cases, it is straightforward to extend our analytical results.

For example, when $f(p_{jt}) = \ln p_{jt}$, we can obtain our identification results under a modified assumption about the structure of costs. The optimal price in such demand systems takes the multiplicative form $p_{jt} = \mu_{jt} \cdot mc_{jt}$, where μ_{jt} is a markup that reflects demand parameters and (possibly) demand shocks. Assume that log marginal costs are linear in characteristics, such that $\ln mc_{jt} = x_{jt}\gamma + \eta_{jt}$. As in Section IB, consider the augmented exogenous characteristics \tilde{x} to include a full set of dummy variables for products and markets. The probability limit ($T \rightarrow \infty$) of the OLS estimate of α obtained from a regression of h on $\ln p$ and \tilde{x} is given by

$$(A2) \quad \alpha^{OLS} = \alpha + \frac{\text{cov}[\ln \mu, \Delta \xi]}{\text{var}[\ln p^*]} + \frac{\text{cov}[\Delta \eta, \Delta \xi]}{\text{var}[\ln p^*]}.$$

This expression is analogous to equation (1). Therefore, the results developed in this paper extend in a straightforward manner.

We provide some typical examples below for single-product firms with Bertrand competition. We then show how multiproduct firms and other models of competition fit within the framework of Section I.

A1. *Nested Logit Demand*

Following the exposition of Cardell (1997), let the firms be grouped into $g = 0, 1, \dots, G$ mutually exclusive and exhaustive sets, and denote the set of firms in group g as \mathcal{J}_g . An outside good, indexed by $j = 0$, is the only member of group 0. Then the left-hand side of equation (1) takes the form

$$h^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2) \equiv \ln(s_{jt}) - \ln(s_{0t}) - \sigma \ln(\bar{s}_{j|g,t}),$$

where $\bar{s}_{j|g,t} = \frac{s_{jt}}{\sum_{j \in \mathcal{J}_g} s_{jt}}$ is the market share of firm j within its group. Shares are obtained by dividing quantities by the market size M_t . The market size and group identities are contained in \mathbf{w}_t . The parameter $\sigma \in [0, 1)$ is the only element of θ_2 , and

it determines the extent to which consumers substitute disproportionately among firms within the same group. If $\sigma = 0$, then the logit model obtains.

For single-product firms, the first-order condition for profit maximization $p_{jt} - mc_{jt} = -\frac{1}{dq_{jt}/dp_{jt}}q_{jt}$ can be expressed as

$$(A3) \quad \mu_{jt} = -\frac{1}{ds_{jt}/dp_{jt}}s_{jt}.$$

For equation (3) to hold, it must be that $\lambda^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2) = -\alpha \mu_{jt}$. We can solve for $\lambda^{(j,t)}$ in the nested logit model by taking the derivatives of s with respect to p to obtain

$$(A4) \quad \lambda^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2) = \frac{1}{\frac{1}{1-\sigma} - s_{jt} - \frac{\sigma}{1-\sigma} \bar{s}_{j|g,t}}.$$

Thus, $\lambda^{(j,t)}$ can be expressed as a function of \mathbf{q}_t and \mathbf{w}_t and θ_2 . In our third application, we use the nested logit model to estimate bounds on the structural parameters (Section IVC).

A2. Random Coefficients Logit Demand

In our application in Section IVA, we develop the underlying indirect utility model of the random coefficients logit model, following Berry (1994) and Nevo (2000, 2001). Here, we provide some additional results using the notation from that section.

The probability with which consumer i selects product j ($j \neq 0$) is

$$(A5) \quad \rho_{ijt}(\boldsymbol{\delta}_t, \mathbf{w}_t; \theta_2) \equiv \frac{\exp\{\delta_{jt} + \phi_{ijt}(\mathbf{w}_t; \theta_2)\}}{1 + \sum_{k=1}^J \exp\{\delta_{kt} + \phi_{ikt}(\mathbf{w}_t; \theta_2)\}},$$

which is obtained under the assumption that ϵ is distributed iid type 1 extreme value. Equation (30) aggregates choice probabilities across consumers. This calculation of expected market shares converges to observed shares as $N_t \rightarrow \infty$. In implementation, equation (30) is often approximated by summing across a number of simulated consumers, with each simulated consumer being characterized by a set of demographics $\{D_i, \nu_i\}$.

Consider a candidate parameter vector $\tilde{\theta}_2$ that includes Π and Σ . Given $\tilde{\theta}_2$, Berry, Levinsohn, and Pakes (1995) prove that a contraction mapping recovers the $J \times 1$ vector $\tilde{\boldsymbol{\delta}}_t(\mathbf{s}_t, \mathbf{w}_t, \tilde{\theta}_2)$ such that the choice probabilities implied by the model match the market shares observed in the data. This “mean valuation” vector is equivalent to the vector \mathbf{h}_t in our notation. These vectors can be stacked to obtain the full $JT \times 1$ vector \mathbf{H} in a single procedure.

The supply restriction from equation (3) is satisfied when (multiproduct) firms compete by setting prices, following our more general results for differentiated-products Bertrand in Appendix A5. For example, for the special case with single-product

firms and no random coefficient on price ($\alpha = \alpha_i \forall i$), the Bertrand-Nash equilibrium yields

$$(A6) \quad \lambda^{(j,t)}(\mathbf{s}_t, \mathbf{w}_t; \theta_2) = \frac{s_{jt}}{\frac{1}{N_t} \sum_i \rho_{ijt} (1 - \rho_{ijt})},$$

where the denominators integrate over the (product of) consumer-specific choice probabilities. From an econometric standpoint, λ is free from the price parameter α because it depends only on market shares and consumer-specific choice probabilities, $\rho_{ijt}(\delta_t, \mathbf{w}_t; \theta_2)$. As discussed above, δ_t can be obtained as a function of shares.

More complicated versions of $\lambda^{(j,t)}$ can be constructed numerically; however, this step is not necessary as estimation can proceed by implementing the covariance restriction directly using the method of moments. Confirming the restrictions on h and λ ensures identification of the price coefficient and other linear parameters, conditional on $\tilde{\theta}_2$. The identification of θ_2 is a distinct issue that has received a great deal of attention from theoretical and applied research (e.g., Waldfogel 2003; Romeo 2013; Berry and Haile 2014; Gandhi and Houde 2023; Miller and Weinberg 2017). We demonstrate how to estimate these parameters using additional covariance restrictions in the application in Section IVA.

A3. Constant Elasticity Demand

With the modified demand assumption of equation (A1), the constant elasticity of substitution (CES) demand model of Dixit and Stiglitz (1977) can be incorporated:

$$h^{(j,t)}(\mathbf{q}_t, \mathbf{w}_t; \theta_2) \equiv \ln(q_{jt}/\bar{q}_t) = \beta + \alpha \ln(p_{jt}/\Pi_t) + \xi_{jt},$$

where \bar{q}_t is an observed demand shifter, Π_t is a price index, and α provides the constant elasticity of demand. In our notation, \bar{q}_t and Π_t are contained in \mathbf{w}_t .

This model is often used in empirical research on international trade and firm productivity (e.g., De Loecker 2011; Doraszelski and Jaumandreeu 2013). Due to the constant elasticity, profit maximization and uncorrelatedness imply $\text{cov}[p, \xi] = 0$, and OLS produces unbiased estimates of the demand parameters when marginal costs are constant.³⁵ Indeed, this is an excellent illustration of our basic argument: so long as the data-generating process is sufficiently well understood, it is possible to characterize the bias of OLS estimates.

A4. Other Demand Systems

The demand assumption in equation (1) accommodates many rich demand systems. Consider the linear demand system, $q_{jt} = \beta_j + \sum_k \alpha_{jk} p_k + \xi_{jt}$, which

³⁵The international trade literature following Feenstra (1994) considers nonconstant marginal costs, which requires an additional restriction. See Section IVB for an extension of our methodology to nonconstant marginal costs.

sometimes appears in identification proofs (e.g., Nevo 1998) but is seldom applied empirically due to the large number of price coefficients.

The system can be formulated such that $h^{(j,t)}(\mathbf{q}, \mathbf{w}; \theta_2) \equiv q_{jt} - \sum_{k \neq j} \alpha_{jk} p_{kt}$. In this demand system, other prices (p_{kt}) are elements of \mathbf{w}_t , and the cross-product price coefficients $\alpha_{jk} (k \neq j)$ are elements of θ_2 . In addition to the own-product uncorrelatedness restrictions that could identify α_{jj} , one could impose cross-product covariance restrictions to identify $\alpha_{jk} (k \neq j)$. We discuss these cross-product covariance restrictions in the first application (Section IVA). A similar approach could be used with the almost ideal demand system of Deaton and Muellbauer (1980).

A5. Multiproduct Firms with Bertrand Competition

We illustrate how our framework more generally incorporates multiproduct firms with the case of Bertrand pricing. For this setting, we assume that the derivatives $\partial q_j / \partial p_k$ exist and that \mathbf{D} , the $J \times J$ matrix of derivatives $\partial q_j / \partial p_k \forall j, k \in \mathcal{J}$, is invertible. For matrix elements, let j index rows and k index columns. The market subscript, t , is omitted to simplify notation.

We begin by establishing properties of demand using the restriction $h^{(j)}(\mathbf{q}, \mathbf{w}; \theta_2) = \alpha p_j + x_j \beta + \xi_j$ from equation (1). Taking the derivative with respect to p_j (holding fixed \mathbf{x} , ξ , and $p_k \forall k \neq j$), we obtain

$$(A7) \quad \sum_k \frac{\partial h^{(j)}}{\partial q_k} \frac{\partial q_k}{\partial p_j} = \alpha.$$

Similarly, we obtain $\sum_k \frac{\partial h^{(j)}}{\partial q_k} \frac{\partial q_k}{\partial p_\ell} = 0 \forall \ell \neq j$. These restrictions on demand admit the expression

$$(A8) \quad \mathcal{H} \mathbf{D} = \alpha \mathbf{I} \Rightarrow \mathbf{D} = \alpha \mathcal{H}^{-1}.$$

where \mathcal{H} denotes the $J \times J$ matrix of derivatives $\partial h^{(j)} / \partial q_k \forall j, k \in \mathcal{J}$, \mathbf{D} is defined as above, and \mathbf{I} is a $J \times J$ identity matrix. Because $h^{(j)}$ is a known function of \mathbf{q} , \mathbf{w} , and θ_2 , its derivatives with respect to q_k and, thus, each element of \mathcal{H} can be calculated, and \mathcal{H}^{-1} can be solved for. We let $\mathbf{A} \equiv \mathcal{H}^{-1}$ such that \mathbf{A} denotes the matrix of demand derivatives up to the scalar α .

We now turn to the supply side of the model. Let K^m denote the set of products owned by multiproduct firm m . When the firm sets prices on each of its products to maximize joint profits, there are $|K^m|$ first-order conditions, which can be expressed as

$$\sum_{k \in K^m} \frac{\partial q_k}{\partial p_j} (p_k - mc_k) = -q_j \quad \forall j \in K^m.$$

Let \mathbf{D}^m denote the matrix of derivatives $\partial q_j / \partial p_k \forall j, k \in K^m$, and let \mathbf{q}^m and $\boldsymbol{\mu}^m$ denote the stacked vector of q_j and $\mu_j = p_j - mc_j$ for $j \in K^m$. Stacking the first-order conditions yields

$$(A9) \quad \mathbf{D}^m \boldsymbol{\mu}^m = -\mathbf{q}^m,$$




FIGURE A1. SUPPLY RELATIONSHIP

Notes: The figure plots an illustrative example of demand (D), marginal costs (MC), and the supply relationship described in the paper (SR). The supply relationship can be interpreted as the opportunity cost to the firm of selling an additional unit. The opportunity cost is the sum of the marginal cost and the inframarginal losses of lowering price. The equilibrium price (P^*) is determined by the intersection of D and SR .

and solving for markups, μ^m , we obtain

$$(A10) \quad \mu^m = -(\mathbf{D}^m)^{-1} \mathbf{q}^m.$$

Equation (3) requires the construction of λ such that $\mu^m = -\frac{1}{\alpha} \lambda^m$. In the case of multiproduct firms with Bertrand competition, it immediately follows that

$$(A11) \quad \lambda^m = \alpha (\mathbf{D}^m)^{-1} \mathbf{q}^m.$$

Following the conditions on demand above, we have $\mathbf{D} = \alpha \mathbf{A}$, and taking the corresponding element-by-element minor, we have $\mathbf{D}^m = \alpha \mathbf{A}^m$. This yields

$$(A12) \quad \lambda^m = \alpha (\alpha \mathbf{A}^m)^{-1} \mathbf{q}^m,$$

$$(A13) \quad \lambda^m = (\mathbf{A}^m)^{-1} \mathbf{q}^m.$$

Because \mathbf{A}^m is a known function of \mathbf{q} , \mathbf{w} , and θ_2 , λ^m can be constructed of the same arguments. Thus, we see that multiproduct Bertrand fits in the class of models specified by equation (3).

A6. Alternative Models of Competition

Our restriction on additive markups from equation (3) applies to a broad set of competitive assumptions. Consider, for example, Nash competition among

profit-maximizing firms that have a single choice variable, a , and constant marginal costs. The individual firm's objective function is

$$\max_{a_j | a_i, i \neq j} [p_j(a) - mc_j] q_j(a).$$

This generalized model of Nash competition nests Bertrand ($a = p$) and Cournot ($a = q$). The first-order condition, holding fixed the actions of the other firms, is given by

$$p_j(a) = mc_j - \frac{p'_j(a)}{q'_j(a)} q_j(a).$$

In equilibrium, we obtain the structural decomposition $p = mc + \mu$, where μ incorporates the structure of demand and its parameters. This decomposition provides a restriction on how prices move with demand shocks, aiding identification. Using restrictions about demand, such as those imposed by equation (1), one can construct the appropriate form of $\lambda^{(j,i)}(\cdot)$ and solve for the price coefficient. Related first-order conditions can be obtained in other contexts, such as consistent conjectures.

Bresnahan (1982) refers to the above equation as the "supply relation" and notes that it generalizes to many different forms of conduct. Figure A1 plots the supply relationship along with the demand curve for an illustrative setting. The supply relationship lies above the marginal cost curve, and the difference is given by the inframarginal loss in revenue for selling an additional unit (i.e., the gap between price and marginal revenue). As the inframarginal loss has an opportunity cost interpretation, the supply relation can be conceptualized as the sum of the marginal cost curve and the firm's opportunity cost curve, with the latter incorporating any market power that the firm has. The equilibrium price is determined by the intersection of demand and the supply relationship. This is equivalent to the equilibrium price that obtains if the firm sets the price to equate marginal revenue and marginal cost.

APPENDIX B. EMPIRICAL VARIATION IN PRICES AND QUANTITIES

In Section IC, we state that the empirical variation that identifies the price parameter (α) under a covariance restriction relates to the relative variation in (transformed) quantities and prices. In this Appendix, we provide a numerical example for illustrative purposes, using the monopoly model of equations (18) and (19).

We assume that demand and marginal cost are given by

$$q_{jt} = 60 + \alpha p_{jt} + \Delta \xi_{jt} \text{ and } mc_{jt} = 20 + \Delta \eta_{jt}.$$

We consider four values of α : -0.2 , -0.5 , -1.0 , and -1.5 . As the slope of the inverse demand relationship is $1/\alpha$ and that of the inverse supply relationship is $-1/\alpha$, the more negative values for α generate flatter inverse demand/supply relationships. We let $\Delta \eta \sim 2.5 \times N(0, 1)$ and let $\Delta \xi \sim \frac{2.5}{\alpha} \times N(0, 1)$, with α affecting the support

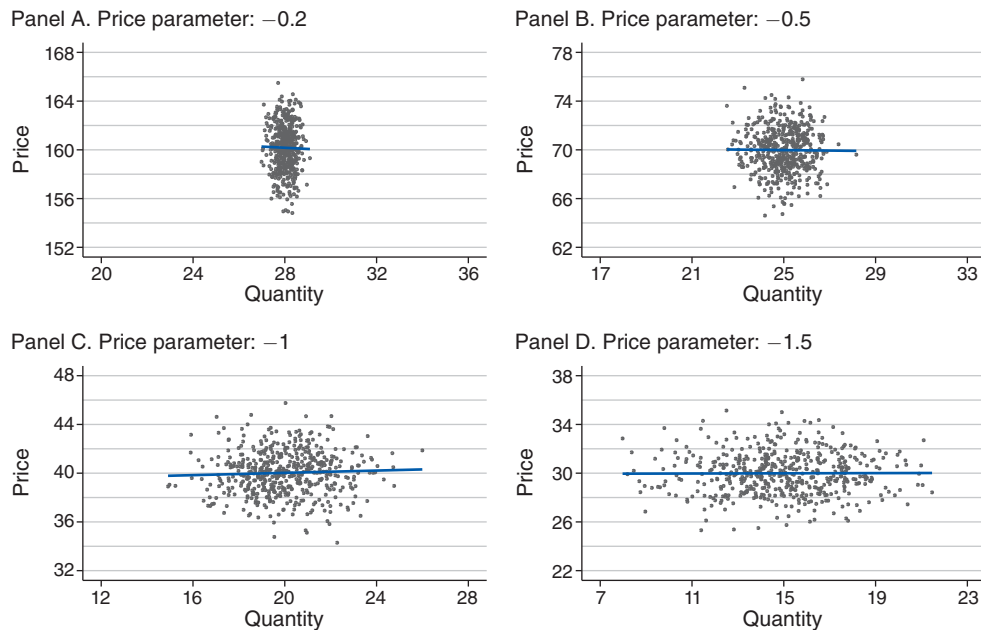


FIGURE B1. RELATIVE VARIATION IN PRICES AND QUANTITIES WITH DIFFERENT PRICE PARAMETERS

Notes: This figure displays equilibrium prices and quantities with four different values of the price parameter. The supports of $\Delta\xi$ and $\Delta\eta$ are selected so that the supply-side and demand-side variation is balanced and that the equilibrium price-quantity pairs form a cloud with no apparent slope. The line in each figure indicates the slope obtained by OLS regression.

of $\Delta\xi$ so that both variables have the same support if measured in the same units. We then take 500 draws on these demand and cost shocks for each of the four price parameters and compute the equilibrium prices and quantities.

Figure B1 shows the results. The four panels correspond to the four values of α . The panels have comparable scales (16 units by 16 units) and are recentered along the x - and y -axes. In each, the equilibrium price-quantity pairs form a cloud with no apparent slope. The reason is that supports of $\Delta\xi$ and $\Delta\eta$ are selected so that the supply-side and demand-side variation is balanced. (Figure 1 illustrates how the clouds would slope down if the supply shocks dominate and up if the demand shocks dominate.) The four panels in Figure B1 illustrate that a more negative value of α leads to greater variation in quantities relative to variation in prices. The reason is that the inverse demand and supply relationships are flatter, and uncorrelated shifts in flatter inverse supply and demand relationships produce more variation in quantities, all else equal. Intuitively, then, it should be possible to compare the relative variance of quantity and price to learn about the price coefficient. Proposition 3 formalizes this result.

For the monopoly model that we use for this numerical example—which features linear demand and constant marginal cost—the approximation provided in Proposition 3 is exact and simplifies to $\alpha^{CR} = -\sqrt{\text{var}[q]}/\sqrt{\text{var}[p]}$. Calculating the implied estimate for each scatterplot, we obtain -0.19 , -0.50 , -0.98 , and -1.54 , which are close to the respective values of the parameters used to generate the data.

APPENDIX C. SUPPLY-SIDE MISSPECIFICATION

To illustrate how supply-side misspecification can affect the performance of the estimators, we simulate duopoly markets in which the standard assumption of Bertrand price competition may not match the data-generating process.³⁶ We assume that the demand system is logit, providing consumers with a differentiated discrete choice, and we allow them to select an outside option in addition to a product from each firm. The quantity demanded of firm j in market t is

$$q_{jt} = \frac{\exp(2 - p_{jt} + \Delta\xi_{jt})}{1 + \sum_{k=j,i} \exp(2 - p_{kt} + \Delta\xi_{kt})}.$$

On the supply side, marginal costs are $c_{kt} = \Delta\eta_{kt}$ ($k = j, i$). Firm j sets price to maximize $\pi_j + \kappa \pi_i$, and likewise for firm i , where $\kappa \in [0, 1]$ is a conduct parameter (e.g., Miller and Weinberg 2017). The first-order conditions take the form

$$\begin{bmatrix} p_j \\ p_i \end{bmatrix} = \begin{bmatrix} c_j \\ c_i \end{bmatrix} - \left[\begin{pmatrix} 1 & \kappa \\ \kappa & 1 \end{pmatrix} \circ \left(\frac{\partial q}{\partial p} \right)^T \right]^{-1} \begin{bmatrix} q_j \\ q_i \end{bmatrix},$$

where $\partial q / \partial p$ is a matrix of demand derivatives and \circ denotes element-by-element multiplication. The model nests Bertrand competition ($\kappa = 0$) and joint price-setting behavior ($\kappa = 1$), as well as capturing (non-microfounded) intermediate cases.

We generate data with different conduct parameters: $\kappa \in \{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$. For each specification, we simulate datasets with 400 observations (200 markets \times 2 firms), and estimate the model under the (erroneous) assumption of Bertrand price competition ($\kappa = 0$), thus generating supply-side misspecification. We then estimate the model using the covariance restriction approach assuming $\text{cov}[\Delta\xi, \Delta\eta] = 0$, using $\Delta\eta$ as an (observed) excluded instrument for demand and using $\Delta\xi$ as an (observed) excluded instrument for supply. Across all specifications, $\Delta\xi \sim U(0, 0.5)$ and $\Delta\eta \sim U(0, 0.5)$.³⁷

Table C1 displays the results. As expected, supply-side misspecification can introduce bias into the covariance restriction approach. The bias does not appear to be meaningful for modest values of κ (i.e., 0.6 or less). When the true nature of conduct is $\kappa = 1$ (joint price setting) but we assume Bertrand price competition, the bias is -3.8 percent. Likewise, the demand-side instruments (IV-2), which invoke the formal assumption about conduct in estimation, perform worse when the true κ is farther from the assumed value. The demand-side instruments perform poorly when the true conduct is $\kappa = 1$, with a mean bias of over 20 percent. By contrast, supply-side instruments do not use a formal assumption about conduct in estimation and provide consistent estimates across the specifications (IV-1). Consistent with the earlier simulations, the three-stage estimator outperforms IV-1 when conduct is correctly specified ($\kappa = 0$).

³⁶ Another form of misspecification could arise if prices or quantities are measured with error, in which case the demand and cost residuals might be correlated even if the underlying shocks are uncorrelated.

³⁷ We note that these are not mean zero, but it does not matter in this case. It is simply a normalization.

TABLE C1—SMALL-SAMPLE PROPERTIES: SUPPLY-SIDE MISSPECIFICATION

Estimation method	$\kappa = 0.0$	$\kappa = 0.2$	$\kappa = 0.4$	$\kappa = 0.6$	$\kappa = 0.8$	$\kappa = 1.0$
	(1)	(2)	(3)	(4)	(5)	(6)
Covariance restrictions	−1.001 (0.050)	−1.002 (0.052)	−1.000 (0.053)	−1.003 (0.054)	−1.016 (0.053)	−1.038 (0.051)
IV-1: Supply shifters	−1.002 (0.076)	−1.000 (0.077)	−1.001 (0.077)	−1.001 (0.076)	−1.001 (0.073)	−1.002 (0.071)
IV-2: Demand shifters	−1.015 (0.153)	−1.017 (0.155)	−1.012 (0.159)	−1.025 (0.178)	−1.082 (0.213)	−1.220 (0.298)
IV-1: First-stage F -statistic	1,079.9	1,335.3	1,424.9	1,277.8	1,027.1	801.9
IV-2: First-stage F -statistic	99.0	108.4	111.4	100.8	77.8	50.8

Notes: Results are based on 10,000 simulations of 200 duopoly markets for each specification. The demand curve is $h_{jt} = 2 - p_{jt} + \Delta\xi_{jt}$, so that $\alpha = -1$, and marginal costs are $c_{jt} = \Delta\eta_{jt}$. Demand is logit: $h(q_{jt}) = \ln(q_{jt}) - \ln(q_{0t})$, where q_{0t} is consumption of the outside good. IV-1 estimates are calculated using 2SLS with marginal costs ($\Delta\eta$) as an instrument in the demand equation. Analogously, IV-2 estimates are calculated using 2SLS with demand shocks ($\Delta\xi$) as an instrument in the supply relationship. Across all specifications, $\Delta\xi \sim U(0, 0.5)$ and $\Delta\eta \sim U(0, 0.5)$. The data-generating process varies in the nature of competition across specifications, indexed by the conduct parameter κ . The coefficients are estimated under the (misspecified) assumption of Bertrand price competition ($\kappa = 0$).

These results illustrate a key trade-off to the econometrician: if the supply-side assumptions are to be maintained, then covariance restrictions can offer better precision relative to instrument-based approaches. However, supply-side instruments are robust to misspecification of firm conduct, whereas covariance restrictions are not.

We note that the covariance restriction approach, which uses both demand-side and supply-side variation, is not as susceptible to misspecification bias as demand-side instruments in our simulations. The estimator appears to place greater weight on the source of variation with more power. In specification (6), the mean coefficient of -1.038 is much closer to the supply-shifter mean of -1.002 than the demand-shifter mean of -1.220 . Indeed, it is approximately equal to the IV-1 and IV-2 estimates weighted by the square root of the respective F -statistics. By placing greater weight on supply-side shocks as the demand-side instruments degrade, the covariance restriction approach appears to partially mitigate the bias from model misspecification.

REFERENCES

- Aguirregabiria, Victor, and Chun-Yu Ho.** 2012. "A Dynamic Oligopoly Game of the US Airline Industry: Estimation and Policy Experiments." *Journal of Econometrics* 168 (1): 156–73.
- Backus, Matthew, Christopher Conlon, and Michael Sinkinson.** 2021. "Common Ownership and Competition in the Ready-to-Eat Cereal Industry." Unpublished.
- Berry, Steven T.** 1994. "Estimating Discrete-Choice Models of Product Differentiation." *RAND Journal of Economics* 25 (2): 242–62.
- Berry, Steven, and Philip A. Haile.** 2014. "Identification in Differentiated Products Markets Using Market Level Data." *Econometrica* 82 (5): 1749–97.
- Berry, Steven T., and Philip A. Haile.** 2024. "Nonparametric Identification of Differentiated Products Demand Using Micro Data." *Econometrica* 92 (4): 1135–62.
- Berry, Steven, James Levinsohn, and Ariel Pakes.** 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): 841–90.
- Bound, John, David A. Jaeger, and Regina A. Baker.** 1995. "Problems with Instrumental Variables Estimation When the Correlation between the Instruments and the Endogenous Explanatory Variable Is Weak." *Journal of the American Statistical Association* 90 (430): 443–50.

- Bresnahan, Timothy F.** 1982. "The Oligopoly Solution Concept Is Identified." *Economics Letters* 10 (1–2): 87–92.
- Broda, Christian, and David E. Weinstein.** 2006. "Globalization and the Gains from Variety." *Quarterly Journal of Economics* 121 (2): 541–85.
- Broda, Christian, and David E. Weinstein.** 2010. "Product Creation and Destruction: Evidence and Price Implications." *American Economic Review* 100 (3): 691–723.
- Cardell, N. Scott.** 1997. "Variance Components Structures for the Extreme-Value and Logistic Distributions with Applications to Models of Heterogeneity." *Econometric Theory* 13 (2): 185–213.
- Chevalier, Judith A., Anil K. Kashyap, and Peter E. Rossi.** 2003. "Why Don't Prices Rise during Periods of Peak Demand? Evidence from Scanner Data." *American Economic Review* 93 (1): 15–37.
- Cho, Sungjin, Gong Lee, John Rust, and MengKai Yu.** 2018. "Optimal Dynamic Hotel Pricing." Unpublished.
- Conlon, Christopher, and Jeff Gortmaker.** 2020. "Best Practices for Differentiated Products Demand Estimation with PyBLP." *RAND Journal of Economics* 51 (4): 1108–61.
- De Loecker, Jan.** 2011. "Product Differentiation, Multiproduct Firms, and Estimating the Impact of Trade Liberalization on Productivity." *Econometrica* 79 (5): 1407–51.
- De los Santos, Babur, Daniel P. O'Brien, and Matthijs R. Wildenbeest.** 2024. "Agency Pricing and Bargaining: Evidence from the E-book Market." Unpublished.
- Deaton, Angus, and John Muellbauer.** 1980. "An Almost Ideal Demand System." *American Economic Review* 70 (3): 312–26.
- Dixit, Avinash K., and Joseph E. Stiglitz.** 1977. "Monopolistic Competition and Optimum Product Diversity." *American Economic Review* 67 (3): 297–308.
- Döpper, Hendrik, Alexander MacKay, Nathan Miller, and Joel Stiebale.** Forthcoming. "Rising Markups and Consumer Preferences." *Journal of Political Economy*.
- Doraszelski, Ulrich, and Jordi Jaumandreu.** 2013. "R&D and Productivity: Estimating Endogenous Productivity." *Review of Economic Studies* 80 (4): 1338–83.
- Dubé, Jean-Pierre, Jeremy T. Fox, and Che-Lin Su.** 2012. "Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation." *Econometrica* 80 (5): 2231–67.
- Feenstra, Robert C.** 1994. "New Product Varieties and the Measurement of International Prices." *American Economic Review* 84 (1): 157–77.
- Fisher, Franklin M.** 1963. "Uncorrelated Disturbances and Identifiability Criteria." *International Economic Review* 4 (2): 134–52.
- Fisher, Franklin M.** 1965. "Near-Identifiability and the Variances of the Disturbance Terms." *Econometrica* 33 (2): 409–19.
- Fowle, Meredith, Mar Reguant, and Stephen P. Ryan.** 2016. "Market-Based Emissions Regulation and Industry Dynamics." *Journal of Political Economy* 124 (1): 249–302.
- Frisch, Ragnar.** 1933. *Pitfalls in the Statistical Construction of Demand and Supply Curves*. Leipzig: Hans Buske, Verlag.
- Gandhi, Amit, and Jean-François Houde.** 2023. "Measuring Substitution Patterns in Differentiated Products Industries." Unpublished.
- Gandhi, Amit, and Aviv Nevo.** 2021. "Empirical Models of Demand and Supply in Differentiated Products Industries." In *Handbook of Industrial Organization*, Vol. 4, edited by Kate Ho, Ali Hortaçsu, and Alessandro Lizzeri, 63–139. Amsterdam: Elsevier.
- Grieco, Paul L. E., Charles Murry, and Ali Yurukoglu.** 2023. "The Evolution of Market Power in the US Automobile Industry." *Quarterly Journal of Economics* 139 (2): 1201–53.
- Hansen, Lars Peter.** 1982. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica* 50 (4): 1029–54.
- Hausman, Jerry A.** 1978. "Specification Tests in Econometrics." *Econometrica* 46 (6): 1251–71.
- Hausman, Jerry A.** 1996. "Valuation of New Goods under Perfect and Imperfect Competition." In *Economics of New Goods, Studies in Income and Wealth*, Vol. 58, edited by Timothy F. Bresnahan and Robert J. Gordon. Chicago, IL: NBER.
- Hausman, Jerry A., and William E. Taylor.** 1983. "Identification in Linear Simultaneous Equations Models with Covariance Restrictions: An Instrumental Variables Interpretation." *Econometrica* 51 (5): 1527–49.
- Hausman, Jerry A., Whitney K. Newey, and William E. Taylor.** 1987. "Efficient Estimation and Identification of Simultaneous Equation Models with Covariance Restrictions." *Econometrica* 55 (4): 849–74.

- Hendel, Igal, and Aviv Nevo.** 2013. "Intertemporal Price Discrimination in Storable Goods Markets." *American Economic Review* 103 (7): 2722–51.
- Keane, Michael, and Timothy Neal.** 2024. "A Practical Guide to Weak Instruments." *Annual Review of Economics* 16: 185–212.
- Knittel, Christopher R., and Konstantinos Metaxoglou.** 2014. "Estimation of Random-Coefficients Demand Models: Two Empiricists' Perspective." *Review of Economics and Statistics* 96 (1): 34–59.
- Koopmans, Tjalling C.** 1949. "Identification Problems in Economic Model Construction." *Econometrica* 17 (2): 125–44.
- Koopmans, Tjalling C., Herman Rubin, and Roy B. Leipnik.** 1950. "Measuring the Equation Systems of Dynamic Economics." In *Statistical Inference in Dynamic Economic Models (Cowles Commission Monograph 10)*, edited by Tjalling C. Koopmans, 53–237. New York, NY: John Wiley and Sons.
- Leamer, Edward E.** 1981. "Is It a Demand Curve, or Is It a Supply Curve? Partial Identification through Inequality Constraints." *Review of Economics and Statistics* 63 (3): 319–27.
- Leontief, Wassily.** 1929. "Ein Versuch Zur Statistischen Analyse von Angebot und Nachfrage." *Weltwirtschaftliches Archiv* 30 (1): 1–53.
- Li, Sophia, Joe Mazur, Yongjoon Park, James Roberts, Andrew Sweeting, and Jun Zhang.** 2022. "Repositioning and Market Power after Airline Mergers." *RAND Journal of Economics* 53 (1): 166–99.
- MacKay, Alexander, and Nathan H. Miller.** 2025. *Data and Code for: "Estimating Models of Supply and Demand: Instruments and Covariance Restrictions."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. <http://doi.org/10.3886/E198865V1>.
- Matzkin, Rosa L.** 2016. "On Independence Conditions in Nonseparable Models: Observable and Unobservable Instruments." *Journal of Econometrics* 191 (2): 302–11.
- Miller, Nathan H., and Matthew C. Weinberg.** 2017. "Understanding the Prices Effects of the Miller-Coors Joint Venture." *Econometrica* 85 (6): 1763–91.
- Miller, Nathan H., Matthew Osborne, and Gloria Sheu.** 2017. "Pass-through in a Concentrated Industry: Empirical Evidence and Regulatory Implications." *RAND Journal of Economics* 48 (1): 69–93.
- Nevo, Aviv.** 1998. "Identification of the Oligopoly Solution Concept in a Differentiated-Products Industry." *Economics Letters* 59 (3): 391–95.
- Nevo, Aviv.** 2000. "A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand." *Journal of Economics and Management Strategy* 9 (4): 513–48.
- Nevo, Aviv.** 2001. "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica* 69 (2): 307–42.
- Nevo, Aviv, and Adam M. Rosen.** 2012. "Identification with Imperfect Instruments." *Review of Economics and Statistics* 94 (3): 659–71.
- Petterson, Marco Stenborg, David Seim, and Jesse M. Shapiro.** 2023. "Bounds on a Slope from Size Restrictions on Economic Shocks." *American Economic Journal: Microeconomics* 15 (3): 552–72.
- Romeo, Charles.** 2013. "Filling Out the Instrument Set in Mixed Logit Demand Systems for Aggregate Data." Unpublished.
- Rothenberg, Thomas J.** 1971. "Identification in Parametric Models." *Econometrica* 39 (3): 577–91.
- Schultz, Henry.** 1928. *Statistical Laws of Demand and Supply*. Chicago, IL: University of Chicago Press.
- Soderbery, Anson.** 2015. "Estimating Import Supply and Demand Elasticities: Analysis and Implications." *Journal of International Economics* 96 (1): 1–17.
- Staiger, Douglas, and James H. Stock.** 1997. "Instrumental Variables with Weak Instruments." *Econometrica* 65 (3): 557–86.
- Stock, James H., and Motohiro Yogo.** 2005. "Testing for Weak Instruments in Linear IV Regression." In *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, edited by Donald W. K. Andrews and James H. Stock, 80–108. Cambridge, MA: Cambridge University Press.
- Thomadsen, Raphael.** 2005. "The Effect of Ownership Structure on Prices in Geographically Differentiated Industries." *RAND Journal of Economics* 36 (4): 908–29.
- Villas-Boas, Sofia Berto.** 2007. "Vertical Relationships between Manufacturers and Retailers: Inference with Limited Data." *Review of Economic Studies* 74 (2): 625–52.
- Waldfogel, Joel.** 2003. "Preference Externalities: An Empirical Study of Who Benefits Whom in Differentiated-Product Markets." *RAND Journal of Economics* 34 (3): 557–68.

- Wegge, Leon L.** 1965. "Identifiability Criteria for a System of Equations as a Whole." *Australian Journal of Statistics* 7 (3): 67–77.
- Williams, Kevin R.** 2022. "The Welfare Effects of Dynamic Pricing: Evidence from Airline Markets." *Econometrica* 90 (2): 831–58.
- Wooldridge, Jeffrey M.** 2010. *Econometric Analysis of Cross Sectional and Panel Data*, 2nd ed. Cambridge, MA: MIT Press.
- Zoutman, Floris T., Evelina Gavrilova, and Arnt O. Hopland.** 2018. "Estimating Both Supply and Demand Elasticities Using Variation in a Single Tax Rate." *Econometrica* 86 (2): 763–71.