

Nonparametric Identification of Differentiated Products Demand Using Micro Data*

Steven T. Berry
Yale University

Philip A. Haile
Yale University

January 2023

Abstract

We examine identification of differentiated products demand when one has “micro data” linking the characteristics and choices of individual consumers. Our model nests standard specifications featuring rich observed and unobserved consumer heterogeneity as well as product/market-level unobservables that introduce the problem of econometric endogeneity. Previous work establishes identification of such models using market-level data and instruments for all prices and quantities. Micro data provides a panel structure that facilitates richer demand specifications and reduces requirements on both the number and types of instrumental variables. We address identification of demand in the standard case in which non-price product characteristics are assumed exogenous, but also cover identification of demand elasticities and other key features when these product characteristics are endogenous and not instrumented. We discuss implications of these results for applied work.

*Early versions of this work were presented in the working paper “Nonparametric Identification of Multinomial Choice Demand Models with Heterogeneous Consumers,” first circulated in 2007 and superseded by the present paper. We thank Jesse Shapiro, Suk Joon Son, and numerous seminar participants for helpful comments. Miho Hong and Jaewon Lee provided capable research assistance.

1 Introduction

Demand systems for differentiated products are central to many questions in economics. In practice it is common to estimate demand using data on the characteristics and choices of many individual consumers within each market. This setting is often referred to as “micro data,” in contrast to another common case in which only market-level outcomes are observed.¹ At an intuitive level, the panel structure of micro data seems to offer more information than market-level data alone. But in what precise sense does micro data help? How significant are the advantages of micro data? What specific kinds of variation within and across markets are helpful, and how?

This paper considers nonparametric identification of demand using micro data. Our nonparametric consumer-level demand model substantially generalizes standard parametric models used in a large literature building on Berry, Levinsohn, and Pakes (1995, 2004). Micro data provides a panel structure, with many consumers in each of many markets. A key benefit is that unobservables at the level of the product \times market remain fixed as consumers’ attributes and choices (quantities demanded) vary within a given market. We show that identification can be obtained by combining this clean “within” variation with cross-market variation in choice characteristics, market characteristics, prices, and instruments for prices (only). Compared to settings with only market-level data, this both allows a more general demand model and substantially reduces demands on instrumental variables.

Although we focus exclusively on identification, our aim is to inform the practice and evaluation of empirical work. The celebrated “credibility revolution” in applied microeconomics has redoubled attention to identification obtained through quasi-experimental variation, such as that arising through instrumental variables, geographic boundaries, or repeated observations within a single economic unit. Identification of demand presents challenges that are absent in much of empirical economics (see, e.g., Berry and Haile (2021)). Nonetheless, we show that these same types of variation allow identification of demand systems exhibiting rich consumer heterogeneity and endogeneity. Nonparametric identification results do not eliminate concerns about the impact of parametric assumptions relied on in practice. However, they address the important question of whether such assumptions can be viewed properly as finite-sample approximations rather than essential maintained hypotheses. Identification results can also clarify which assumptions may be most difficult to relax, reveal essential sources of variation, point to specific roles that functional forms may play in practice, offer assurance that robustness analysis is

¹See Berry and Haile (2021) for a discussion of other forms of data, including consumer panels and hybrids such as that in Petrin (2002).

possible, and potentially lead to new (parametric or nonparametric) estimation approaches.

Our most important message for applied work is that micro data has a high marginal value over market-level data alone. Availability of instrumental variables is the most important and challenging requirement for identification of demand, and micro data can substantially reduce both the number and types of instruments needed. Berry and Haile (2014) showed that with market-level data, nonparametric identification typically requires instruments for all quantities and prices. There, the so-called “BLP instruments” (i.e., exogenous characteristics of competing products) play a crucial role as instruments for quantities.² In contrast, here we find that with sufficiently rich micro data the only essential instruments are those for prices. This cuts the number of required instruments in half and avoids the necessary reliance on BLP instruments. This in turn permits a more flexible model of how non-price product characteristics affect demand and avoids the necessity that at least some such characteristics be exogenous. Micro data also opens the possible use of additional classes of instruments.

We also show that it is often possible to identify the *ceteris paribus* effects of prices on quantities demanded—critically, e.g., own- and cross-price demand elasticities—when observed non-price characteristics of products/markets are endogenous and not themselves instrumented. This requires that instruments for prices remain valid when conditioning on the endogenous non-price observables.³ We show that standard instruments (or variations thereon) can satisfy this requirement under many models of endogeneity. Our analysis of instruments for this case makes elementary use of causal graphs, which provide attractive tools for evaluating the necessary exclusion condition. Endogenous product characteristics are an important concern in the applied literature on differentiated products, and it can be difficult to find instruments for all such characteristics. Thus, our findings expand the range of applications in which primary features of interest can be identified despite these concerns.

Our model and setting incorporate several key features. First, as in the large empirical literature building on Berry (1994) and Berry, Levinsohn, and Pakes (1995, 2004), we emphasize the role of market-level demand shocks (unobservables at the level of the product \times market) that result in the econometric endogeneity of prices. Explicit accounting for these demand shocks is essential to the identification of policy-relevant features such as demand elasticities and equilibrium counterfactuals (see Berry and Haile (2021)). This drives our

²Instruments for quantities are what have sometimes been referred to informally as instruments for the “nonlinear parameters” in the applied literature using random coefficients discrete choice models. Berry and Haile (2021) provide additional discussion.

³This is related to well-known results regarding endogenous controls in regression models.

focus on market-level endogeneity and cross-market data, differentiating our work from much of the prior research on the identification of choice models with micro data.⁴ To our knowledge, the examination of identification with market-level data in Berry and Haile (2014) offers the only prior nonparametric identification results applicable to the workhorse models of this large empirical literature.

Second, the panel structure of consumers-within-markets is essential to the questions we ask. It is what distinguishes micro data from market-level data. This panel structure is responsible for the reduction in the number of needed instrumental variables, as well as the elimination of restrictions on the way product-level observables enter demand. These features contrast with the setting and model in Berry and Haile (2014). There the demand system (especially if combined with a model of supply) also connects to nonparametric simultaneous equations models, as studied by, e.g., Benkard and Berry (2006), Matzkin (2008, 2015), Blundell, Kristensen, and Matzkin (2013, 2020), and Berry and Haile (2018). However, the panel structure essential to the present paper is absent in all of that prior work.

Third, our model avoids requirements that consumer-level observables be exogenous, or that certain consumer observables be linked exclusively to the desirability of specific products. The latter requirement (often in combination with large support and exogeneity assumptions) is widely used in “special regressor” approaches to identification of consumer-level discrete choice models,⁵ but is often difficult to motivate in practice. More natural are situations in which multiple consumer-level observables interact to alter tastes for all goods. As a simple example—one illustrating a broader interpretation of “demand”—consider a discrete choice model of expressive voting in a two-party (“R” vs. “D”) election, applied to survey data matching individual reported votes to voter sociodemographics.⁶ Although voter-specific measures like age, income,

⁴This includes prior work on identification of discrete choice models allowing market-level demand shocks only through composite “error” terms—one for each choice—representing all latent heterogeneity (e.g., Lewbel (2000)). Explicit modeling of demand shocks also makes clear that the strong functional form assumptions permitting application of control function approaches (Blundell and Matzkin (2014)) generally fail, even in standard parametric models. See Berry and Haile (2014, 2021) for additional discussion of these issues.

⁵See the review by Lewbel (2014) and references therein. A very early version of this paper, (Berry and Haile (2010)), featured an example of such an approach. In practice, geographic distances are often modeled as providing consumer-level variation exclusive to each product. But even these are inherently restricted to lie on a 2-dimensional surface in \mathbb{R}_+^J , since the underlying consumer heterogeneity reflects only consumer locations.

⁶Advertising, rather than price, often plays the role of the endogenous choice characteristic—one whose effects are sometimes of primary interest. See, e.g., Gerber (1998) and Gordon and Hartmann (2013).

gender, race, and education may provide rich variation in preferences between the two parties (and the outside option to abstain), no such measure is naturally associated exclusively with the attractiveness of a single option.

Fourth, although we initially emphasize discrete choice demand, this is not essential. The primitive feature of interest in our analysis is a demand function mapping observables (at the level of market, products, and consumer) and a vector of market-level demand shocks to expected quantities demanded. This can allow continuous demand as well as departures from common assumptions regarding consumers' full information or rationality.

Of course, our results do require some structure, including conditions on sources of variation. In addition to instruments (for prices) satisfying standard conditions, we rely on three important assumptions. One is a nonparametric index restriction on the way market-level demand shocks and some observed consumer attributes enter the model.⁷ The second is injectivity of the mappings that link observed consumer attributes to choice probabilities. Below we connect these requirements to canonical specifications from the literature.

Finally, we require sufficient variation in the consumer observables to satisfy a “common choice probability” condition that we believe is new to the literature. This condition requires that the number of observed consumer attributes be at least as large as the number of products, and that they have sufficient independent variation. However, it contrasts with a standard “large support” condition, which would require that variation in such observables drive certain choice probabilities arbitrarily close to zero or one. Whereas large support would imply our common choice probability condition, the latter allows a broad range of cases where choice probabilities are never close to one or zero.⁸ An attractive feature of the common choice probability condition is that it is verifiable; i.e., its satisfaction or failure is identified.

Our assumptions on the dimension and variation of the observed consumer attributes are tightly related to the generality of our demand model which, e.g., places few restrictions substitution patterns or how price effects vary across products. Of course, some settings—particularly those with a large number of choices—may lack the dimension of variation we require for the most flexible

⁷In contrast to related conditions in Berry and Haile (2014, 2018) or Matzkin (2008), here each index depends on observed consumer attributes rather than observed product characteristics (which are fixed within markets), and there is no requirement that these observables be exogenous.

⁸In the voting example, our condition would require a vote share vector—say 0.4 for R and 0.4 for D (the remainder abstaining)—such that in every market (e.g., metro area) with the same pair of candidates and equal values of any metro-level observables, there is a combination of individual-level sociodemographic measures that generates this conditional vote share. The level of education etc. required to match the given vote share might be higher (and perhaps income lower, etc.) in an unobservably conservative market.

models. This motivates our exploration (in an appendix) of trade-offs between common modeling assumptions and the types/dimension of variation sufficient for identification. For example, we discuss semiparametric restrictions that can allow identification (even with a large choice set) with a single consumer-level observed attribute having only binary support.

Our results are relevant to a large empirical literature exploiting micro data to estimate demand. A classic example is McFadden’s study of transportation demand (McFadden, Talvitie, and Associates (1977)), where each consumer’s preferences over different modes of transport are affected by her available mode-specific commute times and other factors. This example illustrates an essential feature of the type of micro data considered here: consumer-specific observables that alter the relative attractiveness of different options. Consumer distances to different options have been used in a number of applications, including those involving demand for hospitals, retail outlets, residential locations, or schools, as in the examples of Capps, Dranove, and Satterthwaite (2003), Burda, Harding, and Hausman (2015), Bayer, Keohane, and Timmins (2009), and Neilson (2021). More broadly, observable consumer-level attributes that shift tastes for products might include income, sociodemographic measures, or other proxies for idiosyncratic preferences. For example, income and family size have been modeled as shifting preferences for cars (Goldberg (1995), Petrin (2002)); race, education, and birth state have been modeled as shifting preferences for residential location (Diamond (2016)). Other prominent examples include applications to demand for grocery products (Akerberg (2003)), newspapers (Gentzkow and Shapiro (2010)), neighborhoods (Bayer, Ferreira, and McMillan (2007)), and schools (Hom (2018)).⁹ An important feature of many examples, reflected by our model, is that the typical consumer-level observable cannot be tied exclusively to a single good.

In what follows, section 2 sets up our model of multinomial choice demand. Section 3 connects this model to random coefficients random utility specifications widely used in practice. We present our identification results in section 4. We discuss some key implications for applied work in section 5 before concluding in section 6. Appendix A discusses variations on our baseline model, including continuous demand and examples in which key assumptions in the text can be relaxed by strengthening others. Appendix B provides an examination of the proper excludability of standard instruments for prices when non-price observables are endogenous and uninstrumented.

⁹Here we cite only a small representative handful of papers out of a selection that spans many topics and many years. Recent development of commercial consumer-level data sets suggests the potential for micro data to play an even larger role in the future.

2 Model and Features of Interest

We consider choice among J goods (“products”) and an outside option (“good 0”) by consumers i in “markets” t . Formally, a market is defined by:

- a price vector $P_t = (P_{1t}, \dots, P_{Jt})$;
- a set of additional observables X_t ;
- a vector $\Xi_t = (\Xi_{1t}, \dots, \Xi_{Jt})$ of unobservables;
- a distribution $F_{YZ}(\cdot; t)$ of consumer observables $(Y_{it}, Z_{it}) \in \mathbb{R}^H \times \mathbb{R}^J$ with support $\Omega(X_t)$, for some $H \geq 0$.

For clarity, we write random variables (all of which have t among their subscript indices) in uppercase and their realizations in lowercase. The variables (P_t, X_t, Ξ_t) are common to all consumers in a given market.¹⁰ We distinguish between P_t and X_t due to the particular interest in how demand responds to prices and the typical focus on endogeneity of prices. However, we have not yet made the standard assumption that X_t is exogenous—e.g., independent or mean independent of the demand shocks Ξ_t . We will see below that identification of demand elasticities and other key features of demand can often be obtained without such an assumption (or additional instruments for X_t).¹¹

Although X_t will typically include observable product characteristics, it may also include other factors defining markets.¹² For example, consumers might be partitioned into “markets” based on a combination of geography, time, product availability, and demographics (average or individual-level) included in X_t . In contrast, observables varying across consumers within a market are represented by Y_{it} and Z_{it} . We make a distinction between Y_{it} and Z_{it} in order to isolate our requirements on consumer-level data. Key conditions, made precise below, are that consumer observables provide variation of dimension at least J (hence, $Z_{it} \in \mathbb{R}^J$) and that changes in Z_{it} alter the

¹⁰The assumption that all consumers in a market face the same prices and product characteristics is standard and may influence how markets are defined in practice. This rules out some forms of price discrimination.

¹¹Alternatively, when instruments are available for endogenous components of X_t , our results generalize immediately by expanding P_t to include these endogenous characteristics.

¹²Because X_t could include indicators for product availability, our treatment of J as fixed is without loss. With additional assumptions, variation in the number of goods available can be valuable; e.g., data from markets with J available goods could be used to predict outcomes in markets with more or fewer goods. X_t could also include product fixed effects. Such fixed effects generally do not address the endogeneity challenges central to identification of demand (see, e.g., Berry and Haile (2021)).

relative attractiveness of different goods. We do not require the additional consumer observables Y_{it} but can accommodate them in an unrestricted way; and conditioning on an appropriate value of Y_{it} can weaken some assumptions. Although our requirements on Z_{it} permit the case in which each component Z_{ijt} exclusively affects the attractiveness of good j , we will not require this. Nor will we require independence (full, conditional, or mean independence) between (Y_{it}, Z_{it}) and Ξ_t .

The choice environment of consumer i in market t is then represented by

$$C_{it} = (Z_{it}, Y_{it}, P_t, X_t, \Xi_t).$$

Let \mathcal{C} denote the support of C_{it} . A basic primitive characterizing consumer behavior in this setting is a distribution of decision rules for each $c_{it} \in \mathcal{C}$.¹³ As usual, heterogeneity in decision rules (i.e., nondegeneracy of the distribution) within a given choice environment may reflect a variety of factors, including latent preference heterogeneity across consumers, shocks to individual preferences, latent variation in consideration sets, or stochastic elements of choice (e.g., optimization error).¹⁴

2.1 Demand and Conditional Demand

The choice made by consumer i is represented by $Q_{it} = (Q_{i1t}, \dots, Q_{iJt})$, where Q_{ijt} denotes the quantity (here, 0 or 1) of good j purchased. Given C_{it} , a distribution of decision rules is fully characterized by the conditional cumulative joint distribution function $F_Q(q|C_{it}) = E[1\{Q_{it} \leq q\} | C_{it}]$. In the case of discrete choice, this distribution can be represented without loss by the structural choice probabilities

$$\delta(C_{it}) = (\delta_1(C_{it}), \dots, \delta_J(C_{it})) = E[Q_{it} | C_{it}]. \quad (1)$$

¹³Under additional conditions a distribution of decision rules can be represented as the result of utility maximization. See, e.g., Mas-Colell, Whinston, and Green (1995), Block and Marschak (1960), Falmagne (1978), and McFadden (2005). We will not require such conditions or consider a utility-based representation. A related issue is identification of welfare effects. Standard results allow construction of valid measures of aggregate welfare changes from a known demand system in the absence of income effects. Bhattacharya (2018) provides such results for discrete choice settings when income effects are present.

¹⁴For many purposes, one need not take a stand on the interpretation of this randomness, since the economic questions of interest involve changes to the arguments of demand functions, not to the functions themselves. This covers the canonical motivation for demand estimation: quantifying responses to *ceteris paribus* price changes. However, for some questions—e.g., those involving information interventions or requiring identification of cardinal utilities—the interpretation becomes important. See Barseghyan, Coughlin, Molinari, and Teitelbaum (2021) for a recent contribution on this topic.

Given the measure of consumers in each choice environment, the mapping \mathcal{s} fully characterizes consumer demand. We will therefore consider identification of the *demand* mapping \mathcal{s} on \mathcal{C} .

However, it is useful to also consider identification of the *conditional demand* functions

$$\bar{\mathcal{s}}(Z_{it}, Y_{it}, P_t; t) = \mathcal{s}(Z_{it}, Y_{it}, P_t, x_t, \xi_t)$$

on

$$\mathcal{C}(x_t, \xi_t) = \text{supp}(Z_{it}, Y_{it}, P_t) | \{X_t = x_t, \Xi_t = \xi_t\}$$

for each market t . The function $\bar{\mathcal{s}}(Z_{it}, Y_{it}, P_t; t)$ is simply the demand function \mathcal{s} when (X_t, Ξ_t) are fixed at the values (x_t, ξ_t) realized in market t . Because Ξ_t is unobserved and prices are fixed within each market, identification of $\bar{\mathcal{s}}(Z_{it}, Y_{it}, P_t; t)$ is nontrivial. However, this mapping fully characterizes the responses of demand (at all combinations of (Z_{it}, Y_{it})) to counterfactual *ceteris paribus* price variation, holding X_t and Ξ_t fixed at their realized values in market t . Thus, knowledge of $\bar{\mathcal{s}}(\cdot; t)$ for each market t suffices for many purposes motivating demand estimation in practice.

Notably, $\bar{\mathcal{s}}(\cdot; t)$ fully determines the own- and cross-price demand elasticities for all goods in market t . One implication is that $\bar{\mathcal{s}}(\cdot; t)$ is the feature of \mathcal{s} needed to discriminate between alternative models of firm competition (e.g., Berry and Haile (2014), Backus, Conlon, and Sinkinson (2021), Duarte, Magnolfi, Sølvesten, and Sullivan (2021)). And, given a model of supply, $\bar{\mathcal{s}}(\cdot; t)$ suffices to identify firm markups and marginal costs, following Berry, Levinsohn, and Pakes (1995) and Berry and Haile (2014); to decompose the sources of firms’ market power, as in Nevo (2001); to determine equilibrium outcomes under a counterfactual tax, tariff, subsidy, or exchange rate (e.g., Anderson, de Palma, and Kreider (2001), Nakamura and Zerom (2010), Decarolis, Polyakova, and Ryan (2020)); or to determine equilibrium “unilateral effects” of a merger (e.g., Nevo (2000), Miller and Sheu (2021)). Furthermore, $\bar{\mathcal{s}}(\cdot; t)$ alone determines the “diversion ratios” (e.g., Conlon and Mortimer (2021)) that often play a central role in the practice of antitrust merger review.

Of course, because the functions $\bar{\mathcal{s}}(\cdot; t)$ are defined with fixed values of (X_t, Ξ_t) , they do not suffice for answering all questions—in particular, those requiring knowledge of *ceteris paribus* effects of X_t on demand.¹⁵ However, by avoiding the need to separate the effects of X_t and Ξ_t , identification of $\bar{\mathcal{s}}(\cdot; t)$ in each market t can often be obtained without requiring exogeneity of X_t . This can be important when exogeneity is in doubt and one lacks the additional

¹⁵In some cases, such effects may be of direct interest—e.g., to infer willingness to pay for certain product features. In other cases, such effects are inputs to determination of demand under counterfactual product offerings or entry. Thus, while knowledge of $\bar{\mathcal{s}}(\cdot; t)$ in all markets suffices in a large fraction of applications, knowledge of \mathcal{s} is required for others.

instruments that would allow treating endogenous elements of X_t as we treat prices P_t below.

2.2 Core Assumptions

So far we have implicitly made three significant assumptions:

- (i) all latent heterogeneity across markets can be represented by a J -vector Ξ_t ;
- (ii) conditional on X_t , the support of (Y_{it}, Z_{it}) is the same in all markets; and
- (iii) the consumer-level observables Z_{it} have dimension J .

The first is important but standard (when market-level unobservables are acknowledged). The second seems mild for many applications and can be relaxed at the cost of more cumbersome exposition. The third will be more easily satisfied in applications with modest J , although some modern micro data sets can offer dozens or even hundreds of consumer-level observables. We focus on J -dimensional Z_{it} to explore the extent to which micro data can eliminate the need to instrument for J endogenous quantities. Appendix A illustrates how our requirements on the dimensions of Z_{it} and excluded instruments can be relaxed by strengthening other conditions.

We will also rely on Assumptions 1–4 below. These serve a key element of our strategy for demonstrating the gains from micro data: inferring variation across markets in the demand shock vector Ξ_t using variation in the value of the vector Z_{it} required to produce particular choice probabilities (conditional on other observables).

Assumption 1 (Index). $s(C_{it}) = \sigma(\gamma(Z_{it}, Y_{it}, X_t, \Xi_t), Y_{it}, P_t, X_t)$, with $\gamma(Z_{it}, Y_{it}, X_t, \Xi_t) = (\gamma_1(Z_{it}, Y_{it}, X_t, \Xi_t), \dots, \gamma_J(Z_{it}, Y_{it}, X_t, \Xi_t)) \in \mathbb{R}^J$.

Assumption 2 (Invertible Demand). $\sigma(\cdot, Y_{it}, P_t, X_t)$ is injective on the support of $\gamma(Z_{it}, Y_{it}, X_t, \Xi_t) | (Y_{it}, P_t, X_t)$.

Assumption 3 (Injective Index). $\gamma(\cdot, Y_{it}, X_t, \Xi_t)$ is injective on the support of $Z_{it} | (Y_{it}, X_t)$.

Assumption 4 (Separable Index). $\gamma_j(Z_{it}, Y_{it}, X_t, \Xi_t) = \Gamma_j(Z_{it}, Y_{it}, X_t) + \Xi_{jt}$ for all j .

As we illustrate below, these assumptions generalize standard specifications in the literature. Given (Y_{it}, P_t, X_t) , Assumption 1 requires that Z_{it} and Ξ_t affect choices only through indices $(\gamma_1(Z_{it}, Y_{it}, X_t, \Xi_t), \dots, \gamma_J(Z_{it}, Y_{it}, X_t, \Xi_t))$ that exclude P_t . This is a type of weak separability assumption. Observe that X_t and Y_{it} can affect demand both directly and through the indices, and that the indices themselves enter the function σ in fully flexible form. Assumption 2 requires that the choice probability function σ be “invertible” with respect

to the index vector—that, holding (Y_{it}, P_t, X_t) fixed, distinct index vectors map to distinct choice probabilities. Berry, Gandhi, and Haile (2013) provide sufficient conditions for invertibility of a demand system, which are natural here when each $\gamma_j(Z_{it}, Y_{it}, X_t, \Xi_t)$ can be interpreted as a (here, consumer-specific) quality index for good j .¹⁶ Assumption 3 requires injectivity of the index function γ with respect to the vector Z_{it} . This generalizes common utility-based specifications of demand while avoiding the common requirement that each Z_{ijt} affect the utility of good j exclusively. Assumption 4 requires each Ξ_{jt} to enter the index $\gamma_j(Z_{it}, Y_{it}, X_t, \Xi_t)$ additively. This provides a sense in which variation in Z_{it} has comparable effects across markets: a unit change in $\Gamma_j(Z_{it}, Y_{it}, X_t)$ has the same effect on demand as a unit change in Ξ_{jt} .¹⁷

Finally, as documented in Assumption 5, we will assume continuously distributed consumer-level observables Z_{it} .¹⁸ This allows transparent exploration of the potential gains from micro data, using calculus and moment equalities. Absent appropriate restrictions on the dimensionality of other elements of the model, it should not be surprising that continuous variation will be required for nonparametric point identification. Indeed, below we will require continuously distributed instruments as well. In practice, of course, one may often rely on at least some instruments or consumer-level observables with discrete (even binary) support. As in other types of empirical models, in such cases parametric forms used in estimation will typically fill the gaps left by the lack of (or limits on) continuous variation (see section A.1.4 for an illustration).

Assumption 5 (Support). *supp $Z_{it}|(Y_{it}, X_t)$ is open and connected.*

¹⁶Other general sufficient conditions for injectivity can be found in, e.g., Palais (1959), Gale and Nikaido (1965), and Parthasarathy (1983).

¹⁷We show in Appendix A that Assumption 4 can be dropped by strengthening other conditions.

¹⁸Given continuously distributed Z_{it} , the assumptions of open and connected support are technical conditions simplifying the arguments below.

2.3 A Useful Representation of the Index

Although we have thus far written the index vector

$$\gamma(Z_{it}, Y_{it}, X_t, \Xi_t) \equiv \Gamma(Z_{it}, Y_{it}, X_t) + \Xi_t$$

in a way that maximizes clarity about our core assumptions, for demonstrating identification (particularly in the case of endogenous X_t) it will be convenient to define

$$g(Z_{it}, Y_{it}, X_t) = \Gamma(Z_{it}, Y_{it}, X_t) + E[\Xi_t | X_t]$$

and

$$h(X_t, \Xi_t) = \Xi_t - E[\Xi_t | X_t], \quad (2)$$

so that the index can be rewritten as

$$\gamma(Z_{it}, Y_{it}, X_t, \Xi_t) = g(Z_{it}, Y_{it}, X_t) + h(X_t, \Xi_t). \quad (3)$$

Here we have simply let $g(Z_{it}, Y_{it}, X_t)$ absorb the mean of Ξ_t conditional on X_t , leaving the residualized structural error vector $h(X_t, \Xi_t)$. Observe that

$$E[h(X_t, \Xi_t) | X_t] = 0 \quad (4)$$

by construction.

With this notation, we have

$$s(C_{it}) = \sigma(g(Z_{it}, Y_{it}, X_t) + h(X_t, \Xi_t), Y_{it}, P_t, X_t) \quad (5)$$

and

$$\bar{s}(Z_{it}, Y_{it}, P_t; t) = \sigma(g(Z_{it}, Y_{it}, x_t) + h(x_t, \xi_t), Y_{it}, P_t, x_t) \quad (6)$$

We henceforth work with this representation of the demand and conditional demand functions.

2.4 Technical Conditions

Let \mathcal{X} denote the support of X_t . For $x \in \mathcal{X}$, let $\mathcal{Y}(x)$ denote the support of $Y_{it} | \{X_t = x\}$ and, for $y \in \mathcal{Y}(x)$, let $\mathcal{Z}(y, x) \subset \mathbb{R}^J$ denote the support of $Z_{it} | \{Y_{it} = y, X_t = x\}$. In parts (i)–(ii) of Assumption 6 we assume conditions permitting our applications of calculus and continuity arguments below. Part (iii) strengthens the injectivity requirements of Assumptions 2 and 3 slightly

by requiring that the Jacobian matrices $\partial g(z, y, x)/\partial z$ and $\partial \sigma(\gamma, y, p, x)/\partial \gamma$ be nonsingular almost surely.¹⁹

Assumption 6 (Technical Conditions). *For all $x \in \mathcal{X}$ and $y \in \mathcal{Y}(x)$,*

- (i) $g(z, y, x)$ is uniformly continuous in z on $\mathcal{Z}(y, x)$ and continuously differentiable with respect to z on $\mathcal{Z}(y, x)$;
- (ii) $\sigma(\gamma, y, p, x)$ is continuously differentiable with respect to γ for all $(\gamma, p) \in \text{supp}(\gamma(Z_{it}, Y_{it}, X_t, \Xi_t), P_t) | \{Y_{it} = y, X_t = x\})$; and
- (iii) $\partial g(z, y, x)/\partial z$ and $\partial \sigma(\gamma, y, p, x)/\partial \gamma$ are nonsingular almost surely on $\mathcal{Z}(y, x)$ and $\text{supp}(\gamma(Z_{it}, Y_{it}, X_t, \Xi_t), P_t) | \{Y_{it} = y, X_t = x\}$, respectively.

2.5 Normalization

The model requires two types of normalizations before the identification question can be properly posed. The first reflects the fact that the latent demand shocks have no natural location. Thus, we set $E[\Xi_t] = 0$ without loss. The second reflects the fact that any injective transformation of the index vector $\gamma(Z_{it}, Y_{it}, X_t, \Xi_t)$ can be reversed by appropriate modification of the unknown function σ . For example, take arbitrary $A(X_t) : \mathcal{X} \rightarrow \mathbb{R}^J$ and $B(X_t) : \mathcal{X} \rightarrow \mathbb{R}^{J \times J}$ ($B(x)$ invertible at all x). By letting

$$\begin{aligned} \tilde{\gamma}(Z_{it}, Y_{it}, X_t, \Xi_t) &= A(X_t) + B(X_t)\gamma(Z_{it}, Y_{it}, X_t, \Xi_t) \\ \tilde{\sigma}(\tilde{\gamma}(Z_{it}, Y_{it}, X_t, \Xi_t), Y_{it}, P_t, X_t) &= \sigma(B(X_t)^{-1}(\tilde{\gamma}(Z_{it}, Y_{it}, X_t, \Xi_t) - A(X_t)), Y_{it}, P_t, X_t) \end{aligned}$$

one obtains a new representation of the same distribution of decision rules (and thus same demand), the new one satisfying our assumptions whenever the original does. We must choose a single representation of demand before exploring whether the observables allow identification.²⁰

¹⁹Although we state Assumption 6 with the quantifier “for all $y \in \mathcal{Y}(x)$,” our arguments require these properties only at the arbitrary point $y^0(x)$ selected below. Observe also that, given Assumption 5 and part (i) of Assumption 6, the injectivity of $g(\cdot, Y_{it}, X_t)$ required by Assumption 3 implies (by invariance of domain) that the image $g(\mathcal{O}, y, x)$ of any open set $\mathcal{O} \subseteq \mathcal{Z}(y, x)$ is open. An implication is that even without part (iii) there could be no nonempty open set $\mathcal{O} \subseteq \mathcal{Z}(y, x)$ on which $\partial g(z, y, x)/\partial z$ was singular, as $g(\mathcal{O}, y, x)$ would then be a nonempty open subset of \mathbb{R}^J , contradicting Sard’s theorem. A similar observation applies to $\partial \sigma(\gamma, y, p, x)/\partial \gamma$.

²⁰Like location and scale normalizations of utility functions, our normalizations place no restriction on the demand function s or the conditional demand functions $\bar{s}(\cdot; t)$. However, our example illustrates an inherent ambiguity in the interpretation of how a given variable alters preferences. For example, in terms of consumer behavior (e.g., demand), there is no difference between a change in Z_{ijt} (all else fixed) that makes good j more desirable and a change in Z_{ijt} that makes all other goods (including the outside good) less desirable. In practice, this ambiguity is often resolved with *a priori* exclusion assumptions—e.g., an assumption that Z_{ijt} affects only the utility obtained from good j . Such assumptions could only aid identification. See the additional discussion in Appendix A.

To do this, for each x we take an arbitrary $(z^0(x), y^0(x))$ from the support of $(Z_{it}, Y_{it}) | \{X_t = x\}$. We then select the representation of demand in which

$$g(z^0(x), y^0(x), x) = 0 \quad \forall x. \quad (7)$$

and

$$\left. \frac{\partial g(z, y^0(x), x)}{\partial z} \right|_{z=z^0(x)} = I \quad \forall x, \quad (8)$$

where I denotes the J -dimensional identity matrix.

In the example above this choice of normalization is equivalent to taking

$$B(x) = \left[\left. \frac{\partial g(z, y^0(x), x)}{\partial z} \right]^{-1} \right|_{z=z^0(x)}$$

and

$$A(x) = -B(x)g(z^0(x), y^0(x), x)$$

at each x , then dropping the tildes from the transformed model.

3 Example

The literature includes many examples of parametric specifications that are special cases of our model. A canonical discrete choice demand model is generated from a random coefficients random utility specification²¹ like

$$u_{ijt} = x_{jt}\beta_{ijt} - \alpha_{it}p_{jt} + \xi_{jt} + \epsilon_{ijt}, \quad (9)$$

where u_{ijt} is consumer i 's conditional indirect utility from good j in market t . The idiosyncratic taste shock ϵ_{ijt} is usually specified as a draw from a type-1 extreme value or normal distribution. Components k of the random coefficient vector β_{ijt} are often specified as

$$\beta_{ijt}^{(k)} = \beta_{0j}^{(k)} + \sum_{\ell=1}^L \beta_{zj}^{(k,\ell)} z_{i\ell t} + \beta_{\nu j}^{(k)} \nu_{it}^{(k)}, \quad (10)$$

where each $z_{i\ell t}$ represents an observable characteristic of individual i , and each $\nu_{it}^{(k)}$ is a random variable with a pre-specified distribution. Often, the

²¹Random coefficients, represented here by terms in (10) and (11), are popular in discrete choice models because they allow for greater (if still limited) flexibility in the substitution patterns permitted by the resulting demand system. We focus directly on identification of a demand system with very flexible substitution patterns. Section 5.1 provides additional discussion.

coefficient on price is also specified as varying with some observed consumer characteristics y_{it} , such as income.²² A typical specification of α_{it} is

$$\ln(\alpha_{it}) = \alpha_0 + \alpha_y y_{it} + \alpha_\nu \nu_{it}^{(0)}. \quad (11)$$

The consumer-level preference shocks $(\nu_{i0t}, \dots, \nu_{iJt}, \epsilon_{i0t}, \dots, \epsilon_{iJt})$ are assumed i.i.d. across consumers and markets.

With (10) and (11), we can rewrite (9) as

$$u_{ijt} = g_j(z_{it}, x_t) + \xi_{jt} + \mu_{ijt}, \quad (12)$$

where

$$g_j(z_{it}, x_t) = \sum_k x_{jt}^{(k)} \sum_{\ell=1}^L \beta_{zj}^{(k,\ell)} z_{i\ell t} \quad (13)$$

$$\mu_{ijt} = \sum_k x_{jt}^{(k)} \left(\beta_{0j}^{(k)} + \beta_{\nu j}^{(k)} \nu_{it}^{(k)} \right) - p_{jt} \exp(\alpha_0 + \alpha_y y_{it} + \alpha_\nu \nu_{it}^{(0)}) + \epsilon_{ijt}. \quad (14)$$

Observe that all effects of z_{it} and ξ_t operate through indices

$$\gamma_j(z_{it}, x_t, \xi_t) = g_j(z_{it}, x_t) + \xi_{jt} \quad j = 1, \dots, J,$$

satisfying our Assumptions 1 and 4. It is easy to show that the resulting choice probabilities satisfy Berry, Gandhi and Haile’s (2013) “connected substitutes” condition with respect to the vector of indices $(\gamma_1(z_{it}, x_t, \xi_t), \dots, \gamma_J(z_{it}, x_t, \xi_t))$; therefore, the injectivity of demand required by Assumption 2 holds. Our core assumptions require $L \geq J$.²³ Invertibility of the linear mapping $g(z_{it}, x_t) = (g_1(z_{it}, x_t), \dots, g_J(z_{it}, x_t))$ in z_{it} (Assumption 3) might then be assumed as a primitive condition of the model or derived from other conditions.

This example connects our nonparametric model to a large number of applications. Of course, our model generalizes the example substantially. It does not require linear utility functions,²⁴ parametric distributional assumptions, or even a representation of demand through utility maximization. Even within the linear random coefficients random utility discrete choice paradigm, our

²²Our model would permit y_{it} to affect the random coefficients $\beta_{ijt}^{(k)}$ as well, reintroducing y_{it} as an argument of each g_j defined in (13). Exclusion of z_{it} from α_{it} is one way to satisfy the weak separability requirement of Assumption 1 and may be interpreted as defining which variables serve as z_{it} . However, more general specifications consistent with our model would allow, e.g., arbitrary interactions between p_{jt} , y_{it} , x_t , and the index $g_j(z_{it}, y_{it}, x_t) + \xi_{jt}$.

²³If $L > J$, we can combine the “extra” components of Z_{it} with income to redefine the partition of consumer observables as (Y_{it}, Z_{it}) with $Z_{it} \in \mathbb{R}^J$.

²⁴Thus, it would allow generalization of the nonparametric random utility model in Allen and Rehbeck (2019) to incorporate market-level demand shocks, flexible heterogeneity in tastes for product characteristics, and nonadditive product-level taste shocks.

model would allow the joint distribution of $(\nu_{i0t}, \dots, \nu_{iJt}, \epsilon_{i0t}, \dots, \epsilon_{iJt})$ to depend on $(g(z_{it}, y_{it}, x_t) + \xi_t, y_{it}, x_t, p_t)$. More generally, consumer heterogeneity in our model is not limited to a finite vector of shocks entering with particular functional forms: utilities could be specified as nonparametric random functions of $(g(z_{it}, y_{it}, x_t) + \xi_t, y_{it}, x_t, p_t)$.

Finally, observe that the example above lacks features sometimes relied on in results showing identification of discrete choice models: aside from the absence of individual characteristics that exclusively affect the utility from one choice j , this model does not exhibit independence between the “error term” $(\xi_{jt} + \mu_{ijt})$ in (12) and any of the observables z_{it}, x_t, p_t .²⁵

4 Identification

We consider identification of the demand system

$$\mathcal{D}(C_{it})$$

and the conditional demand systems

$$\bar{\mathcal{D}}(Z_{it}, Y_{it}, P_t; t).$$

The observables comprise the market index t , the variables $(Q_{it}, Z_{it}, Y_{it}, P_t, X_t)$, and a vector of instruments W_t discussed further below. As usual, to consider identification we treat the population joint distribution of the observables as known. Loosely, we may view this as the result of observing $(Q_{it}, Z_{it}, Y_{it}, P_t, X_t)$ for many consumers i in each of many markets t .²⁶ These observables imply observability of choice probabilities conditional on $(Z_{it}, Y_{it}, P_t, X_t)$ in each market t . Of course, this implies observability of market-level choice probabilities (market shares) as well.

Because our arguments do not require variation in Y_{it} , in much of what follows we will fix Y_{it} (conditional on X_t) at $y^0(X_t)$. We proceed in three steps. First, in section 4.1 we demonstrate identification of the function $g(\cdot, y^0(x), x)$ at each $x \in \mathcal{X}$. Second, in section 4.2 we use this result to link latent market-level variation in $h(X_t, \Xi_t)$ to variation in the observed value of Z_{it} required to produce a given conditional choice probability in each market. In particular, given instruments for prices, we show that the realized values $h(x_t, \xi_t)$ can be

²⁵Even aside from ξ_{jt} , the variables x_{jt} and p_{jt} enter the composite error μ_{ijt} . Furthermore, x_{jt} and p_{jt} may be correlated with changes in the distribution of z_{it} across markets, introducing dependence between z_{it} and μ_{ijt} .

²⁶Note that, unlike the case of a consumer panel (see, e.g., Berry and Haile (2021)), the population of consumers is different for each t .

pinned down in every market, making identification of the conditional demand systems $\bar{y}(\cdot; t)$ in each market straightforward. Finally, in section 4.3 we show that \bar{y} is also identified when one adds the usual assumption that X_t is exogenous. Thus, after the initial setup and lemmas, the main results themselves follow relatively easily.

Before proceeding, we provide some key definitions and observations. For $(p, x, \xi) \in \text{supp}(P_t, X_t, \Xi_t)$ let

$$\mathcal{S}(p, x, \xi) = \sigma(g(\mathcal{Z}(y^0(x), x), y^0(x), x) + h(x, \xi), y^0(x), p, x)). \quad (15)$$

Thus, $\mathcal{S}(p, x, \xi)$ denotes the support of choice probabilities in any market t for which $P_t = p, X_t = x$, and $\Xi_t = \xi$ (holding $Y_{it} = y^0(x)$).²⁷ By Assumptions 2 and 3, for each $s \in \mathcal{S}(p, x, \xi)$ there must be a unique $z^* \in \mathcal{Z}(y^0(x), x)$ such that $\sigma(g(z^*, y^0(x), x) + h(x, \xi), y^0(x), p, x)) = s$. So for $(p, x, \xi) \in \text{supp}(P_t, X_t, \Xi_t)$ and $s \in \mathcal{S}(p, x, \xi)$, we define the function

$$z^*(s; p, x, \xi)$$

implicitly by

$$\sigma(g(z^*(s; p, x, \xi), y^0(x), x) + h(x, \xi), y^0(x), p, x)) = s. \quad (16)$$

These definitions lead to two observations that play key roles in what follows. First, in each market t the set $\mathcal{S}(p_t, x_t, \xi_t)$ and the values of $z^*(s; p_t, x_t, \xi_t)$ for all $s \in \mathcal{S}(p_t, x_t, \xi_t)$ are observed, even though the value of the argument ξ_t is not. That is, if the choice probability s is in the support of those within market t , there must be some consumer attribute vector for which this choice probability is implied. This attribute vector is unique under our assumptions, is observable, and is equal to $z^*(s; p_t, x_t, \xi_t)$ by definition. Second, by the invertibility of σ (Assumption 2), we have

$$g(z^*(s; p, x, \xi), y^0(x), x) + h(x, \xi) = \sigma^{-1}(s; y^0(x), p, x) \quad (17)$$

for all $(p, x, \xi) \in \text{supp}(P_t, X_t, \Xi_t)$ and $s \in \mathcal{S}(p, x, \xi)$.

4.1 Initial Steps

Let $\|\cdot\|$ denote the Euclidean norm. We will require the following nondegeneracy condition.

²⁷Note that because $\mathcal{Z}(y^0(x), x)$ is open, continuity and injectivity of σ with respect to the index vector and of the index function with respect to Z_{it} imply (by invariance of domain) that $\mathcal{S}(p, x, \xi)$ is open.

Assumption 7 (Nondegeneracy). *For each $x \in \mathcal{X}$, there exists (possibly unknown) $p \in \text{supp } P_t | \{X_t = x\}$ such that $\text{supp } \Xi_t | \{P_t = p, X_t = x\}$ contains an open subset of \mathbb{R}^J .*

Assumption 7 requires continuously distributed Ξ_t but is otherwise mild, ruling out trivial cases in which conditioning on (P_t, X_t) indirectly fixes Ξ_t as well. Such cases are ruled out by standard models of supply, where prices respond to continuous cost shifters or markup shifters (observed or unobserved), allowing the same equilibrium price vector p to arise under different realizations of Ξ_t . A key implication follows from the definition (2): for each $x \in \mathcal{X}$ there exist $\epsilon > 0$ and $p \in \text{supp } P_t | \{X_t = x\}$ such that for any $d \in \mathbb{R}^J$ satisfying $\|d\| < \epsilon$, $\text{supp } \Xi_t | \{P_t = p, X_t = x\}$ contains vectors ξ and ξ' satisfying $h(x, \xi) - h(x, \xi') = d$. This is exploited to prove Lemma 1, which demonstrates how local variation across markets in the latent Ξ_t will produce local variation in observed values of $z^*(s; p_t, x_t, \xi_t)$ for those markets.

Lemma 1. *Let Assumptions 1–7 hold. For each $x \in \mathcal{X}$, there exist $p \in \text{supp } P_t | \{X_t = x\}$ and $\Delta > 0$ such that for all z and z' in $\mathcal{Z}(y^0(x), x)$ satisfying*

$$\|z' - z\| < \Delta, \quad (18)$$

there exist a choice probability vector s and vectors ξ and ξ' in $\text{supp } \Xi_t | \{P_t = p, X_t = x\}$ such that $z = z^(s; p, x, \xi)$ and $z' = z^*(s; p, x, \xi')$. Furthermore, such (Δ, p) are identified.*

Proof. See Appendix C □

With this result in hand, Lemma 2 demonstrates that one can use equation (17) to relate partial derivatives of $g(z, y(x), x)$ at any point z to those at nearby points z' by examining the change in consumer characteristics required to create a given change in the vector of choice probabilities.²⁸

Lemma 2. *Let Assumptions 1–7 hold. Then for every $x \in \mathcal{X}$ there exists a known $\Delta > 0$ such that for almost all $z, z' \in \mathcal{Z}(y^0(x), x)$ satisfying (18) the matrix $\left[\frac{\partial g(z, y^0(x), x)}{\partial z} \right]^{-1} \left[\frac{\partial g(z', y^0(x), x)}{\partial z} \right]$ is identified.*

Proof. Given any $x \in \mathcal{X}$, take a (known) (p, Δ) as in Lemma 1. Consider markets t and t' in which $P_t = P_{t'} = p$ but, for some choice probability vector \hat{s} ,

$$z = z^*(\hat{s}; p, x, \xi_t) \neq z' = z^*(\hat{s}; p, x, \xi_{t'}), \quad (19)$$

²⁸Abusing notation to simplify key expressions, below we write $\frac{\partial g(z, y^0(x), x)}{\partial z}$ to represent the Jacobian matrix $\left. \frac{\partial g(\hat{z}, y^0(x), x)}{\partial \hat{z}} \right|_{\hat{z}=z}$. Similarly, we write $\frac{\partial g(z', y^0(x), x)}{\partial z}$ to represent $\left. \frac{\partial g(\hat{z}, y^0(x), x)}{\partial \hat{z}} \right|_{\hat{z}=z'}$.

revealing that $\xi_t \neq \xi_{t'}$. Lemma 1 ensures that such t, t' , and \hat{s} exist for all $z, z' \in \mathcal{Z}(y^0(x), x)$ satisfying (18). And although ξ_t and $\xi_{t'}$ are latent, the identities of markets t and t' satisfying (19) are observed, as are the associated values of \hat{s} , $z^*(\hat{s}; p, x, \xi_t)$, and $z^*(\hat{s}; p, x, \xi_{t'})$. Differentiating (17) with respect to the share vector within these two markets, we obtain

$$\frac{\partial g(z, y^0(x), x)}{\partial z} \frac{\partial z^*(\hat{s}; p, x, \xi_t)}{\partial s} = \frac{\partial \sigma^{-1}(\hat{s}; y^0(x), p, x)}{\partial s} \quad (20)$$

and

$$\frac{\partial g(z', y^0(x), x)}{\partial z} \frac{\partial z^*(\hat{s}; p, x, \xi_{t'})}{\partial s} = \frac{\partial \sigma^{-1}(\hat{s}; y^0(x), p, x)}{\partial s}. \quad (21)$$

Thus, recalling Assumption 6, for almost all such (z, z') we have

$$\left[\frac{\partial g(z', y^0(x), x)}{\partial z} \right]^{-1} \frac{\partial g(z, y^0(x), x)}{\partial z} = \frac{\partial z^*(\hat{s}; p, x, \xi_{t'})}{\partial s} \left[\frac{\partial z^*(\hat{s}; p, x, \xi_t)}{\partial s} \right]^{-1}.$$

The matrices on the right-hand side are observed. \square

This leads us to the main result of this section, obtained by connecting (for each value of x) the matrix products $\left[\frac{\partial g(z, y^0(x), x)}{\partial z} \right]^{-1} \left[\frac{\partial g(z', y^0(x), x)}{\partial z} \right]$ identified in Lemma 2 to the known (normalized) value of the matrix $\left[\frac{\partial g(z, y^0(x), x)}{\partial z} \right]$ at $z = z^0(x)$.

Lemma 3. *Under Assumptions 1–7, $g(\cdot, y^0(x), x)$ is identified on $\mathcal{Z}(y^0(x), x)$ for all $x \in \mathcal{X}$.*

Proof. For $\epsilon > 0$, let $\mathcal{B}(b, \epsilon)$ denote an open ball in \mathbb{R}^J of radius ϵ , centered at b . Take any $x \in \mathcal{X}$ and associated $\Delta > 0$ as in Lemma 2. For each vector of integers $\tau \in \mathbb{Z}^J$, define the set

$$\mathcal{B}_\tau = \mathcal{Z}(y^0(x), x) \cap \mathcal{B}\left(z^0(x) + \frac{\tau\Delta}{J}, \frac{\Delta}{2}\right).$$

By construction, all z and z' in any given set \mathcal{B}_τ satisfy (C.2). So by Lemma 2, the value of

$$\left[\partial g(z, y^0(x), x) / \partial z \right]^{-1} \left[\partial g(z', y^0(x), x) / \partial z \right] \quad (22)$$

is known for almost all z and z' in any set \mathcal{B}_τ . Because $\cup_{\tau \in \mathbb{Z}^J} \mathcal{B}_\tau$ forms an open cover of $\mathcal{Z}(y^0(x), x)$, given any $z \in \mathcal{Z}(y^0(x), x)$ there exists a simple chain of open sets \mathcal{B}_τ in $\mathcal{Z}(y^0(x), x)$ linking the point $z^0(x)$ to z .²⁹ Thus,

$$\left[\partial g(z, y^0(x), x) / \partial z \right]^{-1} \left[\partial g(z^0(x), y^0(x), x) / \partial z \right]$$

²⁹See, e.g., van Mill (2002, Lemma 1.5.21).

is known for almost all $z \in \mathcal{Z}(y^0(x), x)$. With the normalization (8) and the continuity of $\partial g(z, y^0(x), x)/\partial z$ with respect to z , the result then follows from the fundamental theorem of calculus for line integrals and the boundary condition (7). \square

Before moving to identification of conditional demand, we pause to point out that our constructive identification of $g(\cdot, y^0(x), x)$ used only a single price vector p at each value of x —that required by Assumption 7. In typical models of supply this condition would hold for almost all price vectors in the support of $P_t|\{X_t = x\}$. In addition to providing falsifiable restrictions, this indicates a form of redundancy that would typically be exploited by estimators used in practice. Similarly, our proof of Lemma 3 used, for each $z \in \mathcal{Z}(y^0(x), x)$, only one of infinitely many paths between z^0 and z ; integrating along any such path must yield the same function $g(\cdot, y^0(x), x)$ at each x .

4.2 Identification of Conditional Demand

We demonstrate identification of the conditional demand functions $\bar{y}(\cdot; t)$ under two additional conditions. The first is a requirement of sufficient variation in the consumer-level observables Z_{it} .

Assumption 8 (Common Choice Probability). *For each $x \in \mathcal{X}$, there exists a choice probability vector $s^*(x)$ such that $s^*(x) \in \mathcal{S}(p, x, \xi)$ for all $(p, \xi) \in \text{supp}(P_t, \Xi_t) |\{X_t = x\}$.*

Assumption 8 requires that, at each $x \in \mathcal{X}$, there exist some choice probability vector $s^*(x)$ that is common to all markets—that

$$\bigcap_{(p, \xi) \in \text{supp}(P_t, \Xi_t) |\{X_t = x\}} \mathcal{S}(p, x, \xi)$$

be nonempty. The nondegeneracy of each set $\mathcal{S}(p_t, x_t, \xi_t)$ (recall (15)) reflects variation in Z_{it} across its support. Assumption 8 requires enough variation in Z_{it} that for some $s^*(x)$ we have $s^*(x) \in \mathcal{S}(p_t, x, \xi_t)$ for all (p_t, ξ_t) in their support conditional on $X_t = x$.

The strength of this assumption depends on the joint support of (P_t, Ξ_t) given $\{X_t = x\}$ and on the relative impacts of (Z_{it}, Ξ_t, P_t) on choice behavior. Observe that P_{jt} and Ξ_{jt} typically will have opposing impacts and will be positively dependent conditional on X_t under equilibrium pricing behavior; thus, large support for $g(Z_{it}, y^0(x), x)$ may not be required even if Ξ_t were to have large support. Indeed, we can contrast our assumption with a requirement of special regressors with large support: the latter would imply that *every* interior choice probability vector s is a common choice probability for all x ;

we require only a single common choice probability at each x . Because choice probabilities conditional on (Z_{it}, Y_{it}) are observable in all markets, Assumption 8 is verifiable.³⁰ And, because the choice of each $y^0(x)$ was arbitrary, it implies that we require only existence (for each x) of one such $y^0(x) \in \mathcal{Y}$ such that Assumption 8 holds.³¹ Importantly for what follows, the values of any common choice probability vectors $s^*(x)$ may be treated as known.

Our second requirement is existence of instruments for prices satisfying the standard nonparametric IV conditions.

Assumption 9 (Instruments for Prices).

- (i) $E[h_j(X_t, \Xi_{jt})|X_t, W_t] = E[h_j(X_t, \Xi_{jt})|X_t]$ almost surely for all $j = 1, \dots, J$;
- (ii) In the class of functions $\Psi(X_t, P_t)$ with finite expectation, $E[\Psi(X_t, P_t)|X_t, W_t] = 0$ almost surely implies $\Psi(X_t, P_t) = 0$ almost surely.

Part (i) of Assumption 9 is the exclusion restriction, requiring that variation in W_t not alter the mean of the latent $h(X_t, \Xi_t)$ conditional on X_t . Recall that $E[h(X_t, \Xi_t)|X_t] = 0$ by construction; thus part (i) implies

$$E[h_j(X_t, \Xi_{jt})|X_t, W_t] = 0 \quad \text{a.s. for all } j. \quad (23)$$

This is true regardless of whether X_t itself is exogenous. Of course, one must be particularly cautious about satisfaction of part (i) when X_t is thought to be endogenous. We discuss this further in section 5.3 and Appendix B. The relevance requirement of part (ii) is a standard completeness condition. This is the nonparametric analog of the rank condition required for identification of linear regression models. For example, Newey and Powell (2003) have shown that under mean independence (the analog of (23) here), completeness is necessary and sufficient for identification in separable nonparametric regression.³² The following result demonstrates that, given existence of a common choice probability vector s^* , the same instrumental variables conditions suffice here to allow identification of $h_j(x_t, \xi_{jt})$ for all j and t .

Lemma 4. *Under Assumptions 1–9, the scalar $h_j(x_t, \xi_{jt})$ is identified for all j and t .*

Proof. In each market t , taking $x = x_t, p = p_t, \xi = \xi_t$ and $s = s^*(x_t)$ in equation (17), we have

$$g(z^*(s^*(x_t); p_t, x_t, \xi_t), y^0(x_t), x_t) = \sigma^{-1}(s^*(x_t); y^0(x_t), p_t, x_t) - h(x_t, \xi_t).$$

³⁰See Berry and Haile (2018) for a formal definition of verifiability.

³¹When more than one such value $y^0(x)$ exists, or when there is more than one common choice probability vector $s^*(x)$, this introduces additional falsifiable restrictions.

³²See also, e.g., Florens and Rolin (1990), Chernozhukov and Hansen (2005), and Severini and Tripathi (2006).

Thus, for all t and each $j = 1, \dots, J$,

$$g_{jt} = f_j(x_t, p_t) - e_{jt} \quad (24)$$

where we have defined $e_{jt} \equiv h_j(x_t, \xi_{jt})$, $f_j(x_t, p_t) \equiv \sigma_j^{-1}(s^*(x_t); y^0(x_t), p_t, x_t)$, and $g_{jt} \equiv g_j(z^*(s^*(x_t); p_t, x_t, \xi_t), y^0(x_t), x_t)$. By Lemma 3 each g_{jt} on the left side of (24) is known (recall that the values of each $z^*(s^*(x_t); p_t, x_t, \xi_t)$ are observable, even though the value of each ξ_t is not). Thus, for each j this equation takes the form of a standard separable nonparametric regression model with exogenous regressors X_t , endogenous regressors P_t , and additive structural errors E_{jt} . By (23), we have $E[E_{jt}|X_t, W_t] = 0$ a.s. So under the completeness condition (part (ii) of Assumption 9) identification of each function f_j follows immediately from the identification result (Proposition 2.1) of Newey and Powell (2003). This implies identification of each e_{jt} (i.e., each $h_j(x_t, \xi_{jt})$) as well. \square

Identification of the conditional demand functions $\bar{\beta}(\cdot; t)$ now follows easily.

Theorem 1. *Under Assumptions 1–9, $\bar{\beta}(\cdot; t)$ is identified on $\mathcal{C}(x_t, \xi_t)$ for all t .*

Proof. Recall that

$$\begin{aligned} \bar{\beta}(Z_{it}, Y_{it}, P_t; t) &= s(Z_{it}, Y_{it}, P_t, x_t, \xi_t) \\ &= \sigma(g(Z_{it}, Y_{it}, x_t) + h(x_t, \xi_t), Y_{it}, P_t, x_t) \\ &= E[Q_{it}|Z_{it}, Y_{it}, P_t, x_t, h(x_t, \xi_t)]. \end{aligned}$$

Because $Q_{it}, Z_{it}, Y_{it}, P_t, X_t$ are observed and each $h(x_t, \xi_t)$ is known, the result follows. \square

We emphasize that although the conditional demand functions $\bar{\beta}(\cdot; t)$ are indexed by t , this merely stands in for the values of X_t and $h(X_t, \Xi_t)$. Within a single market, there is no price variation. However, Lemma 4 allows us to utilize information from all markets with given values of X_t and $h(X_t, \Xi_t)$ to reveal how price variation affects demand at all $(Z_{it}, Y_{it}, P_t, X_t, h(X_t, \Xi_t))$ in their joint support.

4.3 Identification of Demand

As discussed already, knowledge of the conditional demand functions suffices for a large fraction of the questions motivating demand estimation, but not all. In particular, it is not sufficient to answer questions concerning effects of X_t on demand or other counterfactual outcomes when X_t changes holding Ξ_t fixed. Addressing such questions will require separating the impacts of X_t and Ξ_t . This can be done by adding the standard assumption that X_t is exogenous.

Assumption 10 (Exogenous Product Characteristics). $E[\Xi_t|X_t] = 0$.

When Assumption 10 holds, the definition (2) implies

$$h(X_t, \Xi_t) = \Xi_t.$$

This has two important implications. First, when Assumption 10 is maintained, the IV exclusion condition (part (i) of Assumption 9) softens to require instruments W_t that are exogenous conditional on exogenous (rather than endogenous) X_t . Second, Lemma 4 now implies that each realization ξ_t of the demand shock vector is identified. Recalling that

$$s(C_{it}) = E[Q_{it}|Z_{it}, Y_{it}, P_t, X_t, \Xi_t],$$

identification of s follows immediately from the facts that $(Q_{it}, Z_{it}, Y_{it}, P_t, X_t)$ are observed and all realizations of Ξ_t are now known.

Theorem 2. *Under Assumptions 1–10, s is identified on \mathcal{C} .*

5 Lessons for Applied Work

Although the study of identification is formally a theoretical exercise, a primary motivation is to provide guidance for the practice and evaluation of demand estimation in applied work. Here we discuss some key messages.

5.1 The Incremental Value of Micro Data

The most important practical lesson from our results is that the marginal value of micro data is high. The specific benefits of micro data concern some of the most significant challenges to identification of demand when one has only market-level data: (i) the need to instrument for all prices and quantities, and (ii) the nonparametric functional form and exogeneity conditions that allow some of these IV requirements (in particular, the proper excludability of BLP instruments) to be satisfied.

The gains from micro data reflect the fact that consumer-level observables create within-market variation in consumers' choice problems. Such variation is similar in some ways to that which can be generated by instruments for quantities (see footnote 2 and Berry and Haile (2014)). In particular, it can pin down key aspects of consumer substitution patterns. From (5), we see that $\frac{\partial s}{\partial Z_{it}} = \frac{\partial \sigma}{\partial \gamma} \frac{\partial g}{\partial Z_{it}}$. Since $\frac{\partial s}{\partial Z_{it}}$ is observed, identification of g (which we demonstrated without instruments) implies identification of the derivatives of demand with respect to the index vector γ (and, thus, with respect to the vector of demand shocks Ξ_t). In standard parametric models like the example of

section 3, these substitution patterns—and those with respect to prices—are determined (conditional on observables) by the joint distribution of the random coefficients and product-level taste shocks. Our nonparametric model, of course, allows more flexible substitution patterns, and our results show that micro data allows their identification without any instruments—a stark contrast to the case of market-level data alone.

Critically, however, the reason micro-data variation is free from confounding effects of demand shocks is not an assumption of exogeneity—indeed, Z_{it} was *not* assumed to be independent or mean independent of Ξ_t . Rather, this follows from the fact that market-level demand shocks Ξ_t do not vary within a market. This has a strong connection to the “within” identification of slope parameters in panel data models with fixed effects.

Thus, researchers should prefer micro data and seek it out whenever possible. Collecting reliable micro data will sometimes be difficult, and in some cases only limited forms of micro data may be available. But even when the setting and assumptions permit use of BLP instruments—or when the micro data available are more limited than we have assumed to explore fully nonparametric identification—variation from micro data can be a powerful addition. This message is consistent, for example, with the findings in the empirical literature (e.g., Petrin (2002), Berry, Levinsohn, and Pakes (2004)) that the addition of even limited forms of individual-level data can result in much more precise estimates than those obtained with market-level data alone.

5.2 The Necessity of Cross-Market Variation

Although within-market variation accounts for the advantages of micro data, cross-market variation remains essential. The proof of Lemma 2, for example, relied on variation in Ξ_t across markets in the key steps toward identification of g and, thus, of the substitution patterns discussed in section 5.1. More fundamentally, the demand system (5) depends on arguments (X_t, P_t, Ξ_t) that have no variation within a market. Thus, without strong additional restrictions, data from a single market cannot reveal anything about the effects of prices or product characteristics on demand.³³

This observation serves as a caution. As a practical matter, a parametric

³³For example, with data from a single market τ , one could set $\Xi_\tau = 0$ without loss, assume that X_t has no effect on demand, and assume that demand follows the multinomial logit model with money-metric mean utilities $g_j(Z_{it}) - p_{jt}$ for each good j . One can then fit the observed conditional choice probabilities $s_\tau(z_{i\tau})$ in market τ perfectly by setting $g(z_{i\tau}) = \sigma_{MNL}^{-1}(s_\tau(z_{i\tau})) + p_\tau$, where σ_{MNL}^{-1} is the inverse share function for the multinomial logit. The fitted model’s implications regarding effects of X_t and P_t on demand, of course, reflect only the *a priori* assumptions, and neither these nor other arbitrary assumptions can be ruled out using data on a single market.

specification of demand may allow estimation using data from only one market, exploiting a combination of functional form restrictions and cross-product variation in prices and product characteristics.³⁴ In some cases, as in the classic work of McFadden, Talvitie, and Associates (1977), only a single market is available for study. However, identification in such cases—indeed, the ability to rule out any arbitrary model of how prices (and other product characteristics) affect demand—will be possible only through additional *a priori* restrictions that could be relaxed in a multi-market setting.

5.3 A Focus on Instruments

Given sufficiently rich micro data, our main requirement for identification is a set of valid instruments for prices. Candidate instruments include most of those typically relied upon in the case of market-level data: these include cost shifters, proxies for cost shifters (e.g., “Hausman instruments”), and exogenous shifters of market structure. Micro data can also make available a related category of candidate instruments: market-level observables (e.g., average demographics) that alter equilibrium markups, the “Waldfoegel instruments.”³⁵ With micro data, one can directly account for the impacts of individual-specific demographics, so it can be reasonable to assume that market-level demographics are properly excluded from the demand mapping \mathcal{D} .

An important and subtle question is whether the required IV exclusion condition (part (i) of Assumption 9) will hold when X_t is endogenous. In Appendix B we find that, depending on the instrument and model of endogeneity, conditioning on endogenous X_t can (i) render otherwise-valid instruments for prices invalid; (ii) render otherwise-invalid instruments valid; or (iii) have no effect on instrument validity. In some cases where conditioning on endogenous X_t causes a violation of the exclusion condition, the problem can be “fixed” through natural timing assumptions. Overall, this analysis reveals that one must carefully examine the exclusion condition when X_t is endogenous. As Appendix B illustrates, causal graphs offer a tool for such examination that is simultaneously transparent, robust, and formal.

Absent from the discussion above are the BLP instruments. The characteristics X_{-jt} of products competing with good j have direct effects on demand for good j and (in standard models) on good j ’s markup. The BLP instruments are thus relevant shifters of prices and in our model they are not needed as exogenous shifters of quantities. Of course, the excludability of X_{-jt} requires

³⁴Absent additional restrictions on our nonparametric model, such cross-product variation does not contribute to identification of demand.

³⁵See Waldfoegel (2003) as well as Gentzkow and Shapiro (2010), Fan (2013), and Li, Hartmann, and Amano (2020).

not only their exogeneity but also a restriction on the way they enter demand (Berry and Haile (2014, 2021)). Appendix A shows how such restrictions can allow BLP instruments for prices in the micro data case.

5.4 What Does Not Follow

Nonparametric identification results demonstrate a particular sense in which parametric assumptions are not essential, but this does not mean that parametric (or other) assumptions relied on in practice can be ignored. Functional form restrictions can constrain the answers to key questions, and empirical demand research should continue to explore the sensitivity of estimates to functional form choices. Likewise, it remains important to explore new (parametric, semiparametric, or nonparametric) estimation approaches. Our nonparametric identification results ensure that such explorations are possible and may even suggest new estimation strategies.

We also emphasize that our sufficient conditions for nonparametric identification should not be viewed as necessary conditions for demand estimation in practice, but should rather guide our thinking about the strength of the available data and empirical results. Nonparametric identification of economic models (even regression models) relies on assumptions—index assumptions, separability assumptions, completeness conditions, support conditions, monotonicity conditions, or other shape restrictions—that will often (perhaps typically) fall short of full satisfaction in practice. Conditions for nonparametric identification are not a hurdle but an ideal—a point of reference that can guide our quest for and aid our assessment of the best available empirical evidence.

6 Conclusion

Since Berry, Levinsohn, and Pakes (1995), there has been an explosion of interest in empirical demand models that incorporate both flexible substitution patterns and explicit treatment of the demand shocks that introduce endogeneity. Understandably, this development has been accompanied by questions about identification of these models. Our results offer a reassurance that identification follows from traditional sources of quasi-experimental variation in the form of instrumental variables and panel-style within-market variation. This reassurance is particularly important because of the wide relevance of these models to economic questions and the special identification challenges arising in the case of demand (see Berry and Haile (2021)).

Furthermore, identification of these models is not fragile. It does not rely on “identification-at-infinity” arguments; it is not limited to particular types of settings (e.g., random utility discrete choice); one can substitute one type

of variation for another (e.g., replacing instruments for quantities with micro-data variation), depending on the type of data available; and one can relax many key conditions by strengthening others. Thus, although this is a case where identification results come well after an extensive empirical literature has already developed, the nonparametric foundation for this literature is strong. Of course, not all applications will offer the combinations of micro data and instrumental variables permitting nonparametric identification. But even when such data limitations lead to greater reliance on functional form restrictions, our results shed light on the roles such assumptions will play.

Appendices

A Extensions, Variations, and Robustness

A.1 Alternative Modeling Assumptions

Here we explore several variations on our baseline model and associated identification conditions. This includes straightforward extensions to models of continuous or mixed discrete-continuous demand. Our results here show that identification is robust in the sense that a relaxation of one condition for identification often can be accommodated by strengthening another. And an understanding of these trade-offs will often be helpful to both producers and consumers of research relying on demand estimates. A full exploration of such trade-offs describes an entire research agenda. However, here we illustrate some possibilities that enlarge the set of potential instruments, reduce the number of required instruments, allow a nonseparable index structure, or reduce the required dimensionality of the micro data.

For simplicity, our discussion will consider the traditional case in which X_t is exogenous, focusing then on identification of demand rather than conditional demand. Recall that in this case we have $h(X_t, \Xi_t) = \Xi_t$. Given our focus on the role of Z_{it} , we will also fix and suppress any additional consumer-level observables Y_{it} in what follows.

A.1.1 Strengthening the Index Structure

The model used by Berry and Haile (2014) to study identification with market-level data restricted the way some elements of X_t enter: for each good j , one element of X_{jt} affects demand only through the j th element of an index vector. Such a restriction is common in practice, and adding it here can allow the use of BLP instruments for prices.³⁶

To illustrate this as simply as possible, partition X_t as $(X_t^{(1)}, X_t^{(2)})$, where $X_t^{(1)} = (X_{1t}^{(1)}, \dots, X_{Jt}^{(1)}) \in \mathbb{R}^J$. Suppose demand takes the form

$$s(Z_{it}, P_t, X_t, \Xi_t) = \sigma\left(\gamma(Z_{it}, X_t, \Xi_t), P_t, X_t^{(2)}\right), \quad (\text{A.1})$$

where for $j = 1, \dots, J$

$$\gamma_j(Z_{it}, X_t, \Xi_t) = g_j(Z_{ijt}, X_t^{(2)}) + \eta_j(X_{jt}^{(1)}, X_t^{(2)}) + \Xi_{jt}. \quad (\text{A.2})$$

³⁶As suggested in section 5.3, the key issue is proper excludability of these instruments, not their relevance.

Here we (a) restrict $X_t^{(1)}$ to enter only through the index vector; (b) associate the j th components of Z_{it} and $X_t^{(1)}$ exclusively with the j th element of the index vector; and (c) impose additive separability between Z_{ijt} and $X_{jt}^{(1)}$ within each index.³⁷ Many specifications in the literature satisfy these requirements, typically with additional restrictions such as linear substitution between Z_{ijt} and $X_{jt}^{(1)}$. We will also strengthen the common choice probability condition to require existence of a common choice probability vector $s^*(X_t)$ that does not vary with $X_t^{(1)}$.³⁸ For simplicity we also assume here that, for each $x^{(2)}$ in the support of $X_t^{(2)}$ there is some point $z^0(x^{(2)})$ common to $\mathcal{Z}(x, x^{(2)})$ for all x in the support of $X_t^{(1)}|\{X_t^{(2)} = x^{(2)}\}$.

For the remainder of this section we will condition on $X_t^{(2)}$, suppress it from the notation, and let X_t represent $X_t^{(1)}$.³⁹ With this more restrictive model, we can simplify our normalizations. First, because adding a constant κ_j to g_j and subtracting the same constant from η_j would leave the model unchanged, we take an arbitrary $x^0 \in \mathcal{X}$ and set

$$\eta_j(x_j^0) = 0 \quad \forall j. \quad (\text{A.3})$$

Even with (A.3) (and our maintained $E[\Xi_t] = 0$), it remains true that any linear (or other injective) transformation of the index γ_j could offset by an appropriate adjustment to the function σ , yielding multiple representations of the same demand system (recall the related observation in section 2.5). Thus, without loss, we normalize the location and scale of each index γ_j by setting $g_j(z_j^0) = 0$ and $\frac{\partial g_j(z_j^0)}{\partial z_j} = 1$.

The arguments in Lemmas 1–3 will now demonstrate identification of each function g_j . At the common choice probability vector s^* , the inverted demand system takes the form of equations

$$g_j(z_j^*(s^*)) + \eta_j(x_{jt}) + \xi_{jt} = \sigma_j^{-1}(s^*; p_t)$$

for each j . Writing the j th equation as

$$g_j(z_j^*(s^*)) = -\eta_j(x_{jt}) + \sigma_j^{-1}(s^*; p_t) - \xi_{jt}, \quad (\text{A.4})$$

³⁷Exclusivity of $X_{jt}^{(1)}$ to the index γ_j is essential to the point we illustrate here, and this is most natural when exclusivity of each Z_{ijt} differentiates the elements of the index vector. As in our more general model, however, the elements of $\gamma(Z_{it}, X_t, \Xi_t)$ need not be linked to particular goods.

³⁸Formally, we assume that for each $x^{(2)} \in \text{supp } X_t^{(2)}$, there exists a choice probability vector $s^*(x^{(2)})$ such that for all $x^{(1)} \in \text{supp } X_t^{(1)}|\{X_t^{(2)} = x^{(2)}\}$, $s^*(x^{(2)}) \in \mathcal{S}(p, (x^{(1)}, x^{(2)}), \xi)$ for all $(p, \xi) \in \text{supp } (P_t, \Xi_t)|\{X_t = (x^{(1)}, x^{(2)})\}$.

³⁹Conditioning on $X_t^{(2)}$ treats it fully flexibly, since the same argument can be applied at each value of $X_t^{(2)}$.

we obtain a nonparametric regression equation with RHS variables x_{jt} and p_t . In this equation x_{-jt} is excluded, offering $J - 1$ potential instruments for the endogenous prices p_t . Thus, one additional instrument—e.g., a scalar market-level cost shifter or Waldfoegel instrument—would yield enough instruments to obtain identification of the unknown RHS functions and the “residuals” ξ_{jt} . Once the demand shocks are identified, identification of demand follows immediately.

Many variations on this structure are possible. For example, as in many empirical specifications, one might assume that p_{jt} enters demand only through the j^{th} index. This can lead to a regression equation (the analog of (A.4)) of the form

$$g_j(z_j^*(s^*)) = -\eta_j(x_{jt}, p_{jt}) + \sigma_j^{-1}(s^*) - \xi_{jt}.$$

Now only one instrument for price is necessary. For example, the BLP instruments can overidentify demand.

A.1.2 A Nonparametric Special Regressor

A different approach is to assume that the demand system of interest is generated by a random utility model with conditional indirect utilities of the form

$$U_{ijt} = g_j(Z_{ijt}) + \Xi_{jt} + \mathcal{E}_{ijt},$$

where \mathcal{E}_{ijt} is a scalar random variable whose nonparametric distribution depends on X_{jt} and P_{jt} (equation (14) gives a parametric example). In this case, our Lemma 3 demonstrates identification of each function $g_j(\cdot)$ up to a normalization of utilities.

Adding the assumption of independence between Z_{ijt} and \mathcal{E}_{ijt} then turns $g_j(Z_{ijt})$ into a known special regressor. Under a further (and typically very strong) large support assumption on $g_j(Z_j)$, standard arguments demonstrate identification of the marginal distribution of $(\Xi_{jt} + \mathcal{E}_{ijt})|(X_t, P_t)$. This is not sufficient to identify demand. However, one can use these marginal distributions to define a nonparametric IV regression equation for each choice j , where the LHS is a conditional mean and Ξ_{jt} appears on the RHS as an additive structural error.⁴⁰ In each of these equations the prices and characteristics of goods $k \neq j$ are excluded. Identification of these equations identifies all demand shocks, and identification of demand then follows as in Theorem 2. Thus, in this framework one needs only one instrument for price, and exogenous characteristics of competing goods (BLP IVs) would be available as instruments.

⁴⁰See the early working paper version of this paper, Berry and Haile (2010).

A.1.3 A Nonseparable Index

Our results in the text exploited additive separability of the index functions γ_j that enter the fully flexible demand mapping σ (i.e., Assumption 4). Here we show one way that separability can be dropped.⁴¹ Suppose that for each j the index $\gamma_j(Z_{it}, Y_{it}, X_t, \Xi_t)$ takes the nonseparable form $\gamma_j(Z_{ijt}, Y_{it}, X_t, \Xi_{jt})$, where γ_j is strictly increasing in both Z_{ijt} and Ξ_{jt} . We again fix $Y_{it} = y^0(X_t)$ and condition on X_t , suppressing them in the notation. Then if s^* is a common choice probability vector, for every market t there exists a vector z_t^* such that for all j ,

$$\sigma_j(\gamma^*, p_t) = s_j^*,$$

where $\gamma^* = (\gamma_1(z_{1t}^*, \xi_{jt}), \dots, \gamma_J(z_{Jt}^*, \xi_{Jt}))$.

By injectivity of σ , we have

$$\gamma_j(z_{jt}^*, \xi_{jt}) = \sigma_j^{-1}(s_j^*; p_t)$$

for all j and t . Since γ_j is strictly increasing in both arguments, we can take the partial inverse of both sides, yielding

$$z_{jt}^* = \gamma_j^{-1}(\sigma_j^{-1}(s_j^*; p_t); \xi_{jt}),$$

where γ_j^{-1} is strictly decreasing in ξ_{jt} . Rewriting this as

$$z_{jt}^* = \psi(p_t, \xi_{jt}),$$

we have an equation taking the form of the nonparametric nonseparable regression model considered by Chernozhukov and Hansen (2005), who showed (their Theorem 4) identification given instruments for prices that are independent of ξ_{jt} and satisfy an appropriate completeness condition.

A.1.4 Semiparametric Models

Moving further in the direction of parametric models commonly used in practice is one way to reduce both the required dimensionality of consumer attributes and the number of required instruments. Indeed, in some cases a single instrument and a single binary-valued consumer observable Z_{it} can suffice.

Consider first a semi-parametric nested logit model where inverse demand in market t , given z_{it} , is

$$g_j(z_{it}) + \xi_{jt} = \ln(s_{jt}(z_{it})/s_{0t}(z_{it})) - \theta \ln(s_{j/n,t}(z_{it})) + \alpha p_{jt}. \quad (\text{A.5})$$

⁴¹In the context of market-level data, Berry and Haile (2014) include related results relaxing additive separability. See also Matzkin (2015) and Blundell, Kristensen, and Matzkin (2020). None of these covers the panel structure of the micro data setting we consider here.

Here we have conditioned on X_t and suppressed it from the notation.⁴² We have also let $s_{jt}(z_{it})$ denote good j 's (observable) choice probability in market t conditional on z_{it} , with $s_{j/n,t}(z_{it})$ denoting its within-nest conditional choice probability. The scalar θ denotes the usual “nesting parameter.” Here we consider a one-dimensional binary Z_{it} taking values 0 and 1.

As with the standard representation of most parametric models of inverse demand, the nested logit model embeds normalizations of the indices and demand function analogous to our choices of $A(x)$ and $B(x)$ in section 2.5. However, we must still normalize the location of either Ξ_{jt} or g_j for each j to pose the identification question. Here we will set $g_j(0) = 0$ for all j , breaking with our prior convention by leaving each $E[\Xi_{jt}]$ free.

Here (A.5) implies the two equations

$$g_j(1) + \xi_{jt} = \ln(s_{jt}(1)/s_{0t}(1)) - \theta \ln(s_{j/n,t}(1)) + \alpha p_{jt} \quad (\text{A.6})$$

$$g_j(0) + \xi_{jt} = \ln(s_{jt}(0)/s_{0t}(0)) - \theta \ln(s_{j/n,t}(0)) + \alpha p_{jt} \quad (\text{A.7})$$

for every market t and product j . Differencing these equations (and recalling that $g_j(0) = 0$), we obtain

$$g_j(1) = \ln\left(\frac{s_{jt}(1)}{s_{0t}(1)}\right) - \ln\left(\frac{s_{jt}(0)}{s_{0t}(0)}\right) - \theta [\ln(s_{j/n,t}(1)) - \ln(s_{j/n,t}(0))]. \quad (\text{A.8})$$

This yields one equation in the two unknowns, $g_j(1)$ and θ . As with our fully nonparametric model, within-market variation does not suffice for identification on its own.

However, moving now to a different market, t' , where the observed choice probabilities are different (perhaps because $\xi_t \neq \xi_{t'}$), we have

$$g_j(1) = \ln\left(\frac{s_{jt'}(1)}{s_{0t'}(1)}\right) - \ln\left(\frac{s_{jt'}(0)}{s_{0t'}(0)}\right) - \theta [\ln(s_{j/n,t'}(1)) - \ln(s_{j/n,t'}(0))]. \quad (\text{A.9})$$

Given minimal variation in choice probabilities across markets, ensuring that

$$\ln(s_{j/n,t}(1)) - \ln(s_{j/n,t}(0)) \neq \ln(s_{j/n,t'}(1)) - \ln(s_{j/n,t'}(0)),$$

(A.8) and (A.9) can be solved for the two unknowns. Identification of the remaining parameter α can then be obtained from the “regression” equation (obtained from (A.5))

$$g_j(z_{it}) = \ln(s_{jt}(z_{it})/s_{0t}(z_{it})) - \theta \ln(s_{j/n,t}(z_{it})) + \alpha p_{jt} - \xi_{jt} \quad (\text{A.10})$$

⁴²This again treats X_t fully flexibly. Here, let conditional indirect utilities take the form

$$u_{ijt} = u(x_t, g_j(z_{it}, x_t) + \xi_{jt} - \alpha(x_t) p_{jt} + \mu_{ijt}(x_t)),$$

where u is strictly increasing in its second argument, $\alpha(x_t)$ is arbitrary, and $\mu_{ijt}(x_t)$ is a stochastic component taking the standard composite nested-logit form at each x_t . The identification argument sketched here may be repeated at each value of x_t in its support.

using a single excluded instrument for price—e.g., an excluded exogenous market-level cost shifter or markup shifter that affects all prices. The intercept in this regression equation will be $E[\Xi_{jt}]$.

Although this example involves a model that is more flexible than nested logit models typically estimated in practice, it moves a considerable distance from our fully nonparametric model. But it makes clear that additional structure can further reduce the dimension of the required exogenous variation. Here one can obtain identification with a single instrument and a scalar binary individual-level observable z_{it} . This compares to the usual requirement of two instruments in the fully parametric nested logit when one has only market-level data (see Berry (1994)). Thus, as in the fully nonparametric case, micro data cuts the number of required instruments by half.

Other semiparametric models can offer more intermediate points in the set of feasible trade-offs between the flexibility of the model and the dimension of exogenous variation needed for identification. Indeed, it is easy to see how the example here generalizes to settings where Z_{it} has more than two points of support. Each additional point in the support of Z_{it} adds one additional equation but also one additional unknown $g_j(z)$. Thus, cross-market variation remains necessary (recall section 5.2). However, when we turn to cross-market variation, as the number of markets and number of points in the support of Z_{it} increase, we obtain an increasing number of restrictions that can allow identification of more complex semi-parametric models, including cases where the inverse share function is described by multiple scalar parameters. Semi-parametric generalizations of BLP-style random coefficient logit models offer one class of examples.

A.2 Beyond Discrete Choice

Although we have focused on the case in which the consumer-level quantities Q_{ijt} take the particular form implied by a discrete choice model, nothing in our proofs requires this. In other settings, the demand function s defined in (1) may simply be reinterpreted as the expected vector of quantities demanded conditional on $(X_t, P_t, \Xi_t, Z_{it}, Y_{it})$.⁴³ Applying our results to continuous demand is therefore just a matter of verifying the suitability of our assumptions.⁴⁴

As one possibility, consider a “mixed CES” model of continuous choice,

⁴³Note that the demand faced by firms in market t is the expectation of this expected demand over the joint distribution of (Z_{it}, Y_{it}) in the market.

⁴⁴Berry, Gandhi, and Haile (2013) describe a broad class of continuous choice models that can satisfy the key injectivity property of Assumption 2. These can include mixed continuous/discrete settings, where individual consumers may purchase zero or any positive quantity of each good.

similar to the model in Adao, Costinot, and Donaldson (2017), with $J + 1$ products. Here we reintroduce Y_{it} to denote consumer i 's income, measured in units of the numeraire good 0. Each consumer i in market t has utility over consumption vectors $q \in \mathbb{R}_+^{J+1}$ given by

$$u(q; z_{it}, x_t, p_t, \xi_t) = \left(\sum_{j=0}^J \phi_{ijt} q_j^\rho \right)^{1/\rho},$$

where $\rho \in (0, 1)$ is a parameter and each ϕ_{ijt} represents idiosyncratic preferences of consumer i . Normalizing $\phi_{i0t} = 1$, let

$$\phi_{ijt} = \exp[(1 - \rho)(g_j(z_{it}, x_t) + \xi_{jt} + x_{jt}\beta_{it})], \quad j = 1, \dots, J,$$

where β_{it} is a random vector with distribution F representing consumer-level preferences for product characteristics. With $p_{0t} = 1$, familiar CES algebra shows that Marshallian demands are

$$q_{ijt} = \frac{y_{it} \exp(g_j(z_{it}, x_t) + \xi_{jt} + x_{jt}\beta_{it} - \alpha \ln(p_{jt}))}{1 + \left[\sum_{k=1}^J \exp(g_k(z_{it}, x_t) + \xi_{kt} + x_{kt}\beta_{it} - \alpha \ln(p_{kt})) \right]}, \quad (\text{A.11})$$

where $\alpha = 1/(1 - \rho)$. Equation (A.11) resembles a choice probability for a random coefficients logit model, although the quantities q_{it} here take on continuous values and do not sum to one. It is easy to show that our Assumptions 1–4 are satisfied for the expected CES demand functions, which take the form

$$\sigma_t(g(z_{it}, x_t) + \xi_t, y_{it}, x_t, p_t) = E[Q_{it} | z_{it}, y_{it}, p_t, x_t, \xi_t],$$

where the j th component of $E[Q_{it} | z_{it}, y_{it}, x_t, p_t, \xi_t]$ is

$$\int \frac{y_{it} \exp(g_j(z_{it}, x_t) + \xi_{jt} + x_{jt}\beta_{it} - \alpha \ln(p_{jt}))}{1 + \left[\sum_{k=1}^J \exp(g_k(z_{it}, x_t) + \xi_{kt} + x_{kt}\beta_{it} - \alpha \ln(p_{kt})) \right]} dF(\beta_{it}).$$

B Instruments When X_t is Endogenous

In section 5.3 we discussed several categories of instruments W_t commonly relied upon to provide exogenous variation in prices. Here we examine the question of when such instruments remain properly excluded when conditioning on observables X_t that are not (mean) independent of Ξ_t . Such instruments are required for Theorem 1 to apply when Theorem 2 does not, allowing identification of conditional demand without requiring exogeneity of X_t (or instruments for X_t). Unconditional independence is not sufficient (or necessary) for conditional independence. Indeed, it is well known that conditioning on an

endogenous “control” variable can (e.g., in the case of linear regression) lead to violation of independence conditions required for identification.

In what follows we suppress the market subscripts t on the random variables X_t, P_t, W_t, Ξ_t , etc. Our discussion will utilize graphical causal models, with the d-separation theorem providing the key criterion for assessing the independence between W and Ξ conditional on X .⁴⁵ Our use of these tools is elementary, and a graphical approach is not essential—causal graphs represent information implied by a fully specified model, standard change-of-variables formulas, and Bayes’ rule. However, the graphical approach allows transparent treatment of many possible economic examples inducing a smaller number of canonical dependence structures. It also can be highly clarifying when one ventures beyond the simplest cases. Following the literature on graphical causal models, we focus on full conditional independence,

$$W \perp\!\!\!\perp \Xi \mid X, \tag{B.1}$$

which of course implies the conditional mean independence required by Theorem 1.

Below we first discuss several causal graphs (and motivating economic examples) that “work”—i.e., that imply (B.1). We then discuss the main type of structure that does not work—i.e., in which (B.1) fails despite unconditional independence between W and Ξ . We will see that each type of instrument discussed in section 5.3 can remain valid under several models of endogenous X . However, each type of instrument can also fail; in particular, (B.1) will fail despite unconditional independence between W and Ξ when firms choose X in ways that depend on both W and Ξ (or their ancestors). However, in many of these situations, a natural timing assumption can yield a new set of valid instruments for prices.

B.1 Graphs that Work

B.1.1 Fully Exogenous Instruments

The simplest cases arise when the instruments W satisfy


$$W \perp\!\!\!\perp (X, \Xi). \tag{B.2}$$

The conditional independence condition (B.1) is then immediate, regardless of any dependence between X and Ξ . Although formal analysis is unnecessary in this case, it is also easily illustrated to build toward less obvious cases.

⁴⁵See, e.g., Pearl (2009) and Pearl, Glymour, and Jewell (2016), including references therein. Throughout we maintain the standard assumption that nodes in a causal directed acyclic graph are independent of their nondescendants conditional on their parents.

For example, suppose X is chosen by firms with knowledge of Ξ , so that X is endogenous in the same sense that prices are. Given (B.2), one obtains the causal graph shown in Figure 1.⁴⁶ The conditional independence condition (B.1) then can also be seen to follow immediately by the d-separation criterion. We would reach the same conclusion if the direction of causation between Ξ and X is reversed—e.g., if X is chosen without knowledge of Ξ but the distribution of Ξ changes with the choice of X . Taking the classic example of demand for automobiles, a manufacturer’s choice to offer a fuel efficient hybrid sedan may imply a very different set of relevant unobserved characteristics than had a pickup truck or luxury SUV been offered instead.⁴⁷

Figure 1



Most of the instrument types discussed in section 5.3 can satisfy (B.2). For example, W could represent exogenous (independent of Ξ) cost shifters such as input price shocks, realized after X is chosen and not affected by X . These might be shocks to import tariffs; shipping costs; retailer costs (e.g., rents, wages); demand shifters in other markets served by the same firms (if they face upward sloping marginal costs); or prices of manufacturing inputs. One can also obtain this structure when W represents exogenous shifters of markups. Mergers (full or partial) that are independent of Ξ and leave product offerings unchanged offer one possibility. Another is cross-market variation in the distribution $F_{YZ}(\cdot|t)$ (or other aggregate demographic measure at the market or regional level), as long as this variation is independent of Ξ (as required generally for the validity of Waldfoel instruments) and X .

B.1.2 Instruments Caused by X


Independence between X and W is not required. For example, consider the case in which X is chosen with knowledge of Ξ and the choice of X affects

⁴⁶We assume throughout that prices and quantities are not among the ancestors of (X, W, Ξ) . This is implied by standard assumptions that consumers take X and Ξ as given when making purchase decisions, that W does not respond to prices or quantities, and that prices are not chosen before X .

⁴⁷A similar structure is obtained when dependence between Ξ and X reflects a common cause.

W . We then obtain a causal graph in Figure 2, where (B.1) is again easily confirmed by the d-separation criterion. This causal structure allows additional examples of cost shifters beyond those discussed above. For example, suppose X represents product characteristics affecting the level of labor skill (or quality of another input) required in production, while W is the producer’s average wage. Alternatively, if producers have market power in input markets, input prices W would be affected by firms’ choices of product characteristics X . Models in which the the direction of causation between X and Ξ in Figure 2 reverses will lead to the same conclusions.


Figure 2



B.1.3 X Caused by Instruments

In some cases, the conditional independence condition (B.1) can hold even when X is affected by W . Consider the causal graph in Figure 3, where (B.1) is easily confirmed by d-separation. As an example motivating this structure, suppose W is a product-level cost of producing a product feature measured by X , the latter chosen by firms with knowledge of W but before Ξ is known.

Figure 3




B.1.4 Hausman Instruments

When X is not independent of Ξ , Hausman instruments generally cannot yield the trivial case in which (B.2) holds. This is because prices set by a firm in some “other” market depend on the product characteristics in that market, and the firm’s product characteristics are typically (highly) correlated across markets. Nonetheless, Hausman instruments can remain valid in some cases.

Figure 4 illustrates one such case. Here L represents latent marginal cost shifters. We let X_{-t} and Ξ_{-t} represent the non-price observables and demand shocks of “other markets,” both of which (along with L) affect the Hausman instruments W (prices in those markets).⁴⁸ The absence of an edge linking Ξ and Ξ_{-t} reflects an essential assumption justifying Hausman instruments in general (i.e., even when X is exogenous), as does the absence of an edge directly linking L and Ξ . Here W and Ξ are not independent. But the conditional independence condition (B.1) is satisfied.

Figure 4




The causal structure of Figure 4 is consistent with a fully specified model in which X is chosen in each market with knowledge of the latent cost shocks but before Ξ is realized. The direction of causality between X and L (and similarly between X_{-t} and L) is not important to this conclusion. However, as demonstrated in the following section, the direction of causality between Ξ and X is typically critical in the case of Hausman instruments.

B.2 Graphs that Don’t Work: X is a Collider

The conditional independence condition (B.1) fails when both W and Ξ affect X . This is illustrated in Figure 5. Here X is a collider in the (undirected) path between W and Ξ . Thus, although Ξ and W are independent, (B.1) fails.

Figure 5




⁴⁸Following standard convention, we use a dashed bidirectional edge to represent dependence between X and X_{-t} arising from unmodeled common causes.

This structure arises whenever firms' choices of X depend on both W and Ξ (or their ancestors). An example is when W is a cost shifter affecting firms' choices of X , the latter also chosen with knowledge of Ξ . Another example is when W is a market-level demographic measure or market structure measure (e.g., product ownership matrix) that, along with Ξ , influences firms' choices of X .⁴⁹

As suggested already, we can also obtain this type of structure with Hausman instruments. Figure 6 illustrates. This graph again represents models in which prices and product characteristics X are chosen with knowledge of the demand shocks Ξ and the latent cost shifters L . Here, regardless of presence or direction of causation between X and L (the graph shows one possibility), X is a collider on an unblocked path between Ξ and W .

Figure 6



Thus, just as there are cases in which each type of instrument discussed in section 5.3 remains valid when conditioning on endogenous characteristics X , there are other important cases in which (B.1) will fail. In such situations, identification will require different instruments for prices. In many cases such instruments can be constructed under natural timing assumptions. This is a topic we take up in the final section of this appendix.⁵⁰

B.3 Avoiding Colliders: Sequential Timing

The previous section describes a class of situations in which candidate instruments that would be properly excluded unconditional on X would fail to be properly excluded conditional on X . A leading case is that of cost shifters (e.g., input prices) that, along with Ξ (or its ancestors), partially determine firms' choices of product characteristics X . However, in such cases one may be able to obtain valid instruments by exploiting the (typical) sequential timing of a firm's decisions. For example, physical characteristics of new automobiles

⁴⁹A similar structure arises if the dependence between Ξ and X reflects a latent common cause.

⁵⁰We also note that when X is a collider, W provides a candidate instrument for X .


sold in year τ will reflect design choices made well in advance—in particular, before the input costs for year- τ production are fully known. Pricing in year τ , on the other hand, will typically take place after those costs are known. Such timing is common to many markets. And, as in other contexts, the temporal separation of observable choices can offer an identification strategy.⁵¹ Here, for example, even if product characteristics are chosen in response to demand shocks and expected input costs, current-period *innovations* to input costs can offer candidate instruments for prices.

To illustrate, we introduce a time superscript τ to all random variables. Let M^τ denote a vector of period- τ input prices and suppose M^τ follows

$$M^\tau = \Phi(M^{\tau-1}) + W^\tau, \quad (\text{B.3})$$

where Φ is a possibly unknown function and $W^\tau \perp\!\!\!\perp (\Xi^\tau, X^\tau, M^{\tau-1})$. Given observability of $(M^\tau, M^{\tau-1})$ in all markets, each vector of period- τ innovations W^τ is identified. Now suppose that X^τ is chosen by firms in period $\tau - 1$, whereas prices for period τ are chosen in period τ . The causal graph in Figure 7 illustrates key features of such a model.⁵²

Figure 7



Here endogeneity of X^τ reflects its selection with knowledge of $\Xi^{\tau-1}$, the latter correlated with Ξ^τ . Neither the contemporaneous cost shifters M^τ nor the lagged cost shifters $M^{\tau-1}$ can serve as instruments for prices conditional on X^τ : X^τ would be a collider, as in the previous section. However, the period- τ innovation W^τ can serve as the instrument. Because W^τ alters period- τ marginal cost, it is relevant for the determination of P^τ , conditional on

⁵¹Familiar examples in IO include strategies used by Olley and Pakes (1996), Akerberg, Caves, and Frazer (2015), and others in the literature on estimation of production functions.

⁵²The presence (or direction) of an edge from X^τ to Ξ^τ is not important to the argument. Likewise, although we show the case in which $\Xi^{\tau-1}$ is a cause of Ξ^τ , the same conclusion is reached if dependence between $\Xi^{\tau-1}$ and Ξ^τ reflects unmodeled common causes.

X^τ . And, by the d-separation criterion, we see that W^τ is independent of Ξ^τ conditional on X^τ . Indeed, W^τ here is an example of a “fully exogenous instrument,” as discussed in section B.1.1. Ultimately, the innovation W^τ is simply a cost shifter that is independent of all else. The important insight, however, is that natural timing assumptions can allow such fully independent cost shifters to be constructed from measures like input prices that themselves are not independent of Ξ^τ conditional on X^τ . Similar arguments can allow construction of valid instruments from observed markup shifters (e.g., market-level demographics) whose lagged values affect firms’ choices of X . Indeed, one may simply reinterpret M^τ above as a period- τ markup shifter.

C Proof of Lemma 1

Fix a value of $x \in \mathcal{X}$. By Assumption 7, there exist $p \in \text{supp } P_t | \{X_t = x\}$ and $\epsilon > 0$ such that for any z and z' in $\mathcal{Z}(y^0(x), x)$ for which

$$\|g(z', y^0(x), x) - g(z, y^0(x), x)\| < \epsilon, \quad (\text{C.1})$$

there exist ξ and ξ' in $\text{supp } \Xi_t | \{P_t = p, X_t = x\}$ such that

$$h(x, \xi) - h(x, \xi') = g(z', y^0(x), x) - g(z, y^0(x), x),$$

i.e., $\gamma(z', y^0(x), x, \xi') = \gamma(z, y^0(x), x, \xi)$. Taking

$$s = \sigma(\gamma(z', y^0(x), x, \xi'), y^0(x), p, x) = \sigma(\gamma(z, y^0(x), x, \xi), y^0(x), p, x),$$

the definition (16) implies that

$$z = z^*(s; p, x, \xi) \quad \text{and} \quad z' = z^*(s; p, x, \xi').$$

By uniform continuity of $g(\cdot, y^0(x), x)$, there exists $\Delta > 0$ such that (C.1) holds whenever

$$\|z' - z\| < \Delta. \quad (\text{C.2})$$

To see that all such (Δ, p) are identified, observe that (Δ, p) meet the requirement if and only if for all z and z' in $\mathcal{Z}(y^0(x), x)$ that satisfy (C.2), there exist a share vector s and markets t and t' such that

$$z = z^*(s; p, x, \xi_t) \quad \text{and} \quad z' = z^*(s; p, x, \xi_{t'}).$$

For any (Δ, p) , satisfaction of (C.2) is observable, as are the values of $z^*(s; p, x, \xi_\tau)$ for all $s \in \mathcal{S}(x, p, \xi_\tau)$ in every market τ such that $P_\tau = p$. Thus, satisfaction of this condition is observable. \square

References

- ACKERBERG, D. (2003): “Advertising Learning and Customer Choice in Experience Good Markets: A Structural Empirical Examination,” *International Economic Review*, 44, 1007–1040.
- ACKERBERG, D. A., K. CAVES, AND G. FRAZER (2015): “Identification Properties of Recent Production Function Estimators,” *Econometrica*, 83, 2411–2451.
- ADAO, R., A. COSTINOT, AND D. DONALDSON (2017): “Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade,” *American Economic Review*, 107, 633–89.
- ALLEN, R. AND J. REHBECK (2019): “Identification with Additively Separable Heterogeneity,” *Econometrica*, 87, 1021–1054.
- ANDERSON, S. P., A. DE PALMA, AND B. KREIDER (2001): “Tax Incidence in Differentiated Product Oligopoly,” *Journal of Public Economics*, 81, 173–192.
- BACKUS, M., C. CONLON, AND M. SINKINSON (2021): “Common Ownership and Competition in the Ready-To-Eat Cereal Industry,” working paper, NYU-Stern.
- BARSEGHYAN, L., M. COUGHLIN, F. MOLINARI, AND J. C. TEITELBAUM (2021): “Heterogeneous Choice Sets and Preferences,” *Econometrica*, 89, 2015–2048.
- BAYER, P., F. FERREIRA, AND R. McMILLAN (2007): “A Unified Framework for Measuring Preferences for Schools and Neighborhoods,” *Journal of Political Economy*, 115, 588–638.
- BAYER, P., N. KEOHANE, AND C. TIMMINS (2009): “Migration and Hedonic valuation: The Case of Air Quality,” *Journal of Environmental Economics and Management*, 58, 1 – 14.
- BENKARD, L. AND S. T. BERRY (2006): “On the Nonparametric Identification of Nonlinear Simultaneous Equations Models: Comment on Brown (1983) and Roehrig (1988),” *Econometrica*, 74, 1429–1440.
- BERRY, S. (1994): “Estimating Discrete Choice Models of Product Differentiation,” *RAND Journal of Economics*, 23, 242–262.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 60, 889–917.

- (2004): “Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Vehicle Market,” *Journal of Political Economy*, 112, 68–105.
- BERRY, S. T., A. GANDHI, AND P. A. HAILE (2013): “Connected Substitutes and Invertibility of Demand,” *Econometrica*, 81, 2087–2111.
- BERRY, S. T. AND P. A. HAILE (2010): “Nonparametric Identification of Multinomial Choice Demand Models with Heterogeneous Consumers,” Discussion Paper 1718, Cowles Foundation, Yale University.
- (2014): “Identification in Differentiated Products Markets Using Market Level Data,” *Econometrica*, 82, 1749–1797.
- (2018): “Nonparametric Identification of Simultaneous Equations Models with a Residual Index Structure,” *Econometrica*, 86, 289–315.
- (2021): “Foundations of Demand Estimation,” in *Handbook of Industrial Organization*, ed. by K. Ho, A. Hortaçsu, and A. Lizzeri, Elsevier.
- BHATTACHARYA, D. (2018): “Empirical Welfare Analysis for Discrete Choice: Some General Results,” *Quantitative Economics*, 9, 571–615.
- BLOCK, H. AND J. MARSCHAK (1960): “Random Orderings and Stochastic Theories of Responses,” in *Contributions to Probability and Statistics: Essays in Honor of Harold Hotelling*, ed. by I. Olkin, S. Ghurye, W. Hoeffding, W. G. Mado, and H. B. Mann, Stanford University Press, 97–132.
- BLUNDELL, R., D. KRISTENSEN, AND R. L. MATZKIN (2013): “Control Functions and Simultaneous Equations Methods,” *American Economic Review, Papers and Proceedings*, 103, 563–569.
- (2020): “Individual Counterfactuals with Multidimensional Unobserved Heterogeneity,” Cemmap working paper cwp60/17, University College, London.
- BLUNDELL, R. AND R. L. MATZKIN (2014): “Control Functions in Nonseparable Simultaneous Equations Models,” *Quantitative Economics*, 5, 271–295.
- BURDA, M., M. HARDING, AND J. HAUSMAN (2015): “A Bayesian Mixed Logit-Probit Model for Multinomial Choice,” *Journal of Applied Econometrics*, 30, 353–376.
- CAPPS, C., D. DRANOVE, AND M. SATTERTHWAITTE (2003): “Competition and Market Power in Option Demand Markets,” *RAND Journal of Economics*, 34, 737–763.

- CHERNOZHUKOV, V. AND C. HANSEN (2005): “An IV Model of Quantile Treatment Effects,” *Econometrica*, 73, 245–261.
- CONLON, C. AND J. H. MORTIMER (2021): “Empirical Properties of Diversion Ratios,” *RAND Journal of Economics*.
- DECAROLIS, F., M. POLYAKOVA, AND S. P. RYAN (2020): “Subsidy Design in Privately Provided Social Insurance: Lessons from Medicare Part D,” *Journal of Political Economy*, 128, 1712–1752.
- DIAMOND, R. (2016): “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980–2000,” *The American Economic Review*, 106, 479–524.
- DUARTE, M., L. MAGNOLFI, M. SØLVSTEN, AND C. SULLIVAN (2021): “Testing Firm Conduct,” working paper, UW-Madison.
- FALMAGNE, J.-C. (1978): “A Representation Theorem for Finite Random Scale Systems,” *Journal of Mathematical Psychology*, 18, 52–72.
- FAN, Y. (2013): “Ownership Consolidation and Product Characteristics: A Study of the U.S. Daily Newspaper Market,” *American Economic Review*, 103, 1598–1628.
- FLORENS, J.-P., M. M. AND J. ROLIN (1990): *Elements of Bayesian Statistics*, Marcel Decker.
- GALE, D. AND H. NIKAIDO (1965): “The Jacobian Matrix and Global Univalence of Mappings,” *Mathematische Annalen*, 159, 81–93.
- GENTZKOW, M. AND J. SHAPIRO (2010): “What Drives Media Slant? Evidence from U.S. Newspapers,” *Econometrica*, 78, 35–71.
- GERBER, A. (1998): “Estimating the Effect of Campaign Spending on Senate Election Outcomes Using Instrumental Variables,” *The American Political Science Review*, 92, 401–411.
- GOLDBERG, P. K. (1995): “Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry,” *Econometrica*, 63, 891–951.
- GORDON, B. AND W. R. HARTMANN (2013): “Advertising Effects in Presidential Elections,” *Marketing Science*, 32, 19–35.
- HOM, M. (2018): “School Choice, Segregation and Access to Quality Schools: Evidence from Arizona,” working paper, Yale University.

- LEWBEL, A. (2000): “Semiparametric Qualitative Response Model Estimation with Unknown Heteroscedasticity or Instrumental Variables,” *Journal of Econometrics*, 97, 145–177.
- (2014): “An Overview of the Special Regressor Method,” in *The Oxford Handbook of Applied Nonparametric and Semiparametric Econometrics and Statistics*, ed. by J. S. Racine, L. Su, and A. Ullah, Oxford University Press, 38–62.
- LI, X., W. HARTMANN, AND T. AMANO (2020): “Preference Externality Estimators: A Comparison of Border Approaches and IVs,” Working paper, Stanford University.
- MAS-COLELL, A., M. D. WHINSTON, AND J. R. GREEN (1995): *Microeconomic Theory*, Oxford University Press.
- MATZKIN, R. L. (2008): “Identification in Nonparametric Simultaneous Equations,” *Econometrica*, 76, 945–978.
- (2015): “Estimation of Nonparametric Models with Simultaneity,” *Econometrica*, 83, 1–66.
- MCFADDEN, D. (2005): “Revealed Stochastic Preference: A Synthesis,” *Economic Theory*, 26, 245–264.
- MCFADDEN, D., A. TALVITIE, AND ASSOCIATES (1977): *Demand Model Estimation and Validation*, Berkeley CA: Institute of Transportation Studies.
- MILLER, N. AND G. SHEU (2021): “Quantitative Methods for Evaluating the Unilateral Effects of Mergers,” *Review of Industrial Organization*, 58, 143–177.
- NAKAMURA, E. AND D. ZEROM (2010): “Accounting for Incomplete Pass-Through,” *Review of Economic Studies*, 77, 1192–1230.
- NEILSON, C. (2021): “Targeted Vouchers, Competition Among Schools, and the Academic Achievement of Poor Students,” Working paper, Princeton University.
- NEVO, A. (2000): “Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry,” *RAND Journal of Economics*, 31, 395–421.
- (2001): “Measuring Market Power in the Ready-to-Eat Cereal Industry,” *Econometrica*, 69, 307–42.

- NEWKEY, W. K. AND J. L. POWELL (2003): “Instrumental Variable Estimation in Nonparametric Models,” *Econometrica*, 71, 1565–1578.
- OLLEY, S. G. AND A. PAKES (1996): “The Dynamics of Productivity in the Telecommunications Equipment Industry,” *Econometrica*, 64, 1263–97.
- PALAIS, R. S. (1959): “Natural Operations on Differential Forms,” *Transactions of the American Mathematical Society*, 92, 125–141.
- PARTHASARATHY, T. (1983): *On Global Univalence Theorems*, Springer.
- PEARL, J. (2009): *Causality*, Cambridge University Press, second ed.
- PEARL, J., M. GLYMOUR, AND N. JEWELL (2016): *Causal Inference in Statistics: A Primer*, Wiley.
- PETRIN, A. (2002): “Quantifying the Benefits of New Products: The Case of the Minivan,” *JPE*, 110, 705–729.
- SEVERINI, T. A. AND G. TRIPATHI (2006): “Some Identification Issues in Nonparametric Linear Models with Endogenous Regressors,” *Econometric Theory*, 22, 258–278.
- VAN MILL, J. (2002): *The Infinite-Dimensional Topology of Function Spaces*, Elsevier.
- WALDFOGEL, J. (2003): “Preference Externalities: An Empirical Study of Who Benefits Whom in Differentiated-Product Markets,” *RAND Journal of Economics*, 34, 557–568.