

Choice Models and Permutation Invariance: Demand Estimation in Differentiated Products Markets

Amandeep Singh^{*1}, Ye Liu^{†1}, and Hema Yoganarasimhan^{‡1}

¹University of Washington

February 15, 2024

Abstract

Choice modeling is at the core of understanding how changes to the competitive landscape affect consumer choices and reshape market equilibria. In this paper, we propose a fundamental characterization of choice functions that encompasses a wide variety of extant choice models. We demonstrate how non-parametric estimators like neural nets can easily approximate such functionals and overcome the curse of dimensionality that is inherent in the non-parametric estimation of choice functions. We demonstrate through extensive simulations that our proposed functionals can flexibly capture underlying consumer behavior in a completely data-driven fashion and outperform traditional parametric models. As demand settings often exhibit endogenous features, we extend our framework to incorporate estimation under endogenous features. Further, we also describe a formal inference procedure to construct valid confidence intervals on objects of interest like price elasticity. Finally, to assess the practical applicability of our estimator, we utilize a real-world dataset from [Berry et al. \(1995\)](#). Our empirical analysis confirms that the estimator generates realistic and comparable own- and cross-price elasticities that are consistent with the observations reported in the existing literature.

Keywords: Choice Models, Demand Estimation, Permutation Invariance, Set Functions, Neural networks

*Email: amdeep@uw.edu

†Email: yeliu@uw.edu

‡Email: hemay@uw.edu

1 Introduction

Demand estimation is a critical component in fields of marketing, operations, and economics, enabling practitioners to model consumer choice behavior and understand how consumers react to changes in a market. This understanding helps policymakers and businesses to make informed decisions, whether it be about launching new products, adjusting pricing strategies, or analyzing the consequences of mergers (Nevo, 2000; Petrin, 2002; Nevo, 2003; Wollmann, 2018). Over the years, various approaches, both parametric and non-parametric, have been developed to address the complexities inherent in demand estimation.

Parametric methods, based on logit or probit assumptions, often require strong assumptions about the underlying choice process, limiting their ability to capture the true complexity of consumer preferences. As such, the substantive and policy (counterfactual) implications from such models could be largely biased and/or misleading. Nevertheless, they have remained popular because of a few reasons – (1) simplicity, interpretability, and scalability to a large product set, (2) the ability to model counterfactual demand in situations involving new product introductions or mergers, and (3) the ability to handle endogenous product features.

Non-parametric methods, on the other hand, offer a more flexible approach to demand estimation, allowing for more nuanced representations of consumer preferences without restrictive assumptions about the underlying distributions of observables and unobservables. However, despite their potential advantages, non-parametric approaches face significant challenges that have limited their widespread adoption in practice. Firstly, a common limitation across all non-parametric demand estimation methods is the “curse of dimensionality,” where the computational complexity of estimating choice functions increases exponentially with the number of products. Additionally, a large subset of these models, specifically those that are ‘black box’ in nature and lack additional economic structure, are unable to perform counterfactual predictions, which are crucial for important tasks such as policy simulations. Furthermore, most non-parametric methods cannot properly handle endogenous product features. This inability to account for endogeneity and scalability constraints, along with the drawback of limited counterfactual analysis, significantly undermines their utility in practice, where such capabilities are often the primary objective of demand estimation.

In this paper, we make significant strides in bridging the gap between the flexibility of non-parametric methods and the tractability of parametric models. We achieve this by introducing a fundamental characterization of choice models. Our work begins by considering a broad set of choice functions, encompassing most existing choice modeling approaches in the empirical literature. We demonstrate that most choice functions exhibit specific symmetry properties. We then leverage recent advances in computer science and mathematics literature (for instance, Han et al. (2019); Zaheer et al. (2017); Wagstaff et al. (2019)) to characterize these functions. Our characterization allows us to build on the strengths of the non-parametric approaches (e.g., not assuming a model of consumer behavior) and overcome the challenges associated with them. First, it addresses the curse of dimensionality in choice systems and enables the flexible estimation of choice functions via non-parametric estimators. By leveraging the inherent permutation-invariant structure of choice models, our characterization can close the estimator’s generalization gap (for instance, see Sannai and Imaizumi (2019) in the context of neural networks) by a factor of $\sqrt{J!}$ (where J is the number of products). Second, we recognize that real-world demand systems often contain unobserved demand

shocks correlated with observable product features, such as prices. These shocks can lead to endogeneity issues, which can bias the estimated choice functions if not properly addressed. To tackle this challenge, we extend our framework to accommodate endogeneity. Third, our proposed approach successfully estimates counterfactual demand for scenarios that involve changes in the product set (e.g., new product introduction, mergers, product exit). Finally, we build upon recent advances in automatic debiased machine learning and provide an inference procedure for constructing valid confidence intervals on objects of interest, such as the average effect of price.

We demonstrate the effectiveness of our proposed approach using a series of numerical simulations. We consider a variety of data-generating processes in our simulations – multinomial logit model with linear utility, random coefficients logit with linear utility, random coefficient logit with non-linear utilities, and a setting where some consumers have inattention (i.e., ignore certain products in the market). Across all these scenarios, we show that our approach can predict market shares, own-, and cross-price elasticities with relatively high accuracy, even though we do not make any assumptions about consumer behavior. Further, even when the underlying DGP is complex, our model can recover market shares and elasticities similar to oracle estimators (that are assumed to know the true DGP). This allows researchers and managers to use our approach in general-purpose situations without making ad-hoc assumptions about consumer behavior.

Next, we consider counterfactual analyses and empirically show that our model can generate accurate counterfactuals in the case of new product introductions. Finally, we showcase the performance of the automatic debiasing procedure and show that we can provide consistent inference and confidence intervals over the average effect of price on demand across all the products. In sum, our extensive numerical simulations cover a wide range of DGPs and show that our approach – (1) is able to accurately predict market shares and price elasticities for a variety of scenarios, (2) can generate realistic counterfactual predictions in cases where the product set changes, and (3) can provide inference over economic objects of interest.

Finally, to showcase the effectiveness and applicability of our approach to real-life datasets, we use the [Berry et al. \(1995\)](#) automobile dataset and estimate the price elasticities using our non-parametric estimator (while correcting for price endogeneity). The results of our analysis align with existing literature and demonstrate the practical utility of our approach, underscoring its potential for adoption in real-world demand estimation settings. In summary, given the theoretical, practical, and computational advantages of our approach, we expect it to be easily applicable to a wide variety of demand estimation problems in both research and practice.

2 Literature Review

2.1 Parametric Discrete Choice Models

Discrete choice models play a crucial role in various fields, including economics, marketing, and operations management, as they describe decision-making processes when individuals face multiple alternatives. These models typically rely on a random utility maximization assumption, which can be traced back to [Thurstone \(1927\)](#) and [Marschak \(1959\)](#).

Over time, various choice models have emerged under different specifications for the density of unob-

served utility, following the general framework in [Marschak \(1959\)](#). The (Multinomial) logit model, first proposed by [Luce \(1959\)](#), is widely used for capturing systematic taste variance based on observed characteristics across alternatives. However, the logit model assumes that error terms are independent of each other and of the characteristics of the alternative. Additionally, the independent extreme value distribution results in the irrelevant alternatives (IIA) property, implying proportional substitution across alternatives.

To enable more flexible substitution patterns, Generalized Extreme Value (GEV) models were introduced, allowing for correlated unobserved utility across alternatives. The most commonly used GEV model is the nested logit model ([Train et al., 1987](#); [Forinash and Koppelman, 1993](#)). In nested logit models, the unobserved error of alternatives within a nest is specified as correlated, while the marginal distribution of each unobserved error remains as a univariate extreme value.

The mixed logit model ([McFadden and Train, 2000](#)), also known as random coefficient logit (RCL), is an even more versatile model that allows for randomness in both unobserved factors and coefficients of observed characteristics. The model was first applied by [Boyd and Mellman \(1980\)](#) and [Cardell and Dunbar \(1980\)](#), with random coefficients typically specified as normal or lognormal ([Ben-Akiva et al., 1993](#); [Mehndiratta, 1996](#); [Revelt and Train, 1998](#)) but also other distributions such as uniform and triangular ([Greene and Hensher, 2003](#); [Train, 2001](#)). When random coefficients follow a mixture distribution, the model becomes the well-known mixed logit latent class model.

Aside from logit family models (MNL, GEV, mixed logit) that have an extreme value distributed random component, the probit model [Hausman and Wise \(1978\)](#) assumes a jointly normal distribution for the error term. This allows for any pattern of substitution and can handle random taste variation. [Blanchet et al. \(2016\)](#) proposed a model where substitution between alternatives is a state transition in a Markov chain, which can approximate MNL, probit, and mixed logit models.

As the decision space grows, assuming that all alternatives are considered becomes unrealistic, leading to research in consideration sets and consumer search ([Honka et al., 2019](#); [Jiang et al., 2021](#)). These models typically impose strong parametric assumptions on the decision rule, including whether decisions are simultaneous or sequential, the stopping rule for searches, the size of the consideration set, and the functional form of the match value.

Further, real-world demand systems often contain unobserved demand shocks correlated with observable product features, such as prices. These shocks can lead to endogeneity issues, which can bias the estimated parameters if not properly addressed. Thus, various approaches have been proposed to address these issues. For instance, [Berry et al. \(1995\)](#) proposed a generalized method of moments-based estimator to estimate a random-coefficients logit model of demand using instrumental variables. Similarly, [Petrin and Train \(2010\)](#) demonstrated the application of control functions to resolve endogeneity in the random coefficient logit model of demand.

Two primary concerns related to discrete choice models are the need for correct model specification and the distribution of unobserved factors. To address these issues, nonparametric demand estimation models have been developed.

2.2 Nonparametric Demand Estimation

While parametric discrete choice models make assumptions about the distribution of unobserved variables and specify functional forms for utilities, recent research has developed more flexible semi-parametric and nonparametric approaches for demand estimation. These newer methods ease some of the restrictive assumptions and improve the computational efficiency of choice models while retaining some structure.

Early semi-parametric work (Manski, 1987; Lewbel, 2000; Honoré and Kyriazidou, 2000; Abrevaya, 2000) focused on relaxing the distribution of random shocks in individual-level binary choice models. More recent research (Khan et al., 2021; Shi et al., 2018; Pakes and Porter, 2022) has extended this approach to relax the parametric assumptions on random shocks in individual-level multinomial choice models, allowing for increased flexibility, such as individual-level fixed effects. Fox and Gandhi (2016); Briesch et al. (2010); Allen and Rehbeck (2019); Fosgerau and Kristensen (2021); Chitla et al. (2022); Lu et al. (2023); Wang (2023) concentrate on nonparametric identification and estimation of distributions of heterogeneous unobservables, like random coefficients, in various demand models, relaxing assumptions about heterogeneity distribution. More recent work has looked at keeping the indirect utility or choice probability functions largely unspecified while retaining parametric error terms. For instance (Bentz and Merunka, 2000; Wang et al., 2020; Han et al., 2022; Sifringer et al., 2020; Wong and Farooq, 2021; Aouad and Désir, 2022) use neural networks to characterize the indirect utility while retaining the logit error structure. Our work, in comparison, provides a completely flexible mapping from observed product and consumer characteristics to observed demand without relying on any assumptions regarding the choice making process.

Extant research in nonparametric methods aims to completely avoid any parametric assumptions to remove any potential source of misspecification. Berry and Haile (2014, 2020) presented the identification results for nonparametric estimation of demand from market-level and individual-level data, respectively. Studies by Hausman and Newey (2016); Blundell et al. (2017); Chen and Christensen (2018) focus on individual-level data. In particular, Hausman and Newey (2016) employs a nonparametric approach to estimate consumer surplus bounds, while Blundell et al. (2017) introduces a method for consistently estimating demand functions with nonseparable unobserved taste heterogeneity, subject to the shape restriction imposed by the Slutsky inequality. Chen and Christensen (2018) concentrates on nonparametric instrumental variables and inference in individual-level data. This paper, alongside Compiani (2022) and Tebaldi et al. (2023), examines market-level data. Compiani (2022) develops a nonparametric method based on Bernstein polynomials, drawing on the identification result of Berry and Haile (2014). Meanwhile, Tebaldi et al. (2023) proposes a technique for deriving nonparametric bounds on demand counterfactuals and applies it to California’s health insurance market.

The common challenge across all extant nonparametric demand estimation work has been that the computational complexity of estimating nonparametric demand models increases exponentially with the number of products. For instance, Compiani (2022) could estimate their nonparametric model for just two products. Similarly, Cai et al. (2022) needed data on more than 100,000 markets to reasonably estimate demand with 50 products. To put this in perspective, economic datasets are much smaller; for instance, the dataset in Berry et al. (1995) had 150 products across only 20 markets. This curse of dimensionality has been the primary obstacle preventing the widespread adoption of nonparametric methods in demand estimation. To

overcome this barrier, our work demonstrates how to exploit the inherent permutation invariant structure of choice functions to break this curse of dimensionality and flexibly estimate demand in markets with large a number of products. In addition, our work is also related to recent work in marketing [Wei and Jiang \(2022\)](#) exploring the use of neural networks to estimate parameters of structural models.

Finally, our work leverages recent work in mathematics and computer science (for instance, [Han et al. \(2019\)](#); [Zaheer et al. \(2017\)](#); [Wagstaff et al. \(2019\)](#)), which investigate the universal approximation of symmetric and antisymmetric functions, offering fundamental characterizations for functions defined on sets. We build on this literature to characterize a general class of choice functions in demand systems.

3 Theory

3.1 Choice Models

In this section, we provide a general characterization of consumer choice functions. In particular, we focus on a scenario where researchers have access only to aggregate market-level demand data, while individual-level choices and characteristics remain unobserved. Aggregate demand models have been extensively studied in marketing and economics ([Berry et al., 1995](#); [Besanko et al., 1998](#); [Sudhir, 2001](#); [Chintagunta, 2001](#); [Albuquerque and Bronnenberg, 2009](#); [Compiani, 2022](#)), and are particularly useful when individual-level demand data is not available.

Suppose consumers in a market t face an offer set \mathcal{S}_t that can comprise any subset of J_t distinct products ($\{1, 2, \dots, J_t\}$). We use u_{ijt} to represent the index tuple $\{X_{jt}, p_{jt}, I_{it}, \varepsilon_{ijt}\}$, where $X_{jt} \in \mathbb{C}^d$ denotes d non-price features belonging to some countable universe \mathbb{C}^d ; $p_{jt} \in \mathbb{C}$ denotes the price of the product; $I_{it} \in \mathbb{C}^l$ denotes demographics of consumer i in market t , we assume there are l features and belong to some countable universe \mathbb{C}^l , and ε_{ijt} denotes random idiosyncratic components pertinent to consumer i for product j in market t that are not unobservable to the researcher but observable to consumers.

Definition 1 (Choice Function). Given the offer set $\mathcal{S}_t \subset \{1, 2, 3, \dots, J_t\}$, we define a function $\pi : \{u_{ijt} : j \in \mathcal{S}_t\} \rightarrow \mathbb{R}^{|\mathcal{S}_t|}$ that maps a set of index tuples $\{u_{ijt}\}_{j \in \mathcal{S}_t}$ to a $|\mathcal{S}_t|$ -dimensional probability vector. Each element in the $\pi(\cdot)$ vector represents the probability of consumer i choosing product j in market t .

Here we present a very general characterization of choice functions that maps the observable and unobservable components of product and individual characteristics to observed choices through some choice function π . Note that, traditionally, u_{ijt} is a scalar that represents utility in choice models. However, in our framework, u_{ijt} does not necessarily represent utility. Further, we have not yet imposed any assumption on π , i.e., how consumers make choices.

We now specify a set of assumptions on the model and data-generating process below.

Assumption 1 (Exogeneity). The unobserved error term ε_{ijt} is independent and identically distributed (i.i.d.) across all products. This can be expressed as follows:

$$\mathbb{P}(\varepsilon_{ijt} \mid X_{.t}, p_{.t}) = \mathbb{P}(\varepsilon_{ijt})$$

This assumption implies that the error term ε_{ijt} is not correlated with any of the observed variables $X_{\cdot t}$ and $p_{\cdot t}$. As such, it precludes the possibility of endogenous prices and/or marketing-mix variables, as is common in observational data. We start with the basic case with exogenous covariates in this section and later in §3.2, we relax this assumption and allow for endogenous covariates.

Assumption 2 (Identity Independence). For any product $j \in \mathcal{S}_t$ and any market t , we assume the choice function π does not depend on the identity of the product (jt). That is:

$$\pi_{ijt}(\{u_{ikt}\}_{k \in \mathcal{S}_t}) = \pi_{ijt}(u_{ijt}, \{u_{ikt}\}_{k \in \mathcal{S}_t, k \neq j}) = \pi(u_{ijt}, \{u_{ikt}\}_{k \in \mathcal{S}_t, k \neq j})$$

This assumption implies two things: first, the functional form of the choice probability for different products and markets is the same; second, for any market-level heterogeneity (e.g., in the distribution $F_t(I_{it}, \varepsilon_{ijt})$), we can include them in X_{jt} as features. Intuitively, this assumption suggests that conditional on product and consumer features and the unobserved error term, the choice probabilities are not functions of the identities of the products themselves.

Assumption 3 (Permutation Invariance). The choice function π is invariant under any permutation function $\sigma_j(\cdot)$ that rearranges the indices of the competitors of product j , such that:

$$\pi_{ijt} = \pi(u_{ijt}, \{u_{i\sigma_j(k)t}\}_{k \in \mathcal{S}_t, k \neq j})$$

In this assumption, we state that the choice function for product j is invariant to all permutations of its competitors. This implies that the individual's choice for product j is not affected by the order or identity of the other products in the market, and it only depends on the set of competitors' characteristics.

Since researchers only observe aggregate data, we next define the aggregate demand function. In aggregate demand settings, individual-level choices are not observable and only aggregate demand is observable. It is often the case that the market-specific individual features are not observable and are assumed to be exogenously drawn from some distribution $\mathcal{F}(m_t)$, where m_t represents the market-level characteristics. For the sake of notional simplicity, we let m_t to be the same across all markets. One can easily incorporate market-specific user demographics in the choice function. Thus the demand of product j in market t denoted by π_{jt} can be expressed as follows:

$$\pi_{jt} = \int \int \pi_{ijt}(\{u_{ikt}\}_{k \in \mathcal{S}_t}) d\mathcal{F}(m_t) d\mathcal{G}(\varepsilon_{ijt}), \quad (1)$$

where $\mathcal{G}(\varepsilon_{ijt})$ denotes the CDF of unobserved errors ε_{ijt} . Since u_{ijt} is determined by $\{X_{jt}, p_{jt}, I_{it}, \varepsilon_{ijt}\}$ and $I_{it}, \varepsilon_{ijt}$ are integrated out in a market. Hence, we can express π_{jt} as a function of only the observable product characteristics –

$$\pi_{jt} = g(X_{jt}, p_{jt}, \{X_{kt}, p_{kt}\}_{k \in \mathcal{S}, k \neq j}). \quad (2)$$

Lemma 1. *For any choice function that satisfies Assumption 1 and 3, the aggregate demand function is permutation invariant.*

This permutation invariance of the aggregate demand function exists because, under the exogeneity

assumption, the aggregate demand function is simply the sum (or integral) of individual choice functions that are themselves invariant to permutation. Hence, changes to the order of competitors have no impact on the aggregated result. When the assumption of exogeneity is not satisfied, the aggregate demand function does not retain the permutation invariance, even though the individual-level choice function exhibits permutation invariance. We will return to this issue in §3.2.

Our Assumptions 2 (identity independence) and 3 (permutation invariance) are fairly standard in the choice modeling literature, although they might not always be explicitly stated as such. Table 1 summarizes models that satisfy these assumptions. Please see Web Appendix A for detailed derivations of how these models satisfy permutation invariance.

Table 1: Choice Models Satisfying Identity Independence and Permutation Invariance

Choice Model	Literature
Multinomial Logit Model	McFadden et al. (1973)
Mixed Logit Model	McFadden and Train (2000)
Nested Logit Model	Train et al. (1987)
Random Coefficients Nested Logit	Grigolon and Verboven (2014)
Generalized Extreme Value (GEV) Model	Train (2009)
Probit Model	Hausman and Wise (1978)
Latent Class Logit Model	Kamakura and Russell (1989)
Markov Chain Choice Model	Blanchet et al. (2016)
Customer Inattention Based Models	Goeree (2008) , Turlo et al. (2023) , Compiani (2022) and Joo (2023)
Customer Search Models	Mehta et al. (2003)

Theorem 1. For any offer set $\mathcal{S}_t \subset \{1, 2, 3, \dots, J_t\}$, if a choice function $\pi : \{u_{ijt} : j \in \mathcal{S}_t\} \rightarrow \mathbb{R}^{|\mathcal{S}_t|}$ where u_{ijt} represents the index tuple $\{X_{jt}, p_{jt}, I_{it}, \varepsilon_{ijt}\}$ satisfies Assumption 1, 2 and 3, then there exists suitable ρ , ϕ_1 and ϕ_2 such that

$$\pi_{jt} = \rho(\phi_1(X_{jt}, p_{jt}) + \sum_{k \neq j, k \in \mathcal{S}_t} \phi_2(X_{kt}, p_{kt})),$$

Proof: See Web Appendix B.

This result is the generalization of the results shown in [Zaheer et al. \(2017\)](#) and can be shown following similar arguments. The above result is very powerful and has two important takeaways: (i) The input space of the choice function does not grow with the number of products in the assortment. Rather, the input space of the choice function (i.e., ϕ_1 and ϕ_2) grows only as a function of the number of features of the products in consideration, and (ii) the same transformations (ρ , ϕ_1 , and ϕ_2) remain valid for all offer sets, denoted by \mathcal{S} , irrespective of their size. This property allows us to easily simulate the demand and entry of new products or changes in market structure, as one does with traditional parametric models. As an example, assuming v_{jt} is the utility of product j at market t , for the multinomial logit model one possible set of transformations could be $\phi_1(v_{jt}) = \begin{bmatrix} \exp(v_{jt}) \\ 0 \end{bmatrix}$ and $\phi_2(v_{kt}) = \begin{bmatrix} 0 \\ \exp(v_{kt}) \end{bmatrix}$ that generate two-dimensional vectors, and the

function $\rho \left(\begin{bmatrix} \phi_1(v_{jt}) \\ \sum_{k \neq j, k \in \mathcal{S}_t} \phi_2(v_{kt}) \end{bmatrix} \right) = \frac{\phi_1(v_{jt})}{\phi_1(v_{jt}) + \sum_{k \neq j, k \in \mathcal{S}_t} \phi_2(v_{kt})}$ operates on these vectors. ¹

3.2 Endogenous Covariates

In this section, we relax the exogeneity assumption and handle the potential endogeneity issue that is commonplace in demand settings. Note that, when the price (or other product characteristics or market-mix variables, such as promotions, correlate with unobserved variables (ε_{ijt}), Assumption 1 (exogeneity) is compromised. As a result, it becomes infeasible to integrate out ε_{ijt} in the aggregate demand function, as we did in Equation (1). This means that the aggregate demand function loses its property of permutation invariance with respect to the observable characteristics of competitors. To address this, we build on the approach developed in [Petrin and Train \(2010\)](#) to allow for endogenous observable features. Without loss of generality, we assume that price $p_{jt} \in \mathbb{C}$ is the endogenous variable and all other characteristics of the product $X_{jt} \in \mathbb{C}^d$ are exogenous variables. i.e.,

$$\mathbb{E}[p_{jt} \cdot \varepsilon_{ijt}] \neq 0 \quad \text{and} \quad \mathbb{E}[X_{jt} \cdot \varepsilon_{ijt}] = 0.$$

Given valid instruments IV_{jt} , we can express p_{jt} as

$$p_{jt} = \gamma(X_{jt}, IV_{jt}) + \mu_{jt}. \quad (3)$$

At this point, no specific assumptions are made regarding the function γ . However, in the subsequent inference section, we will discuss that the estimator of γ must be estimable at $n^{-1/2}$ in order to construct valid confidence intervals. Next, to address the issue of price endogeneity, we impose a mild restriction on the space of choice functions we consider.

Assumption 4 (Linear Separability). The unobserved product characteristics can be expressed as the sum of an endogenous (CF) and exogenous component

$$\varepsilon_{ijt} = CF(\mu_{jt}; \lambda) + \tilde{\varepsilon}_{ijt}, \quad (4)$$

where $\mathbb{E}[p_{jt} \cdot \tilde{\varepsilon}_{ijt}] = 0$.

This assumption implies that, after controlling for μ_{jt} using the control function CF , the endogenous variable p_{jt} is uncorrelated with the error term ε_{ijt} in the model, thus it becomes exogenous. Then, we can re-write the index tuple u_{ijt} as

$$u_{ijt} = \{X_{jt}, p_{jt}, CF(\mu_{jt}; \lambda) + \tilde{\varepsilon}_{ijt}\}, \quad (5)$$

such that $\mathbb{E}[\tilde{\varepsilon}_{ijt} | (X_{jt}, p_{jt}, \mu_{jt})] = 0$

$$\frac{\rho \left(\begin{bmatrix} \phi_1(v_{jt}) \\ \sum_{k \neq j} \phi_2(v_{kt}) \end{bmatrix} \right)}{\phi_1(v_{jt}) + \sum_{k \neq j} \phi_2(v_{kt})} = \rho \left(\begin{bmatrix} \exp(v_{jt}) \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ \sum_{k \neq j} \exp(v_{kt}) \end{bmatrix} \right) = \rho \left(\begin{bmatrix} \exp(v_{jt}) \\ \sum_{k \neq j} \exp(v_{kt}) \end{bmatrix} \right) =$$

Theorem 2. For any offer set $\mathcal{S}_t \subset \{1, 2, 3, \dots, J_t\}$, if a choice function $\pi : \{u_{ijt} : j \in \mathcal{S}_t\} \rightarrow \mathbb{R}^{|\mathcal{S}_t|}$ where u_{ijt} represents the index tuple $\{X_{jt}, p_{jt}, I_{it}, \varepsilon_{ijt}\}$ satisfies Assumptions 2 to 4. Then under the condition of knowing the true function (γ_0) of γ , there exists suitable ρ, ϕ_1 and ϕ_2 such that

$$\pi_{jt} = \rho(\phi_1(X_{jt}, p_{jt}, \mu_{jt}(\gamma_0)) + \sum_{k \neq j, k \in \mathcal{S}} \phi_2(X_{kt}, p_{jt}, \mu_{kt}(\gamma_0))),$$

The result follows straightforwardly from the observation that after controlling for $CF(\mu_{jt}; \lambda)$ the unobservable component $\tilde{\varepsilon}$ is exogenous. This implies the aggregate demand function is invariant under any permutation applied to the competitors of product j . The result demonstrates that endogeneity can be addressed by using the residuals from Equation (4) as an additional set of features along with observable product characteristics.

3.3 Inference

This paper aims to estimate choice functions flexibly using non-parametric estimators. However, often in social science contexts, researchers and managers are also interested in conducting inference over some economic objects. Note that because non-parametric regression functions are estimated at a slower rate compared to parametric regressions, it is often infeasible to construct confidence intervals directly on the estimated $\hat{\pi}$. However, it is generally possible to perform inference and construct valid confidence intervals for specific economic objects that are functions of π . In this section, we will provide an example of one such important economic object and demonstrate how to construct valid confidence intervals for it. This will be done by leveraging the recent advances in automatic debiased machine learning as shown in the works of [Ichimura and Newey \(2022\)](#); [Chernozhukov et al. \(2022b,a, 2021\)](#), and others. However, unlike existing automatic debiased machine learning setups we also have to account for an additional first-stage estimator $\hat{\gamma}$.

In demand estimation, researchers are often interested in estimating the average effect of a price change on the demand for a product, as it can significantly influence market dynamics, pricing strategies, and regulatory decisions. To proceed with our analysis, let $w_{jt} = (y_{jt}, p_{jt}, X_{jt}, \{p_{kt}, X_{kt}\}_{k \neq j})$ and $z_{jt} = (p_{jt}, X_{jt}, \{p_{kt}, X_{kt}\}_{k \neq j})$ represent the variables associated with product j in market t . Here, $p_{jt} \in \mathbb{C}$ denotes the observed prices, $X_{jt} \in \mathbb{C}^{d-1}$ represents other product characteristics, and $y_{jt} \in \mathbb{R}$ refers to the observed demand for product j in market t , such as market shares or log shares. Note that either the observed price (p_{jt}) or other characteristics (X_{jt}) could be endogenous. For simplicity and without loss of generality, we focus on p_{jt} as the endogenous variable in the following analysis.

The average effect of a price change² can be expressed as the difference between the demand function $\pi_{jt}(\cdot; \gamma)$ evaluated at the original price p_{jt} and at the price incremented by Δp_{jt} , given by the following expression:

$$m(w_{jt}, \pi(\cdot; \gamma)) = \pi(p_{jt} + \Delta p_{jt}, X_{jt}, \{p_{kt}, X_{kt}\}_{k \neq j}; \gamma) - \pi(p_{jt}, X_{jt}, \{p_{kt}, X_{kt}\}_{k \neq j}; \gamma).$$

²The expression for the average effect of a price change can be adapted to represent average price elasticity by placing the known and fixed value of Δp_{jt} in the denominator.

The parameter of interest, θ_0 , is the expected value of this price change effect over the true population distribution³ of w_{jt} , which can be calculated as:

$$\theta_0 = \mathbb{E}[m(w_{jt}, \pi(\cdot; \gamma))] = \mathbb{E}[\pi(p_{jt} + \Delta p_{jt}, X_{jt}, \{X_{kt}\}_{k \neq j}; \gamma) - \pi(p_{jt}, X_{jt}, \{X_{kt}\}_{k \neq j}; \gamma)].$$

In summary, the average effect of a price change on demand, denoted by θ_0 , is calculated by evaluating the difference between the demand function at the original price and at the price incremented by Δp_{jt} , and then computing the expected value of this difference.

In practice, we estimate θ_0 by computing its empirical analog using the estimated demand function $\hat{\pi}$ and first-stage estimator $\hat{\gamma}$, i.e.,

$$\hat{\theta} = \frac{1}{n} \sum_{t=1}^n m(w_{jt}, \hat{\pi}(z_{jt}; \hat{\gamma})), \quad (6)$$

where n is the number of observations. When parametric methods are employed to estimate $\hat{\pi}$ and $\hat{\gamma}$, the estimator for $\hat{\theta}$ is generally \sqrt{n} -consistent, assuming that the model is correctly specified. However, \sqrt{n} -consistency may not hold when non-parametric estimators are used, particularly if the first-order bias does not vanish at a rate of \sqrt{n} . Irrespective of the method used to estimate π , this is often the case, as flexible estimation of π always requires some form of regularization and/or model selection. Debiasing techniques are required to mitigate the effects of regularization and/or model selection when learning flexible demand models. These approaches can help improve the performance of the estimator and facilitate valid inference with $\hat{\theta}$. We therefore adapt recent debiasing techniques developed in recent automatic debiased machine learning literature (see [Chernozhukov et al. \(2022b\)](#)). Specifically, we will focus on problems where there exists a square-integrable random variable $\alpha_0(z)$ such that $\forall \|\gamma - \gamma_0\|$ small enough –

$$\begin{aligned} \mathbb{E}[m(w_{jt}, \pi(z_{jt}; \gamma))] &= \mathbb{E}[\alpha_0(z_{jt})\pi(z_{jt}; \gamma)] \\ \forall \pi \text{ with } \mathbb{E}[\pi_{jt}(z_{jt}; \gamma)^2] &< \infty \end{aligned}$$

By the Riesz representation theorem, the existence of such $\alpha_0(z_{jt})$ is equivalent to $\mathbb{E}[m(w_{jt}, \pi(z_{jt}; \gamma))]$ being a mean square continuous functional of $\pi(z_{jt}; \gamma)$. Henceforth, we refer to $\alpha_0(z)$ as Riesz representer (or RR). [Newey \(1994\)](#) shows that the mean square continuity of $\mathbb{E}[m(w_{jt}, \pi_{jt}(z_{jt}; \gamma))]$ is equivalent to the semiparametric efficiency bound of θ_0 being finite. Thus, our approach focuses on regular functionals. Similar uses of the Riesz representation theorem can be found in [Ai and Chen \(2007\)](#), [Ackerberg et al. \(2014\)](#), [Hirshberg and Wager \(2020\)](#), and [Chernozhukov et al. \(2022b\)](#) among others. The debiasing term in this case takes the form $\alpha(z_{jt})(y_{jt} - \pi(z_{jt}; \gamma))$. To see that, consider the score $m(w_{jt}, \pi(z_{jt}; \gamma)) + \alpha(z_{jt})(y_{jt} - \pi(z_{jt}; \gamma)) - \theta_0$. It satisfies the following mixed bias property:

$$\begin{aligned} \mathbb{E}[m(w_{jt}, \pi(z_{jt}; \gamma)) + \alpha(z_{jt})(y_{jt} - \pi(z_{jt}; \gamma)) - \theta_0] \\ = -\mathbb{E}[(\alpha(z_{jt}) - \alpha_0(z_{jt}))(\pi(z_{jt}) - y_{jt})]. \end{aligned}$$

This property implies double robustness ([Robins et al., 1994](#); [Funk et al., 2011](#)) of the score. That is, if either

³We assume the data reflects the true population.

$\alpha(z_{jt})$ is correctly estimated, which would mean $\alpha(z_{jt}) - \alpha_0(z_{jt}) = 0$, or $\pi(z_{jt})$ is correctly estimated, implying $\pi(z_{jt}) - y_{jt} = 0$, then the term $(\alpha(z_{jt}) - \alpha_0(z_{jt}))(\pi(z_{jt}) - y_{jt})$ will be zero. This results in the score going to zero, thereby making the estimator consistent for θ_0 . A debiased machine learning estimator of θ_0 can be constructed from this score and first-stage learners $\hat{\pi}$ and $\hat{\alpha}$. Let $\mathbb{E}_n[\cdot]$ denote the empirical expectation over a sample of size n , i.e., $\mathbb{E}_n[x_i] = \frac{1}{n} \sum_{i=1}^n x_i$. We consider:

$$\hat{\theta} = \mathbb{E}_n[m(w_{jt}; \hat{\pi}) + \hat{\alpha}(z_{jt})(y_{jt} - \hat{\pi}(z_{jt}))].$$

The mixed bias property implies that the bias of this estimator will vanish at a rate equal to the product of the mean-square convergence rates of $\hat{\alpha}$ and $\hat{\pi}$. Therefore, in cases where the demand function π can be estimated very well, the rate requirements on $\hat{\alpha}$ will be less strict, and vice versa. More notably, whenever the product of the mean-square convergence rates of $\hat{\alpha}$ and \hat{f} is larger than \sqrt{n} , we have that $\sqrt{n}(\hat{\theta} - \theta_0)$ converges in distribution to centered normal law $N(0, \mathbb{E}[\psi_0(w_{jt})^2])$, where

$$\psi_0(w_{jt}) := m(w_{jt}; \pi_0) + \alpha_0(z_{jt})(y_{jt} - \pi_0(z_{jt})) - \theta_0,$$

as proven formally in Theorem 3 of [Chernozhukov et al. \(2022b\)](#). Results in [Newey \(1994\)](#) imply that $\mathbb{E}[\psi_0(w_i)^2]$ is a semiparametric efficient variance bound for θ_0 , and therefore the estimator achieves this bound.

Theorem 3. [[Chernozhukov et al. \(2021\)](#)] One can view the Riesz representer as the minimizer of the loss function:

$$\begin{aligned} \alpha_0 &= \arg \min_{\alpha} \mathbb{E} [(\alpha(z_{jt}) - \alpha_0(z_{jt}))^2] \\ &= \arg \min_{\alpha} \mathbb{E} [\alpha(z_{jt})^2 - 2\alpha_0(z_{jt})\alpha(z_{jt}) + \alpha_0(z_{jt})^2] \\ &= \arg \min_{\alpha} \mathbb{E} [\alpha(z_{jt})^2 - 2m(w_{jt}; \alpha)], \end{aligned}$$

In our earlier discussions, we employed the moment function of π , whereas in Theorem 3, we focus on the moment function of α . This shift is justified by the Riesz Representation Theorem, which implies $\mathbb{E}[m(w_{jt}; \pi)] = \mathbb{E}[\alpha_0(z_{jt})\pi(z_{jt})]$. Given that π can represent any function, substituting α for π is permissible, thereby validating the transition from the second to the third line in Theorem 3. We use the above theorem to flexibly estimate the RR. The advantage of this approach is that it eliminates the need to derive an analytical form for the RR estimator, allowing it to be addressed as a simple computational optimization problem.

Theorem 4. [[Chernozhukov et al. \(2021\)](#)] Let δ_n be an upper bound on the critical radius ([Wainwright \(2019\)](#)) of the function spaces:

$$\begin{aligned} &\{z \mapsto \zeta(\alpha(z) - \alpha_0(z)) : \alpha \in \mathcal{A}_n, \zeta \in [0, 1]\} \text{ and} \\ &\{w \mapsto \zeta(m(w; \alpha) - m(w; \alpha_0)) : \alpha \in \mathcal{A}_n, \zeta \in [0, 1]\} \end{aligned}$$

and suppose that for all f in the spaces above: $\|f\|_{\infty} \leq 1$. Suppose, furthermore, that m satisfies the

mean-squared continuity property:

$$\mathbb{E} \left[(m(w; \alpha) - m(w; \alpha'))^2 \right] \leq M \|\alpha - \alpha'\|_2^2$$

for all $\alpha, \alpha' \in \mathcal{A}_n$ and some $M \geq 1$. Then for some universal constant C , we have that w.p. $1 - \zeta$:

$$\|\hat{\alpha} - \alpha_0\|_2^2 \leq C \left(\delta_n^2 M + \frac{M \log(1/\zeta)}{n} + \inf_{\alpha_* \in \mathcal{A}_n} \|\alpha_* - \alpha_0\|_2^2 \right)$$

The critical radius has been widely studied in various function spaces, such as high-dimensional linear functions, neural networks, and superficial regression trees, often showing $\delta_n = O(d_n n^{-1/2})$, where d_n represents the effective dimensions of the hypothesis spaces (Chernozhukov et al. (2021)). In our research, we focus on applying Theorem 3 from an application standpoint to neural networks.

To that end, we make the following assumptions.

Assumption 5. (i) $\alpha_0(z)$ is bounded, (ii) $\forall \|\gamma - \gamma_0\|$ small enough, $\mathbb{E}[(y - \pi_0(z_{jt}; \gamma))^2 | z_{jt}]$ is bounded, and (iii) $\mathbb{E}[m(w_{jt}, \pi_0(z_{jt}; \gamma_0))^2] < \infty$.

These assumptions are standard regularity conditions used in the automatic machine learning literature.

Assumption 6. i) $\forall \|\gamma - \gamma_0\|$ small enough $\|\hat{\pi}(\cdot; \gamma) - \pi_0(\cdot; \gamma)\| \xrightarrow{p} 0$ and $\|\hat{\alpha} - \alpha_0\| \xrightarrow{p} 0$; ii) $\sqrt{n} \|\hat{\alpha} - \alpha_0\| \xrightarrow{p} 0$; iii) $\hat{\alpha}$ is bounded; (iv) $\sqrt{n} \|\hat{\gamma} - \gamma_0\| \xrightarrow{p} 0$

Intuitively these assumptions mean that (i) the estimator of both π and α should be consistent for values of γ in a close enough neighborhood of γ_0 . Further, it requires that the product of mean square error of $\hat{\alpha}$ and mean square error of π should vanish at \sqrt{n} -rate. This can be achieved if both these terms converge at least at $n^{-1/4}$ rate. Finally, we also assume that the first stage estimator $\hat{\gamma}$ is estimable at $n^{-1/2}$ rate. This limits the class of functions one can use to estimate γ .

Assumption 7. $m(w, \pi)$ is linear in π and there is $C > 0$ such that

$$|E[m(w, \pi) - \theta_0 + \alpha_0(z)(y - \pi(z; \gamma))]| \leq C \|\pi - \pi_0\|^2$$

Proposition 1. *If Assumptions 5-7 are satisfied then for $V = E[\{m(w, \pi_0(z; \gamma_0)) - \theta_0 + \alpha_0(z)(y - \pi_0(z; \gamma_0))\}^2]$,*

$$\sqrt{n}(\hat{\theta} - \theta_0) \xrightarrow{D} N(0, V), \hat{V} \xrightarrow{p} V.$$

We show the proof in Web Appendix B. This theorem shows that if $\hat{\gamma}$ is estimable at a fast enough rate one can still construct valid confidence intervals for $\hat{\theta}$. This result can be shown following similar arguments as in Chernozhukov et al. (2022a). Finally, we note that while the above arguments focus on the estimation of the average effect of a price change on demand, we can follow the same arguments to derive inference results for other economic quantities of interest, e.g., the effect of changing some product features on demand.

4 Estimation Procedure

Based on theoretical results presented above, we now outline an estimation procedure for both the choice function (π) and the average effects of price changes (θ).

Consider a dataset, where $\{y_t, z_t, IV_t\}_{t=1}^n$ are independently and identically distributed. Here, y_t is the vector of market shares in market t , z_t is $J \times (d + 1)$ matrix of product features and IV_t is the J -dimension vector of instrumental variables.

- Stage 0 (Data partition): We randomly split the observed markets into L folds such that the data $D_l := \{y_t, z_t\}_{t \in l}$, where l denotes the l^{th} partition. Note that all the observations for one market are always in one fold.
- Stage 1 (Estimate $\hat{\gamma}$): For each fold l , we estimate γ_l by regressing the endogenous variable on the exogenous instruments on the left out data $D_l^c := \{y_t, z_t\}_{t \notin l}$. We then use the cross-fitting technique, same as [Chernozhukov et al. \(2021\)](#), to calculate the residual $\hat{\mu}_l$ of fold l with estimated $\hat{\gamma}_l$ on D_l^c .
- Stage 2 (Estimate $\hat{\pi}$ and $\hat{\theta}$):
 - Stage 2a (Estimation): In the second stage, for each fold l , we estimate both the choice function ($\hat{\pi}$) and the Riesz estimator ($\hat{\alpha}$) on the left out data $D_l^c := \{y_t, z_t\}_{t \notin l}$

$$\hat{\pi}_l = \arg \min_{\pi \in \mathcal{F}} \frac{1}{\sum_{t \in D_l^c} J_t} \sum_{t \in D_l^c} \sum_{j \in J_t} [(y_{jt} - \pi(z_{jt}; \hat{\gamma}))^2] \quad (7)$$

$$\hat{\alpha}_l = \arg \min_{\alpha \in \mathcal{A}} \frac{1}{\sum_{t \in D_l^c} J_t} \sum_{t \in D_l^c} \sum_{j \in J_t} [\alpha(z_{jt})^2 - 2m(w_{jt}; \alpha)]. \quad (8)$$

Based on Theorem 2, instead of directly estimating the function π , we decompose the estimation into three sub-components: ρ , ϕ_1 , and ϕ_2 . Specifically, for each component of our model (ϕ_1 , ϕ_2 , and ρ), we use a standard 3-layer neural network⁴, and this is implemented without further hyperparameter tuning. We implement ReLU activation function at each layer as it is standard in feedforward designs due to simplicity and computation efficiency in gradients. Figure 1 illustrates how we pass the data to the neural networks. Specifically, we pass the focal product's characteristics (price p_{jt} and other product features X_{jt}) and the residuals ($\hat{\mu}_{jt}$) estimated from the first stage regression to the ϕ_1 . In parallel, we pass all the products' characteristics of the other products in the same market (p_{kt} and X_{kt}) and the corresponding residuals ($\hat{\mu}_{kt}$) to the same ϕ_2 , and then sum the output up. The output of ϕ_1 and ϕ_2 have the same data structure (e.g., a 64-dimension vector). Next, we pass the summation of the output of ϕ_1 and ϕ_2 to a third neural network ρ . The output of ρ is a scalar which represents the market share of the focal product jt .

⁴For ϕ_1 and ϕ_2 , each of the three layers consists of 64 neurons, with the output vector also featuring 64 neurons. For ρ , the layers are configured with 300, 100, and 64 neurons for the first, second, and third layers, respectively.




Figure 1: Illustration of Neural Network Architecture

We use the same neural network structure (with the three sub-components same as ϕ_1 , ϕ_2 , and ρ) to estimate α . The only difference is that the loss function of α is not based on the difference between the observed and the predicted market share as in Equation 7. Instead, the loss function is based on the squared difference between α and the moment function of α as stated in Equation 8 and Theorem 3.

- Stage 2b (Cross-fitting): Now we again use the cross-fitting technique to reduce the bias when estimating $\hat{\theta}$. Specifically, we use the estimators ($\hat{\pi}$ and $\hat{\alpha}$) estimated on D_l^c to estimate the $\hat{\theta}_l$ of l . By applying cross-fitting, we ensure that the nuisance functions and the parameters are estimated on separate, non-overlapping datasets. This approach diminishes the risk of overfitting and enhances the robustness of our estimation. And finally, to estimate $\hat{\theta}$, we randomly select one observation t^* in each market t and average it out across all folds. Thus the estimator for θ_0 and its variance can be given as follows:

$$\hat{\theta} = \frac{1}{n} \sum_{\ell=1}^L \sum_{t \in D_\ell^c} \{m(w_{t^*}, \hat{\pi}_\ell) + \hat{\alpha}_\ell(z_{t^*})(y_{t^*} - \hat{\pi}_\ell(z_{t^*}; \hat{\gamma}))\} \quad (9)$$

$$\hat{V} = \frac{1}{n} \sum_{\ell=1}^L \sum_{t \in D_\ell^c} \hat{\psi}_{t^* \ell}^2, \quad \hat{\psi}_{t^* \ell} = m(w_{t^*}, \hat{\pi}_\ell) - \hat{\theta} + \hat{\alpha}_\ell(z_{t^*})(y_{t^*} - \hat{\pi}_\ell(z_{t^*}; \hat{\gamma})) \quad (10)$$

5 Numerical Experiments

We now present a series of simulation studies that establish the numerical performance of our approach. First, in §5.1, we examine the predictive performance of our model on a series of models, including stylized discrete choice models with linear utilities as well as more general models that allow non-linear utilities

and realistic consumer behaviors such as inattention. Next, in §5.2, we present numerical experiments that demonstrate our approach’s ability to simulate counterfactuals. Finally, in §5.3, we demonstrate the applicability of the inference procedure proposed in §3.3.

5.1 Predictive Performance

In this section, we show that our approach can recover demand and elasticity estimates for a wide variety of settings without the knowledge of the underlying choice model and/or making parametric assumptions on consumers’ behaviors. Before describing the simulations, we first describe the metrics used for comparing the performance of different estimators and the benchmark estimation approaches used.

First, to assess the predictive performance of our approach, we focus on three quantities of interest:

- Market share ($\hat{\pi}_{jt}$)
- Own price elasticity ($\frac{\partial \hat{\pi}_{jt} / \pi_{jt}}{\partial p_{jt} / p_{jt}}$)
- Cross-price elasticity ($\frac{\partial \hat{\pi}_{jt} / \pi_{jt}}{\partial p_{k \neq jt} / p_{k \neq jt}}$)

In each simulation, we compare the performance of our model with the predictive performance of four baseline models:

- Multinomial Logit model (MNL)
- Random Coefficient Logit model (RCL)
- A standard neural network-based non-parametric method (NP).

Here we simply use all the features of all the products in the market as input and predict a market share vector without using the permutation invariance property or correcting for endogeneity. This is similar to the approaches used by [Gabel and Timoshenko \(2022\)](#) and [Cai et al. \(2022\)](#), which simply use a large neural network for demand prediction. For this baseline, we tune the hyperparameters of the neural network, including the number of layers, number of nodes in each layer, learning rate, and the number of epochs using 5-fold cross-validation for each data generation. We detail the space of hyperparameters in [Web Appendix C](#). We also apply the ReLU activation for each layer.

- A Mean Predictor (MP).

Here we predict a uniform market share for all observed products within a market, excluding the outside option. This predicted share is set to the average of all observed market shares in that market⁵. This serves as a baseline estimator as it does not account for individual product characteristics, providing a benchmark for the simplest prediction scenario. Additionally, comparing the performance of other models with MP also enables us to quantify the natural decrease in the MAE and RMSE when the number of products increases because the magnitude of the market shares decreases as the number of products increases⁶.

⁵For instance, consider a market with three observed products having market shares of 0.2, 0.3, and 0.4, respectively. This implies the outside option holds a market share of 0.1. In such a scenario, the MP would predict the market share for each of the three products to be 0.3, which is the average $((0.2 + 0.3 + 0.4) / 3)$.

⁶For instance, in a market with 100 products, the MAE and RMSE for any estimator are expected to be quite low. Comparing with the MP in such situations helps establish a baseline for MAE or RMSE.

5.1.1 Multinomial Logit and Random Coefficient Logit with Linear Utility

We first consider the two standard Data Generating Processes (DGPs) used in the demand estimation literature that use linear utility-based choice models – (1) Multinomial Logit model, and (2) the Random Coefficient Logit model. For both cases, we consider a setting with 10 products ($J = 10$), 100 markets ($T = 100$), one price feature, and 10 non-price features ($d = 10$). We define the utility u_{ijt} that consumer i in market t derives from product j as the following linear function:

$$u_{ijt} = \alpha_i p_{jt} + \beta_i X_{jt} + \varepsilon_{ijt}, \quad (11)$$

where ε_{ijt} represents an independently and identically distributed (iid) Type-I extreme value across products and consumers. $X_{jt} \in \mathbb{C}^d$ denotes the non-price features of the product. α_i, β_i are the model coefficients, which are kept constant for all consumers in the MNL, while in the RCL, they are normally distributed across consumers. The probability distribution of features and coefficients used are shown in Web Appendix D. Also, the mean utility from the outside option is normalized to 0.

We denote the market share of product j in market t generated from MNL by π_{jt}^{MNL} and the market share generated from RCL by π_{jt}^{RCL} . For each market, we generate the market shares of each product by simulating $N = 10,000$ individual choices and aggregating by each market as shown below.

$$\pi_{jt}^{MNL} = \frac{1}{N} \sum_i^N 1(u_{ijt} = \max_{k \in \mathcal{S}_t}(u_{ikt})) \quad (12)$$

$$\pi_{jt}^{RCL} = \frac{1}{N} \sum_i^N \frac{\exp(\alpha_i p_{jt} + \beta_i X_{jt})}{1 + \sum_{k \in \mathcal{S}_t} \exp(\alpha_i p_{kt} + \beta_i X_{kt})} \quad (13)$$

Note that for MNL, instead of simulating each individual’s choice probability, we simulate each individual’s choice based on the utility maximization principle. This approach ensures that when we use MNL (true model) for estimation, it does not reproduce the data perfectly.

For each DGP, we split the generated data into training data (80%) and test data (20%). We use the training data for estimation (both our model and the benchmark models described above).⁷ For the predicted market share, we present all the model results and comparisons on the test data. For the predicted own- and cross-elasticities, we present all the model results and comparisons on the training data.⁸

Tables 2a shows the Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) in the predicted market share ($\hat{\pi}_{jt}$) for our approach as well as the baseline models. We see that when the true model is MNL,

⁷In our simulations, we train both our model and NP with the log market shares ($\log(\pi_{jt})$). The performance metrics reported for predicted market shares are computed based on the exponential values of the predicted log market shares, bringing these metrics back to market shares (π_{jt}). The performance metrics reported for elasticities are computed based on the relative change of the predicted log market shares ($\partial \log(\pi_{jt})$) divided by the percentage change of the price ($\partial p_{jt}/p_{jt}$ for own-elasticity and $\partial p_{kt}/p_{kt}$ for cross-elasticity), which is equivalent to the elasticity calculated directly using the market share ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{jt}/p_{jt}}$ for own-elasticity and $\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{k \neq jt}/p_{k \neq jt}}$ for cross-elasticity).

⁸The reason we only use test data to report predicted market share accuracy is to demonstrate the model’s predictive performance on unseen data. In contrast, we use training data to report accuracy in elasticities to mimic the real empirical setting where we use full data to estimate elasticity.

Table 2: Baseline Predictive Performance: Market Shares, Own-Elasticity, and Cross-Elasticity

(a) Market Shares ($\hat{\pi}_{jt}$)

#	True Model	J	T	d	Our Model		MNL		RCL		NP		MP		No. Obs.
					MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	
0	MNL	5	100	10	0.0534	0.0834	0.0078	0.0105	0.0082	0.0052	0.1269	0.2364	0.2312	0.2220	2000
1	MNL	10	20	10	0.0585	0.1086	0.0040	0.0131	0.0089	0.0134	0.1129	0.2191	0.1365	0.2948	800
2	MNL	10	100	10	0.0333	0.0591	0.0044	0.0039	0.0026	0.0053	0.1181	0.1717	0.1422	0.1503	4000
3	MNL	10	200	10	0.0307	0.1346	0.0032	0.0102	0.0034	0.0197	0.1096	0.2170	0.1416	0.2106	8000
4	MNL	20	100	10	0.0194	0.0765	0.0015	0.0077	0.0023	0.0068	0.0707	0.2242	0.0768	0.2201	8000
5	RCL	5	100	10	0.0240	0.0314	0.0307	0.0382	0.0033	0.0042	0.0456	0.0583	0.0538	0.0656	2000
6	RCL	10	20	10	0.0206	0.0281	0.0270	0.0343	0.0034	0.0044	0.0540	0.0612	0.0418	0.0525	800
7	RCL	10	100	10	0.0171	0.0231	0.0262	0.0326	0.0025	0.0033	0.0458	0.0583	0.0413	0.0514	4000
8	RCL	10	200	10	0.0141	0.0187	0.0252	0.0318	0.0032	0.0039	0.0431	0.0559	0.0412	0.0513	8000
9	RCL	20	100	10	0.0099	0.0140	0.0262	0.0281	0.0018	0.0024	0.0390	0.0489	0.0276	0.0354	8000

(b) Own-Elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{jt}/p_{jt}}$)

#	True Model	J	T	d	Our Model		MNL		RCL		NP		No. Obs
					MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	
0	MNL	5	100	10	0.2588	1.1815	0.1414	1.1232	0.1757	1.1322	0.4554	1.2115	8000
1	MNL	10	20	10	0.3523	1.3557	0.1967	1.2970	0.2585	1.3118	1.0057	1.5316	3200
2	MNL	10	100	10	0.3346	1.4327	0.2066	1.3851	0.2150	1.3863	0.9266	1.5857	16000
3	MNL	10	200	10	0.3245	1.4131	0.2007	1.3842	0.2357	1.3876	0.8266	1.5768	32000
4	MNL	20	100	10	0.4146	1.6596	0.3305	1.6203	0.3570	1.6229	1.0151	1.8353	32000
5	RCL	5	100	10	0.1189	0.2310	0.1474	0.3735	0.0125	0.0305	0.1802	0.3145	8000
6	RCL	10	20	10	0.1799	0.3729	0.2039	0.5326	0.0365	0.1196	0.3768	0.3254	3200
7	RCL	10	100	10	0.1498	0.2862	0.2154	0.5531	0.0224	0.0685	0.2987	0.3643	16000
8	RCL	10	200	10	0.1209	0.2416	0.2188	0.5512	0.0233	0.0732	0.2464	0.4241	32000
9	RCL	20	100	10	0.1658	0.3533	1.4591	1.7099	0.0429	0.1319	0.4555	0.4741	32000

(c) Cross-Elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{k \neq jt}/p_{k \neq jt}}$)

#	True Model	J	T	d	Our Model		MNL		RCL		NP		No. Obs
					MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	
0	MNL	5	100	10	0.1349	0.8115	0.0107	0.6780	0.0118	0.6807	0.1968	0.9442	32000
1	MNL	10	20	10	0.0649	0.6742	0.0043	0.5326	0.0059	0.5424	0.1527	0.7123	28800
2	MNL	10	100	10	0.0862	0.6104	0.0043	0.5142	0.0049	0.5143	0.1885	0.7488	144000
3	MNL	10	200	10	0.0901	0.6075	0.0042	0.5140	0.0047	0.5153	0.2110	0.7831	288000
4	MNL	20	100	10	0.0482	0.4665	0.0015	0.4151	0.0016	0.4157	0.1270	0.5303	608000
5	RCL	5	100	10	0.0293	0.0571	0.0492	0.0551	0.0030	0.0070	0.0617	0.0960	32000
6	RCL	10	20	10	0.0261	0.0435	0.0332	0.0455	0.0039	0.0143	0.0972	0.0791	28800
7	RCL	10	100	10	0.0257	0.0493	0.0324	0.0447	0.0028	0.0090	0.0795	0.1124	144000
8	RCL	10	200	10	0.0213	0.0417	0.0353	0.0455	0.0035	0.0097	0.0745	0.1277	288000
9	RCL	20	100	10	0.0220	0.0431	0.0204	0.0359	0.0022	0.0103	0.0715	0.0890	608000

Note: Table 2a, 2b, 2c present the baseline predictive performance for predicted market shares, own- and cross-elasticities of our model and four baseline models. J, T, and d represent the number of products, non-price features, and markets, respectively. NP denotes a benchmark non-parametric method, which is a standard neural network. MP denotes the mean predictor. The Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) of predicted market shares for each scenario (i.e., true model, J, d, T) are computed using the test data from 20 iterations of data generation, while the MAE and RMSE of own- and cross-elasticities are computed based on the training data from 20 iterations of data generation. The column titled “No. Obs.” indicates the total number of observations for each metric. Specifically, the number of observations for market share ($\hat{\pi}_{jt}$) is calculated based on $T \times J \times 20\%$ (the portion of test data) $\times 20$ (the number of draws of simulations). The number of observations for own-elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{jt}/p_{jt}}$) is calculated based on $T \times J \times 80\%$ (the portion of training data) $\times 20$ (the number of draws of simulations). The number of observations for cross-elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{k \neq jt}/p_{k \neq jt}}$) is calculated based on $T \times J \times (J - 1) \times 80\%$ (the portion of training data) $\times 20$ (the number of draws of simulations).

our model cannot beat RCL or MNL, which is as expected; but the error of our model is quite close to the true model. When the true model is RCL, our model can beat MNL consistently and the performance of our model is also close to the true model. Importantly, we find that our model consistently outperforms the benchmark Non-Parametric (NP) method in all data generation processes. This is despite extensive hyperparameter tuning for the NP method.

There are two key reasons why our approach outperforms the standard neural network-based non-parametric method, especially as the number of products increases. First, unlike the standard neural network, our method can circumvent the curse of dimensionality that arises with the increase in the number of products. The standard neural network uses the stacked product features (d -dimension non-price features X_{jt} and price p_{jt}) as input and has $(J \times (d + 1)) \times h_1$ parameters in the input layer, where h_1 denotes the size of the first hidden layer.⁹ In contrast, our model only uses product features as the input (with dimension $d + 1$); see Figure 1. Therefore, the parameters for our model do not scale with the number of products. Thus, as the number of products increases, our method is able to exploit this information to improve its performance, whereas the standard NP is unable to do so. Second, our model can leverage product-level market-share data more effectively. Note that one observation in the standard NN consists of one market, whereas one observation for our method consists of one product in a market. Hence, given data on T markets, the number of samples available for the NP method is T , whereas the sample available for our model is $T \times J$. Together, these strengths of our approach lead to significantly better performance compared to a naive neural network.

We observe both sources of performance improvement in the numerical simulation results in Table 2a. As we vary the number of products (5, 10, and 20), the MAE of the predicted market shares from our model decreases monotonically (see simulation numbers 0, 2, and 4 for MNL, and 5, 7, and 9 for RCL). In contrast, the performance of the benchmark non-parametric estimator deteriorates as the number of products increases and becomes even worse than the Mean Prediction for the 20 product cases (simulation numbers 4 and 9). These findings demonstrate the overfitting problems and the curse of dimensionality issues discussed above. Indeed, this limitation has also been theoretically established by Sannai and Imaizumi (2019), who showed that for neural networks that do not take into account the inherent invariance structure, the generalization gap increases in proportion to the number of possible permutations, which is $\sqrt{J!}$ in our case. Further, we examine the performance of our model and the other benchmarks by varying the number of markets (20, 100, and 200) while keeping the number of products constant at 10; see simulations 1, 2, and 3 for MNL and 6, 7, and 8 for RCL. Although both our model and the NP method show improved performance with more markets, the non-parametric estimator is more adversely affected by a decrease in market numbers due to a more significant reduction in its sample size. This is particularly problematic in scenarios with one market, as the NP method becomes infeasible for estimation for only one sample.

Finally, in Tables 2b and 2c, we show the predictive performance of own-elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{jt}/p_{jt}}$) and cross-elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{k \neq jt}/p_{k \neq jt}}$), respectively. Again, we find that our model consistently outperforms the NP method in all scenarios for both own- and cross-elasticity predictions. When the true model is MNL, our model un-

⁹In this section, we consider only cases where there are no correlated unobservables. When there are potential endogeneity concerns, then we can include a residual μ_{jt} estimated from a first-stage regression, and then size of the input layer becomes $J \times (d + 2)$. We consider settings with endogeneity in Web Appendix §E and in the experiments on inference in §5.3.

derperforms compared to RCL, given that RCL inherently captures the substitution pattern among products in MNL¹⁰. When the true model is RCL, our model is the closest one to the true model. Unlike the market share predictions, the accuracy of own-elasticity predictions does not exhibit a clear monotonical improvement as the number of products increases. We observe a similar pattern even for the true model—the accuracy of own-elasticity decreases as the number of products increases. This suggests that it is not a deficiency of our model but due to the inherent complexities in estimating own-elasticities in markets with many products. In the prediction of cross-elasticity, our model shows a monotonical improvement in accuracy with an increase in the number of products. Additionally, for both own- and cross-elasticities, as the number of markets increases, the performance of our model is better due to the increase in the sample size.

So far, in the above simulations, we did not consider any endogenous explanatory variables. As discussed in §3.2, our approach can easily account for endogeneity following the estimation steps in §4. For interested readers, we present a set of numerical experiments with endogeneous explanatory variables in Web Appendix §E. The key takeaway from this analysis is that ignoring endogeneity can lead to significant biases in the estimates of own- and cross-price elasticities. Therefore, in the application to real data, we take care of endogeneity carefully; please see §6 for further details.

5.1.2 Random Coefficient Logit with Non-Linear Utility

In the previous section, we focused on linear utility specifications and standard choice behaviors. Recent literature has highlighted that the oversight of non-linear relationships between features and utilities can introduce biases in the estimates (Allenby et al., 2004). Conversely, non-parametric estimators are adept at capturing these non-linear patterns directly from the data. As a result, there has been a growing trend towards the adoption of non-linear utility functions. Therefore, we now focus on data generated from a random coefficient logit model with non-linear transformations applied to observable features. We consider a case with two product features – price and a non-price feature x . We apply a non-linear transformation $g(x)$. Following Bakhitov and Singh (2022), we consider two functions for $g(x)$:

a. $\log()$: $g(x) = \log(|16x - 8| + 1)\text{sign}(x - 0.5)$

b. $\sin()$: $g(x) = \sin(x)$

The log transformation is common in empirical studies, to capture a diminishing sensitivity of a feature on the market share. The sine transformation captures periodic or cyclical effects. For example, when the feature represents the time of year, normalized from 0 to 1, then using sine transformation can effectively capture the seasonal variations in consumer preferences.

The utility that consumer i in market t derives from product j then has the following non-linear form:

$$u_{ijt} = \alpha_i p_{jt} + \beta_i g(X_{jt}) + \varepsilon_{ijt}. \quad (14)$$

The marketshares then follow a similar structure to that from Equation (13).

¹⁰In an extreme case, when the variance of random coefficients is zero, RCL is equivalent to MNL.

As before, we generate data using this model and estimate the marketshares and elasticities using both our approach and the baseline models. When estimating the baseline MNL and RCL models, we assume that the researcher does not have knowledge of the non-linearities in the utility function and hence uses the simple linear utility(s) in their estimation (as shown in Equation (11)). The results for the predicted market shares ($\hat{\pi}_{jt}$), own-price elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{jt}/p_{jt}}$) and cross-elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{k \neq jt}/p_{k \neq jt}}$) are presented in Tables 3a, 3b, and 3c, respectively.

Regarding the MAE of predicted market shares, our model surpasses the RCL model by a factor of 8X and 4X across transformations (a) and (b), respectively. Similarly, considering the MAE of predicted own-elasticity in transformations (a) and (b), our model outperforms RCL by factors of 20X and 2.5X, respectively. For the MAE of predicted cross-elasticity, our model is 2X and 1.5X superior to RCL across transformations (a) and (b), respectively. It’s worth noting that while our model consistently outperforms the NP method across metrics, the NP method still shows better performance than both RCL and MNL in terms of MAE and RMSE for estimated own-elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{jt}/p_{jt}}$), underscoring the strengths of neural network-driven approaches in navigating non-linearities.

Table 3: Predictive Performance in Non-Linear Utility: Market Shares, Own-Elasticity, and Cross-Elasticity

(a) Market Shares ($\hat{\pi}_{jt}$)

#	True Model	Our model		MNL		RCL		NP		Mean		No. Obs.
		MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	
0	RCL-log()	0.0025	0.0063	0.0358	0.0361	0.0213	0.0309	0.0588	0.1235	0.0836	0.1401	200
1	RCL-sin()	0.0029	0.0046	0.0281	0.0340	0.0102	0.0172	0.0315	0.0449	0.0388	0.0527	200

(b) Own-Elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{jt}/p_{jt}}$)

#	True Model	Our Model		MNL		RCL		NP		No. Obs
		MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	
0	RCL-log()	0.0566	0.1523	5.4278	2.8249	1.1961	2.1135	0.6119	0.9085	16000
1	RCL-sin()	0.0609	0.2820	0.6229	1.1350	0.1777	0.4246	0.4057	1.0787	16000

(c) Cross-Elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{k \neq jt}/p_{k \neq jt}}$)

#	True Model	Our Model		MNL		RCL		NP		No. Obs
		MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	
0	RCL-log()	0.0150	0.0543	0.0436	0.1389	0.0372	0.4414	0.2552	0.5444	144000
1	RCL-sin()	0.0226	0.1047	0.0471	0.1794	0.0354	0.1751	0.1448	0.3357	144000

Note: This table presents the results when we add non-linear transformation in data generation. We generate using the Random Coefficient Logit (RCL) model, with 10 products and 100 markets, while only considering a single non-linearly transformed feature, which is the price.

5.1.3 Models with Consumer Inattention and Consideration Set Formation

Finally, we consider a scenario where consumers do not pay attention to all the products and/or are not fully informed of all the alternatives in the choice set. Recent literature has shown that this is often the case in many empirical settings (Goeree, 2008; Gabaix, 2019; Honka et al., 2019; Abaluck et al., 2020; Compiani, 2022). However, such cases violate a standard assumption of the choice model: that consumers are informed and consider all options when they make purchase decisions. In some parametric models (e.g., Van Nierop

et al., 2010), this issue is managed by constructing a consumer-level consideration set. However, consideration sets are usually unobserved in data; so these approaches often require assumptions on how consideration sets are formed, which might not always be appropriate or reflective of actual consumer behavior. Another way to manage this issue is to model search costs, i.e., allow consumers to ignore certain products because search is costly (e.g., Weitzman, 1978; Mehta et al., 2003; Hortaçsu and Syverson, 2004; Kim et al., 2010). Similarly, it also requires researchers to specify how search cost enters the utility function and decision process. In contrast to these models, our approach refrains from making any parametric assumptions, which allows for a potentially more flexible representation of consumer behavior in the case of consumer attention. To demonstrate how our model can capture the inattentive behavior, we look at a scenario where consumers are inattentive and deviate from the traditional random coefficient logit model.

We consider a simple model of inattention where a portion of consumers $(1 - \frac{1}{1+p_{ht}})$ ignore the product with the highest price, following Compiani (2022). In other words, the consideration set of $1 - \frac{1}{1+p_{ht}}$ consumers in market t excludes the highest-price product h . So when the price increases, the portion of inattentive consumers increases. Suppose there is only one feature, price, then, the choice probability of the highest price product h in the market t is:

$$\frac{1}{1 + p_{ht}} \frac{\exp(\alpha_i p_{ht})}{\sum_k \exp(\alpha_i p_{kt})}.$$

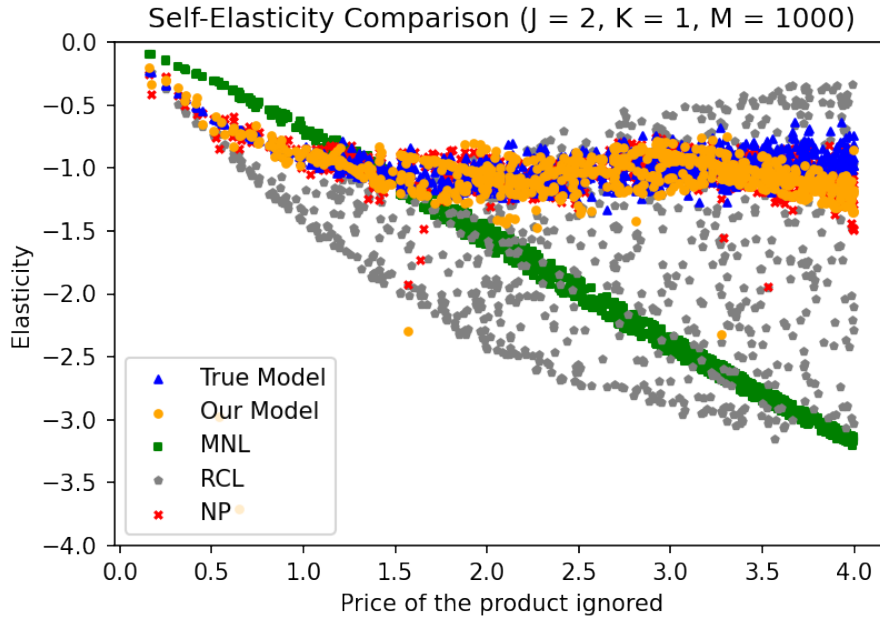
The choice probability of other products $j \neq h$ are given by:

$$\frac{1}{1 + p_{ht}} \frac{\exp(\alpha_i p_{jt})}{\sum_k \exp(\alpha_i p_{kt})} + \left(1 - \frac{1}{1 + p_{ht}}\right) \frac{\exp(\alpha_i p_{jt})}{\sum_{k \neq ht} \exp(\alpha_i p_{kt})}.$$

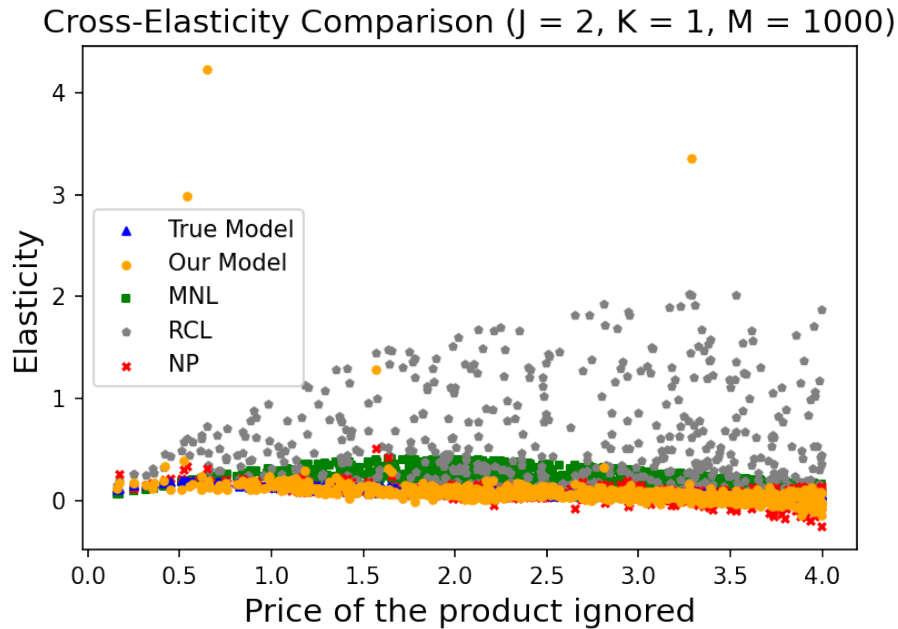
We first present detailed results on the own- and cross-elasticity for the two-product case in Figure 2. We consider the number of markets to be 1,000 so that we can observe enough variance in our data. Figure 2a shows how the estimated own-elasticity for the highest-priced product (which is ignored by a subset of consumers), for a range of prices. Note that in our model, when the price is higher, the portion of inattentive consumers is higher. Thus when we change the price, the change in market share is smaller than the case without inattention. While our model and the fully NP model are able to capture this pattern, both the parametric models (MNL and RCL) are unable to do so. Figure 2b shows how the estimated cross-elasticity of the other products vary with the price of the highest-priced product. Similarly, due to the ignorance of inattention, both MNL and RCL overestimate the magnitude of the elasticity of the other product. In contrast, our model and the fully NP are able to capture the true cross-price elasticity and are close to the true model.

Next, we show a more comprehensive set of results for all three metrics (market-shares, own-, and cross-price elasticities) when there are more products (2, 5, and 10) and fewer markets (100) in Table 4a, 4b, and 4c. We find that our approach consistently outperforms RCL and MNL, as expected. Further, the performance of our model improves as the number of products increases while that of the NP model monotonically decreases with the number of products (for the reasons discussed in §5.1.1).

In summary, we find that our approach adapts well even as the underlying model of consumer behavior



(a) Own-elasticity ($\frac{\partial \pi_{jt} / \pi_{jt}}{\partial p_{jt} / p_{jt}}$) in consumer inattention



(b) Cross-elasticity ($\frac{\partial \pi_{jt} / \pi_{jt}}{\partial p_{k \neq jt} / p_{k \neq jt}}$) in consumer inattention

Figure 2: Elasticity Effects in Consumer Inattention

Note: Figure 2 illustrates how different models perform when there is $1 - \frac{1}{1+p_{ht}}$ consumers who ignore the product with the highest price. We simulate market shares in the case of 2 products, 1000 markets, and 1 feature (price). Due to the existence of inattention, when the price is higher, the portion of inattentive consumers is higher. Thus when we change the price, the change in market share is smaller than the case without inattention.

Table 4: Consumer Inattention - Model Performance of Predicted Market Shares, Own-Elasticity, and Cross-Elasticity

(a) Market Shares ($\hat{\pi}_{jt}$)

#	J	Our Model		MNL		RCL		NP		Mean		No. Obs
		MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	
0	2	0.0047	0.0198	0.0590	0.0753	0.0316	0.0439	0.0076	0.0273	0.1840	0.2044	40
1	5	0.0064	0.0139	0.0203	0.0252	0.0178	0.0226	0.0250	0.0345	0.0780	0.1022	100
2	10	0.0033	0.0068	0.0137	0.0166	0.0072	0.0100	0.0258	0.0365	0.0492	0.0656	200

(b) Own-Elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{jt}/p_{jt}}$)

#	J	Our Model		MNL		RCL		NP		No. Obs
		MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	
0	2	0.0609	0.5190	0.6758	1.5698	0.3929	1.4598	0.0573	0.8804	3200
1	5	0.0978	1.8579	0.6917	1.9614	0.3753	1.9782	0.4288	1.9546	8000
2	10	0.0708	2.4273	0.8306	2.1031	0.1842	2.2787	0.5464	2.1450	16000

(c) Cross-Elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{k \neq jt}/p_{k \neq jt}}$)

#	J	Our Model		MNL		RCL		NP		No. Obs
		MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	
0	2	0.0419	4.3553	0.1911	8.8197	0.3350	8.5791	0.0378	5.6814	3200
1	5	0.0486	4.7745	0.0827	5.8357	0.0503	5.8404	0.1897	5.5601	32000
2	10	0.0212	3.8765	0.0476	4.0961	0.0146	4.1113	0.1676	4.1198	144000

Note: These tables present the MAE and RMSE for predicted market shares, own-elasticity, and cross-elasticity under consumer inattention across scenarios with 2, 5, and 10 products. The market share and elasticity are simulated assuming a portion of consumers ($1 - \frac{1}{1+p_{j_{ht}}}$), j_{ht} is the index of the highest-price product in market t ignore the product with the highest price. Each scenario is fixed at 100 markets and 1 feature (price) to maintain consistency with the RCL baseline.

changes without the need to impose any specific assumptions on consumer decision-making.

5.2 Counterfactual Analysis

In general, a key advantage of parametric models like MNL and RCL, or a structural approach to consumer behavior is their ability to predict outcomes in counterfactual scenarios outside the distribution of the data used for estimation. In contrast, fully non-parametric approaches like the naive neural network model that simply uses all product features as inputs cannot be used for counterfactual predictions¹¹. By leveraging the choice invariance property, our method adds some structure to the neural network architecture; and as a result, accommodates certain types of counterfactual models. In particular, we focus on a specific type of counterfactual that is of interest to firms and policy-makers – demand estimation when the choice set changes; e.g., through the introduction of a new product, the exit of an existing product, or the merger of two firms (Nevo, 2000; Petrin, 2002; Nevo, 2003; Wollmann, 2018). Our model can easily handle such counterfactuals since the choice function is specified as a function of the focal product features and the features of a set of competing products (see Theorem 1 and Figure 1). In contrast, estimating counterfactual demand when the choice set changes is infeasible with a standard neural network estimator due to its structural constraints on the input space. A change in the choice set would result in the change of size of the input vector, making such estimation infeasible.

To showcase the capability of our model to estimate counterfactuals, we consider a setting where a new product is introduced to the market. For comparison, we will only consider an MNL and RCL estimator since it is infeasible for the NP method to estimate the counterfactual. We use two data generation processes – Multinomial Logit (MNL) and Random Coefficient Logit (RCL), the same as §5.1.1 with 100 markets, 10 products (and one outside option) and 10 features. For each of these two cases, we consider a counterfactual where an 11th product is introduced to each market. The observable characteristics of the new product are simulated from the same distribution as other products.

In Table 5, we present the error in the estimated market share of all products when a new product is introduced to the market. We find that our model does quite well. When the true model is MNL, both MNL and RCL outperform our method, similar to what we find in §5.1.1. Our model also outperforms the MNL when the underlying data generation process is RCL, and produces results comparable to the true model. Overall, these results suggest that our approach can be used for counterfactual estimation, which is often a key focus of many substantive studies in marketing and economics.

5.3 Inference and Coverage Analysis

We now demonstrate the performance of the debiasing and inference procedure discussed in §3.3. The objective is to demonstrate the validity of the estimated confidence intervals. To this end, we estimate the average effect of a 1% change in own price on demand over all products ($\hat{\theta}$) and compute the corresponding

¹¹Both our model and fully NP can do counterfactuals when shocks only result in the change in features $\{X_{jt}, P_{jt}\}$. For example, consider a choice model where ranking is a feature. In such a scenario, researchers would like to see how the demand would change if the ranking policy is changed. To estimate this counterfactual scenario, one would update the value assigned to the ranking feature within a model. Therefore, both our model and fully NP are capable of doing such analysis, offering insights into how changes in specific features could impact demand.

Table 5: New Product Demand Estimation - Predicted Market Share ($\hat{\pi}_{jt}$)

True Model	Our Model		MNL		RCL		No. Obs.
	MAE	RMSE	MAE	RMSE	MAE	RMSE	
MNL	0.0234	0.0644	0.0041	0.0095	0.0023	0.0045	22,000
RCL	0.0186	0.0145	0.0265	0.0331	0.0023	0.0031	22,000

Note: This table presents the MAE and RMSE of predicted market shares of all products in the market when a new product enters. We simulate market shares as in our baseline scenario (10 products and one outside option, 100 markets, 10 non-price features) when an eleventh product is introduced.

confidence intervals of this effect. The difference between this section and §5.1.1 lies in both the estimators and the methods. In terms of estimators, in §5.1.1, we predict the market share ($\hat{\pi}_{jt}$), own-elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{jt}/p_{jt}}$) and cross-elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{k \neq jt}/p_{k \neq jt}}$) for individual products. In contrast, the object of interest in this section is the *average* effect of price on demand across all products ($\hat{\theta}$). As a result, in §5.1.1, we did not use the debiasing techniques that we apply here. It is important to emphasize that, in our approach, constructing a confidence interval is viable only for aggregate measures, not for individual observations.


To simulate the data, we consider a random coefficient logit model of demand with 3 products across 100 markets. We set the true model parameters to be $\beta_{ik} \sim \mathcal{N}(1, 0.5)$, $\alpha_i \sim \mathcal{N}(-1, 0.5)$. The effect of a 1% increase in a product’s price is given by

$$\theta_0 = \mathbb{E}[m(w_{jt}, \pi)] = \mathbb{E}[\pi(p_{jt} * (1.01), X_{jt}, \{p_{kt}, X_{kt}\}_{k \neq j}) - \pi(p_{jt}, X_{jt}, \{p_{kt}, X_{kt}\}_{k \neq j})],$$

As discussed earlier, one way to estimate this effect is to compute the sample analog of this using the estimated $\hat{\pi}$, such that $\hat{\theta} = \frac{1}{n} \sum_{i=1}^n m(w_{jt}, \hat{\pi})$. However, as we pointed out earlier, the distribution of $\hat{\theta}$ might not be asymptotically normal. To demonstrate this, in Figure 3a, we display the histogram of the estimated effect across 50 random samples by using the plug-in method. We standardize the estimates by subtracting the mean and then dividing by the standard deviation and plot them against the standard normal distribution. As one can observe the distribution appears multi-modal and deviates from a standard normal distribution. Next, we use our proposed debiased estimator and plot the standardized estimates across 50 samples of draws in Figure 3b. The resultant distribution with the debiased estimator is much closer to a standard normal. Finally, we calculate the 95% confidence intervals using our debiased estimator across 50 random draws. In Table 6, we report the bias (mean absolute error from all draws) and the coverage i.e., the percentage of times the true parameter is covered in the estimated confidence intervals. We find that bias across both data-generating processes (RCL and MNL) is notably low, reflecting only a -0.0001 difference from the true effect. The coverage rate of the 95% confidence interval in our corrected model is 90%, indicating good coverage. This shows that our debiased estimator can be used to conduct valid inference in finite samples.

6 Empirical Data Analysis: US Automobile Data (1971 - 1990)

In this section, we apply our model to a real-world dataset. We use the “US Automobile Data (1971 – 1990)” from Berry et al. (1995). The dataset features cars in the US market from 1971 to 1990, with each



(a) Distribution of Estimated Average Effect of Price Change with Plug-in

(b) Distribution of Estimated Average Effect of Price Change with the Debiased Estimator

Figure 3: Distribution of Estimated Average Effect of Price Change

Note: The figure shows the distribution of the standardized plug-in and debiased estimators of the average effect of 1% change in price on demand. To simulate the data, we consider 3 products across 100 markets with 5 non-price features using RCL for 50 samples. For each sample, we set the true model parameters to be $\beta_{ik} \sim \mathcal{N}(1, 0.5)$, $\alpha_i \sim \mathcal{N}(-1, 0.5)$. Figure 3a displays the distribution of the estimated effect with the Plug-in method and Figure 3b shows the result when employing the debiased estimator.

Table 6: Inference Coverage Analysis

True Model	True Effect	Bias	95% CI Cov.
RCL	-0.0013	-0.0001	90%
MNL	-0.0016	-0.0001	90%

Note: This table presents the bias and coverage rate of 95% confidence interval using our debiased estimator of the average effect of 1% change in price on demand. To simulate the data, we consider 3 products across 100 markets with 5 non-price features using RCL and MNL for 50 samples of draws. For each RCL sample, we set the true model parameters to be $\beta_{ik} \sim \mathcal{N}(1, 0.5)$, $\alpha_i \sim \mathcal{N}(-1, 0.5)$. For each MNL sample, we set the true model parameters to be $\beta_{ik} = 1$, $\alpha_i = -1$.

year regarded as a market. The number of cars varies from 86 to 150 each year. For each car, the dataset provides information such as the car’s name, the manufacturing company, factory region, market share, price, and four exogenous car characteristics: horsepower, space, mileage per dollar, and the presence of an air conditioning device.

Even though the dataset is relatively small, it presents three key challenges that make it difficult to use naive non-parametric estimators: (i) the dataset features markets with more than 100 products and only 20 markets in total, (ii) the product assortment in each year or market varies; (iii) the feature “price” has the endogeneity issue, which was not considered in numerical experiments in §5. In this section, we demonstrate the use of our estimator, which is capable of effectively addressing such challenges posed in real-world datasets.

For each component of our model (ϕ_1 , ϕ_2 , and ρ), we use a standard 3-layer neural network, and this is implemented without further hyperparameter tuning. We implement ReLU activation function at each layer. This architecture is the same as the one we used in our numerical experiments. For comparison, we replicate the random coefficient logit model (with only the demand side) used by [Berry et al. \(1995\)](#) using the Python package `pyblp` ([Conlon and Gortmaker, 2020](#)). In our replication, we allow for heterogeneity in random coefficients across all variables. Our findings show that the estimates obtained from our model




Figure 4: Elasticity Estimation Comparison

Figure 4 presents the estimated own-elasticity $(\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{jt}/p_{jt}})$ of our model without IV and with IV. The x-axis represents the price of the focal product, while the y-axis shows the products' own-elasticities. Each point corresponds to a product in a market, resulting in 2,217 observations. We report the estimated elasticity based on the same price variation used in the BLP paper (a 1,000-dollar change). Although we cannot ascertain the true value of own-elasticity, it is widely accepted that own-elasticity should generally be negative for most, if not all, products. In Figure 4a, we observe that the majority of low-priced products (priced below 5,000 dollars) exhibit positive estimated own-elasticity, demonstrating the existence of the endogeneity.


are comparable to the random coefficient logit estimation presented in [Berry et al. \(1995\)](#).

We estimate our model both without and with consideration of endogeneity. To address endogeneity, we utilize three sets of IVs – (i) the sum of characteristics of all car models, excluding the product in focus, produced by the same firm in the same year; (ii) the sum of characteristics of all car models, excluding the product in focus, produced by rival firms in the same year; and (iii) cost shifters, which encompass the wage and exchange rate prevalent in the year and region where the factory is located. The utilization of traditional BLP-style instruments, as discussed by [Gandhi and Houde \(2019\)](#), can be problematic due to their relative weakness, often resulting in considerable bias in the estimation of parameters. These issues are significantly exacerbated in non-parametric models. Thus, to counter potential concerns related to weak instruments, we employ a machine-learning-based IV methodology (MLIV) as proposed by [Singh et al. \(2020\)](#). We detail the estimation procedure and results using BLP style IVs in Web Appendix F.

In Figure 4, we present the estimated own-elasticity $(\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{jt}/p_{jt}})$ of our model without IV and with IV. The x-axis represents the price of the focal product, while the y-axis shows the product's own-elasticity. Each point corresponds to a product in a market, resulting in 2,217 observations. We report the estimated elasticity based on the same price variation used in the BLP paper (a 1,000-dollar change). In Figure 4a, we observe that the majority of low-priced products (priced below 6,000 dollars) exhibit positive estimated own-elasticity, demonstrating the existence of the endogeneity. We also notice that this issue is attenuated when we correct for endogeneity (see Figure 4), i.e., which suggests that our approach is able to handle situations with endogenous features.

We report the own-elasticity $(\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{jt}/p_{jt}})$ and cross-elasticity $(\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{k \neq jt}/p_{k \neq jt}})$ estimated in our model and random coefficient logit model with a sample of 13 cars in the 1990 market in Tables 7 and 8. The sample of 13 cars is the same as the one reported in [Berry et al. \(1995\)](#). Overall, our results are very similar and

comparable to [Berry et al. \(1995\)](#). We also plot the distributions of the estimated own-elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{jt}/p_{jt}}$) and cross-elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{k \neq jt}/p_{k \neq jt}}$) obtained from our model and the BLP model in [Figure 5](#). The filled areas in the violin plots represent the complete range of the elasticities, while the text labels next to the line indicate the mean values. The estimated mean own- and cross-elasticities appear to be similar between our model and the BLP model, though our model exhibits a larger RMSE in the estimated elasticity values compared to the BLP model.



(a) Own-Elasticity Estimation (Our Model vs. BLP Model) (b) Cross-Elasticity Estimation (Our Model vs. BLP Model)

Figure 5: Elasticity Estimation Comparison

Note: [Figure 5](#) illustrates the distributions of the estimated own- and cross-elasticities obtained from our model and the BLP model. The filled areas in the violin plots represent the complete range of the elasticities, while the text labels indicate the mean values.

We further estimate the average own-elasticity ($\hat{\theta}$) for high-priced, medium-priced, and low-priced cars and construct a confidence interval for each category using our inference procedure. We present our result in [Table 9](#). Both our model and the BLP model indicate that the average own-elasticity is highest (in terms of the absolute value of own-elasticity) for high-priced cars and lowest for low-priced cars. Moreover, the 95% confidence intervals for all three categories do not include zero. This also demonstrates the efficiency of our model even when there is a limited sample of only 20 observations.

The empirical analysis demonstrates the applicability and effectiveness of our model in a real-world setting, addressing challenges such as limited sample size, variability in product assortments, and endogeneity. The comparable results with established econometric models, such as BLP model, help in validating the robustness and reliability of our approach.

7 Conclusion

Choice models are fundamental in understanding consumer behavior and informing business decisions. Over the years, various methods, both parametric and non-parametric, have been developed to represent consumer behavior. While parametric methods, such as logit or probit-based models, are favored for their simplicity and interpretability, their restrictive assumptions can limit their ability to fully capture consumer preferences' intricacies. On the other hand, non-parametric methods offer a more flexible approach, but

	Acura Legend	BMW 735i	Buick Century	Cadillac Seville	Chevy Cavalier	Ford Escort	Ford Taurus	Honda Accord	Lexus LS400	Lincoln Town Car	Mazda 323	Nissan Maxima	Nissan Sentra
Acura Legend	-5.6060	0.1993	0.2198	0.2221	0.0632	0.2317	0.2199	0.2337	0.2144	0.2187	0.2595	0.2354	0.2595
BMW 735i	0.4095	-6.1528	0.3653	0.4093	0.0352	0.3525	0.3655	0.3807	0.4120	0.4084	0.3547	0.4161	0.3547
Buick Century	0.1234	0.1020	-6.1023	0.1213	0.0376	0.1294	0.1215	0.1205	0.1175	0.1191	0.1400	0.1299	0.1400
Cadillac Seville	0.2895	0.2179	0.2631	-7.2896	0.0647	0.2692	0.2631	0.2566	0.2810	0.2818	0.2892	0.2849	0.2892
Chevy Cavalier	0.0142	-0.0013	0.0202	0.0167	-1.3447	0.0171	0.0202	0.0353	0.0126	0.0167	0.0354	0.0291	0.0355
Ford Escort	0.0410	0.0245	0.0353	0.0413	-0.0148	-1.8519	0.0353	0.0520	0.0392	0.0413	0.0494	0.0513	0.0494
Ford Taurus	0.1166	0.0914	0.1188	0.1160	0.0183	0.1199	-6.1473	0.1258	0.1066	0.1157	0.1451	0.1290	0.1451
Honda Accord	0.0975	0.0647	0.0968	0.1006	-0.0003	0.0945	0.0969	-5.7438	0.0914	0.1006	0.1166	0.1140	0.1165
Lexus LS400	0.3357	0.2606	0.3136	0.3325	0.0822	0.3235	0.3137	0.3126	-6.8791	0.3271	0.3495	0.3348	0.3494
Lincoln Town Car	0.2681	0.2310	0.2548	0.2656	0.0713	0.2663	0.2548	0.2648	0.2586	-5.3996	0.3009	0.2719	0.3009
Mazda 323	0.0361	0.0212	0.0272	0.0326	-0.0127	0.0249	0.0272	0.0363	0.0323	0.0326	-2.6589	0.0404	0.0357
Nissan Maxima	0.1579	0.1367	0.1589	0.1555	0.0425	0.1670	0.1589	0.1689	0.1484	0.1534	0.1884	-7.2216	0.1884
Nissan Sentra	0.0386	0.0239	0.0304	0.0384	-0.0202	0.0294	0.0304	0.0496	0.0375	0.0383	0.0439	0.0470	-1.8754

Table 7: Estimated own- and cross-elasticities of a sample of automobile data using our method

Note: This table presents the estimated own- and cross-elasticity of a sample of 13 cars in the 1990 market using our model. The selected cars are the same as Berry et al. (1995) reports. Each entry with row index i and column index j gives the percentage change in demand divided by the percentage change in price (based on \$1,000 change in the price of i).

	Acura Legend	BMW 735i	Buick Century	Cadillac Seville	Chevy Cavalier	Ford Escort	Ford Taurus	Honda Accord	Lexus LS400	Lincoln Town Car	Mazda 323	Nissan Maxima	Nissan Sentra
Acura Legend	-5.4677	0.0489	0.0205	0.1029	0.0221	0.0220	0.0143	0.1477	0.1503	0.0273	0.0013	0.1359	0.0039
BMW 735i	0.1267	-9.8502	0.0156	0.1058	0.0122	0.0121	0.0057	0.1313	0.1546	0.0278	0.0006	0.1375	0.0022
Buick Century	0.0184	0.0054	-5.1978	0.0124	0.1153	0.1043	0.1475	0.1982	0.0165	0.0800	0.0076	0.0379	0.0174
Cadillac Seville	0.1293	0.0513	0.0175	-6.6819	0.0151	0.0150	0.0099	0.1393	0.1576	0.0271	0.0008	0.1406	0.0027
Chevy Cavalier	0.0131	0.0028	0.0766	0.0071	-3.1163	0.1421	0.0849	0.2608	0.0086	0.0404	0.0100	0.0395	0.0241
Ford Escort	0.0137	0.0029	0.0726	0.0074	0.1487	-3.0590	0.0603	0.2781	0.0090	0.0263	0.0106	0.0419	0.0258
Ford Taurus	0.0048	0.0007	0.0554	0.0027	0.0479	0.0326	-4.0258	0.0727	0.0017	0.1779	0.0026	0.0122	0.0057
Honda Accord	0.0387	0.0133	0.0582	0.0291	0.1151	0.1173	0.0568	-4.3399	0.0409	0.0297	0.0081	0.0618	0.0200
Lexus LS400	0.1350	0.0536	0.0166	0.1126	0.0130	0.0130	0.0045	0.1400	-7.4316	0.0243	0.0006	0.1464	0.0024
Lincoln Town Car	0.0087	0.0034	0.0286	0.0069	0.0217	0.0135	0.1693	0.0362	0.0086	-5.6139	0.0011	0.0123	0.0024
Mazda 323	0.0114	0.0020	0.0743	0.0056	0.1476	0.1494	0.0679	0.2723	0.0063	0.0304	-2.8631	0.0390	0.0254
Nissan Maxima	0.1008	0.0393	0.0314	0.0831	0.0493	0.0500	0.0271	0.1749	0.1209	0.0286	0.0033	-4.7872	0.0086
Nissan Sentra	0.0140	0.0031	0.0707	0.0078	0.1471	0.1504	0.0617	0.2763	0.0095	0.0271	0.0105	0.0422	-3.1799

Table 8: Estimated own and cross-elasticities of a sample of automobile data using the BLP model

Note: This table presents the estimated own- and cross-elasticity of a sample of 13 cars in the 1990 market using the BLP model. The selected cars are the same as Berry et al. (1995) reports. Each entry with row index i and column index j gives the percentage change in demand divided by the percentage change in price (based on \$1,000 change in the price of i).

	BLP Model	Our Model	No. Obs.
	Mean Estimate	Mean Estimate (95% Confidence Interval)	
High	-5.6705	-4.1922 (-6.2204, -2.1641)	20
Medium	-3.7354	-3.7215 (-4.9260, -2.5169)	20
Low	-2.9174	-2.0697 (-3.1980, -0.9415)	20

Table 9: Estimates of Average Own-Elasticity

Note: This table presents the estimated average own-elasticity for cars across various price categories. We categorize cars with a price over \$20k as “high-priced”, cars priced between \$8k and \$20k as “medium-priced”, and all other cars as “low-priced”. For each category, we randomly select one car of the category from each market as one observation. For the BLP model, we calculate the average own-elasticity of the sampled cars as the mean estimate. For our model, we estimate the average own-elasticity and construct the confidence interval following our inference procedure.

they often suffer from the “curse of dimensionality”, where the complexity of estimating choice functions escalates exponentially with an increase in the number of products.

In this paper, we propose a fundamental characterization of choice models that combines the tractability of traditional choice models and the flexibility of non-parametric estimators. This characterization specifically tackles the challenge of high dimensionality in choice systems and facilitates flexible estimation of choice functions. Through extensive simulations, we validate the efficacy of our model, demonstrating its superior ability to capture a range of consumer behaviors that traditional choice models fail to capture. We also show how to address the endogeneity issue and estimate counterfactuals in our characterization. Furthermore, leveraging the recent strides in the automatic debiased machine learning literature, we offer an inference procedure that constructs confidence intervals on relevant objects, such as price elasticities. Finally, we apply our method to the automobile dataset from [Berry et al. \(1995\)](#). Our empirical analysis affirms that our model produces results that align well with the extant literature.

Our paper opens many avenues for future research. We focus on using neural network-based estimators. However, estimators, such as Gaussian processes and Gradient boosting-based estimators can be adopted to estimate the proposed functionals. Also, we believe more experimentation needs to be done on the neural network design side.

Competing Interests Declaration

Author(s) have no competing interests to declare.

References

- J. Abaluck, G. Compiani, and F. Zhang. A method to estimate discrete choice models that is robust to consumer search. Technical report, National Bureau of Economic Research, 2020.
- J. Abrevaya. Rank estimation of a generalized fixed-effects regression model. *Journal of Econometrics*, 95(1):1–23, 2000.
- D. Ackerberg, X. Chen, J. Hahn, and Z. Liao. Asymptotic efficiency of semiparametric two-step gmm. *Review of Economic Studies*, 81(3):919–943, 2014.

- C. Ai and X. Chen. Estimation of possibly misspecified semiparametric conditional moment restriction models with different conditioning variables. *Journal of Econometrics*, 141(1):5–43, 2007.
- P. Albuquerque and B. J. Bronnenberg. Estimating demand heterogeneity using aggregated data: an application to the frozen pizza category. *Marketing Science*, 28(2):356–372, 2009.
- R. Allen and J. Rehbeck. Identification with additively separable heterogeneity. *Econometrica*, 87(3):1021–1054, 2019.
- G. M. Allenby, T. S. Shively, S. Yang, and M. J. Garratt. A choice model for packaged goods: Dealing with discrete quantities and quantity discounts. *Marketing Science*, 23(1):95–108, 2004.
- A. Aouad and A. Désir. Representing random utility choice models with neural networks. *arXiv preprint arXiv:2207.12877*, 2022.
- E. Bakhitov and A. Singh. Causal gradient boosting: Boosted instrumental variable regression. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 604–605, 2022.
- M. Ben-Akiva, D. Bolduc, and M. Bradley. Estimation of travel choice models with randomly distributed values of time. *Transportation Research Record*, 1413:88–97, 1993.
- Y. Bentz and D. Merunka. Neural networks and the multinomial logit for brand choice modelling: a hybrid approach. *Journal of Forecasting*, 19(3):177–200, 2000.
- S. Berry, J. Levinsohn, and A. Pakes. Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, pages 841–890, 1995.
- S. T. Berry and P. A. Haile. Identification in differentiated products markets using market level data. *Econometrica*, 82(5):1749–1797, 2014.
- S. T. Berry and P. A. Haile. Nonparametric identification of differentiated products demand using micro data. Technical report, National Bureau of Economic Research, 2020.
- D. Besanko, S. Gupta, and D. Jain. Logit demand estimation under competitive pricing behavior: An equilibrium framework. *Management Science*, 44(11-part-1):1533–1547, 1998.
- J. Blanchet, G. Gallego, and V. Goyal. A markov chain approximation to choice modeling. *Operations Research*, 64(4):886–905, 2016.
- R. Blundell, J. Horowitz, and M. Parey. Nonparametric estimation of a nonseparable demand function under the slusky inequality restriction. *Review of Economics and Statistics*, 99(2):291–304, 2017.
- J. H. Boyd and R. E. Mellman. The effect of fuel economy standards on the us automotive market: an hedonic demand analysis. *Transportation Research Part A: General*, 14(5-6):367–378, 1980.
- R. A. Briesch, P. K. Chintagunta, and R. L. Matzkin. Nonparametric discrete choice models with unobserved heterogeneity. *Journal of Business & Economic Statistics*, 28(2):291–307, 2010.
- Z. Cai, H. Wang, K. Talluri, and X. Li. Deep learning for choice modeling. *arXiv preprint arXiv:2208.09325*, 2022.
- N. S. Cardell and F. C. Dunbar. Measuring the societal impacts of automobile downsizing. *Transportation Research Part A: General*, 14(5-6):423–434, 1980.

- S. Chatterjee and J. Jafarov. Prediction error of cross-validated lasso. *arXiv preprint arXiv:1502.06291*, 2015.
- X. Chen and T. M. Christensen. Optimal sup-norm rates and uniform inference on nonlinear functionals of nonparametric iv regression. *Quantitative Economics*, 9(1):39–84, 2018.
- V. Chernozhukov, W. K. Newey, V. Quintas-Martinez, and V. Syrgkanis. Automatic debiased machine learning via neural nets for generalized linear regression. *arXiv preprint arXiv:2104.14737*, 2021.
- V. Chernozhukov, J. C. Escanciano, H. Ichimura, W. K. Newey, and J. M. Robins. Locally robust semiparametric estimation. *Econometrica*, 90(4):1501–1535, 2022a.
- V. Chernozhukov, W. K. Newey, and R. Singh. Automatic debiased machine learning of causal and structural effects. *Econometrica*, 90(3):967–1027, 2022b.
- P. K. Chintagunta. Endogeneity and heterogeneity in a probit demand model: Estimation using aggregate data. *Marketing Science*, 20(4):442–456, 2001.
- S. Chitla, S. Jagabathula, and A. Venkataraman. Nonparametric demand estimation in the presence of unobserved factors. *Available at SSRN*, 2022.
- G. Compiani. Market counterfactuals and the specification of multiproduct demand: A nonparametric approach. *Quantitative Economics*, 13(2):545–591, 2022.
- C. Conlon and J. Gortmaker. Best practices for differentiated products demand estimation with PyBLP. *The RAND Journal of Economics*, 51(4):1108–1161, 2020.
- C. V. Forinash and F. S. Koppelman. Application and interpretation of nested logit models of intercity mode choice. *Transportation research record*, (1413), 1993.
- M. Fosgerau and D. Kristensen. Identification of a class of index models: A topological approach. *The Econometrics Journal*, 24(1):121–133, 2021.
- J. T. Fox and A. Gandhi. Nonparametric identification and estimation of random coefficients in multinomial choice models. *The RAND Journal of Economics*, 47(1):118–139, 2016.
- M. J. Funk, D. Westreich, C. Wiesen, T. Stürmer, M. A. Brookhart, and M. Davidian. Doubly robust estimation of causal effects. *American journal of epidemiology*, 173(7):761–767, 2011.
- X. Gabaix. Behavioral inattention. In *Handbook of behavioral economics: Applications and foundations 1*, volume 2, pages 261–343. Elsevier, 2019.
- S. Gabel and A. Timoshenko. Product choice with large assortments: A scalable deep-learning model. *Management Science*, 68(3):1808–1827, 2022.
- A. Gandhi and J.-F. Houde. Measuring substitution patterns in differentiated-products industries. *NBER Working paper*, (w26375), 2019.
- M. S. Goeree. Limited information and advertising in the us personal computer industry. *Econometrica*, 76(5):1017–1074, 2008.
- W. H. Greene and D. A. Hensher. A latent class model for discrete choice analysis: contrasts with mixed logit. *Transportation Research Part B: Methodological*, 37(8):681–698, 2003.

- L. Grigolon and F. Verboven. Nested logit or random coefficients logit? a comparison of alternative discrete choice models of product differentiation. *Review of Economics and Statistics*, 96(5):916–935, 2014.
- J. Han, Y. Li, L. Lin, J. Lu, J. Zhang, and L. Zhang. Universal approximation of symmetric and anti-symmetric functions. *arXiv preprint arXiv:1912.01765*, 2019.
- Y. Han, F. C. Pereira, M. Ben-Akiva, and C. Zengras. A neural-embedded discrete choice model: Learning taste representation with strengthened interpretability. *Transportation Research Part B: Methodological*, 163:166–186, 2022.
- J. A. Hausman and W. K. Newey. Individual heterogeneity and average welfare. *Econometrica*, 84(3):1225–1248, 2016.
- J. A. Hausman and D. A. Wise. A conditional probit model for qualitative choice: Discrete decisions recognizing interdependence and heterogeneous preferences. *Econometrica: Journal of the econometric society*, pages 403–426, 1978.
- D. A. Hirshberg and S. Wager. Debaised inference of average partial effects in single-index models: Comment on wooldridge and zhu. *Journal of Business & Economic Statistics*, 38(1):19–24, 2020.
- E. Honka, A. Hortaçsu, and M. Wildenbeest. Empirical search and consideration sets. In *Handbook of the Economics of Marketing*, volume 1, pages 193–257. Elsevier, 2019.
- B. E. Honoré and E. Kyriazidou. Panel data discrete choice models with lagged dependent variables. *Econometrica*, 68(4):839–874, 2000.
- A. Hortaçsu and C. Syverson. Product differentiation, search costs, and competition in the mutual fund industry: A case study of s&p 500 index funds. *The Quarterly journal of economics*, 119(2):403–456, 2004.
- H. Ichimura and W. K. Newey. The influence function of semiparametric estimators. *Quantitative Economics*, 13(1):29–61, 2022.
- Z. Jiang, T. Chan, H. Che, and Y. Wang. Consumer search and purchase: An empirical investigation of retargeting based on consumer online behaviors. *Marketing Science*, 40(2):219–240, 2021.
- J. Joo. Rational inattention as an empirical framework for discrete choice and consumer-welfare evaluation. *Journal of Marketing Research*, 60(2):278–298, 2023.
- W. A. Kamakura and G. J. Russell. A probabilistic choice model for market segmentation and elasticity structure. *Journal of marketing research*, 26(4):379–390, 1989.
- S. Khan, F. Ouyang, and E. Tamer. Inference on semiparametric multinomial response models. *Quantitative Economics*, 12(3):743–777, 2021.
- J. B. Kim, P. Albuquerque, and B. J. Bronnenberg. Online demand under limited consumer search. *Marketing science*, 29(6):1001–1023, 2010.
- A. Lewbel. Semiparametric qualitative response model estimation with unknown heteroscedasticity or instrumental variables. *Journal of econometrics*, 97(1):145–177, 2000.
- Z. Lu, X. Shi, and J. Tao. Semi-nonparametric estimation of random coefficients logit model for aggregate

- demand. *Journal of Econometrics*, 2023.
- R. D. Luce. On the possible psychophysical laws. *Psychological review*, 66(2):81, 1959.
- C. F. Manski. Semiparametric analysis of random effects linear models from binary panel data. *Econometrica: Journal of the Econometric Society*, pages 357–362, 1987.
- J. Marschak. Binary choice constraints on random utility indicators. 1959.
- D. McFadden and K. Train. Mixed mnl models for discrete response. *Journal of applied Econometrics*, 15(5):447–470, 2000.
- D. McFadden et al. Conditional logit analysis of qualitative choice behavior. 1973.
- S. R. Mehandiratta. *Time-of-day effects in inter-city business travel*. University of California, Berkeley, 1996.
- N. Mehta, S. Rajiv, and K. Srinivasan. Price uncertainty and consumer search: A structural model of consideration set formation. *Marketing science*, 22(1):58–84, 2003.
- A. Nevo. Mergers with differentiated products: The case of the ready-to-eat cereal industry. *The RAND Journal of Economics*, pages 395–421, 2000.
- A. Nevo. New products, quality changes, and welfare measures computed from estimated demand systems. *Review of Economics and statistics*, 85(2):266–275, 2003.
- W. K. Newey. The asymptotic variance of semiparametric estimators. *Econometrica: Journal of the Econometric Society*, pages 1349–1382, 1994.
- A. Pakes and J. Porter. Moment inequalities for multinomial choice with fixed effects. *Quantitative Economics*, Forthcoming, 2022.
- A. Petrin. Quantifying the benefits of new products: The case of the minivan. *Journal of political Economy*, 110(4):705–729, 2002.
- A. Petrin and K. Train. A control function approach to endogeneity in consumer choice models. *Journal of marketing research*, 47(1):3–13, 2010.
- D. Revelt and K. Train. Mixed logit with repeated choices: households’ choices of appliance efficiency level. *Review of economics and statistics*, 80(4):647–657, 1998.
- J. M. Robins, A. Rotnitzky, and L. P. Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association*, 89(427):846–866, 1994.
- A. Sannai and M. Imaizumi. Improved generalization bound of permutation invariant deep neural networks. 2019.
- X. Shi, M. Shum, and W. Song. Estimating semi-parametric panel multinomial choice models using cyclic monotonicity. *Econometrica*, 86(2):737–761, 2018.
- B. Sifringer, V. Lurkin, and A. Alahi. Enhancing discrete choice models with representation learning. *Transportation Research Part B: Methodological*, 140:236–261, 2020.
- A. Singh, K. Hosanagar, and A. Gandhi. Machine learning instrument variables for causal inference. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 835–836, 2020.

- K. Sudhir. Competitive pricing behavior in the auto market: A structural analysis. *Marketing Science*, 20(1):42–60, 2001.
- P. Tebaldi, A. Torgovitsky, and H. Yang. Nonparametric estimates of demand in the california health insurance exchange. *Econometrica*, 91(1):107–146, 2023.
- L. L. Thurstone. A law of comparative judgment. *Psychological review*, 34(4):273, 1927.
- K. Train. A comparison of hierarchical bayes and maximum simulated likelihood for mixed logit. *University of California, Berkeley*, pages 1–13, 2001.
- K. E. Train. *Discrete choice methods with simulation*. Cambridge university press, 2009.
- K. E. Train, D. L. McFadden, and M. Ben-Akiva. The demand for local telephone service: A fully discrete model of residential calling patterns and service choices. *The RAND Journal of Economics*, pages 109–123, 1987.
- S. Turlo, M. Fina, J. Kasinger, A. Laghaie, and T. Otter. Discrete choice in marketing through the lens of rational inattention. 2023.
- E. Van Nierop, B. Bronnenberg, R. Paap, M. Wedel, and P. H. Franses. Retrieving unobserved consideration sets from household panel data. *Journal of Marketing Research*, 47(1):63–74, 2010.
- E. Wagstaff, F. Fuchs, M. Engelcke, I. Posner, and M. A. Osborne. On the limitations of representing functions on sets. In *International Conference on Machine Learning*, pages 6487–6494. PMLR, 2019.
- M. J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019.
- A. Wang. Sieve blp: A semi-nonparametric model of demand for differentiated products. *Journal of Econometrics*, 235(2):325–351, 2023.
- S. Wang, Q. Wang, and J. Zhao. Deep neural networks for choice analysis: Extracting complete economic information for interpretation. *Transportation Research Part C: Emerging Technologies*, 118:102701, 2020.
- Y. Wei and Z. Jiang. Estimating parameters of structural models using neural networks. *USC Marshall School of Business Research Paper*, 2022.
- M. Weitzman. *Optimal search for the best alternative*, volume 78. Department of Energy, 1978.
- T. G. Wollmann. Trucks without bailouts: Equilibrium product characteristics for commercial vehicles. *American Economic Review*, 108(6):1364–1406, 2018.
- M. Wong and B. Farooq. Reslogit: A residual neural network logit model for data-driven choice modelling. *Transportation Research Part C: Emerging Technologies*, 126:103050, 2021.
- M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. R. Salakhutdinov, and A. J. Smola. Deep sets. *Advances in neural information processing systems*, 30, 2017.

Appendices

A Identity Independence and Permutation Invariance of Traditional Choice Models

In this Web Appendix, we show why choice models listed in Table 1 satisfy the identity independence and the permutation invariance assumption. Identity independence is achieved when all the information about a product utilized in the model is included in its features. To show that these models satisfy the permutation invariance assumption, we first permute any two competitor products associated with a focal product. We then illustrate that such permutations do not alter the demand for the focal product, thereby establishing the models' compliance with the permutation invariance assumption.

Since all models require parametric assumptions about the utility function, for the sake of simplicity in notation, we assume that the utility of a product is a parametric function of its observed features X_{jt} and price p_{jt} , determined by a set of parameters β unless otherwise specified. Mathematically, we express the utility as:

$$v_{jt} = f(X_{jt}, p_{jt}; \beta)$$

A.1 MNL

Assuming there are J products, the market share of each product under the MNL model is

$$\pi_{jt} = g(X_{jt}, p_{jt}, \{X_{kt}, p_{kt}\}_{k \in \mathcal{S}, k \neq j}) = \frac{\exp(f(X_{jt}, p_{jt}; \beta))}{\exp(f(X_{jt}, p_{jt}; \beta) + \sum_{k \in \mathcal{S}, k \neq j} \exp(f(X_{kt}, p_{kt}; \beta)))}. \quad (\text{A1})$$

If we permute two products k_1 and k_2 in the set $\mathcal{S} \setminus \{j\}$, the only change is in the order of terms in the summation in the denominator of the market share. Since addition is commutative, the sum $\sum_{k \in \mathcal{S}, k \neq j} \exp(f(X_{kt}, p_{kt}; \beta))$ remains the same regardless of the order of k_1 and k_2 . Therefore, the value of π_{jt} remains unchanged, demonstrating that the MNL model satisfies permutation invariance.

A.2 Mixed Logit

Assuming there are J products, the market share of each product under the RCL model is

$$\pi_{jt} = g(X_{jt}, \{X_{kt}\}_{k \in \mathcal{S}, k \neq j}) = \int \frac{\exp(f(X_{jt}, p_{jt}; \beta_i))}{\exp(f(X_{jt}, p_{jt}; \beta_i) + \sum_{k \in \mathcal{S}, k \neq j} \exp(f(X_{kt}, p_{kt}; \beta_i)))} dF(\beta_i) \quad (\text{A2})$$

, where $F(\beta_i)$ is the CDF of random coefficient β_i . Similar to MNL, if we permute the position of any other two products $k_1, k_2 \in \mathcal{S} \setminus \{j\}$, it only affects the sequence of the summation in the denominator, thus the demand of product j remains the same. Therefore, it satisfies the permutation invariance.

A.3 Nested Logit

In a nested logit model, products are grouped into nests to account for the correlation in unobserved factors among products within the same nest. Consider a scenario where there are J products categorized into L nests. The market share of a product j in a nest l is calculated taking into account this nesting structure. A key aspect of showing permutation invariance in the Nested Logit model is to include the nest affiliation of each product as a feature. This is represented by an L -dimensional 0-1 vector $N_{jt} \in \mathbb{R}^{1 \times L}$, which denotes

the affiliation of product jt to one of the L nests. Each element of the vector corresponds to a specific nest, indicating the product's membership in that nest. Specifically,

$$\pi_{jt} = g(X_{jt}, p_{jt}, N_{jt}, \{X_{kt}, p_{kt}, N_{kt}\}_{k \in S, k \neq j}). \quad (\text{A3})$$

Then, the Nested Logit model can be represented as:

$$\pi_{jt} = P(jt|l)P(l) \quad (\text{A4})$$

$$P(j|l) = \frac{\exp\left(\frac{v_{jt}}{N_{jt}\sigma}\right)}{\sum_{k \in \mathcal{N}_l} \exp\left(\frac{v_{kt}}{N_{kt}\sigma}\right)} \quad (\text{A5})$$

$$P(l) = \frac{\exp\left(N_l\sigma \log\left(\sum_{k \in \mathcal{N}_l} \exp\left(\frac{v_{kt}}{N_{kt}\sigma}\right)\right)\right)}{\sum_{m=1}^L \exp\left(N_m\sigma \log\left(\sum_{k \in \mathcal{N}_m} \exp\left(\frac{v_{kt}}{N_m\sigma}\right)\right)\right)} \quad (\text{A6})$$

Where:

- \mathcal{N}_l is the set of products in nest l .
- $\sigma \in R^L$ is an L -dimensional vector where the l th element represents the scale parameter for nest l which captures the correlation in unobserved utilities among products in the same nest. Therefore, the scale parameter of product jt , which belongs to nest l , can be expressed as $N_{jt}\sigma$.
- N_m denotes a binary vector of dimension L , in which the m th element is one and all other elements are zero. Therefore, $N_m\sigma$ is the scale parameter of the nest m .

It is intuitive that when we swap the position of any two products, the probability of choosing a nest ($P(l)$) would not change because both the denominator and the numerator only depend on the nest affiliations. Moreover, $P(j|l)$ also satisfies permutation invariance because of the additive nature of the denominator in its formula, where the sum over all products in nest l remains constant despite any permutation of product positions within the nest. The permutation invariance property could be easily extended to that random coefficient nested logit following the same logic as random coefficient logit.

A.4 Latent Class Logit Model

In a latent class logit model, products can be divided into multiple latent classes and each latent class has a distinct set of parameters. Assume there are c latent classes, the utility of a product in class c is denoted by $v_{jt,c} = f(X_{jt}, p_{jt}; \beta_c)$. The probability of one product belonging to a class c is determined by the linear combination of Z_{jt} with parameters α_c .

$$\begin{aligned} \pi_{jt} &= g(X_{jt}, p_{jt}, Z_{jt}, \{X_{kt}, p_{kt}, Z_{kt}\}_{k \in S, k \neq j}) \\ &= \sum_c \frac{\exp(f(X_{jt}, p_{jt}; \beta_c)) \exp(\alpha_c Z_{jt})}{\exp(f(X_{jt}, p_{jt}; \beta_c)) + \sum_{k \in S, k \neq j} \exp(f(X_{kt}, p_{kt}; \beta_c)) \sum_c \exp(\alpha_c Z_{jt})} \end{aligned} \quad (\text{A7})$$

Similarly, the $\sum_{k \in S, k \neq j} \exp(f(X_{kt}, p_{kt}; \beta_c))$ is commutative thus the demand remains the same when we alter the order of any two products except jt . Therefore, latent class models satisfy the permutation invariance.

A.5 Consumer Inattention and Search Model

In this part, we specifically demonstrate the permutation invariance property of the consumer inattention model that we used in the numerical experiments. Our analysis could easily be extended to other models based on consumer inattention in the literature, such as [Goeree \(2008\)](#), [Turlo et al. \(2023\)](#), [Compiani \(2022\)](#) and [Joo \(2023\)](#), as well as search models, such as [Mehta et al. \(2003\)](#).

In our simulation, we generate the market shares of products by considering a segment of consumers who ignore the highest-priced product. Here we generalize the specification by assuming the portion of the inattentive consumers is a function of observed features of the highest-priced product $1 - s(X_{ht})$, where h denotes the highest-priced product. A consumer forms their consideration set with the products they pay attention to. For products in their consideration set, they use the logit function as the decision rule. We allow for each consumer to have a set of random coefficients. Then the choice probability of the highest-priced product becomes

$$\pi_{ht} = g(X_{ht}, \{X_{kt}\}_{k \in \mathcal{S}, k \neq h}) = \frac{s(X_{ht}) \exp(\beta_i X_{ht})}{f(X_{jt}, p_{jt}; \beta_i) + \sum_{k \in \mathcal{S}, k \neq h} f(X_{kt}, p_{kt}; \beta_i)} \quad (\text{A8})$$

When we permute the order of any two competitor products, the choice probability of the highest-price product remains the same. The choice probability of other products can be expressed as

$$\begin{aligned} \pi_{jt} &= g(X_{jt}, \{X_{kt}\}_{k \in \mathcal{S}, k \neq j}) \\ &= \frac{s(X_{ht}) \exp(f(X_{kt}, p_{jt}; \beta_i))}{\exp(f(X_{kt}, p_{jt}; \beta_i)) + \sum_{k \in \mathcal{S}, k \neq j} f(X_{kt}, p_{kt}; \beta_i)} + \frac{(1 - s(X_{ht})) \exp(f(X_{kt}, p_{jt}; \beta_i))}{\exp(f(X_{kt}, p_{jt}; \beta_i)) + \sum_{k \in \mathcal{S}, k \neq j, h} f(X_{kt}, p_{kt}; \beta_i)} \end{aligned} \quad (\text{A9})$$

Similar to the highest-priced product, the choice probability for other products remains unaffected when the order of competitor products is altered. This is because the feature values of the most expensive product X_h do not change even though the index h may change. Therefore, consumer inattention models are also permutation invariant.

A.6 Other Models

We have illustrated that the Multinomial Logit model, Nested Logit model, Mixed Logit model, and Latent Class Logit model all exhibit the property of permutation invariance. It becomes apparent that other models, which employ a similar random utility framework but differentiate in the probability distribution of the error term, likewise possess this permutation invariance property.

For example, this permutation invariance property extends Generalized Extreme Value (GEV) models ([Train, 2009](#)), where the error term follows the generalized extreme value distribution, and the Probit model ([Hausman and Wise, 1978](#); [Chintagunta, 2001](#)), where the error term follows a normal distribution. Although there is no straightforward closed-form representation of choice probability for these models, the choice probability can be expressed as

$$\pi_{ijt} = \Pr(u_{ijt} \geq u_{ikt}, \forall k \neq j, k \in \mathcal{S}_t), \quad (\text{A10})$$

where u_{ijt} represents the consumer i 's utility of product j in market t . Since u_{ikt} for any $k \in \mathcal{S}_t$ is parameterized by its own feature (X_{kt}), permuting any two $k \neq j, k \in \mathcal{S}_t$, does not affect the choice probability of π_{ijt} as well as the aggregate demand π_{jt} . Indeed, this shows that any models building on the random utility framework with the decision rule as stated in Equation A10 are permutation invariant.

The permutation invariance property also extends to other recently developed choice models, such as Markov Chain Choice model ([Blanchet et al., 2016](#)). The choice probability of a product jt is modeled as a

Markov Chain, which consists of two parts: arrival probability λ_{jt} and transition probability $p_{jk,t}$, defined as

$$\lambda_{jt} = f(j, \mathcal{S}_t), \quad (\text{A11})$$

$$p_{jk,t} = \begin{cases} 1, & \text{if } j = 0 \text{ and } k = 0, \\ \frac{f(j, \mathcal{S}_t \setminus \{j\}) - f(k, \mathcal{S}_t)}{f(j, \mathcal{S}_t)}, & \text{if } j, k \in \mathcal{S}_t, k \neq j, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{A12})$$

Given that both λ_{jt} and $p_{jk,t}$ only rely on jt , \mathcal{S}_t and $\mathcal{S}_t \setminus \{j\}$, permuting the order of any two products that are not the focal product jt in the market t does not change both arrival probability and transition probability. Therefore, the choice probability does not change.

B Proof of Main Results

Theorem 1. For any offer set $\mathcal{S}_t \subset \{1, 2, 3, \dots, J_t\}$, if a choice function $\pi : \{u_{ijt} : j \in \mathcal{S}_t\} \rightarrow \mathbb{R}^{|\mathcal{S}_t|}$ where u_{ijt} represents the index tuple $\{X_{jt}, p_{jt}, I_{it}, \varepsilon_{ijt}\}$ satisfies Assumption 1, 2 and 3, then there exists suitable ρ , ϕ_1 and ϕ_2 such that

$$\pi_{jt} = \rho(\phi_1(X_{jt}, p_{jt}) + \sum_{k \neq j, k \in \mathcal{S}_t} \phi_2(X_{kt}, p_{kt})),$$

Proof. The sufficiency follows by observing that the function $\pi_{jt} = \rho(\phi_1(X_{jt}, p_{jt}) + \sum_{k \in \mathcal{S}_t \setminus \{j\}} \phi_2(X_{kt}, p_{kt}))$ satisfies assumption 2 and 3. To prove necessity, first consider $\mathbb{E} = \{2n \mid n \in \mathbb{N}\}$ and $\mathbb{O} = \{2n+1 \mid n \in \mathbb{N}\}$ as the set of all even and odd natural numbers, respectively. Next, to show that all choice functions satisfying assumptions 2 (identity independence) and 3 (permutation invariance) can be decomposed in the above manner, we begin by noting that there must be a mapping from the elements to the set of even number and odd numbers respectively, since the elements belong to a countable universe \mathbb{C}^k . Let these mappings be denoted by $c^e : \mathbb{C}^k \rightarrow \mathbb{E}$ and $c^o : \mathbb{C}^k \rightarrow \mathbb{O}$. Now if we let $\phi_1(X_{jt}) = 4^{-c^e(X_{jt}, p_{jt})}$ and $\phi_2(X_{kt}, p_{kt}) = 4^{-c^o(X_{kt}, p_{kt})}$ then $\phi_1(X_{jt}, p_{jt}) + \sum_{k \in \mathcal{S}_t \setminus \{j\}} \phi_2(X_{kt}, p_{kt})$ constitutes a unique representation for every product j and competing assortment $\mathcal{S}_t \setminus \{j\}$. Now a function $\rho : \mathbb{R} \rightarrow \mathbb{R}$ can always be constructed such that $\pi_{jt} = \rho(\phi_1(X_{jt}, p_{jt}) + \sum_{k \in \mathcal{S}_t \setminus \{j\}} \phi_2(X_{kt}, p_{kt}))$.

Proposition 1. *If Assumptions 5-7 are satisfied then for $V = E[\{m(w, \pi_0(z; \gamma_0)) - \theta_0 + \alpha_0(z)(y - \pi_0(z; \gamma_0))\}^2]$,*

$$\sqrt{n}(\hat{\theta} - \theta_0) \xrightarrow{D} N(0, V), \hat{V} \xrightarrow{P} V.$$

Proof. To show the asymptotic normality we will first verify the Assumptions 1-3 of [Chernozhukov et al. \(2022a\)](#), from now on CEINR, with $g(w, \pi(z; \gamma), \theta) = m(w, \pi(z; \gamma)) - \theta$ and $\phi(w, \pi(z; \gamma), \alpha(z), \theta) = \alpha(z) \cdot (y - \pi(z; \gamma))$. Using Taylor series expansion, Assumption 6 and $\|\hat{\pi}(z; \gamma) - \pi_0(z; \gamma)\| \xrightarrow{P} 0$ we have,

$$\begin{aligned}
& \int \|g(w, \hat{\pi}(z_i; \hat{\gamma}), \theta_0) - g(w, \pi_0(z_i; \gamma_0), \theta_0)\|^2 \mathcal{P}_0(dw) \\
&= \int \|m(w, \hat{\pi}(z_i; \hat{\gamma})) - m(w, \pi_0(z_i; \gamma_0))\|^2 \mathcal{P}_0(dw) \\
&\leq C \int \|\hat{\pi}(z_i; \hat{\gamma}) - \pi_0(z_i; \gamma_0)\|^2 \mathcal{P}_0(dw) \\
&\leq C \int \|\hat{\pi}(z_i; \hat{\gamma}) - \hat{\pi}(z_i; \gamma_0) + \hat{\pi}(z_i; \gamma_0) - \pi_0(z_i; \gamma_0)\|^2 \mathcal{P}_0(dw)
\end{aligned}$$

By the triangle inequality

$$\begin{aligned}
&\leq C \int \|\hat{\pi}(z_i; \hat{\gamma}) - \hat{\pi}(z_i; \gamma_0)\|^2 \mathcal{P}_0(dw) \\
&\quad + C \int \|\hat{\pi}(z_i; \gamma_0) - \pi_0(z_i; \gamma_0)\|^2 \mathcal{P}_0(dw) \\
&\quad + C \int \|\hat{\pi}(z_i; \hat{\gamma}) - \hat{\pi}(z_i; \gamma_0)\| \|\hat{\pi}(z_i; \gamma_0) - \pi_0(z_i; \gamma_0)\| \mathcal{P}_0(dw) \xrightarrow{P} 0 \quad (\text{A13})
\end{aligned}$$

The first term converges in probability to 0 by Taylor series expansion.

Also by Assumption 5 i) and ii), and as just showed $\|\hat{\pi}(z; \hat{\gamma}) - \pi_0(z; \gamma_0)\| \xrightarrow{P} 0$,

$$\begin{aligned}
\int \|\phi(w, \hat{\pi}(z; \hat{\gamma}), \alpha_0, \theta_0) - \phi(w, \pi_0(z; \gamma_0), \alpha_0, \theta_0)\|^2 \mathcal{P}_0(dw) &= \int \|\alpha_0(z)(\pi_0(z; \gamma_0) - \hat{\pi}(z; \hat{\gamma}))\|^2 \mathcal{P}_0(dw) \\
&\leq C \int \|(\pi_0(z; \gamma_0) - \hat{\pi}(z; \hat{\gamma}))\|^2 \mathcal{P}_0(dw) \\
&\leq C \|\hat{\pi}(z; \hat{\gamma}) - \pi_0(z; \gamma_0)\|^2 \xrightarrow{P} 0
\end{aligned} \quad (\text{A14})$$

Also by Assumption 5 i) and $\|\hat{\alpha} - \alpha_0\| \xrightarrow{P} 0$, we have,

$$\begin{aligned}
\int \|\phi(w, \pi_0(z; \gamma_0), \hat{\alpha}, \tilde{\theta}) - \phi(w, \pi_0(z; \gamma_0), \alpha_0, \theta_0)\|^2 \mathcal{P}_0(dw) &= \int \|(\hat{\alpha}(z) - \alpha_0(z))(y - \pi_0(z; \gamma_0))\|^2 \mathcal{P}_0(dw) \\
&\leq C \int \|\hat{\alpha} - \alpha_0\|^2 \mathcal{P}_0(dw) \\
&\leq C \|\hat{\alpha} - \alpha_0\|^2 \xrightarrow{P} 0
\end{aligned} \quad (\text{A15})$$

This satisfies Assumption 1 parts i), ii), and iii) of CEINR.

Next, consider

$$\begin{aligned}
\hat{\Delta}(w) &:= \phi(w, \hat{\pi}(z; \hat{\gamma}), \hat{\alpha}, \tilde{\theta}) - \phi(w, f_0, \hat{\alpha}, \tilde{\theta}) - \phi(w, \hat{\pi}(z; \hat{\gamma}), \alpha_0, \theta_0) + \phi(w, f_0, \alpha_0, \theta_0) \\
&= -[\hat{\alpha}(z) - \alpha_0(z)][\hat{\pi}(z; \hat{\gamma}) - \pi_0(z; \gamma)].
\end{aligned}$$

Then by the Cauchy-Schwartz inequality, and Assumptions 6 i) and ii)

$$\begin{aligned}
E \left[\hat{\Delta}(w) \right] &= \int - [\hat{\alpha}(z) - \alpha_0(z)] [(\hat{\pi}(z; \hat{\gamma}) - \pi(z; \gamma))] \mathcal{P}_0(dz) \\
&\leq \|\hat{\alpha} - \alpha_0\| \|(\hat{\pi}(z; \hat{\gamma}) - \pi(z; \gamma))\| = o_p \left(\frac{1}{\sqrt{n}} \right)
\end{aligned} \tag{A16}$$

Also since $\hat{\alpha}(z)$ and $\alpha(z)$ is bounded, we have

$$\begin{aligned}
\int \left\| \hat{\Delta}(w) \right\|^2 \mathcal{P}_0(dw) &= \int [\hat{\alpha}(z) - \alpha_0(z)]^2 [(\hat{\pi}(z; \hat{\gamma}) - \pi_0(z))]^2 \mathcal{P}_0(dz) \\
&\leq C \int [(\hat{\pi} - \pi_0(z))]^2 \mathcal{P}_0(dz) \xrightarrow{p} 0
\end{aligned} \tag{A17}$$

Thus Equation A16 and Equation A17 verify Assumption 2 i) of CEINR.

Next Assumption 3 of CEINR is satisfied through Assumption 7. Thus Assumptions 1-3 of CEINR are satisfied. Thus asymptotic normality follows by Lemma 15 of CEINR and the Lindberg-Levy central limit theorem.

Finally, we know $\theta \xrightarrow{p} \theta_0$. And thus we have, $\int \left\| g \left(w, \hat{\pi}(z; \hat{\gamma}), \hat{\theta} \right) - g \left(w, \hat{\pi}(z; \hat{\gamma}), \theta_0 \right) \right\|^2 \mathcal{P}_0(dw) \xrightarrow{p} 0$

To get the second conclusion, we need to show that \hat{V} is a consistent estimator of V . To show this, we closely follow Chernozhukov et al. (2021). Let $\psi_i = \psi_0(w_i)$ and consider

$$\hat{V} = \frac{1}{n} \sum_{i=1}^n \hat{\psi}_i^2 = \frac{1}{n} \sum_{i=1}^n (\hat{\psi}_i - \psi_i)^2 + \frac{2}{n} \sum_{i=1}^n (\hat{\psi}_i - \psi_i) \psi_i + \frac{1}{n} \sum_{i=1}^n \psi_i^2$$

hence, by re-arranging the terms and Cauchy-Schwarz inequality,

$$\hat{V} - V = \frac{1}{n} \sum_{i=1}^n (\hat{\psi}_i - \psi_i)^2 + \frac{2}{n} \sum_{i=1}^n (\hat{\psi}_i - \psi_i) \psi_i \leq \frac{1}{n} \sum_{i=1}^n (\hat{\psi}_i - \psi_i)^2 + 2 \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{\psi}_i - \psi_i)^2} \sqrt{\frac{1}{n} \sum_{i=1}^n \psi_i^2}.$$

Using the triangle inequality, for $i \in I_\ell$,

$$\left(\hat{\psi}_i - \psi_i \right)^2 \leq C \sum_{j=1}^4 R_{ij} = C \sum_{j=1}^3 R_{ij} + o_p(1)$$

where

$$\begin{aligned}
R_{i1} &= [m(w_i, \hat{\pi}_\ell(z_i; \hat{\gamma}_\ell)) - m(w_i, \pi_0(z_i; \gamma_0))]^2, \\
R_{i2} &= \hat{\alpha}_\ell^2(z_i) [\hat{\pi}_\ell(z_i; \hat{\gamma}_\ell) - \pi_0(z_i; \gamma_0)]^2, \\
R_{i3} &= [\hat{\alpha}_\ell(z_i) - \alpha_0(z_i)]^2 [y_i - \pi_0(z_i; \gamma_0)]^2, \\
R_{i4} &= \left(\hat{\theta} - \theta_0 \right)^2.
\end{aligned}$$

We already showed consistency, so $R_{i4} \xrightarrow{p} 0$.

Let $I_{-\ell}$ denote observations not in I_ℓ . By Markov's inequality, for some $\delta > 0$,

$$\mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n (\hat{\psi}_i - \psi_i)^2 > \delta \right) \leq \frac{\mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (\hat{\psi}_i - \psi_i)^2 \right]}{\delta}$$

Note that the cross-fitting allows us to write

$$\mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (\hat{\psi}_i - \psi_i)^2 \right] \leq \mathbb{E} \left[\frac{C}{n} \sum_{\ell=1}^L \sum_{i \in I_\ell} \sum_{j=1}^3 R_{ij} \right] + o_p(1) = C \sum_{\ell=1}^L \frac{n_\ell}{n} \sum_{j=1}^3 \mathbb{E} [\mathbb{E} [R_{ij} \mid I_{-\ell}]] + o_p(1).$$

We already showed,

$$\mathbb{E} [R_{i1} \mid I_{-\ell}] = \int [m(w_i, \hat{\gamma}_\ell) - m(w_i, \gamma_0)]^2 F_0(dw) = o_p(1)$$

Next by triangle inequality, we have

$$\begin{aligned} \mathbb{E} [R_{i2} \mid \mathcal{W}_{-l}] &= \int \hat{\alpha}_l^2 [\hat{\pi}_l(z_i; \hat{\gamma}_l) - \pi_0(z_i; \gamma_0)]^2 F_0(dz) \\ &= \int [\hat{\alpha}_l + \alpha_0 - \alpha_0]^2 [\hat{\pi}_l(z_i; \hat{\gamma}_l) - \pi_0(z_i; \gamma_0)]^2 F_0(dz) \\ &\leq \int [\hat{\alpha}_l - \alpha_0]^2 [\hat{\pi}_l(z_i; \hat{\gamma}_l) - \pi_0(z_i; \gamma_0)]^2 F_0(dz) \\ &\quad + \int [\alpha_0]^2 [\hat{\pi}_l(z_i; \hat{\gamma}_l) - \pi_0(z_i; \gamma_0)]^2 F_0(dz) \\ &\leq O_p(1) \int [\hat{\pi}_l(z_i; \hat{\gamma}_l) - \pi_0(z_i; \gamma_0)]^2 F_0(dz) \xrightarrow{p} 0 \end{aligned}$$

Finally, we have

$$\begin{aligned} \mathbb{E} [R_{i3} \mid I_{-\ell}] &= \mathbb{E} \left[\mathbb{E} \left[[\hat{\alpha}_\ell(z_i) - \alpha_0(z_i)]^2 [y_i - \pi_0(z_i; \gamma_0)]^2 \mid z_i, I_{-\ell} \right] \mid I_{-\ell} \right] \\ &= \mathbb{E} \left[[\hat{\alpha}_\ell(Z_i) - \alpha_0(Z_i)]^2 \mathbb{E} \left[[y_i - \pi_0(z_i; \gamma_0)]^2 \mid Z_i \right] \mid I_{-\ell} \right] \\ &\leq C \|\hat{\alpha}_\ell - \alpha_0\|^2 \xrightarrow{p} 0. \end{aligned}$$

As a result,

$$\frac{1}{n} \sum_{i=1}^n (\hat{\psi}_i - \psi_i)^2 \xrightarrow{p} 0$$

Thus, we have $\hat{V} \xrightarrow{p} V$

Also by Assumption 9 and iterated expectations

$$\begin{aligned} \mathbb{E} [R_{i3} \mid \mathcal{W}_{-\ell}] &\leq \int \{\hat{\alpha}_\ell(z) - \bar{\alpha}(z)\}^2 \mathbb{E} [(y - \bar{\pi}(z))^2 \mid Z = z] F_Z(dz) \\ &\leq C \int \{\hat{\alpha}_\ell(z) - \bar{\alpha}(z)\}^2 F_z(dz) = C \|\hat{\alpha}_\ell - \bar{\alpha}\|^2 = o_p(1). \end{aligned}$$

□

C Hyperparameter Space for Tuning Non-parametric Estimator Benchmark

Hyperparameter	Space
Number of hidden layers	[3, 4, 5]
Number of nodes in each layer	[64, 128, 256]
Learning rate	[1e-2, 1e-3, 1e-4]
Number of epochs	[1, 2, 4]

Table A1: Hyperparameter Space for Tuning Non-parametric Estimator Benchmark

D Distribution of Features and Coefficients in Numerical Experiments

	MNL	RCL
p_{jt}	$U[0, 4]$	$U[0, 4]$
X_{jt}	$N(0, 1)$	$N(0, 1)$

Table A2: Distribution of Features

	MNL	RCL
α_i	-1	$N(-1, 1)$
β_{ik}	1	$N(\mu_{\beta_k}, 1)$

Notes : $\mu_{\beta_k} \sim N(0, 1/2d)$

Table A3: Distribution of Coefficients

E An Numerical Experiment of Endogeneity

In §3.2, we show how our model handles endogeneity in theory. Here, we provide a simulation that demonstrates the performance of our model handling endogeneity.

Let the utility u_{ijt} that consumer i in market t derives from product j as the following linear function

$$u_{ijt} = \alpha_i p_{jt} + \beta_i X_{jt} + \gamma_i \mu_{jt} + \varepsilon_{ijt}, \quad (\text{A18})$$

where μ_{jt} is the unobservable that correlates p_{jt} and ε_{ijt} is i.i.d. Type-I extreme value distributed. Specifically, without loss of generality, we specify the correlation between p_{jt} and μ_{jt} as

$$p_{jt} = \delta_1 IV_{jt} + \delta_2 \mu_{jt}, \quad (\text{A19})$$

where IV_{jt} is the exogenous instrumental variable.

In this simulation, we first generate X_{jt} and μ_{jt} , following the same distribution as the X_{jt} in our baseline model, and IV_{jt} , following the same distribution as our as p_{jt} in our baseline model, as discussed in Web Appendix D. We next generate p_{jt} using Equation A19. Finally, we generate market shares assuming that the true model is RCL. For simplicity, we let $\delta_1 = \delta_2 = 1$. In addition, α_i and β_i follow the same

distribution as the baseline simulations and γ_i follows the same distribution as β_i , as stated in Web Appendix D. We consider a case with 10 products, 100 markets, and 10 non-price features (in addition to price), the same as our baseline simulation in §5.1.1. In the estimation, we let μ_{jt} be an unobservable. We use the OLS linear regression as our first-stage regressor.

We consider two cases for comparison:

- **Exogeneous Benchmark:** Uses the exact same DGP but treat μ_{jt} as observables to researchers in estimation. Since the coefficients of price and market shares are the same as the endogeneity case, true elasticities are the same for these two DGPs. This gives a benchmark performance with the same data under the assumption that endogeneity is not present.
- **Ignoring Endogeneity:** Directly trains the model using only the observed features X_{jt} and p_{jt} without considering the endogeneity problem. Note that, even though the endogeneity problem is ignored, this does not mean the predictive performance in market shares would be low for this case because the model could be overfitted. However, the elasticities will be biased when the endogeneity is ignored.

Following the routine in our main text, we simulate 20 times for the same DGP with different parameters and features and report the predictive performance on market share ($\hat{\pi}_{jt}$), own price elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{jt}/p_{jt}}$), and cross-price elasticity ($\frac{\partial \hat{\pi}_{jt}/\pi_{jt}}{\partial p_{k \neq jt}/p_{k \neq jt}}$). We report the performance of our model in Table A4. When we use the control function to correct for endogeneity (Row 1: Our Method), our model underperforms slightly compared to the case where μ_{jt} is treated as an observable (Row 2: Exogenous Benchmark), on all three metrics. However, when we ignore the endogeneity issue and apply our model (Row 3: Ignore Endogeneity), although the predictive performance of market shares is not bad, the estimation of own- and cross-elasticities are significantly biased, with the MAEs being 10–25 times higher than the cases where we account for endogeneity. This demonstrates both the importance of accounting for endogeneity as well as the effectiveness of our method in handling this problem in real settings.

	Market shares		Own-elasticity		Cross-elasticity	
	MAE	RMSE	MAE	RMSE	MAE	RMSE
Our Method	0.0256	0.0194	0.2307	0.3516	0.0669	0.1033
Exogeneous Benchmark	0.0254	0.0195	0.2290	0.3469	0.0657	0.1039
Ignore Endogeneity	0.0263	0.0202	2.3832	5.1901	1.9751	4.4306

Table A4: Model Performance of Endogeneity Case

F Details in Adopting the “MLIV” Method

Following Singh et al. (2020), we perform the steps below to construct the machine-learning-based IV (MLIV) and use them to estimate $\hat{\gamma}$ to control for endogeneity in prices.

- **Step 1: Data Partition** We randomly split the data set, \mathcal{D} , into three separate partitions of markets, each denoted as D_l . Each market is exclusively assigned to only one partition. For each partition, we define its complement set, D_l^c , as the subset of data in \mathcal{D} that is not included in D_l .
- **Step 2: Cross-fitting** For each partition l , we first estimate a linear regression model on the complement data set, D_l^c , using the Lasso method with hyperparameters tuned by 3-fold cross-validation. As discussed in section 3.3, we need the estimator of γ to converge at $n^{-1/2}$ rate, a similar result

that bounds the prediction error of the lasso estimator has been established in [Chatterjee and Jafarov \(2015\)](#). Then, we use this trained model to predict the outcomes (prices) of the D_t . We denote the fitted value as \hat{f}_t , which is essentially the MLIV.

- **Step 3: First-stage Regression** We estimate the first-stage estimator γ and residual μ using the MLIV as the only predictor following step 1 in Section 4.

As a supplement to our main result, we also run our model using non-machine learning-based IVs. Similar to Figure 4 in the main text, we present the estimated own-elasticity of our model without IV, with BLP-style IVs, with differentiation IVs, and with MLIV in Figure A1. In Figure A1b, even when IVs are applied, the persistence of many positive own-elasticities suggests the weakness of the BLP style IVs. Furthermore, we apply the differentiation IVs ([Gandhi and Houde, 2019](#)), which use exogenous measures of differentiation and provide a more robust instrument compared to the conventional BLP IVs. As one can see from Figure A1c, the use of differentiation IV provides a more realistic estimation of own-elasticities, strengthening the issue of weak instruments of the BLP style IVs. We also include the distributions of the estimated own- and cross-elasticities obtained from our model using different sets of IVs in Figure A2.





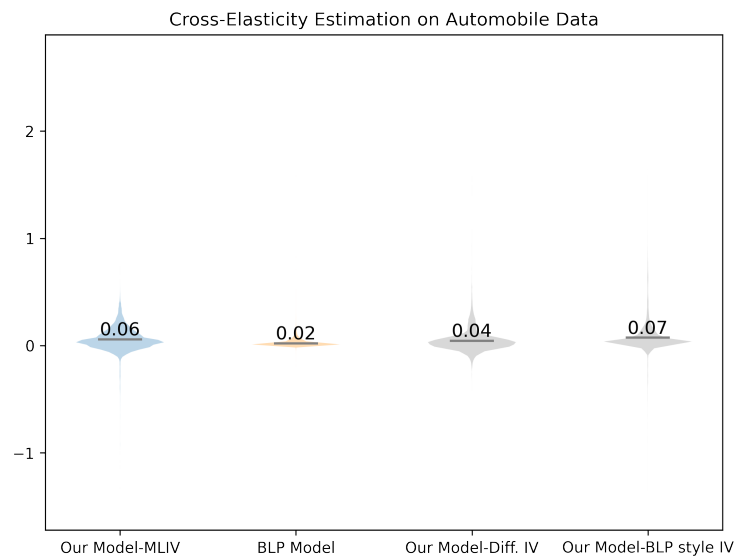
Figure A1: Elasticity Estimation Comparison

Figure A1 presents the estimated own-elasticity of our model without IV, with BLP Style IVs, with differentiation IVs and with MLIV. The x-axis represents the price of the focal product, while the y-axis shows the product’s own-elasticity. Each point corresponds to a product in a market, resulting in 2,217 observations. We report the estimated elasticity based on the same price variation used in the BLP paper (a 1,000-dollar change).

In addition, we also perform a weak instrument test on both BLP Style IVs and the MLIV and report the F-statistics and p-value in Table A5. Both BLP Style IVs and MLIV pass the weak instrument tests.



(a) Own-Elasticity Estimation (Our Model vs. BLP Model)



(b) Cross-Elasticity Estimation (Our Model vs. BLP Model)

Figure A2: Elasticity Estimation Comparison

Note: Figure A2 illustrates the distributions of the estimated own- and cross-elasticities obtained from our model (using different sets of IVs) and the BLP model. The filled areas in the violin plots represent the complete range of the elasticities, while the text labels indicate the mean values.

	F-statistic	P-value
BLP Style IVs	241.5	< 1e-8
MLIV	280.9	<1e-8

Table A5: Weak Instrument Test