Nonparametric Demand Estimation in Differentiated Products Markets*

Click here for the latest version

Giovanni Compiani†

January 3, 2018

Abstract

I develop and apply a nonparametric approach to estimate demand in differentiated products markets. Estimating demand flexibly is key to addressing many questions in economics that hinge on the shape—and notably the curvature—of market demand functions. My approach applies to standard discrete choice settings, but accommodates a broader range of consumer behaviors and preferences, including complementarities across goods, consumer inattention, and consumer loss aversion. Further, no distributional assumptions are made on the unobservables and only limited functional form restrictions are imposed. Using California grocery store data, I apply my approach to perform two counterfactual exercises: quantifying the pass-through of a tax, and assessing how much the multi-product nature of sellers contributes to markups. In both cases, I find that estimating demand flexibly has a significant impact on the results relative to a standard random coefficients discrete choice model, and I highlight how the outcomes relate to the estimated shape of the demand functions.

Keywords: Nonparametric demand estimation, Incomplete tax pass-through, Multi-product firm

JEL codes: L1, L66

*I am very grateful to Steven Berry, Philip Haile, Pinelopi Goldberg and Xiaohong Chen for their guidance and support, as well as Jason Abaluck, Daniel Ackerberg, Abigail Adams, Nikhil Agarwal, Judith Chevalier, Jan De Loecker, Jean-Pierre Dubé, Ying Fan, Sharat Ganapati, Mitsuru Igami, Myrto Kalouptsidi, Adam Kapor, Jakub Kastl, Yuichi Kitamura, Lorenzo Magnolfi, Eduardo Morales, Camilla Roncoroni, Marcelo Sant'Anna, Paul Scott, and participants of the Yale IO Prospectus Workshop, 2017 Young Economists Symposium, EARIE 2017, and the Marketing Recruiting Seminar at UC Berkeley-Haas for helpful comments and conversations. The empirical results in this paper were calculated (or derived) based on data from The Nielsen Company (US), LLC and marketing databases provided by the Kilts Center for Marketing Data Center at the University of Chicago Booth School of Business. All errors and omissions are my own.

†email: giovanni.compiani@yale.edu

1 Introduction

Many areas of economics study questions that hinge on the shape of the demand functions for given products. Examples include investigating the sources of market power,[1] evaluating the effect of a tax or subsidy,[2] merger analysis,[3] assessing the impact of a new product being introduced into the market,[4] understanding the drivers of the well-documented incomplete pass-through of cost shocks and exchange-rate shocks to downstream prices,[5] and determining whether firms play a game with strategic complements or substitutes.[6] Given a model of supply, the answers to these questions crucially depend on the level, the slope and often the curvature of the demand functions. Therefore, if the chosen demand model is not rich enough, the results could turn out to be driven by the convenient, but often arbitrary, restrictions embedded in the model, rather than by the true underlying economic forces. This motivates using demand estimation methods that rely on minimal parametric assumptions.

In this paper, I propose a nonparametric approach to estimate demand in differentiated products markets based on aggregate data.[7] In such settings, a standard practice is to posit a random coefficients discrete choice logit model[8] and estimate it using the methodology developed by Berry, Levinsohn, and Pakes (1995) (henceforth BLP). The approach in BLP accomplishes two crucial goals. First, unlike simpler alternatives such as plain logit models, it generates reasonable substitution patterns. Second, it accounts for price endogeneity. However, one limitation of this type of models is that they rely on a number of parametric restrictions for tractability. Specifically, a functional form is assumed for utility, and the distributions of both the random coefficients and the idiosyncratic taste shocks are taken to belong to some parametric family. In contrast to this, my approach does not make any

[1] E.g., Berry, Levinsohn, and Pakes (1995) and Nevo (2001).

[2] See, e.g., Bulow and Pfleiderer (1983) and Weyl and Fabinger (2013).

[3] E.g., Nevo (2000) and Capps, Dranove, and Satterthwaite (2003).

[4] E.g., Petrin (2002).

[5] E.g., Nakamura and Zerom (2010) and Goldberg and Hellerstein (2013).

[6] See Bulow, Geanakoplos, and Klemperer (1985).

[7] By differentiated products markets, I mean markets in which consumers face a range of options that are differentiated in ways that are both observed and unobserved to the researcher. In the presence of unobserved heterogeneity, all the variables that are chosen by firms after observing consumer preferences - e.g. prices in many models - are endogenous. The need to deal with endogeneity is a key feature of the estimation approach developed in this paper.

[8] Throughout the paper, we use the terms “random coefficients (discrete choice) logit model” and “mixed logit model” interchangeably.

distributional assumptions on the unobservables and imposes only limited functional form restrictions. Instead, it leverages a range of constraints, such as monotonicity of demand in certain variables and properties of the derivatives of demand, that are grounded in economic theory.

In addition, by directly targeting the demand functions as opposed to the underlying utility parameters, my approach relaxes several assumptions on consumer behavior and preferences that are embedded in BLP-type models. The latter models assume that each consumer picks the product yielding the highest (indirect) utility among all the available options. This implies, among other things, that the goods are substitutes to each other,[9] that consumers are aware of all products and their characteristics,[10] and that each consumer buys at most one unit of a single product.[11] Further, because the estimation approach requires data on market shares, one has to take a stand on the size of the market.[12] In contrast to this, my approach allows for a much broader range of consumer behaviors and preferences, including complementarities across goods, consumer inattention, and consumer loss aversion. Moreover, it applies to settings in which the researcher does not wish to commit to a definition of market size. I illustrate the wide applicability and the feasibility of the approach through various simulations in Section 4. The results suggest that the method works well in moderate samples and that it is able to capture the shape of the own- and cross-price elasticities in a number of settings where standard random coefficients logit models fail. Further, the penalty one has to pay in terms of increased standard errors relative to a correctly specified parametric model appears to be small.

The proposed estimation approach builds on recent results by Berry and Haile (2014) (hence-

---

[9] Gentzkow (2007) develops a parametric demand model that allows for complementarities across goods and applies it to the market for news. Given the relatively small number of options available to consumers, pursuing a nonparametric approach seems feasible in this industry and we view this as a promising avenue for future research.

[10] Goeree (2008) uses a combination of market-level and micro data to estimate a BLP-type model where consumers are allowed to ignore some of the available products. The model specifies the inattention probability as a parametric function of advertising and other variables. Relative to Goeree (2008), this paper allows for more general forms on inattention. Specifically, any model that satisfies the connected substitutes conditions in Assumption 2 (as well as the index restriction in Assumption 1) is permitted. Section 4.2 presents simulation results from one such model. A recent paper by Abaluck and Adams (2017) obtains identification of both utility and consideration probabilities in a class of models with inattentive consumers facing exogenous prices.

[11] A few studies, including Hendel (1999) and Dubé (2004), estimate models of "multiple discreteness", where agents buy multiple units of multiple products. However, these papers typically rely on individual-level data rather than aggregate data. The same applies to papers that model discrete/continuous choices, such as Dubin and McFadden (1984) for the case of electric appliances and electricity.

[12] For instance, BLP take the market size to be the number of households in the US in a given year, while Nevo (2001) sets the potential number of servings of cereal to one per capita per day and uses the total number of potential servings in a city/quarter as a measure of market size.

forth BH). BH show that one can identify the structural demand functions—and therefore address most economic questions of interest[13]—while imposing relatively mild restrictions on consumer behavior and preferences. While this type of results constitutes a necessary first step for any estimation approach, one may wonder whether it is possible to translate the identification arguments in BH into an estimation method that can be implemented on the type of data available to economists. To the best of my knowledge, this paper is the first to do so.[14] Specifically, I propose approximating the demand functions using Bernstein polynomials, which make it easy to enforce a number of economic constraints in the estimation routine. Some of these constraints—notably the exchangeability restriction discussed in Section 3.2—are especially helpful in dealing with the curse of dimensionality that is inherent in nonparametric estimation. In order to obtain valid standard errors, I rely on advances in the econometrics literature on nonparametric instrumental variables regression—specifically Chen and Pouzo (2015) and Chen and Christensen (forthcoming) (henceforth, CC)—and I provide primitive conditions for the case where the objects of interest are price elasticities and counterfactual prices.

In the second part of the paper, I apply the approach to perform two counterfactual exercises using data from California grocery stores. The first is to quantify the pass-through of a tax into retail prices. It is well known that the pass-through of a tax hinges on the curvature of demand.[15] Therefore, flexibly capturing the shape of the demand function is crucial to accurately assessing the effect of a tax on prices and quantities, which motivates pursuing a nonparametric approach. The second counterfactual concerns the role played by the multi-product nature of retailers in driving up markups (the "portfolio effect" in the terminology of Nevo (2001)). A firm simultaneously pricing multiple goods is able to internalize the effects of competition between those goods, pushing prices upwards. The magnitude of this effect once again depends on the shape of the demand curve faced by the retailer.

I focus on sales of fresh strawberries. While this is a fairly small product category, it has features that make it especially suitable for cleanly illustrating how different demand estimation methods affect the outcomes of interest. Specifically, due to the high perishability of strawberries, I can safely abstract from dynamic considerations on both the demand and the supply side. Further, scanner data provides an abundance of store-level observations, which is helpful given the large number of parameters involved in nonparametric estimation.

[13]As pointed out by BH, an important exception is the assessment of individual welfare changes, which is typically precluded with aggregate data unless one commits to a parametric form for utility.

[14]Souza-Rodrigues (2014) proposes a nonparametric estimation approach for a class of models that includes binary demand. However, extension to the case with multiple inside goods does not appear to be trivial.

[15]See, e.g., Bulow and Pfleiderer (1983) and Weyl and Fabinger (2013).

Finally, the data I use comes in the form of quantities sold,[16] which allows me to illustrate two of the above-mentioned advantages of the proposed approach, namely the fact that it can accommodate continuous choice and that it does not require the researcher to make assumptions on market size. I compare the results obtained using my approach to those given by a standard random coefficients logit model and find substantial differences. Notably, mixed logit significantly over-estimates the pass-through of a per-unit tax on organic strawberries in a representative market relative to the nonparametric approach. This reflects the fact that the mixed logit own-price elasticity grows (in absolute value) with price much more slowly than the nonparametric elasticity. A retailer finds it optimal to increase prices by a greater extent when faced with a more slowly-growing elasticity function. Similarly, in the second counterfactual, mixed logit substantially over-estimates the effect of multi-product pricing for both organic and non-organic strawberries.

This paper contributes to the vast literature on models of demand in differentiated products markets pioneered by BLP. As mentioned above, the estimation approach I propose builds on the recent nonparametric identification results in BH. We emphasize that the present paper, as well as BH, focus on the case where the researcher has access to market-level data, typically in the form of shares or quantities, prices, product characteristics and other market-level covariates. This is in contrast to studies that are based on consumer-level data, such as Goldberg (1995), and Berry, Levinsohn, and Pakes (2004). Recently, Berry and Haile (2010) have provided conditions for the nonparametric identification in this "micro data" setting, which opens up an interesting avenue for future research on nonparametric estimation of those models.

Second, the paper is related to the large literature on incomplete pass-through[17] and, particularly, the papers that adopt a structural approach to decompose the different sources of incompleteness. For instance, Goldberg and Hellerstein (2008), Nakamura and Zerom (2010) and Goldberg and Hellerstein (2013) estimate BLP-type models to assess the contribution of markup adjustment in generating incomplete pass-through[18] against competing explanations (i.e. nominal rigidities and the presence of costs not affected by the shocks). The present paper contributes to that literature by providing a method to evaluate markups that

[16]The raw data comes in the form of number of units sold at the UPC level, but I aggregate it into two categories (organic and non-organic) by computing total pounds sold in each category. For more on this, see Section 5.1.

[17]The literature on estimating pass-through is large and we do not attempt to provide an exhaustive list of references. We just mention an interesting recent paper by Atkin and Donaldson (2015) which estimates the pass-through of wholesale prices into retail prices, and uses this to quantify how the gains from falling international trade barriers vary geographically within developing countries.

[18]Specifically, Goldberg and Hellerstein (2008) and Goldberg and Hellerstein (2013) focus on exchange rate pass-through, while Nakamura and Zerom (2010) consider cost pass-through.

relaxes a number of restrictions on consumer behavior and preferences. The results I obtain suggest that estimating demand flexibly may decrease or increase the estimates of markup adjustment depending on the product. Notably, for the fresh organic strawberry market, the estimated markup adjustment is significantly higher and thus the tax pass-through is more incomplete relative to what is predicted by a more restrictive model.[19] [20]

Third, the paper relates to the literature investigating the sources of market power, notably Nevo (2001). Once again, I offer a more flexible method to disentangle different components of market power, and I quantify the role of one such component, the portfolio effect, in the California market for fresh strawberries.

The rest of the paper is organized as follows. Section 2 presents the general model and summarizes the nonparametric identification results from BH. Section 3 discusses the proposed nonparametric estimation approach, with a special emphasis on how to impose a range of constraints from economic theory. Section 4 presents the results of several Monte Carlo simulations. Section 5 applies the methodology to data from California grocery stores to assess the pass-through of a tax and evaluate the effect of multi-product retailers on markups. Section 6 concludes the main text. Appendix A provides some details on Bernstein polynomials. Appendix B discusses several economic constraints and shows how to enforce them in estimation. Appendix C contains all the assumptions and proofs for the inference results. Finally, Appendix D provides two possible micro-foundations for the demand model estimated in the empirical application.

2 Model and Identification

The general model we consider is the same as that in BH. In this section, we summarize the main features of the model as well as the key identification result. In a given market t, there is a continuum of consumers choosing from the set $\mathcal{J} \equiv \{1, ..., J\}$. Each market t is defined by the choice set $\mathcal{J}$ and by a collection of characteristics χ_t specific to the market and/or products. The set χ_t is partitioned as follows:

$$\chi_t \equiv (x_t, p_t, \xi_t),$$

[19] This comparison is performed assuming a Bertrand-Nash model for the supply side. Note that alternative supply models (e.g. Cournot) tend to deliver lower levels of pass-through. However, they do not appear to be realistic in many industries, including the one we consider in this paper.

[20] For non-organic strawberries, I find that mixed logit over-estimates markup adjustment—and thus under-estimates pass-through—relative to the nonparametric approach, but the two confidence intervals overlap.

where x_t is a vector of exogenous observable characteristics (e.g. exogenous product characteristics or market-level income), $p_t \equiv (p_{1t}, ..., p_{Jt})$ are observable endogenous characteristics (typically, market prices) and $\xi_t \equiv (\xi_{1t}, ..., \xi_{Jt})$ represent unobservables potentially correlated with p_t (e.g. unobserved product quality). Let $\mathcal{X}$ denote the support of χ_t.

Next, we define the structural demand system

$$\sigma : \mathcal{X} \to \Delta^J,$$

where Δ^J is the unit $J-$simplex. The function σ gives, for every market t, the vector s_t of shares for the J goods. We emphasize that this formulation of the model is general enough to allow for different interpretations of shares. The vector s_t could simply be the vector of choice probabilities (market shares) for the inside goods in a standard discrete choice model. However, s_t could also represent a vector of "artificial shares," e.g. a transformation of the vector of quantities sold in the market to the unit simplex. This case arises whenever the data does not come in the form of shares, but quantities, and the researcher does not want to take a stand on what the market size is. Indeed, the empirical application in Section 5 fits into this framework. We also define

$$\sigma_0 (\chi_t) \equiv 1 - \sum_{j=1}^{J} \sigma_j (\chi_t),$$

for every market t, where $\sigma_j (\chi_t)$ is the $j-$th element of $\sigma (\chi_t)$. In a standard discrete choice setting, σ_0 corresponds to the share of the outside option, but again this interpretation is not required for the results stated below.

Following BH, we impose two key conditions on the share functions σ that ensure that demand is invertible.[21] The first is an index restriction.

Assumption 1. *Let the exogenous observables be partitioned as* $x_t = \left(x_t^{(1)}, x_t^{(2)}\right)$, *where* $x_t^{(1)} \equiv \left(x_{1t}^{(1)}, ..., x_{Jt}^{(1)}\right)$, $x_{jt}^{(1)} \in \mathbb{R}$ *for* $j \in \mathcal{J}$, *and define the linear indices*

$$\delta_{jt} = x_{jt}^{(1)} \beta_j + \xi_{jt}, \quad j = 1, ..., J$$

For every market t,

$$\sigma (\chi_t) = \sigma \left(\delta_t, p_t, x_t^{(2)}\right)$$

where $\delta_t \equiv (\delta_{1t}, ..., \delta_{Jt})$.

[21] Here we present the identification conditions for a generic demand system. More primitive sufficient conditions tailored to our empirical setting are given in Appendix D.

Assumption 1 requires that, for $j = 1, ..., J$, $x_{jt}^{(1)}$ and ξ_{jt} affect consumer choice only through the linear index δ_{jt}. In other words, $x_{jt}^{(1)}$ and ξ_{jt} are assumed to be perfect substitutes. On the other hand, $x_t^{(2)}$ is allowed to enter the share function in an unrestricted fashion. We emphasize that Assumption 1 is stronger than what is needed for identification of the demand system. Specifically, as shown in Appendix B of BH, both the linearity of δ_{jt} in $x_{jt}^{(1)}$ and its separability in the unobservable ξ_{jt} can be relaxed.[22] However, this stronger assumption simplifies the estimation procedure in that it leads to a *separable* nonparametric regression model. Given that this is the first attempt at estimating demand nonparametrically for this class of models, maintaining Assumption 1 appears to be a reasonable compromise.[23]

The second condition is what BH call "connected substitutes assumption."[24]

Assumption 2. *(i)* $\frac{\partial}{\partial \delta_{jt}} \sigma_k \left(\delta_t, p_t, x_t^{(2)}\right) \leq 0$ *for all* $j > 0, k \neq j$ *and all* $\left(\delta_t, p_t, x_t^{(2)}\right)$*; (ii) for every* $\mathcal{K} \subseteq \mathcal{J}$ *and every* $\left(\delta_t, p_t, x_t^{(2)}\right)$*, there exist a* $k \in \mathcal{K}$ *and a* $j \notin \mathcal{K}$ *such that* $\frac{\partial}{\partial \delta_{kt}} \sigma_j \left(\delta_t, p_t, x_t^{(2)}\right) < 0$.

Assumption 2 requires the goods to be weak gross substitutes in δ_t in terms of the demand system σ; in addition, part (ii) requires some degree of strict substitutability. We emphasize that, because the model does not require that s_t be interpreted as a vector of market shares, Assumptions 1 and 2 must only hold under *some* transformation of the demand system. As a result, although Assumption 2 seems to rule out complementary goods, the model does accommodate some forms of complementarities, as illustrated in Section 4.4.

By Theorem 1 in Berry, Gandhi, and Haile (2013), Assumptions 1 and 2 ensure that demand is invertible, i.e. that, for any triplet of vectors s_t, p_t and $x_t^{(2)}$, there exists at most one vector δ_t such that $s_t = \sigma\left(\delta_t, p_t, x_t^{(2)}\right)$. This means that we can write

$$\delta_{jt} = \sigma_j^{-1}\left(s_t, p_t, x_t^{(2)}\right), \quad j = 1, ..., J. \tag{1}$$

[22] What is critical turns out to be the strict monotonicity in ξ_{jt}.

[23] In the absence of separability of δ_{jt} in ξ_{jt}, one could think of applying existing estimation approaches for nonseparable regression models with endogeneity (e.g. Chernozhukov and Hansen (2006), Chen and Pouzo (2009), Chen and Pouzo (2012) and Chen and Pouzo (2015)).

[24] Here we use a strengthened version of this assumption that also requires differentiability of σ (see Assumption 3* in Berry, Gandhi, and Haile (2013)). While the stronger version is not needed for identification, it will be helpful because it implies properties of the Jacobian matrix of σ that can be imposed in estimation. Further, note that, unlike Assumption 2 in BH, we only require σ to satisfy the connected substitutes assumption in δ, but not in p. This is because connected substitutes in δ is sufficient for identification of the demand system. BH use connected substitutes in p to show that one can discriminate between different oligopoly models on the supply side, but we do not pursue this here.

To obtain identification, we impose two additional instrumental variable (IV) restrictions from BH.

Assumption 3. $\mathbb{E}(\xi_j|X,Z) = 0$ *a.s.-*(X,Z)*, for* $j \in \mathcal{J}$ *and for a random vector* $Z = (Z_1, ..., Z_J)$ *of instruments.*

Assumption 4. *For all functions* $B(\cdot,\cdot,\cdot)$ *with finite expectation, if* $\mathbb{E}\left(B\left(S,P,X^{(2)}\right)|X,Z\right) = 0$ *a.s.-*(X,Z)*, then* $B\left(S,P,X^{(2)}\right) = 0$ *a.s.-*$(S,P,X^{(2)})$*.*

Assumption 3 imposes exogeneity of the observed characteristics x_t. In addition, it requires a vector of (exogenous) instruments z_t excluded from the share function (e.g. cost shifters). Intuitively, the vector $x_t^{(1)}$ serves as "included" instruments for the market shares s_t, while z_t is instrumenting for prices p_t. Assumption 4 constitutes a completeness condition on the joint distribution of (s_t, p_t, x_t, z_t) with respect to (s_t, p_t). In words, it requires the exogenous variables (x_t, z_t) to shift the distribution of the endogenous variables (s_t, p_t) to a sufficient extent. It is a nonparametric analog of the standard rank condition in linear IV models.[25]

BH show that, under the maintained assumptions, the demand system is identified, which we formalize in the following result.

Theorem 1. *Under Assumptions 1, 2, 3 and 4, the structural demand system* σ *is point-identified.*

Theorem 1 implies that the own-price and cross-price elasticity functions are identified. Therefore, in combination with a model of supply, it allows one to address a wide range of economic questions, including evaluation of markups, predicting equilibrium responses to a policy (e.g. a tax), and testing hypotheses on consumer preferences or behavior (e.g. testing the presence of income effects). As pointed out by BH (Section 4.2), one important exception is evaluation of individual consumer welfare, which can be performed with aggregate data only by committing to a parametric functional form for utility. However, identification of the demand system σ is sufficient to pin down some welfare measures, such as the aggregate change in consumer surplus due to a change in prices.

[25] BH also show that identification of the demand model can be achieved without the completeness condition under additional structure. We do not pursue this here.

3 Nonparametric Estimation

3.1 Setup and asymptotic results

Our proposed estimation approach is based on equation (1). Consistent with the goal of avoiding arbitrary functional form and distributional restrictions, we propose a nonparametric approach. Specifically, we rewrite (1) as

$$x_{jt}^{(1)} = \sigma_j^{-1}\left(s_t, p_t, x_t^{(2)}\right) - \xi_{jt} \quad j = 1, ..., J, \tag{2}$$

where we use the normalization $\beta_j = 1$ as in BH. Equation (2), coupled with the IV exogeneity restriction in Assumption 3

$$\mathbb{E}\left(\xi_j | X, Z\right) = 0 \;\; a.s. - (X, Z), \quad j \in \mathcal{J}$$

suggests estimating σ_j^{-1} using nonparametric instrumental variables (NPIV) methods.[26]

Some additional notation is needed to formalize this. We denote by T the sample size, i.e. the number of markets in the data. Let Σ be the space of functions to which σ^{-1} belongs and let $\psi_{M_j}^{(j)}(\cdot) \equiv \left(\psi_{1,M_j}^{(j)}(\cdot), ..., \psi_{M_j,M_j}^{(j)}(\cdot)\right)'$ be the basis functions used to approximate σ_j^{-1} for $j \in \mathcal{J}$.[27] Note that, since we pursue a nonparametric approach, we will let M_j go to infinity with the sample size for all j; thus, although we suppress it in the notation, M_j is a function of T. Then, we let $\Sigma_T \equiv \left\{\left(\tilde{\sigma}_1^{-1}, ..., \tilde{\sigma}_J^{-1}\right) : \tilde{\sigma}_j^{-1} = \pi_j' \psi_{M_j}^{(j)}(\cdot), \; j \in \mathcal{J}\right\}$ be the sieve space for Σ and note that Σ_T depends on the sample size through $\{M_j\}_{j \in \mathcal{J}}$. Next, we denote by $a_{K_j}^{(j)}(\cdot) \equiv \left(a_{1,K_j}^{(j)}(\cdot), ..., a_{K_j,K_j}^{(j)}(\cdot)\right)'$ the basis functions used to approximate the instrument space for good j's equation, and we let $A_{(j)} = \left(a_{K_j}^{(j)}(x_1, z_1), ..., a_{K_j}^{(j)}(x_T, z_T)\right)'$ for $j \in \mathcal{J}$. Again, we suppress the dependence of $\{K_j\}_{j \in \mathcal{J}}$ on the sample size. Further, we require that $K_j \geq M_j$ for all j, which corresponds to the usual requirement in parametric instrumental variable models that the number of instruments be at least as large as the number of endogenous variables. Finally, we let $r_{jt}\left(s_t, p_t, x_t, z_t; \tilde{\sigma}_j^{-1}\right) \equiv \left(x_{jt}^{(1)} - \tilde{\sigma}_j^{-1}\left(s_t, p_t, x_t^{(2)}\right)\right) \times a_{K_j}^{(j)}(x_t, z_t)$.

[26] The literature on NPIV methods is vast and we refer the reader to recent surveys, such as Horowitz (2011) and Chen and Qiu (2016).

[27] In the simulations of Section 4, as well as in the application of Section 5, I use Bernstein polynomials to approximate each of the unknown functions. However, the inference result in Theorem 2 below does not depend on this choice, hence the general notation used in the first part of this section.

Then, the estimation problem is as follows[28]

$$\min_{\tilde{\sigma}^{-1}\in\Sigma_T}\sum_{j=1}^{J}\left[\sum_{t=1}^{T}r_{jt}\left(s_t,p_t,x_t,z_t;\tilde{\sigma}_j^{-1}\right)\right]'\left(A_j'A_j\right)^{-}\left[\sum_{t=1}^{T}r_{jt}\left(s_t,p_t,x_t,z_t;\tilde{\sigma}_j^{-1}\right)\right], \tag{3}$$

The solution $\hat{\sigma}^{-1}$ to (3) minimizes a quadratic form in the terms $\{r_{jt}(\cdot), j\in\mathcal{J}, t=1,...,T\}$, i.e. the implied regression residuals interacted with the instruments. Note that the objective function in (3) is convex. Thus, if the set Σ_T is also convex, readily available algorithms are guaranteed to converge to the global minimizer.

Moreover, we can leverage recent advances in the NPIV literature to perform inference. Specifically, I rely on results in Chen and Christensen (forthcoming) (henceforth, CC) to obtain asymptotically valid standard errors for functionals of the demand system.[29] I now state a result for general scalar functionals, which is a slight modification of Theorem D.1 in CC.[30] For conciseness, the regularity conditions needed for the result, the definition of the estimator for the variance of the functional, and the proof of the theorem, are postponed to Appendix C.

Theorem 2. *Let the assumptions of Theorem 1 hold. Let f be a scalar functional of the demand system and $\hat{v}_T(f)$ be the estimator of the standard deviation of $f\left(\hat{\sigma}^{-1}\right)$ defined in* (13) *in Appendix C. In addition, let Assumptions 5, 6, 7 and 8 in Appendix C hold. Then,*

$$\sqrt{T}\frac{\left(f\left(\hat{\sigma}^{-1}\right)-f\left(\sigma^{-1}\right)\right)}{\hat{v}_T(f)}\xrightarrow{d}N(0,1).$$

Proof. See Appendix C. □

Theorem 2 yields asymptotically valid standard errors for scalar functionals of the demand system. These, in turn, may be used to construct confidence intervals for the functionals and to test hypothesis on the demand functions.

We now provide more primitive sufficient conditions for Theorem 2 for two functionals of interest: price elasticities and equilibrium prices. In both cases, we assume $J=2$, which corresponds to the model estimated in the empirical application of Section 5. Further, for simplicity, we focus on the case where there are no additional exogenous covariates $x^{(2)}$ in

[28]In equation (3), we let $\tilde{A}^{-}$ denote the Moore-Penrose inverse of a matrix $\tilde{A}$.

[29]See also Chen and Pouzo (2015).

[30]Note that CC consider inference on functionals of an *unconstrained* sieve estimator of the unknown regression function, whereas our model features a range of economic constraints. However, since imposing constraints can only (weakly) reduce the variance of the estimators in large samples, the proof strategy in CC may be used to obtain valid standard errors for our setting.

the inverse of the demand system (2), although $x^{(2)}$ could be included following Section 3.3 of CC. We state the results here and postpone the full presentation of the assumptions, as well as the proof, to Appendix C

Theorem 3. *Let the assumptions of Theorem 1 hold. Let f_ϵ be the own-price elasticity functional defined in* (24) *in Appendix C, let $\hat{\upsilon}_T(f_\epsilon)$ denote the estimator of the standard deviation of $f_\epsilon(\hat{\sigma}^{-1})$ based on* (13)*, and let Assumptions 5, 6(iii), 7 and 9 from Appendix C hold. Then,*

$$\sqrt{T}\frac{\left(f_\epsilon\left(\hat{\sigma}^{-1}\right)-f_\epsilon\left(\sigma^{-1}\right)\right)}{\hat{\upsilon}_T\left(f_\epsilon\right)} \xrightarrow{d} N\left(0,1\right).$$

Proof. See Appendix C. □

Theorem 3 establishes the asymptotic distribution of the own-price elasticity for good 1. An analogous argument holds for the own-price elasticity of good 2 and for the cross-price elasticities.

Next, we state a result establishing the asymptotic distribution of the equilibrium price for good 1. Again, the case of good 2 follows immediately.

Theorem 4. *Let the assumptions of Theorem 1 hold. Let f_{p_1} be the equilibrium price functional defined in* (26) *in Appendix C, let $\hat{\upsilon}_T(f_{p_1})$ denote the estimator of the standard deviation of $f_{p_1}(\hat{\sigma}^{-1})$ based on* (13)*, and let Assumptions 5, 6(iii), 7 and 10 from Appendix C hold. Then,*

$$\sqrt{T}\frac{\left(f_{p_1}\left(\hat{\sigma}^{-1}\right)-f_{p_1}\left(\sigma^{-1}\right)\right)}{\hat{\upsilon}_T\left(f_{p_1}\right)} \xrightarrow{d} N\left(0,1\right).$$

Proof. See Appendix C. □

In Section 5, we apply Theorem 4 to obtain confidence intervals for equilibrium prices under two counterfactual scenarios, i.e. the levying of a tax and the switch from monopoly to duopoly.

3.2 Constraints

We conclude this section with a discussion of the curse of dimensionality that is inherent in nonparametric estimation, and of ways to tackle the issue. Note that each of the unknown functions σ_j^{-1} has $2J+n_{x^{(2)}}$ arguments, where $n_{x^{(2)}}$ denotes the number of variables included in $x^{(2)}$. Therefore, the number of parameters to estimate grows quickly with the number of goods and/or the number of characteristics included in $x^{(2)}$, and it will typically be much larger than in conventional parametric models (in the hundreds or even thousands). While this is an objective limitation of the approach, we argue that there are a number of factors alleviating the problem. First, as illustrated by the extensive simulations in Section 4, the own- and cross-price elasticity functions are precisely estimated with moderate sample sizes (3,000 observations) for the $J = 2$ case. Given that many economic questions hinge on functionals of the elasticities, this suggests that it is possible to obtain informative confidence intervals for several quantities of interest. Second, very large market-level data sets are increasingly available to researchers, including the Nielsen scanner data which is used in Section 5. This provides further reassurance on the viability of data-intensive nonparametric methods. Lastly, we show below that imposing constraints from economic theory substantially aids nonparametric estimation. Some constraints (e.g. exchangeability) directly reduce the number of parameters to be estimated. This simplification is often dramatic, especially as the number of goods increases. Other restrictions (e.g. monotonicity) do not affect the number of parameters, but play a role in disciplining the estimation routine, as illustrated in Section 4.

Imposing constraints in model (2) is complicated by the fact that economic theory gives us restrictions on the demand system σ, but what is targeted by the estimation routine is σ^{-1}. Therefore, one contribution of the paper is to translate constraints on the demand system σ into constraints on its inverse σ^{-1}, and show that the latter can be enforced in a computationally feasible way. Specifically, we propose to estimate the functions σ_j^{-1} in (2) using Bernstein polynomials, which we find to be very convenient for imposing economic restrictions.[31]

In this paper, we consider a number of constraints, including exchangeability of the demand functions, monotonicity and lack of income effects. We emphasize that this is not an exhaustive list, and one may wish to impose additional constraints in a given application.[32] Conversely, not all constraints discussed in this paper need to be enforced simultaneously in

[31] See Appendix A for a more formal discussion of Bernstein polynomials.

[32] Note that one restriction that could yield a substantial reduction in the number of parameters is the constraint that $x^{(2)}$ enter the demand functions through the indices δ. Specifically, each demand function goes from having $2J + n_{x^{(2)}}$ to $2J$ arguments. This, in turn, means that the number of Bernstein coefficients for each demand function goes from $m^{2J+n_{x^{(2)}}}$ to m^{2J}, where for simplicity we assume the degree m of the polynomials is the same for all arguments. Restricting the way in which characteristics enter the demand function is typically hard to motivate on economic grounds and therefore this type of constraints is somewhat at odds with the general spirit of the paper. However, such restrictions might constitute an appealing compromise in settings where the number of characteristics and/or goods is relatively high and dimension reduction becomes a necessity.

order to make the approach feasible.[33]

In the remainder of this section, we focus on an exchangeability constraint, which leads to a dramatic reduction in the number of parameters to estimate and thus is very helpful in tackling the curse of dimensionality. In Appendix B, we consider other constraints that one might be willing to impose and show how to enforce them in estimation. In order to define exchangeability, let $\pi : \{1, ..., J\} \to \{1, ..., J\}$ be any permutation and let $x^{(2)} = \left(x_1^{(2)}, ..., x_J^{(2)}\right)$, i.e. we assume that $x^{(2)}$ is a vector of product-specific characteristics.[34] Then, we say the structural demand system σ is exchangeable if

$$\sigma_j \left(\delta, p, x^{(2)}\right) = \sigma_{\pi(j)} \left(\delta_{\pi(1)}, ..., \delta_{\pi(J)}, p_{\pi(1)}, ..., p_{\pi(J)}, x_{\pi(1)}^{(2)}, ..., x_{\pi(J)}^{(2)}\right), \tag{4}$$

for $j = 1, ..., J$. In words, the structural demand system is exchangeable if the demand functions do not depend on the identity of the products, but only on their attributes $\left(\delta, p, x^{(2)}\right)$.[35] For instance, for $J = 3$, exchangeability implies that $\sigma_1 = \sigma_2 = \sigma_3$, and that

$$\sigma_1 \left(\delta_1, \underline{\delta}, \overline{\delta}, p_1, \underline{p}, \overline{p}, x_1^{(2)}, \underline{x}^{(2)}, \overline{x}^{(2)}\right) = \sigma_1 \left(\delta_1, \overline{\delta}, \underline{\delta}, p_1, \overline{p}, \underline{p}, x_1^{(2)}, \overline{x}^{(2)}, \underline{x}^{(2)}\right)$$

for all $\left(\delta_1, \overline{\delta}, \underline{\delta}, p_1, \overline{p}, \underline{p}, x_1^{(2)}, \overline{x}^{(2)}, \underline{x}^{(2)}\right)$. One may be willing to impose exchangeability when it seems reasonable to rule out systematic discrepancies between the demands for different products. This assumption is often implicitly made in discrete choice models. For example, in a standard random coefficient logit model without brand fixed-effects, if the distribution of the random coefficients is the same across goods, then exchangeability is satisfied.[36]

Moreover, one may allow for additional flexibility by allowing the intercepts of the δ indices to vary across goods. This preserves the advantages of exchangeability in terms of dimension reduction, which we discuss below, while simultaneously allowing each unobservable to have a different mean. Relative to existing methods, this is no more restrictive than standard random coefficient logit models with brand fixed-effects and the same distribution of random coefficients across goods.

---

[33] For example, in the empirical application in Section 5, we do not assume lack of income effects.

[34] This need not be the case. For instance, $x^{(2)}$ could be a vector of market-level variables. In these settings, we say the demand system is exchangeable if $\sigma_j \left(\delta, p, x^{(2)}\right) = \sigma_{\pi(j)} \left(\delta_{\pi(1)}, ..., \delta_{\pi(J)}, p_{\pi(1)}, ..., p_{\pi(J)}, x^{(2)}\right)$, which requires $x^{(2)}$ to affect the demand of each good in the same way. Of course, the case where $x^{(2)}$ includes both market-level and product-specific variables can be handled similarly at the cost of additional notation.

[35] For simplicity, here I consider the extreme case of exchangeability across all goods $1, ..., J$. However, one could also think of imposing exchangeability only within a subset of the goods, e.g. the set of goods produced by one company. The arguments in this section would then apply to the subset of products on which the restriction is imposed.

[36] This also uses the fact that the idiosyncratic logit shocks are iid - and thus exchangeable - across goods.

Table 1: Number of parameters with and without exchangeability

J	exchangeability	no exchangeability
3	324	729
4	900	6,561
5	2,025	59,049
10	27,225	3.4bn

Note: Tensor product of univariate Bernstein polynomials of degree 2. $n_{x^{(2)}}$ is assumed to be zero.

Imposing exchangeability on the demand system σ is facilitated by the following result.

Lemma 1. *If σ is exchangeable, then σ^{-1} is also exchangeable.*

Proof. See Appendix B.3. □

Lemma 1 implies that we can directly impose exchangeability of the target functions σ^{-1}. This can be achieved by simply requiring that the Bernstein coefficients be the same for all goods (up to appropriate rearrangements of the arguments of the functions) and imposing that the value of each function be invariant to certain permutations of its arguments. By the approximation properties of Bernstein polynomials (see Appendix A), the latter may be conveniently enforced through *linear* restrictions on the Bernstein coefficients.

Exchangeability is especially helpful in that it dramatically reduces the number of parameters to be estimated. To illustrate this point, we show in Table 1 how the number of parameters grows with J depending on whether we do or do not impose exchangeability.[37] While the number of parameters grows large with J in both cases, the curse of dimensionality is much worse if we do not impose exchangeability—indeed to the point where estimation becomes quickly computationally intractable.

4 Monte Carlo Simulations

In this section, we present the results from a range of Monte Carlo simulations. The goal is twofold. On the one hand, we illustrate that the estimation procedure works well in moderate sample sizes. In doing this, we also highlight the importance of imposing constraints, which motivates the work in Section 3. On the other hand, we show how the general model from

[37] In Table 1, the degree of the Bernstein basis is taken to be 2. Further, each structural demand function σ_j is assumed to depend on the $2J$ arguments (δ, p) only. Of course, the number of parameters would be even higher if one were to use a higher degree for the Bernstein basis and/or include additional covariates.

Section 2 may be applied to a variety of settings which include—but are not limited to—standard discrete choice. All simulations are for the case with $J = 2$ number of goods and $T = 3,000$ number of markets. This might appear to be high, but it is much lower than the sample sizes available in data sets such as the Nielsen scanner data, even for very narrowly defined product categories. Each set of results below is based on 200 Monte Carlo repetitions.

We compare the performance of the estimated procedure to that of standard methods. Specifically, we take as a benchmark a random coefficient logit model with normal random coefficients. We refer to this model as BLP. In order to summarize the results, we plot the own- and cross-price elasticities as a function of price, along with point-wise confidence bands. We choose to plot elasticity functions as that is what is needed for many counterfactuals of interest. In each plot, all market-level variables different from the own-price (i.e. the competitor's price, the exogenous shifters $[x_1, x_2]$ and the unobserved quality levels $[\xi_1, \xi_2]$) are fixed at their median values.

4.1 Correctly specified BLP model

The first simulation is a random coefficients logit model with normal random coefficients. This means that the BLP procedure is correctly specified and therefore we expect it to perform well. On the other hand, we expect the nonparametric approach to yield larger standard errors, due to the fact that it does not rely on any parametric assumptions. Thus, comparing the relative performance of the two should shed some light on how large a cost one has to pay for not committing to a parametric structure when that happens to be correct. In the simulation, the utility that consumer i derives from good j takes the form

$$u_{ij} = \alpha_i p_j + \beta x_j + \xi_j + \epsilon_{ij}$$

where ϵ_{ij} is independently and identically distributed (iid) extreme value across goods and consumers, α_i is distributed $N\left(-1, 0.15^2\right)$ iid across consumers, and we set $\beta = 1$. The exogenous shifters x_j are drawn from a uniform $[0, 2]$ distribution,[38] whereas the unobserved quality indices ξ_j are distributed normally with mean 1 and standard deviation 0.15. We draw excluded instruments z_j from a uniform [0,1] distribution and generate prices according to

$$p_j = 2\left(z_j + \eta_j\right) + \xi_j,$$

[38]Note that we drop the superscript on x_j, since in the simulations we only have one scalar exogenous shifter for each good, i.e. there is no $x^{(2)}$. This applies to all the simulations in this section.

where η_j is uniform [0,0.1]. Note that, while for simplicity we do not generate the prices from the supply first order conditions, the definition of prices above is such that they are positively correlated with both the excluded instruments (consistent with their interpretation as cost shifters) and the unobserved quality (consistent with what would typically happen in equilibrium).

Given this data generating process (dgp), we run our proposed estimation procedure. We impose the following constraints from Section 3 and Appendix B: exchangeability, diagonal dominance of $\mathbb{J}_{\sigma}^{\delta}$ and monotonicity of σ^{-1}. We then compare our results with those obtained by applying BLP. Figures 1 and 2 show the own- and cross-price elasticity functions for good 1, respectively, together with confidence bands for BLP and our proposed procedure (henceforth shortened as NPD, for "Non-Parametric Demand"). Both the NPD and the BLP confidence bands contain the true elasticity functions. As expected, the NPD confidence band is larger than the BLP one for the cross-price elasticity; however, they are still informative. On the other hand, the NPD and the BLP confidence bands for the own-price elasticity appear to be comparable. Overall, we take this as suggestive that the penalty one pays when ignoring correct parametric assumptions is not substantial.

Figure 1: BLP model: Own-price elasticity function

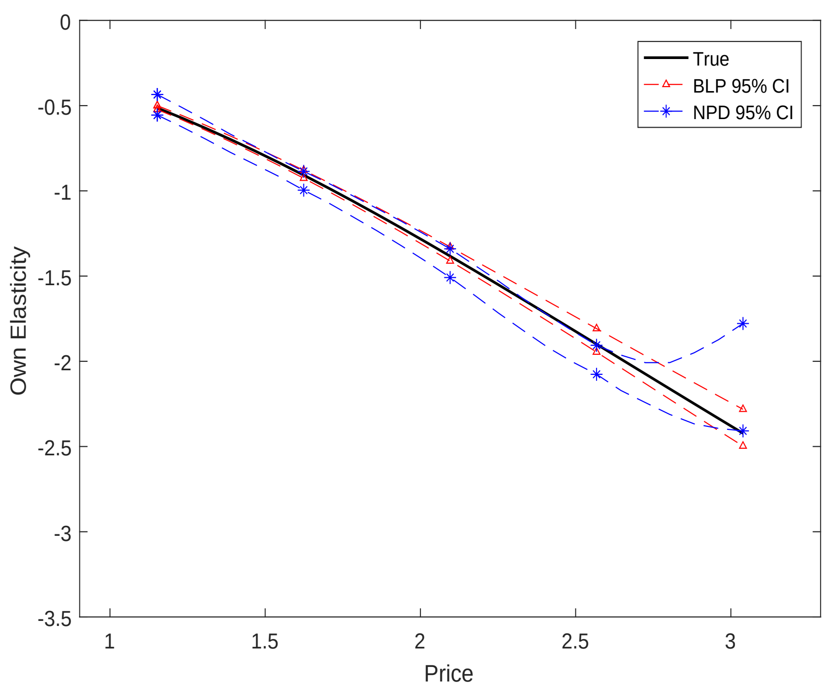

Figure 2: BLP model: Cross-price elasticity function

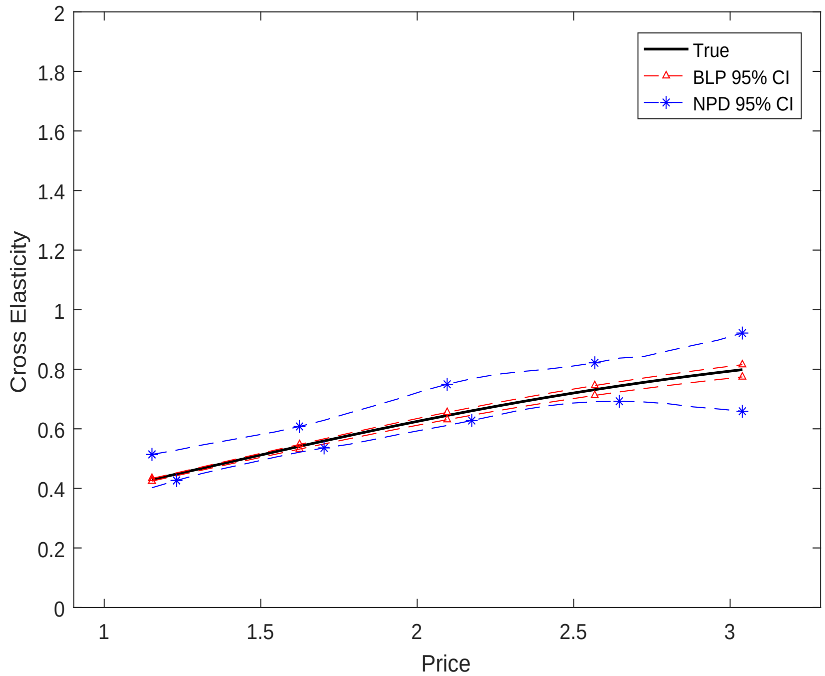

4.2 Inattention

Next, we consider a discrete choice model with inattention. In any given market, we assume a fraction of consumers ignore good 1 and therefore maximize their utility over good 2 and the outside option only. On the other hand, the remaining consumers consider all goods. We take the fraction of inattentive consumers to be $1 - \Phi(3 - p_1)$, where Φ is the standard normal cdf. This implies that, as the price of good 1 increases, more consumers will ignore good 1, which is consistent with the idea that consumers might pay more attention to cheaper products (e.g. if goods are filtered in on-line shopping or products that are on sale are advertised in supermarkets). Except for the presence of inattentive consumers, the simulation design is the same as in Section 4.1. In estimation, we impose the following constraints: monotonicity of σ^{-1}, and diagonal dominance and symmetry of $\mathbb{J}^{\delta}_{\sigma}$.[39] Note that we do not impose exchangeability, since the demand function for good 1 is now different from that of good 2 due to the presence of inattentive consumers. Accordingly, in the BLP procedure, we allow different constants for the two goods.

Figures 3 and 4 show the results for good 1. The nonparametric method captures the shape

[39]See Appendix B for a discussion of these constraints.

of both the own- and the cross-price elasticity functions, whereas BLP is off the mark, underestimating the own-price elasticity and overestimating the cross-price elasticity.

Figure 3: Inattention: Own-price elasticity function

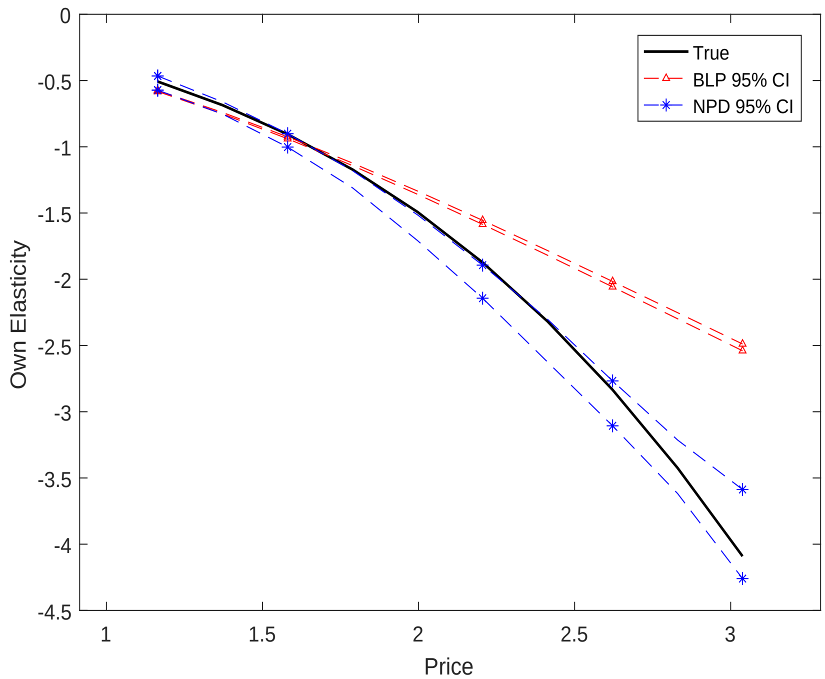

Figure 4: Inattention: Cross-price elasticity function

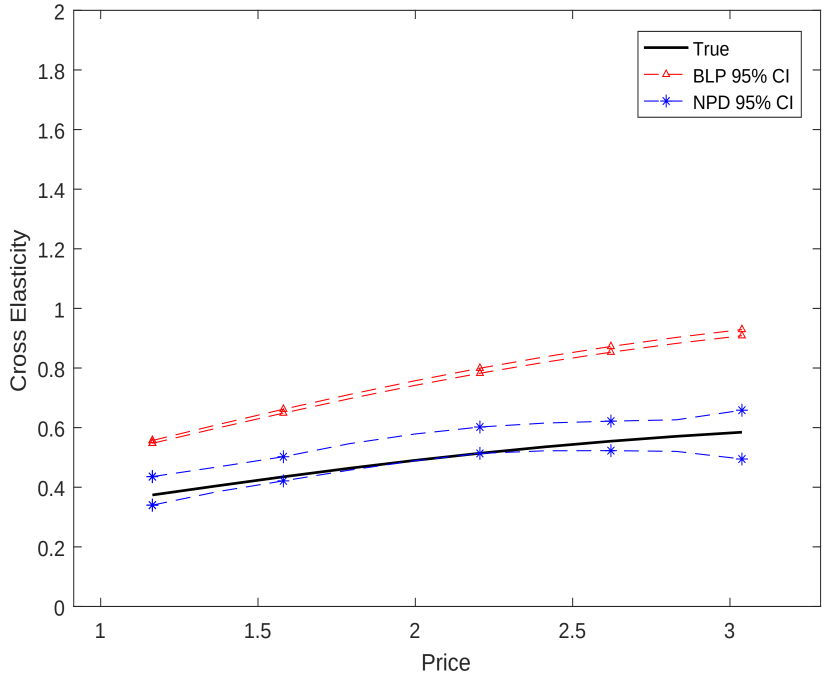

4.3 Loss aversion

Another type of consumer behavior allowed by the NPD model is one where the consumer disutility from paying a price for good j increases with how much more expensive j is relative to another good. I refer to this as "loss aversion" since it may be viewed as one instance of the pattern studied by Tversky and Kahneman (1991) that goes by this name. Specifically, I let the indirect utility from good 1 depend negatively not only on p_1 but also on $p_1 - p_2$, and similarly for good 2. The underlying idea is that if good j is more expensive than good k, paying p_j will be perceived as a loss and therefore will lead to a greater disutility than if good j were cheaper than good k. I set the coefficient on the price difference to -0.15; the simulation design is otherwise the same as that in Section 4.1. As in the previous simulations, we compare the performance of NPD with that of a mixed logit model. In this case, the latter is misspecified in that it only allows p_1, but not $p_1 - p_2$ to enter the utility of good 1, and similarly for good 2. In the NPD estimation, we impose the following constraints: monotonicity of σ^{-1}, diagonal dominance of $\mathbb{J}^{\delta}_{\sigma}$ and exchangeability.[40]

Figures 5 and 6 show the own- and cross-price elasticity functions, respectively. While NPD

[40] See Section 3.2 and Appendix B for a discussion of these constraints.

is on target, BLP tends to underestimate the magnitude of both and the discrepancy grows with price.

Figure 5: Loss aversion: Own-price elasticity function

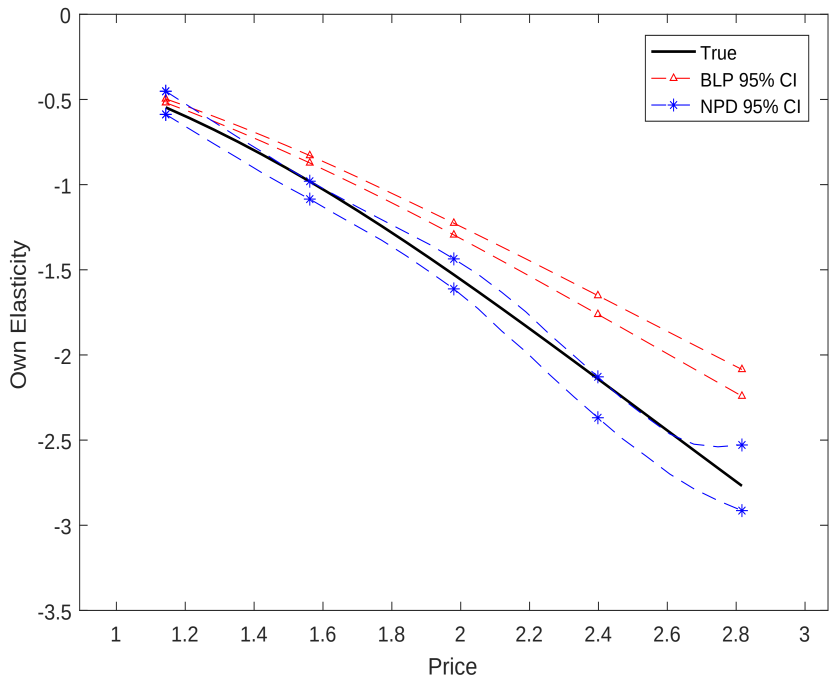

Figure 6: Loss aversion: Cross-price elasticity function

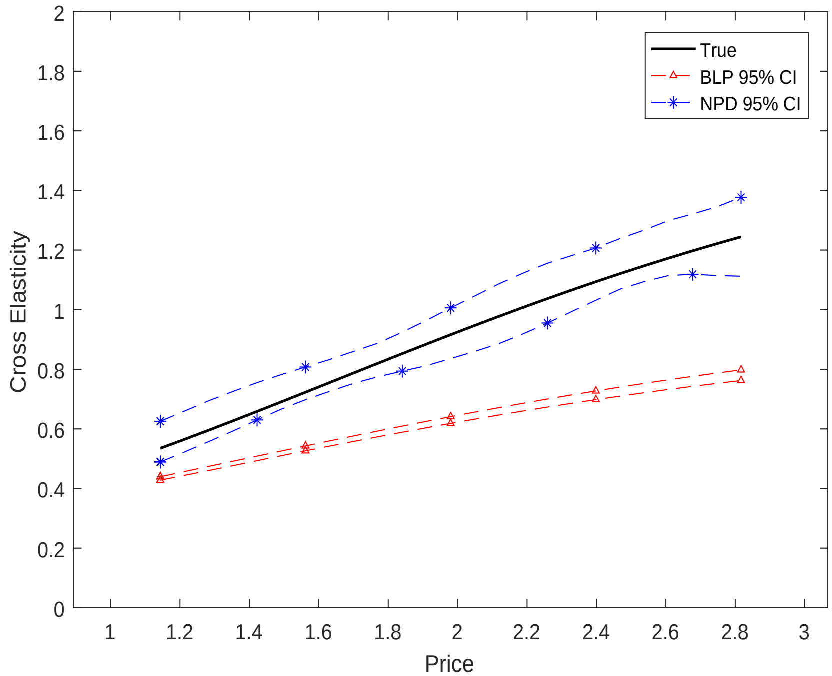

4.4 Complementary goods

We now consider a model where good 1 and 2 are not substitutes, but complements. We generate the exogenous covariates and prices as in the previous two simulations,[41] but we now let market *quantities* be as follows

$$q_j(\delta, p) \equiv 10 \frac{\delta_j}{p_j^2 p_k} \quad j = 1, 2; \quad k \neq j.$$

Note that q_j decreases with p_k and thus the two goods are complements. Now define the function σ_j as

$$\sigma_j(\delta, p) = \frac{q_j(\delta, p)}{1 + q_1(\delta, p) + q_2(\delta, p)}$$

Unlike in standard discrete choice settings, here σ_j does not correspond to the market share function of good j. Instead, it is simply a transformation of the quantities yielding a demand system that satisfies the connected substitutes assumption.[42] In the NPD estimation,

[41] One difference is that we now take the mean of ξ_1 and ξ_2 to be 2 instead of 1 in order to obtain shares that are not too close to zero.

[42] See also Example 1 in Berry, Gandhi, and Haile (2013).

we impose the following constraints: monotonicity of σ^{-1}, diagonal dominance of $\mathbb{J}_{\sigma}^{\delta}$ and exchangeability.[43]

Figures 7 and 8 show the results for good 1. Again, NPD captures the shape of the elasticity functions well. Specifically, note that the cross-price elasticity is slightly negative given that good 1 and good 2 are complements. On the other hand, the BLP confidence bands are mostly off target, consistent with the fact that a discrete choice model is not well-suited for markets with complementarities. In particular, BLP largely over-estimates the magnitude of the own-price elasticity and forces the cross-price elasticity to be positive.

Figure 7: Complements: Own-price elasticity function

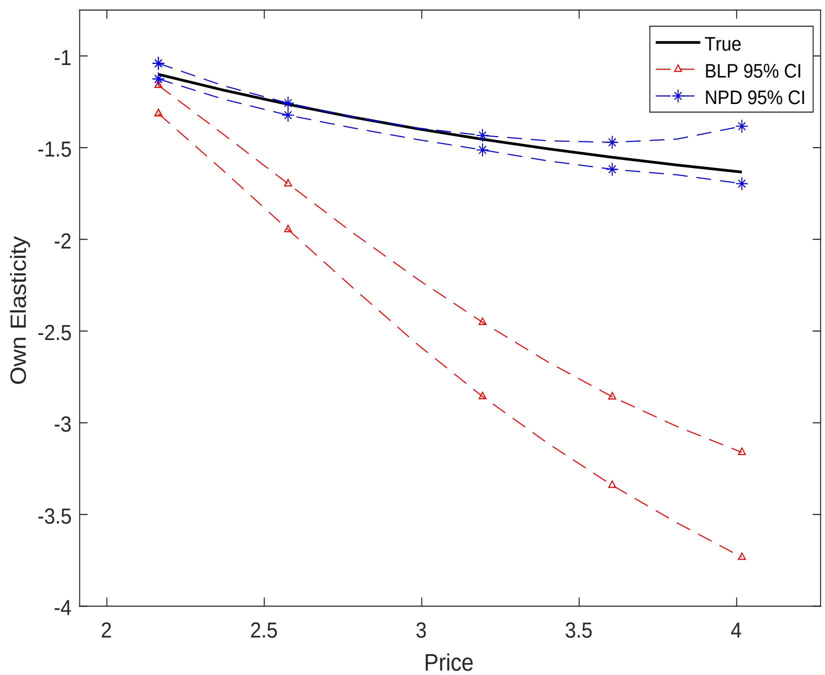

[43]See Section 3.2 and Appendix B for a discussion of these constraints.

Figure 8: Complements: Cross-price elasticity function

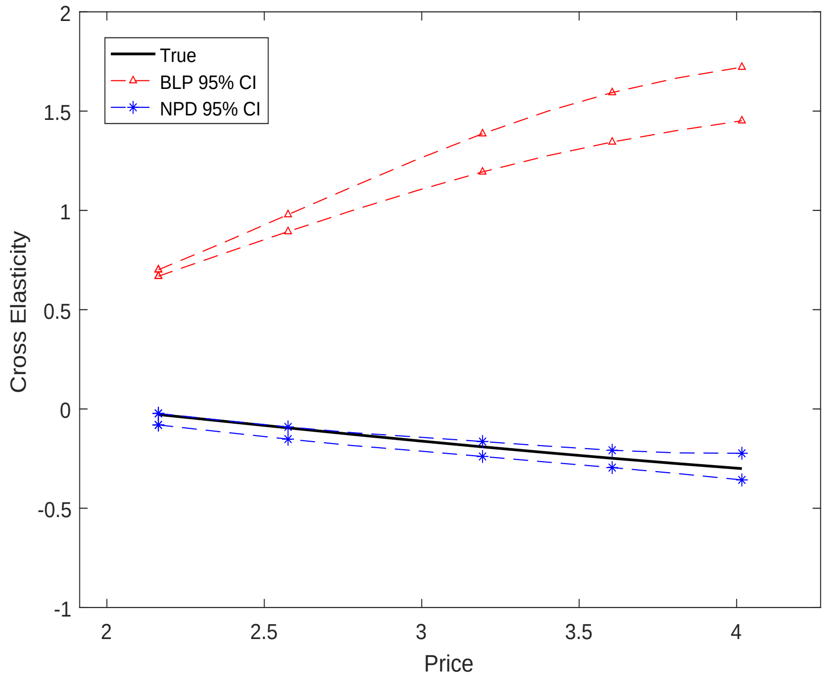

4.5 Role played by the constraints

We conclude this section with one simulation illustrating how the constraints discussed in Section 3 and Appendix B affect the performance of the NPD estimator. The exogenous covariates and prices are draw as in Section 4.4, whereas the market shares are given by a simple logit model with price coefficient equal to -1 and coefficient on the index δ_j equal to 0.5. We then compare the results delivered by NPD while imposing all of the constraints from Section 3 and Appendix B to those obtained while not imposing any constraints. As shown in Figures 9 and 10, imposing the constraints substantially reduces the width of the confidence bands on the elasticity functions.

Figure 9: Own-price elasticity function

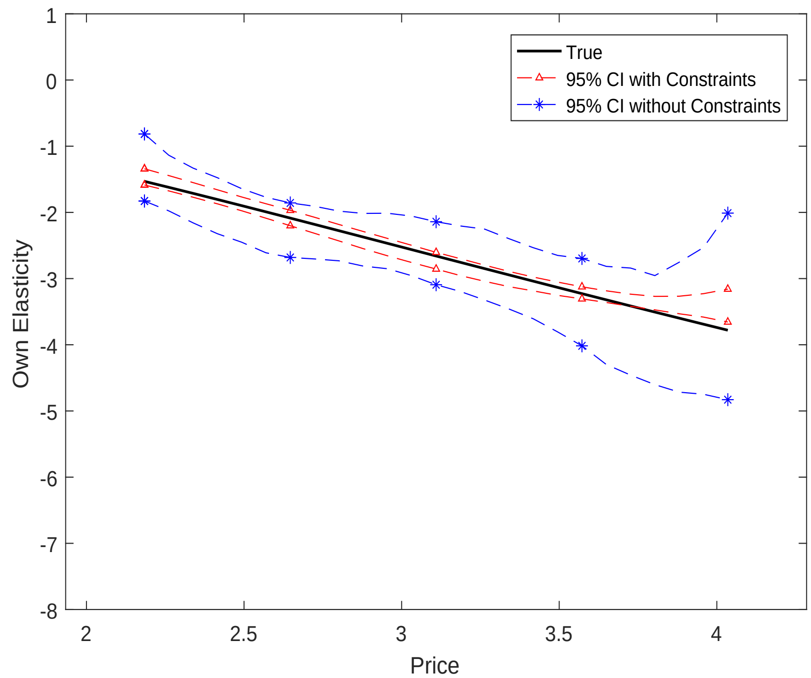

Figure 10: Cross price elasticity function

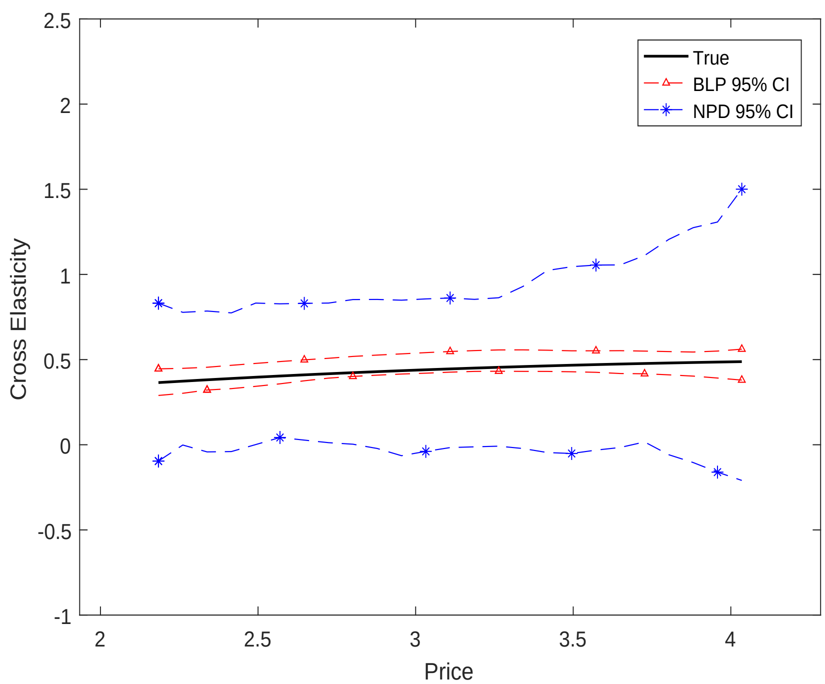

5 Application to Tax Pass-Through and Multi-Product Firm Pricing

In this section, we apply the proposed nonparametric procedure to perform two counterfactual exercises using data from California grocery stores. The first is to quantify the pass-through of a tax into retail prices. It is well-known that the extent to which a tax is passed through to consumers hinges on the curvature of demand.[44] Therefore, flexibly capturing the shape of the demand function is crucial to accurately assessing the effect of a tax on prices and quantities, which motivates pursuing a nonparametric approach.

The second counterfactual concerns the role played by the multi-product nature of retailers in driving up markups. Specifically, a firm simultaneously pricing multiple goods is able to internalize the competition that would occur if those goods were sold by different firms, which pushes prices upwards.[45] Quantifying the magnitude of this effect is ultimately an empirical question which again depends on the shape of the demand functions.

For both counterfactuals, we compare the results obtained with the nonparametric approach to those given by a standard mixed logit model. We find substantial differences and we highlight how these discrepancies relate to the shapes of the estimated demand functions.

5.1 Data

We use data on sales of fresh fruit at stores in California from 2014.[46] Specifically, we focus on strawberries, and look at how consumers choose between organic strawberries, non-organic strawberries and other fresh fruit, which we label as the outside option. This choice is motivated by several considerations. First, there is an abundance of data available for US grocery stores, which is especially important in nonparametric estimation, given the large number of parameters to be estimated. Second, we focus on fresh produce because its perishability implies that we may reasonably abstract from dynamic considerations on both the demand and the supply side. Strawberries, in particular, belong to the category of non-climacteric fruits,[47] which means that they cannot be artificially ripened using ethylene (unlike climac-

[44] See, e.g., Weyl and Fabinger (2013).

[45] This is one of the determinants of markups considered by Nevo (2001) in his analysis of the ready-to-eat cereal industry. On the other hand, when multiple products are sold by the same firm, prices might be driven down by, among other things, economies of scale and loss-leader pricing (see, e.g., Lal and Matutes (1994), Lal and Villas-Boas (1998) and Chevalier, Kashyap, and Rossi (2003)).

[46] Price and quantity data come from the Nielsen scanner dataset.

[47] See, e.g., Knee (2002).

teric fruits, such as bananas). This limits the ability of retailers as well as consumers to stockpile, which further motivates ignoring dynamic considerations in the model. Thus, we are able to keep the setup simple and cleanly compare how different demand models affect the counterfactuals of interest. Finally, the focus on California is motivated by the sheer size of the state, which implies there are over 80,000 observations available for an individual year.

We take a market to be a week/store combination.[48] For each market, the most granular unit of observation in the Nielsen data is a UPC (i.e. a specific bar code). Due to the large number of UPCs, we choose to aggregate according to whether they bear or do not bear the USDA Organic Seal. When this information is missing, we assume the UPC is non-organic. The aggregate quantities are obtained by simply summing the quantities for the individual UPCs, whereas for prices we take a weighted average where the weights are determined by the yearly share of sales that a given UPC has in that supermarket. Similarly, we aggregate across UPCs for selected non-strawberry fruits.[49] Specifically, we focus on the top four non-strawberry fruits according to Produce for Better Health Foundation (2015) in terms of per capita consumption nationwide, i.e. bananas, apples, other berries and oranges. For each of these fruits, we compute a price index (across UPCs) following the same procedure we used for strawberries. These fruit-level price indices are then aggregated even further into a single price index using weights that are proportional to the per capita eatings of each fruit and are normalized to sum to one.

Besides prices, we include the following demand shifters: (i) the availability of non-strawberry fresh fruit in the given week at the state level; (ii) the percentage of yearly organic lettuce sales over total yearly lettuce sales in a given store, as a measure of the local consumers' taste for organic products; and (iii) income.[50]

We instrument for strawberry prices as well as for the price index for the outside option using a combination of shipping-point spot prices and Hausman IV, jointly denoted by Z. Spot prices are a proxy, albeit noisy, for the wholesale prices faced by retailers.[51] Regarding Hausman instruments, we take the mean price of strawberries and the mean price index

[48]We will use the terms "store" and "retailer" interchangeably.

[49]In this case, however, we do not distinguish between organic and non-organic fruits.

[50]Data on income at the zip-code level is downloaded from the Internal Revenue Service website: https://www.irs.gov/uac/soi-tax-stats-individual-income-tax-statistics-zip-code-data-soi.

[51]Spot prices for strawberries are obtained from the US Department of Agriculture website: http://cat.marketnews.usda.gov/cat/index.html. The data reports spot prices for the following shipping points: California, Texas, Florida, North Carolina, and Mexico. In absence of information on where each supermarket sources their strawberries, we take a simple average of the prices at the various shipping points in any given week.

Table 2: Descriptive statistics

	Mean	Median	Min	Max
Quantity non-organic	660.38	504.00	4.00	5,729.00
Quantity organic	112.28	64.00	1.00	2,647.00
Price non-organic	3.19	3.03	0.93	5.99
Price organic	4.24	3.99	1.24	6.99
Spot non-organic	1.69	1.42	0.99	3.40
Spot organic	2.33	2.17	1.25	4.88
Quantity other fruit (per capita)	0.83	0.82	0.60	1.08
Share organic lettuce	0.08	0.06	0.00	0.41
Income	82.40	72.61	33.44	405.09

Note: Prices in dollars per pound. Quantities in pounds. Income in thousands of dollars per household.

for the outside option, respectively, across the Californian supermarkets that are not in the same marketing area[52] as a given store. Excluding supermarkets in the same marketing area is meant to alleviate the usual concerns about Hausman instruments, i.e. that likely spatial correlation in the unobserved quality of the products might induce a violation of the exogeneity assumption.

The resulting dataset has 86,562 markets. Table 2 reports descriptive statistics for each variable and Figure 11 shows the price pattern for a typical store over time (the plot refers to organic strawberries). Both the retail price and the spot price exhibit strong seasonality. Moreover, the retail price sometimes displays a pattern in which it drops and then jumps back up to the initial level. This is typical of supermarket prices given the prevalence of periodic sales. However, in the case of strawberries, this pattern is much less marked than for other items, such as packaged goods. Therefore, we do not explicitly model sales in what follows.[53]

[52]Here we follow the Nielsen partition of the United States into Designated Marketing Areas.

[53]Inventory is often invoked as a justification for sales in models of retail. However, because strawberries are so perishable, it is unlikely that inventory plays a first-order role in driving the retailer's pricing behavior.

Figure 11: Price patterns

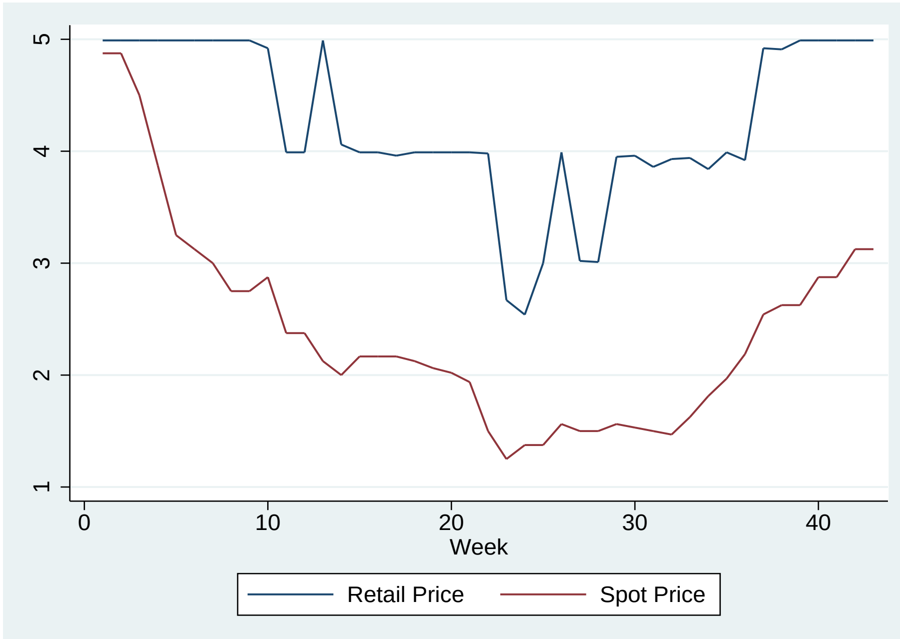

Note: Prices in dollars per pound for organic strawberries sold at a representative store.

5.2 Model

Let 0,1 and 2 denote non-strawberry fresh fruit, non-organic strawberries and organic strawberries, respectively. We take the following model to the data

$$\begin{aligned} s_1 &= \sigma_1 \left(\delta_{strawb}, \delta_{org}, p_0, p_1, p_2, y\right) \\ s_2 &= \sigma_2 \left(\delta_{strawb}, \delta_{org}, p_0, p_1, p_2, y\right) \\ \delta_{strawb} &= \beta_{strawb} - \beta_{out} x_{out} + \xi_{strawb} \\ \delta_{org} &= \beta_{org} + \beta_{lett} x_{lett} + \xi_{org} \end{aligned} \tag{5}$$

In the display above, s_i denotes the share of product i, defined as the quantity of i divided by the total quantity across the three products, x_{lett} denotes the share of organic lettuce sales over total lettuce sales at the store level over the year, x_{out} denotes total sales of non-strawberry fruits at all stores included in the Nielsen dataset in a given week, y denotes income, and $(\xi_{strawb}, \xi_{org})$ denote unobserved shocks at the store/week level, respectively. These unobservables could include, among other things, shocks to the quality of produce at the store/week level,[54] variation in advertising and/or display across stores and time, and taste shocks idiosyncratic to a given store's customer base (possibly varying over time). To the extent that these factors are take into account by the store when pricing produce, (p_0, p_1, p_2) will be endogenous. In contrast, we assume that the demand shifters (x_{out}, x_{lett}) are mean independent of $(\xi_{strawb}, \xi_{org})$. Regarding x_{out}, this is a proxy for the total supply of non-strawberry fruits in California in a given week. As such, we view this as a purely supply-side variable that shifts demand for strawberries inwards by increasing the richness of the outside option,[55] but is independent of store-level shocks.[56] As for x_{lett}, this is meant to approximate the taste for organic products of a given store's customer base. One plausible violation of exogeneity for this variable would arise if consumers with a stronger preference for organic products (e.g. wealthy consumers) tended to go to stores that sell better-quality organic produce (e.g. Whole Foods). This could induce positive correlation between x_{lett} and ξ_{org}. However, note that the counterfactuals we are interested in are robust to some forms of endogeneity arising through this channel. Specifically, consistent with the potential endogeneity concern mentioned above, we could relax the exogeneity assumption $\mathbb{E}\left(\xi_{org}|x_{out}, x_{lett}, z\right) = 0$ to $\mathbb{E}\left(\xi_{org}|x_{out}, x_{lett}, z\right) = \beta_{\xi} x_{lett}$, and the identification argument for σ would still be valid. Under this weaker assumption, the estimated coefficient on x_{lett} would reflect both the actual consumer preferences for organic produce and the fact that the (unobserved) quality of organic produce correlates with x_{lett}. This, however, is not an issue for our purposes given that our counterfactuals only rely on how the demand functions change with prices. The estimation results presented in Section 5.3 assume that the structural demand functions σ_1 and σ_2 are the same, i.e. we impose the exchangeability assumption discussed in Section 3.2.[57] Further, we normalize $\beta_{out} = \beta_{lett} = 1$ and $\beta_{strawb} = 0$ without loss.

We compare the nonparametric approach to a standard parametric model of demand. Specifically, we consider the following mixed logit model:

$$\begin{aligned} u_{i,1} &= \beta_1 + \left(\beta_{p,i} + \beta_y y\right) p_1 + \beta_{p,0} p_0 + \beta_{out}^{par} x_{out} + \xi_1 + \epsilon_{i,1} \\ u_{i,2} &= \beta_2 + \left(\beta_{p,i} + \beta_y y\right) p_2 + \beta_{p,0} p_0 + \beta_{out}^{par} x_{out} + \beta_{lett}^{par} x_{lett} + \xi_2 + \epsilon_{i,2} \end{aligned} \tag{6}$$

[54] Note that we do not include time dummies. This is motivated by the fact that (i) more than 90% of all strawberries produced in the US are grown in California (United States Department of Agriculture (2017)), and (ii) strawberries are harvested in California essentially year-round. Thus, the assumption that the quality of strawberries *sold* in California does not systematically vary over time seems reasonable to a first order.

[55] For example, in the summer many fresh fruits (e.g, Georgia peaches) are in season, which tends to increase the appeal of the outside option relative to strawberries.

[56] The variable x_{out} would be endogenous if the quality of strawberries systematically varied with the harvesting patterns of other fresh fruits. However, as motivated in footnote 54, we abstract from this.

[57] In Appendix D, we provide conditions on the consumer primitives that yield exchangeability of the demand functions.

where $(\epsilon_{i,norg}, \epsilon_{i,org})$ are iid extreme value shocks, (ξ_1, ξ_2) represent unobserved quality of non-organic and organic strawberries, respectively, and the price coefficient $\beta_{p,i}$ is allowed to vary across consumers.[58]

Note one important difference between model (5) and model (6). The latter specifies the indirect utility from each good and thus imposes the implicit (and unrealistic) assumption that each consumer makes a discrete choice between one unit of non-organic strawberries, one unit of organic strawberries, and one unit of non-strawberry fruits. On the other hand, model (5) allows for a broader range of consumer behaviors, including continuous choice, as we show in Appendix D.2. This is one of the advantages of targeting the structural demand function directly as opposed to the underlying utility parameters.

Finally, a remark is in order concerning the market shares s_1 and s_2 in (5). As mentioned in the introduction, the nonparametric procedure applies not only to standard cases where data on market shares is available, but also to settings in which the demand data comes in the form of quantities. The Nielsen data on strawberry sales corresponds to the latter case. In order to implement standard parametric procedures such as (6), one would typically have to define the relevant markets. Given a measure of market size, one would then convert the observed quantities into market shares.[59] However, the nonparametric procedure allows us to circumvent this. The reason is that the identification argument does not require the pair (s_1, s_2) to represent actual market shares. Instead, (s_1, s_2) is allowed to be any transformation of the observed quantities, provided that it satisfies the key conditions in Assumptions 1 and 2.[60] The consumer-level model in Appendix D.2 provides one such transformation. On the other hand, model (6) requires interpreting (s_1, s_2) as the actual fraction of consumers in the market choosing non-organic strawberries and organic strawberries versus non-strawberry fruit, respectively.

[58]We chose this specification over one where the price coefficient is normally distributed (as in BLP) because we found that it made it much easier to impose non-negativity constraints on the marginal costs. Given that we are imposing such constraints in the nonparametric procedure, we wish to impose them in the mixed logit estimation as well for a fair comparison.

[59]For instance, in model (6), we define market shares as sales of organic or non-organic-strawberries over total sales of fresh fruit (i.e. strawberries and other fruit) in the store/week. This implicitly assumes that each consumer chooses at least one unit of fresh fruit. Another possibility would be to define the catchment area of each store, obtain an estimate of the total number of consumers and make an assumption about how much fresh fruit each consumer could potentially buy (this is the strategy adopted by, e.g., Nevo (2001)). This would give a measure of total market size, which could be used to convert the observed quantities into shares. We do not pursue this option here because the data does not provide direct information on the catchment area of each store.

[60]Note, however, that performing counterfactuals that involve the supply side requires mapping the shares back into quantities, given that it is the latter that enter the firm's maximization problem. For the counterfactuals in Section 5.4, I assume that the quantities are proportional to the shares, just like in standard discrete choice. This is not needed, but it makes the comparison with the mixed logit model more transparent.

Table 3: First-stage regressions

	Non-organic		Organic	
	Price	Share	Price	Share
Spot price	0.50**	−0.23**	0.42**	−0.04**
Availability other fruit	−0.71**	0.02	−0.49**	−0.16**
Share organic lettuce	0.04**	−0.02**	−0.07**	0.04**
Hausman other fruit	−0.00	−0.01**	−0.05**	0.01**
Income	0.00	−0.00	−0.01**	0.03**
R^2	0.45	0.31	0.52	0.20

Note: Model includes marketing area dummies. ** denotes significance at the 95% level.

5.3 Estimation Results

We first present the results of the first-stage regressions in Table 3. As expected, the retail prices significantly increase with the spot prices. Further, the share of organic strawberries increases with the share of organic lettuce, consistent with the interpretation of the latter as taste for organic products. This also explains why the coefficient on share of organic lettuce in the regression with share of non-organic strawberries as dependent variable is negative. As expected, the share of organic strawberries decreases with the availability of other fruit. We would expect the same pattern for non-organic strawberries, but the coefficient happens to be non-significant.

We now turn to structural estimation. First, we present the results from the mixed logit model in Table 4. We take the random coefficient on price to have a two-point distribution. Intuitively, this means that we allow consumers to be of two different types depending on their price sensitivity. The last two columns report the coefficients for each of the two types. All coefficients have the expected signs and are significant, except for the coefficient on the price of other fruit, which is insignificant. In particular, note that the coefficient on the availability of other fruit is negative as that variable is a proxy for the richness of the outside option. On the other hand, the coefficient on the share of organic lettuce is positive, consistent with the fact that this is supposed to capture the local consumers' taste for organic products.

We now turn to the nonparametric estimation results. Because this procedure involves a number of parameters too large to report, we choose to show the own-price elasticity functions for a representative market in Figures 12 and 13. All plots have own price on the horizontal axis. Further, we show the corresponding functions estimated using mixed logit for comparison. As in Section 4, we label the mixed logit procedure by BLP and the non-parametric approach by NPD. Asymptotically valid standard errors for the nonparametric

Table 4: Mixed logit estimation results

Variable	Type I	Type II
Price	-11.10^{**}	-24.93^{**}
Price$\times$Income	0.42^{**}	
Price other fruit	0.26	
Other fruit	-0.85^{**}	
Organic lettuce share	0.36^{**}	
Fraction of consumers	0.53^{**}	0.47^{**}

Note: Model includes product dummies. ** denotes significance at the 95% level based on asymptotically valid standard errors.

elasticities at each value of the own price are obtained by applying Theorem 3. While for non-organic strawberries the elasticities estimated by BLP and NPD are relatively close and the shape of the elasticity functions are similar, the two estimation procedures yield markedly different results in the case of organic strawberries. Specifically, BLP underestimates both the magnitude of the elasticity and the rate at which the elasticity function grows (in absolute value) with own price.

Figure 12: Non-organic strawberries: own-price elasticity function

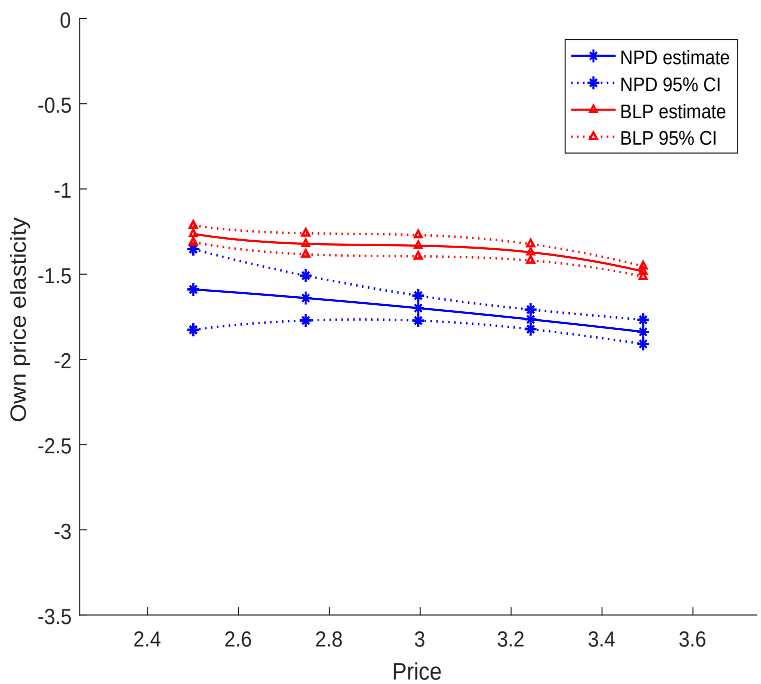

Figure 13: Organic strawberries: own-price elasticity function

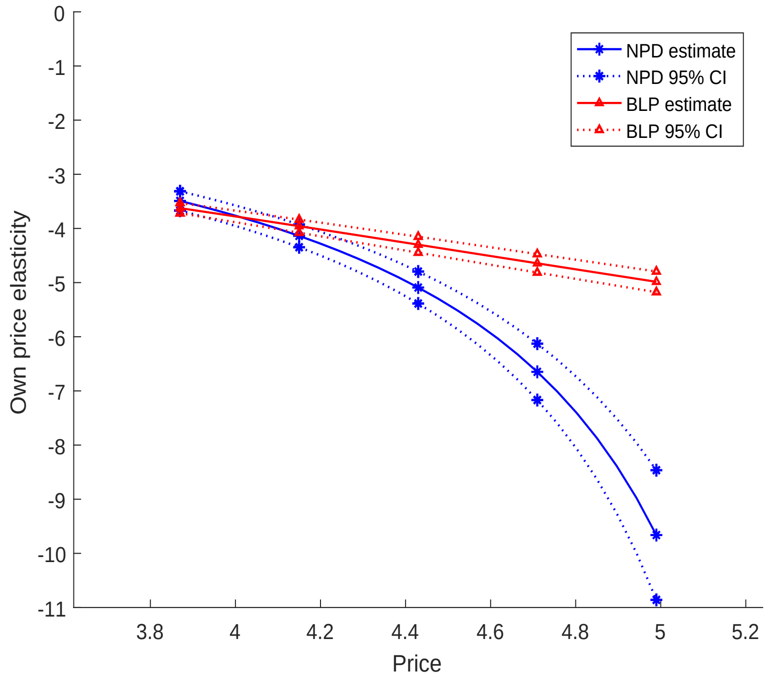

5.4 Counterfactuals

We use the estimation results to address two counterfactual questions. In either case, all variables (except for own price) are set to the same values as in Figures 12 and 13, so that it is possible to see how the shapes of the demand and elasticity functions in the plots affect the outcomes of interest. Further, we set the price of non-organic strawberries to \$2.99 per pound and that of organic strawberries to \$4.43 per pound.[61]

First, we consider the effects of a per-unit tax on prices and quantities. As argued in Weyl and Fabinger (2013), the equilibrium outcomes are not affected by whether the tax is nominally levied on the consumers or on the retailer. This is true for a variety of models of supply, including monopoly. Therefore, without loss of generality, we may assume the tax is nominally levied on consumers in the form of a sales tax. We consider two scenarios: a tax on non-organic only and a tax on organic strawberries only.[62] In each case, we set the tax at 25% of the price for that good.

[61]These values correspond to the mid-points of the horizontal axes in Figures 12 and 13, which in turn cover the inter-quartile range of the corresponding distributions.

[62]Obviously, one could think of other counterfactuals that involve simultaneously taxing both products.

Table 5: Specific tax on non-organic strawberries

	Absolute		Percentage	
	NPD	BLP	NPD	BLP
Δ Price	1.53	1.18	2.04	1.57
	(0.28)	(0.10)	(0.37)	(0.13)
Δ Quantity	−378.69	−302.57		
	(56.44)	(27.54)		

Note: Prices per pound. Quantities in pounds. The first two columns report the absolute changes in prices and quantities. The last two columns report the percentage changes in prices relative to the tax. Values refer to a representative market.

Table 6: Specific tax on organic strawberries

	Absolute		Percentage	
	NPD	BLP	NPD	BLP
Δ Price	0.90	1.18	0.81	1.06
	(0.10)	(0.01)	(0.09)	(0.01)
Δ Quantity	−67.76	−55.14		
	(6.63)	(0.89)		

Note: Prices per pound. Quantities in pounds. The first two columns report the absolute changes in prices and quantities. The last two columns report the percentage changes in prices relative to the tax. Values refer to a representative market.

Tables 5 and 6 report the effect of the tax on prices (in dollars per pound) and quantities (in pounds) for a representative market. Asymptotically valid standard errors for the nonparametric estimates are obtained by applying Theorem 4. Note that, for non-organic strawberries, the effects of the tax on price and quantity estimated by BLP and NPD are not significantly different. This is consistent with the fact that the two elasticity functions—and crucially their slopes—are fairly similar (Figure 12). On the other hand, the pass-through of the tax estimated by BLP is much larger than the one estimated by NPD. According to BLP, the price of organic strawberries increases by more than the amount of the tax in equilibrium, whereas the NPD model predicts that only 81% of the tax is passed through to the retail price. This is consistent with the fact that the NPD elasticity function is much steeper at the initial price, i.e. that the own-price elasticity grows in absolute value much faster with price (Figure 13). A retailer facing such a demand function has a stronger incentive to not increase price too much in response to the tax relative to a retailer facing a flatter elasticity function. In other words, the shape of the demand function estimated nonparametrically is such that it is optimal for the retailer to pass the tax through to consumers to a lesser extent, compared to the scenario where the retailer faces the demand function estimated by BLP.

Table 7: Role of multi-product pricing

	Absolute		Percentage	
	NPD	BLP	NPD	BLP
Non-organic	0.51 (0.18)	1.39 (0.02)	0.26 (0.01)	0.46 (0.01)
Organic	1.06 (0.06)	2.66 (0.02)	0.48 (0.01)	0.75 (0.01)

Note: Prices per pound. The first two columns report the difference between the observed prices and the optimal prices chosen by two competing retailers. The last two columns report the corresponding differences as a percentage of markups. All values refer to a representative market.

As a second counterfactual experiment, we quantify the "portfolio effect". Specifically, we ask what prices would be charged if, in each market, there were two competing retailers, one selling organic strawberries and the other selling non-organic strawberries. We assume the two retailers compete on prices, compute the resulting equilibrium and compare it to the observed prices which, according to our model, are jointly chosen by the (multi-product) retailer to maximize profits. In Table 7 we report the difference between the observed prices and the prices that would arise in the counterfactual world with two single-product retailers. For both organic and non-organic strawberries, BLP overestimates the portfolio effect relative to NPD. Notably, BLP attributes 46% of the markup for non-organic strawberries and 75% of the markup for organic strawberries to the presence of a two-product retailer, whereas the corresponding figures for NPD are 26% and 48%. This is consistent with the differences in the levels of the elasticity functions estimated by BLP and NPD in Figures 12 and 13. Since BLP estimates lower (in absolute value) own-price elasticities, the backed-out marginal costs tend to be lower than the NPD ones. This, in turn, tends to decrease the counterfactual prices in the single-product retailers regime, which implies that BLP will over-estimate the portfolio effect.[63] If similar results were to be found in markets with more than one multi-product firm,[64] it would be of interest to investigate the way in which such results affect the conclusions reached about competing sources of market power, notably collusion among firms.

[63]For fixed marginal costs, the fact that the BLP elasticities are lower in absolute value tends to *increase* the equilibrium prices, which goes in the opposite direction. Given the results in Table 7, it appears that this effect is dominated by the impact of the elasticities on the estimated marginal costs.

[64]The ready-to-eat cereal industry studied by Nevo (2001) is one such example.

6 Conclusion

In this paper, we develop and apply a nonparametric approach to estimate demand in differentiated products markets. Our proposed methodology relaxes several arguably arbitrary restrictions on consumer behavior and preferences that are embedded in standard discrete choice models. Instead, we pursue a nonparametric approach that directly targets the demand functions and leverages a number of constraints from economic theory. Further, we provide primitive conditions sufficient to obtain valid standard errors for quantities of interest. In the second part of the paper, we apply our approach to quantify the pass-through of a tax and assess the effect of multi-product retailers on prices.

This paper is part of a broader research agenda that aims at applying flexible demand estimation procedures to a range of economic questions and industries, above and beyond the illustrations presented here. For instance, our approach could be used to determine whether online news and print news are complements or substitutes,[65] and to assess the role of advertising in US presidential elections.[66] Both of these examples involve estimating demand systems with a low number of options,[67] which suggests that our approach could be successfully applied to those settings.

[65] See Gentzkow (2007).

[66] See Gordon and Hartmann (2016).

[67] Regarding the news market example, Gentzkow (2007) estimates a model with three inside goods. Regarding the US presidential elections example, two parties capture the vast majority of the votes.

Appendix A: Bernstein Polynomials

For a positive integer m, the Bernstein basis function is defined as

$$b_{v,m}(u) = \binom{m}{v} u^v (1-u)^{m-v},$$

where $v = 0, 1, ..., m$ and $u \in [0, 1]$. The integer m is called the degree of the Bernstein basis. In order to approximate a univariate function on the unit interval, one may take a linear combination of the Bernstein basis functions

$$\sum_{v=0}^{m} \theta_{v,m} b_{v,m}(u),$$

for some coefficients $(\theta_{v,m})_{v=0}^{m}$. Similarly, for a function of N variables living in the $[0,1]^N$ hyper-cube, one may use a polynomial of the form

$$\sum_{v_1=0}^{m} \cdots \sum_{v_N=0}^{m} \theta_{v_1,\cdots,v_N,m} b_{v_1,m}(u_1) \cdots b_{v_N,m}(u_N)$$

Note that here we are assuming that the order m is the same for each dimension $n = 1, ..., N$. This is not needed, but we only discuss this case for notational convenience.

Historically, Bernstein polynomials were introduced to approximate an arbitrary function g by a sequence of smooth functions. This is motivated by the following result.[68]

Lemma 2. *Let g be a bounded real-valued function on $[0,1]^N$ and define*

$$B_m[g] = \sum_{v_1=0}^{m} \cdots \sum_{v_N=0}^{m} g\left(\frac{v_1}{m}, \cdots, \frac{v_N}{m}\right) b_{v_1,m}(u_1) \cdots b_{v_N,m}(u_N)$$

Then,

$$\sup_{\mathbf{u} \in [0,1]^N} |B_m[g](\mathbf{u}) - g(\mathbf{u})| \to 0$$

as $m \to \infty$.

This means that, for an appropriate choice of the coefficients, the sequence of Bernstein polynomials provide a uniformly good approximation to any bounded function on the unit hyper-cube as the degree m increases. Specifically, the approximation in Lemma 2 is such that the coefficient on the $b_{v_1,m}(u_1) \cdots b_{v_N,m}(u_N)$ term corresponds to the target function evaluated at $\left[\frac{v_1}{m}, \cdots, \frac{v_N}{m}\right]$, for $v_i = 0, ..., m$ and $i = 1, ..., N$.

One important implication of this result is that, for large m, any property satisfied by the target function g at the grid points $\left\{\left\{\frac{v_1}{m}, \cdots, \frac{v_N}{m}\right\}_{v_i=0}^{m}\right\}_{i=1}^{N}$ should be inherited by the corresponding Bernstein coefficients in order for the resulting approximation to be uniformly good. This gives us *necessary* conditions on the Bernstein coefficients for large m.

To fix ideas, consider the following simple example. Suppose the target function $g : [0,1] \to \mathbb{R}$ is nondecreas-

[68] See, e.g., Chapter 2 of Gal (2008).

ing and that we approximate it using

$$\hat{g}(u) = \sum_{v=0}^{m} \theta_{v,m} b_{v,m}(u) \qquad u \in [0,1].$$

Then for large m, the coefficients $(\theta_{v,m})_{v=0}^{m}$ must satisfy $\theta_{0,m} \leq \theta_{1,m} \leq \cdots \leq \theta_{m,m}$ in order for $\hat{g}$ to be uniformly close to g. To see this, suppose by contradiction that $\theta_{j,m} > \theta_{k,m}$ for some $j < k$ and large m. Then, by Lemma 2, $\hat{g}$ is close to a function h such that $h\left(\frac{j}{m}\right) > h\left(\frac{k}{m}\right)$, i.e. a function that is *not* monotonically nondecreasing. In other words, a monotonicity restriction on the target g implies that the Bernstein coefficients must satisfy intuitive - and in this case linear - monotonicity constraints for large m. The same logic applies to any other assumptions we might be willing to impose on g. This is a powerful tool in the context of demand estimation because economic theory provides us with several restrictions on the structural demand function σ (and therefore on its inverse σ^{-1}).

Appendix B: Additional Constraints

In this appendix, we consider several constraints that one might be willing to impose in estimation besides the exchangeability restrictions discussed in Section 3.2, and we show how to enforce them in estimation in a computationally tractable way. Because these constraints are defined conditional on any given value of $x^{(2)}$, we drop this from notation for notational convenience.

B.1 Symmetry

Let $\mathbb{J}_{\sigma}^{p}(\delta,p)$ denote the Jacobian matrix of σ with respect to p:

$$\mathbb{J}_{\sigma}^{p}(\delta,p) = \begin{bmatrix} \frac{\partial}{\partial p_1}\sigma_1(\delta,p) & \cdots & \frac{\partial}{\partial p_J}\sigma_1(\delta,p) \\ \vdots & \ddots & \vdots \\ \frac{\partial}{\partial p_1}\sigma_J(\delta,p) & \cdots & \frac{\partial}{\partial p_J}\sigma_J(\delta,p) \end{bmatrix}$$

This matrix is the Jacobian of the Marshallian demand system. If we assume that there are no income effects, it coincides with the Jacobian of the Hicksian demand by Slutsky equation and therefore it must be symmetric.
Similarly, let $\mathbb{J}_{\sigma}^{\delta}(\delta,p)$ denote the Jacobian matrix of σ with respect to δ:

$$\mathbb{J}_{\sigma}^{\delta}(\delta,p) = \begin{bmatrix} \frac{\partial}{\partial \delta_1}\sigma_1(\delta,p) & \cdots & \frac{\partial}{\partial \delta_J}\sigma_1(\delta,p) \\ \vdots & \ddots & \vdots \\ \frac{\partial}{\partial \delta_1}\sigma_J(\delta,p) & \cdots & \frac{\partial}{\partial \delta_J}\sigma_J(\delta,p) \end{bmatrix}$$

In a discrete choice model where δ_j is interpreted as a quality index for good j, if one assumes that, for all j, δ_j enters the utility of good j linearly (and does not enter the utility of the other goods), then $\mathbb{J}_{\sigma}^{\delta}(\delta,p)$ must be symmetric.

Conveniently, symmetry of $\mathbb{J}_{\sigma}^{\delta}(\delta,p)$ implies linear constraints on the Bernstein coefficients. To see this, note that by the implicit function theorem, for every (δ, p) and for $s = \sigma(\delta, p)$,

$$\mathbb{J}^{s}_{\sigma^{-1}}(s, p) = \left[\mathbb{J}^{\delta}_{\sigma}(\delta, p)\right]^{-1} \tag{7}$$

Because the inverse of a symmetric matrix is symmetric, symmetry of $\mathbb{J}^{\delta}_{\sigma}(\delta, p)$ implies symmetry of $\mathbb{J}^{s}_{\sigma^{-1}}(s, p)$. This, in turn, imposes *linear* constraints on the Bernstein coefficients as the degree of the approximation goes to infinity.[69]

On the other hand, it appears that symmetry of $\mathbb{J}^{p}_{\sigma}$ requires nonlinear constraints. This is because, by the implicit function theorem, for every (δ, p) and for $s = \sigma(\delta, p)$,

$$\mathbb{J}^{p}_{\sigma}(\delta, p) = -\left[\mathbb{J}^{s}_{\sigma^{-1}}(s, p)\right]^{-1} \mathbb{J}^{p}_{\sigma^{-1}}(s, p) \tag{8}$$

which shows that $\mathbb{J}^{p}_{\sigma}$ is a nonlinear function of the derivatives of σ^{-1} and therefore of the Bernstein coefficients. In estimation, we found it convenient to rewrite (8) as

$$\mathbb{J}^{s}_{\sigma^{-1}}(s, p)\, \mathbb{J}^{p}_{\sigma}(\delta, p) = -\mathbb{J}^{p}_{\sigma^{-1}}(s, p)$$

We then express $\mathbb{J}^{s}_{\sigma^{-1}}$ and $\mathbb{J}^{p}_{\sigma^{-1}}$ as linear combinations of the Bernstein polynomials and introduce extra parameters (call them γ) for the entries of $\mathbb{J}^{p}_{\sigma}$. In this way, we obtain a set of nonlinear constraints that are linear in the Bernstein coefficients θ, given γ, and linear in γ, given θ.[70]

B.2 Additional Properties of the Jacobian of Demand

The matrix $\mathbb{J}^{\delta}_{\sigma}(\delta, p)$ has a number of additional features that we might want to impose in estimation. First, the weak substitutability in Assumption 2(i) requires the off-diagonal elements to be non-positive. Further, it follows from Remark 2 of Berry, Gandhi, and Haile (2013) that the diagonal elements must be positive.[71]

Moreover, $\mathbb{J}^{\delta}_{\sigma}(\delta, p)$ belongs to the class of M-matrices, which are the object of a vast literature in linear algebra.[72] One of the most common definitions of this class is as follows.

Definition 1. A square real matrix is called an M-matrix if (i) it of the form $A = \alpha I - P$, where all entries of P are non-negative; (ii) A is nonsingular and A^{-1} is entry-wise non-negative.

[69] More precisely, the difference between two (appropriately chosen) Bernstein coefficients approximates the change in the function σ_j^{-1} given by a change in s_k. Thus, we would in principle need to divide by the distance between the grid points associated with the two coefficients in order to obtain the derivative of σ_j^{-1} with respect to s_k. However, because we are interested in comparisons between derivatives and the grid points are equidistant, the increments in the denominator cancel out. Therefore, we are left with inequalities involving simple differences of the Bernstein coefficients.

[70] This is helpful especially when it comes to writing the analytic gradient of the constraints to input in the optimization problem.

[71] This is simply the requirement that the structural demand of product j increase in the index δ_j. While a very reasonable condition, it is not needed for identification, but rather it follows from the sufficient conditions given in Section 2.

[72] See, e.g., Plemmons (1977).

We now formalize the aforementioned result, which is a simple corollary of Theorem 2 in Berry, Gandhi, and Haile (2013).

Lemma 3. *Let Assumptions 1 and 2 hold. Then* $\mathbb{J}^{\delta}_{\sigma}(\delta, p)$ *is an M-matrix for all* (δ, p).

Proof. See Section B.3. □

The linear algebra literature provides several properties of M-matrices. However, it is not *a priori* clear how to impose these properties in estimation, since we estimate σ^{-1} rather than σ itself. The Jacobian of the function we estimate is

$$\mathbb{J}^{s}_{\sigma^{-1}}(s, p) = \begin{bmatrix} \frac{\partial}{\partial s_1}\sigma_1^{-1}(s,p) & \cdots & \frac{\partial}{\partial s_J}\sigma_1^{-1}(s,p) \\ \vdots & \ddots & \vdots \\ \frac{\partial}{\partial s_1}\sigma_J^{-1}(s,p) & \cdots & \frac{\partial}{\partial s_J}\sigma_J^{-1}(s,p) \end{bmatrix}$$

Recall that, by the implicit function theorem, we have that, for every (δ, p) and for $s = \sigma(\delta, p)$,

$$\mathbb{J}^{s}_{\sigma^{-1}}(s, p) = \left[\mathbb{J}^{\delta}_{\sigma}(\delta, p)\right]^{-1}$$

Therefore, $\mathbb{J}^{s}_{\sigma^{-1}}(s, p)$ is the inverse of an M-matrix or, in the jargon used in the linear algebra literature, an *inverse M-matrix*. Fortunately, inverse M-matrices have also been widely studied.[73] Thus, we may borrow results from that literature to impose conditions on the Bernstein coefficients for σ^{-1} that must hold in order for $\mathbb{J}^{\delta}_{\sigma}(\delta, p)$ to be an M-matrix.

First, it follows from part (ii) of Definition 1 that $\mathbb{J}^{s}_{\sigma^{-1}}(s, p)$ must have non-negative elements for all (s, p). This means that, for every j, σ_j^{-1} must be increasing in s_k for all k. As discussed in Appendix A, monotonicity is very easy to impose in estimation, given that it reduces to a collection of linear inequalities on the Bernstein coefficients.

Second, under Assumption 2, $\mathbb{J}^{\delta}_{\sigma}$ satisfies a property called column diagonal dominance. The economic content of this property is that the (positive) effect of δ_j on the share of good j is larger than the combined (negative) effect of δ_j on the shares of all other goods, in absolute value. A few definitions are necessary to formalize this point.

Definition 2. An I−by−I matrix $A = (a_{ij})$ is (weakly) diagonally dominant of its rows if

$$|a_{ii}| \geq \sum_{j \neq i} |a_{ij}|,$$

for $i = 1, ..., I$.

Definition 3. An I−by−I matrix $A = (a_{ij})$ is (weakly) diagonally dominant of its row entries if

$$|a_{ii}| \geq |a_{ij}|,$$

for $i = 1, ..., I$ and $j \neq i$.

Column diagonal dominance and column entry diagonal dominance are defined analogously. By Theorem 3.2 of McDonald, Neumann, Schneider, and Tsatsomeros (1995), if an M-matrix M is weakly diagonally

[73] See, e.g., Johnson and Smith (2011).

dominant matrix of its columns, then $(M)^{-1}$ is weakly diagonally dominant of its row entries.[74] This immediately implies the following result.

Lemma 4. *Fix (δ, p) and let $s = \sigma(\delta, p)$. If $\mathbb{J}_{\sigma}^{\delta}$ is diagonally dominant of its columns, then $\frac{\partial}{\partial s_j}\sigma_j^{-1}(s,p) \geq \frac{\partial}{\partial s_k}\sigma_j^{-1}(s,p)$ for all j and all $k \neq j$.*

Lemma 4 translates the assumption that $\mathbb{J}_{\sigma}^{\delta}$ is diagonally dominant of its columns into linear inequalities involving the derivatives of σ^{-1}. Therefore it follows from the same argument used for symmetry that diagonal dominance may be imposed through linear constraints on the Bernstein coefficients.

B.3 Proofs

Proof of Lemma 1. Let $\pi : \{1, ..., J\} \to \{1, ..., J\}$ be any permutation and let $\tilde{\pi}$ denote the function that, for any J-vector y returns the reshuffled version of y obtained by permuting its subscripts according to π, i.e.

$$\tilde{\pi}([y_1, ..., y_J]) = \left[y_{\pi(1)}, ..., y_{\pi(J)}\right]$$

Then, we can rewrite the definition of exchangeability for a generic function $g(y, w)$ of $2J$ arguments as

$$\tilde{\pi}(g(y, w)) = g(\tilde{\pi}(y), \tilde{\pi}(w)).$$

Now take any (δ, p) and let $s = \sigma(\delta, p)$. We can invert the demand system at (δ, p) to obtain

$$\delta = \sigma^{-1}(s, p) \tag{9}$$

By exchangeability of σ, we also have

$$\tilde{\pi}(s) = \sigma(\tilde{\pi}(\delta), \tilde{\pi}(p))$$

Inverting this demand system, we obtain

$$\tilde{\pi}(\delta) = \sigma^{-1}(\tilde{\pi}(s), \tilde{\pi}(p)) \tag{10}$$

Combining (9) and (10),

$$\tilde{\pi}\left(\sigma^{-1}(s, p)\right) = \sigma^{-1}(\tilde{\pi}(s), \tilde{\pi}(p))$$

which shows that σ^{-1} is exchangeable.

□

Proof of Lemma 3. Under Assumptions 1 and 2, Theorem 2 in Berry, Gandhi, and Haile (2013) implies that $\mathbb{J}_{\sigma}^{\delta}(\delta, p)$ is a P-matrix for every (δ, p), where a P-matrix is a square matrix such that all of its principal minors are strictly positive. Next, by the weak substitutability imposed by Assumption 2, $\mathbb{J}_{\sigma}^{\delta}(\delta, p)$ is also a Z-matrix, where a Z-matrix is a matrix with non-positive off-diagonal entries. Finally, since a Z-matrix which is also a P-matrix is an M-matrix,[75] the result follows.

□

[74] To reconcile Theorem 3.2 of McDonald, Neumann, Schneider, and Tsatsomeros (1995) and Definition 3, recall that an inverse M-matrix has non-negative entries.

[75] See, e.g., 8.148 in Seber (2007).

Appendix C: Inference Results

This appendix contains all the assumptions and proofs for the inference results stated in the main text.

C.1 Setup

We first introduce some notation that is used throughout this appendix. We denote by $\mathcal{S}, \mathcal{P}, \mathcal{Z}, \Xi$ the support of S, P, Z, ξ, respectively. Also, we let $W \equiv (X, Z)$ denote the exogenous variables and $\mathcal{W}$ denote its support. Similarly, we let $Y \equiv \left(S, P, X^{(2)}\right)$ denote the (endogenous and exogenous) regressors and $\mathcal{Y}$ denote its support. For every $y \in \mathcal{S} \times \mathcal{P}$, let $h_0(y) \equiv [h_{0,1}(y), ..., h_{0,J}(y)]' \equiv \left[\sigma_1^{-1}(y), ...\sigma_J^{-1}(y)\right]'$, so that the estimating equations become

$$x_j = h_{0,j}(y) + \xi_j, \quad j \in \mathcal{J}. \tag{11}$$

We assume that $h_0 \in \mathcal{H}$, where $\mathcal{H}$ is the Hölder ball of smoothness r, and we endow it with the norm $||\cdot||_\infty$ defined by $||h||_\infty \equiv \max_{j \in \mathcal{J}} ||h_j||_\infty$ for a function $h = [h_1, ...h_J]$.
Further, we let $\left\{\psi^{(i)}_{1,M_i}, ..., \psi^{(i)}_{M_i,M_i}\right\}$ be the collection of basis functions used to approximate $h_{0,i}$ for $j \in \mathcal{J}$, and let $M = \sum_{j=1}^J M_j$ be the dimension of the overall sieve space for h. Similarly, we let $\left\{a^{(i)}_{1,K_i}, ..., a^{(i)}_{K_i,K_i}\right\}$ be the collection of basis functions used to approximate the instrument space for $h_{0,i}$, and let and $K = \sum_{j=1}^J K_j$.

Next, for $j \in \mathcal{J}$, letting $\text{diag}(mat_1, ..., mat_J) \equiv \begin{bmatrix} mat_1 & 0_{d_{1,r} \times d_{2,c}} & \cdots & 0_{d_{1,r} \times d_{J,c}} \\ 0_{d_{2,r} \times d_{1,c}} & mat_2 & \cdots & 0_{d_{2,r} \times d_{J,c}} \\ \vdots & \vdots & \ddots & \vdots \\ 0_{d_{J,r} \times d_{1,c}} & 0_{d_{J,r} \times d_{2,c}} & \cdots & mat_J \end{bmatrix}$ for matrices $mat_j \in$

$\mathbb{R}^{d_{j,r}\times d_{j,c}}$ with $i\in\mathcal{J}$, we define

$$\psi^{(i)}_{M_i}(y)=\left(\psi^{(i)}_{1,M_i}(y),...,\psi^{(i)}_{M_i,M_i}(y)\right)' \qquad M_i-\text{by}-1$$
$$\psi_M(y)=\text{diag}\left(\psi^{(1)}_{M_1}(y),...,\psi^{(J)}_{M_J}(y)\right) \qquad M-\text{by}-J$$
$$\Psi_{(i)}=\left(\psi^{(i)}_{M_i}(y_1),...,\psi^{(i)}_{M_i}(y_T)\right)' \qquad T-\text{by}-M_i$$
$$a^{(i)}_{K_i}(w)=\left(a^{(i)}_{1,K_i}(w),...,a^{(i)}_{K_i,K_i}(w)\right)' \qquad K_i-\text{by}-1$$
$$a_K(w)=\text{diag}\left(a^{(1)}_{K_1}(w),...,a^{(J)}_{K_J}(w)\right) \qquad K-\text{by}-J$$
$$A_{(i)}=\left(a^{(i)}_{K_i}(w_1),...,a^{(i)}_{K_i}(w_T)\right)' \qquad T-\text{by}-K_i$$
$$A=\text{diag}\left(A_{(1)},...,A_{(J)}\right) \qquad JT-\text{by}-K$$
$$L_i=\mathbb{E}\left(a^{(i)}_{K_i}(W_t)\,\psi^{(i)}_{M_i}(Y_t)'\right) \qquad K_i-\text{by}-M_i$$
$$L=\text{diag}\left(L_1,...,L_J\right) \qquad K-\text{by}-M$$
$$\hat{L}_i=\frac{A'_{(i)}\Psi_{(i)}}{T} \qquad K_i-\text{by}-M_i$$
$$\hat{L}=\text{diag}\left(\hat{L}_1,...,\hat{L}_J\right) \qquad K-\text{by}-M$$
$$G_{A,i}=\mathbb{E}\left(a^{(i)}_{K_i}(W_t)\,a^{(i)}_{K_i}(W_t)'\right) \qquad K_i-\text{by}-K_i$$
$$G_A=\text{diag}\left(G_{A,1},...,G_{A,J}\right) \qquad K-\text{by}-K$$
$$\hat{G}_{A,i}=\frac{A'_{(i)}A_{(i)}}{T} \qquad K_i-\text{by}-K_i$$
$$\hat{G}_A=\text{diag}\left(\hat{G}_{A,1},...,\hat{G}_{A,J}\right) \qquad K-\text{by}-K$$

$$G_{\psi,i}=\mathbb{E}\left(\psi^{(i)}_{M_i}(Y_t)\,\psi^{(i)}_{M_i}(Y_t)'\right) \qquad M_i-\text{by}-M_i$$
$$G_\psi=\text{diag}\left(G_{\psi,1},...,G_{\psi,J}\right) \qquad M-\text{by}-M$$
$$X_{(i)}=\left(x_{i1},...,x_{iT}\right)' \qquad T-\text{by}-1$$
$$X=\left(X'_{(1)},...,X'_{(J)}\right)' \qquad JT-\text{by}-1$$

Also, we let

$$\Omega_{jj}=\mathbb{E}\left(\xi^2_{j,t}a^{(j)}_{K_j}(W_t)\,a^{(j)}_{K_j}(W_t)'\right) \qquad K_j-\text{by}-K_j$$
$$\Omega_{jk}=\Omega'_{kj}=\mathbb{E}\left(\xi_{j,t}\xi_{k,t}a^{(j)}_{K_j}(W_t)\,a^{(k)}_{K_k}(W_t)'\right) \qquad K_j-\text{by}-K_k$$
$$\Omega=\begin{bmatrix}\Omega_{11} & \Omega_{12} & \cdots & \Omega_{1J}\\ \Omega_{21} & \Omega_{22} & \cdots & \Omega_{2J}\\ \vdots & \vdots & \ddots & \cdots\\ \Omega_{J1} & \Omega_{J2} & \cdots & \Omega_{JJ}\end{bmatrix} \qquad K-\text{by}-K$$

and, similarly,

$$\hat{\Omega}_{jj} = \frac{1}{T}\sum_{t=1}^{T}\hat{\xi}_{j,t}^{2} a_{K_j}^{(j)}(w_t)\, a_{K_j}^{(j)}(w_t)' \qquad K_j - \text{by} - K_j$$

$$\hat{\Omega}_{jk} = \hat{\Omega}_{kj}' = \frac{1}{T}\sum_{t=1}^{T}\hat{\xi}_{j,t}\hat{\xi}_{j,t} a_{K_j}^{(j)}(w_t)\, a_{K_k}^{(k)}(w_t)' \qquad K_j - \text{by} - K_k$$

$$\hat{\Omega} = \begin{bmatrix} \hat{\Omega}_{11} & \hat{\Omega}_{12} & \cdots & \hat{\Omega}_{1J} \\ \hat{\Omega}_{21} & \hat{\Omega}_{22} & \cdots & \hat{\Omega}_{2J} \\ \vdots & \vdots & \ddots & \cdots \\ \hat{\Omega}_{J1} & \hat{\Omega}_{J2} & \cdots & \hat{\Omega}_{JJ} \end{bmatrix} \qquad K - \text{by} - K$$

where $\hat{\xi}_{j,t} = x_{j,t} - \hat{h}_j(y_t)$.
For $j \in \mathcal{J}$, we define

$$\zeta_{A,i} \equiv \sup_{w\in\mathcal{W}} \left|\left| G_{A,i}^{-1/2} a_{K_i}^{(i)}(w) \right|\right| \qquad \zeta_{\psi,i} \equiv \sup_{y\in\mathcal{Y}} \left|\left| G_{\psi,i}^{-1/2} \psi_{M_i}^{(i)}(y) \right|\right| \qquad \zeta_i \equiv \zeta_{A,i} \vee \zeta_{\psi,i}$$

and let $\zeta \equiv \max_{j\in\mathcal{J}} \zeta_j$. As in CC, we use the following sieve measure of ill-posedness, for $i \in \mathcal{J}$,

$$\tau_{M_i}^{(i)} \equiv \sup_{h_i\in\Psi_{M_i}:h_i\neq 0} \left(\frac{\mathbb{E}\left[(h_i(Y))^2\right]}{\mathbb{E}\left[(\mathbb{E}[h_i(Y)|W])^2\right]} \right)^{1/2}$$

where $\Psi_{M_i} \equiv clsp\left\{\psi_{M_i}^{(i)}\right\}$ and we let $\tau_M \equiv \max_{j\in\mathcal{J}} \tau_{M_j}^{(j)}$.
For every $2J-$vector of integers $\tilde{\alpha}$ and function $g : \mathcal{Y} \mapsto \mathbb{R}$, we let $|\tilde{\alpha}| \equiv \sum_{j=1}^{2J} \tilde{\alpha}_j$ and $\partial^{\tilde{\alpha}} g \equiv \frac{\partial^{|\tilde{\alpha}|} g}{\partial^{\tilde{\alpha}_1} s_1 \cdots \partial^{\tilde{\alpha}_J} s_J \partial^{\tilde{\alpha}_{J+1}} p_1 \cdots \partial^{\tilde{\alpha}_{2J}} p_J}$.
Similarly, for $h = [h_1, ..., h_J] : \mathcal{Y} \mapsto \mathbb{R}^J$, we let $\partial^{\tilde{\alpha}} h \equiv \left[\partial^{\tilde{\alpha}} h_1, \cdots, \partial^{\tilde{\alpha}} h_J\right]$.
The (unconstrained) sieve NPIV estimator $\hat{h}_i$ has the following closed form

$$\hat{h}_i(y) \quad = \quad \psi_M^{(i)}(y)'\,\hat{\theta}_i$$

for

$$\hat{\theta}_i \quad = \quad \left[\Psi_{(i)}' A_{(i)} \left(A_{(i)}' A_{(i)}\right)^{-} A_{(i)}' \Psi_{(i)}\right]^{-} \Psi_{(i)}' A_{(i)} \left(A_{(i)}' A_{(i)}\right)^{-} A_{(i)}' X_{(i)}$$

We write this in a more compact form as

$$\hat{\theta}_i \quad = \quad \frac{1}{T}\left[\hat{L}_i' \hat{G}_{A,i}^{-} \hat{L}_i\right]^{-} \hat{L}_i' \hat{G}_{A,i}^{-} A_{(i)}' X_{(i)}$$

Stacking the two estimators, we write

$$\hat{\theta} = \left(\hat{\theta}_1', \cdots, \hat{\theta}_2'\right)' = \frac{1}{T}\left[\hat{L}' \hat{G}_A^{-} \hat{L}\right]^{-} \hat{L}' \hat{G}_A^{-} A' X$$

and

$$\hat{h}(y) = \psi_M(y)'\,\hat{\theta}$$

Next, letting $H_{0,j} \equiv (h_{0,j}(y_1), ..., h_{0,j}(y_T))'$ and $H_0 \equiv \left(H'_{0,1}, \cdots, H'_{0,J}\right)'$, we define

$$\tilde{\theta} = \frac{1}{T}\left[\hat{L}'\hat{G}_A^{-}\hat{L}\right]^{-}\hat{L}'\hat{G}_A^{-}A'H_0 \tag{12}$$

and let

$$\tilde{h}(y) = \psi_M(y)'\tilde{\theta}$$

Given a functional $f : \mathcal{H} \mapsto \mathbb{R}$ and $h \in \mathcal{H}$, we let $vec_{g,J,j}$ be the column $J-$vector valued function that returns all zeros except for the $j-$th element, where it returns the function g. Further, we let

$$Df(h)\left[\psi^{(j)}_{M_j}\right] \equiv \left(Df(h)\left[vec_{\psi^{(j)}_{1,M_j},J,j}\right], ..., Df(h)\left[vec_{\psi^{(j)}_{M_j,M_j},J,j}\right]\right)' \qquad M_j - \text{by} - 1$$

$$Df(h)[\psi_M] \equiv \left(Df(h)\left[\psi^{(1)}_{M_1}\right]', ...Df(h)\left[\psi^{(J)}_{M_J}\right]'\right)' \qquad M - \text{by} - 1$$

and, for $(h, v) \in \mathcal{H} \times \mathcal{H}$, we let $Df(h)[v] \equiv \left.\frac{\partial f(h+\tau v)}{\partial \tau}\right|_{\tau=0}$ denote the pathwise derivative of f at h in the direction v.

Finally, we let

$$v_T^2(f) = Df(h_0)[\psi_M]'\left(S'G_A^{-1}S\right)^{-1}S'G_A^{-1}\Omega G_A^{-1}S\left(S'G_A^{-1}S\right)^{-1}Df(h_0)[\psi_M]$$

denote the sieve variance for the estimator $f\left(\hat{h}\right)$ of the functional f, and let the sieve variance estimator be

$$\hat{v}_T^2(f) = Df\left(\hat{h}\right)[\psi_M]'\left(\hat{L}'\hat{G}_A^{-1}\hat{L}\right)^{-1}\hat{L}'\hat{G}_A^{-1}\hat{\Omega}\hat{G}_A^{-1}\hat{L}\left(\hat{L}'\hat{G}_A^{-1}\hat{L}\right)^{-1}Df\left(\hat{h}\right)[\psi_M] \tag{13}$$

Because the functionals of interest are defined for fixed $(\overline{s}, \overline{p})$, they will typically be irregular, i.e. $v_T^2(f) \nearrow \infty$ as $T \to \infty$. Therefore, we tailor the proofs to this case.

C.2 Theorem 2: General Irregular Functionals

The following is largely based on the proof of Theorem D.1 in CC.

We make the following assumptions.

Assumption 5. *For all $j, k \in \mathcal{J}, j \neq k$:*
(i) $\sup_{w\in\mathcal{W}} \mathbb{E}\left(\xi_j^2|w\right) \leq \overline{\sigma}^2 < \infty$;
(ii) $\inf_{w\in\mathcal{W}} \mathbb{E}\left(\xi_j^2|w\right) \geq \underline{\sigma}^2 > 0$;
(iii) $\sup_{w\in\mathcal{W}} \mathbb{E}\left(|\xi_j\xi_k||w\right) \leq \overline{\sigma}_{cov} < \infty$;
(iv) $\sup_{w\in\mathcal{W}} \mathbb{E}\left[\xi_j^2\mathbb{I}\left\{\sum_{i=1}^J|\xi_i| > \ell(T)\right\}|w\right] = o(1)$ *for any positive sequence* $\ell(T) \nearrow \infty$;
(v) $\mathbb{E}\left(|\xi_j|^{2+\gamma^{(1)}}\right) < \infty$ *for some* $\gamma^{(1)} > 0$;
(vi) $\mathbb{E}\left(|\xi_j\xi_k|^{1+\gamma^{(2)}}\right) < \infty$ *for some* $\gamma^{(2)} > 0$.

Assumption 6. *(i)* $\tau_M\zeta\sqrt{M(\log M)/T} = o(1)$;
(ii) $\zeta^{\left(2+\gamma^{(1)}\right)/\gamma^{(1)}}\sqrt{(\log K)/T} = o(1)$ *and* $\zeta^{\left(1+\gamma^{(2)}\right)/\gamma^{(2)}}\sqrt{(\log K)/T} = o(1)$*, where* $\gamma^{(1)}, \gamma^{(2)} > 0$ *are defined in*

Assumption 5(v)-5(vi);
(iii) $K \asymp M$.
Assumption 7. *The basis used for the instrument spaces is the same across all goods, i.e.* $K_j = K_k$ *and* $a^{(j)}_{K_j}(\cdot) = a^{(k)}_{K_k}(\cdot)$ *for all* $j, k \in \mathcal{J}$*.*
Assumption 8. *Let* $\mathcal{H}_T \subset \mathcal{H}$ *be a sequence of neighborhoods of* h_0 *with* $\hat{h}, \tilde{h} \in \mathcal{H}_T$ *wpa1 and assume* $v_T(f) > 0$ *for every* T*. Further, assume that:*

(i) $v \mapsto Df(h_0)[v]$ *is a linear functional and there exists* α *with* $|\alpha| \geq 0$ *s.t.* $|Df(h_0)[h - h_0]| \lesssim ||\partial^\alpha h - \partial^\alpha h_0||_\infty$ *for all* $h \in \mathcal{H}_T$*;*

There are α_1, α_2 *with* $|\alpha_1|, |\alpha_2| \geq 0$ *s.t.*

(ii) $\left| f\left(\hat{h}\right) - f(h_0) - Df(h_0)[\hat{h} - h_0] \right| \lesssim ||\partial^{\alpha_1}\hat{h} - \partial^{\alpha_1}h_0||_\infty ||\partial^{\alpha_2}\hat{h} - \partial^{\alpha_2}h_0||_\infty$*;*

(iii) $\frac{\sqrt{T}}{\sigma_T(f)}\left(||\partial^{\alpha_1}\hat{h} - \partial^{\alpha_1}h_0||_\infty ||\partial^{\alpha_2}\hat{h} - \partial^{\alpha_2}h_0||_\infty + ||\partial^{\alpha}\tilde{h} - \partial^{\alpha}h_0||_\infty\right) = O_p(\eta_T)$ *for a nonnegative sequence* η_T *such that* $\eta_T = o(1)$*;*

(iv) $\frac{1}{v_T(f)}\left|\left| \left(Df\left(\hat{h}\right)[\psi_M]' - Df(h_0)[\psi_M]'\right)\left(G_A^{-1/2}S\right)^-_l \right|\right| = o_p(1)$*.*

Discussion of assumptions. Assumption 5 corresponds to Assumption 2 in CC, modified to account for the fact that my model has two main equations and thus two error terms. Assumption 6 restricts the growth rate of M and K; part 6(i) corresponds to the condition imposed by CC in Theorem D.1, while part 6(ii) is similar to Assumption 3(iii) in CC. Assumption 7 is stronger than necessary but we impose for simplicity. Assumption 8 corresponds to the sufficient conditions in Remark 4.1 of CC.

We now provide a proof of Theorem 2.

Proof of Theorem 2. We prove that

$$\sqrt{T}\frac{\left(f\left(\hat{h}\right) - f(h_0)\right)}{v_T(f)} \xrightarrow{d} N(0,1) \tag{14}$$

The result then follows from Lemma 5 below. By Assumption 8(ii),

$$\sqrt{T}\frac{\left(f\left(\hat{h}\right) - f(h_0)\right)}{v_T(f)} = \sqrt{T}\frac{Df(h_0)[\hat{h} - h_0]}{v_T(f)} + o_p\left(\underbrace{\frac{\sqrt{T}}{v_T(f)}||\partial^{\alpha_1}\hat{h} - \partial^{\alpha_1}h_0||_\infty ||\partial^{\alpha_2}\hat{h} - \partial^{\alpha_2}h_0||_\infty}_{c_T}\right)$$

By Assumption 8(iii), $c_T = o_p(1)$ and therefore,

$$\sqrt{T}\frac{\left(f\left(\hat{h}\right) - f(h_0)\right)}{v_T(f)} = \sqrt{T}\frac{Df(h_0)[\hat{h} - h_0]}{v_T(f)} + o_p(1) \tag{15}$$

Further, by Assumption 8(i)

$$Df(h_0)[\hat{h} - h_0] = Df(h_0)[\hat{h} - \tilde{h}] + Df(h_0)[\tilde{h} - h_0] \tag{16}$$

and

$$Df\left(h_0\right)[\tilde{h}-h_0] \lesssim ||\partial^{\alpha}\tilde{h}-\partial^{\alpha}h_0||_{\infty} \tag{17}$$

By (17) and Assumption 8(iii),

$$\sqrt{T}\frac{Df\left(h_0\right)[\tilde{h}-h_0]}{v_T\left(f\right)} = o_p\left(1\right) \tag{18}$$

Combining (15), (16) and (18), we obtain

$$\sqrt{T}\frac{\left(f\left(\hat{h}\right)-f\left(h_0\right)\right)}{v_T\left(f\right)} \quad = \quad \sqrt{T}\frac{Df\left(h_0\right)[\hat{h}-\tilde{h}]}{v_T\left(f\right)} + o_p\left(1\right) \tag{19}$$

We define

$$R_T\left(w\right) = \frac{Df\left(h_0\right)[\psi_M]'\left(S'G_A^{-1}S\right)^{-1}S'G_A^{-1}a_K\left(w\right)}{v_T\left(f\right)}$$

and note that $\mathbb{E}\left[\left(R_T\left(W_t\right)\cdot[\xi_{1,t},\cdots,\xi_{J,t}]'\right)^2\right]=1$. Then,

$$\begin{aligned}
\sqrt{T}\frac{Df\left(h_0\right)[\hat{h}-\tilde{h}]}{v_T\left(f\right)} &= \frac{1}{\sqrt{T}}\sum_{t=1}^{T}R_T\left(w_t\right)\cdot[\xi_{1,t},\cdots,\xi_{J,t}]' \\
&+\frac{Df\left(h_0\right)[\psi_M]'\left(\left(\hat{L}'\hat{G}_A^{-}\hat{L}\right)^{-}\hat{L}'\hat{G}_A^{-}-\left(S'G_A^{-1}S\right)^{-1}S'G_A^{-1}\right)\left(A'\xi/\sqrt{T}\right)}{v_T\left(f\right)} \\
&\equiv T_1+T_2
\end{aligned}$$

where $\xi\equiv[\xi_{1,1},\cdots,\xi_{1,T},\xi_{2,1},\cdots,\xi_{2,T},\cdots,\xi_{J,1},\cdots,\xi_{J,T}]'$.
We show that $T_1\xrightarrow{d}N\left(0,1\right)$ by the Lindeberg-Feller theorem. The Lindeberg condition requires that, for every $\epsilon>0$,

$$C_{0,T}\equiv\mathbb{E}\left[\left(R_T\left(W\right)\cdot[\xi_1,\cdots,\xi_J]'\right)^2\underbrace{\mathbb{I}\left\{\left|R_T\left(W\right)\cdot[\xi_1\cdots\xi_J]'\right|>\epsilon\sqrt{T}\right\}}_{Q_T(W,\xi)}\right]=o\left(1\right) \tag{20}$$

To show that this condition holds, note that

$$\begin{aligned}
R_T\left(w_t\right)\cdot[\xi_{1,t}\cdots\xi_{J,t}]' \quad &= \quad \sum_{i=1}^{J}\frac{Df\left(h_0\right)\left[\psi_{M_i}^{(i)}\right]'\left(L_i'G_{A,i}^{-1}L_i\right)^{-1}L_i'G_{A,i}^{-1}a_{K_i}^{(i)}\left(w_t\right)}{v_T\left(f\right)}\xi_{i,t} \\
&\equiv \quad \sum_{i=1}^{J}R_T^{(i)}\left(w_t\right)\xi_{i,t}
\end{aligned}$$

Now, for $i \in \mathcal{J}$,

$$\begin{aligned}
\left|R_T^{(i)}\left(w_t\right)\right| &\leq \frac{\left|\left|Df\left(h_0\right)\left[\psi_{M_i}^{(i)}\right]\left(L_i' G_{A,i}^{-1} L_i\right)^{-1} L_i' G_{A,i}^{-1/2}\right|\right|}{v_T\left(f\right)} \times \sup_{w \in \mathcal{W}} \left|\left|G_{A,i}^{-1/2} a_{K_i}^{(i)}\left(w\right)\right|\right| \\
&\equiv \lambda_i\left(T\right) \times \zeta_{A,i}
\end{aligned}$$

by the Cauchy-Schwarz inequality and thus

$$\left|\sum_{j=1}^{J} R_T^{(j)}\left(w_t\right)\xi_{j,t}\right| \leq \sum_{j=1}^{J}\left|\xi_{j,t}\right| \ \max_i\left[\lambda_i\left(T\right) \times \zeta_{A,i}\right] \tag{21}$$

Equation (21) implies

$$Q_T\left(w,\xi\right) \leq \mathbb{I}\left\{\sum_{j=1}^{J}\left|\xi_j\right| > \frac{\epsilon\sqrt{T}}{\max_i\left[\lambda_i\left(T\right) \times \zeta_{A,i}\right]}\right\} \equiv \overline{Q}_T\left(\xi\right)$$

for all $w \in \mathcal{W}$ and all $\xi \in \Xi$. Therefore, also using Cauchy-Schwarz and the law of iterated expectations, we have

$$\begin{aligned}
C_{0,T} &\leq \mathbb{E}\left[\sum_{j=1}^{J}\left(R_T^{(j)}\left(W\right)\right)^2 \times \sum_{j=1}^{J}\xi_j^2 \times \overline{Q}_T\left(\xi\right)\right] \\
&\leq \sum_{j=1}^{J}\mathbb{E}\left[\left(R_T^{(j)}\left(W\right)\right)^2\right]\sum_{j=1}^{J}\sup_{w \in \mathcal{W}}\mathbb{E}\left[\xi_j^2 \times \overline{Q}_T\left(\xi\right)|w\right]
\end{aligned}$$

Now, note that, for $i \in \mathcal{J}$,

$$\limsup_{T\to\infty}\mathbb{E}\left[\left(R_T^{(i)}\left(W\right)\right)^2\right] = \limsup_{T\to\infty}\lambda_i\left(T\right) < \infty$$

where the inequality follows from Lemma 8 below. Further, $\sup_{w\in\mathcal{W}}\mathbb{E}\left[\xi_i^2\overline{Q}_T\left(\xi\right)|w\right] = o\left(1\right)$ by Assumption 5(iv) and the fact that, by Assumption 6(i) and Lemma 8, $\frac{\sqrt{T}}{\max_i\left[\lambda_i(T)\zeta_{A,i}\right]} \nearrow \infty$. Therefore, $C_{0,T} = o\left(1\right)$, the Lindeberg condition is verified, and $T_1 \xrightarrow{d} N\left(0,1\right)$.

Next, for T_2, we have

$$\begin{aligned}
\left|T_2\right| &\leq v_T\left(f\right)^{-1}\left|\left|Df\left(h_0\right)\left[\psi_M\right]'\left(G_A^{-1/2}S\right)_l^{-}\right|\right| \left|\left|G_A^{-1/2}S\left\{\left(\hat{G}_A^{-1/2}\hat{L}\right)_l^{-}\hat{G}_A^{-1/2}G_A^{1/2} - \left(G_A^{-1/2}S\right)_l^{-}\right\}\right|\right| \left|\left|G_A^{-1/2}A'\xi/\sqrt{T}\right|\right| \\
&= \left[\left(\lambda_1\left(T\right)\right)^2 + \left(\lambda_2\left(T\right)\right)^2\right]^{1/2} \left|\left|G_A^{-1/2}S\left\{\left(\hat{G}_A^{-1/2}\hat{L}\right)_l^{-}\hat{G}_A^{-1/2}G_A^{1/2} - \left(G_A^{-1/2}S\right)_l^{-}\right\}\right|\right| \left|\left|G_A^{-1/2}A'\xi/\sqrt{T}\right|\right| \\
&\leq \left[\left(\lambda_1\left(T\right)\right)^2 + \left(\lambda_2\left(T\right)\right)^2\right]^{1/2} \max_{j\in\mathcal{J}}\left|\left|G_{A,i}^{-1/2}L_i\left\{\left(\hat{G}_{A,i}^{-1/2}\hat{L}_i\right)_l^{-}\hat{G}_{A,i}^{-1/2}G_{A,i}^{1/2} - \left(G_{A,i}^{-1/2}L_i\right)_l^{-}\right\}\right|\right| \left|\left|G_A^{-1/2}A'\xi/\sqrt{T}\right|\right| \\
&= O_p\left(\max_{j\in\mathcal{J}}\left[\tau_{M_i}^{(i)}\zeta_i\sqrt{M_i\log M_i/T}\right]\right) \\
&= o_p\left(1\right).
\end{aligned}$$

The first inequality follows from some algebra and the Cauchy-Schwarz inequality, the second inequality holds by the definition of matrix norm,the second equality is by Lemmas A.1, F.8 and F.10(c) in CC, Lemma 8 below and Assumption 6(iii), and the last step follows from Assumption 6(i). This completes the proof of (14).

□

Lemma 5. *Let $||\hat{h}-h_0||_\infty = o_p(1)$ and let Assumptions 5(i), 5(ii), 5(iii), 5(v), 5(vi), 6, 7, 8(iv) hold. Then*

$$\left|\frac{\hat{v}_T(f)}{v_T(f)} - 1\right| = o_p(1). \tag{22}$$

Proof. Following the proof of Lemma G.2 in CC, we write

$$\frac{\hat{v}_T^2(f)}{v_T^2(f)} - 1 = \frac{(\hat{\gamma}_T - \gamma_T)'\,\Omega^o\,(\hat{\gamma}_T + \gamma_T)}{v_T^2(f)} + \frac{\hat{\gamma}_T'\left(\hat{\Omega}^o - \Omega^o\right)\hat{\gamma}_T}{v_T^2(f)} \equiv T_1 + T_2$$

where

$$\hat{\Omega}^o = G_A^{-1/2}\hat{\Omega}G_A^{-1/2} \qquad \hat{\gamma}_T = G_A^{1/2}\hat{G}_A^{-1}\hat{L}\left(\hat{L}'\hat{G}_A^{-1}\hat{L}\right)^{-1} Df\left(\hat{h}\right)[\psi_M]$$

$$\Omega^o = G_A^{-1/2}\Omega G_A^{-1/2} \qquad \gamma_T = G_A^{-1/2}S\left(S'G_A^{-1}S\right)^{-1} Df(h_0)[\psi_M]$$

and note that $\frac{||\gamma_T||^2}{v_T^2(f)} = \sum_{j=1}^{J}\lambda_j(T)^2$.
We consider T_1 and T_2 in turn. Note that

$$\begin{aligned}\frac{||\hat{\gamma}_T - \gamma_T||}{v_T(f)} &= \frac{1}{v_T(f)}\left|\left|Df\left(\hat{h}\right)[\psi_M]'\left(\hat{G}_A^{-1/2}\hat{L}\right)_l^- \hat{G}_A^{-1/2}G_A^{1/2} - Df(h_0)[\psi_M]'\left(G_A^{-1/2}S\right)_l^-\right|\right| \\ &\le \frac{1}{v_T(f)}\left|\left|Df\left(\hat{h}\right)[\psi_M]'\left(G_A^{-1/2}S\right)_l^-\right|\right| \times \left|\left|G_A^{-1/2}S\left\{\left(\hat{G}_A^{-1/2}\hat{L}\right)_l^- \hat{G}_A^{-1/2}G_A^{1/2} - \left(G_A^{-1/2}S\right)_l^-\right\}\right|\right| \\ &+ \frac{1}{v_T(f)}\left|\left|\left(Df\left(\hat{h}\right)[\psi_M]' - Df(h_0)[\psi_M]'\right)\left(G_A^{-1/2}S\right)_l^-\right|\right| \equiv T_1^{(1)} \times T_1^{(2)} + T_1^{(3)}\end{aligned}$$

Now,

$$\begin{aligned}T_1^{(1)} &\le \frac{1}{v_T(f)}\left|\left|\left(Df\left(\hat{h}\right)[\psi_M]' - Df(h_0)[\psi_M]'\right)\left(G_A^{-1/2}S\right)_l^-\right|\right| + \frac{J}{v_T(f)}\max_{i\in\mathcal{J}}\left|\left|Df(h_0)\left[\psi_{M_i}^{(i)}\right]'\left(G_{A,i}^{-1/2}L_i\right)_l^-\right|\right| \\ &= O_p(1)\end{aligned}$$

where the last step follows from Assumption 8(iv) and Lemma 8. Further, $T_1^{(2)} = o_p(1)$ by Lemmas F.10(c) and A.1 in CC and Assumption 6(i),and $T_1^{(3)} = o_p(1)$ by Assumption 8(iv). This implies that

$$\frac{||\hat{\gamma}_T - \gamma_T||}{v_T(f)} = o_p(1). \tag{23}$$

Therefore, by Cauchy-Schwarz

$$|T_1| \le \frac{||\hat{\gamma}_T - \gamma_T||}{v_T(T)} \times ||\Omega^o|| \times \frac{||\hat{\gamma}_T + \gamma_T||}{v_T(T)} \le \frac{||\hat{\gamma}_T - \gamma_T||}{v_T(T)} \times ||\Omega^o|| \times \left(\frac{||\hat{\gamma}_T - \gamma_T|| + 2||\gamma_T||}{v_T(T)}\right) = o_p(1)$$

where in the last step we also use Lemma 8 and the fact that $||\Omega^o|| < \infty$ by Assumption 5(i), 5(ii), 5(iii).

Turning to $|T_2|$, note that

$$\begin{aligned}
|T_2| &\leq \frac{||\hat{\gamma}_T||}{v_T(T)} \times ||\hat{\Omega}^o - \Omega^o|| \times \frac{||\hat{\gamma}_T||}{v_T(T)} \\
&\leq \frac{||\hat{\gamma}_T - \gamma_T|| + ||\gamma_T||}{v_T(T)} \times ||\hat{\Omega}^o - \Omega^o|| \times \frac{||\hat{\gamma}_T - \gamma_T|| + ||\gamma_T||}{v_T(T)} \\
&= O_p(1) \times ||\hat{\Omega}^o - \Omega^o|| \times O_p(1)
\end{aligned}$$

where the last step follows again from Lemma 8 and (23). We complete the proof by showing that $||\hat{\Omega}^o - \Omega^o|| = O_p(1)$. Note that

$$\Omega^o = \begin{bmatrix} \Omega^o_{11} & \Omega^o_{12} & \cdots & \Omega^o_{1J} \\ \Omega^o_{21} & \Omega^o_{22} & \cdots & \Omega^o_{2J} \\ \vdots & \vdots & \ddots & \cdots \\ \Omega^o_{J1} & \Omega^o_{J2} & \cdots & \Omega^o_{JJ} \end{bmatrix} \qquad \hat{\Omega}^o = \begin{bmatrix} \hat{\Omega_{11}}^o & \hat{\Omega_{12}}^o & \cdots & \hat{\Omega_{1J}}^o \\ \hat{\Omega_{21}}^o & \hat{\Omega_{22}}^o & \cdots & \hat{\Omega_{2J}}^o \\ \vdots & \vdots & \ddots & \cdots \\ \hat{\Omega_{J1}}^o & \hat{\Omega_{J2}}^o & \cdots & \hat{\Omega_{JJ}}^o \end{bmatrix}$$

where, for $j, k \in \mathcal{J}$,

$$\Omega^o_{jk} = G^{-1/2}_{A,j} \Omega_{jk} G^{-1/2}_{A,k} \qquad \hat{\Omega}^o_{jk} = G^{-1/2}_{A,j} \hat{\Omega}_{jk} G^{-1/2}_{A,k}$$

Using this notation, we have that, for some $v = [v_1' \cdots v_J']'$, with $v_i \in \mathbb{R}^{K_i}$, $j \in \mathcal{J}$, and $||v|| = 1$,

$$\begin{aligned}
||\hat{\Omega}^o - \Omega^o|| &= \sum_{j=1}^{J} v_j' \left(\hat{\Omega}^o_{jj} - \Omega^o_{jj}\right) v_j + 2 \sum_{j=1}^{J} \sum_{k=1}^{j-1} v_j' \left(\hat{\Omega}^o_{jk} - \Omega^o_{jk}\right) v_k \\
&\leq J \max_{j \in \mathcal{J}} ||\hat{\Omega}^o_{jj} - \Omega^o_{jj}|| + 2J^2 \max_{j,k \in \mathcal{J}, j \neq k} ||\hat{\Omega}^o_{jk} - \Omega^o_{jk}||
\end{aligned}$$

where each step uses the definition of matrix norm and the second step also uses Cauchy-Schwarz and the fact that $||v_j|| \leq 1$ for all j. Now, $||\hat{\Omega}^o_{ii} - \Omega^o_{ii}|| = O_p(1)$ for $i \in \mathcal{J}$ by Lemma G.3 in CC.For the third term, note that, by the triangle inequality, for all $j, k \in \mathcal{J}, j \neq k$,

$$\begin{aligned}
||\hat{\Omega}^o_{jk} - \Omega^o_{jk}|| \leq & \left|\left| G^{-1/2}_{A,j} \left[\frac{1}{T} \sum_{t=1}^{T} \xi_{j,t} \xi_{k,t} a^{(j)}_{K_j}(w_t) \, a^{(j)}_{K_j}(w_t)' - \mathbb{E}\left(\xi_j \xi_k a^{(j)}_{K_j}(W) \, a^{(j)}_{K_j}(Z)' \right) \right] G^{-1/2}_{A,j} \right|\right| \\
&+ \left|\left| G^{-1/2}_{A,j} \frac{1}{T} \sum_{t=1}^{T} \left[\left(\hat{\xi}_{j,t} - \xi_{j,t} \right) \xi_{k,t} a^{(j)}_{K_j}(w_t) \, a^{(j)}_{K_j}(w_t)' \right] G^{-1/2}_{A,j} \right|\right| \\
&+ \left|\left| G^{-1/2}_{A,j} \frac{1}{T} \sum_{t=1}^{T} \left[\left(\hat{\xi}_{j,t} - \xi_{j,t} \right) \left(\hat{\xi}_{k,t} - \xi_{k,t} \right) a^{(j)}_{K_j}(w_t) \, a^{(j)}_{K_j}(w_t)' \right] G^{-1/2}_{A,j} \right|\right| \\
&+ \left|\left| G^{-1/2}_{A,j} \frac{1}{T} \sum_{t=1}^{T} \left[\xi_{1,t} \left(\hat{\xi}_{2,t} - \xi_{2,t} \right) a^{(j)}_{K_j}(w_t) \, a^{(j)}_{K_j}(w_t)' \right] G^{-1/2}_{A,j} \right|\right| \\
&\equiv ||T_{\Omega,1}|| + ||T_{\Omega,2}|| + ||T_{\Omega,3}|| + ||T_{\Omega,4}||
\end{aligned}$$

where we use the fact that $G_{A,j} = G_{A,k}$ and $a^{(j)}_{K_j} = a^{(k)}_{K_k}$ for all $j, k \in \mathcal{J}$ by Assumption 7. Using Lemma 9 below, we obtain $||T_{\Omega,1}|| = O_p(1)$. Further, $||T_{\Omega,2}|| = O_p(1)$ by $\left(\hat{\xi}_{j,t} - \xi_{j,t} \right) \xi_{k,t} \leq ||\hat{h}_j - h_{0,j}||_\infty \left(1 + \xi^2_{k,t} \right)$ and Lemma F.7 in CC.Similarly, $||T_{\Omega,4}|| = O_p(1)$. Finally, $||T_{\Omega,3}|| = O_p(1)$ by $\left(\hat{\xi}_{j,t} - \xi_{j,t} \right) \left(\hat{\xi}_{k,t} - \xi_{k,t} \right) \leq ||\hat{h}_j - h_{0,j}||^2_\infty$ and Lemma F.7 in CC.

□

Remark 1. *Note that I do not impose Assumption 4(i) in CC. This is because the assumption is automatically satisfied if the basis functions used for the sieve space and those used for the instrument space are both Riesz bases for the conditional expectation operator. I follow CC in assuming that this is the case.*

C.3 Theorem 3: Price elasticity functionals

We now focus on the case where the functional f is the own-price price elasticity of good 1 at a fixed $(\overline{s}, \overline{p}) \equiv (\overline{s}_1, \overline{s}_2, \overline{p}_1, \overline{p}_2)$ and Bernstein polynomials are used for both the sieve space and the instrument space. The goal is to provide sufficient, lower-level conditions for Theorem 2. Analogous arguments hold for the own-price elasticity of good 2 and for the cross price elasticities.
The functional of interest takes the form

$$f_\epsilon(h_0) = -\frac{p_1}{s_1} \frac{\frac{\partial h_{0,2}(\overline{s},\overline{p})}{\partial s_2} \frac{\partial h_{0,1}(\overline{s},\overline{p})}{\partial p_1} - \frac{\partial h_{0,1}(\overline{s},\overline{p})}{\partial s_2} \frac{\partial h_{0,2}(\overline{s},\overline{p})}{\partial p_1}}{\frac{\partial h_{0,1}(\overline{s},\overline{p})}{\partial s_1} \frac{\partial h_{0,2}(\overline{s},\overline{p})}{\partial s_2} - \frac{\partial h_{0,1}(\overline{s},\overline{p})}{\partial s_2} \frac{\partial h_{0,2}(\overline{s},\overline{p})}{\partial s_1}} \equiv -\frac{p_1}{s_1} \frac{N_1 - N_2}{D_1 - D_2} \tag{24}$$

We make the following assumptions.
Assumption 9. *(i) P has bounded support and (P, S) have densities bounded away from 0 and ∞;*
(ii) The basis used for both the sieve space and the instrument space is tensor-product Bernstein polynomials. Further, for the sieve space, the univariate Bernstein polynomials all have the same degree $M^{1/4}$;
(iii) $h_0 = [h_{0,1}, h_{0,2}]$ where $h_{0,1}$ and $h_{0,2}$ belong to the Hölder ball of smoothness $r \geq 4$ and finite radius L, and the order of the tensor-product Bernstein polynomials used for the sieve space is greater than r;
(iv) $M^{\frac{2+\gamma^{(1)}}{2\gamma^{(1)}}} \sqrt{\frac{\log T}{T}} = o(1)$ and $M^{\frac{1+\gamma^{(2)}}{2\gamma^{(2)}}} \sqrt{\frac{\log T}{T}} = o(1)$, where $\gamma^{(1)}, \gamma^{(2)} > 0$ are defined in Assumption 5(v)-5(vi);
(v) $\frac{\sqrt{T}}{v_T(f_\epsilon)} \times \left(M^{\frac{3-r}{4}} + \tau_M^2 M^{9/4} \frac{\log M}{T} \right) = o(1)$.

Discussion of Assumptions. Assumptions 9(i), 9(iii) and 9(iv) are regularity conditions sufficient to apply the sup-norm rate results in CC. 9(ii) is made for simplicity but it is not necessary. 9(v) corresponds to the second part of Assumption CS(v) in CC and is used to verify Assumption 8. More concrete sufficient conditions for Assumptions 9(iv) and 9(v) may be provided in specific settings. For example, Lemma 6 below gives sufficient conditions for the mildly ill-posed case.[76]

We now provide a proof of Theorem 3.

Proof of Theorem 3. We prove the statement by showing that the assumptions of Theorem 2 hold. Assumptions 5, 6(iii) and 7 are maintained. Assumption 6(i) is implied by Assumption 9(v), Lemma 10 and the fact that $\zeta = O_p\left(\sqrt{M}\right)$ for Bernstein polynomials.[77] Similarly, Assumption 6(ii) is implied by Assumption 9(iv) and $\zeta = O_p\left(\sqrt{M}\right)$.
We now verify Assumption 8. In what follows, unless otherwise specified, it is assumed that the arguments of all functions are $(\overline{s}, \overline{p})$ and the dependence is suppressed for notational convenience.

[76] See CC (p.15) for a formal definition of mild and severe ill-posedness.

[77] This follows from the fact that Bernstein polynomials are a special case of B-splines (see, e.g., Remark 4.7 on p.188 of Schumaker (2007).).

8(*i*) The pathwise derivative of f_ϵ in the direction $v \equiv (v_1, v_2)' \in \mathcal{H}$ is

$$Df_\epsilon(h_0)[v] \equiv \left.\frac{\partial f_\epsilon(h_0+\tau v)}{\partial \tau}\right|_{\tau=0} = \frac{p_1}{s_1}\left(C_1\frac{\partial v_2}{\partial s_2} + C_2\frac{\partial v_1}{\partial s_2} + C_3\frac{\partial v_1}{\partial p_1} + C_4\frac{\partial v_2}{\partial p_1} + C_5\frac{\partial v_2}{\partial s_1} + C_6\frac{\partial v_1}{\partial s_1}\right) \tag{25}$$

where

$$C_1 = -\frac{(D_1-D_2)\frac{\partial h_{0,1}}{\partial p_1} - (N_1-N_2)\frac{\partial h_{0,1}}{\partial s_1}}{(D_1-D_2)^2} \qquad C_2 = -\frac{-(D_1-D_2)\frac{\partial h_{0,2}}{\partial p_1} + (N_1-N_2)\frac{\partial h_{0,2}}{\partial s_1}}{(D_1-D_2)^2}$$

$$C_3 = -\frac{\frac{\partial h_{0,2}}{\partial s_2}}{(D_1-D_2)} \qquad C_4 = \frac{\frac{\partial h_{0,1}}{\partial s_2}}{(D_1-D_2)}$$

$$C_5 = -\frac{(N_1-N_2)\frac{\partial h_{0,1}}{\partial s_2}}{(D_1-D_2)^2} \qquad C_6 = \frac{(N_1-N_2)\frac{\partial h_{0,2}}{\partial s_2}}{(D_1-D_2)^2}$$

Therefore, $Df_\epsilon(h_0) : \mathcal{H} \mapsto \mathbb{R}$ is a linear functional.
Next, note that, for any $h = [h_1, h_2] \in \mathcal{H}_T$,

$$\begin{aligned}\left|\frac{\partial h_1}{\partial s_1} - \frac{\partial h_{0,1}}{\partial s_1}\right| &\leq \int_{-\infty}^{\overline{s}_2}\int_{-\infty}^{\overline{p}_1}\left|\frac{\partial^3 h_1}{\partial s_1 \partial s_2 \partial p_1}\left(\overline{s}_1, \underline{s}_2, \underline{p}_1, \overline{p}_2\right) - \frac{\partial^3 h_{0,1}}{\partial s_1 \partial s_2 \partial p_1}\left(\overline{s}_1, \underline{s}_2, \underline{p}_1, \overline{p}_2\right)\right| d\underline{s}_2 d\underline{p}_1 \\ &\leq constant\left\|\frac{\partial^3 h_1}{\partial s_1 \partial s_2 \partial p_1} - \frac{\partial^3 h_{0,1}}{\partial s_1 \partial s_2 \partial p_1}\right\|_\infty\end{aligned}$$

where the first inequality follows from the triangle inequality and the fundamental theorem of calculus, and the second inequality follows from assumption $9(i)$ and the fact that the support of (S_1, S_2) is the unit simplex and thus trivially bounded. By a similar argument, we can bound all the other derivatives in (25) and write

$$\begin{aligned}Df_\epsilon(h_0)[h-h_0] &\leq constant \times \max\left\{\left\|\frac{\partial^3 h_1}{\partial s_1 \partial s_2 \partial p_1} - \frac{\partial^3 h_{0,1}}{\partial s_1 \partial s_2 \partial p_1}\right\|_\infty, \left\|\frac{\partial^3 h_2}{\partial s_1 \partial s_2 \partial p_1} - \frac{\partial^3 h_{0,2}}{\partial s_1 \partial s_2 \partial p_1}\right\|_\infty\right\} \\ &\equiv constant\left\|\frac{\partial^3 h}{\partial s_1 \partial s_2 \partial p_1} - \frac{\partial^3 h_0}{\partial s_1 \partial s_2 \partial p_1}\right\|_\infty\end{aligned}$$

which shows that Assumption $8(i)$ holds with $\alpha = [1, 1, 1, 0]$.

8(*ii*) By the mean value theorem,

$$\begin{aligned}f_\epsilon\left(\hat{h}\right) - f_\epsilon(h_0) = \frac{p_1}{s_1}\Bigg[&\tilde{C}_1\left(\frac{\partial \hat{h}_2}{\partial s_2} - \frac{\partial h_{0,2}}{\partial s_2}\right) + \tilde{C}_2\left(\frac{\partial \hat{h}_1}{\partial s_2} - \frac{\partial h_{0,1}}{\partial s_2}\right) + \tilde{C}_3\left(\frac{\partial \hat{h}_1}{\partial p_1} - \frac{\partial h_{0,1}}{\partial p_1}\right) + \tilde{C}_4\left(\frac{\partial \hat{h}_2}{\partial p_1} - \frac{\partial h_{0,2}}{\partial p_1}\right) \\ &+ \tilde{C}_5\left(\frac{\partial \hat{h}_2}{\partial s_1} - \frac{\partial h_{0,2}}{\partial s_1}\right) + \tilde{C}_6\left(\frac{\partial \hat{h}_1}{\partial s_1} - \frac{\partial h_{0,1}}{\partial s_1}\right)\Bigg]\end{aligned}$$

$$\tilde{C}_1 = -\frac{\left(\tilde{D}_1 - \tilde{D}_2\right)\frac{\partial \tilde{h}_1}{\partial p_1} - \left(\tilde{N}_1 - \tilde{N}_2\right)\frac{\partial \tilde{h}_1}{\partial s_1}}{\left(\tilde{D}_1 - \tilde{D}_2\right)^2} \qquad \tilde{C}_2 = -\frac{-\left(\tilde{D}_1 - \tilde{D}_2\right)\frac{\partial \tilde{h}_2}{\partial p_1} + \left(\tilde{N}_1 - \tilde{N}_2\right)\frac{\partial \tilde{h}_2}{\partial s_1}}{\left(\tilde{D}_1 - \tilde{D}_2\right)^2}$$

$$\tilde{C}_3 = -\frac{\frac{\partial \tilde{h}_2}{\partial s_2}}{\left(\tilde{D}_1 - \tilde{D}_2\right)} \qquad \tilde{C}_4 = \frac{\frac{\partial \tilde{h}_1}{\partial s_2}}{\left(\tilde{D}_1 - \tilde{D}_2\right)}$$

$$\tilde{C}_5 = -\frac{\left(\tilde{N}_1 - \tilde{N}_2\right)\frac{\partial \tilde{h}_1}{\partial s_2}}{\left(\tilde{D}_1 - \tilde{D}_2\right)^2} \qquad \tilde{C}_6 = \frac{\left(\tilde{N}_1 - \tilde{N}_2\right)\frac{\partial \tilde{h}_2}{\partial s_2}}{\left(\tilde{D}_1 - \tilde{D}_2\right)^2}$$

where $\left[\frac{\partial \tilde{h}_1}{\partial p_1}, \frac{\partial \tilde{h}_1}{\partial s_1}, \frac{\partial \tilde{h}_1}{\partial s_2}, \frac{\partial \tilde{h}_2}{\partial p_1}, \frac{\partial \tilde{h}_2}{\partial s_1}, \frac{\partial \tilde{h}_2}{\partial s_2}\right]$ lies on the line segment between $\left[\frac{\partial h_{0,1}}{\partial p_1}, \frac{\partial h_{0,1}}{\partial s_1}, \frac{\partial h_{0,1}}{\partial s_2}, \frac{\partial h_{0,2}}{\partial p_1}, \frac{\partial h_{0,2}}{\partial s_1}, \frac{\partial h_{0,2}}{\partial s_2}\right]$ and $\left[\frac{\partial \hat{h}_1}{\partial p_1}, \frac{\partial \hat{h}_1}{\partial s_1}, \frac{\partial \hat{h}_1}{\partial s_2}, \frac{\partial \hat{h}_2}{\partial p_1}, \frac{\partial \hat{h}_2}{\partial s_1}, \frac{\partial \hat{h}_2}{\partial s_2}\right]$ and $\tilde{N}_1, \tilde{N}_2, \tilde{D}_1, \tilde{D}_2$ are defined accordingly. Therefore, after some algebra, we obtain

$$\begin{aligned}\left|f_\epsilon\left(\hat{h}\right) - f_\epsilon\left(h_0\right) - Df_\epsilon\left(h_0\right)\left[\hat{h} - h_0\right]\right| \leq\ & F_1\left|\frac{\partial \hat{h}_2}{\partial s_2} - \frac{\partial h_{0,2}}{\partial s_2}\right| + F_2\left|\frac{\partial \hat{h}_1}{\partial s_2} - \frac{\partial h_{0,1}}{\partial s_2}\right| + F_3\left|\frac{\partial \hat{h}_1}{\partial p_1} - \frac{\partial h_{0,1}}{\partial p_1}\right| \\ +\ & F_4\left|\frac{\partial \hat{h}_2}{\partial p_1} - \frac{\partial h_{0,2}}{\partial p_1}\right| + F_5\left|\frac{\partial \hat{h}_2}{\partial s_1} - \frac{\partial h_{0,2}}{\partial s_1}\right| + F_6\left|\frac{\partial \hat{h}_1}{\partial s_1} - \frac{\partial h_{0,1}}{\partial s_1}\right|\end{aligned}$$

where $(F_i)_{i=1}^6$ are linear combinations (with finite coefficients) of $||\partial^{\tilde{\alpha}}\hat{h} - \partial^{\tilde{\alpha}}h_0||_\infty$ for vectors $\tilde{\alpha}$ with $|\tilde{\alpha}| = 1$. Thus

$$\left|f_\epsilon\left(\hat{h}\right) - f_\epsilon\left(h_0\right) - Df_\epsilon\left(h_0\right)\left[\hat{h} - h_0\right]\right| \leq constant||\partial^{\alpha_1}\hat{h} - \partial^{\alpha_1}h_0||_\infty||\partial^{\alpha_2}\hat{h} - \partial^{\alpha_2}h_0||_\infty$$

for some α_1, α_2 with $|\alpha_1| = |\alpha_2| = 1$.

8(iii) Given the choice of $\alpha, \alpha_1, \alpha_2$ above and by Corollary 3.1 in CC, we have

$$\begin{aligned}\left|\left|\partial^{\alpha_1}\hat{h} - \partial^{\alpha_1}h_0\right|\right|_\infty\left|\left|\partial^{\alpha_2}\hat{h} - \partial^{\alpha_2}h_0\right|\right|_\infty + \left|\left|\partial^{\alpha}\tilde{h} - \partial^{\alpha}h_0\right|\right|_\infty =\ & O_p\left(\left[M^{\frac{1-r}{4}} + \tau_M M^{3/4}\sqrt{\log M/T}\right]^2\right) \\ +\ & O_p\left(M^{\frac{3-r}{4}}\right)\end{aligned}$$

Thus, Assumption 8(iii) is implied by Assumption 9(v) and Lemma 10.

8(iv) By Remark 4.1 in CC, a sufficient condition for Assumption 8(iv) is

$$T_{iv,\epsilon} \equiv \frac{\tau_M\sqrt{\sum_{m=1}^M\left(Df_\epsilon\left(\hat{h}\right)\left[\left(G_\psi^{-1/2}\psi_M\right)_m\right] - Df_\epsilon\left(h_0\right)\left[\left(G_\psi^{-1/2}\psi_M\right)_m\right]\right)^2}}{v_T\left(f_\epsilon\right)} = O_p\left(1\right)$$

where $\left(G_\psi^{-1/2}\psi_M\right)_m$ denotes the $m-$th row of the matrix $G_\psi^{-1/2}\psi_M$. Note that, after some algebra, we can write $Df_\epsilon\left(\hat{h}\right)\left[\left(G_\psi^{-1/2}\psi_M\right)_m\right] - Df_\epsilon\left(h_0\right)\left[\left(G_\psi^{-1/2}\psi_M\right)_m\right]$ for every m as the linear combination of terms, where each term is the difference between a first-order partial derivative of $\hat{h}_i$ and the same derivative of $h_{0,i}$ for some $i \in \{1,2\}$, and each coefficient is a first-order partial derivative of an element of $\left(G_\psi^{-1/2}\psi_M\right)_m$.

Therefore, we can write

$$
\begin{aligned}
T_{iv,\epsilon} &\lesssim \frac{\tau_M}{\upsilon_T(f_\epsilon)} \sqrt{M^{3/2}\zeta^2 \times \left[O_p\left(M^{(1-r)/4} + \tau_M M^{3/4}\sqrt{\log M/T}\right)\right]^2} \\
&= O_p\left(\frac{\sqrt{T}}{\upsilon_T(f_\epsilon)} \times \left[\frac{\tau_M M}{\sqrt{T}} M^{(2-r)/4} + \frac{\tau_M^2 M^2 \sqrt{\log M}}{T}\right]\right) = O_p(1)
\end{aligned}
$$

where the first step uses Corollary 3.1 in CC and the proof of Lemma 10, the second step uses the fact that, under the maintained assumptions, $\zeta = O\left(\sqrt{M}\right)$, and the last step follows from Assumption 9(v) (which implies $\frac{\tau_M M}{\sqrt{T}} = o(1)$).

□

Lemma 6. *Let Assumptions 9(i) and 9(iii) hold, and let $(\upsilon_T(f_\epsilon))^2 \asymp \tau_M^2 \sum_{m=1}^M m^a$ for $a \leq 0$. Further, assume that $\tau_M \asymp M^{\varsigma/2}$ for $\varsigma \geq 0, a+\varsigma+1 > 0$.*[78] *Then, Assumptions 9(iv) and 9(v) are satisfied if $M \asymp T^\rho$ with*

$$
\rho \in \left(\frac{2}{r-3+2(a+\varsigma+1)}, \min\left\{\frac{2}{2\varsigma-2a+7}, \frac{\gamma^{(1)}}{2+\gamma^{(1)}}, \frac{\gamma^{(2)}}{1+\gamma^{(2)}}\right\}\right)
$$

Further, M may be chosen to satisfy the latter condition if $r+4a-8>0$ and $\gamma^{(i)}(r+2a+2\varsigma-3)-4>0$ for $i \in \{1,2\}$.

Proof. As shown in CC, under the maintained assumptions, $(\upsilon_T(f_\epsilon))^2 \asymp M^{a+\varsigma+1}$. The result follows by inspection.

□

C.4 Theorem 4: Equilibrium price functionals

We now specialize Theorem 2 to the case where the functional f is the equilibrium price of good 1 in a market characterized by marginal costs $\overline{mc} \equiv (\overline{mc}_1, \overline{mc}_2)$ and indices $\overline{\delta} \equiv \left(\overline{\delta}_1, \overline{\delta}_2\right)$. I let $f_p \equiv [f_{p_1}, f_{p_2}] : \mathcal{H} \mapsto \mathbb{R}^2$ denote the functional that returns the equilibrium prices, so that the goal is to obtain the asymptotic distribution of the sieve estimator $f_{p_1}\left(\hat{h}\right)$. An analogous argument holds for the price of good 2. As in the rest of the appendix, I let $h_0 = [h_{0,1}, h_{0,2}]$ denote the inverse of the demand system σ_0. Further, I use $h_0^{-1} = \left[h_{0,1}^{-1}, h_{0,2}^{-1}\right] = [\sigma_{0,1}, \sigma_{0,2}]$ to denote the demand system itself. The equilibrium prices $\overline{p} \equiv (\overline{p}_1, \overline{p}_2) \equiv [f_{p_1}(h_0), f_{p_2}(h_0)]$ solve the firm's first-order conditions

$$
\begin{bmatrix} g_1\left(\overline{\delta}, \overline{p}, \overline{mc}, h_0\right) \\ g_2\left(\overline{\delta}, \overline{p}, \overline{mc}, h_0\right) \end{bmatrix} \equiv -\left[\left(\mathbb{J}_{h_0}^s\right)^{-1} \mathbb{J}_{h_0}^p\right]' \begin{bmatrix} \overline{p}_1 - \overline{mc}_1 \\ \overline{p}_2 - \overline{mc}_2 \end{bmatrix} + \begin{bmatrix} h_{0,1}^{-1}\left(\overline{\delta}, \overline{p}\right) \\ h_{0,2}^{-1}\left(\overline{\delta}, \overline{p}\right) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \tag{26}
$$

where

$$
\mathbb{J}_{h_0}^s \equiv \begin{bmatrix} \frac{\partial h_{0,1}\left(h_0^{-1}\left(\overline{\delta},\overline{p}\right),\overline{p}\right)}{\partial s_1} & \frac{\partial h_{0,1}\left(h_0^{-1}\left(\overline{\delta},\overline{p}\right),\overline{p}\right)}{\partial s_2} \\ \frac{\partial h_{0,2}\left(h_0^{-1}\left(\overline{\delta},\overline{p}\right),\overline{p}\right)}{\partial s_1} & \frac{\partial h_{0,2}\left(h_0^{-1}\left(\overline{\delta},\overline{p}\right),\overline{p}\right)}{\partial s_2} \end{bmatrix} \qquad \mathbb{J}_{h_0}^p \equiv \begin{bmatrix} \frac{\partial h_{0,1}\left(h_0^{-1}\left(\overline{\delta},\overline{p}\right),\overline{p}\right)}{\partial p_1} & \frac{\partial h_{0,1}\left(h_0^{-1}\left(\overline{\delta},\overline{p}\right),\overline{p}\right)}{\partial p_2} \\ \frac{\partial h_{0,2}\left(h_0^{-1}\left(\overline{\delta},\overline{p}\right),\overline{p}\right)}{\partial p_1} & \frac{\partial h_{0,2}\left(h_0^{-1}\left(\overline{\delta},\overline{p}\right),\overline{p}\right)}{\partial p_2} \end{bmatrix}
$$

[78] This corresponds to the "mildly ill-posed" case discussed by CC in Corollary 5.1. CC also provide sufficient conditions for the maintained assumption on the rate of divergence of $\upsilon_T(f)$.

We make the following assumptions.

Assumption 10. *(i) P has bounded support and (P,S) have densities bounded away from 0 and ∞;*
(ii) The basis used for both the sieve space and the instrument space is tensor-product Bernstein polynomials. Further, for the sieve space, the univariate Bernstein polynomials all have the same degree $M^{1/4}$;
(iii) $h_0 = [h_{0,1}, h_{0,2}]$ where $h_{0,1}$ and $h_{0,2}$ belong to the Hölder ball of smoothness $r \geq 5$ and finite radius L, and the order of the tensor-product Bernstein polynomials used for the sieve space is greater than r;
(iv) $M^{\frac{2+\gamma^{(1)}}{2\gamma^{(1)}}}\sqrt{\frac{\log T}{T}} = o(1)$ and $M^{\frac{1+\gamma^{(2)}}{2\gamma^{(2)}}}\sqrt{\frac{\log T}{T}} = o(1)$, where $\gamma^{(1)}, \gamma^{(2)} > 0$ are defined in Assumption 5(v)-5(vi);
(v) $\frac{\sqrt{T}}{v_T(f_{p_1})} \times \left(M^{\frac{4-r}{4}} + \tau_M^2 M^{9/4}\frac{\log M}{T}\right) = o(1)$.

Discussion of Assumptions. Assumptions 10(i), 10(iii) and 10(iv) are regularity conditions sufficient to apply the sup-norm rate results in CC. 10(ii) is made for simplicity but it is not necessary. 10(v) corresponds to the second part of Assumption CS(v) in CC and is used to verify Assumption 8. More concrete sufficient conditions for Assumptions 10(iv) and 10(v) may be provided in specific settings. For example, Lemma 7 below gives sufficient conditions for the mildly ill-posed case.

We now provide a proof of Theorem 4.

Proof of Theorem 4. We prove the statement by showing that the assumptions of Theorem 2 hold. Assumptions 5, 6(iii) and 7 are maintained. Assumption 6(i) is implied by Assumption 10(v) , Lemma 10 and the fact that $\zeta = O_p\left(\sqrt{M}\right)$ for Bernstein polynomials.[79] Similarly, Assumption 6(ii) is implied by Assumption 10(iv) and $\zeta = O_p\left(\sqrt{M}\right)$.
We now verify Assumption 8.

8(i) Applying the implicit function theorem to (26),

$$Df_p(h)[v] = -\left[\begin{array}{cc} \frac{\partial g_1(\overline{\delta},\overline{p},\overline{mc},h+\tau v)}{\partial p_1} & \frac{\partial g_1(\overline{\delta},\overline{p},\overline{mc},h+\tau v)}{\partial p_2} \\ \frac{\partial g_2(\overline{\delta},\overline{p},\overline{mc},h+\tau v)}{\partial p_1} & \frac{\partial g_2(\overline{\delta},\overline{p},\overline{mc},h+\tau v)}{\partial p_2} \end{array}\right]^{-1} \left[\begin{array}{c} \frac{\partial g_1(\overline{\delta},\overline{p},\overline{mc},h+\tau v)}{\partial \tau} \\ \frac{\partial g_2(\overline{\delta},\overline{p},\overline{mc},h+\tau v)}{\partial \tau} \end{array}\right]\Bigg|_{\tau=0} \equiv -\left(\mathbb{J}_g^p\right)^{-1}\mathbb{J}_g^\tau\Big|_{\tau=0} \qquad (27)$$

for all $h, v \in \mathcal{H}$. Now, note that $\mathbb{J}_g^p\big|_{\tau=0}$ does not depend on v, and that $\mathbb{J}_g^\tau\big|_{\tau=0}$ is a linear function of $v\left(h^{-1}\left(\overline{\delta},\overline{p}\right),\overline{p}\right)$ and its first derivatives, with coefficients that depend on derivatives of h of order 2 or lower, i.e. we can write

$$Df_{p_1}(h)[v] = \sum_{\tilde{\alpha}:|\tilde{\alpha}|\leq 1}\sum_{j=1}^{2} C_{\tilde{\alpha},j}\left(\overline{\delta},\overline{mc},\left\{\partial^\beta h : |\beta| \leq 2\right\}\right) \times \partial^{\tilde{\alpha}} v_j\left(h^{-1}\left(\overline{\delta},\overline{p}\right),\overline{p}\right) \qquad (28)$$

for real-valued functionals $C_{\tilde{\alpha},j}$. This shows that $Df_p(h_0)[v]$ is linear. Further, by the fundamental theorem of calculus, following an argument analogous to that in part 4(i) of the proof of Theorem 3, we obtain

$$|Df_{p_1}(h_0)[h - h_0]| \leq constant \left\|\frac{\partial^4 h}{\partial s_1 \partial s_2 \partial p_1 \partial p_2} - \frac{\partial^4 h_0}{\partial s_1 \partial s_2 \partial p_1 \partial p_2}\right\|_\infty$$

for all $h \in \mathcal{H}$. Therefore, Assumption 8(i) holds with $\alpha = [1,1,1,1]$.

[79] This follows from the fact that Bernstein polynomials are a special case of B-splines.

8(*ii*) As in part 4(ii) of the proof of Theorem 3, by the mean value theorem, we obtain

$$\left| f_{p_1}\left(\hat{h}\right) - f_{p_1}(h_0) - Df_{p_1}(h_0)[\hat{h} - h_0] \right| \leq$$

$$\sum_{\tilde{\alpha}:|\tilde{\alpha}|\leq 1} \sum_{j=1}^{2} \left[C_{\tilde{\alpha},j}\left(\bar{\delta}, \overline{mc}, \left\{\partial^\beta \hat{h} : |\beta| \leq 2\right\}\right) - C_{\tilde{\alpha},j}\left(\bar{\delta}, \overline{mc}, \left\{\partial^\beta h_0 : |\beta| \leq 2\right\}\right)\right] \times \left|\left|\partial^{\tilde{\alpha}} \hat{h}_j - \partial^{\tilde{\alpha}} h_{0,j}\right|\right|_\infty$$

Since, each of the $C_{\tilde{\alpha},j}\left(\bar{\delta}, \overline{mc}, \left\{\partial^\beta \hat{h} : |\beta| \leq 2\right\}\right) - C_{\tilde{\alpha},j}\left(\bar{\delta}, \overline{mc}, \left\{\partial^\beta h_0 : |\beta| \leq 2\right\}\right)$ terms may be bounded, after some algebra, by a linear combination of $\left\{||\partial^\beta \hat{h} - \partial^\beta h_0||_\infty : |\beta| \leq 2\right\}$, Assumption 8($ii$) holds with $|\alpha_1| = 1, |\alpha_2| = 2$.

8(*iii*) Given the choice of $\alpha, \alpha_1, \alpha_2$ above and by Corollary 3.1 in CC, we have

$$\begin{aligned} \left|\left|\partial^{\alpha_1}\hat{h} - \partial^{\alpha_1} h_0\right|\right|_\infty \left|\left|\partial^{\alpha_2}\hat{h} - \partial^{\alpha_2} h_0\right|\right|_\infty + \left|\left|\partial^{\alpha}\tilde{h} - \partial^{\alpha} h_0\right|\right|_\infty &= O_p\left(M^{\frac{3-2r}{4}} + \tau_M M^{\frac{5-r}{4}}\sqrt{\frac{\log M}{T}} + \tau_M^2 M^{\frac{7}{4}}\frac{\log M}{T}\right) \\ &+ O_p\left(M^{\frac{4-r}{4}}\right) \end{aligned}$$

Thus, Assumption 8(iii) is implied by Assumption 10(v) and Lemma 10.

8(*iv*) By Remark 4.1 in CC, a sufficient condition for Assumption 8(iv) is

$$T_{iv,p} \equiv \frac{\tau_M \sqrt{\sum_{m=1}^M \left(Df_{p_1}\left(\hat{h}\right)\left[\left(G_\psi^{-1/2}\psi_M\right)_m\right] - Df_{p_1}(h_0)\left[\left(G_\psi^{-1/2}\psi_M\right)_m\right]\right)^2}}{v_T(f_{p_1})} = O_p(1)$$

where $\left(G_\psi^{-1/2}\psi_M\right)_m$ denotes the $m-$th row of the matrix $G_\psi^{-1/2}\psi_M$. Note that, after some algebra, we can write $Df_{p_1}\left(\hat{h}\right)\left[\left(G_\psi^{-1/2}\psi_M\right)_m\right] - Df_{p_1}(h_0)\left[\left(G_\psi^{-1/2}\psi_M\right)_m\right]$ for every m as the linear combination of terms, where each term is the difference between a partial derivative of $\hat{h}_i$ of order at most 2 and the same derivative of $h_{0,i}$ for some $i \in \{1, 2\}$, and each coefficient is a partial derivative of an element of $\left(G_\psi^{-1/2}\psi_M\right)_m$ of order at most 1. Therefore, we can write

$$\begin{aligned} T_{iv,p} &\lesssim \frac{\tau_M}{v_T(f_{p_1})}\sqrt{M^{3/2}\zeta^2 \times \left[O_p\left(M^{(2-r)/4} + \tau_M M\sqrt{\log M/T}\right)\right]^2} \\ &= O_p\left(\frac{\sqrt{T}}{v_T(f_{p_1})} \times \left[\frac{\tau_M M}{\sqrt{T}} M^{(3-r)/4} + \frac{\tau_M^2 M^{\frac{9}{4}}\sqrt{\log M}}{T}\right]\right) = O_p(1) \end{aligned}$$

where the first step uses Corollary 3.1 in CC and the proof of Lemma 10, the second step uses the fact that, under the maintained assumptions, $\zeta = O\left(\sqrt{M}\right)$, and the last step follows from Assumption 10(v) (which implies $\frac{\tau_M M}{\sqrt{T}} = o(1)$).

□

Lemma 7. *Let Assumptions 10(i) and 10(iii) hold, and let* $(v_T(f_\epsilon))^2 \asymp \tau_M^2 \sum_{m=1}^M m^a$ *for* $a \leq 0$*. Further, assume that* $\tau_M \asymp M^{\varsigma/2}$ *for* $\varsigma \geq 0, a + \varsigma + 1 > 0$.[80] *Then, Assumptions 10(iv) and 10(v) are satisfied if*

[80] This corresponds to the "mildly ill-posed" case discussed by CC in Corollary 5.1. CC also provide sufficient conditions for the maintained assumption on the rate of divergence of $v_T(f)$.

$M \asymp T^{\rho}$ *with*

$$\rho \in \left(\frac{2}{r - 4 + 2\left(a + \varsigma + 1\right)}, \min\left\{ \frac{2}{2\varsigma - 2a + 7}, \frac{\gamma^{(1)}}{2 + \gamma^{(1)}}, \frac{\gamma^{(2)}}{1 + \gamma^{(2)}} \right\} \right)$$

Further, M may be chosen to satisfy the latter condition if $r + 4a - 9 > 0$ and $\gamma^{(i)}\left(r + 2a + 2\varsigma - 4\right) - 4 > 0$ for $i \in \{1, 2\}$.

Proof. As shown in CC, under the maintained assumptions, $\left(v_T\left(f_{\epsilon}\right)\right)^2 \asymp M^{a+\varsigma+1}$. The result follows by inspection.

□

C.5 Supplementary lemmas and proofs

Lemma 8. *For $i \in \mathcal{J}$, let $\lambda_i\left(T\right) \equiv \frac{\left\|\left|Df(h_0)\left[\psi^{(i)}_{M_i}\right]\left(L_i' G^{-1}_{A,i} L_i\right)^{-1} L_i' G^{-1/2}_{A,i}\right|\right\|}{v_T(f)}$ and let Assumption 5(ii) hold. Then,* $\limsup_{T\to\infty} \lambda_i\left(T\right) < \infty$.

Proof. Note that

$$\begin{aligned}
v_T^2\left(f\right) &= \sum_{i=1}^{J} Df\left(h_0\right)[\psi^{(i)}_{M_i}]' \left(L_i{}' G^{-1}_{A,i} L_i\right)^{-1} L_i{}' G^{-1}_{A,i} \Omega_{ii} G^{-1}_{A,i} L_i \left(L_i{}' G^{-1}_{A,i} L_i\right)^{-1} Df\left(h_0\right)[\psi^{(i)}_{M_i}] \\
&+ 2\sum_{j=1}^{J}\sum_{k=1}^{j-1} Df\left(h_0\right)[\psi^{(j)}_{M_j}]' \left(L_j{}' G^{-1}_{A,j} L_j\right)^{-1} L_j{}' G^{-1}_{A,j} \Omega_{jk} G^{-1}_{A,k} L_k \left(L_k{}' G^{-1}_{A,k} L_k\right)^{-1} Df\left(h_0\right)[\psi^{(k)}_{M_k}] \\
&\equiv \sum_{i=1}^{J} \sigma^2_{T,j} + 2\sum_{j=1}^{J}\sum_{k=1}^{j-1} \sigma_{T,j,k}
\end{aligned}$$

Further, by Assumption 5(ii)

$$\left\|\left| Df\left(h_0\right)\left[\psi^{(i)}_{M_i}\right] \left(L_i' G^{-1}_{A,i} L_i\right)^{-1} L_i' G^{-1/2}_{A,i} \right|\right\|^2 \le \underline{\sigma}^{-2} \sigma^2_{T,i}$$

for $i \in \mathcal{J}$. Therefore, we can write

$$\left[\lambda_i\left(T\right)\right]^2 \quad \le \quad \frac{\underline{\sigma}^{-2}\sigma^2_{T,i}}{\sum_{i=1}^{J} \sigma^2_{T,j} + 2\sum_{j=1}^{J}\sum_{k=1}^{j-1} \sigma_{T,j,k}}$$

from which the result follows.

□

Lemma 9. *Let Assumptions 5(iii), 5(vi) and 6(ii) hold. Then $||T_{\Omega,1}|| = O_p\left(1\right)$.*

Proof. The proof follows that of Lemma 3.1 in Chen and Christensen (2015) with minor changes. Let

$C_T \asymp \zeta^{\left(1+\gamma^{(2)}\right)/\gamma^{(2)}}$ be a sequence of positive numbers with $\gamma^{(2)}$ defined in Assumption 5(vi), and let

$$T_{\Omega,1}^{(1)} \equiv \frac{1}{T}\sum_{t=1}^{T}\left(\Xi_{1,t} - \mathbb{E}\left[\Xi_1\right]\right) \qquad\qquad T_{\Omega,1}^{(2)} \equiv \frac{1}{T}\sum_{t=1}^{T}\left(\Xi_{2,t} - \mathbb{E}\left[\Xi_2\right]\right)$$

where

$$\begin{aligned}
\Xi_{1,t} &\equiv \xi_{j,t}\xi_{k,t}G_{A,j}^{-1/2}a_{K_j}^{(k)}\left(w_t\right)a_{K_j}^{(j)}\left(w_t\right)'G_{A,j}^{-1/2}\mathbb{I}\left\{||\xi_{j,t}\xi_{k,t}G_{A,j}^{-1/2}a_{K_j}^{(j)}\left(w_t\right)a_{K_j}^{(j)}\left(w_t\right)'G_{A,j}^{-1/2}|| \le C_T^2\right\}\\
\Xi_{2,t} &\equiv \xi_{j,t}\xi_{k,t}G_{A,j}^{-1/2}a_{K_j}^{(j)}\left(w_t\right)a_{K_j}^{(j)}\left(w_t\right)'G_{A,j}^{-1/2}\mathbb{I}\left\{||\xi_{j,t}\xi_{k,t}G_{A,j}^{-1/2}a_{K_j}^{(j)}\left(w_t\right)a_{K_j}^{(j)}\left(w_t\right)'G_{A,j}^{-1/2}|| > C_T^2\right\}
\end{aligned}$$

Note that $T_{\Omega,1} = T_{\Omega,1}^{(1)} + T_{\Omega,1}^{(2)}$, so that $||T_{\Omega,1}^{(1)}|| = O_p\left(1\right)$ and $||T_{\Omega,1}^{(2)}|| = O_p\left(1\right)$ imply the statement of the lemma.
Control of $||T_{\Omega,1}^{(1)}||$: By definition, $||\Xi_{1,t}|| \le C_T^2$ and thus, by the triangle inequality and Jensen's inequality ($||\cdot||$ is convex), we have $||\Xi_{1,t} - \mathbb{E}\left(\Xi_1\right)|| \le 2C_T^2$. Further,

$$\begin{aligned}
\mathbb{E}\left[\Xi_1 - \mathbb{E}\left(\Xi_1\right)\right]^2 &\le\\
\mathbb{E}\left[\xi_j^2\xi_k^2||G_{A,j}^{-1/2}a_{K_j}^{(j)}||^2G_{A,j}^{-1/2}a_{K_j}^{(j)}\left(W\right)a_{K_j}^{(j)}\left(W\right)'G_{A,j}^{-1/2}\mathbb{I}\left\{||\xi_j\xi_kG_{A,j}^{-1/2}a_{K_j}^{(j)}\left(W\right)a_{K_j}^{(j)}\left(W\right)'G_{A,j}^{-1/2}|| \le C_T^2\right\}\right] &\le\\
C_T^2\mathbb{E}\left[|\xi_j\xi_k|G_{A,j}^{-1/2}a_{K_j}^{(j)}\left(W\right)a_{K_j}^{(j)}\left(W\right)'G_{A,j}^{-1/2}\mathbb{I}\left\{||\xi_j\xi_kG_{A,j}^{-1/2}a_{K_j}^{(j)}\left(W\right)a_{K_j}^{(j)}\left(W\right)'G_{A,j}^{-1/2}|| \le C_T^2\right\}\right] &\le\\
C_T^2\mathbb{E}\left[\mathbb{E}\left(|\xi_j\xi_k||W\right)G_{A,j}^{-1/2}a_{K_j}^{(j)}\left(W\right)a_{K_j}^{(j)}\left(W\right)'G_{A,j}^{-1/2}\right] &\lesssim\\
C_T^2\mathbb{E}\left[G_{A,j}^{-1/2}a_{K_j}^{(j)}\left(W\right)a_{K_j}^{(j)}\left(W\right)'G_{A,j}^{-1/2}\right] = C_T^2I_{K_j}&
\end{aligned}$$

where the inequalities are in the sense of positive semi-definite matrices. The first inequality follows from $\mathbb{E}\left[\Xi_1 - \mathbb{E}\left(\Xi_1\right)\right]^2 \le \mathbb{E}\left(\Xi_1\right)^2$ (and the fact that matrix multiplication is commutative), the second follows from the fact that $||G_{A,j}^{-1/2}a_{K_j}^{(j)}\left(W\right)a_{K_j}^{(j)}\left(W\right)'G_{A,j}^{-1/2}|| = ||G_{A,j}^{-1/2}a_{K_j}^{(j)}\left(W\right)||^2$, the third from the law of iterated expectations, the fourth follows from Assumption 5(iii), and the last step is because $G_{A,j}^{-1/2}a_{K_j}^{(j)}$ is an orthonormal basis. Then, Corollary 4.1 in Chen and Christensen (2015) yields $||T_{\Omega,1}^{(1)}|| = O_p\left(C_T\sqrt{(\log K)/T}\right) = O_p\left(1\right)$, where the last step uses Assumption 6(ii).
Control of $||T_{\Omega,1}^{(2)}||$: Since $||\Xi_{2,t}|| \le \zeta^2|\xi_{j,t}\xi_{k,t}|\mathbb{I}\left\{|\xi_{j,t}\xi_{k,t}| \ge C_T^2/\zeta^2\right\}$,by the triangle inequality and Jensen's inequality ($||\cdot||$ is convex), we have

$$\mathbb{E}\left[||T_{\Omega,1}^{(2)}||\right] \le 2\zeta^2\mathbb{E}\left[|\xi_j\xi_k|\mathbb{I}\left\{|\xi_j\xi_k| \ge C_T^2/\zeta^2\right\}\right] \le \frac{2\zeta^{2(1+\gamma^{(2)})}}{C_T^{2\gamma^{(2)}}}\mathbb{E}\left[|\xi_j\xi_k|^{1+\gamma^{(2)}}\mathbb{I}\left\{|\xi_j\xi_k| \ge C_T^2/\zeta^2\right\}\right] = o\left(1\right)$$

where the last step follows from Assumption 5(vi), the fact that $C_T^2/\zeta^2 \asymp \zeta^{2/\gamma^{(2)}} \to \infty$ and that $\zeta^{(1+\gamma^{(2)})}/C_T^{\gamma^{(2)}} \asymp 1$. Thus, $||T_{\Omega,1}^{(2)}|| = O_p\left(1\right)$ by Markov's inequality.

□

Lemma 10. *Let Assumption 9(i)-9(iii) hold. Then, for* $f \in \{f_\epsilon, f_{p_1}\}$

$$\left[v_T\left(f\right)\right]^2 \lesssim \tau_M^2M^{5/2}$$

Proof. We prove this for $f = f_\epsilon$. The proof for $f = f_{p_1}$ is identical. As shown in CC,[81] the maintained

[81] See also Chen and Pouzo (2015).

assumption imply $[v_T(f_\epsilon)]^2 \lesssim \tau_M^2 \sum_{m=1}^{M} \left(Df_\epsilon(h_0)\left[\left(G_\psi^{-1/2}\psi_M\right)_m\right]\right)^2$, where $\left(G_\psi^{-1/2}\psi_M\right)_m$ denotes the $m-$th row of the $M-$by$-2-$valued function $G_\psi^{-1/2}\psi_M$. Next, note that

$$\left|Df_\epsilon(h_0)\left[\left(G_\psi^{-1/2}\psi_M\right)_m\right]\right| \lesssim \max_{\tilde{\alpha}:|\tilde{\alpha}|=1} ||\partial^{\tilde{\alpha}}\left(G_\psi^{-1/2}\psi_M\right)_m||_\infty \leq 4M^{1/4} \sup_{(y)\in\mathcal{S}\times\mathcal{P}} ||G_\psi^{-1/2}\psi_M(y)||_\infty \leq 4M^{1/4}\zeta$$

where we use the fact that the first derivatives of Bernstein polynomials can be written in terms of lower order polynomials, and that the latter can in turn be written as a linear combination of higher order polynomials.Finally, using $\zeta = O\left(\sqrt{M}\right)$, we obtain the result.

□

Appendix D: Microfoundation of the Empirical Model

In this appendix, we show how to map the model used in the empirical application to the general BH framework outlined in Section 2. Specifically, we outline two models of consumer behavior that yield the demand system in equation (5) and prove that the system is indeed invertible. It should be emphasized that these are only two out of many models that are compatible with (5) and invertibility, and that the estimation procedure does not rely on any of the parametric restrictions embedded in either model.[82]

D.1 Model 1

We first consider a standard discrete choice model. While the model is clearly at odds with the fact that consumers buying fresh fruit face an (at least partially) continuous choice, this serves as a building block for the more realistic model discussed in Section D.2. Moreover, given the prevalence of discrete choice models in the literature, it provides a connection between the demand system in (5) and a more familiar setup.

We assume that consumers face a discrete choice between one unit (say, one pound) of non-organic strawberries, one unit of organic strawberries and one unit of other fresh fruit. Consumer i's indirect utilities for each of these goods are, respectively

$$\begin{aligned} u_{i1} &= \theta_s \delta_{strawb} + \alpha_i p_1 + \epsilon_{i1} \\ u_{i2} &= \theta_s \delta_{strawb} + \theta_{org}\delta_{org} + \alpha_i p_2 + \epsilon_{i2} \\ u_{i0} &= \theta_{0,org}\delta_{org} + \theta_{0,out} x_{out} + \alpha_i p_0 + \epsilon_{i0} \end{aligned} \tag{29}$$

where

$$\begin{aligned} \delta_{strawb} &= \xi_{strawb} \\ \delta_{org} &= \theta_{lett} x_{lett} + \xi_{org} \end{aligned}$$

[82]For instance, while Model 1 below assumes that prices enter linearly in utilities, this restriction is not needed for identification or estimation, given that we do not impose symmetry of the Jacobian of demand with respect to price.

and p_1, p_2, p_0 denote the prices of non-organic strawberries, organic strawberries and the price index for other fresh fruit, respectively. We interpret δ_{strawb} as the mean quality of all strawberries in the market and δ_{org} as the mean utility for organic products (including - but not limited to - organic strawberries). Because the outside option of buying other fresh fruit includes organic produce (e.g. organic apples), we let δ_{org} enter u_{i0}. We use $(\xi_{strawb}, \xi_{org})$ to denote the unobserved components of the mean utilities and $(\epsilon_{i2}, \epsilon_{i2})$ to denote taste shocks idiosyncratic to consumer i. Unlike BLP, we will not make any parametric assumptions on $(\epsilon_{i2}, \epsilon_{i2})$, nor on the distribution of the price coefficient α_i. In particular, note that the correlation structure of the vector $(\epsilon_{i2}, \epsilon_{i2}, \alpha_i)$ is unrestricted, which allows for patterns such as the fact that wealthier consumers may have a stronger preference for organic produce. However, if we impose exchangeability of the demand system in estimation, we need to assume that the distributions of ϵ_{i1} and ϵ_{i2} are the same in order to ensure that $\sigma_1 = \sigma_2$ in (5), as will be clear from the derivations below. Although not made explicit in the notation, the distribution of α_i is allowed to depend on other covariates such as mean income y^{inc} in the market.

Now we show that the demand system generated by the model above can be inverted for the mean utilities $(\delta_{strawb}, \delta_{org})$ under the following very mild assumption (as well as the standard exogeneity and completeness assumptions discussed in Section 2).

Assumption 11. *The coefficients $\theta_s, \theta_{org}, \theta_{0,org}, \theta_{0,out}$ and θ_{lett} are non-zero.*

Note that Assumption 11 is very mild. It is satisfied if (i) consumers care about the quality of strawberries $(\theta_s > 0)$ and organic produce $(\theta_{org}, \theta_{0,org} > 0)$, as well as the availability of non-strawberry fruit $\theta_{0,out} > 0$, when purchasing fresh fruit; and (ii) the variable x_{lett} is indeed a proxy for taste for organic produce $(\theta_{lett} > 0)$.

Lemma 11. *Under Assumptions 3, 4 and 11, the functions σ_1 and σ_2 in (5) are point-identified.*

Proof. Since utility is ordinal, we can subtract $\theta_{0,org}\delta_{org} + \theta_{0,out}x_{out} + \alpha_i p_0$ from each equation in (29) and write

$$\begin{aligned} u_{i1} &= \tilde{\delta}_1 - \theta_{0,out}x_{out} + \alpha_i (p_1 - p_0) + \epsilon_{i1} \\ u_{i2} &= \tilde{\delta}_2 - \theta_{0,out}x_{out} + \alpha_i (p_2 - p_0) + \epsilon_{i2} \\ u_{i0} &= \epsilon_{i0}, \end{aligned} \tag{30}$$

where

$$\begin{aligned} \tilde{\delta}_1 &\equiv \theta_s \delta_{strawb} - \theta_{0,org}\delta_{org} \\ \tilde{\delta}_2 &\equiv \theta_s \delta_{strawb} + (\theta_{org} - \theta_{0,org})\,\delta_{org} \end{aligned}$$

Using (30), we can write the demand system as

$$s = \tilde{\sigma}\left(\tilde{\delta}_1 - \theta_{0,out}x_{out}, \tilde{\delta}_2 - \theta_{0,out}x_{out}, p, y^{inc}\right), \tag{31}$$

where $p \equiv (p_0, p_1, p_2)$, $s \equiv (s_1, s_2)'$ is the vector of market shares and $\tilde{\sigma}$ is a function from $\mathbb{R}^2 \times \mathbb{R}^4_+$ to the unit 2-simplex. Next, by Theorem 1 of Berry, Gandhi, and Haile (2013), we can invert the system in (31) for the mean utility levels as follows

$$\begin{aligned} \tilde{\delta}_1 &= \tilde{\sigma}_1^{-1}\left(s, p, y^{inc}\right) + \theta_{0,out}x_{out} \\ \tilde{\delta}_2 &= \tilde{\sigma}_2^{-1}\left(s, p, y^{inc}\right) + \theta_{0,out}x_{out}, \end{aligned} \tag{32}$$

where $\tilde{\sigma}_k^{-1}$ denotes the $k-$th element of the inverse, $\tilde{\sigma}^{-1}$, of $\tilde{\sigma}$. We now show that there is a one-to-one mapping between $(\delta_{strawb}, \delta_{org})$ and $\left(\tilde{\delta}_1, \tilde{\delta}_2\right)$, which means that we can invert the system for the original demand indices. Letting $\delta \equiv (\delta_{strawb}, \delta_{org})'$ and $\tilde{\delta} \equiv \left(\tilde{\delta}_1, \tilde{\delta}_2\right)'$, we have

$$\tilde{\delta} = A\delta,$$

where

$$A \equiv \begin{bmatrix} \theta_s & -\theta_{0,org} \\ \theta_s & \theta_{org} - \theta_{0,org} \end{bmatrix}$$

Since $\det(A) = \theta_s \theta_{org} \neq 0$ under Assumption 11, we can rewrite (32) as

$$\delta = A^{-1}\tilde{\sigma}^{-1}\left(s, p, y^{inc}\right) + A^{-1} \cdot [1 \quad 1]' \times \theta_{0,out} x_{out} \tag{33}$$

or, equivalently,

$$\begin{aligned} \delta_{strawb} &= \sigma_1^{-1}\left(s, p, y^{inc}\right) + \theta_1 x_{out} \\ \delta_{org} &= \sigma_2^{-1}\left(s, p, y^{inc}\right) + \theta_2 x_{out}, \end{aligned} \tag{34}$$

for functions $\sigma_i^{-1} : \Delta^2 \times \mathbb{R}_+^4 \to \mathbb{R}^2$, $i = 1, 2$, where Δ^2 denotes the unit 2-simplex. Now we derive expressions for the coefficients θ_1 and θ_2 in terms of the model primitives. Note that

$$A^{-1} = \frac{1}{\theta_{org}} \begin{bmatrix} \frac{\theta_{org} - \theta_{0,org}}{\theta_s} & \frac{\theta_{0,org}}{\theta_s} \\ -1 & 1 \end{bmatrix}$$

and thus

$$A^{-1} \cdot [1 \quad 1]' = \begin{bmatrix} \frac{1}{\theta_s} & 0 \end{bmatrix}',$$

i.e.

$$\begin{aligned} \theta_1 &= \frac{\theta_{0,out}}{\theta_s} \\ \theta_2 &= 0 \end{aligned}$$

Plugging this into (34) and using the definitions on δ_{strawb} and δ_{org}, we obtain

$$\begin{aligned} \xi_{strawb} &= \sigma_1^{-1}\left(s, p, y^{inc}\right) + \frac{\theta_{0,out}}{\theta_s} x_{out} \\ \theta_{lett} x_{lett} + \xi_{org} &= \sigma_2^{-1}\left(s, p, y^{inc}\right) \end{aligned} \tag{35}$$

The final step is to show that we can identify the system in (35), given the instruments available. Because we are free to normalize the scale of ξ_{strawb} and ξ_{org} in the display above, we can divide the first equation of (35) by $\frac{\theta_{0,out}}{\theta_s}$ and the second equation by (35) by θ_{lett} without loss,[83] and rearrange terms as follows

$$-x_{out} = \sigma_1^{-1}\left(s, p, y^{inc}\right) - \xi_{strawb} \tag{36}$$

$$x_{lett} = \sigma_2^{-1}\left(s, p, y^{inc}\right) - \xi_{org}, \tag{37}$$

[83]These divisions are well-defined operations as $\frac{\theta_{0,out}}{\theta_s}$ and θ_{lett} are nonzero by Assumption 11.

Equations (36) and (37) are in the same form as Equation (6) in BH and thus we can follow their argument to show that σ_1 and σ_2 are identified.

$\square$

D.2 Model 2

We now turn to a model of continuous choice that appears to be a closer approximation to the behavior of consumers buying fresh fruit. Let consumer i face the following maximization problem

$$\begin{aligned} &\max_{q_0,q_1,q_2} U_i\left(q_0,q_1,q_2\right) \\ \text{s.t.}\quad & p_0q_0+p_1q_1+p_2q_2 \leq y_i^{inc} \end{aligned} \tag{38}$$

where y_i^{inc} denotes the income consumer i allocates to fresh fruit, q_0 is the quantity of non-strawberry fresh fruit, q_1 is the quantity of non-organic strawberries and q_2 is the quantity of organic strawberries, and similarly for prices p_0, p_1, p_2. One could think of y_i^{inc} as being the outcome of a higher-level optimization problem in which the consumer chooses how to allocate total income across different product categories, including fresh fruit. Assume U_i takes the Cobb-Douglas form

$$U_i\left(q_0,q_1,q_2\right) = q_0^{\delta_0\epsilon_{i,0}} q_1^{\delta_1\epsilon_{i,1}} q_2^{\delta_2\epsilon_{i,2}},$$

for positive $\delta \equiv (\delta_0, \delta_1, \delta_2)$ and $\epsilon_i \equiv (\epsilon_{i,0}, \epsilon_{i,1}, \epsilon_{i,2})$. Then, the optimal quantities chosen by the consumer are

$$q_j^*\left(\delta, p, y_i^{inc}, \epsilon_i\right) = \frac{y_i^{inc}}{p_j} \cdot \frac{\delta_j\epsilon_{i,j}}{\sum_{k=0}^{2}\delta_k\epsilon_{i,k}} \qquad j=0,1,2 \tag{39}$$

where $\delta \equiv (\delta_0, \delta_1, \delta_2)$ and $p \equiv (p_0, p_1, p_2)$. Now assume that

$$\begin{aligned} \delta_0 &= \gamma_{org}^{\theta_{0,org}} \tilde{x}_{out}^{\theta_{0,out}} \\ \delta_1 &= \gamma_{strawb}^{\theta_s} \\ \delta_2 &= \gamma_{strawb}^{\theta_s}\gamma_{org}^{\theta_{org}} \end{aligned}$$

where

$$\begin{aligned} \gamma_{strawb} &\equiv \exp\left\{\delta_{strawb}\right\} \\ \gamma_{org} &\equiv \exp\left\{\delta_{org}\right\} \\ \tilde{x}_{out} &\equiv \exp\left\{x_{out}\right\} \end{aligned}$$

We can then re-write (39) as

$$q_j^*\left(\tilde{\delta}, p, y_i^{inc}, \epsilon_i\right) = \frac{y_i^{inc}}{p_j} \cdot \frac{\tilde{\delta}_j\epsilon_{i,j}}{\sum_{k=0}^{2}\tilde{\delta}_k\epsilon_{i,k}} \qquad j=0,1,2 \tag{40}$$

where

$$\begin{aligned}
\tilde{\delta}_0 &\equiv 1 \\
\tilde{\delta}_1 &\equiv \gamma_{strawb}^{\theta_s} \gamma_{org}^{-\theta_{0,org}} \tilde{x}_{out}^{-\theta_{0,out}} \\
\tilde{\delta}_2 &\equiv \gamma_{strawb}^{\theta_s} \gamma_{org}^{\theta_{org}-\theta_{0,org}} \tilde{x}_{out}^{-\theta_{0,out}}
\end{aligned}$$

and $\tilde{\delta} \equiv \left(\tilde{\delta}_0, \tilde{\delta}_1, \tilde{\delta}_2\right)$.

Next, let $F_{Y,\epsilon}$ denote the joint distribution of income and ϵ in the market, and define[84]

$$Q_j^*\left(\tilde{\delta}, p\right) = \int q_j^*\left(\tilde{\delta}, p, y, \epsilon\right) dF_{Y,\epsilon}(y, \epsilon) \qquad j = 0, 1, 2$$

$Q_j^*\left(\tilde{\delta}, p\right)$ is the model counterpart to the market-level quantity Q_j observed in the data.

The last step is to show that there exists a mapping of quantities into (artificial) market shares such that the resulting demand system is invertible. For $j = 0, 1, 2$, define

$$\tilde{\tilde{\sigma}}_j\left(\tilde{\delta}, p\right) = \frac{Q_j^*\left(\tilde{\delta}, p\right)}{\sum_{k=0}^{2} Q_k^*\left(\tilde{\delta}, p\right)}$$

and

$$s_j = \frac{Q_j}{\sum_{k=0}^{2} Q_k}$$

Then, equating observed shares to their model counterparts, we obtain the system

$$s = \tilde{\tilde{\sigma}}\left(\tilde{\delta}, p\right) \tag{41}$$

where $s \equiv (s_0, s_1, s_2)'$ and $\tilde{\tilde{\sigma}}\left(\tilde{\delta}, p\right) \equiv \left(\tilde{\tilde{\sigma}}_0\left(\tilde{\delta}, p\right), \tilde{\tilde{\sigma}}_1\left(\tilde{\delta}, p\right), \tilde{\tilde{\sigma}}_2\left(\tilde{\delta}, p\right)\right)'$.

Because $\tilde{\tilde{\sigma}}_j$ is strictly decreasing in $\tilde{\delta}_k$ for all j and all $k > 0, k \neq j$, by Theorem 1 in Berry, Gandhi, and Haile (2013), we can invert (41) as follows

$$\tilde{\delta} = \tilde{\tilde{\sigma}}^{-1}(s, p)$$

and, taking logs, we can write

$$\begin{aligned}
\theta_s \delta_{strawb} - \theta_{0,org} \delta_{org} &= \tilde{\sigma_1}^{-1}(s, p) + \theta_{0,out} x_{out} \\
\theta_s \delta_{strawb} + (\theta_{org} - \theta_{0,org}) \delta_{org} &= \tilde{\sigma}_2^{-1}(s, p) + \theta_{0,out} x_{out}
\end{aligned} \tag{42}$$

where $\tilde{\sigma}_j^{-1}(s, p) \equiv \log\left(\tilde{\tilde{\sigma}}_j^{-1}(s, p)\right)$ for $j = 1, 2$.

Note that (42) has the exact same form as (32). Therefore, we can use the argument in Section D.1 to show that the demand system is identified.

[84] Note that the function Q_j also depends on $F_{Y,\epsilon}$, but the dependence is suppressed for notational convenience.

References

ABALUCK, J., AND A. ADAMS (2017): “What Do Consumers Consider Before They Choose? Identification from Asymmetric Demand Responses,” IFS Working Paper W17/09.

ATKIN, D., AND D. DONALDSON (2015): “Whos Getting Globalized? The Size and Implications of Intra-national Trade Costs,” NBER Working Paper No. 21439.

BERRY, S., A. GANDHI, AND P. HAILE (2013): “Connected substitutes and invertibility of demand,” *Econometrica*, 81, 2087–2111.

BERRY, S., AND P. HAILE (2010): “Nonparametric Identification of Multinomial Choice Demand Models with Heterogeneous Consumers,” Working Paper (Yale University).

——— (2014): “Identification in differentiated products markets using market level data,” *Econometrica*, 82, 1749–1797.

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63, 841–890.

——— (2004): “Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market,” *Journal of Political Economy*, 112, 68–105.

BULOW, J. I., J. D. GEANAKOPLOS, AND P. D. KLEMPERER (1985): “Multimarket Oligopoly: Strategic Substitutes and Complements,” *The Journal of Political Economy*, 93, 488–511.

BULOW, J. I., AND P. PFLEIDERER (1983): “A Note on the Effect of Cost Changes on Prices,” *The Journal of Political Economy*, 91, 182185.

CAPPS, C., D. DRANOVE, AND M. SATTERTHWAITE (2003): “Competition and Market Power in Option Demand Markets,” *RAND Journal of Economics*, 34, 737–763.

CHEN, X., AND T. CHRISTENSEN (2015): “Optimal uniform convergence rates and asymptotic normality for series estimators under weak dependence and weak conditions,” *Journal of Econometrics*, 188, 447–465.

——— (forthcoming): “Optimal Sup-norm Rates and Uniform Inference on Nonlinear Functionals of Nonparametric IV Regression,” *Quantitative Economics*.

CHEN, X., AND D. POUZO (2009): “Efficient estimation of semiparametric conditional moment models with possibly nonsmooth residuals,” *Journal of Econometrics*, 152, 46–60.

——— (2012): "Estimation of nonparametric conditional moment models with possibly nonsmooth generalized residuals," *Econometrica*, 80, 277321.

——— (2015): "Sieve Wald and QLR Inferences on Semi/Nonparametric Conditional Moment Models," *Econometrica*, 83, 1013–1079.

Chen, X., and Y. J. J. Qiu (2016): "Methods for Nonparametric and Semiparametric Regressions with Endogeneity: A Gentle Guide," *Annual Review of Economics*, 8, 259–290.

Chernozhukov, V., and C. Hansen (2006): "Instrumental quantile regression inference for structural and treatment effect models," *Journal of Econometrics*, 132, 491–525.

Chevalier, J. A., A. K. Kashyap, and P. E. Rossi (2003): "Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data," *American Economic Review*, 93, 15–37.

Dubé, J.-P. (2004): "Multiple Discreteness and Product Differentiation: Demand for Carbonated Soft Drinks," *Marketing Science*, 23, 66–81.

Dubin, J. A., and D. L. McFadden (1984): "An Econometric Analysis of Residential Electric Appliance Holdings and Consumptions," *Econometrica*, 52, 345–362.

Gal, S. (2008): *Shape-Preserving Approximation by Real and Complex Polynomials*. Birkhauser, Boston.

Gentzkow, M. (2007): "Valuing New Goods in a Model with Complementarity: Online Newspapers," *American Economic Review*, 97, 713.

Goeree, M. S. (2008): "Limited Information and advertising in the U.S. personal computer industry," *Econometrica*, 76, 1017–1074.

Goldberg, P. K. (1995): "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry," *Econometrica*, 63, 891–951.

Goldberg, P. K., and R. Hellerstein (2008): "A Structural Approach to Explaining Incomplete Exchange-Rate Pass-Through and Pricing-to-Market," *American Economic Rewview: Papers and Proceedings*, 98, 423–429.

——— (2013): "A Structural Approach to Identifying the Sources of Local Currency Price Stability," *Review of Economic Studies*, 80, 175–210.

Gordon, B., and W. Hartmann (2016): "Advertising competition in presidential elections," *Quantitative Marketing and Economics*, 14, 1–40.

Hendel, I. (1999): "Estimating Multiple-Discrete Choice Models: An Application to Computerization Returns," *The Review of Economic Studies*, 66, 423–446.

Horowitz, J. (2011): "Applied nonparametric instrumental variable estimation," *Econometrica*, 79, 347–394.

Johnson, C. R., and R. L. Smith (2011): "Inverse M-matrices, II," *Linear Algebra and its Applications*, 435, 953983.

Knee, M. (2002): *Fruit Quality and Its Biological Basis*. CRC Press.

Lal, R., and C. Matutes (1994): "Retail Pricing and Advertising Strategies," *Journal of Business*, 67, 345–370.

Lal, R., and J. M. Villas-Boas (1998): "Price Promotions and Trade Deals with Multiproduct Retailers," *Management Science*, 44, 935–949.

McDonald, J., M. Neumann, H. Schneider, and M. Tsatsomeros (1995): "Inverse M-Matrix Inequalities and Generalized Ultrametric Matrices," *Linear Algebra and Its Applications*, 22, 321–341.

Nakamura, E., and D. Zerom (2010): "Accounting for Incomplete Pass-Through," *Review of Economic Studies*, 77, 1192–1230.

Nevo, A. (2000): "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry," *The RAND Journal of Economics*, 31, 395–421.

——— (2001): "Measuring market power in the ready-to-eat cereal industry," *Econometrica*, 69, 307–642.

Petrin, A. (2002): "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 110(4), 705–729.

Plemmons, R. (1977): "M-matrix characterizations," *Linear Algebra and its Applications*, 18, 175–188.

Produce for Better Health Foundation (2015): *State of the plate: 2015 study on America's consumption of fruit and vegetables*. Produce for Better Health Foundation, Web http://www.PBHFoundation.org.

Schumaker, L. (2007): *Spline Functions: Basic Theory*. Cambridge Mathematical Library, Cambridge.

Seber, G. A. (2007): *A Matrix Handbook for Statisticians*. Wiley, New York.

Souza-Rodrigues, E. A. (2014): "Nonparametric Estimation of a Generalized Regression Model with Group Effects," working paper, University of Toronto, Department of Economics.

Tversky, A., and D. Kahneman (1991): "Loss Aversion in Riskless Choice: A Reference-Dependent Model," *The Quarterly Journal of Economics*, 106, 1039–1061.

United States Department of Agriculture (2017): *Noncitrus Fruits and Nuts 2016 Summary*.

Weyl, E. G., and M. Fabinger (2013): "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition," *Journal of Political Economy*, 121, 528–583.