

Large Market Asymptotics for Differentiated Product Demand Estimators with Economic Models of Supply

Timothy B. Armstrong*

Yale University

July 3, 2013

Abstract

IO economists often estimate demand for differentiated products using data sets with a small number of large markets. This paper addresses the question of consistency and asymptotic distributions of IV estimates as the number of products increases in some commonly used models of demand under conditions on economic primitives. I show that, in a Bertrand-Nash equilibrium, product characteristics lose their identifying power as price instruments in the limit in many of these models, leading to inconsistent estimates. The reason is that product characteristic instruments achieve identification through correlation with markups, and, depending on the model of demand, the supply side can constrain markups to converge to a constant quickly relative to sampling error. I find that product characteristic instruments can yield consistent estimates in many of the cases I consider, but care must be taken in modeling demand and choosing instruments. A monte carlo study confirms that the asymptotic results are relevant in market sizes of practical importance.

1 Introduction

The simultaneous determination of quantity and prices is a classic problem in demand estimation. A common solution in markets with differentiated products is to use characteristics

*Email: timothy.armstrong@yale.edu. Thanks to Han Hong and Liran Einav for guidance and many useful discussions, and Ariel Pakes, John Lazarev, Steve Berry, Phil Haile and participants at seminars at Stanford and Yale and the 2010 Econometric Society World Congress for useful comments and criticism. All remaining errors are my own.

of competing products as a source of exogenous variation in prices. The idea is that a firm facing stiffer competition will set lower markups, so, as long as they are independent of demand shocks, characteristics of competing products will be valid instruments for a product's price. The use of characteristics of competing products as price instruments is common in empirical studies of differentiated product markets, and goes back at least to Bresnahan (1987) and Berry, Levinsohn, and Pakes (1995) (hereafter, BLP). I refer to instruments of this form as product characteristic instruments or, in reference to the latter paper, "BLP instruments."

Many empirical studies of markets with differentiated products, including those that use the BLP instruments, use data on a relatively small number of markets, each with many products, to estimate demand elasticities. For example, in their application to automobile demand, BLP use data on 20 markets, each with about 100 products.¹ For demand models where the number of parameters grows with the number of product characteristics rather than the number of products, one might expect a small number of markets with a large number of products to give good estimates of demand.

This paper uses asymptotic approximations where asymptotics are taken in the number of products per market to examine the behavior of IV estimators of demand in large market settings, with a focus on product characteristic based instruments. Since the BLP instruments are correlated with prices only through equilibrium markups, their validity in this setting depends crucially on the nature of competition in markets with many products. If the dependence of markups on characteristics of other products disappears as the number of products increases and does so quickly enough, the BLP instruments will lose power in large markets and estimates based on them will be inconsistent when asymptotics are taken with respect to the number of products per market. The results in this paper use the asymptotic behavior of equilibrium markups to determine when this is the case.

I find that, in several commonly used models, the dependence of prices on product characteristic instruments through markups disappears at a fast enough rate that the BLP instruments lead to inconsistent estimates when asymptotics are taken in the number of products per market. In particular, this is the case with the logit and random coefficients logit models in a large market setting with many small firms and Bertrand competition.

Despite this negative result, I show that, in many other settings, the dependence of

¹In addition to BLP and other papers on the automobile industry (e.g. Petrin, 2002), examples of industries with data that fit this description include personal computers (e.g. Eizenberg, 2011; Bresnahan, Stern, and Trajtenberg, 1997), LCD televisions (e.g. Conlon, 2012) and other consumer goods sold at a national level.

markups on product characteristics remains or decreases slowly enough that the BLP instruments lead to consistent estimates under large market asymptotics. Under certain conditions, the BLP instruments lead to consistent estimates under logit demand with a small number of asymmetric firms, each with a large number of products. The asymptotic dependence of the markup on product characteristics can also be obtained in a setting with many small firms if the dimension of the space of observed product characteristics increases with the number of products (but with enough restrictions that the dimension of the parameter space does not). I illustrate this point with a nested logit model with many nests. Finally, considering the case with many large markets, the BLP instruments will lead to consistent estimates in certain cases if the number of markets increases quickly enough relative to the number of products in each market.

While this paper focuses on product characteristic instruments, in many situations, instruments will be available that shift marginal costs directly. Since cost shifter instruments do not need to enter the markup to be valid, they do not suffer from the issues brought up here, although it is still good practice to test for identification when using these instruments.

The results described above show that, while the BLP instruments provide consistent estimates under large market asymptotics in certain settings, care must be taken in choosing a demand specification that allows these instruments to have power in a fully specified model of supply and demand, particularly if the supply side is being used elsewhere in the analysis. The asymptotic results in this paper, along with monte carlos with prices generated from the supply side model, can be used as a guide to specifying a model of supply and demand that does not suffer from these issues. See section 4 for further discussion of how these issues can be diagnosed, and section 5 for monte carlo results based on the random coefficients logit model.

It should be emphasized that the inconsistency results in this paper are not just an issue of weak instruments (see, for example, Stock, Wright, and Yogo, 2002). Rather, the issue is that the supply side model may constrain product characteristic instruments to have very little identifying power. If this is the case and a researcher performs a test for identification and finds that these instruments are highly correlated with price, this must also be taken as evidence that the model is misspecified in ways that will likely lead to incorrect conclusions in policy counterfactuals based on the supply side model.² Performing specification tests based on supply side moments can help in diagnosing these issues, but these still have the

²The same issue will arise if the researcher performs inference in some other weak instrument robust way. Unless the supply side is misspecified, any weak instrument robust test will have power close to its size.

possibility of giving the wrong answer because of sampling error.

While it is good practice to use weak instrument robust methods, the most thorough way of guarding against the issues brought up in this paper is to confirm that it is possible for the BLP instruments to perform well when prices are generated from the supply side model being used elsewhere in the analysis and primitives are generated from a reasonable data generating process. This can be done through asymptotic approximations and monte carlos before even looking at the data (except, perhaps, to calibrate the data generating processes of product characteristics in the monte carlos). Only by taking these steps can one fully guard against the possibility of performing an analysis based on internally inconsistent assumptions. See section 4 for further discussion of how these issues can be diagnosed.

It should also be emphasized that, while this paper focuses on asymptotic approximations for their generality and tractability, the fundamental issue that the supply side may constrain product characteristics to perform poorly as price instruments is not a mere product of asymptotics, and can be elucidated in a finite sample framework through monte carlos. For a given sample size and data generating process for the model primitives, one can generate one data set with prices from a Bertrand equilibrium and another data set with markups set to a constant. If BLP instruments do not perform noticeably better in the first case (on average over monte carlo repetitions), one should look for different instruments or a different model of supply and demand. See section 5.1 for an illustration of this approach.

1.1 Related Literature

To my knowledge, the only other paper that considers asymptotics in the number of products per market in this setting is Berry, Linton, and Pakes (2004). In order to focus on other questions involving error from simulation based estimators and sampling error in product shares, those authors abstract from identification issues and from the supply side by placing high level conditions that assume that the instruments strongly identify the model under large market asymptotics. In contrast, the present paper asks which models of supply and demand allow product characteristic instruments to have power in a large market setting (and abstracts from the questions of simulation error and sampling error in market shares). In the cases where the present paper gives a negative answer to this question, the high level conditions for strong identification in Berry, Linton, and Pakes (2004) do not hold, and the present paper takes the further step of showing that this leads to inconsistent estimates (in contrast to the consistency results obtained by Berry, Linton, and Pakes, 2004, when their identification conditions hold).

More broadly, others have proposed different approaches to modeling and estimating demand in large markets, including Bajari and Benkard (2005) and Pinkse, Slade, and Brett (2002). While the present paper focuses on the demand models and estimators proposed in BLP, one could make a similar point about taking into account the implications of an equilibrium model for the behavior of estimators in other settings where one deals with a small number of distinct venues in which many agents interact. Performing this type of analysis in other settings is expected to be a useful topic for future research.

This paper is also related to the literature on weak instruments (see Stock, Wright, and Yogo, 2002, for a survey of this literature). That literature uses sequences of underlying distributions in which the correlation of instruments with endogenous variables shrinks with the sample size to get asymptotic approximations that better approximate finite sample distributions. This paper shows that such sequences arise endogenously from equilibrium prices when asymptotics are taken in the number of products per market in a certain class of models. To my knowledge, this is the first result in which weak instruments arise endogenously from equilibrium behavior in a sequence of pricing games. The fact that the power of instruments is tied to demand through equilibrium pricing in a supply side model can be used to inform decisions about how to model demand.

While the present paper focuses on product characteristic instruments, other papers have dealt with the optimal use of other instruments when they are available, such as variables that shift marginal cost directly. In contemporaneous work, Reynaert and Verboven (2012) focus on settings where cost instruments are available and consider approximations to optimal functions of cost shifters and product characteristics based on the assumption of perfect competition (a setting where BLP instruments have no power). Their focus also differs from the present paper in that they focus on estimation of the distribution of the random coefficients (σ in the notation below), while the present paper focuses on the price parameter (α in the notation below). Romeo (2010) proposes other instruments for models similar to those considered here.

In addition to the literatures on weak instruments and on estimation of discrete choice models of demand, this paper relates to theoretical results on oligopoly pricing in markets where demand is characterized by a discrete choice model. Gabaix, Laibson, Li, Li, Resnick, and de Vries (2013) consider the limiting behavior of prices in large markets in a similar setting to the present paper, but focus on different questions, leading to a different formulation and different results (for example, Gabaix, Laibson, Li, Li, Resnick, and de Vries achieve more generality in other directions by restricting attention to symmetric firms, while the

present paper deals with instruments that attempt to exploit observed asymmetry between firms). Existence of equilibrium in some of the pricing games I consider follows from arguments in Caplin and Nalebuff (1991), Vives (2001) and Konovalov and Sandor (2010) or similar methods. There is also a literature examining how restrictions on demand elasticities in discrete choice models place restrictions on the possible outcomes of empirical applications (see, among others, Bajari and Benkard, 2003; Akerberg and Rysman, 2005). While some of the findings of this paper add to this body of literature, the main focus is on implications for the identifying strength of BLP instruments.

1.2 Plan for Paper

The paper is organized as follows. Section 2 describes the class of models being studied. Section 3 derives the asymptotic behavior of equilibrium prices and IV estimates in some commonly used models of supply and demand. Section 4 presents recommendations for diagnosing whether the supply side constrains BLP instruments to have poor power. Section 5 provides a monte carlo study. Section 6 concludes. Proofs and additional results are in a supplementary appendix.

2 The Model

This section describes the class of models and estimators considered in this paper and defines some notation that will be used later. The models and much of the notation follow BLP and Berry (1994).

The researcher observes data from a single market with J products labeled 1 through J and an outside good labeled 0, and M consumers. Each product j has a price p_j and a vector of other characteristics observed to the researcher, $x_j \in \mathbb{R}^K$, and an unobserved variable ξ_j , which can be interpreted as a combination of unobserved product characteristics and aggregate preference shocks. In addition, each individual consumer i has consumer specific unobserved components of demand ε_{ij} and ζ_i , which are iid across consumers. In what follows, x_j is assumed to contain a constant.

Consumer i 's utility for the j th product is given by $u_{ij} = u(x_j, p_j, \xi_j, \varepsilon_{ij}, \zeta_i)$ for some function u . Each consumer buys the product for which utility is the highest, and no consumer buys more than one product. Rather than individual purchasing decisions, we observe aggregate market shares, including the proportion of consumers who make no purchase (the share of the outside good). These come from aggregating purchasing decisions over the ε

and ζ s of all consumers. It is assumed that the number of consumers is large enough that sampling variation in market shares from realizations of the ε s and ζ s can be ignored, so that the market share of good j , $s_j(x, \xi, p)$, is equal to the population probability of choosing good j conditional on x , ξ , and p : $s_j(x, \xi, p) = E_{\varepsilon, \zeta} I(u_{ij} > u_{ik} \text{ all } k \neq j)$.

In the models considered here, utility can be written in the following form for some parameters (α, β, σ) : $u_{ij} = x'_j \beta - \alpha p_j + \xi_j + g_i(\varepsilon_{ij}, \zeta_i, x_j, p_j, \sigma)$. Thus, utility can be separated into a linear part that does not depend on individual preferences, and a function that may depend on individual preferences, and the unobserved aggregate demand shock only enters the first part. The linear part is denoted by $\delta_j \equiv x'_j \beta - \alpha p_j + \xi_j$. Since shares only depend on ξ through δ , we can write them as $s_j(\delta, x, p, \sigma)$.

On the producer side, there are F firms labeled 1 through F . Firm f produces the set of goods $\mathcal{F}_f \subseteq \{1, \dots, J\}$. Profits of firm f are given by $\sum_{k \in \mathcal{F}_f} p_k \cdot M s_k(x, p, \xi) - C_f(\{M \cdot s_k(x, p, \xi)\}_{k \in \mathcal{F}_f})$ where M is the number of consumers and C_f is firm f 's cost function. Firms play a Nash-Bertrand equilibrium in prices, and rearranging the first order conditions for an interior best response gives

$$\sum_{k \in \mathcal{F}_f} (p_k - MC_k) \frac{\partial}{\partial p_j} s_k(x, p, \xi) + s_j(x, p, \xi) = 0 \quad (1)$$

for each product j owned by firm f , where MC_k is the marginal cost of product k .³ For single product firms, this simplifies to

$$p_j = MC_j - \frac{s_j(x, p, \xi)}{\frac{d}{dp_j} s_j(x, p, \xi)}. \quad (2)$$

When new products are added to the demand system, the equilibrium price will change, so that the equilibrium price and share of good j depend on the size of the market J . That is, even though x , MC , and ξ are sequences, prices and market shares will be triangular arrays, so that a more precise notation for the equilibrium price of good j would be $p_{j,J}$. To avoid extra subscripts, I use p_j to denote the price of good j when the context is clear.⁴

³While the models of sections 3.2, 3.3 and 3.4 have a unique equilibrium, results showing whether the pricing game in the random coefficients logit model has an equilibrium (or a unique one) in the general setting of section 3.1 are, to my knowledge, not available in the literature. The results in that section hold for any sequence of equilibria so long as such a sequence exists, and do not impose uniqueness.

⁴While the entry decision is not explicitly modeled here, one could think of the market size J as being endogenously determined by firms' decisions of whether to pay a fixed cost of entering the market. If the number of consumers M is large relative to the entry cost, more firms will enter, and one can think of asymptotics in J as arising from asymptotics in M in a two stage model with endogenous entry (note that

Finally, in all of the models below, I assume that the vector of unobserved demand shocks ξ is mean independent of observed product characteristics: $E(\xi|x) = 0$. This assumption is the exclusion restriction that provides the basis for the BLP instruments.

2.1 Estimation

In the models considered here, $s(\delta, x, p, \sigma)$ is invertible in its first argument (see Berry, 1994, BLP). Letting $\delta(s, x, p, \sigma)$ denote the inverse with respect to the first argument, this leads to the equation

$$\delta_j(s, x, p, \sigma) = x'_j \beta - \alpha p_j + \xi_j, \quad (3)$$

which can potentially be used to estimate the demand parameters σ , β and α . However, the parameter σ enters into a function with shares, which are endogenous. In addition, prices may be correlated with the unobserved ξ through at least two channels. First, ξ_j enters the markup $s(\delta, x, p, \sigma)/\frac{d}{dp_j}s(\delta, x, p, \sigma)$. Second, ξ_j may be correlated with unobserved components of marginal costs if goods that are more desirable in unobserved ways are also more expensive to make in unobserved ways.

To overcome this endogeneity problem, one needs instruments that are uncorrelated with ξ , and shift prices enough to identify α and σ . This paper focuses on instruments based on characteristics of other products, a common solution in the IO literature that goes back at least to BLP and Bresnahan (1987). Since we are assuming that ξ_j is independent of the observed characteristics of all products, functions of characteristics of other products will satisfy the instrumental variables exclusion restriction. Since characteristics of other products enter price through the markup $s(\delta, x, p, \sigma)/\frac{d}{dp_j}s(\delta, x, p, \sigma)$, they have the potential to shift prices enough to consistently estimate the model. Suppose that we use some vector valued function $h_j(x_{-j})$ as excluded instruments. The parameter estimates minimize the GMM criterion function

$$\left\| \frac{1}{J} \sum_{j=1}^J (\delta_j(s, x, p, \sigma) - x'_j \beta + \alpha p_j) z_j \right\|_{W_J} \quad (4)$$

where $z_j = (x_j, h_j(x_{-j}))'$ and W_J is a positive definite weighting matrix. Following com-

this interpretation requires that the information structure of the entry game is such that the exogeneity assumptions on x and ξ described below hold conditional on entry, which can be ensured by assuming that entry decisions are made before firms observe x , ξ and marginal costs).

mon terminology, this paper refers to instruments of this form as product characteristic instruments or BLP instruments.

3 Large Market Asymptotics

I now turn to the question of the asymptotic behavior of demand estimates, particularly those based on product characteristic instruments, under large market asymptotics. The results depend on the model of demand and on how products are added (how new products enter utility, whether the number of firms increases, whether firms are asymmetric, whether the number of markets increases as well as the number of products, etc.). Reflecting the different possible results, this section is organized into subsections considering these cases.

To get an idea of the results in this paper and how they are shown, consider first the simple logit model with many small firms in a single market. Consumer i 's utility for product j takes the form $u_{i,j} = x'_j\beta - \alpha p_j + \varepsilon_{i,j} + \xi_j$ where $\varepsilon_{i,j}$ is distributed extreme value independently across products and consumers. This leads to shares taking the form $s_j(x, p, \xi) = \frac{\exp(x'_j\beta - \alpha p_j + \xi_j)}{\sum_k \exp(x'_k\beta - \alpha p_k + \xi_k)}$, which can be inverted to get (normalizing the mean utility of the outside good 0 to zero) $\log s_j - \log s_0 = x'_j\beta - \alpha p_j + \xi_j$. The derivative of firm j 's share with respect to j 's price is $\frac{d}{dp_j} s_j(x, p, \xi) = -\alpha s_j(x, p, \xi)(1 - s_j(x, p, \xi))$, which gives the Bertrand pricing formula, Equation (2), as $p_j = MC_j + \frac{1}{\alpha(1-s_j)}$.

It is clear from this formula that, as long as the share of each firm converges to zero, markups of all products will converge to $1/\alpha$. If the markup were exactly equal to $1/\alpha$, product characteristic instruments would yield inconsistent estimates, since they must be correlated with markups to have identifying power. If the convergence of the markup $\frac{1}{\alpha(1-s_j)}$ to $1/\alpha$ is fast enough, this will also be true for the actual sequence of prices. In particular, it is sufficient for $\frac{1}{\alpha(1-s_j)}$ to converge to $1/\alpha$ at a faster than \sqrt{J} rate (with the appropriate uniformity conditions). The basic intuition for this is that instrumental variables estimators involve sample means that converge at a \sqrt{J} rate, and that the identifying power of instruments must shrink at a slower rate for consistency (this result is well known in the literature on weak instruments, and shows up in the distinction between “weak” and “semi-strong” asymptotics). In the simple logit model, since shares shrink at a $1/J$ rate (assuming they are roughly equal), the markup does indeed converge to $1/\alpha$ at a faster than \sqrt{J} rate, leading to inconsistency for estimators that use the BLP instruments.

The next subsection states these results formally for the random coefficients logit, which includes the simple logit described above as a special case. The general strategy for this and

most other models in this section is to find some limiting form of the markup that leads to some hypothetical prices p_j^* in a limiting model, then to show that actual prices p_j converge to p_j^* at a faster than \sqrt{J} rate. For the simple logit model with many small firms, this holds for $p_j^* = MC_j + 1/\alpha$, which leads to inconsistency of BLP instrument based estimators. In other cases, such as the nested logit model with many nests considered in section 3.2, markups do depend on product characteristics in the limiting model with p_j replaced by p_j^* , so that the BLP instruments retain identifying strength under large market asymptotics.⁵

3.1 Random Coefficients Logit with Many Small Firms

The random coefficients logit model, used by BLP, generalizes the simple logit, allowing for more general forms of consumer heterogeneity through random coefficients on the observed product characteristics x . Consumer i 's utility for product j takes the form

$$u_{ij} = x_j' \beta - \alpha p_j + \xi_j + \sum_k x_{jk} \zeta_{ik} + \varepsilon_{ij} \equiv \delta_j + \sum_k x_{jk} \zeta_{ik} + \varepsilon_{ij}$$

where ζ_{ik} is a random coefficient on product k . This specification assumes that there is no random coefficient on price. This is done for tractability and similar results will likely hold with a more general specification (in practice, including a random coefficient on price can be important because of differences in price sensitivity among consumers). This section considers asymptotics in which the number of products increases with a single product per firm. The single product firm assumption is made for simplicity, and the results will be similar as long as products are added by increasing the number of firms rather than the number of products per firm (see section B.3 of the supplementary appendix). It is also assumed that the dimension of the random coefficients ζ is fixed, and that new products differ only in drawing new characteristics (x_j and ξ_j) and $\varepsilon_{i,j}$ terms. As shown in sections 3.2 and 3.3, increasing the dimension of the random coefficients or the number of products per firm can lead to dramatically different results.

Shares can be obtained by integrating the logit shares for fixed ζ , which gives, letting P_ζ be the probability measure of the random coefficients, $s_j = \int \tilde{s}_j(\delta, \zeta) dP_\zeta(\zeta)$ where

⁵For the general specification of the random coefficients logit model in section 3.1 with many small firms, the results are stated for estimators of the linear parameters (α, β) with the parameter σ determining the random coefficients treated as known. Since the BLP instruments lead to inconsistent estimates in that setting, the (negative) message is essentially the same: even with σ known, product characteristic instruments do not give consistent estimates. However, extending consistency results such as those in sections 3.2, 3.3 and 3.4 to more general specifications of random coefficients is left for future research.

$\tilde{s}_j(\delta, \zeta) = \frac{\exp(\delta_j + \sum_k x_{jk}\zeta_k)}{\sum_\ell \exp(\delta_\ell + \sum_k x_{\ell k}\zeta_k)}$. Differentiating under the integral and using the formulas for logit elasticities for fixed ζ gives $\frac{ds_j}{dp_j} = -\alpha \int \tilde{s}_j(\delta, \zeta)(1 - \tilde{s}_j(\delta, \zeta))dP_\zeta(\zeta)$ so that the Bertrand markup is $p_j - MC_j = \frac{\int \tilde{s}_j(\delta, \zeta) dP_\zeta(\zeta)}{\alpha \int \tilde{s}_j(\delta, \zeta)(1 - \tilde{s}_j(\delta, \zeta))dP_\zeta(\zeta)}$. If the tails of ζ are thin enough, this can be shown to approach $1/\alpha$ quickly by truncating the integral and using bounds on the logit shares for ζ fixed, then arguing as in the simple logit model.

Theorem 1. *In the random coefficients model with single product firms and no random coefficient on price, suppose that product characteristics, marginal costs and prices are bounded with probability one, and that (x_j, ξ_j, MC_j) is iid across j . Suppose that the distribution of ζ is normal, or the tails of the distribution of ζ are bounded by the tails of a normal random variable. Let $z_j = (x_j, h(x_{j-L}, \dots, x_{j+M}))$ for some function h with finite variance. Let $(\hat{\alpha}, \hat{\beta})$ be the estimates that minimize (4) for this choice of z_j with σ fixed at its true value, and let $(\hat{\alpha}^*, \hat{\beta}^*)$ be obtained in the same way from data with p_j replaced by $p_j^* = MC_j + 1/\alpha$. Then $(\hat{\alpha}, \hat{\beta})$ is inconsistent and $\|(\hat{\alpha}, \hat{\beta}) - (\hat{\alpha}^*, \hat{\beta}^*)\| \xrightarrow{P} 0$.*

Theorem 1 shows that product characteristics lead to inconsistent estimates of α and β even if the nonlinear parameter σ is known and used in estimation. Note also that Theorem 1 makes the high level assumption that prices do not increase without bound when products are added. This will hold under the remaining primitive conditions of the theorem if ζ has bounded support (see section C of the supplementary appendix), and can be expected to hold for many other choices of the distribution of random coefficients. The support condition on product characteristics can be relaxed as well (also shown in section C of the supplementary appendix for the case of bounded ζ).

3.2 Nested Logit with Many Nests

Despite the seemingly rich specification of consumer heterogeneity in the random coefficients model considered above, the asymptotic behavior of markups is the same as in the simple logit model. The reason for this is that, with each new product, a new idiosyncratic $\varepsilon_{i,j}$ term is added, but the distribution of the other random coefficients remains the same, which leads to a limiting markup that is driven entirely by the $\varepsilon_{i,j}$ terms, with the random coefficients playing no role at all.⁶

One way of getting a more interesting limiting model is to add new random coefficients ζ in addition to adding new idiosyncratic terms ε . Increasing the dimension of ζ in a completely

⁶One can also consider a “pure characteristics” model with no $\varepsilon_{i,j}$ ’s (see Berry and Pakes, 2007). As shown in section B.2 of the supplementary appendix for a special case of this model, removing the $\varepsilon_{i,j}$ terms and adding products in an obvious way does not solve the problem.

unrestricted way, one would end up with an increasing number of parameters, which leads to its own problems. Section B.1 of the supplementary appendix considers the nested logit model, a special case of the random coefficients logit model in which products are placed in groups, and the random coefficients are associated with group indicator variables. The nested logit model places enough structure on the random coefficients that the number of groups can be increased without increasing the number of parameters that need to be estimated. The results in section B.1 show that under asymptotics where the dimension of ζ is increased by adding more groups, BLP instruments can have power asymptotically even with single product firms.

It should be emphasized that the nested logit model is used for tractability and is intended to illustrate the point that one can reverse the negative results of section 3.1 through a specification where the dimension of the random coefficients increases with the number of products while the dimension of the parameter space stays fixed. One could likely achieve a similar goal through other specifications of random coefficients (for example, starting with the nested logit model with many nests and adding some random coefficients to continuous variables).

3.3 Logit with Large Firms

Now consider the logit model with multiproduct firms, with the number of firms F fixed and asymptotics taken in the number of products per firm. The own and cross price elasticities can be shown to take the form $\frac{d}{dp_j} s_j(x, p, \xi) = -\alpha s_j(x, p, \xi)(1 - s_j(x, p, \xi))$ and $\frac{d}{dp_k} s_j(x, p, \xi) = \alpha s_j(x, p, \xi) s_k(x, p, \xi)$ respectively. Plugging these in to the equilibrium pricing equations (1) and rearranging gives the markup of product j produced by firm f as $p_j - MC_j = \frac{1}{\alpha} + \sum_{k \in \mathcal{F}_f} (p_k - MC_k) s_k$, so that markups are constant within a firm (see Konovalov and Sandor, 2010). Letting b_f be the common markup for firm f , this gives a system of equations that define b_f for $f = 1, \dots, F$. Rearranging and plugging in the formula for shares yields

$$b_f = \frac{1}{\alpha} \frac{\sum_{k \in \mathcal{F}_f} \exp(x'_k \beta - \alpha MC_k + \xi_k - \alpha b_f)}{\sum_{h \neq f} \sum_{k \in \mathcal{F}_h} \exp(x'_k \beta - \alpha MC_k + \xi_k - \alpha b_h)} = \frac{1}{\alpha} \frac{\exp(-\alpha b_f) \hat{\pi}_f \bar{r}_f}{\sum_{h \neq f} \exp(-\alpha b_h) \hat{\pi}_h \bar{r}_h} \quad (5)$$

where $\hat{\pi}_f \equiv |\mathcal{F}_f|/J$ is the proportion of products produced by firm f and \bar{r}_f is an average of the characteristics of firm f 's products given by $\bar{r}_f \equiv \frac{1}{|\mathcal{F}_f|} \sum_{k \in \mathcal{F}_f} \exp(x'_k \beta - \alpha MC_k + \xi_k)$.

Under a law of large numbers, \bar{r}_f will converge to some $\mu_{r,f}$ for each firm f . Assuming that $\hat{\pi}_f$ also converges to some π_f for each firm f , this suggests that equilibrium prices will

be determined asymptotically by the solution to equation (5) with \bar{r}_f and $\hat{\pi}_f$ replaced by $\mu_{r,f}$ and π_f . This is formalized in the following theorem.

Theorem 2. *In the simple logit model with asymptotics in the number of products per firm, suppose that (x_j, ξ_j, MC_j) is independent across all j and identically distributed within each firm with finite variance. Let $\mu_{r,f} = E \exp(x'_j \beta - \alpha MC_j + \xi_j)$ for j produced by firm f , and suppose $\hat{\pi}_f \rightarrow \pi_f$ for some π_f each f . Let (b_1^*, \dots, b_F^*) be the unique solution to (5) with $\hat{\pi}_f \bar{r}_f$ replaced by $\pi_f \mu_{r,f}$ for each firm f , and let $p_j^* = MC_j + b_f$ for product j produced by firm f . Let $z_j = \left(x_j, \frac{1}{J} \sum_{k \in \mathcal{F}_f} \tilde{h}(x_k) \right)$ for some finite variance function \tilde{h} for product j owned by firm f , and let $(\hat{\alpha}, \hat{\beta})$ be the estimators defined by (4) with these instruments, and let $(\hat{\alpha}^*, \hat{\beta}^*)$ be defined in the same way, but with p_j^* replacing p_j . Then, if $\frac{1}{J} \sum_{j=1}^J E z_j(x'_j, p_j^*)$ converges to a positive definite matrix, $(\hat{\alpha}, \hat{\beta})'$ will be consistent and asymptotically normal with the same asymptotic distribution as $(\hat{\alpha}^*, \hat{\beta}^*)'$.*

Theorem 2 shows that product characteristic instruments can have identifying power in this setting if they exploit variation in $\pi_f \mu_{r,f}$ across firms. Thus, BLP instruments can have power through variation across firms, but not within firms.

3.4 Many Large Markets

According to the results of section 3.1, the BLP instruments lose power in a single large market (or any fixed number of large markets) with many firms fast enough that estimates based on them are inconsistent. In contrast, BLP instruments will typically provide enough variation to consistently estimate these models if the market size is bounded, and asymptotics are taken in the number of markets. This section considers the intermediate case where both the number of products and markets are allowed to go to infinity. To simplify the analysis, attention is restricted to the simple logit model with no random coefficients.

Some additional notation is needed to describe the results with many markets. We consider N markets, with J_i products in market i . Notation is otherwise the same as described in section 2, except that prices, product characteristics, etc. are now indexed by the market i as well as the product j , so that $p_{i,j}$ denotes the price of product i in market j (as before, the dependence of $p_{i,j}$ on the total market size J_i is suppressed in the notation). Let $\bar{J} = \frac{1}{N} \sum_{i=1}^N J_i$ be the average number of products per market. Under the asymptotics in this section, each J_i (and therefore \bar{J}) increases with the number of markets N , but the dependence of \bar{J} and the J_i s on N is suppressed in the notation.

Let $v_N = \frac{1}{N} \sum_{i=1}^N (J_i - \bar{J})^2 / \bar{J}^2$. Here, the J_i s are nonrandom, but v_N can be thought of as the normalized sample variance of the J_i s. The following theorem derives the asymptotic behavior of estimates based on the BLP instruments in the case where v_N is bounded away from zero.

Theorem 3. *In the simple logit model with many large markets, suppose that (x_j, ξ_j, MC_j) is bounded and iid across j . Let $(\hat{\alpha}, \hat{\beta})$ be IV estimates with instrument vector $z_{i,j} = (x_{i,j}, \frac{1}{J} \sum_{k=1}^{J_i} \tilde{h}(x_{i,k}))$ for some finite variance function \tilde{h} . Suppose that $v_N \rightarrow v$ for some $v > 0$. Then, if $N/\bar{J} \rightarrow \infty$, $(\hat{\alpha}, \hat{\beta})$ will be consistent and asymptotically normal, with a $\sqrt{N/\bar{J}}$ rate of convergence and asymptotic variance given in the supplementary appendix. If $N/\bar{J} \rightarrow c$ for some c , then $(\hat{\alpha}, \hat{\beta})$ will be inconsistent, and will follow a weak instrument asymptotic distribution given in the supplementary appendix. The weak instrument asymptotic coincides with what would be obtained with markups equal to $1/\alpha$ in the case where $c = 0$.*

Theorem 3 states that, as long as there is sufficient variation in the number of products per market, the BLP instruments will use this variation to obtain consistent estimates at a rate $\sqrt{N/\bar{J}}$ as long as $N/\bar{J} \rightarrow \infty$. If $N/\bar{J} \rightarrow c$ for some finite c , the results give “weak instrument” asymptotics, in which the estimates do not converge and follow a nonstandard asymptotic distribution. Note that Theorem 3 shows that the identifying power of BLP instruments in this setting relies on variation in market size (v_N must not converge to zero). If one is not comfortable using this variation to identify demand (for example, because of issues raised by Akerberg and Rysman, 2005), one will want to look for other instruments or a different specification.

The weak instrument formula in Theorem 3 is one way of assessing whether, in the simple logit model, the number of markets is large enough relative to the number of products that the negative results in section 3.1 can be circumvented without appealing to something more like the positive results in section 3.2 or 3.3. This approach appears to be very conservative in situations where richer specifications of random coefficients or firms with many products provide additional variation, so a negative outcome of such an analysis should be taken as a warning, but not as evidence that nothing can be learned with BLP instruments.

4 Diagnosis and Recommendations for Empirical Practice

The negative results in section 3.1 are a combination of two potential problems. First, BLP instruments may lack power. Second, the supply side may be misspecified. More specifically, the results in section 3.1 state that, for certain specifications of supply and demand, the correct specification of the supply side constrains the BLP instruments to have essentially no power, so such a specification of supply and demand will necessarily suffer from problem one (lack of power) or problem two (misspecification). Thus, a careful diagnosis will first look at whether there exists a reasonable data generating process for which the analysis will suffer from neither of these problems. This is discussed below in section 4.1. One can then proceed to address the issues of whether the data at hand are consistent with instrument power and correct model specification through statistical tests. Sections 4.2, 4.3 and 4.4 review some tests from the literature and propose additional tests tailored to demand estimation with BLP instruments.

4.1 Verifying Internal Consistency of Supply with Identification of Demand

The asymptotic results in section 3 can be used as a guide to specifying a model of demand that is internally consistent with BLP instruments having identifying strength with a given model of supply. To examine these issues in finite samples for a particular data set, I recommend the following procedure. First, obtain point estimates of all parameters, including parameters that determine the joint distribution of demand and marginal cost unobservables, where marginal costs are constrained to be independent of characteristics of other products.⁷ Next, generate monte carlo data sets with this estimated data generating process for the model primitives. For each monte carlo repetition, generate one data set where prices are generated from the supply side model, and another data set where prices are set equal to marginal costs plus $1/\alpha$. Perform the planned analysis (estimates or confidence intervals for demand elasticities, counterfactual prices, etc.) for both data sets. If the performance in the data set with prices generated from the actual model is good enough relative to the data set with constant markups or relative to some other standards set by the researcher, then

⁷If possible, one should perform this analysis for all parameter values not rejected by a first stage weak instrument robust test, rather than just the point estimates.

the researcher can conclude that the model does not constrain BLP instruments to perform poorly.

4.2 Tests for Identification

The null of lack of identification can be tested for the estimators considered here using general methods for testing identification in GMM, such as the test proposed by Wright (2003). While this test does not address the possibility of “weak” instruments that retain some power, it will not reject with probability greater than its size under sequences such as those in section 3.1 in which instruments have essentially no power. One can also guard against size distortion from weak identification by using weak identification robust tests (see, e.g., Kleibergen, 2005).

While useful for detecting identification issues, a test for identification as in Wright (2003) may find that BLP instruments are strong even in a setting where the supply side model constrains them to have negligible identifying power for the sample size at hand. In what follows, I describe a test statistic for identification for product characteristic instruments that is designed to find evidence of identification only through correlation with the markup allowed by the supply side. Let $z_{i,j} = (x_{i,j}, h(x_{i,-j}))$ be the instruments being considered for estimation, where h is a \mathbb{R}^{d_h} valued function of product characteristic instruments.

The test is based on a sample analog of the concentration parameter, a measure of the strength of identification given by

$$\mu^2(\pi, \sigma_\eta) = \pi' \sum_{i=1}^N \sum_{j=1}^{J_i} \bar{z}_{i,j} \bar{z}'_{i,j} \pi / \sigma_\eta^2 \quad (6)$$

where $\bar{z}_{i,j} = h(x_{i,-j}) - \left[\sum_{i=1}^N \sum_{j=1}^{J_i} h(x_{i,-j}) x'_{i,j} \right] \left[\sum_{i=1}^N \sum_{j=1}^{J_i} x_{i,j} x'_{i,j} \right]^{-1} x_{i,j}$, π is the coefficient of $h(x_{i,-j})$ in the population projection of $p_{i,j}$ on $(x'_{i,j}, h(x_{i,-j})')$, and σ_η is the variance of the residual in the projection of $p_{i,j}$ on $(x'_{i,j}, h(x_{i,-j})')$. Under weak instrument asymptotics, the limiting distribution of the first stage F statistic is $1/d_h$ times a limiting noncentral chi square distribution with noncentrality parameter μ^2 , and can be used to diagnose weak identification through inference on μ^2 (see, e.g. Stock and Yogo, 2002).

To form an analogue of the first stage F statistic that only finds support for identification by BLP instruments through correctly specified channels, note that, if $h(x_{i,-j})$ is correlated with price only through the markup, then the population first stage coefficient π will be equal to the coefficient of $h(x_{i,-j})$ in the population best linear predictor for the markup. Thus,

for a given parameter vector (α, β, σ) , one can estimate π using the regression of the markup $(\frac{d}{dp} s_{i,j}(x, p, \xi, \alpha, \beta, \sigma))^{s_{i,j}}$ in the single product firm case) on $(x'_{i,j}, h(x_{i,-j}))'$. Plugging this estimate of π and an estimate of σ_η from the residuals of the first stage regression for $p_{i,j}$ into the formula for μ^2/d_h gives a version of the first stage F statistic tailored to the BLP instruments. (It should be noted, however, that this test estimates π and σ_η from two different regressions, and so cannot be obtained from the F statistic from any one regression.)

Results in Stock and Yogo (2002) suggest that one should worry about weak identification leading to a relatively large bias for estimates of α when this statistic is below 10, although this is based on a somewhat different setting (in particular, one that would correspond to an estimate with σ known in the present setting).⁸ See section D of the supplementary appendix for monte carlos based on this test statistic in the full model with $\hat{\sigma}$ estimated. The test can be performed for markups derived from several plausible values of (α, β, σ) (ideally, over a grid of all values not rejected by an initial weak instrument robust confidence region) to assess the plausibility of strong or weak identification. Given that this test only detects weak identification of the linear part of the model, it should be supplemented with a full test for lack of identification.

4.3 Tests of Overidentifying Restrictions

After finding evidence that the BLP instruments have identifying strength, it is a good idea to perform tests to ensure that the supply side model is consistent with the assumption that identification comes through correlation with the markup. This can be done by including supply side moments and testing overidentifying restrictions. Namely, if the instruments $z_{i,j} = (x_{i,j}, h_j(x_{i,-j}))$ are used, and characteristics of other products $h_j(x_{i,-j})$ are “BLP instruments” rather than cost instruments, they should not be included in the cost equation. Thus, standard overidentification tests with supply side moments using BLP instruments can be used to detect cases where the identifying power of BLP instruments constrains them to be misspecified. Since failing to reject with such a test may be the result of type II error, such a test should be performed only after the analysis described in section 4.1 has indicated that the specification should not be thrown out a priori.

⁸The critical values in Stock and Yogo (2002) based on the limiting distribution of the usual first stage F statistic also differ from the corresponding test based on the statistic proposed here because a different estimate of π is used. Given that the estimate of π used here uses additional information about the instruments only entering the markup, it seems likely that the distribution will be less dispersed in most cases, making the Stock and Yogo (2002) critical values more conservative, although verification of this conjecture is left for future research.

4.4 An Additional Diagnostic for Spurious Results from BLP Instruments

A good sanity check for results based on BLP instruments is to simulate price data from a model where these instruments should have no identifying power, and reestimate the model with this data. In particular, one can simulate prices $p_{i,j}^* = \widehat{MC}_{i,j} + 1/\hat{\alpha}$ from a model with constant markups, with marginal costs and all other parameters estimated from the original data set, and form estimates $(\hat{\beta}^*, \hat{\alpha}^*, \hat{\sigma}^*)$ and standard errors for these estimates. If the new estimates and standard errors are close to the ones obtained from the original data, the researcher should be suspicious that the original results were spurious since, if those results were correct, it would mean that one could generate similarly precise estimates from an unidentified model. Formally, it can be shown through an analysis similar to the one in section 3 that, if the model constrains markups to be close to constant relative to sampling error, these simulated $p_{i,j}^*$ s will, by construction, be close enough to the original prices that the new estimates will be close to the original estimates relative to sampling error. The question of how close the sets of estimates should be to conclude that the model is misspecified could be addressed formally using a statistical test, although this is left for future research.

5 Monte Carlo

This section presents the results of a monte carlo study of the random coefficients logit model of section 3.1. The performance of the BLP instruments and of cost instruments is examined over a range of specifications for the number of markets, the number of products per market and variation in the number of products per market. Prices are generated from a Bertrand equilibrium with 10 products per firm. See the supplementary appendix for details of the data generating process, and results for additional specifications in which firm size varies as well.⁹

Tables 1 and 2 show the results for BLP instruments and cost instruments for several monte carlo designs. Results are reported for estimates of the price coefficient α and for a nominal level .05 two sided test for α . The monte carlo results appear consistent with the overall result that BLP instruments perform poorly for this demand specification when the number of products is large enough relative to the number of markets, and that cost

⁹See also the contemporaneous work of Skrainka (2012) for additional monte carlo results for BLP and cost instruments with prices generated from equilibrium play.

markets	products per market	median bias of $\hat{\alpha}$	median absolute deviation from α_0	rejection prob. at true α	power of test of $\alpha = 0$
1	60	-0.3698	0.6691	0.0783	0.1004
1	100	-0.3648	0.7177	0.1211	0.1381
3	20	-0.1053	0.3358	0.0390	0.3980
3	varied	-0.0494	0.2966	0.1523	0.4649
3	60	-0.2186	0.5827	0.0200	0.1040
3	100	-0.2525	0.6383	0.1351	0.1762
20	20	0.0039	0.1140	0.0390	0.9880
20	varied	-0.0014	0.1001	0.0400	0.9960
20	60	0.0130	0.2345	0.0230	0.7710
20	100	-0.0379	0.3154	0.0200	0.4560

Table 1: Monte Carlo Results for BLP Instruments

markets	products per market	median bias of $\hat{\alpha}$	median absolute deviation from α_0	rejection prob. at true α	power of test of $\alpha = 0$
1	60	-0.0247	0.1749	0.1130	0.6710
1	100	-0.0196	0.1358	0.0762	0.7623
3	20	-0.0262	0.1837	0.1002	0.6092
3	varied	-0.0122	0.1000	0.0852	0.7916
3	60	-0.0102	0.1007	0.0661	0.7768
3	100	-0.0054	0.0767	0.0662	0.8175
20	20	0.0065	0.0663	0.0220	0.7840
20	varied	-0.0008	0.0385	0.0522	0.8554
20	60	-0.0023	0.0365	0.0641	0.8707
20	100	-0.0027	0.0298	0.0481	0.8826

Table 2: Monte Carlo Results for Cost Instruments

instruments do not suffer from these issues. Estimates that use BLP instruments have substantial bias and variability when the number of products per market is large and the number of markets small (as measured by median bias and median absolute deviation from the true value), and perform better when the number of products per market is small. In contrast, cost instruments lead to relatively precise estimates in large market settings.

5.1 When is the Single Large Market Limiting Model a Good Approximation?

The results of Section 3.1 show that, under asymptotics where the number of products increases with firm size and the number of markets fixed, IV estimators do no better than they would with constant markups. To address how well this limiting model approximates Bertrand equilibrium with the monte carlo data generating processes used in this section, I simulate from the same data generating process for x , ξ and MC , but set markups to $1/\alpha$ for all products and compare estimates based on these data sets to the previously reported estimates computed from data sets with Bertrand prices. Table 3 reports the results of applying the same BLP instrument based estimators used in Table 1 to the monte carlo data sets with constant markups.

markets	products per market	median bias	median absolute deviation from α_0	rejection prob. at true α	power of test of $\alpha = 0$
1	60	-0.3112	0.6440	0.0521	0.0922
1	100	-0.3156	0.6748	0.1117	0.1368
3	20	-0.3069	0.7160	0.0150	0.0520
3	varied	-0.3049	0.7559	0.0090	0.0560
3	60	-0.3540	0.7290	0.0120	0.0460
3	100	-0.3341	0.7353	0.0250	0.0581
20	20	-0.3292	0.7525	0.0100	0.0430
20	varied	-0.3570	0.8237	0.0090	0.0500
20	60	-0.3387	0.8265	0.1533	0.1814
20	100	-0.3454	0.7592	0.0090	0.0470

Table 3: Monte Carlo Results for BLP Instruments with Constant Markup

The results show that, while the limiting model gives a pessimistic description of the behavior of BLP instrument based estimates for some of the cases considered, in other cases it is accurate enough that one would worry about applying the BLP instruments. With a single market, BLP instruments do not appear to perform noticeably better under Bertrand

pricing than in the limiting model in any of the monte carlo designs. With 3 markets and 100 products, the median bias and median absolute deviation of the estimate of α are only slightly better in the true model than they are with a constant markup, and the size distortion in the two sided test for α is actually worse. With smaller firms (reported in the appendix), the picture is even worse, with estimates of α doing essentially no better than in the constant markup model for most monte carlo designs, although the results are better when the number of products per firm varies (also reported in the appendix).

6 Conclusion

This paper derives asymptotic approximations for differentiated products demand estimators when the number of products is large. The question of whether product characteristic instruments have nontrivial identifying power is addressed through asymptotic correlation with equilibrium markups derived from a full model of supply and demand. The results show that certain supply and demand models constrain these instruments to have trivial power under large market asymptotics, and should therefore be avoided in cases where these asymptotic results are relevant. Other asymptotic settings (demand models and ways of adding products) are shown to lead to consistent estimates and standard asymptotic distributions under large market asymptotics. The results can be used as a guide to forming a model of demand in a large market setting that is consistent with finding identification through product characteristic instruments and variation in the markups. A monte carlo study shows that the asymptotic results are a good description of finite sample settings of practical importance.

References

- ACKERBERG, D. A., AND M. RYSMAN (2005): “Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects,” *The RAND Journal of Economics*, 36(4), 771–788.
- BAJARI, P., AND C. L. BENKARD (2005): “Demand estimation with heterogeneous consumers and unobserved product characteristics: A hedonic approach,” *Journal of Political Economy*, 113(6), 1239–1276.

- BAJARI, P., AND L. BENKARD (2003): “Discrete Choice Models as Structural Models of Demand: Some Economic Implications of Common Approaches,” .
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63(4), 841–890.
- BERRY, S., O. B. LINTON, AND A. PAKES (2004): “Limit Theorems for Estimating the Parameters of Differentiated Product Demand Systems,” *The Review of Economic Studies*, 71(3), 613–654.
- BERRY, S., AND A. PAKES (2007): “The Pure Characteristics Demand Model,” *International Economic Review*, 48(4), 1193–1225.
- BERRY, S. T. (1994): “Estimating Discrete-Choice Models of Product Differentiation,” *The RAND Journal of Economics*, 25(2), 242–262.
- BRESNAHAN, T. F. (1987): “Competition and Collusion in the American Automobile Industry: The 1955 Price War,” *The Journal of Industrial Economics*, 35(4), 457–482.
- BRESNAHAN, T. F., S. STERN, AND M. TRAJTENBERG (1997): “Market Segmentation and the Sources of Rents from Innovation: Personal Computers in the Late 1980s,” *The RAND Journal of Economics*, 28, S17–S44.
- CAPLIN, A., AND B. NALEBUFF (1991): “Aggregation and Imperfect Competition: On the Existence of Equilibrium,” *Econometrica*, 59(1), 25–59.
- CONLON, C. T. (2012): “A Dynamic Model of Prices and Margins in the LCD TV Industry,” *Manuscript, Columbia University*.
- DEVROYE, L. (1981): “Laws of the Iterated Logarithm for Order Statistics of Uniform Spacings,” *The Annals of Probability*, 9(5), 860–867.
- EIZENBERG, A. (2011): “Upstream Innovation and Product Variety in the U.S. Home PC Market,” SSRN Scholarly Paper ID 1760828, Social Science Research Network, Rochester, NY.
- GABAIX, X., D. LAIBSON, D. LI, H. LI, S. RESNICK, AND C. G. DE VRIES (2013): “The Impact of Competition on Prices with Numerous Firms,” Discussion paper.

- KLEIBERGEN, F. (2005): “Testing Parameters in GMM without Assuming That They Are Identified,” *Econometrica*, 73(4), 1103–1123.
- KONOVALOV, A., AND Z. SANDOR (2010): “On price equilibrium with multi-product firms,” *Economic Theory*, 44(2), 271–292.
- PETRIN, A. (2002): “Quantifying the Benefits of New Products: The Case of the Minivan,” *Journal of Political Economy*, 110(4), 705–729.
- PINKSE, J., M. E. SLADE, AND C. BRETT (2002): “Spatial Price Competition: A Semi-parametric Approach,” *Econometrica*, 70(3), 1111–1153.
- REYNAERT, M., AND F. VERBOVEN (2012): “Improving the Performance of Random Coefficients Demand Models: The Role of Optimal Instruments,” Discussion paper.
- ROMEO, C. (2010): “Filling Out the Instrument Set in Mixed Logit Demand Systems for Aggregate Data,” Discussion paper.
- SKRAINKA, B. S. (2012): “A Large Scale Study of the Small Sample Performance of Random Coefficient Models of Demand,” SSRN Scholarly Paper ID 1942627, Social Science Research Network, Rochester, NY.
- STAIGER, D., AND J. H. STOCK (1997): “Instrumental Variables Regression with Weak Instruments,” *Econometrica*, 65(3), 557–586.
- STOCK, J. H., AND J. H. WRIGHT (2000): “GMM with Weak Identification,” *Econometrica*, 68(5), 1055–1096.
- STOCK, J. H., J. H. WRIGHT, AND M. YOGO (2002): “A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments,” *Journal of Business and Economic Statistics*, 20(4), 518–529.
- STOCK, J. H., AND M. YOGO (2002): “Testing for Weak Instruments in Linear IV Regression,” *National Bureau of Economic Research Technical Working Paper Series*, No. 284.
- VIVES, X. (2001): *Oligopoly pricing*. MIT Press.
- WRIGHT, J. H. (2003): “Detecting Lack of Identification in Gmm,” *Econometric Theory*, 19(02), 322–330.