

Market Counterfactuals and the Specification of Multi-Product Demand: A Nonparametric Approach*

Giovanni Compiani[†]

October 16, 2021

Abstract

Demand estimates are essential for addressing a wide range of positive and normative questions in economics that are known to depend on the shape—and notably the curvature—of the true demand functions. The existing frontier approaches, while allowing flexible substitution patterns, typically require the researcher to commit to a parametric specification. An open question is whether these a priori restrictions are likely to significantly affect the results. To address this, I develop a nonparametric approach to estimation of demand for differentiated products, which I then apply to California supermarket data. While the approach subsumes workhorse models such as mixed logit, it allows consumer behaviors and preferences beyond standard discrete choice, including continuous choices, complementarities across goods, and consumer inattention. When considering a tax on one good, the nonparametric approach predicts a much lower pass-through than a standard mixed logit model. However, when assessing the market power of a multi-product firm relative to that of a single-product firm, the models give similar results. I also illustrate how the nonparametric approach may be used to guide the choice among parametric specifications.

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

JEL codes: L1, L66

*I am very grateful to Steve Berry, Phil Haile, Penny Goldberg and Xiaohong Chen for their guidance and support, as well as Daniel Akerberg, Nikhil Agarwal, Matteo Benetton, Judith Chevalier, Tim Christensen, Jan De Loecker, Ying Fan, Sharat Ganapati, Amit Gandhi, Mitsuru Igami, Adam Kapor, Yuichi Kitamura, Lorenzo Magnolfi, Camilla Roncoroni, Marcelo Sant’Anna, Paul Scott, three anonymous referees, and participants of EARIE 2017, ASSA 2019, and seminars at MIT, UCL, EIEF, NYU, Stanford, LSE, Mannheim, Tilburg, Berkeley, UCLA, Chicago, Irvine for helpful comments and conversations. Mengsi Gao provided excellent research assistance. 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. This paper was previously circulated under the title “Nonparametric Demand Estimation in Differentiated Products Markets.” All errors and omissions are my own.

[†]University of Chicago. Email: Giovanni.Compiani@chicagobooth.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, evaluating the effect of a tax or subsidy, merger simulation, assessing the impact of a new product being introduced into the market, understanding the drivers of the well-documented incomplete pass-through of cost and exchange-rate shocks to downstream prices, and determining whether firms play a game with strategic complements or substitutes.¹ 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 flexible 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. Addressing this concern requires an approach that relaxes the parametric assumptions, thus providing results that may be used as a benchmark.

To this end, I propose a nonparametric approach to estimate demand in differentiated products markets based on aggregate data.² Specifically, I focus on markets in which consumers face a range of options that are differentiated in ways that are both observed and unobserved to the researcher. Importantly, the presence of unobserved heterogeneity at the product or market level implies that all the variables that are chosen by firms after observing consumer preferences—e.g. prices in many models—are endogenous. A vast literature in industrial organization and other fields has focused on the empirical analysis of this type of markets. The current frontier approach is to posit a random coefficients discrete choice logit model³ and estimate it using the methodology developed by Berry et al. (1995) (henceforth BLP).⁴ While the methodology in BLP accomplishes the crucial goals of generating

¹References include Berry et al. (1995) and Nevo (2001) for the study of market power, Bulow and Pfleiderer (1983) and Weyl and Fabinger (2013) for the effect of taxes and subsidies, Nevo (2000) and Capps et al. (2003) for merger simulation, Petrin (2002) for the analysis of new products, Nakamura and Zerom (2010) and Goldberg and Hellerstein (2013) for incomplete pass-through, and Bulow et al. (1985) for strategic complementarity and substitution.

²Souza-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. Pinkse and Slade (2004) estimate a semiparametric demand system in which price enters linearly and the price effects are flexible functions of the distance between products in characteristics space.

³Throughout the paper, I use the terms “random coefficients logit model” and “mixed logit model” interchangeably.

⁴Another influential approach to demand estimation is the Almost Ideal Demand System (AIDS) pioneered by Deaton and Muellbauer (1980). I choose to compare my approach to BLP-type models and not AIDS-type models, because the latter restrict the role of the unobserved heterogeneity in a way that is at odds with the differentiated products markets literature from the last twenty years. Specifically, Deaton and Muellbauer (1980) use their model of consumer behavior to obtain a demand equation only involving

reasonable substitution patterns and allowing for price endogeneity, it relies on a number of parametric assumptions which may affect the results of counterfactual exercises. For example, while it is well known that the pass-through of a tax depends on the curvature—i.e. the second derivatives—of the demand functions, it is not a priori clear whether BLP is flexible enough to capture these features of the true demand system. In contrast, the approach proposed in this paper does not rely on any distributional assumptions and imposes only limited functional form restrictions. For instance, it is not necessary to assume that the idiosyncratic taste shocks or the random coefficients on product characteristics in the utility function belong to a parametric family of distributions. Instead, I leverage a range of constraints—such as monotonicity of demand in certain variables and properties of the derivatives of demand—that are grounded in consumer theory.

In addition, by directly targeting the demand functions as opposed to the 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,⁵ that consumers are aware of all products and their characteristics,^{6,7} and that each consumer buys at most one unit of a single product.⁸ In contrast to this, my approach allows for a broader range

observables and add an additively separable error term to carry out estimation (equation (15) in their paper). This implies that the unobservables do not have an immediate “structural” interpretation (such as product quality not captured by the data). One consequence is that the standard arguments used to motivate the issue of (price) endogeneity, as well as justify the instrumental variables solution to it, do not typically apply in the AIDS framework.

⁵Gentzkow (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 I view this as a promising avenue for future research.

⁶Goeree (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 assumptions in Section 2 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.

⁷One could conceivably use a BLP-type model to estimate consumer preferences on data generated by inattentive consumers. Whether the BLP functional form is flexible enough in such contexts is an open question that depends on the object of interest. The simulation evidence presented in this paper suggests that a BLP-type model tends to under-estimate own-price elasticities and over-estimate cross-price elasticities for one pattern of inattention.

⁸A 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,

of consumer behaviors and preferences, including complementarities across goods, consumer inattention, and multiple discrete or continuous choices. On the other hand, directly targeting the demand functions means that the distribution of preference heterogeneity across consumers cannot be recovered without further restrictions. While this precludes quantifying individual consumer welfare, many other questions can still be addressed, including evaluating markups, predicting equilibrium prices and quantities after a policy change and testing features of the demand functions (e.g., income effects).

In practice, I propose approximating the (inverse of the) demand system using Bernstein polynomials, which make it easy to enforce a number of economic constraints in the estimation routine. Computationally, the objective function to be minimized is convex in the parameters; thus, if the constraints are also convex, standard algorithms are guaranteed to converge to the global optimum. In order to show validity of the standard errors, I adapt proofs from recent work in econometrics on nonparametric instrumental variables regression and I provide primitive conditions for the case where the objects of interest are price elasticities and (counterfactual) equilibrium prices.

As with many nonparametric estimators, one limitation of my approach is that the number of parameters tends to increase quickly with the number of goods and/or covariates. This does preclude dealing with markets featuring many products and as such the proposed method represents a first step towards developing a widely applicable nonparametric estimator. As a contribution in that direction, I show that one can partially mitigate the curse of dimensionality by imposing micro-founded restrictions on the demand functions while preserving most of the flexibility of the nonparametric approach. Specifically, I consider (i) an exchangeability restriction; and (ii) constraints on the way covariates and prices enter the demand system. Both (i) and (ii) substantially reduce the number of parameters relative to the most general model. In particular, (ii) highlights that there is a trade-off between functional form restrictions and severity of the curse of dimensionality. In practice, this means that a researcher can—to a certain extent—tailor the model to the specifics of her setting by choosing how many assumptions to impose. For example, if the sample size is moderate, a researcher might choose to assume more in terms of functional form to contain the number of parameters, while still avoiding several assumptions on the distribution of the unobservables and consumer behavior relative to a standard discrete choice model. On the other hand, with larger samples, the researcher might be able to relax some of these functional form restrictions.

Besides requiring a nonparametric approach, the assessment of how counterfactual outcomes are affected by parametric restrictions necessitates an amount of data sufficient to

such as Dubin and McFadden (1984).

obtain informative results in the more flexible model. To this end, I leverage a large sample of store/week-level quantities and prices from Nielsen. Specifically, I focus on strawberry sales in California, which allows me to keep the number of goods low and thus avoid the curse of dimensionality. In addition, given the perishability of the product, I am able to reasonably abstract from dynamic considerations and perform a clean comparison between static demand models. Of course, this is a small product category, but the increasing availability of large data sets suggests that it might be possible to apply nonparametric approaches such as that proposed here to a much broader class of settings.

I consider two counterfactual exercises. The first is to quantify the pass-through of a tax into retail prices. Comparing the results to those given by a standard mixed logit model, I find that the nonparametrically estimated tax pass-through is significantly lower than that delivered by mixed logit for organic strawberries. I relate this to the fact that the nonparametric own-price elasticity for that good increases in absolute value much faster with own-price, which provides an incentive for the retailer to contain the price increase after the tax, all else equal. 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, e.g., Nevo (2001)). In this case, a mixed logit model with product-specific fixed effects matches the nonparametric results very closely. This is not the case for mixed logit models with fewer fixed effects, suggesting that the proposed approach may be used to guide the choice among competing parametric specifications.

Related literature. This paper contributes to the vast literature on models of demand in differentiated products markets pioneered by BLP. In particular, a recent paper by Berry and Haile (2014) (henceforth BH) shows that most of the parametric assumptions imposed by BLP are not needed for identification of the demand functions, i.e. that these restrictions are not necessary to uniquely pin down the demand functions in the hypothetical scenario in which the researcher has access to data on the entire relevant population. While I build on the *identification* result in BH, I focus on a distinct set of issues pertaining nonparametric *estimation*. Other papers developing flexible estimation approaches to demand estimation include Bajari et al. (2007), Fox et al. (2011), Fox et al. (2012), Fox et al. (2016) and Fox and Gandhi (2016). The goal in these papers is to recover the distribution of random coefficients in discrete choice settings, whereas I directly target the structural demand function. On the one hand, this allows for a broader range of consumer behaviors; on the other, as discussed above, it faces a curse of dimensionality. A recent paper by Tebaldi et al. (2019) proposes a method to obtain nonparametric bounds on demand counterfactuals and applies it to the California health insurance market, but does not develop inference procedures.

It should be emphasized that the present paper focuses 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 et al. (2004),⁹ and, for more recent nonparametric approaches, Hausman and Newey (2016), Blundell et al. (2017) and Chen and Christensen (2018).

Second, the paper is related to the large literature on incomplete pass-through¹⁰ 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 how much of the incomplete pass-through is explained by sellers adjusting their markups.¹¹ The present paper contributes to this literature by providing a method to evaluate markups that relaxes a number of restrictions on consumer behavior and preferences. In my empirical setting, I estimate a significantly larger reduction in markups after the tax—and thus a more incomplete pass-through—for the organic product relative to what is predicted by a more restrictive parametric model.¹²

Third, the paper relates to the literature investigating the sources of market power based on demand estimates, notably Nevo (2001).¹³ Once again, I offer a more flexible method to disentangle and quantify the different components of market power. In my empirical setting, I find that a mixed logit model with product fixed effects matches the nonparametric results very closely, suggesting that standard parametric models might be sufficient to address this type of questions.

Overview. 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

⁹Of course, any method based on market-level data may be immediately applied to consumer-level data by simply aggregating the latter at the market level. However, recent work by Berry and Haile (2010) shows that within-market variation makes it possible to identify demand under weaker conditions relative to the case where only market-level data is available. This opens an interesting avenue for future research on nonparametric estimation of demand based on individual-level data.

¹⁰The literature on estimating pass-through is large and I do not attempt to provide an exhaustive list of references. Here, I 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.

¹¹Specifically, Goldberg and Hellerstein (2008) and Goldberg and Hellerstein (2013) focus on exchange rate pass-through, while Nakamura and Zerom (2010) consider cost pass-through. Competing explanations for incomplete pass-through considered in these papers are nominal rigidities and the presence of costs not affected by the shocks.

¹²For 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.

¹³Another approach to studying market power is based on estimates of the firm production function (de Loecker (2011), de Loecker and Warzynski (2012)).

the proposed nonparametric estimation approach. Section 4 presents the results of several Monte Carlo simulations. Section 5 contains the empirical application. Section 6 concludes. All proofs, additional simulations, and more details on the empirical application are presented in the appendices.

2 Model and Identification

The general model I consider is the same as that in BH. In this section, I summarize the main features of the model. In a given market t , there is a continuum of consumers choosing from the set $\mathcal{J} \equiv \{1, \dots, J\}$. 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 vector χ_t is partitioned as follows:

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

where x_t is a vector of exogenous observable characteristics (e.g. exogenous product characteristics or market-level income), $p_t \equiv (p_{1t}, \dots, p_{Jt})$ are observable endogenous characteristics (typically, market prices) and $\xi_t \equiv (\xi_{1t}, \dots, \xi_{Jt})$ represent unobservables potentially correlated with p_t (e.g. unobserved product quality).¹⁴

Next, I define the structural demand system

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

where \mathcal{X} denotes the support of χ_t and 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. I 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. For example, this case arises when the goods are complements to each other and the interpretation of market shares as fractions of consumers preferring one good over all others does not apply.¹⁵ I 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.

¹⁴While the leading case in the demand estimation literature is that in which only one attribute per product (price) is endogenous, the framework allows for p_{jt} to be a vector, as long as appropriate instruments are available.

¹⁵See Example 1 in Berry et al. (2013) and the simulation in Section 4.3.

Next, following BH, I restrict the way in which some of the variables in X enter demand. Specifically, I partition x_t as $(x_t^{(1)}, x_t^{(2)})$, where $x_t^{(1)} \equiv (x_{1t}^{(1)}, \dots, x_{Jt}^{(1)})$, $x_{jt}^{(1)} \in \mathbb{R}$ for $j \in \mathcal{J}$, and define the *linear indices*

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

Then, for every market t , I assume that

$$\sigma(\chi_t) = \sigma(\delta_t, p_t, x_t^{(2)}) \tag{1}$$

where $\delta_t \equiv (\delta_{1t}, \dots, \delta_{Jt})$.¹⁶ Equation (1) requires that, for $j = 1, \dots, J$, $x_{jt}^{(1)}$ 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. In a standard BLP-type discrete choice setting, a simple sufficient condition is that $x_{jt}^{(1)}$ enters good j 's indirect utility with a non-random coefficient. On the other hand, $x_t^{(2)}$ is allowed to enter the share function in an unrestricted fashion.¹⁷ Two remarks about the restriction in (1) are in order. First, while (1) requires that the dimension of $x^{(1)}$ be equal to J , it is possible to include more than one covariate in each linear index. In fact, this is one of the strategies for dimension reduction suggested in Section 3.2. Second, $x_{jt}^{(1)}$ could be a characteristic of good j , but it need not be. The model allows for the case where $x^{(1)}$ includes market-level demand shifters that are not necessarily product specific, as long as the dimension of $x^{(1)}$ is at least J . This is illustrated in the application of Section 5, where $x^{(1)}$ consists of variables that shift consumer preferences for strawberries but do not represent product characteristics.

Throughout the paper, I assume sufficient conditions for the structural demand system σ to be point-identified, which I record in the next assumption.

Assumption 1. *There exist price instruments $Z = (Z_1, \dots, Z_J)$, excluded from the demand system, such that $\mathbb{E}(\xi_j | X, Z) = 0$ a.s. $-(X, Z)$ for $j \in \mathcal{J}$. Further, the additional sufficient conditions for point-identification of σ in BH hold.*

The moment conditions in Assumption 1, which will motivate the estimation strategy, require the unobservables ξ to be mean-independent of the price instruments and exogenous

¹⁶As shown in Appendix B of BH, what is critical for identification is the strict monotonicity of δ_{jt} in ξ_{jt} . Both its linearity in $x_{jt}^{(1)}$ and its additive separability in ξ_{jt} can be relaxed. However, the assumption in (1) simplifies the estimation procedure in that it leads to an additively separable nonparametric regression model. Given that this is the first attempt at estimating demand nonparametrically for this class of models, maintaining (1) appears to be a reasonable compromise. Footnote 20 further elaborates on the non-separable case.

¹⁷Indeed, the case where $x_t^{(2)}$ does not enter the model at all is allowed.

characteristics X . I refer the reader to BH for a detailed discussion of the additional assumptions that suffice for identification. Under these conditions, BH show that the demand system in (1) can be inverted as follows:

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

where I use the normalization $\beta_j = 1$, which is available since the unobservables ξ_{jt} have no natural scale. This is a set of equations each featuring one additively separable scalar unobservable, which makes it more amenable to estimation than the original demand system. The goal will be to flexibly estimate the functions σ_j^{-1} , or functionals thereof. Note that the distribution of ξ_{jt} will not be directly targeted in estimation; instead, ξ_{jt} will be simply recovered as a residual from (2) given an estimate of σ_j^{-1} . Thus, no restrictions will be placed on the distribution of ξ_{jt} (besides regularity conditions needed for the asymptotic results discussed below).

Before turning to the proposed estimation method, it is worth pausing to highlight the breadth of the model. First, unlike BLP-type approaches, the model places no parametric assumptions on the distribution of consumer preferences and only limited functional form restrictions (i.e., the index assumption). Second, by directly targeting the demand functions as opposed to the utility parameters, the approach can be applied to consumer behaviors and preferences beyond standard discrete choice settings, including complementarities across goods, consumer inattention, and multiple discrete or continuous choices. Section 4 illustrates this through several simulations.

Because the model does not fully specify a functional form for utility, it is not possible to recover the distribution of preferences across consumers without further restrictions. However, this is often not needed. Indeed, knowledge of the market demand functions, possibly in combination with a model of supply, is sufficient to address many questions of interest, including evaluation of markups, predicting equilibrium responses to a policy (e.g., a tax), testing hypotheses on consumer preferences or behavior (e.g., testing for the presence of income effects), and even aggregate welfare analysis.¹⁸

¹⁸For example, the aggregate change in consumer surplus due to a change in prices can be computed given knowledge of demand alone, under the assumptions spelled out in McFadden (1981) and Small and Rosen (1981). On the other hand, as pointed out by BH (Section 4.2), one important exception is the evaluation of *individual* consumer welfare, which can be performed with aggregate data only by committing to a parametric functional form for utility.

3 Nonparametric Estimation

3.1 Setup and asymptotic results

The key idea behind my estimation strategy is to combine the inverted demand system in (2) with the IV exogeneity restriction, $\mathbb{E}(\xi_j|X, Z) = 0$, to estimate σ_j^{-1} using nonparametric instrumental variables (NPIV) methods.^{19/20} In particular, I approximate the functions σ_j^{-1} via the method of sieves, i.e. using a sequence of models whose dimension grows with the sample size. For instance, in the case of polynomial approximations, the degree of the polynomials increases with the sample size. Therefore, the approach does not require one to assume any functional form asymptotically, which guards against misspecification bias. Implementing the procedure is straightforward in that, in practice, it amounts to estimating a (large) parametric model. On the other hand, proving theoretical properties of the estimator—e.g. establishing the validity of the standard errors for price elasticities—is more complicated due to the fact that the unknown parameter is an entire function as opposed to a finite-dimensional object. Specifically, I cannot rely on standard results from parametric models and I need to adapt recent results from the econometrics literature on NPIV.

Some additional notation is needed to formalize the approach. I denote by T the sample size, i.e. the number of markets in the data. While T grows to infinity asymptotically, the number of goods J is fixed. Let Σ be the space of functions to which σ^{-1} belongs²¹ and let $\psi_{M_j}^{(j)}(\cdot) \equiv \left(\psi_{1, M_j}^{(j)}(\cdot), \dots, \psi_{M_j, M_j}^{(j)}(\cdot)\right)'$ be the basis functions used to approximate σ_j^{-1} for $j \in \mathcal{J}$.²² Note that, although I suppress it in the notation, M_j grows with T for all j . Let Σ_T be the resulting sieve space for Σ . Next, I denote by $a_{K_j}^{(j)}(\cdot) \equiv \left(a_{1, K_j}^{(j)}(\cdot), \dots, a_{K_j, K_j}^{(j)}(\cdot)\right)'$ the basis functions used to approximate the instrument space for good j 's equation, and I let $A_{(j)} = \left(a_{K_j}^{(j)}(x_1, z_1), \dots, a_{K_j}^{(j)}(x_T, z_T)\right)'$ for $j \in \mathcal{J}$. Again, I suppress the dependence

¹⁹The literature on NPIV methods is vast and I refer the reader to recent surveys, such as Horowitz (2011) and Chen and Qiu (2016).

²⁰In the absence of separability of δ_{jt} in ξ_{jt} , following Appendix B of BH, we can write $\xi_{jt} = g_j(s_t, p_t, x_{jt}^{(1)}, x^{(2)})$. This is the same as (2) except that $x_{jt}^{(1)}$ now enters the unknown function g_j . In terms of implementation, the sieve estimator proposed here could be easily adapted to this setting by including $x_{jt}^{(1)}$ as an extra exogenous argument of the unknown function (see also Section 3.3 of Chen and Christensen (2018)). Of course, this would increase the dimensionality of the problem, which is why we focus on the linear index case here.

²¹This class of functions will be formally defined in Appendix A.

²²In 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 1 below does not depend on this choice, hence the general notation used in the first part of this section.

of $\{K_j\}_{j \in \mathcal{J}}$ on the sample size. I 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, I let $r_{jt}(s_t, p_t, x_t, z_t; \tilde{\sigma}_j^{-1}) \equiv \left(x_{jt}^{(1)} - \tilde{\sigma}_j^{-1}(s_t, p_t, x_t^{(2)})\right) \times a_{K_j}^{(j)}(x_t, z_t)$. Then, the estimator solves the following GMM program²³

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

where \tilde{A}^{-} denotes the Moore-Penrose inverse of a matrix \tilde{A} . The solution $\hat{\sigma}^{-1}$ to (3) minimizes a quadratic form in the terms $\{r_{jt}(\cdot), j \in \mathcal{J}, t = 1, \dots, T\}$, i.e. the implied regression residuals interacted with the instruments. When Σ_T is chosen to be a linear sieve (e.g. polynomials, splines, wavelets), the approximation to σ_j^{-1} will be of the form $\tilde{\sigma}_j^{-1} = \theta'_j \psi_{M_j}^{(j)}(\cdot)$ for $j \in \mathcal{J}$. This, in turn, implies that (3) will be a convex program in the coefficients θ , for which readily available algorithms are guaranteed to converge to the global minimizer. In contrast, BLP-type models typically require minimizing non-convex functions and thus off-the-shelf solvers are in general not guaranteed to converge to the global minimum (Knittel and Metaxoglou (2014)). One caveat to the above is that, if one wants to impose non-convex constraints on θ , the optimization problem will become harder. One such constraint is the symmetry of the Jacobian of demand with respect to prices (see Appendix C.2). On the other hand, several other constraints, including monotonicity and the exchangeability constraint considered in Section 3.2, are linear and thus can be handled with off-the-shelf convex optimization methods. For the case with linear constraints, I recommend using the Matlab package CVX (see Grant and Boyd (2008), Grant and Boyd (2014)), whereas in the presence of nonlinear—and possibly nonconvex—constraints I found Knitro to perform well (R. H. Byrd and Waltz (2006)).

I now state a result that yields asymptotically valid standard errors for generic functionals of the demand system in the i.i.d. case.²⁴ In turn, this may be used to obtain confidence intervals for quantities of interest, such as price elasticities, and establish the distribution of test statistics under a null hypothesis on consumer behavior (e.g., lack of income effects), thus yielding critical values that can be used to test the null. The result adapts Theorem D.1 in Chen and Christensen (2017) (henceforth, CC). Note that CC consider a model

²³This could also be viewed as a minimum-distance estimator in which the conditional expectation is estimated via series least squares and the weighting matrix is taken to be the identity matrix (see, e.g., p.5568 of Chen (2007)). If there is correlation in the error terms across goods, one may want to use a non-diagonal weighting matrix in order to improve efficiency.

²⁴Consistency of the estimator in the sup-norm follows directly from Theorem 3.1 of Chen and Christensen (2018).

with only one equation and one unknown function, whereas the setting here involves J equations, each with a distinct unknown function and error term ξ_j . This requires imposing additional (mild) restrictions on the covariance matrix of the errors and modifying the proof accordingly. For conciseness, the formal assumptions needed for the result and the definition of the estimator for the variance of the functional are postponed to Appendix A. In words, the assumptions restrict: (i) the distribution of the error terms by way of standard bounded moment conditions; (ii) the rate at which the dimension of the approximation to the unknown functions grows with the sample size; and (iii) the rate of convergence of the nonparametric estimator for the demand functions and their derivatives. The restriction in (iii), which is formalized in Assumption 7 in Appendix A, is high-level and I provide more primitive sufficient conditions for two special cases of interest in Theorems 2 and 3 below. Also, following CC, I consider inference on functionals of an unconstrained sieve estimator of the unknown functions. Under the assumption that the true demand functions satisfy the inequality constraints strictly, the constrained and unconstrained estimators will coincide asymptotically and thus the unconstrained standard errors will be valid in large samples. Recent papers by Chernozhukov et al. (2015) and Freyberger and Reeves (2018) develop inference procedures for constrained estimators that could be applied to our model (see also Chetverikov et al. (2018)). I leave comparing different nonparametric methods for future research.

Theorem 1. *Let f be a scalar functional of the demand system and $\hat{v}_T(f)$ be the estimator of the standard deviation of $f(\hat{\sigma}^{-1})$ defined in (10) in Appendix A. In addition, let Assumptions 1, 2, 3, 4, 5, 6 and 7 in Appendix A hold. Then,*

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

Proof. See online Appendix B. □

Importantly, the standard deviation $\hat{v}_T(f)$ in the statement of the theorem is allowed to grow to infinity with the sample size, implying that the result covers the scenario in which the functional f is estimable at a rate slower than the parametric rate \sqrt{T} . This will typically be the case when, as in the empirical analysis of Section 5, the functionals of interest are defined for a fixed market, as opposed to being averages across markets.

Theorem 1 applies to a wide class of functionals f and estimators $\hat{\sigma}^{-1}$. I now specialize the result to two functionals—price elasticities and equilibrium prices—that are key inputs for many (counterfactual) questions in industrial organization. Also, consistent with the empirical application of Section 5, I assume that $J = 2$ and that the unknown functions are

approximated via Bernstein polynomials. I state the results here and again postpone the full presentation of the assumptions, as well as the proofs, to Appendix A. In words, Theorems 2 and 3 replace the high-level Assumption 7 in Theorem 1 with sufficient conditions on the smoothness of the unknown functions, the support of the endogenous variables, and the growth rate of the sieve approximation. These are standard assumptions in the NPIV literature.

Theorem 2. *Let f_ϵ be the own-price elasticity functional defined in (11) in Appendix A, let $\hat{v}_T(f_\epsilon)$ denote the estimator of the standard deviation of $f_\epsilon(\hat{\sigma}^{-1})$ based on (10), and let Assumptions 1, 2, 3, 4(iii), 5, 6 and 8 from Appendix A hold. Then,*

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

Proof. See Appendix A. □

Next, I state a result establishing the asymptotic distribution of equilibrium prices.

Theorem 3. *Let f_{p_1} be the equilibrium price functional defined in (15) in Appendix A, let $\hat{v}_T(f_{p_1})$ denote the estimator of the standard deviation of $f_{p_1}(\hat{\sigma}^{-1})$ based on (10), and let Assumptions 1, 2, 3, 4(iii), 5, 6 and 9 from Appendix A hold. Then,*

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

Proof. See Appendix A. □

In the empirical application of Section 5, I apply Theorem 2 to obtain confidence intervals for own- and cross-price elasticities, and Theorem 3 to obtain confidence intervals for equilibrium prices under two counterfactual scenarios.

Lemmas 3–6 in Appendix A provide even more concrete restrictions for the “mildly ill-posed case,” i.e. the scenario where a measure of the degree of endogeneity in the nonparametric problem grows polynomially with the dimension of the sieve space.²⁵ In particular, the lemmas show that the assumptions of Theorems 2 and 3 can be satisfied by letting the dimension of the sieve space grow polynomially in the sample size.

Next, I provide examples of demand models satisfying the restrictions on existence of moments and smoothness that are required by Theorems 2 and 3. Regarding the first class

²⁵See Blundell et al. (2007) for a formal definition of the measure of ill-posedness and CC for a discussion of its estimation.

of restrictions, it suffices for the distribution of $\xi \equiv (\xi_1, \dots, \xi_J)$ to be non-degenerate and have some finite moment of order higher than two conditional on all values of (X, Z) . This requirement is satisfied, for instance, if ξ has a non-degenerate normal distribution with mean zero conditional on every value of (X, Z) . Turning to the smoothness restrictions, they are satisfied by any demand system that is infinitely differentiable (although the restrictions are weaker). Most models used in empirical work meet this requirement, including BLP, limited consumer information models such as Goeree (2008), and constant elasticity demand models along the lines of, e.g., Example 1 in Berry et al. (2013). Simulation results for special cases of each of these models are presented in Section 4.

3.2 Constraints

I conclude this section with a discussion of the curse of dimensionality that is inherent in nonparametric estimation. 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. While breaking the curse of dimensionality is outside the scope of this paper, I show that the issue can be partially mitigated by imposing micro-founded constraints on the estimated demand functions, including exchangeability,²⁶ index restrictions and monotonicity. I emphasize that this is not an exhaustive list, and one may wish to impose additional constraints in a given application. Conversely, not all constraints discussed in this paper need to be enforced simultaneously in order to make the approach feasible.

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, I propose to estimate the functions σ_j^{-1} using Bernstein polynomials, which are convenient for imposing economic restrictions due to their approximation properties. For a positive integer m , the Bernstein basis functions are defined as

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

where $v = 0, 1, \dots, m$ and $u \in [0, 1]$. The integer m is called the degree of the Bernstein basis.

²⁶Similar exchangeability restrictions are discussed in Pakes (1994), Berry et al. (1995) and Gandhi and Houde (2017) in relation to optimal instruments.

For our purposes, the following result from the approximation literature will be important (see, e.g., Chapter 2 of Gal (2008)).

Result 1. *Let g be a real-valued function that is continuous on $[0, 1]^N$ and define*

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

Then,

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

as $m \rightarrow \infty$.

This means that, for an appropriate choice of the coefficients, Bernstein polynomials provide a uniformly good approximation to any continuous function on the unit hyper-cube as the degree m increases. Specifically, the approximation in Result 1 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 the grid of points $\left(\frac{v_1}{m}, \dots, \frac{v_N}{m} \right)$, for $v_i = 0, \dots, m$ and $i = 1, \dots, N$. This yields an intuitive approach to imposing restrictions on the Bernstein estimator.

To illustrate, consider the special case in which the degree m is two and g is a function of $N = 2$ arguments. Since there are three terms for each argument (degree 0, 1 and 2), the tensor-product approximation is a linear combination of nine terms in total. Let θ denote the coefficients on these nine terms. Result 1 says that, as m grows to infinity, a good approximation will be such that the coefficients θ are equal to the true value of g at a grid of points over the $[0, 1] \times [0, 1]$ square. Arranging the coefficients in a matrix, we can write

$$\begin{bmatrix} \theta_{11} & \theta_{12} & \theta_{13} \\ \theta_{21} & \theta_{22} & \theta_{23} \\ \theta_{31} & \theta_{32} & \theta_{33} \end{bmatrix} = \begin{bmatrix} g(0, 0) & g(0, 0.5) & g(0, 1) \\ g(0.5, 0) & g(0.5, 0.5) & g(0.5, 1) \\ g(1, 0) & g(1, 0.5) & g(1, 1) \end{bmatrix} \quad (4)$$

This is helpful because it allows us to immediately translate restrictions on g into restrictions on θ . For instance, the constraint that g be weakly increasing in its first argument leads to the inequalities $\theta_{1i} \leq \theta_{2i} \leq \theta_{3i}$ for $i = 1, 2, 3$.

Using this argument, one can impose a number of constraints on σ^{-1} that are motivated by economic theory. First, it can be shown that the Jacobian of σ^{-1} with respect to s belongs to the class of inverse M-matrices,²⁷ which in turn implies the following lemma.

²⁷A square real matrix A is called an M-matrix if (i) it is 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. A matrix B is called an inverse M-matrix if it is the inverse of an M-matrix.

Lemma 1. *Let Assumption 2 in BH hold. Then, for all (s, p) , (i) $\frac{\partial \sigma_j^{-1}}{\partial s_k}(s, p) \geq 0$ for all j and k . If further $\left| \frac{\partial \sigma_j}{\partial \delta_j}(\delta, p) \right| \geq \sum_{k \neq j} \left| \frac{\partial \sigma_k}{\partial \delta_j}(\delta, p) \right|$ for all j , then (ii) $\frac{\partial \sigma_j^{-1}}{\partial s_j}(s, p) \geq \frac{\partial \sigma_j^{-1}}{\partial s_k}(s, p)$ for all j and all $k \neq j$. Finally, if in addition the own-price effects $\frac{\partial \sigma_j}{\partial p_j}(\delta, p)$ are negative for all j , then (iii) $\frac{\partial \sigma_j^{-1}}{\partial p_j}(s, p) \geq 0$ for all j .*

Proof. See Appendix C.3. □

This lemma translates properties of the Jacobians of σ into properties of the Jacobians of σ^{-1} . Parts (i) and (iii) yield monotonicity restrictions on σ^{-1} , whereas part (ii) yields a “diagonal dominance” constraint. The assumptions on the Jacobians of σ are mild and hold for many commonly used demand models. In particular, the assumption for part (ii) says that the own- δ effects are larger in magnitude than the sum of the cross- δ effects and the assumption for part (iii) simply requires demand to slope down in own price. Note that the implied constraints are linear in the derivatives of the σ_j^{-1} functions. Thus, since derivatives are linear operators, one can use the approximation property discussed above to enforce the constraints via linear restrictions on the Bernstein coefficients. Section 3.3 discusses how to operationalize this in practice.

Next, I focus on two types of constraints that are especially helpful in alleviating the curse of dimensionality: exchangeability and index restrictions. In order to define exchangeability, let $\pi : \{1, \dots, J\} \mapsto \{1, \dots, J\}$ be a permutation with inverse π^{-1} and, for simplicity, let $x^{(2)} = (x_1^{(2)}, \dots, x_J^{(2)})$, i.e. I assume that $x^{(2)}$ is a vector of product-specific characteristics.²⁸ Also, let $\tilde{n}_{x^{(2)}}$ be the dimension of each $x_j^{(2)}$, so that $n_{x^{(2)}} = J\tilde{n}_{x^{(2)}}$. Then, the structural demand system σ is exchangeable if

$$\sigma_j(\delta, p, x^{(2)}) = \sigma_{\pi(j)}\left(\delta_{\pi^{-1}(1)}, \dots, \delta_{\pi^{-1}(J)}, p_{\pi^{-1}(1)}, \dots, p_{\pi^{-1}(J)}, x_{\pi^{-1}(1)}^{(2)}, \dots, x_{\pi^{-1}(J)}^{(2)}\right), \quad (5)$$

for $j = 1, \dots, J$. In words, this means that the demand functions do not depend on the identity of the products, but only on their attributes $(\delta, p, x^{(2)})$.²⁹ For instance, for $J = 3$,

²⁸ This need not be the case in the general model from Section 2. For instance, $x^{(2)}$ could be a vector of market-level variables. In such settings, I say the demand system is exchangeable if $\sigma_j(\delta, p, x^{(2)}) = \sigma_{\pi(j)}(\delta_{\pi^{-1}(1)}, \dots, \delta_{\pi^{-1}(J)}, p_{\pi^{-1}(1)}, \dots, p_{\pi^{-1}(J)}, x^{(2)})$, 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.

²⁹For simplicity, here I consider the case of exchangeability across all goods $1, \dots, J$. However, one could also think of imposing exchangeability only within a subset of the goods, e.g. the set of goods produced by one firm. The arguments in this section would then apply to the subset of products on which the restriction is imposed.

exchangeability implies that

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

for all $\left(\delta_1, \bar{\delta}, \underline{\delta}, p_1, \bar{p}, \underline{p}, x_1^{(2)}, \bar{x}^{(2)}, \underline{x}^{(2)} \right)$, i.e. the demand for good 1 is the same if we switch the labels for goods 2 and 3. 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.³⁰

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

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

Lemma 2. *(i) If σ is exchangeable, then σ^{-1} is also exchangeable. Moreover, (ii) exchangeability translates into linear equality restrictions on the Bernstein coefficients and thus a reduction in the number of distinct coefficients to estimate.*

Proof. See Appendix C.3. □

Lemma 2 implies that one can directly impose exchangeability on the target functions σ^{-1} . To illustrate, in the example in (4), if the function g is exchangeable Lemma 2 yields the equalities $\theta_{21} = \theta_{12}$, $\theta_{31} = \theta_{13}$ and $\theta_{32} = \theta_{23}$. One can directly embed these restrictions in the estimation routine by minimizing the criterion function over the lower-dimensional space of free parameters (there are six free parameters out of nine in this simple example). A more formal discussion of how to impose exchangeability is provided in Appendix C.1.

Finally, I consider index restrictions. Specifically, suppose we are willing to assume that $x^{(2)}$ enters the demand functions through the indices δ . Then, each demand function goes from having $2J + n_{x^{(2)}}$ to $2J$ arguments, which in a typical scenario reduces the number of parameters from $(m + 1)^{2J + n_{x^{(2)}}}$ to $(m + 1)^{2J}$ where m is the polynomial degree for each

³⁰This also uses the fact that the idiosyncratic taste shocks are typically assumed to be iid—and thus exchangeable—across goods.

Table 1: Number of parameters with and without exchangeability and index restriction on price

J	p in Index		p not in Index	
	Exchangeability	No Exchangeability	Exchangeability	No Exchangeability
3	10	27	405	729
5	45	243	4,455	59,049
7	84	2,187	27,027	4.78×10^6
10	165	59,049	218,790	3.49×10^9

Note: Tensor product of univariate Bernstein polynomials of degree 2 for each argument of the function. $n_{x(2)}$ is assumed to be zero.

argument of the function. Similarly, if we are willing to assume that prices enter the demand functions through the indices δ , each demand function goes from having $2J + n_{x(2)}$ to $J + n_{x(2)}$ arguments, decreasing the number of parameters from $(m + 1)^{2J + n_{x(2)}}$ to $(m + 1)^{J + n_{x(2)}}$. Thus, to a certain extent, it is possible to tailor the approach based on the setting and sample size at hand by choosing how much to assume in terms of functional form. Further, note that, while the index restriction does have bite, including variables in the linear index does not mean that they are restricted to enter the demand functions linearly. As discussed in Section 2, the content of this assumption is that the variables in the index and the unobservables ξ must be perfect substitutes in the “production” of utility. For instance, in a discrete choice model, a sufficient condition is that the variables have nonrandom coefficients, but they are allowed to enter the demand functions in highly nonlinear ways. Additionally, index restrictions do not impose any constraints on the distribution of the unobservables and are thus consistent with the goal of relaxing the arbitrary distributional assumptions often made in estimating demand parametrically.

To illustrate the role played these constraints in alleviating the curse of dimensionality, I show in Table 1 how the number of parameters for each demand function grows with J depending on whether I do or do not impose exchangeability and the index restriction on p . While the dimension of the model grows large with J in both cases, the curse of dimensionality is much more severe when exchangeability or the index restriction are not imposed—indeed to the point where estimation becomes computationally intractable. Thus, 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.

3.3 Implementation of the Estimator

I conclude this section by providing a step-by-step guide to implementing the nonparametric estimator. Consistent with the empirical application in Section 5, I focus on the case with $J = 2$ goods.

1. Choose degree m and let

$$\sum_{s_{\text{own}}=0}^m \sum_{s_{\text{other}}=0}^m \sum_{p_{\text{own}}=0}^m \sum_{p_{\text{other}}=0}^m \theta_{s_{\text{own}}, s_{\text{other}}, p_{\text{own}}, p_{\text{other}}}^{(j)} b_{s_{\text{own}}, m}(s_j) b_{s_{\text{other}}, m}(s_k) b_{p_{\text{own}}, m}(\tilde{p}_j) b_{p_{\text{other}}, m}(\tilde{p}_k)$$

be the Bernstein approximation to the function σ_j^{-1} , where k denotes the good other than j and \tilde{p}_j, \tilde{p}_k denote prices normalized to the $[0, 1]$ interval.

2. Obtain constraints on Bernstein coefficients.

- (a) *Monotonicity.* By parts (i) and (iii) of Lemma 1, σ_j^{-1} is increasing in its first three arguments. Thus,

$$\theta_{s_{\text{own}}, s_{\text{other}}, p_{\text{own}}, p_{\text{other}}}^{(j)} \leq \theta_{s_{\text{own}}+1, s_{\text{other}}, p_{\text{own}}, p_{\text{other}}}^{(j)}$$

for all $s_{\text{own}} = 0, \dots, m-1$ and all $s_{\text{other}}, p_{\text{own}}, p_{\text{other}} = 0, \dots, m$;

$$\theta_{s_{\text{own}}, s_{\text{other}}, p_{\text{own}}, p_{\text{other}}}^{(j)} \leq \theta_{s_{\text{own}}, s_{\text{other}}+1, p_{\text{own}}, p_{\text{other}}}^{(j)}$$

for all $s_{\text{other}} = 0, \dots, m-1$ and all $s_{\text{own}}, p_{\text{own}}, p_{\text{other}} = 0, \dots, m$; and

$$\theta_{s_{\text{own}}, s_{\text{other}}, p_{\text{own}}, p_{\text{other}}}^{(j)} \leq \theta_{s_{\text{own}}, s_{\text{other}}, p_{\text{own}}+1, p_{\text{other}}}^{(j)}$$

for all $p_{\text{own}} = 0, \dots, m-1$ and all $s_{\text{own}}, s_{\text{other}}, p_{\text{other}} = 0, \dots, m$.

- (b) *Diagonal dominance.* Using part (ii) of Lemma 1,

$$\theta_{s_{\text{own}}, s_{\text{other}}+1, p_{\text{own}}, p_{\text{other}}}^{(j)} - \theta_{s_{\text{own}}, s_{\text{other}}, p_{\text{own}}, p_{\text{other}}}^{(j)} \leq \theta_{s_{\text{own}}+1, s_{\text{other}}, p_{\text{own}}, p_{\text{other}}}^{(j)} - \theta_{s_{\text{own}}, s_{\text{other}}, p_{\text{own}}, p_{\text{other}}}^{(j)}$$

which simplifies to

$$\theta_{s_{\text{own}}, s_{\text{other}}+1, p_{\text{own}}, p_{\text{other}}}^{(j)} \leq \theta_{s_{\text{own}}+1, s_{\text{other}}, p_{\text{own}}, p_{\text{other}}}^{(j)}$$

for all $s_{\text{own}}, s_{\text{other}} = 0, \dots, m-1$ and all $p_{\text{own}}, p_{\text{other}} = 0, \dots, m$.

3. Minimize the objective function in (3) with the approximations in step 1 in lieu of $\tilde{\sigma}_j^{-1}$, subject to the constraints in step 2. I recommend using the Matlab package CVX, but

any convex programming solver will work.

4. Plug the estimator of σ in the functional of interest. For derivatives (and thus elasticities) of the demand functions, use the implicit function theorem to write the derivatives of σ in terms of derivatives of σ^{-1} in closed form. For example, $\mathbb{J}_\sigma^p(\delta, p) = -[\mathbb{J}_{\sigma^{-1}}^s(s, p)]^{-1} \mathbb{J}_{\sigma^{-1}}^p(s, p)$, where \mathbb{J}_σ^p is the Jacobian of σ wrt prices and similarly for the other terms.

4 Monte Carlo Simulations

This section presents the results of Monte Carlo simulations. There are three goals. First, I illustrate that the estimation procedure works well with moderate sample sizes—indeed much smaller than the sample size used in the empirical application and other readily available supermarket scanner datasets. Second, I show how the general model from Section 2 may be applied to a variety of settings which include—but are not limited to—standard discrete choice. Finally, I investigate the performance of the estimator as the number of goods increases.

I compare the performance of the nonparametric demand approach (NPD for short) to that of standard methods. Specifically, I take as a benchmark a random coefficient logit model with normal random coefficients. I refer to this model as BLP. In order to summarize the results, I plot the own- and cross-price elasticities as a function of the own price, since these functions are key inputs to many counterfactuals of interest. For instance, the shape of the own-price elasticity function will turn out to play an important role in determining the pass-through rate of a tax in the application of Section 5. In each plot, all market-level variables different from the own-price are fixed at their median values. All simulations are for the case with $J = 2$ number of goods (except for Section 4.4), $T = 3,000$ number of markets. I report 95% intervals based on 200 replications of the estimator. Appendix D presents additional simulation designs in which the sample size is lower ($T = 500$), the number of goods is larger than two, and the index restriction is violated.

4.1 Correctly specified BLP model

First, I generate data from a mixed logit model with normal random coefficients. This means that the BLP procedure is correctly specified and therefore performs well. On the other hand, one would 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 sheds some light on how large a cost one has to pay for not committing to a parametric structure when that happens to be correct.

I generate the utility that consumer i derives from good j as

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

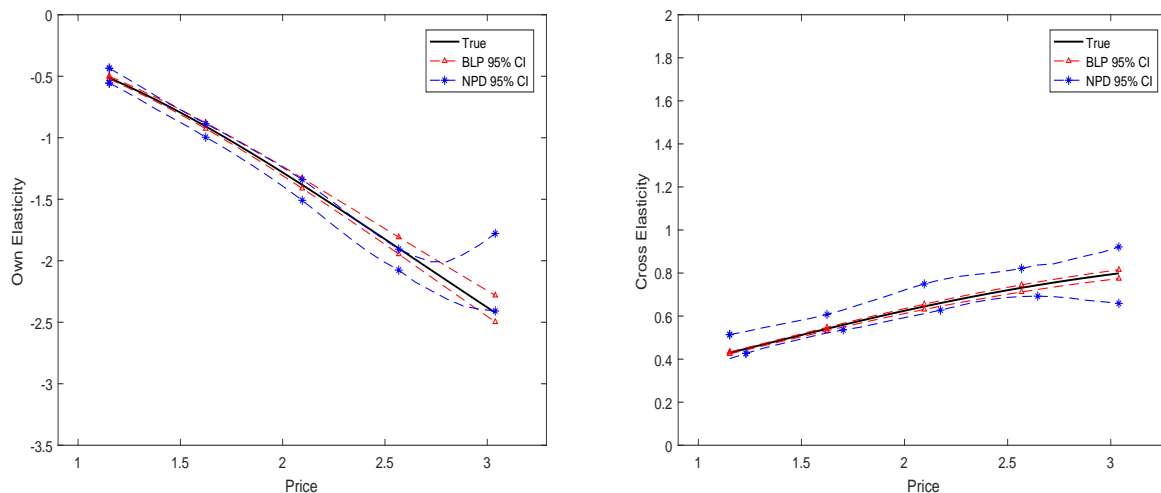
where ϵ_{ij} is independently and identically distributed (iid) extreme value across goods and consumers, α_i is distributed $N(-1, 0.15^2)$ iid across consumers and independent of ϵ_{ij} , and $\beta = 1$. There is an outside option with utility $u_{i0} = \epsilon_{i0}$, where ϵ_{i0} is also extreme-value distributed. The exogenous shifters x_j are drawn from a uniform $[0, 2]$ distribution,³¹ whereas the unobserved quality indices ξ_j are distributed $N(1, 0.15^2)$ iid across goods. Excluded instruments z_j are drawn from a uniform $[0, 1]$ distribution and I generate prices according to $p_j = 2(z_j + \eta_j) + \xi_j$, where η_j is uniform $[0, 0.1]$.³² Note that, letting $\delta_j = \beta x_j + \xi_j$ and using standard properties of the extreme-value distribution, we can write the market share for good j as $s_j = \int \frac{e^{\delta_j + \alpha p_j}}{1 + \sum_{k=1}^J e^{\delta_k + \alpha p_k}} dF_\alpha(\alpha) \equiv \sigma_j(\delta, p)$ where dF_α is the distribution of α . Therefore, this BLP data generating process yields a demand system of the form studied in this paper.

When estimating demand nonparametrically, I impose the constraints in Lemma 1 (diagonal dominance and monotonicity) and I restrict the demand functions for the two goods to be the same. Figure 1 shows the own- and cross-price elasticity functions for good 1, respectively. 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, I take this as suggestive that the penalty one pays when ignoring correct parametric assumptions in finite samples may not be substantial.

³¹Note that I drop the superscript on x_j , since in the simulations there is only one scalar exogenous shifter for each good, i.e. there is no $x^{(2)}$. This applies to all the simulations in this section.

³²Note that, while I do not specify a supply model, 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).

Figure 1: BLP model: Own-price (left) and cross-price (right) elasticity functions



Note: The solid lines are the true elasticity functions, whereas the red lines marked with triangles and the blue lines marked with asterisks correspond to the BLP and NPD 95% intervals, respectively.

One may wonder how robust the nonparametric estimates in Figure 1 are to the choice of the tuning parameter, i.e. the polynomial degree for the Bernstein approximation. Table 2 shows how the estimator for the median own- and cross-price elasticities performs as the tuning parameter changes. While, as expected, the bias tends to decrease and the standard deviation to increase with the polynomial degree, the own- and cross-price elasticities are pinned down reasonably well for a range of tuning parameters. Appendix D.4 provides more results suggesting that this does not just hold for the median levels, but also for the entire elasticity functions.

Table 2: Performance of nonparametric estimator for median own- and cross-price elasticities as the polynomial degree varies


	True	Degree	Bias	S.E.	MSE
Own	-1.339	6	0.079	0.024	0.007
		8	0.035	0.023	0.002
		12	0.003	0.028	0.001
Cross	0.569	6	-0.055	0.009	0.003
		8	-0.059	0.011	0.004
		12	-0.038	0.014	0.002

4.2 Inattention

Next, I consider a discrete choice setting with inattention. In any given market, I assume a fraction of consumers ignore good 1 and therefore maximize their utility over good 2 and the outside option only. The remaining consumers consider all goods. I 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. items that are on sale might have a special display in supermarkets or options might be filtered from cheapest to most expensive on a e-commerce platform). Except for the presence of inattentive consumers, the simulation design is the same as in Section 4.1. In nonparametric estimation, I impose the same constraints as in the previous simulation, except that I do not restrict the two demand functions to be the same, 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, I allow different constants for the two goods.

Figure 2 shows the results for good 1. The nonparametric method captures the shape of both the own- and the cross-price elasticity functions, whereas BLP tends to underestimate the own-price elasticity and overestimate the cross-price elasticity. Intuitively, BLP does not capture the fact that, as the price of good 1 increases, more and more consumers ignore good 1. This results in a BLP own-price elasticity that is too low in absolute value. Similarly, the BLP model does not capture the fact that, as the price of good 2 increases, a fraction of customers will not switch to good 1 because they ignore it. This leads to a BLP cross-price elasticity that is too high.

Figure 2: Inattention: Own-price (left) and cross-price (right) elasticity functions



Note: The solid lines are the true elasticity functions, whereas the red lines marked with triangles and the blue lines marked with asterisks correspond to the BLP and NPD 95% intervals, respectively.

4.3 Complementary goods

I now consider a setting where good 1 and 2 are not substitutes, but complements. I generate the exogenous covariates and prices as in the previous two simulations,³³ but I 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.³⁴ In the NPD estimation, I impose the same constraints as in the simulation of Section 4.1.

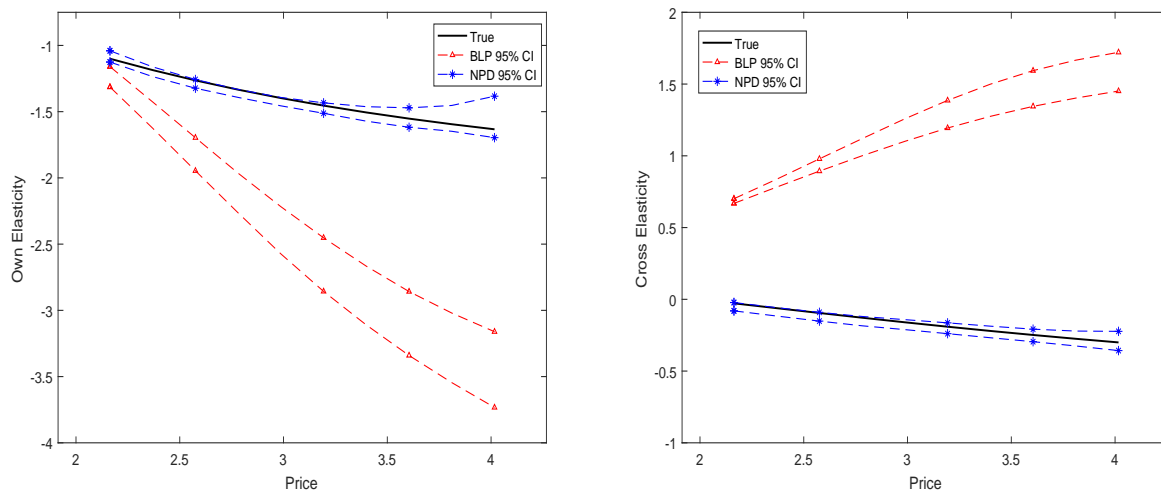
Figure 3 shows the results for good 1. Again, NPD captures the shape of the elasticity functions well. Specifically, note that the cross-price elasticity is slightly negative given that good 1 and good 2 are complements. On the other hand, the BLP confidence bands

³³One difference is that I 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.

³⁴See also Example 1 in Berry et al. (2013).

are mostly off target, consistent with the fact that a discrete choice model is not well-suited to estimate demand for (divisible) complements. In particular, the discrete choice framework implies that the goods are substitutes and thus forces the cross-price elasticity to be positive.

Figure 3: Complements: Own-price (left) and cross-price (right) elasticity functions



Note: The solid lines are the true elasticity functions, whereas the red lines marked with triangles and the blue lines marked with asterisks correspond to the BLP and NPD 95% intervals, respectively.

4.4 $J > 2$ Goods

The simulation designs considered so far featured $J = 2$ goods, which corresponds to the setting of the empirical application. However, researchers are often interested in modeling demand for a larger number of goods. To this end, I now investigate the performance of the estimator as the number of products increases. To alleviate the curse of dimensionality that arises as J grows, I both impose exchangeability and restrict prices (as well as the x attributes) to enter the indices δ in estimation. As discussed in Section 3.2, both these constraints substantially reduce the number of parameters to estimate.

The data is generated from the discrete choice dgp from Section 4.1 with one difference: the coefficients (α, β) on the product attributes (p_j, x_j) are now drawn from a discrete distribution and are correlated.³⁵ Because the product attributes have random coefficients, the

³⁵Specifically, I draw α from the distribution that places equal probabilities on the values -3 and -0.5 and set $\beta = -\alpha$, so that there are two types of consumers, one that places low weights and one that places high weights on the observable product attributes.

Table 3: Performance of estimators for median own- and cross-price elasticities as J varies

	J	True	NPD			BLP		
			Bias	S.E.	MSE	Bias	S.E.	MSE
Own	3	-1.322	-0.017	0.049	0.003	-0.980	0.052	0.963
	5	-1.458	-0.065	0.078	0.010	-1.479	0.089	2.195
	10	-1.559	0.429	0.088	0.191	-0.857	0.137	0.752
Cross	3	0.277	-0.088	0.022	0.008	0.247	0.217	0.108
	5	0.173	-0.015	0.019	0.001	0.050	0.035	0.004
	10	0.091	-0.048	0.012	0.003	0.017	0.050	0.003

Note: Mixed logit dgp with correlated discrete random coefficients. Both the NPD and the BLP model are misspecified.

index restriction is not satisfied (ξ_j does not enter the demand functions in the same way as x_j or p_j) and thus the nonparametric model I estimate is misspecified. In addition, the BLP model is also misspecified in that it incorrectly assumes that the random coefficients are normally distributed and independent of each other. Comparing the performance of the two estimators then illustrates the relative impact of two different types of misspecification: (i) that arising from incorrect distributional assumptions in a parametric model, and (ii) that stemming from incorrectly imposing the index restriction in the proposed nonparametric approach. Table 3 shows that, as the number of goods ranges from 3 to 10, the nonparametric approach consistently outperforms the parametric one in pinning down the cross-price and especially the own-price elasticities. This suggests that even the more restrictive version of the nonparametric estimator that imposes both exchangeability and the index restriction in all the product attributes might be preferable to a parametric model that makes incorrect distributional assumptions on the unobservables. Appendix D.3 further explores the robustness of the nonparametric approach to increasing violations of the index restriction, while Appendix D.5 contains estimates for the entire own- and -cross elasticity functions for the $J > 2$ case.

When the number of goods is higher (in the dozens), one would typically need additional restrictions to make the problem tractable, which is an interesting avenue for future research. As an example, Appendix D.6 provides simulations for a semiparametric model that maintains the conventional logit functional form but is flexible on how prices and covariates enter the logit “wrapper.”³⁶ Under this more restrictive model, the curse of dimensionality in the number of goods is broken and the estimator can be easily scaled to many goods.

³⁶I thank an anonymous referee for this suggestion.

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

In this section, I use the proposed nonparametric procedure to investigate the robustness of two counterfactual exercises to the parametric specification of demand. First, I 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 (see, e.g., Weyl and Fabinger (2013)). Therefore, flexibly capturing the shape of the demand function is crucial to obtaining accurate results.

The second counterfactual concerns the role played by the multi-product nature of retailers in driving up markups. Specifically, a firm simultaneously pricing multiple (substitute) goods is able to internalize the competition that would occur if those goods were sold by different firms, thus pushing prices upwards.³⁷ Quantifying the magnitude of this effect is ultimately an empirical question which again depends on the shape of the demand functions.

5.1 Data

I use data on sales of fresh fruit at stores in California. Specifically, I focus on strawberries, and look at how consumers choose between organic strawberries, non-organic strawberries and other fresh fruit, which I pool together as the outside option. While this is a small product category, it has a few features that make it especially suitable for a clean comparison between different static demand estimation methods. First, given the high perishability of fresh fruit, one 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 (see, e.g., Knee (2002)), which means that they cannot be artificially ripened using ethylene.³⁸ This limits the ability of retailers as well as consumers to stockpile and further motivates ignoring dynamic considerations in the model. Second, while strawberries are harvested in California essentially year-round, other fruits—e.g. peaches—are not, which provides some arguably exogenous supply-side variation in the richness of the outside option relative to the inside goods. Finally, the large number of store/week observations combined with the limited number of goods provide an ideal setting for the first application of a nonparametric—and thus data-intensive—estimation approach.

I instrument for prices using Hausman-type IVs, i.e. the price of the same products in

³⁷This is one of the determinants of markups considered by Nevo (2001) in his analysis of the ready-to-eat cereal industry.

³⁸Unlike climacteric fruits, such as bananas.

nearby markets (see Hausman (1996)). In addition, for the inside goods, I also use shipping-point spot prices, as a proxy for the wholesale prices faced by retailers. Besides prices, I include the following shifters in the demand functions: (i) a proxy for the availability of non-strawberry fruits in any given week; (ii) a measure of consumer tastes for organic produce in any given store; and (iii) income.

Appendix F provides further details on the construction of the dataset, as well as some summary statistics and results for the first-stage regressions.

5.2 Model

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

$$\begin{aligned}
 s_1 &= \sigma_1 (\delta_{str}, \delta_{org}, p_0, p_1, p_2, x^{(2)}) \\
 s_2 &= \sigma_2 (\delta_{str}, \delta_{org}, p_0, p_1, p_2, x^{(2)}) \\
 \delta_{str} &= \beta_{0,str} - \beta_{1,str} x_{str}^{(1)} + \xi_{str} \\
 \delta_{org} &= \beta_{0,org} + \beta_{1,org} x_{org}^{(1)} + \xi_{org}
 \end{aligned} \tag{6}$$

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, $x_{org}^{(1)}$ denotes a measure of taste for organic products,³⁹ $x_{str}^{(1)}$ denotes the availability of other fruit, $x^{(2)}$ denotes income, and (ξ_{str}, ξ_{org}) denote unobserved store/week level shocks for strawberries and organic produce, respectively. In Appendix G, I show that the demand specification (6) is consistent with two classes of models in which consumers optimally choose from the three products, the quality of organic produce $x_{org}^{(1)}$ enters the utility of good 1 and 2 and the variable $x_{str}^{(1)}$ shifts the utility of the outside option. Other micro-foundations are possible and none of the additional restrictions in the models in Appendix G are imposed in estimation. Note that (6) is an example of a model in which the exogenous demand shifters $x_{str}^{(1)}, x_{org}^{(1)}$ are not product-specific characteristics. As mentioned before, this is allowed since the key requirement is that there be at least as many exogenous shifters as there are goods. Further, notice that each of the two demand functions has its own subscript, indicating that I will not impose exchangeability restrictions across products and instead will let the two demands be arbitrarily different functions of their arguments.

The unobservables (ξ_{str}, ξ_{org}) could include, among other things, shocks to the quality of produce at the store/week level, variation in advertising and/or display across stores and

³⁹Specifically, I take the percentage of organic sales over total yearly sales in the lettuce category at the store.

time, and taste shocks idiosyncratic to a given store’s customer base (possibly varying over time). To the extent that these factors are taken into account by the store when pricing produce, the prices (p_0, p_1, p_2) will be endogenous. In contrast, I assume that the demand shifters $(x_{str}^{(1)}, x_{org}^{(1)})$ are mean independent of (ξ_{str}, ξ_{org}) . Regarding $x_{str}^{(1)}$, this is a proxy for the total supply of non-strawberry fruits in California in a given week. As such, I view this as a purely supply-side variable that shifts demand for strawberries inwards by increasing the richness of the outside option,⁴⁰ but is independent of store-level shocks.⁴¹ As for $x_{org}^{(1)}$, 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 $x_{org}^{(1)}$ and ξ_{org} . However, Appendix E shows that many objects of interest, including the counterfactuals in Section 5.4, are robust to certain forms of endogeneity arising through this channel. Note that in model (6), the exogenous shifters $x_{str}^{(1)}$ and $x_{org}^{(1)}$ are not product-specific characteristics, but rather market-level variables. As highlighted in Section 2, the framework of the paper accommodates this.

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

$$\begin{aligned} u_{i,1} &= \beta_1 + (\beta_{p,i} + \beta_{x^{(2)}}x^{(2)})p_1 + \beta_{p,0}p_0 + \beta_{str}^{par}x_{str}^{(1)} + \xi_1 + \epsilon_{i,1} \\ u_{i,2} &= \beta_2 + (\beta_{p,i} + \beta_{x^{(2)}}x^{(2)})p_2 + \beta_{p,0}p_0 + \beta_{str}^{par}x_{str}^{(1)} + \beta_{org}^{par}x_{org}^{(2)} + \xi_2 + \epsilon_{i,2} \end{aligned} \tag{7}$$

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 $\beta_{p,i}$ can take one of two values across consumers.⁴²

Comparing model (6) to model (7) illustrates the flexibility of the approach proposed in this paper. The latter model specifies the indirect utility from each good and thus imposes

⁴⁰For 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.

⁴¹The variable $x_{str}^{(1)}$ would be endogenous if the quality of strawberries sold in California supermarkets systematically varied with the harvesting patterns of other fresh fruits. This does not seem to be a first-order concern given that (i) strawberries are harvested in California essentially year-round; and (ii) more than 90% of all strawberries produced in the US are grown in California (United States Department of Agriculture (2017)).

⁴²Following the original BLP paper, I also estimated a mixed logit model with a normal random coefficient. The coefficients—and more importantly—the counterfactuals in Section 5.4 are very similar across the two specifications. In the paper, I present the two-point distribution because it is slightly more flexible than the one with normal coefficients. Specifically, the former has three parameters for the distribution of the random coefficient (the two values plus the probability of, say, the first value), while the latter has two parameters (the mean and the variance of the normal distribution).

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 other fruits. On the other hand, model (6) allows for a broader range of consumer behaviors, including continuous choice, as I show in Appendix G.2. This is one of the advantages of targeting the structural demand function directly as opposed to the underlying utility parameters.

In order to perform the counterfactual exercises in Section 5.4, I need to take a stand on the supply side. Following the retail literature, I make the assumption that each store acts as a monopolist when choosing strawberry prices. This model of supply is justified if consumers do not compare prices across stores when making their strawberry purchase decisions, which seems to be a reasonable assumption.

5.3 Estimation

In nonparametric estimation, I impose the constraints on the Jacobian of demand based on Lemma 1, but do not impose exchangeability. Thus, I allow the organic and non-organic category to have different demand functions. Further, I choose the degree of the polynomials for the Bernstein approximation based on a two-fold cross-validation procedure.⁴³ Table 4 shows the median own- and cross-price elasticities based on both the parametric and the nonparametric model.⁴⁴ One can see that the nonparametric elasticities tend to be higher than those estimated parametrically.

In order to compare the fit of the nonparametric model relative to the mixed logit model, I follow the same two-fold cross-validation approach used to choose the degree for the Bernstein polynomial approximation. As shown in Table 5, the greater flexibility of the NPD model translates into a lower average MSE.

⁴³See, e.g., Chetverikov and Wilhelm (2017). Specifically, I partition the sample into two subsamples of equal size. Then, I estimate the model using the first subsample and compute the mean squared error (MSE) for the second subsample. I repeat this procedure inverting the role of the two subsamples and use the average of the two MSEs as the criterion for choosing the polynomial degree. I let the polynomial degree vary in the set {6, 8, 10, 12, 14} and find that a polynomial of degree 10 delivers the lowest average MSE.

⁴⁴The standard errors for the elasticities as well as the counterfactual quantities in the next section are valid under the assumption that the unobservables are i.i.d. over time and across stores. Regarding i.i.d.-ness over time, recall that we are in part controlling for seasonal patterns via the availability of other fruit ($x_{str}^{(1)}$); therefore, this amounts to assuming that the remaining demand shocks are not serially correlated. Regarding i.i.d.-ness across stores, one might be concerned that stores belonging to the same chain have correlated demand shocks since they may attract similar customers or have similar quality levels. To alleviate this concern, I re-estimated the model with chain fixed effects and obtained similar results.

Table 4: Estimation results

	NPD		Mixed-logit	
	Non-organic	Organic	Non-organic	Organic
Own-price elasticity	-1.402 (0.032)	-5.503 (0.672)	-0.953 (0.011)	-2.651 (0.039)
Cross-price elasticity	0.699 (0.044)	1.097 (0.177)	0.190 (0.031)	0.959 (0.006)

Note: Median values. Standard errors in parentheses.

Table 5: Two-Fold Cross-Validation Results

	NPD	Mixed Logit
MSE	0.93	2.38

5.4 Counterfactuals

I use the estimates to address two counterfactual questions. First, I consider the effects of a per-unit tax on prices.⁴⁵ In each market, I compute the equilibrium prices when a tax is levied on each of the inside goods individually. I set the tax equal to 25% of the price for the product in that market. As shown in Table 6, the nonparametric approach delivers a higher median tax pass-through in the case of non-organic strawberries relative to the mixed logit model. However, the two confidence intervals overlap. On the contrary, in the case of organic strawberries, the nonparametric model yields a much lower median pass-through (33% of the tax) relative to mixed logit (91%) with no overlap in the confidence intervals. To shed some light on the drivers of this pattern, in Figure 4 I plot uniform confidence bands for the own-price elasticity of the organic product as a function of its price.^{46,47} The own-price elasticity estimated nonparametrically is much steeper than the parametric one. This is consistent with the pass-through results. All else equal, a retailer facing a steeper elasticity function has a stronger incentive to contain the price increase in response to the tax relative to a retailer facing a flatter elasticity function.

⁴⁵As 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, one may assume the tax is nominally levied on consumers in the form of a sales tax.

⁴⁶Own-price on the horizontal axis varies within its interquartile range. I set all other variables at their median levels, except for δ_2 which I set at its 75% percentile. Setting it at its median delivers a similar shape for the elasticity function, but noisier estimates due to the fact that s_2 —which shows up in the denominator of the elasticity—approaches zero as p_2 increases.


⁴⁷The uniform confidence bands are obtained by applying the score bootstrap procedure described in CC.

Table 6: Effect of a specific tax

	NPD	Mixed Logit
Non-organic	0.84 (0.17)	0.53 ($5 \cdot 10^{-3}$)
Organic	0.33 (0.23)	0.91 ($5 \cdot 10^{-4}$)

Note: Median changes in prices as a percentage of the tax. Standard errors in parentheses.

Figure 4: Organic strawberries: Own-price elasticity function



As a second counterfactual experiment, I quantify the “portfolio effect.” Specifically, I 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, instead of a two-product monopolist. I assume the two retailers compete on prices, compute the resulting equilibrium and compare it to the observed (monopoly) prices.⁴⁸ This type of exercise is instrumental to assessing the impact of large retailers on consumer prices. Specifically, it provides a measure of the upwards pressure on prices given by the fact that a retailer selling multiple products is able to partially internalize price competition. On the other hand,

⁴⁸Since this counterfactual exercise amounts to splitting a monopoly into a duopoly, it is related to the literature on merger analysis. See, e.g., Nevo (2000) and Jaffe and Weyl (2013) and references therein.

large retailers might tend to charge lower prices due to, among other things, economies of scale or loss-leader behavior (see , e.g., Lal and Matutes (1994), Lal and Villas-Boas (1998) and Chevalier et al. (2003)). Quantifying these different effects on prices is ultimately an empirical question that requires reliable estimates of demand.

Table 7 reports the difference between the observed prices and the prices that would arise in the counterfactual world with two single-product retailers. The parametric model in (7)—labeled Mixed Logit (I)—and the nonparametric approach yield very similar results. In the median market, both models attribute around 10% and just above 40% of markups to the portfolio effect for non-organic and organic strawberries, respectively. In other words, markups would be 10% to 40% lower in the scenario with two competing single-product retailers. One may wonder how robust this result is to modifications of the parametric specification. To this end, I estimate two additional models—labeled Mixed Logit (II) and (III)—that restrict the constants in model (7) to be the same and to be zero, respectively. Thus, while Mixed Logit (I) allows for product-specific dummies, Mixed Logit (II) only allows for a dummy for the inside goods jointly, and Mixed Logit (III) does not allow for any unobserved systematic differences between the inside goods or between the inside and the outside goods. The two restricted models tend to attribute a larger share of markups to the portfolio effect relative to the more flexible parametric specification or the nonparametric approach. This suggests that allowing for product specific dummies is important in this context and points to a wider use of the approach developed in this paper as a tool for selecting among different possible (parametric) models.

As a further step in this direction, I estimated two additional mixed-logit models: (i) I added a quadratic and a cubic term in own-price (with non-random coefficients) to specification (7); and (ii) I allowed for more flexible heterogeneity in the price sensitivity by letting the price coefficient take three values instead of two. Model (i) gives different point estimates for the counterfactuals, but similar qualitative conclusions,⁴⁹ Model (ii) gives essentially the same results as the parametric baseline in Table 4 since the weight attached to the third value for the price coefficient is estimated to be very close to zero.

⁴⁹Specifically, it yields median pass-through rates of 0.45 and 0.91 for non-organic and organic strawberries, respectively. Thus, this parametric model still over-estimates tax pass-through for the organic category and, if anything, tends to slightly under-estimate pass-through for the non-organic category. Turning to the multi-product pricing counterfactual, the model predicts larger effects relative to the baseline results, with the median markup decreasing by 15% and 49% for non-organic and organic strawberries, respectively.

Table 7: Effect of multi-product pricing

	NPD	Mixed Logit (I)	Mixed Logit (II)	Mixed Logit (III)
Non-organic	0.10 ($3 \cdot 10^{-3}$)	0.08 ($1 \cdot 10^{-3}$)	0.20 ($8 \cdot 10^{-4}$)	0.21 ($2 \cdot 10^{-3}$)
Organic	0.43 ($6 \cdot 10^{-3}$)	0.42 ($2 \cdot 10^{-3}$)	0.54 ($9 \cdot 10^{-4}$)	0.55 ($1 \cdot 10^{-3}$)

Note: Median difference between the observed prices and the optimal prices chosen by two competing retailers as a percentage of markups. Standard errors in parentheses. Mixed Logit (I) refers to the model in (7). Mixed Logit (II) refers to the model in (7) with $\beta_1 = \beta_2$; Mixed Logit (III) refers to the model in (7) with $\beta_1 = \beta_2 = 0$.

6 Conclusion

In this paper, I develop and apply a nonparametric approach to estimate demand in differentiated products markets. The methodology relaxes several arguably arbitrary restrictions on consumer behavior and preferences that are embedded in standard discrete choice models. I achieve this by estimating the demand functions nonparametrically and leveraging a number of constraints from consumer theory. Further, I provide primitive conditions sufficient to obtain valid standard errors for quantities of interest.

I then use the approach as a benchmark to test the robustness of counterfactual outcomes given by standard parametric methods. While I find that a standard model yields a higher tax pass-through for one product relative to the nonparametric approach, an exercise designed to quantify the upward pressure on prices given by the multi-product nature of sellers suggests that a flexible enough parametric model captures the patterns in the data well.

This paper opens several avenues for future research. First, it would be interesting to explore additional ways to tackle the curse of dimensionality and thus enhance the applicability of the approach. For example, in markets with dozens of goods the current methodology would typically be unfeasible. However, if good j is effectively only competing with a handful of other products, then the remaining products' prices and characteristics do not enter good j 's demand function, which would substantially reduce the dimensionality of the model. Therefore, developing a data-driven way of selecting the relevant set of competitors for a given product appears to be a promising line of research. Second, while the counterfactual analysis in this paper suggest that the nonparametric approach may be used to guide the choice among parametric specifications, additional work is required to make this argument formal. In this respect, the statistics literature on focused model selection might provide valuable insights. Finally, it would be interesting to apply the methodology proposed here to a broader range of empirical settings. For instance, a recent paper by Adao et al. (2017) shows that many questions of interest in international trade may be addressed by considering an economy where countries directly exchange factors of production instead of goods. While

their identification argument is nonparametric, they estimate a parametric model in practice. Given that production factors are low dimensional, pursuing a more flexible approach seems feasible in their setting. One could then assess how robust the results are to the maintained parametric assumptions.

References

- Abaluck, J. and A. Adams (2017). What do consumers consider before they choose? identification from asymmetric demand responses. IFS Working Paper W17/09.
- Adao, R., A. Costinot, and D. Donaldson (2017). Nonparametric counterfactual predictions in neoclassical models of international trade. *American Economic Review* 107, 633–689.
- Atkin, D. and D. Donaldson (2015). Who’s getting globalized? the size and implications of intra-national trade costs. NBER Working Paper No. 21439.
- Bajari, P., J. Fox, and S. Ryan (2007). Linear regression estimation of discrete choice models with nonparametric distributions of random coefficients. *AER Papers and Proceedings*, 459–463.
- 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).
- Berry, S. and P. Haile (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.
- Berry, S., J. Levinsohn, and A. Pakes (2004). Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy* 112, 68–105.
- Blundell, R., X. Chen, and D. Kristensen (2007). Semi-nonparametric iv estimation of shape-invariant engel curves. *Econometrica* 75(6), 1613–1669.
- Blundell, R., J. Horowitz, and M. Parey (2017). Nonparametric estimation of a nonseparable demand function under the slutsky inequality restriction. *The Review of Economics and Statistics* 99(2), 291–304.
- 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, 182–185.

- 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. (2007). Large sample sieve estimation of semi-nonparametric models. In J. J. Heckman and E. Leamer (Eds.), *Handbook of Econometrics*, Volume 6. Elsevier.
- 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.
- Chen, X. and T. Christensen (2017). Optimal sup-norm rates and uniform inference on nonlinear functionals of nonparametric iv regression. arXiv:1508.03365.
- Chen, X. and T. Christensen (2018). Optimal sup-norm rates and uniform inference on nonlinear functionals of nonparametric iv regression. *Quantitative Economics* 9, 39–84.
- Chen, X. and D. Pouzo (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., W. Newey, and A. Santos (2015). Constrained conditional moment restriction models. Working Paper (arXiv:1509.06311).
- Chetverikov, D., A. Santos, and A. M. Shaikh (2018). The econometrics of shape restrictions. *Annual Review of Economics* 10, 31–63.
- Chetverikov, D. and D. Wilhelm (2017). Nonparametric instrumental variable estimation under monotonicity. *Econometrica* 85, 1303–1320.
- 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.
- de Loecker, J. (2011). Recovering markups from production data. *International Journal of Industrial Organization* 29, 350—355.
- de Loecker, J. and F. Warzynski (2012). Markups and firm-level export status. *American Economic Review* 102, 2437—2471.
- Deaton, A. and J. Muellbauer (1980, June). An almost ideal demand system. *AER* 70(3), 312–326.

- 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.
- Fox, J. and A. Gandhi (2016). Nonparametric identification and estimation of random coefficients in multinomial choice models. *RAND Journal of Economics* 47(1), 118–139.
- Fox, J., K. il Kim, S. Ryan, and P. Bajari (2011). A simple estimator for the distribution of random coefficients. *Quantitative Economics* 2, 381–418.
- Fox, J., K. il Kim, S. Ryan, and P. Bajari (2012). The random coefficients logit model is identified. *Journal of Econometrics* 166, 204–212.
- Fox, J., K. il Kim, and C. Yang (2016). A simple nonparametric approach to estimating the distribution of random coefficients in structural models. *Journal of Econometrics* 195, 236–254.
- Freyberger, J. and B. Reeves (2018). Inference under shape restrictions. Working Paper.
- Gal, S. (2008). *Shape-Preserving Approximation by Real and Complex Polynomials*. Boston: Birkhauser.
- Gandhi, A. and J. F. Houde (2017). Measuring substitution patterns in differentiated products industries. Working Paper.
- 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 Review: Papers and Proceedings* 98, 423–429.
- Goldberg, P. K. and R. Hellerstein (2013). A structural approach to identifying the sources of local currency price stability. *Review of Economic Studies* 80, 175–210.

- Grant, M. and S. Boyd (2008). Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura (Eds.), *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pp. 95–110. Springer-Verlag Limited. http://stanford.edu/~boyd/graph_dcp.html.
- Grant, M. and S. Boyd (2014, March). CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>.
- Hausman, J. and W. Newey (2016). Individual heterogeneity and average welfare. *Econometrica* 84(3), 1225–1248.
- Hausman, J. A. (1996). Valuation of new goods under perfect and imperfect competition. In T. F. Bresnahan and R. J. Gordon (Eds.), *The Economics of New Goods*, Chapter 5, pp. 209–248. Chicago: University of Chicago Press.
- 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.
- Jaffe, S. and G. Weyl (2013). The first-order approach to merger analysis. *American Economic Journal: Microeconomics* 5, 188–218.
- Knee, M. (2002). *Fruit Quality and Its Biological Basis*. CRC Press.
- Knittel, C. R. and K. Metaxoglou (2014). Estimation of random coefficient demand models: Two empiricists’ perspectives. *The Review of Economic and Statistics* 96(1), 34–59.
- 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.
- McFadden, D. (1981). Econometric models of probabilistic choice. In C. Manski and D. McFadden (Eds.), *Structural Analysis of Discrete Data with Econometric Applications*. Cambridge, MA: MIT Press.

- 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.
- Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica* 69, 307–642.
- Newey, W. (1997). Convergence rates and asymptotic normality for series estimators. *Journal of Econometrics* 79, 147–168.
- Pakes, A. (1994). Dynamic structural models, problems and prospects: mixed continuous discrete controls and market interactions. In J.-J. Laffont and C. Sims (Eds.), *Advances in Econometrics: The Sixth World Congress of the Econometric Society*, Volume 2, pp. 171–260. New York: Cambridge University Press.
- Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy* 110(4), 705–729.
- Pinkse, J. and M. E. Slade (2004). Mergers, brand competition, and the price of a pint. *European Economic Review* 48(3), 617–643.
- 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>.
- R. H. Byrd, J. N. and R. Waltz (2006). Knitro: An integrated package for nonlinear optimization. In G. di Pillo and M. Roma (Eds.), *Large-Scale Nonlinear Optimization*, pp. 35–59. Springer. http://stanford.edu/~boyd/graph_dcp.html.
- Schumaker, L. (2007). *Spline Functions: Basic Theory*. Cambridge: Cambridge Mathematical Library.
- Seber, G. A. (2007). *A Matrix Handbook for Statisticians*. New York: Wiley.
- Small, K. and H. Rosen (1981). Applied welfare economics with discrete choice models. *Econometrica* 49, 105–130.
- Souza-Rodrigues, E. A. (2014). Nonparametric estimation of a generalized regression model with group effects. working paper, University of Toronto, Department of Economics.

- Tebaldi, P., A. Torgovitzky, and H. Yang (2019). Nonparametric estimates of demand in the california health insurance exchange. NBER Working Paper No. 25827.
- 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.

Appendix A presents the assumptions for Theorems 1, 2, 3, and the proofs for Theorems 2 and 3. Online Appendix B contains the proof for Theorem 1 and supplementary results for inference. Online Appendix C discusses additional economic constraints and shows how to enforce them in estimation. Online Appendix D presents the results of additional Monte Carlo simulations. Online Appendix E discusses violations of the exogeneity restriction maintained throughout the paper. Online Appendix F discusses the construction of the data and contains descriptive statistics. Finally, Online Appendix G provides two possible micro-foundations for the demand model estimated in the empirical application.

Appendix A: Inference

This appendix contains all the notation and assumptions for the inference results in Section 3 of the paper, as well as the proofs for Theorem 2 and 3.

A.1 Setup and Notation

For simplicity, we focus on the case where there are no additional exogenous covariates $x^{(2)}$ in the demand system. Accordingly, we drop $x^{(2)}$ and use x to denote what was denoted by $x^{(1)}$ in the main text. As pointed out by CC (Section 3.3), allowing for $x^{(2)}$ is straightforward and does not change anything in the implementation of the estimator.

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 $Y \equiv (S, P)$ denote the arguments of the unknown functions and \mathcal{Y} denote its support. For every $y \in \mathcal{Y}$, let $h_0(y) \equiv [h_{0,1}(y), \dots, h_{0,J}(y)]' \equiv [\sigma_1^{-1}(y), \dots, \sigma_J^{-1}(y)]'$, so that the estimating equations become

$$x_j = h_{0,j}(y) + \xi_j, \quad j \in \mathcal{J}. \quad (8)$$

We assume that, for $j \in \mathcal{J}$, $h_{0,j} \in \mathcal{H}$, where \mathcal{H} is the Hölder ball of smoothness r , and we endow it with the norm $\|\cdot\|_\infty$ defined by $\|h\|_\infty \equiv \max_{j \in \mathcal{J}} \|h_j\|_{1,\infty}$ for a function $h = [h_1, \dots, h_J]$, where $\|h_j\|_{1,\infty}$ denotes the sup-norm for a scalar-valued function h_j . We also let $\|v\|$ denote the Euclidean norm of a vector v , $\|\mathbf{M}\|$ denote the norm of an m_1 -by- m_2 matrix \mathbf{M} defined as $\|\mathbf{M}\| \equiv \sup \{\|\mathbf{M}v\| : v \in \mathbb{R}^{m_2}, \|v\| = 1\}$, and $(\mathbf{M})_l^- \equiv (\mathbf{M}'\mathbf{M})^{-1}\mathbf{M}'$ be the left-inverse of a matrix \mathbf{M} .

Further, we let $\{\psi_{1,M_i}^{(i)}, \dots, \psi_{M_i,M_i}^{(i)}\}$ be the collection of basis functions used to approximate $h_{0,i}$ for $i \in \mathcal{J}$, and let $M = \sum_{j=1}^J M_j$ be the dimension of the overall sieve space for h . Similarly, we let $\{a_{1,K_i}^{(i)}, \dots, a_{K_i,K_i}^{(i)}\}$ be the collection of basis functions used to approximate the instrument space for $h_{0,i}$, and let $K = \sum_{j=1}^J K_j$.

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

with $j \in \mathcal{J}$, we define, for $i \in \mathcal{J}$

$$\psi_{M_i}^{(i)}(y) = \left(\psi_{1,M_i}^{(i)}(y), \dots, \psi_{M_i,M_i}^{(i)}(y) \right)' \quad M_i \text{ - by - } 1$$

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

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

Also, we let, for $j, k \in \mathcal{J}$,

$$\begin{aligned}
\Omega_{jk} &= \Omega'_{kj} = \mathbb{E} \left(\xi_{jt} \xi_{kt} a_{K_j}^{(j)}(W_t) a_{K_k}^{(k)}(W_t)' \right) & K_j - \text{by} - K_k \\
\Omega &= \begin{bmatrix} \Omega_{11} & \Omega_{12} & \cdots & \Omega_{1J} \\ \Omega_{21} & \Omega_{22} & \cdots & \Omega_{2J} \\ \vdots & \vdots & \ddots & \vdots \\ \Omega_{J1} & \Omega_{J2} & \cdots & \Omega_{JJ} \end{bmatrix} & K - \text{by} - K
\end{aligned}$$

and, similarly,

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

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

where $\hat{\xi}_{jt} = x_{jt} - \hat{h}_j(y_t)$.

For $i \in \mathcal{J}$, we define

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

and let $\zeta \equiv \max_{j \in \mathcal{J}} \zeta_j$. The rate at which ζ diverges to infinity with the sample size will play a role in the proofs. When splines (including Bernstein polynomials) are used for a_K and ψ_M , we have $\zeta = O(\sqrt{M})$ (see, e.g., Newey (1997).)

As in CC, we use the following sieve measure of ill-posedness, for $i \in \mathcal{J}$,

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

where $\tilde{\Psi}_{M_i}$ is the closed linear span of $\left\{ \psi_{M_i}^{(i)} \right\}$ and we let $\tau_M \equiv \max_{j \in \mathcal{J}} \tau_{M_j}^{(j)}$. The rate at which τ_M diverges to infinity may be viewed as a measure of how difficult the estimation problem is. In order to formalize this, we will need appropriate notation. Specifically, letting a_T and b_T be sequences of positive numbers, the notation $a_T \lesssim b_T$ means $\limsup_{T \rightarrow \infty} a_T/b_T < \infty$, and the notation $a_T \asymp b_T$ means $a_T \lesssim b_T$ and $b_T \lesssim a_T$. Next, 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 $h = [h_1, \dots, h_J] : \mathcal{Y} \mapsto \mathbb{R}^J$, we let $\partial^{\tilde{\alpha}} h \equiv [\partial^{\tilde{\alpha}} h_1, \dots, \partial^{\tilde{\alpha}} h_J]$.

The (unconstrained) sieve NPIV estimator \hat{h}_i has the following closed form

$$\hat{h}_i(y) = \psi_{M_i}^{(i)}(y)' \hat{\theta}_i$$

for

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

We write this in a more compact form as

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

Stacking the J estimators, we write

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

and

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

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

$$\tilde{\theta} = \frac{1}{T} \left[\hat{L}' \hat{G}_A^- \hat{L} \right]^{-1} \hat{L}' \hat{G}_A^- A' H_0 \quad (9)$$

and let

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

For any functional $f : \mathcal{H} \mapsto \mathbb{R}$ and any $(h, v) \in \mathcal{H} \times \mathcal{H}$, 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 (if it exists). Next, letting $vec_{g,J,j}$ be the column J -vector valued function that returns all zeros except for the j -th element, where it returns the function g , we define

$$\begin{aligned} Df(h) \left[\psi_{M_j}^{(j)} \right] &\equiv \left(Df(h) \left[vec_{\psi_{1,M_j}, J, j}^{(j)} \right], \dots, Df(h) \left[vec_{\psi_{M_j, M_j}, J, j}^{(j)} \right] \right)' & M_j \text{ - by - } 1 \\ Df(h) [\psi_M] &\equiv \left(Df(h) \left[\psi_{M_1}^{(1)} \right]', \dots, Df(h) \left[\psi_{M_J}^{(J)} \right]' \right)' & M \text{ - by - } 1 \end{aligned}$$

Finally, we let

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

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

$$\hat{v}_T^2(f) = Df(\hat{h}) [\psi_M]' \left(\hat{L}' \hat{G}_A^{-1} \hat{L} \right)^{-1} \hat{L}' \hat{G}_A^{-1} \hat{\Omega} \hat{G}_A^{-1} \hat{L} \left(\hat{L}' \hat{G}_A^{-1} \hat{L} \right)^{-1} Df(\hat{h}) [\psi_M] \quad (10)$$

Because the functionals of interest are defined for fixed (\bar{s}, \bar{p}) , they will typically be slower than \sqrt{T} -estimable (or “irregular”), i.e. $v_T^2(f) \nearrow \infty$ as $T \rightarrow \infty$.

A.2 Assumptions for Theorem 1

This section collects the assumptions for Theorem 1. The proof can be found in online Appendix B.

Assumption 2. *The variables (X_t, Z_t, P_t, ξ_t) are independent and identically distributed across markets.*

Assumption 3. *For all $j, k \in \mathcal{J}, j \neq k$:*

- (i) $\sup_{w \in \mathcal{W}} \mathbb{E}(\xi_j^2 | w) \leq \bar{\sigma}^2 < \infty$;
- (ii) $\inf_{w \in \mathcal{W}} \mathbb{E}(\xi_j^2 | w) \geq \underline{\sigma}^2 > 0$;
- (iii) $\sup_{w \in \mathcal{W}} \mathbb{E}(|\xi_j \xi_k| | w) \leq \bar{\sigma}_{cov} < \infty$;
- (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 4. (i) $\tau_M \zeta \sqrt{M(\log M)/T} = o(1)$;

(ii) $\zeta^{(2+\gamma^{(1)})/\gamma^{(1)}} \sqrt{(\log K)/T} = o(1)$ and $\zeta^{(1+\gamma^{(2)})/\gamma^{(2)}} \sqrt{(\log K)/T} = o(1)$, where $\gamma^{(1)}, \gamma^{(2)} > 0$ are defined in Assumption 3(v)-3(vi);

(iii) $K \asymp M$ and $\zeta = O(\sqrt{M})$.

Assumption 5. The basis used for the instrument space is the same across all goods, i.e. $K_j = K_k$ and $a_{K_j}^{(j)}(\cdot) = a_{K_k}^{(k)}(\cdot)$ for all $j, k \in \mathcal{J}$.

Assumption 6. $\|\hat{h} - h_0\|_\infty = o_p(1)$.

Assumption 7. 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 that the sieve variance $v_T(f)$ for the functional f is strictly positive 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) $\left| f(\hat{h}) - f(h_0) - Df(h_0)[\hat{h} - h_0] \right| \lesssim \|\partial^{\alpha_1} \hat{h} - \partial^{\alpha_1} h_0\|_\infty \|\partial^{\alpha_2} \hat{h} - \partial^{\alpha_2} h_0\|_\infty$;

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

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

Discussion of assumptions. Assumption 3 corresponds to Assumption 2 in CC, modified to account for the fact that my model has J equations and J error terms. Assumption 4(i) corresponds to the condition imposed by CC in Theorem D.1, whereas 4(ii) is similar to Assumption 3(iii) in CC. Assumption 4(iii) restricts the growth rates of the sieve spaces for the endogenous variables and the instruments. The requirement that $\zeta = O(\sqrt{M})$ holds, for instance, when splines are used to approximate the unknown functions (see, e.g., Newey (1997)). I impose it since in practice I advocate using Bernstein polynomials, which are a special case of splines. Assumption 5 is not necessary but I impose it for simplicity. Assumption 6 requires \hat{h} to be a consistent estimator. CC provide sufficient conditions for it and characterize the rate of convergence. Assumption 7 corresponds to the sufficient conditions in Remark 4.1 of CC.

A.3 Theorem 2: 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 $(\bar{s}, \bar{p}) \equiv (\bar{s}_1, \bar{s}_2, \bar{p}_1, \bar{p}_2)$ and Bernstein polynomials are used for both the endogenous variables and the instruments. The goal is to provide sufficient, lower-level conditions for Theorem 1. Analogous arguments hold for the own-price elasticity of good 2 and for the cross-price elasticities.

The functional of interest takes the form

$$f_\epsilon(h_0) = -\frac{p_1}{s_1} \frac{\frac{\partial h_{0,2}(\bar{s}, \bar{p})}{\partial s_2} \frac{\partial h_{0,1}(\bar{s}, \bar{p})}{\partial p_1} - \frac{\partial h_{0,1}(\bar{s}, \bar{p})}{\partial s_2} \frac{\partial h_{0,2}(\bar{s}, \bar{p})}{\partial p_1}}{\frac{\partial h_{0,1}(\bar{s}, \bar{p})}{\partial s_1} \frac{\partial h_{0,2}(\bar{s}, \bar{p})}{\partial s_2} - \frac{\partial h_{0,1}(\bar{s}, \bar{p})}{\partial s_2} \frac{\partial h_{0,2}(\bar{s}, \bar{p})}{\partial s_1}} \equiv -\frac{p_1}{s_1} \frac{N_1 - N_2}{D_1 - D_2} \quad (11)$$

Theorem 2 maintains the following assumption.

Assumption 8. (i) P has bounded support and (P, S) have densities bounded away from 0 and ∞ ;
(ii) The basis used for both the endogenous variables and the instruments is tensor-product Bernstein polynomials. Further, the univariate Bernstein polynomials for the endogenous variables all have the same degree $M^{1/4}$;
(iii) The unknown functions $h_0 = [h_{0,1}, h_{0,2}]'$ belong to the Hölder ball of smoothness $r \geq 8$ and finite radius;

- (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 3(v)-3(vi);
- (v) $\frac{\sqrt{T}}{v_T(f_\epsilon)} \times \left(M^{\frac{3-r}{4}} + \frac{\tau_M M^{\frac{9-r}{4}}}{\sqrt{T}} + \tau_M^2 M^3 \frac{\log M}{T} \right) = o(1)$.

Discussion of Assumption 8. Assumptions 8(i), 8(iii) and 8(iv) are conditions needed to apply the sup-norm rate results in CC.⁵⁰ 8(ii) is assumed for simplicity but it is not necessary. 8(v) corresponds to the second part of Assumption CS(v) in CC and is used to verify Assumption 7. More concrete sufficient conditions for Assumptions 8(iv) and 8(v) may be provided in specific settings. For example, Lemma 3 below gives sufficient conditions for the mildly ill-posed case.⁵¹

We now provide a proof of Theorem 2.

Proof of Theorem 2. We prove the statement by showing that the assumptions of Theorem 1 hold. Assumptions 2, 3, 4(iii), 5 and 6 are maintained. Assumption 4(i) is implied by Assumptions 4(iii) and 8(v), and Lemma 10. Similarly, Assumption 4(ii) is implied by Assumptions 4(iii) and 8(iv). We now verify Assumption 7. In what follows, unless otherwise specified, it is assumed that the arguments of all functions are (\bar{s}, \bar{p}) and the dependence is suppressed for notational convenience.

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

$$Df_\epsilon(h_0)[v] \equiv \left. \frac{\partial f_\epsilon(h_0 + \tau v)}{\partial \tau} \right|_{\tau=0} = \frac{p_1}{s_1} \left(C_1 \frac{\partial v_2}{\partial s_2} + C_2 \frac{\partial v_1}{\partial s_2} + C_3 \frac{\partial v_1}{\partial p_1} + C_4 \frac{\partial v_2}{\partial p_1} + C_5 \frac{\partial v_2}{\partial s_1} + C_6 \frac{\partial v_1}{\partial s_1} \right) \quad (12)$$

where

$$\begin{aligned} C_1 &= -\frac{(D_1 - D_2) \frac{\partial h_{0,1}}{\partial p_1} - (N_1 - N_2) \frac{\partial h_{0,1}}{\partial s_1}}{(D_1 - D_2)^2} & C_2 &= -\frac{-(D_1 - D_2) \frac{\partial h_{0,2}}{\partial p_1} + (N_1 - N_2) \frac{\partial h_{0,2}}{\partial s_1}}{(D_1 - D_2)^2} \\ C_3 &= -\frac{\frac{\partial h_{0,2}}{\partial s_2}}{(D_1 - D_2)} & C_4 &= \frac{\frac{\partial h_{0,1}}{\partial s_2}}{(D_1 - D_2)} \\ C_5 &= -\frac{(N_1 - N_2) \frac{\partial h_{0,1}}{\partial s_2}}{(D_1 - D_2)^2} & C_6 &= \frac{(N_1 - N_2) \frac{\partial h_{0,2}}{\partial s_2}}{(D_1 - D_2)^2} \end{aligned}$$

Therefore, $Df_\epsilon(h_0) : \mathcal{H} \mapsto \mathbb{R}$ is a linear functional.

Next, note that, for any $h = [h_1, h_2] \in \mathcal{H}_T$,

$$\begin{aligned} \left| \frac{\partial h_1}{\partial s_1} - \frac{\partial h_{0,1}}{\partial s_1} \right| &\leq \int_{-\infty}^{\bar{s}_2} \int_{-\infty}^{\bar{p}_1} \left| \frac{\partial^3 h_1}{\partial s_1 \partial s_2 \partial p_1} \left(\bar{s}_1, s_2, p_1, \bar{p}_2 \right) - \frac{\partial^3 h_{0,1}}{\partial s_1 \partial s_2 \partial p_1} \left(\bar{s}_1, s_2, p_1, \bar{p}_2 \right) \right| ds_2 dp_1 \\ &\leq \text{constant} \left\| \left\| \frac{\partial^3 h_1}{\partial s_1 \partial s_2 \partial p_1} - \frac{\partial^3 h_{0,1}}{\partial s_1 \partial s_2 \partial p_1} \right\| \right\|_{1, \infty} \end{aligned}$$

where the first inequality follows from the triangle inequality and the fundamental theorem of calculus, and the second inequality follows from assumption 8(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 (12)

⁵⁰CC establish sup-norm rate results for the case where the unknown function is approximated using B-splines, among others. Since Bernstein polynomials are a special case of splines (see, e.g., Schumaker (2007)), their results apply to the setting considered here.

⁵¹See CC (p.15) for a formal definition of mild and severe ill-posedness.

and write

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

which shows that Assumption 7(i) holds with $\alpha = [1, 1, 1, 0]$.

7(ii) By the mean value theorem,

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

$$\begin{aligned} \tilde{C}_1 &= -\frac{(\tilde{D}_1 - \tilde{D}_2) \frac{\partial \tilde{h}_1}{\partial p_1} - (\tilde{N}_1 - \tilde{N}_2) \frac{\partial \tilde{h}_1}{\partial s_1}}{(\tilde{D}_1 - \tilde{D}_2)^2} & \tilde{C}_2 &= -\frac{-(\tilde{D}_1 - \tilde{D}_2) \frac{\partial \tilde{h}_2}{\partial p_1} + (\tilde{N}_1 - \tilde{N}_2) \frac{\partial \tilde{h}_2}{\partial s_1}}{(\tilde{D}_1 - \tilde{D}_2)^2} \\ \tilde{C}_3 &= -\frac{\frac{\partial \tilde{h}_2}{\partial s_2}}{(\tilde{D}_1 - \tilde{D}_2)} & \tilde{C}_4 &= \frac{\frac{\partial \tilde{h}_1}{\partial s_2}}{(\tilde{D}_1 - \tilde{D}_2)} \\ \tilde{C}_5 &= -\frac{(\tilde{N}_1 - \tilde{N}_2) \frac{\partial \tilde{h}_1}{\partial s_2}}{(\tilde{D}_1 - \tilde{D}_2)^2} & \tilde{C}_6 &= \frac{(\tilde{N}_1 - \tilde{N}_2) \frac{\partial \tilde{h}_2}{\partial s_2}}{(\tilde{D}_1 - \tilde{D}_2)^2} \end{aligned}$$

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

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

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

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

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

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

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

$$+ O_p\left(M^{\frac{3-r}{4}}\right)$$

Thus, Assumption 7(iii) is implied by Assumption 8(v).

7(iv) By Remark 4.1 in CC, a sufficient condition for Assumption 7(iv) is

$$T_{iv,\epsilon} \equiv \frac{\tau_M \sqrt{\sum_{m=1}^M \left(Df_\epsilon(\hat{h}) \left[\left(G_\psi^{-1/2} \psi_M \right)_m \right] - Df_\epsilon(h_0) \left[\left(G_\psi^{-1/2} \psi_M \right)_m \right] \right)^2}}{v_T(f_\epsilon)} = o_p(1) \quad (13)$$

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 $Df_\epsilon(\hat{h}) \left[\left(G_\psi^{-1/2} \psi_M \right)_m \right] - Df_\epsilon(h_0) \left[\left(G_\psi^{-1/2} \psi_M \right)_m \right]$ for every m as the linear combination of terms, where each term is the difference between a 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, using Corollary 3.1 in CC and the well-known rate results for splines and their derivatives in, e.g., Newey (1997),

$$T_{iv,\epsilon} = O_p\left(\frac{\sqrt{T}}{v_T(f_\epsilon)} \times \left[\frac{\tau_M M^{(9-r)/4}}{\sqrt{T}} + \frac{\tau_M^2 M^{11/4} \sqrt{\log M}}{T} \right]\right) \quad (14)$$

The conclusion in (13) then follows from Assumption 8(v). □

The next two lemmas provide more primitive sufficient conditions for Assumptions 8(iv) and 8(v). Both lemmas focus on the “mildly ill-posed” in which τ_M grows polynomially in M and show that Assumptions 8(iv) and 8(v) can be satisfied by letting M increase polynomially with T . First, we consider a case where the functional f_ϵ is irregular.

Lemma 3. *Let Assumptions 8(i) and 8(iii) hold. Further, let $(v_T(f_\epsilon))^2 \asymp M^{a+\varsigma+1}$ (i.e. the functional f_ϵ is irregular) and $\tau_M \asymp M^{\varsigma/2}$ for $a \leq 0, \varsigma \geq 0, a + \varsigma + 1 > 0, r + 2a - 7 > 0$. Then, Assumptions 8(iv) and 8(v) are satisfied if $M \asymp T^\rho$ with*

$$\rho \in \left(\frac{2}{r - 3 + 2(a + \varsigma + 1)}, \min \left\{ \frac{1}{\varsigma - a + 5}, \frac{\gamma^{(1)}}{2 + \gamma^{(1)}}, \frac{\gamma^{(2)}}{1 + \gamma^{(2)}} \right\} \right)$$

Further, M may be chosen to satisfy the latter condition if $r + 4a - 11 > 0$ and $\gamma^{(i)}(r + 2a + 2\varsigma - 3) - 4 > 0$ for $i \in \{1, 2\}$.

Proof. The result follows by inspection. □

Next, we consider the case where the functional f_ϵ is regular.

Lemma 4. *Let Assumption 8(i) hold. Further, let $v_T(f_\epsilon) = O(1)$ (i.e. the functional f_ϵ is regular), $\tau_M \asymp M^{\varsigma/2}$ for $\varsigma \geq 0$, and let the unknown functions $h_0 = [h_{0,1}, h_{0,2}]'$ belong to the Hölder ball of smoothness $r \geq 15 + 4\varsigma$ and finite radius. Then, Assumptions 8(iv) and 8(v) are satisfied if $M \asymp T^\rho$ with*

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

Further, M may be chosen to satisfy the latter condition if $\gamma^{(i)}(r-5)-4 > 0$ for $i \in \{1, 2\}$.

Proof. The result follows by inspection. \square

A.4 Theorem 3: Equilibrium price functionals

We now specialize Theorem 1 to the case where the functional f is the equilibrium price of good 1 in a market with two goods characterized by marginal costs $\overline{mc} \equiv (\overline{mc}_1, \overline{mc}_2)$ and indices $\overline{\delta} \equiv (\overline{\delta}_1, \overline{\delta}_2)$. 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 $f_{p_1}(\hat{h})$. An analogous argument holds for the price of good 2. Again, I let $h_0 = [h_{0,1}, h_{0,2}]$ denote the inverse of the demand system σ_0 . Further, I use $h_0^{-1} = [h_{0,1}^{-1}, h_{0,2}^{-1}] = [\sigma_{0,1}, \sigma_{0,2}]$ 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

$$\begin{bmatrix} g_1(\overline{\delta}, \overline{p}, \overline{mc}, h_0) \\ g_2(\overline{\delta}, \overline{p}, \overline{mc}, h_0) \end{bmatrix} \equiv - \left[(\mathbb{J}_{h_0}^s)^{-1} \mathbb{J}_{h_0}^p \right]' \begin{bmatrix} \overline{p}_1 - \overline{mc}_1 \\ \overline{p}_2 - \overline{mc}_2 \end{bmatrix} + \begin{bmatrix} h_{0,1}^{-1}(\overline{\delta}, \overline{p}) \\ h_{0,2}^{-1}(\overline{\delta}, \overline{p}) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad (15)$$

where

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

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 endogenous variables and the instruments 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 $r \geq 9$ and finite radius;
(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 3(v)-3(vi);
(v) $\frac{\sqrt{T}}{v_T(f_{p_1})} \times \left(M^{\frac{4-r}{4}} + \frac{\tau_M M^{\frac{10-r}{4}}}{\sqrt{T}} + \tau_M^2 M^3 \frac{\log M}{T} \right) = o(1)$.

Discussion of Assumptions. Assumptions 9(i), 9(iii) and 9(iv) are conditions needed 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 7. More concrete sufficient conditions for Assumptions 9(iv) and 9(v) may be provided in specific settings. For example, Lemma 5 below gives sufficient conditions for the mildly ill-posed case.

We now provide a proof of Theorem 3.

Proof of Theorem 3. We prove the statement by showing that the assumptions of Theorem 1 hold. Assumptions 2, 3, 4(iii), 5 and 6 are maintained. Assumption 4(i) is implied by Assumptions 4(iii) and 9(v), and Lemma 10. Similarly, Assumption 4(ii) is implied by Assumptions 4(iii) and 9(iv).

We now verify Assumption 7.

7(i) Applying the implicit function theorem to (15),

$$Df_p(h)[v] = - \left[\begin{array}{cc} \frac{\partial g_1(\bar{\delta}, \bar{p}, \bar{m}\bar{c}, h + \tau v)}{\partial p_1} & \frac{\partial g_1(\bar{\delta}, \bar{p}, \bar{m}\bar{c}, h + \tau v)}{\partial p_2} \\ \frac{\partial g_2(\bar{\delta}, \bar{p}, \bar{m}\bar{c}, h + \tau v)}{\partial p_1} & \frac{\partial g_2(\bar{\delta}, \bar{p}, \bar{m}\bar{c}, h + \tau v)}{\partial p_2} \end{array} \right]^{-1} \left[\begin{array}{c} \frac{\partial g_1(\bar{\delta}, \bar{p}, \bar{m}\bar{c}, h + \tau v)}{\partial \tau} \\ \frac{\partial g_2(\bar{\delta}, \bar{p}, \bar{m}\bar{c}, h + \tau v)}{\partial \tau} \end{array} \right] \Big|_{\tau=0} \equiv - (\mathbb{J}_g^p)^{-1} \mathbb{J}_g^\tau \Big|_{\tau=0} \quad (16)$$

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

$$Df_{p_1}(h)[v] = \sum_{\bar{\alpha}: |\bar{\alpha}| \leq 1} \sum_{j=1}^2 C_{\bar{\alpha}, j}(\bar{\delta}, \bar{m}\bar{c}, \{\partial^\beta h : |\beta| \leq 2\}) \times \partial^{\bar{\alpha}} v_j(h^{-1}(\bar{\delta}, \bar{p}), \bar{p}) \quad (17)$$

for real-valued functionals $C_{\bar{\alpha}, j}$. This shows that $Df_p(h_0)[v]$ is linear. Further, by the fundamental theorem of calculus, following an argument analogous to that in the proof of Theorem 2, we obtain

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

for all $h \in \mathcal{H}$. Therefore, Assumption 7(i) holds with $\alpha = [1, 1, 1, 1]$.

7(ii) As in the proof of Theorem 2, by the mean value theorem, we obtain

$$\begin{aligned} & \left| f_{p_1}(\hat{h}) - f_{p_1}(h_0) - Df_{p_1}(h_0)[\hat{h} - h_0] \right| \leq \\ & \sum_{\bar{\alpha}: |\bar{\alpha}| \leq 1} \sum_{j=1}^2 \left[C_{\bar{\alpha}, j}(\bar{\delta}, \bar{m}\bar{c}, \{\partial^\beta \hat{h} : |\beta| \leq 2\}) - C_{\bar{\alpha}, j}(\bar{\delta}, \bar{m}\bar{c}, \{\partial^\beta h_0 : |\beta| \leq 2\}) \right] \times \left\| \partial^{\bar{\alpha}} \hat{h}_j - \partial^{\bar{\alpha}} h_{0,j} \right\|_{1, \infty} \end{aligned}$$

Since, each of the $C_{\bar{\alpha}, j}(\bar{\delta}, \bar{m}\bar{c}, \{\partial^\beta \hat{h} : |\beta| \leq 2\}) - C_{\bar{\alpha}, j}(\bar{\delta}, \bar{m}\bar{c}, \{\partial^\beta h_0 : |\beta| \leq 2\})$ 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 7(ii) holds with $|\alpha_1| = 1, |\alpha_2| = 2$.

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

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

Thus, Assumption 7(iii) is implied by Assumption 9(v).

7(iv) By Remark 4.1 in CC, a sufficient condition for Assumption 7(iv) is

$$T_{iv,p} \equiv \frac{\tau_M \sqrt{\sum_{m=1}^M \left(Df_{p_1}(\hat{h}) \left[\left(G_\psi^{-1/2} \psi_M \right)_m \right] - Df_{p_1}(h_0) \left[\left(G_\psi^{-1/2} \psi_M \right)_m \right] \right)^2}}{v_T(f_{p_1})} = o_p(1) \quad (18)$$

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 $Df_{p_1}(\hat{h}) \left[\left(G_\psi^{-1/2} \psi_M \right)_m \right] - Df_{p_1}(h_0) \left[\left(G_\psi^{-1/2} \psi_M \right)_m \right]$ 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 $(G_\psi^{-1/2} \psi_M)_m$ of order at most 1. Therefore, using Corollary 3.1 in CC and the well-known rate results for splines and their derivatives in, e.g., Newey (1997), we can write

$$T_{iv,p} = O_p \left(\frac{\sqrt{T}}{v_T(f_{p_1})} \times \left[\frac{\tau_M M^{(10-r)/4}}{\sqrt{T}} + \frac{\tau_M^2 M^3 \sqrt{\log M}}{T} \right] \right)$$

The conclusion in (18) then follows from Assumption 9(v). □

Finally, the following lemmas provide more primitive sufficient conditions for Assumptions 9(iv) and 9(v). As for the price elasticity functional, we focus on the mildly ill-posed case and show that Assumptions 9(iv) and 9(v) are satisfied by letting M grow polynomially in T . We first consider the scenario in which the functional f_{p_1} is irregular.

Lemma 5. *Let Assumptions 9(i) and 9(iii) hold. Further, let $(v_T(f_\epsilon))^2 \asymp M^{a+\varsigma+1}$ (i.e. the functional f_{p_1} is irregular) and $\tau_M \asymp M^{\varsigma/2}$ for $a \leq 0, \varsigma \geq 0, a + \varsigma + 1 > 0, r + 2a - 8 > 0$. Then, Assumptions 9(iv) and 9(v) are satisfied if $M \asymp T^\rho$ with*

$$\rho \in \left(\frac{2}{r - 4 + 2(a + \varsigma + 1)}, \min \left\{ \frac{1}{\varsigma - a + 5}, \frac{\gamma^{(1)}}{2 + \gamma^{(1)}}, \frac{\gamma^{(2)}}{1 + \gamma^{(2)}} \right\} \right)$$

Further, M may be chosen to satisfy the latter condition if $r + 4a - 12 > 0$ and $\gamma^{(i)}(r + 2a + 2\varsigma - 4) - 4 > 0$ for $i \in \{1, 2\}$.

Proof. The result follows by inspection. □

Next, we consider the case in which f_{p_1} is regular.

Lemma 6. *Let Assumptions 9(i) hold. Further, let $v_T(f_\epsilon) = O(1)$ (i.e. the functional f_{p_1} is regular), $\tau_M \asymp M^{\varsigma/2}$ for $\varsigma \geq 0$, and let $h_{0,1}$ and $h_{0,2}$ belong to the Hölder ball of smoothness $r > 16 + 4\varsigma$ and finite radius. Then, Assumptions 9(iv) and 9(v) are satisfied if $M \asymp T^\rho$ with*

$$\rho \in \left(\frac{2}{r - 4}, \min \left\{ \frac{1}{2(\varsigma + 3)}, \frac{\gamma^{(1)}}{2 + \gamma^{(1)}}, \frac{\gamma^{(2)}}{1 + \gamma^{(2)}} \right\} \right)$$

Further, M may be chosen to satisfy the latter condition if $\gamma^{(i)}(r - 6) - 4 > 0$ for $i \in \{1, 2\}$.

Proof. The result follows by inspection. □