

Nonparametric Identification of Differentiated Products Demand Using Micro Data*

Steven T. Berry
Yale University

Philip A. Haile
Yale University

May 14, 2024

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.

Keywords: demand, nonparametric identification, differentiated products

*Early versions of this work were circulated under the title “Nonparametric Identification of Multinomial Choice Demand Models with Heterogeneous Consumers.” We thank Jesse Shapiro, Paulo Somaini, Suk Joon Son, and numerous seminar participants for helpful comments. Miho Hong and Jaewon Lee provided capable research assistance. We thank the referees and Co-Editor (Kate Ho) for helpful comments.

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 its advantages? What specific kinds of variation within and across markets are helpful, and how?

This paper explores these questions by examining nonparametric identification in a model substantially generalizing standard demand 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 instrumental variables requirements.

Although we focus exclusively on identification, our aim is to inform the practice and evaluation of empirical work. The celebrated “credibility revolution” in applied economics 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 possible, and potentially lead to new (parametric or nonparametric) estimation approaches.

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

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 model market-level demand shocks (unobservables at the level of the product \times market) that result in the econometric endogeneity of prices. These shocks make identification of demand nontrivial and create a (perhaps counterintuitive) need for non-price sources of variation. Accounting for these demand shocks is essential to the identification of policy-relevant features such as demand elasticities and equilibrium counterfactuals. This drives our focus on market-level endogeneity and cross-market data, differentiating

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.

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 the large empirical literature that motivates our work.

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, that they have large support, or that certain consumer observables be linked exclusively to the desirability of specific products. The last of those requirements is widely used (often in combination with large support and exogeneity assumptions) 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, gender, race, and educa-

⁴This includes prior work on 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 functional form assumptions permitting application of control function approaches (Blundell and Matzkin (2014)) generally fail, even in standard parametric models. See Berry and Haile (2021) for further discussion of the challenges created by these demand shocks.

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.

⁶Candidate spending, 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).

tion 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 structural 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 at least as many observed consumer attributes as goods.

This final requirement is tightly related to our focus on the extent to which micro data can fully eliminate the need to instrument for the endogenous quantities in a flexible nonparametric demand model that, e.g., places few restrictions on substitution patterns or how price effects vary across products. In practice, some settings—particularly those with a large number of goods—may lack the dimension of variation we require for the most flexible models. This motivates our exploration (in the Supplemental 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 choices. Consumer distances to different options have been used in a number of applications, including those involving demand for hospitals, retail outlets, residential

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.

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)).⁸ 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 examines the proper excludability of standard instruments for prices when non-price observables are endogenous and uninstrumented. A Supplemental Appendix discusses variations on our baseline model, including continuous demand and examples in which key assumptions can be relaxed by strengthening others.

2 Model and Features of Interest

We consider choice among J goods/products and an outside option (indexed as 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 $F_Z(\cdot; t)$ of consumer observables $Z_{it} \in \mathbb{R}^H$, $H \geq J$, with support $\mathcal{Z}(X_t)$.

For clarity, we write random variables (all of which have t among their subscript indices) in uppercase and their realizations in lowercase. The variables

⁸Here 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.

(P_t, X_t, Ξ_t) are common to all consumers in a given market.⁹ 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).¹⁰

Although X_t will typically include observable product characteristics (possibly including characteristics of the outside option), it may also include other factors defining markets.¹¹ For example, consumers might be partitioned into “markets” based on a combination of geography, time, product availability, and demographics (e.g., binned age or income) included in X_t . In contrast, observables varying across consumers within a market are represented by Z_{it} . Key conditions, made precise below, are that consumer observables provide variation of dimension at least J (hence $H \geq J$), with J components of Z_{it} entering the model through an index structure. When $H > J$, the “extra” consumer observables can be treated fully flexibly; thus, without further loss, we henceforth assume $H = J$.¹² We emphasize that 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 Z_{it} and Ξ_t .

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

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

Let \mathcal{C} denote the support of C_{it} . A basic primitive characterizing consumer

⁹The 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.

¹⁰Alternatively, when instruments are available for endogenous components of X_t , our results generalize immediately by expanding P_t to include these endogenous characteristics.

¹¹ 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)). With additional assumptions, cross-market variation in the number of goods available could be valuable; e.g., data from markets with J available goods could be used to predict outcomes in markets with more or fewer goods.

¹²For example, one may interpret all that follows as applying conditional on each value of the “extra” consumer observables. Berry and Haile (2023) provide a more explicit and slightly more flexible treatment.

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 $F_Q(q|C_{it}) = \mathbb{E}[1\{Q_{it} \leq q\} | C_{it}]$. In the case of discrete choice, this distribution can be represented without loss by the structural choice probabilities

$$\mathcal{s}(C_{it}) = (\mathcal{s}_1(C_{it}), \dots, \mathcal{s}_J(C_{it})) = \mathbb{E}[Q_{it} | C_{it}]. \quad (1)$$

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

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

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

on

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

for each market t . The function $\bar{\mathcal{s}}(\cdot; t)$ 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

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.

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

Notably, $\bar{\jmath}(\cdot; t)$ fully determines the own- and cross-price demand elasticities for all goods in market t . One implication is that $\bar{\jmath}(\cdot; t)$ is the feature of \jmath needed to discriminate between alternative models of firm competition (e.g., Berry and Haile (2014), Backus, Conlon, and Sinkinson (2021), Duarte, Magnolfi, Sølvssten, and Sullivan (2023)). And, given a model of supply, $\bar{\jmath}(\cdot; t)$ 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, $\bar{\jmath}(\cdot; t)$ 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 $\bar{\jmath}(\cdot; t)$ 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 $\bar{\jmath}(\cdot; t)$ 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 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 made three significant assumptions:

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

The first is important but standard (when market-level unobservables are ac-

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 $\bar{\jmath}(\cdot; t)$ in all markets suffices in a large fraction of applications, knowledge of \jmath is required for others.

knowledged). The second assumption seems mild for most applications,¹⁶ and it can be relaxed at the cost of more cumbersome exposition.¹⁷

The third assumption is both important and restrictive. Although some modern micro data sets can offer dozens or even hundreds of consumer-level observables, others offer a much smaller set of consumer-level measures. Thus, full satisfaction of this requirement will depend on the dimension of the choice set and the richness of the micro data. We focus on J -dimensional Z_{it} to explore the extent to which micro data can fully eliminate the need to instrument for J endogenous quantities in a flexible nonparametric model. The Supplemental Appendix illustrates some of the ways this requirement on the dimension of Z_{it} can be relaxed. For example, we discuss semiparametric models in which a single binary consumer characteristic Z_{it} and a single excluded instrument for price can suffice.

We will also rely on Assumptions 1–5 below, where we let \mathcal{X} and \mathcal{PX} denote the supports of X_t and (P_t, X_t) , respectively. Assumptions 1–3 introduce an index structure that serves 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.

Assumption 1 (Index). $s(C_{it}) = \sigma(\gamma(Z_{it}, X_t, \Xi_t), P_t, X_t)$, where $\gamma(Z_{it}, X_t, \Xi_t) = (\gamma_1(Z_{it}, X_t, \Xi_t), \dots, \gamma_J(Z_{it}, X_t, \Xi_t)) \in \mathbb{R}^J$ and $\gamma_j(Z_{it}, X_t, \Xi_t) = \Gamma_j(Z_{it}, X_t) + \Xi_{jt}$ for all j .

Assumption 2 (Invertible Demand). For all $(p, x) \in \mathcal{PX}$, $\sigma(\cdot, p, x)$ is injective on the support of $\gamma(Z_{it}, x, \Xi_t)$.

Assumption 3 (Injective Index). For all $(x, \xi) \in \text{supp}(X_t, \Xi_t)$, $\gamma(\cdot, x, \xi)$ is injective on $\mathcal{Z}(x)$.

Assumption 1 requires that Z_{it} and Ξ_t affect choices only through a J -vector of indices in which prices are excluded and each Ξ_{jt} is exclusive to the j th component, entering this component of the index in additively separable form.¹⁸ This is an important nonparametric functional form restriction. A key implication is a sense in which variation in Z_{it} has comparable effects across markets: a unit change in $\Gamma_j(Z_{it}, X_t)$ has the same effect on demand

¹⁶Note that the distribution of Z_{it} over this support may vary freely across markets, including in ways that depend on Ξ_t .

¹⁷For example, our results go through without this assumption if we instead let each $\mathcal{Z}(x)$ be a nonempty set (assumed open and connected below) common to the support of Z_{it} in all markets for which $X_t = x$.

¹⁸We show in Berry and Haile (2023) that the separability requirement can be dropped by strengthening other conditions.

as a unit change in Ξ_{jt} .¹⁹ We link this structure (and the remaining core assumptions below) to familiar specifications from the literature in section 3. In the Supplemental Appendix we discuss a more general model permitting the price of good j to enter the j th component of the index vector.

Assumption 2 requires that the choice probability function σ be “invertible” with respect to the index vector—that, holding (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 $\gamma_j(Z_{it}, X_t, \Xi_t)$ can be interpreted as a quality shifter for good j .²⁰ 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.

In addition to the index structure, we assume that both Z_t and Ξ_t have uncountable connected support. This allows transparent exploration of the potential gains from micro data, using calculus and moment equalities. Assumption 5, in addition, rules 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 .

Assumption 4 (Support). *For all $x \in \mathcal{X}$, $\mathcal{Z}(x)$ is open and connected.*

Assumption 5 (Nondegeneracy). *For each $x \in \mathcal{X}$, there exists (possibly unknown) $p \in \text{supp } P_t | \{X_t = x\}$ such that $\text{supp } \Xi_t | \{(P_t = p, X_t = x)\}$ is open and connected.*

Assumption 4’s requirement of continuous variation in Z_{it} is important to our arguments. Absent appropriate restrictions on other elements of the model, it may not be surprising that continuous variation would be required for nonparametric point identification. Indeed, below we will require continuously distributed instruments for prices 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

¹⁹A sufficient condition is that this equivalence hold for the determination of Q_{it} at the individual consumer level. In typical linear random utility discrete choice models (see section 3) this would allow for a random coefficient on each index $\Gamma_j(Z_{it}, X_t) + \Xi_{jt}$ but would otherwise rule out random coefficients on Ξ_{jt} or on terms involving Z_{it} .

²⁰Other general sufficient conditions for injectivity can be found in, e.g., Palais (1959), Gale and Nikaido (1965), and Parthasarathy (1983).

available variation. The Supplemental Appendix considers identification in an example with discrete Z_{it} .

2.3 A Useful Representation of the Index

We have thus far written the index vector

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

in a way that maximizes clarity about our core assumptions. To demonstrate identification while allowing for endogenous X_t , it is more convenient to define

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

and

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

so that the index can be rewritten as

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

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

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

by construction.

With this notation, we have

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

and

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

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

2.4 Technical Conditions

In parts (i) and (ii) of Assumption 6 below we assume smoothness conditions permitting our applications of calculus and continuity arguments below. Parts (iii) and (iv) strengthen the injectivity requirements of Assumptions 2 and 3 slightly by requiring that the Jacobian matrices $\partial g(z, x)/\partial z$ and $\partial \sigma(\gamma, p, x)/\partial \gamma$ be nonsingular almost surely.

Assumption 6 (Technical Conditions). *For all $(p, x) \in \mathcal{PX}$, (i) $g(\cdot, x)$ is uniformly continuous and continuously differentiable; (ii) $\sigma(\cdot, p, x)$ is continuously differentiable; (iii) $\partial g(z, x)/\partial z$ is nonsingular a.s. on $\mathcal{Z}(x)$; and (iv) $\partial \sigma(\gamma, p, x)/\partial \gamma$ is nonsingular a.s. on $\text{supp } \gamma(Z_{it}, X_t, \Xi_t) | \{(P_t, X_t) = (p, x)\}$.*

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 $\mathbb{E}[\Xi_t] = 0$ without loss. The second reflects the fact that arbitrary injective transformations of the index vector $\gamma(Z_{it}, X_t, \Xi_t)$ 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 $B(X_t) : \mathcal{X} \rightarrow \mathbb{R}^{J \times J}$, with $B(x)$ nonsingular at all x . By letting

$$\tilde{\gamma}(Z_{it}, X_t, \Xi_t) = A(X_t) + B(X_t)\gamma(Z_{it}, X_t, \Xi_t) \quad (7)$$

$$\tilde{\sigma}(\tilde{\gamma}(Z_{it}, X_t, \Xi_t), P_t, X_t) = \sigma(B(X_t)^{-1}(\tilde{\gamma}(Z_{it}, X_t, \Xi_t) - A(X_t)), P_t, X_t) \quad (8)$$

one obtains a new representation of the same map from $(Z_{it}, P_t, X_t, \Xi_t)$ to quantities demanded, 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.²¹

To do this, for each $x \in \mathcal{X}$ we take an arbitrary $z^0(x) \in \mathcal{Z}(x)$ such that $\partial \bar{v}(z^0(x), p_t; t) / \partial z$ is nonsingular in some market t for which $X_t = x$, guaranteeing that $\frac{\partial g(z^0(x), x)}{\partial z}$ is nonsingular.²² We then select the representation of demand in which

$$g(z^0(x), x) = 0 \quad (9)$$

and

$$\left. \frac{\partial g(z, x)}{\partial z} \right|_{z=z^0(x)} = I, \quad (10)$$

where I denotes the J -dimensional identity matrix. In the example above this choice of normalization is equivalent to taking

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

and

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

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

²¹Like location and scale normalizations of utility functions, our normalizations place no restriction on the demand function v or the conditional demand functions $\bar{v}(\cdot; t)$. In practice, normalizations are often embedded in parametric functional forms and exclusivity assumptions. See the Supplemental Appendix.

²²By Assumption 6, almost all points in $\mathcal{Z}(x)$ will satisfy this condition; and $\partial \bar{v}(z^0(x), p_t; t) / \partial z$ is observable in every market t for which $X_t = x$.

3 Example

The literature includes many examples of parametric special cases of our model. The canonical discrete choice demand model (e.g., Berry, Levinsohn, and Pakes (2004)) is derived from a random utility specification like

$$\begin{aligned} u_{ijt} &= x_{jt}\beta_{ijt} - \alpha_{it}p_{jt} + \xi_{jt} + \epsilon_{ijt} \quad j = 1, \dots, J \\ u_{i0t} &= \epsilon_{i0t}, \end{aligned} \quad (11)$$

where u_{ijt} is consumer i 's conditional indirect utility from good j in market t . Here, consumer-level heterogeneity is reflected by the idiosyncratic taste shocks $\epsilon_{it} = (\epsilon_{i0t}, \dots, \epsilon_{iJt})$ (e.g., draws from a normal or type-1 extreme value distribution) as well as the random coefficients α_{it} and $\beta_{ijt} \in \mathbb{R}^K$.²³ These random coefficients are often specified as

$$\alpha_{it} = \exp\left(\alpha_0 + \alpha_y y_{it} + \alpha_\nu \nu_{it}^{(0)}\right), \quad (12)$$

$$\beta_{ijt}^{(k)} = \beta_{0j}^{(k)} + \sum_{\ell=1}^L \beta_{zj}^{(k,\ell)} z_{i\ell t} + \beta_{\nu j}^{(k)} \nu_{it}^{(k)} \quad k = 1, \dots, K, \quad (13)$$

where each $z_{i\ell t}$ represents one of L observable characteristic of consumer (or household) i , y_{it} represents consumer/household income—a consumer-level observable beyond those in z_{it} ,²⁴ and each $\nu_{it} = (\nu_{it}^{(0)}, \dots, \nu_{it}^{(K)})$ is a random vector drawn from a pre-specified distribution. The preference shocks $(\nu_{it}, \epsilon_{it})$ are assumed i.i.d. across consumers and markets.

With (12) and (13), we can rewrite (11) as

$$u_{ijt} = g_j(z_{it}, x_t) + \xi_{jt} + \mu_{ijt}, \quad (14)$$

where

$$g_j(z_{it}, x_t) = \sum_k x_{jt}^{(k)} \sum_{\ell=1}^L \beta_{zj}^{(k,\ell)} z_{i\ell t} \quad (15)$$

$$\mu_{ijt} = \sum_k x_{jt}^{(k)} \left(\beta_{0j}^{(k)} + \beta_{\nu j}^{(k)} \nu_{it}^{(k)} \right) - \exp\left(\alpha_0 + \alpha_y y_{it} + \alpha_\nu \nu_{it}^{(0)}\right) p_{jt} + \epsilon_{ijt}. \quad (16)$$

²³Random coefficients are popular in discrete choice models because they can allow demand systems whose substitution patterns are more flexible along certain dimensions. We focus directly on identification of a demand system with very flexible substitution patterns. Section 5.1 provides additional discussion.

²⁴Recall that our analysis suppresses such “extra” consumer observables but treats them fully flexibly. Section S.1 of the Supplemental Appendix discusses a variation of our model allowing prices to interact with the J -dimensional Z_{it} itself.

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 Assumption 1. The injectivity of demand required by Assumption 2 can be confirmed following Berry (1994) under typical distributional assumptions, or Berry, Gandhi, and Haile (2013) under more general conditions. Assumption 3—invertibility of the linear mapping $g(z_{it}, x_t) = (g_1(z_{it}, x_t), \dots, g_J(z_{it}, x_t))$ in a J -dimensional subvector of z_{it} —clearly requires $L \geq J$. The additional conditions required for the assumption might then be assumed directly 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,²⁵ parametric distributional assumptions, or even a representation of demand through utility maximization. Even within the linear random coefficients random utility discrete choice paradigm, our model would allow the joint distribution of $(\nu_{it}, \epsilon_{it})$ to depend on $(g(z_{it}, x_t) + \xi_t, p_t, x_t)$. 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 $(g(z_{it}, x_t) + \xi_t, p_t, x_t)$.

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 (14) and any of the observables z_{it}, p_t, x_t .²⁶

4 Identification

We consider identification of the latent shocks ξ_t , the demand system $d(C_{it})$, and the conditional demand systems $\bar{d}(Z_{it}, P_t; t)$. The observables comprise the market index t , the variables $(Q_{it}, Z_{it}, P_t, X_t)$, and a vector of instruments W_t discussed 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 $(Q_{it}, Z_{it}, P_t, X_t)$ for many markets t and many consumers i

²⁵Thus, 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.

²⁶The variables $(x_{jt}, p_{jt}, \xi_{jt})$ directly enter the composite error $\xi_{jt} + \mu_{ijt}$. 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} .

in each market. These observables imply observability of choice probabilities conditional on (Z_{it}, P_t, X_t) in each market t . Of course, this implies observability of market-level choice probabilities (market shares) as well. We will use the compact notation $s_t(z)$ to denote the observed choice probability in market t conditional on $Z_{it} = z$.

We proceed in three steps. First, in section 4.1 we demonstrate identification of the function $g(\cdot, x)$ at each $x \in \mathcal{X}$. 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 certain conditional choice probabilities in each market. In particular, given instruments for prices, we show that the realized values $h(x_t, \xi_t)$ can be 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 s and each ξ_t are 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

$$\mathcal{S}(p, x, \xi) = \sigma(g(\mathcal{Z}(x), x) + h(x, \xi), p, x),$$

and

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

Thus, $\mathcal{S}(p, x, \xi)$ denotes the support of choice probabilities within any market t for which $P_t = p, X_t = x$, and $\Xi_t = \xi$. $\mathcal{S}(p, x)$ denotes the support (within and across markets) conditional on $\{P_t = p, X_t = x\}$.²⁷

By Assumptions 2 and 3, for each $s \in \mathcal{S}(p, x, \xi)$ there is a unique $z^* \in \mathcal{Z}(x)$ such that $\sigma(g(z^*, x) + h(x, \xi), p, x) = s$. 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

$$\sigma(g(z^*(s; p, x, \xi), x) + h(x, \xi), p, x) = s. \quad (17)$$

We also define the compact notation

$$z_t^*(s) \equiv z^*(s; p_t, x_t, \xi_t). \quad (18)$$

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 $z_t^*(s)$ for

²⁷ Because $\mathcal{Z}(x)$ is open, continuity and injectivity of σ with respect to the index vector and of the index function with respect to Z_{it} imply that $\mathcal{S}(p, x, \xi)$ (and, thus, $\mathcal{S}(p, x)$ as well) is open. Continuity of $\sigma(\cdot)$ with respect to the index vector and continuity of $g(\cdot)$ with respect to Z_{it} imply (recalling Assumption 4) that $\mathcal{S}(p, x)$ is connected.

all $s \in \mathcal{S}(p_t, x_t, \xi_t)$ are directly observed, even though the value of ξ_t is not. Second, by the invertibility of σ (Assumption 2), we have

$$g(z^*(s; p, x, \xi), x) + h(x, \xi) = \sigma^{-1}(s; p, x) \quad (19)$$

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. A key implication of our nondegeneracy condition (Assumption 5) follows from the definition (2): for each $x \in \mathcal{X}$ there exist $p \in \text{supp} P_t | \{X_t = x\}$ and $\epsilon > 0$ 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 $z_t^*(s)$ for those markets.

Lemma 1. *Let Assumptions 1–6 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 $\mathcal{Z}(x)$ satisfying*

$$\|z' - z\| < \Delta, \quad (20)$$

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 (p, Δ) are identified.*

Proof. See Appendix B. □

With this result in hand, Lemma 2 demonstrates that one can use equation (19) to relate partial derivatives of $g(z, x)$ 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 1–6 hold. Then for every $x \in \mathcal{X}$ there exists a known $\Delta > 0$ such that for almost all z and z' in $\mathcal{Z}(x)$ satisfying (20), the matrix $\left[\frac{\partial g(z, x)}{\partial z}\right]^{-1} \left[\frac{\partial g(z', x)}{\partial z}\right]$ 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 $(P_t, X_t) = (P_{t'}, X_{t'}) = (p, x)$ but, for some choice probability vector \hat{s} ,

$$z = z_t^*(\hat{s}) \neq z_{t'}^*(\hat{s}), \quad (21)$$

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

revealing (recall (18)) that $\xi_t \neq \xi_{t'}$. Lemma 1 ensures that such t, t' , and \hat{s} exist for all z and z' in $\mathcal{Z}(x)$ satisfying (20). And although ξ_t and $\xi_{t'}$ are latent, the identities of markets t and t' satisfying (21) are observed, as are the associated values of \hat{s} , $z_t^*(\hat{s})$, and $z_{t'}^*(\hat{s})$. Differentiating (19) with respect to the choice probability vector within these two markets, we obtain

$$\frac{\partial g(z, x)}{\partial z} \frac{\partial z_t^*(\hat{s})}{\partial s} = \frac{\partial \sigma^{-1}(\hat{s}; p, x)}{\partial s} \quad (22)$$

and

$$\frac{\partial g(z', x)}{\partial z} \frac{\partial z_{t'}^*(\hat{s})}{\partial s} = \frac{\partial \sigma^{-1}(\hat{s}; p, x)}{\partial s}. \quad (23)$$

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

$$\left[\frac{\partial g(z', x)}{\partial z} \right]^{-1} \frac{\partial g(z, x)}{\partial z} = \frac{\partial z_{t'}^*(\hat{s})}{\partial s} \left[\frac{\partial z_t^*(\hat{s})}{\partial s} \right]^{-1}.$$

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

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

Lemma 3. *Under Assumptions 1–6, $g(\cdot, x)$ is identified on $\mathcal{Z}(x)$ 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 a point b . Take any $x \in \mathcal{X}$ and associated $\Delta > 0$ as in Lemma 2. For each vector of integers $\iota \in \mathbb{Z}^J$, define the set

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

By construction, all z and z' in any given set \mathcal{B}_ι satisfy (20). So by Lemma 2, the value of $[\partial g(z, x)/\partial z]^{-1} [\partial g(z', x)/\partial z]$ is known for almost all z and z' in every set \mathcal{B}_ι . Because $\cup_{\iota \in \mathbb{Z}^J} \mathcal{B}_\iota$ forms an open cover of $\mathcal{Z}(x)$, given any $z \in \mathcal{Z}(x)$ there exists a simple chain of open sets \mathcal{B}_ι in $\mathcal{Z}(x)$ linking the point $z^0(x)$ to z (see, e.g., van Mill (2002, Lemma 1.5.21)). Thus, $[\partial g(z, x)/\partial z]^{-1} [\partial g(z^0(x), x)/\partial z]$ is known for almost all $z \in \mathcal{Z}(x)$. With the normalization (10) and continuity of $\partial g(z, x)/\partial z$ with respect to z , the result then follows from the fundamental theorem of calculus for line integrals and the boundary condition (9). \square

The following corollary documents implications of Lemma 3 that play an important role in the next steps of our argument.

Corollary 1. *Let Assumptions 1–6 hold. For all $(p, x) \in \mathcal{PX}$,*
(i) $\partial\sigma^{-1}(\cdot; p, x)/\partial s$ is identified on $\mathcal{S}(p, x)$; and
(ii) for all $s^0 \in \mathcal{S}(p, x)$ and $s^1 \in \mathcal{S}(p, x)$,

$$\Omega(p, x, s^1, s^0) \equiv \sigma^{-1}(s^1; p, x) - \sigma^{-1}(s^0; p, x) \quad (24)$$

is identified.

Proof. Take any $(p, x) \in \mathcal{PX}$. For every market t in which $(P_t, X_t) = (p, x)$, (18) and (19) imply

$$\frac{\partial g(z_t^*(s), x)}{\partial z} \frac{\partial z_t^*(s)}{\partial s} = \frac{\partial \sigma^{-1}(s; p, x)}{\partial s}$$

for all $s \in \mathcal{S}(p, x, \xi_t)$. For all such s , $\frac{\partial z_t^*(s)}{\partial s}$ is directly observed and, by Lemma 3, $\frac{\partial g(z_t^*(s), x)}{\partial z}$ is known. Part (i) follows. Because $\mathcal{S}(p, x)$ is open and connected (recall footnote 27), part (ii) follows from part (i) and the fundamental theorem of calculus for line integrals. \square

4.2 Identification of Conditional Demand

We demonstrate identification of the conditional demand functions $\bar{y}(\cdot; t)$ under the additional assumption that there exist instruments W_t for prices satisfying the standard nonparametric IV conditions conditional on X_t .

Assumption 7 (Instruments for Prices).

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

Part (i) of Assumption 7 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 $\mathbb{E}[h(X_t, \Xi_t)|X_t] = 0$ by construction; thus part (i) implies

$$\mathbb{E}[h_j(X_t, \Xi_{jt})|X_t, W_t] = 0 \quad \text{a.s. for all } j. \quad (25)$$

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 A. The relevance requirement, 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 (25) here), completeness is necessary

and sufficient for identification in separable nonparametric regression.²⁹ The following result demonstrates that 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 1–7, the scalar $h_j(x_t, \xi_{jt})$ is identified for all j and t .*

Proof. For each $(p, x) \in \mathcal{PX}$, let $s^0(p, x)$ be an arbitrary point in $\mathcal{S}(p, x)$. For each t , let \tilde{z}_{it} be an arbitrary point in $\mathcal{Z}(x_t)$. By (19), in every market t we have

$$g_j(\tilde{z}_{it}, x_t) = \sigma_j^{-1}(s_t(\tilde{z}_{it}); p_t, x_t) - h_j(x_t, \xi_{jt}) \quad (26)$$

for each $j = 1, \dots, J$. Using (24), rewrite the j th equation as

$$y_{jt} = f_j(x_t, p_t) - e_{jt} \quad (27)$$

where

$$\begin{aligned} y_{jt} &\equiv g_j(\tilde{z}_{it}, x_t) - \Omega_j(p_t, x_t, s_t(\tilde{z}_{it}), s^0(p_t, x_t)), \\ f_j(x_t, p_t) &\equiv \sigma_j^{-1}(s^0(p_t, x_t); p_t, x_t), \\ e_{jt} &\equiv h_j(x_t, \xi_{jt}). \end{aligned}$$

By Lemma 3 and Corollary 1, each y_{jt} on the LHS of (27) is known. Thus, (27) takes the form of a standard separable nonparametric regression model with regressors (X_t, P_t) , an unknown regression function f_j , and additive structural errors E_{jt} (with realizations e_{jt}). By (25), we have $\mathbb{E}[E_{jt}|X_t, W_t] = 0$ almost surely. So under the completeness condition (part (ii) of Assumption 7) identification of f_j follows immediately from Proposition 2.1 in Newey and Powell (2003). This implies identification of each e_{jt} as well. \square

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

Theorem 1. *Under Assumptions 1–7, the conditional demand system $\bar{\alpha}(\cdot; t)$ is identified on $\mathcal{C}(x_t, \xi_t)$ for all t .*

Proof. Recall that

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

For each t , $h(x_t, \xi_t)$ is now known. Q_{it}, Z_{it}, P_t, X_t are observed. \square

²⁹See also, e.g., Florens and Rolin (1990), Chernozhukov and Hansen (2005), and Severini and Tripathi (2006).

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 8 (Exogenous X_t). $\mathbb{E}[\Xi_t|X_t] = 0$.

When Assumption 8 holds, the definition (2) implies

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

This has two important implications. First, when Assumption 8 is maintained, the IV exclusion condition (part (i) of Assumption 7) requires 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

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

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

Theorem 2. *Under Assumptions 1–8, ξ_t is identified for all t , and the demand system s is identified on \mathcal{C} .*

Note that we have required excluded instruments only for prices. This contrasts with the setting of market-level data in Berry and Haile (2014), where instruments for quantities were also required for identification of demand. This difference is one of the gains from the micro data, and is reflected by the absence of choice probabilities (quantities) on the RHS of the “regression” equation (27)—something accomplished using the adjustment factors $\Omega_j(p_t, x_t, s_t(\tilde{z}_{it}), s^0(p_t, x_t))$ identified in Corollary 1.³⁰

³⁰In contrast to early drafts of this paper discussed in Berry and Haile (2016, 2021), identification of these adjustment factors avoids any need to assume that the support of Z_{it} is sufficiently large to ensure existence of choice probability vectors that are common to all markets with the same observables. Rather, the arbitrary $s^0(p, x)$ may be used in all markets t for which $P_t = p$ and $X_t = x$, regardless of whether $s^0(p, x) \in \mathcal{S}(p, x, \xi_t)$. Note that although the use of arbitrary $s^0(p, x)$ and \tilde{z}_{it} in the proof of Lemma 4 might suggest overidentification in that step, this is not the case. Let any continuously differentiable functions $g(\cdot, x)$ and $z_t^*(\cdot)$ (for each t) be given. Then, by construction, any change in the choice of \tilde{z}_{it} merely adds and subtracts the same constant from the LHS of (27) in each market t , leaving this regression equation unchanged. Similarly, any change in the choice of $s^0(p, x)$ merely adds the same function of (p_t, x_t) to each side of (27).

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 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 hold.

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. In particular, it can pin down key aspects of consumer substitution patterns. From (5), we see that $\frac{\partial s_t(z_{it})}{\partial z_{it}} = \frac{\partial \sigma}{\partial \gamma} \frac{\partial g}{\partial z_{it}}$. Since $\frac{\partial s_t(z_{it})}{\partial z_{it}}$ 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 here—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.³¹

This observation serves as a caution. As a practical matter, a parametric 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.³² 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

³¹For 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 choice probability 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 from a single market.

³² Absent additional restrictions on our nonparametric model, such within-market cross-product variation does not contribute to identification of demand.

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.”³³ 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 7) will hold when X_t is endogenous. In Appendix A 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 overcome through natural timing assumptions. While encouraging, this analysis reveals that one must carefully examine the exclusion condition when X_t is endogenous. As Appendix A illustrates, causal graphs offer a tool for such examination that is simultaneously formal, transparent, and robust.

Absent from the discussion above are the BLP instruments. The characteristics X_{-jt} of goods competing with good j have direct effects on demand for good j and (in standard supply models) on good j ’s markup. The BLP instruments are thus relevant shifters of prices, and in our micro data setting they are not needed as exogenous shifters of quantities. However, the excludability of X_{-jt} requires not only their exogeneity but also a restriction on the way they enter demand (Berry and Haile (2014, 2021)). The Supplemental Appendix shows how adding such restrictions to our model can allow the use of BLP instruments for prices.

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 research involving demand estimation should continue to explore sensitivity 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.

³³See Waldfoegel (2003) as well as Gentzkow and Shapiro (2010), Fan (2013), and Li, Hartmann, and Amano (2020).

We also emphasize that our sufficient conditions for nonparametric identification should not be viewed as necessary conditions for demand estimation in practice. Rather, they should 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. With large choice sets, for example, it will often be difficult to fully satisfy our requirements of J excluded instruments for prices and J -dimensional consumer observables. The robustness analysis of the Supplemental Appendix is aimed to provide some insight about the trade-offs between these requirements and additional structure often employed in practice.

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 must be accounted for to pin down policy-relevant features of demand. Understandably, this development has been accompanied by questions about identification of these models. Our results offer a reassurance that identification can be obtained 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 special regressors or “identification-at-infinity” arguments; it is not limited to particular 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 rich 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 Price 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 explore the question of when these instruments remain properly excluded conditional on observables X_t that are not (mean) independent of Ξ_t . Such instruments are needed for Theorem 1 to apply when Theorem 2 does not, allowing identification of conditional demand without exogeneity of X_t (or instruments for X_t). Unconditional independence is not sufficient (or necessary) for conditional independence. Indeed, in other contexts (e.g., linear regression) it is well known that conditioning on an endogenous “control” variable can 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 .³⁴ Our use of these tools is elementary, and a graphical approach is not essential—causal graphs represent information implied by a fully specified economic 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,

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

which implies the conditional mean independence required by Theorem 1.

We first discuss several causal graphs (and motivating economic examples) that “work”—i.e., that imply (A.1). We then discuss the main type of structure that does not work—i.e., where (A.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 . Each type can also fail—in particular, 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.

³⁴See, 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.

A.1 Graphs that Work

A.1.1 Fully Exogenous Instruments


The simplest cases arise when the instruments W satisfy

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

The conditional independence condition (A.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.

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 (A.2), one obtains the causal graph shown in Figure 1.³⁵ The conditional independence condition (A.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.³⁶

Figure 1



Most of the instrument types discussed in section 5.3 can satisfy (A.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

³⁵We 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, that W does not respond to prices or quantities, and that prices are chosen after X .


³⁶A similar structure is obtained when dependence between Ξ and X reflects a common cause.

offerings unchanged offer one possibility. Another is cross-market variation in the distribution $F_Z(\cdot|t)$ (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 Waldfoegel instruments) and X .

A.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 W . We then obtain a causal graph in Figure 2, where (A.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



A.1.3 X Caused by Instruments

In some cases, the conditional independence condition (A.1) can hold even when X is affected by W . Consider the causal graph in Figure 3, where (A.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 , X affects the nature of the unobservables Ξ , and X is chosen with knowledge of W but not Ξ .


Figure 3



A.1.4 Hausman Instruments

When X is not independent of Ξ , Hausman instruments generally cannot yield the trivial case in which (A.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).³⁷ 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 (A.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.

³⁷Following standard convention, we use a dashed bidirectional edge to represent dependence between X and X_{-t} arising from unmodeled common causes.

A.2 Graphs that Don't Work: X is a Collider

The conditional independence condition (A.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, (A.1) fails.


Figure 5



This structure arises when firms' choices of X depend on both W and Ξ , as will be typical if both Ξ and W are known by firms when choosing X . 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 .³⁸

As suggested already, we can also obtain this type of structure with Hausman instruments. Figure 6 illustrates. This graph 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, 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 (A.1) will fail. In such situations, identification will require different instruments for prices. In many cases such

³⁸A similar structure arises if the dependence between Ξ and X (or X and W) reflects a latent common cause.

instruments can be constructed under natural timing assumptions. This is a topic we take up in the final section of this appendix.³⁹

A.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 . 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 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, temporal separation of observable choices can offer an identification strategy.⁴⁰ 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

$$M^\tau = \Phi(M^{\tau-1}) + W^\tau, \tag{A.3}$$

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.⁴¹


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-

³⁹We also note that when X is a collider, W provides a candidate instrument for X .

⁴⁰Familiar 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.

⁴¹The 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.

Figure 7



τ marginal cost, it is relevant for the determination of P^τ , conditional on 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 A.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 . One may simply reinterpret M^τ above as a period- τ markup shifter.

B Proof of Lemma 1

Fix a value of $x \in \mathcal{X}$. By Assumption 5 and the definition (2), there exist $p \in \text{supp } P_t | \{X_t = x\}$ and $\epsilon > 0$ such that for any z and z' in $\mathcal{Z}(x)$ for which

$$\|g(z', x) - g(z, x)\| < \epsilon, \quad (\text{B.1})$$

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

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

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

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

the definition (17) implies that

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

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

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

To see that all such (p, Δ) are identified, recall (18) and observe that (p, Δ) meet the requirement if and only if for all z and z' in $\mathcal{Z}(x)$ that satisfy (B.2), there exist a choice probability vector s and markets t and t' such that

$$z = z_t^*(s) \quad \text{and} \quad z' = z_{t'}^*(s).$$

For any candidate (p, Δ) , satisfaction of (B.2) is observable, as are the values of $z_\tau^*(s)$ for all $s \in \mathcal{S}(p, x, \xi_\tau)$ in every market τ such that $(P_\tau, X_\tau) = (p, x)$. \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.
- (2016): “Identification in Differentiated Products Markets,” *Annual Review of Economics*, 8, 27–52.
- (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.
- (2023): “Nonparametric Identification of Differentiated Products Demand Using Micro Data,” Cowles Foundation Discussion Paper 2357, Yale University.
- 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 (2023): “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,” *Journal of Political Economy*, 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.