Elasticity and Curvature of Discrete Choice Demand Models *

Eugenio J. Miravete† Katja Seim‡ Jeff Thurk§

May 2025

Abstract

We study the determinants of cost pass-through in differentiated product markets. Random utility models of demand, such as mixed logit, place limited restrictions on customer substitution patterns while constraining demand curvature in less known ways. We show that the shape of the distribution of customer preferences determines cost pass-through. Common functional form assumptions for this distribution lead to biased estimates of both pass-through and substitution. We offer a flexible and parsimonious unit-demand specification that accommodates both log-concave demands (incomplete pass-through) and log-convex demands (over-shifted pass-through), for quasi-linear preferences and for preferences that accommodate income effects. Instruments and estimation are straightforward, and Monte Carlo analysis validates our ability to recover the underlying demand curvature. We find large biases from shape restrictions using well-known breakfast cereal and automobile data.

Keywords: Market Power, Substitution, Pass-Through, Demand Curvature.

JEL Codes: C51, D43, L13, L41, L66

* We thank Simon Anderson, Christopher Conlon, Jeremy Fox, Eli Liebman, Monika Mrázová, Philipp Schmidt-Dengler, Stephen Ryan, Paulo Somaini, and Alan Sorensen for valuable discussions. All remaining errors are only our own.

† The University of Texas at Austin, Department of Economics, Austin, TX 78712; and CEPR. E–mail: *eugenio@utexas.edu*; http://www.eugeniomiravete.com

‡ Yale University, Department of Economics and School of Management, New Haven, CT 06520; CEPR and NBER. E–mail: *katja.seim@yale.edu*

§ University of Georgia, Department of Economics, Athens, GA 30602. E–mail: *jeff.thurk@uga.edu*; http://www.jeffthurk.com

1 Introduction

Demand curvature influences cost pass-through and key conclusions to many substantive economic questions, including the ability of digital platforms such as Amazon.com to influence the surplus division between third-party sellers and consumers (Gutierrez, 2022), the welfare implications of uniform pricing in settings ranging from consumer packaged goods (DellaVigna and Gentzkow, 2019) to consumer financial products (Cuesta and Sepúlveda, 2021), or the price effects of cost efficiencies in horizontal mergers (Jaffe and Weyl, 2013; Miller, Remer, Ryan and Sheu, 2015). Demand curvature is also central to the incidence of taxes and exchange rates in non-competitive industries (Weyl and Fabinger, 2013) and to the role of regulation in controlling externalities (Fabra and Reguant, 2014; Miller, Osborne and Sheu, 2017). Other studies show how restrictive functional forms bias pass-through predictions. Bulow and Pfleiderer (1983) famously caution against rigid specifications in tobacco markets, while Froeb, Tschantz and Werden (2005) show constant elasticity of substitution (*CES*) models inflate merger synergies.

We evaluate the conditions under which discrete choice unit-demand models can capture both realistic substitution and pass-through patterns. This model is the workhorse framework within empirical industrial organization (IO) and nests constant elasticity of substitution (*CES*) models commonly used in macroeconomics and international trade. It is well known that the mixed logit (*ML*) model, in particular, can capture realistic substitution patterns across heterogeneous consumers (McFadden and Train, 2000). This flexibility is key to measuring the closeness of competition between products, predicting diversion in response to a merger-induced price change, or identifying collusion among firms. However, understanding the determinants of pass-through (demand curvature) in discrete choice models is less developed, as is the interaction between substitution and pass-through. Berry and Haile (2021), for example, state:

> ...[S]ubstitution patterns drive answers to many questions of interest—e.g., the sizes of markups or outcomes under a counterfactual merger. However, other kinds of counterfactuals can require flexibility in other dimensions. For example, “pass-through” (e.g., of a tariff, tax, or technologically driven reduction in marginal cost) depends critically on second derivatives of demand. It is not clear that a mixed-logit model is very flexible in this dimension.

Motivating Examples. We illustrate the importance of incorporating preference heterogeneity using the well-known simulated ready-to-eat cereal data from Nevo (2000) and US new automobile purchase data from Berry, Levinsohn and Pakes (1999). In both examples, we show that modeling choices have a limited impact on estimated own-price elasticities and a large impact on estimated demand curvatures. As demand curvature plays a key role in determining firm price responses to a change in marginal cost in settings with market power (Cournot, 1838;Weyl and Fabinger, 2013), these examples suggest that careful modeling of preference heterogeneity is an important ingredient in building a model that delivers robust empirical predictions.

We focus on pairs of elasticity and curvature as descriptive statistics of the shape of demand. Although the own-price elasticity is likely familiar to the reader as a simple measure of market power, demand curvature as a simple measure of cost pass-through may not be. In Section 2, we formally define demand curvature and its connection to cost pass-through. Beyond curvature, the observed pass-through rate will also depend on competition and the substitution patterns of the estimated demand, something that we address later in the paper.

Figure 1 presents scatter plots of estimated pairs of own-price elasticity (ε) and curvature (ρ) using multinomial logit (*MNL*) and mixed logit (*ML*) models. Panels A and B summarize estimates of cereal demand using data from Nevo (2000).[1] While the two specifications deliver similar average elasticities, the demand curvature results are very different. This indicates that the two models will predict very different counterfactual pricing equilibria. We observe similar results in Panels (C) and (D) using data on new automobile purchases from Berry et al. (1999).

Foreshadowing later results, the patterns in Figure 1 reflect modeling choices, not data variation. In *MNL* models, demand curvature is determined by market shares before estimation and is truncated at one. Adding flexibility by using nested logit – a popular choice for antitrust – allows estimated demand to decrease demand curvature (i.e., dots move to the left), but curvature remains truncated at one. Adding heterogeneity in non-price preferences similarly decreases demand curvature, while still maintaining truncation at one. Adding heterogeneity in price sensitivity is necessary to accommodate a demand curvature greater than one. Introducing price sensitivity properly is crucial. Specifically, the odd shape of

[1] We estimate each model using Nevo's (2000) original set of Hausman-style price instruments. The full description of the specifications and estimates are reported in Appendix A.

Figure 1: Example Elasticity and Curvature Estimates

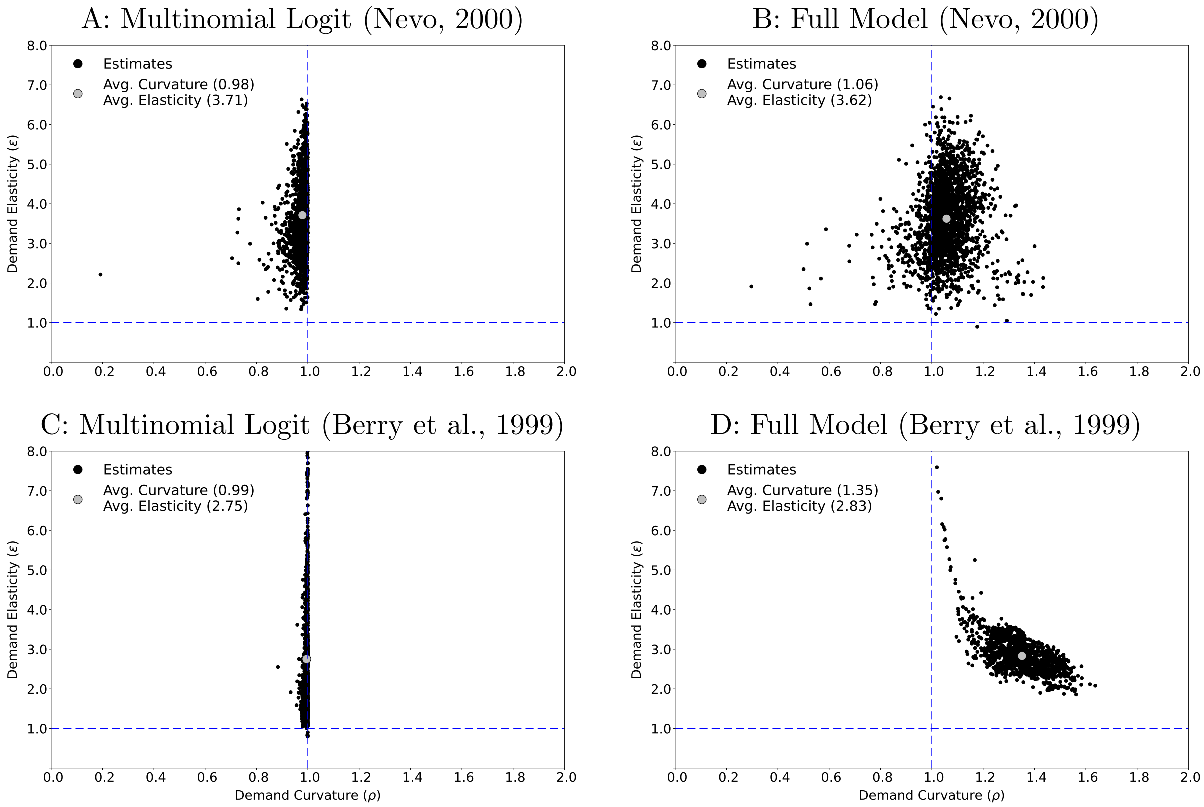

Figure Notes: Black dots represent the estimated own-price elasticity and curvature for a sample product. The gray dot corresponds to the average elasticity and curvature. Demand estimates in the top panels are based on the simulated ready-to-eat cereal data from Nevo (2000), while the bottom panels use data from Berry et al. (1999). All estimates use best practices for mixed logit estimation (Conlon and Gortmaker, 2020).

the dots in Panel (D) reflects the distribution of price sensitivity, which is driven by the assumption of Cobb-Douglas utility and the income distribution.

Contributions. We demonstrate that *ML* discrete-choice unit-demand models can capture both realistic substitution and pass-through patterns. First, we identify how different components of customer preferences influence the shape of mixed-logit (*ML*) demand. We do so by adopting the "demand manifold" approach of Mrázová and Neary (2017) and focus on the set of achievable pairs of demand elasticity and curvature as sufficient statistics for the shape of demand.[2] We are, however, limited to analyzing the demand manifold of a

[2] While Mrázová and Neary (2017) address the behavior of elasticity and curvature for different continuous demand systems (e.g., *CES*, Pollak, translog) in a single-product monopoly model, we instead evaluate how components of mixed-logit demand influence the relationship between elasticity and curvature in a discrete choice framework suitable for differentiated products oligopoly models.

single-product monopolist. Focusing on this case is valuable for highlighting features of the shape of demand that aid in robust empirical work; we consider multi-product oligopoly competition in several empirical case studies.

We show that the specification of preference heterogeneity – the shapes of the mixing distributions of the responsiveness to product attributes and price – determines the set of achievable elasticity-curvature pairs and, therefore, the shape of demand. Heterogeneous tastes over product attributes lower demand curvature and pass-through relative to *MNL*. Recognizing heterogeneous price sensitivity increases demand curvature and pass-through. Here, the skewness of the price mixing distribution plays a vital role. We consider three approaches to specifying idiosyncratic price responsiveness: distributional assumptions for unobserved heterogeneity in price sensitivity, observable consumer heterogeneity via demographic-price interactions, and heterogeneous income effects.

We demonstrate that discrete choice not only nests *CES* demand – the dominant framework in the macro and international trade literature – in the case of continuous demand (Anderson, de Palma and Thisse, 1992) but also in the case of unit-demand. Where *CES* and *ML* differ is in equilibrium pass-through under oligopoly (Head and Mayer, 2025). Since *CES* demand restricts curvature to be greater than one, it generates a pass-through that is always over-shifted and independent of the number of competitors. In contrast, competition with *ML* demand limits pass-through. Hence, while the *ML* framework provides a natural and intuitive foundation for *CES* demand and may generate identical estimates of demand elasticity and curvature, *CES* predicts price responses which are too large; a finding consistent with estimated merger synergies in Froeb et al. (2005).

We provide guidance for empirical work by highlighting the restrictions imposed by common modeling choices for preference heterogeneity, which introduce biases not just into estimates of demand and cost pass-through, but also counterfactual pricing equilibria. Depicting estimated product-level demand in elasticity-curvature space for empirical applications – as in the motivating example in Figure 1 above – helps identify possible restrictions imposed by the chosen preference specification on the shape of demand. We offer an easy and parsimonious way to modulate how correlates of demand heterogeneity, such as consumer demographics, interact with product characteristics and price, thereby increasing the range of feasible pairs of elasticity and demand curvature that the model accommodates, allowing the data patterns to dictate the shape of demand. Our approach nests *ML*, nested logit, and *MNL* models while recovering demand parameters using a standard generalized method of moments estimator. Identification is straightforward and derives from data moments that trace price responses and consumption patterns across distributions of customer demographics.

Monte Carlo simulation demonstrates that our proposed instrumentation strategy recovers the underlying mixing distributions and correctly estimates counterfactual price and welfare effects. Consistent with our motivating examples in Figure 1, misspecification bias in demand curvature can be large even when demand elasticities are largely unaffected. Such biases would have significant implications for the estimated distributional consequences of a particular government policy (e.g., , tariffs, subsidies, health insurance) or of supply chain interruptions and cost shocks. These simulation results also suggest that insights derived from demand manifolds of a single-product monopolist translate to empirical contexts.

As an illustration, we estimate automobile demand using the data from Berry et al. (1999) under alternative assumptions on the strength of income effects of heterogeneity on price sensitivities. We evaluate a hypothetical $1,000 subsidy for consumers who buy a domestically produced vehicle. The average estimated vehicle-level cost pass-through ranges from 0.99 to 1.79, depending on how income effects modulate the price sensitivity distribution. When price sensitivity is inversely proportional to consumer income, subsidy pass-through and predicted price declines of domestic cars are most pronounced, likely overstating the true effectiveness of trade policy. These insights extend naturally to other environments, such as the case of electric vehicle subsidies introduced in the *Inflation Reduction Act of 2022*. This example focuses on the empirical specification of income effects that are important for purchases of large durable goods such as cars. We note, however, that demand curvature and cost pass-through are determined by the shape of the distributions that define customer preferences more generally, whether these are connected to observable demographics or unobservable random taste variation. How to best specify demand heterogeneity in any particular application is a function of the identifying variation available in the data.

Alternative Approaches. We address the role of distributional assumptions in determining the shape of discrete choice demand and provide a tractable empirical approach to modeling demand flexibly for a broad range of empirical settings. Our work complements Compiani (2022), who also focuses on estimating demand flexibly but uses a non-parametric approach. This solution places fewer restrictions on the shape of demand than our environment, but it suffers from the curse of dimensionality. Our treatment of income effects is more general than that of Griffith, Nesheim and O'Connell (2018) and focuses on unit demand models. Alternative demand-side approaches have extended the range of feasible curvatures by adopting a discrete-continuous choice framework, where heterogeneous consumers choose either a budget allocation for a given product (Adão, Costinot and Donaldson, 2017; Björnerstedt and Verboven, 2016; Head and Mayer, 2025) or fractional units of the same product (Anderson and de Palma, 2020; Birchall, Mohapatra and Verboven, 2025).

Finally, our focus on demand specification ignores the effect of supply-side frictions on cost pass-through. For example, menu costs in adjusting price may increase or decrease the pass-through implied by demand curvature under static Nash-Bertrand pricing alone (Conlon and Rao, 2020). Supply-side considerations will depend upon the empirical setting, whereas accurately capturing the shape of demand is a necessary condition for understanding many aspects of modern empirical work, such as analyses of mergers, taxation, tariffs, cost shocks, exchange rates, and price discrimination.

Outline. We introduce the demand manifold framework in Section 2 and characterize the demand manifold for the general *ML* model in Section 3. We show mathematically how features of the mixing distributions used in consumer preferences determine the shape of consumer demand, which we represent through the relationship between elasticity and curvature. We then evaluate the implications of different quasi-linear preference specifications for curvature and elasticity in Section 4. Section 5 addresses estimating and identifying heterogeneity in price sensitivity and non-price characteristics. Here, we describe our proposed instrumentation strategy and investigate its properties in Monte Carlo analyses of both discrete choice models with quasi-linear preferences and choice models with income effects. Section 6 concludes by summarizing our contributions and discussing empirical settings beyond trade policy and electric vehicle subsidies where we think adding demand flexibility is important. Additional results and derivations are reported in the Appendices.

2 A Primer on Demand Manifolds

In this section, we introduce the concept of a demand manifold, a smooth relationship between demand elasticity and curvature consistent with profit maximization. Mrázová and Neary (2017) provide an excellent formal derivation of demand manifolds and their properties for a wide range of continuous demand specifications. We employ demand manifolds to assess the flexibility of alternative preference specifications in the context of discrete-choice demand, highlighting relevant issues that relate to the estimation of mixed-logit demand from an applied perspective.

We begin by discussing the demand manifold for a single-product monopolist. We use this setup in Sections 3 and 4 to illustrate graphically the properties of common discrete-choice demand specifications. Next, we discuss demand sub-convexity, which we impose on the demand systems in these analyses to ensure the existence of well-behaved pricing equilibria and comparative statics. Demand sub-convexity weakly limits the feasible elasticity and curvature combinations by ensuring demand becomes more elastic at higher prices; i.e., *Marshall's Second Law of Demand*.

2.1 Single-Product Monopoly

Consider the case of a single-product monopolist with constant marginal cost c. The monopolist sets the price p that maximizes profits $\Pi(p) = (p - c) \cdot q(p)$ and the following necessary condition holds:

$$\Pi_p(p) = q(p) + (p - c) \cdot q_p(p) = 1 - \frac{p - c}{p} \varepsilon(p) = 0 \iff \varepsilon(p) \equiv -\frac{p \cdot q_p(p)}{q(p)} > 1 \,, \tag{1}$$

where ε denotes the elasticity of demand. Similarly, the sufficient condition for price p to maximize monopoly profits is:

$$\Pi_{pp}(p) = 2q_p(p) + (p - c) \cdot q_{pp}(p) < 0 \iff \rho(p) \equiv \frac{q(p) \cdot q_{pp}(p)}{\left[q_p(p)\right]^2} < 2 \,, \tag{2}$$

with ρ denoting the curvature of inverse demand. Demand can be concave ($\rho < 0$), linear ($\rho = 0$), but not excessively convex ($\rho > 0$), to ensure concavity of the profit function.

Mrázová and Neary (2017) prove that a well-defined smooth equilibrium relationship connecting elasticity ε and curvature ρ exists for continuous demands that are decreasing ($q_p(p) < 0$ and $p_q(q) < 0$) and three times differentiable. This allows us to invert the elasticity in Equation (1), and substituting into Equation (2), we obtain the demand manifold:

$$\rho[\varepsilon(p)] = \frac{q(p)\cdot q_{pp}(p)}{[q_p(p)]^2} = \frac{p^2\cdot q_{pp}(p)}{\left(-\frac{p\cdot q_p(p)}{q(p)}\right)^2\cdot q(p)} = \frac{p^2\cdot q_{pp}(p)}{\varepsilon^2(p)\cdot q(p)}\,. \tag{3}$$

Cournot (1838) first established the connection between demand curvature and pass-through for a monopolist with constant marginal costs:

$$\frac{dp}{dc} = \frac{1}{2-\rho} > 0\,, \tag{4}$$

When the monopolist faces log-concave demand with $\rho < 1$, its pass-through of cost shocks is incomplete, while it is more than complete in the case of log-convex demand with $\rho > 1$. Complete pass-through occurs when $\rho = 1$. Our representation of the manifold in terms of (ε, ρ) therefore directly relates to economic outcomes of interest, namely markups and pass-through, respectively.

2.2 Oligopoly

The monopoly case is helpful to establish the connection between demand curvature and pass-through rate. We will use graphical analysis repeatedly to convey the intuition of how distributional assumptions in discrete choice models affect the relationship between own-elasticity and curvature by plotting demand manifolds corresponding to a single-product monopoly case. In practical applications, firms compete in oligopoly markets selling multiple products, and this graphical representation only approximates the relationship between elasticity and curvature of the residual demand of a particular product, given all other substitution estimates.

In oligopoly markets, the pass-through of a common cost shock depends not only on demand curvature but also on substitution between affected products. To convey the intuition of how substitution affects the pass-through rate for a given product, it is necessary to simplify these substitution patterns. Weyl and Fabinger (2013, §IV.1) focus on the homogeneous product oligopoly version of equation (4):

$$\frac{dp}{dc} = \frac{1}{1+\theta(1-\rho)} > 0\,, \tag{5}$$

where θ is a conduct parameter ranging from $\theta = 0$ for a perfectly competitive industry to $\theta = 1$ for monopoly. The case of $\theta = 1/n$ corresponds to the Cournot solution. As the number of firms increases, the role of demand curvature diminishes, and pass-through gets closer to complete. Genakos and Pagliero (2022) study homogeneous good gasoline markets and find that pass-through increases from 0.4 in monopoly markets to one in markets with four or more competitors. The quantitative importance of competition on pass-through in empirical settings with differentiated goods and more sophisticated substitution patterns is unclear, however.

Weyl and Fabinger (2013, §IV.2) also consider a particular case of symmetric price Bertrand-Nash equilibrium. For this particular case, the conduct parameter has an intuitive interpretation connected to substitution patterns:

$$\theta = 1 + \sum_{j \neq i} \frac{\partial q_j(p)}{\partial p_i} \bigg/ \frac{\partial q_i(p)}{\partial p_i}\,, \tag{6}$$

where the second term on the right-hand side corresponds to the aggregate diversion ratio of Shapiro (1996). If a product i has nearly no close substitutes, the firm can charge higher markups and because $\sum_{j \neq i} \partial q_j(p)/\partial p_i \to 0$ and $\theta \to 1$, its pass-through (5) coincides with the pass-through of a single-product monopolist (4). If, on the other hand, product i has several close substitutes, $\sum_{j \neq i} \partial q_j(p)/\partial p_i > 0$ and $\theta < 1$. The firm therefore faces a more competitive environment, limiting not only its ability to increase price over marginal cost but also its ability to pass any cost increase to consumers.

The monopoly pass-through of equation (4), which ignores substitution effects, could be understood as a rough upper-bound estimate of the pass-through of any oligopoly firm. When the oligopoly equilibrium is not symmetric and different products are sold at different prices, it becomes difficult to make analytical statements about the sensitivity of pass-through to general substitution patterns. We evaluate the quantitative relationship between (4) and (5) in the context of our Monte Carlo study in Appendix E.

2.3 Demand sub-convexity

Demand is said to be sub-convex (super-convex) if $\log[q(p)]$ is concave (convex) in $\log(p)$. In the monopoly manifold examples we consider in Sections 3 and 4, we focus our attention on sub-convex demand or instances when the demand elasticity increases in price; i.e.,

$$\varepsilon_p(p) = \frac{\varepsilon^2(p)}{p} \cdot \left[1 + \frac{1}{\varepsilon(p)} - \rho(p)\right] > 0 \quad \Longleftrightarrow \quad \rho(p) < 1 + \frac{1}{\varepsilon(p)} = \rho(p)^{CES}. \tag{7}$$

Equation (7) establishes a cutoff condition for the curvature of sub-convex demand. For a given price elasticity, this cutoff is the curvature of the Constant Elasticity of Substitution (*CES*) inverse demand, $p(x) = \beta x^{-1/\sigma}$. *CES* demand is the only demand system where a single parameter determines both elasticity and curvature: $\varepsilon^{CES} = \sigma$ and $\rho^{CES} = 1+1/\sigma > 1$. Thus, $\varepsilon_p(p) = 0$, which implies the well-known result that *CES* markups and pass-through are invariant to price.

There is empirical evidence supporting the so-called *Marshall's (1920) Second Law of Demand* of demand becoming more elastic as prices rise.[3] This demand property is key to the equilibrium existence results of oligopoly models with differentiated products, both for single-product firms (Caplin and Nalebuff, 1991a) and for multi-product aggregative games (Nocke and Schutz, 2018). Our analysis below also shows that sub-convexity helps generate well-behaved comparative statics and equilibria: as price rises, the firm no longer has the incentive to continue raising the price and garner increasing markups.

3 Demand Elasticity and Curvature for Discrete Choice Models

We now utilize the demand manifold to explore the relationship between curvature and elasticity in the context of mixed-logit demand, which forms the backbone of much empirical work in IO. We begin by characterizing the demand manifold for arbitrary specifications of preference heterogeneity. Define the indirect utility of consumer i from purchasing product j as:

$$u_{ij} = h_i(d_i, x_j) + f_i(y_i, p_j) + \xi_j + \epsilon_{ij}\,, \qquad i \in \mathcal{I},\ j \in \mathcal{J},\ \epsilon_{ij} \sim \text{EV1}\,, \tag{8}$$

where (x_j, ξ_j) denote the observed and unobserved characteristics of the product j, respectively, p_j its price and y_i the income of the consumer i. Mixed logit allows for heterogeneity in consumers' valuation of the product characteristics x, which we represent via the characteristic sub-function $h_i(d_i, x_j)$. This sub-function takes demographics d_i as an argument reflecting a possible correlation between consumer demographics and taste

[3] This includes evidence on the relationship between markups and the scale of production in macroeconomics (see Mrázová, Neary and Parenti, 2021, and references therein), markup adjustments after trade liberalization (De Loecker, Goldberg, Khandelwal and Pavcnik, 2016), pass-through of exchange rates for coffee and beer in trade (Nakamura and Zerom, 2010; Hellerstein and Goldberg, 2013), as well as tax pass-through in the legal marijuana market (Hollenbeck and Uetake, 2021) and markup adjustments to changes in commodity taxation(Miravete, Seim and Thurk, 2018).

heterogeneity over product characteristics. Lastly, we normalize the value of the outside option to zero.

The sub-function $f_i(y_i, p_j)$ represents how spending on the outside good, $y_i - p_j$, affects indirect utility. The effect of outside good spending varies by individual i, as income varies across consumers and consumers differ in their price sensitivity. To simplify notation, we write:

$$f'_{ij} = \frac{\partial f_i(y_i, p_j)}{\partial p_j}\,, \qquad \text{and} \qquad f''_{ij} = \frac{\partial^2 f_i(y_i, p_j)}{\partial p_j^2}\,. \tag{9}$$

The term f'_{ij} represents the marginal effect of price p_j on consumer i's indirect utility, while f''_{ij} represents how this marginal effect changes with price.

Individual i purchases product j if $u_{ij} \geq u_{ik}\,,\ \forall k \in \{0, 1, \ldots, J\}$. Because of the additive i.i.d. type-I extreme value distribution of ϵ_{ij}, individual i's choice probability of product j is:

$$\mathbb{P}_{ij}(p) = \frac{\exp\big(h_i(d_i, x_j) + f_i(y_i, p_j) + \xi_j\big)}{1 + \sum_{k=1}^{J} \exp\big(h_i(d_i, x_k) + f_i(y_i, p_k) + \xi_k\big)}\,. \tag{10}$$

Notice that individual i makes a dichotomous decision about purchasing product j (i.e., "Buy j" vs. "Buy Something Else"). We therefore consider the individual purchase decision as the outcome of a Bernoulli random process with a probability of success $\mathbb{P}_{ij}$. This probability varies with the vector of prices and characteristics of the different alternatives. The Bernoulli random variable has mean $\mu_{ij} = \mathbb{P}_{ij}$, variance $\sigma^2_{ij} = \mathbb{P}_{ij}(1 - \mathbb{P}_{ij})$, and (non-standardized) skewness $sk_{ij} = \sigma^2_{ij}(1 - 2\mathbb{P}_{ij})$. Aggregating over the measure of heterogeneous individuals summarized by $G(i)$, total demand for product j becomes:

$$Q_j(p) = \int_{i \in \mathcal{I}} \mathbb{P}_{ij}(p)\, dG(i)\,. \tag{11}$$

We can now write the elements defining the demand manifold, elasticity, and curvature of product j, relegating the detailed derivations to Appendix B. The own-price demand elasticity of product j amounts to a scale-free measure that aggregates individual price responses (demand slopes) weighted by their choice variance:

$$\varepsilon_j(p) = -\frac{p_j}{Q_j(p)} \int_{i \in \mathcal{I}} f'_{ij} \cdot \sigma^2_{ij}\, dG(i)\,, \tag{12}$$

Similarly, the demand curvature of our discrete choice model is:

$$\rho_j(p) = \int_{i\in\mathcal{I}} \mu_{ij}\, dG(i) \times \frac{\int f''_{ij}\cdot \sigma^2_{ij}\, dG(i) \;+ \int \left(f'_{ij}\right)^2 \cdot sk_{ij}\, dG(i)}{\left[\int f'_{ij}\cdot \sigma^2_{ij}\, dG(i)\right]^2}\,. \tag{13}$$

Combining elasticity (12) and curvature (13), we obtain the general expression for the mixed logit demand manifold:

$$\rho_j[\varepsilon_j(p)] = \frac{p_j^2}{\varepsilon_j^2(p)\cdot Q_j(p)} \cdot \underbrace{\left[\int f''_{ij}\cdot \sigma^2_{ij}\, dG(i) \;+ \int \left(f'_{ij}\right)^2 \cdot sk_{ij}\, dG(i)\right]}_{\text{Mixing Distributions}}. \tag{14}$$

Discussion. The demand manifold in equation (14) suggests that the mixing distributions (in between brackets) of attribute valuations across the underlying distribution of customers drive the relationship between elasticity and curvature. While these are best understood as primitives of customer demand, the distribution of taste heterogeneity, G, and the sub-functions $f(\cdot)$ and $h(\cdot)$ are typically chosen before and held fixed during estimation. For example, a standard specification of non-price tastes through $h(\cdot)$ is a linear function of customer demographics, and non-observed heterogeneity is captured via a standard normal distribution. These choices implicitly restrict $\{\sigma^2_{ij},\ sk_{ij}\}$, therefore limiting the relationship between elasticity and curvature.

Researchers also often choose quasi-linear preferences, where the pricing sub-function is a linear function of outside good spending, i.e., $f_{ij}(y_i, p_j) = \alpha^\star_i(y_i - p_j)$. Then, heterogeneity in the idiosyncratic price sensitivity $\alpha^\star_i$ for any given elasticity drives demand curvature, generating different implications from the below case of a non-linear sub-function $f_{ij}(y_i, p_j)$, e.g., the polynomial approximations of Griffith et al. (2018). Our setup, however, always consists of discrete choice unit demands that are consistent with utility maximization, i.e., where Roy's identity holds for $q_{ij} = 1$.

4 Demand Manifolds of Common Discrete Choice Models

We provide a graphical illustration in this section of how choices of the distribution for taste heterogeneity G and the pricing sub-function $f(\cdot)$ impact the shape of the demand manifold. We focus on a single-product monopoly as a simplifying example to provide intuition.

4.1 Quasi-linear Preferences

For quasi-linear preferences, the pricing sub-function simplifies to $f_i(y_i, p_j) = \alpha_i^\star(y_i - p_j)$ where $\alpha_i = \alpha + \sigma_p \phi_i$ and $h_i(d_i, x_j) = \beta_i x_j$ where $\beta_i = \beta + \sigma_x \nu_i$. The distribution of price sensitivity has a mean of α with deviations driven by the shape of Φ, the mean-zero distribution of ϕ_i, scaled by σ_p. Similarly, β is the mean valuation of the observed product characteristics, and ν_i captures the idiosyncratic heterogeneity in the valuation, which we assume takes the form of a standard normal random variable scaled by σ_x.

Note that purchase decisions based on indirect utility comparisons do not depend on individual income y_i, which shifts the indirect utility of all products by $\alpha_i^\star y_i$, so there are no income effects. Furthermore, with $f_i(y_i, p_j)$ linear in price, $f'_{ij} = -\alpha_i^\star$ and $f''_{ij} = 0$ so the demand manifold simplifies to:

$$\rho_j[\varepsilon_j(p)] = \frac{p_j^2}{\varepsilon_j^2(p) \cdot Q_j(p)} \cdot \int (\alpha_i^\star)^2 \cdot sk_{ij} \, dG(i) \, . \tag{15}$$

We employ Equation (15) to explore the demand manifolds of several workhorse discrete choice specifications from the empirical literature: *MNL*, *CES*, *ML* with random coefficients on product attributes, and *ML* with a random coefficient on price. The extent and manner in which these specifications introduce flexibility in the preference specification vary, enabling us to demonstrate how the demand model's capacity to accommodate feasible combinations of elasticity and curvature changes as we relax these restrictions.

Multinomial Logit Preferences. In the *MNL* model, there is no unobserved heterogeneity, so $\sigma_p = \sigma_x = 0$; $\alpha_i^\star = \alpha$; and $\beta_i^\star = \beta$. Hence, $\mathbb{P}_{ij} = \mathbb{P}_j = s_j(p)$ is the market share of product j. Elasticity and curvature reduce to:

$$\varepsilon_j(p) = \alpha p_j \big(1 - \mathbb{P}_j\big) \, , \tag{16a}$$

$$\rho_j(p) = \frac{1 - 2\mathbb{P}_j}{1 - \mathbb{P}_j} < 1\,. \tag{16b}$$

Equation (16b) shows that *MNL* demand is concave with negative curvature only in very concentrated markets where the product's market share exceeds 50% of sales.[4] *MNL* restricts demand to be log-concave and $\rho_j(p) < 1$ for all possible prices. Pass-through in *any MNL* demand model is necessarily incomplete regardless of setting and identification strategy. Furthermore, since *MNL* demand curvature (16b) decreases in $\mathbb{P}_j$, pass-through grows arbitrarily close to complete when the product market is atomistic and $\mathbb{P}_j \to 0$. This means the researcher's choice of potential market pins down demand curvature and pass-through before estimation.

The left panel of Figure 2 depicts several demand manifolds for a single-product monopoly *MNL* model. We fix the product attribute to take a value of $X = 1$ and allow consumer valuations for the attribute β to range from $\{\beta, \beta + 1 \ldots, \beta + 5\}$, with $\beta = 1$. For each attribute valuation β, the manifold of the *MNL* is increasing. We set the price response coefficient $\alpha = 0.5$ and consider elasticity-curvature combinations at various price levels. Each manifold is color-coded for prices, ranging from $p_j = 0$ (darkest) to $p_j = 10$ (lightest). Note that higher prices always result in more elastic demands and lower equilibrium markups. Increasing the average valuation of the product attribute, β, to β+1, β+2,. . . , shifts demand manifolds upwards from the base *MNL* manifold in Figure 2. Increasing a product's mean demand, therefore, results in decreases of both curvature and price elasticity for a given price.

Nested Logit. Nested logit – a demand system commonly employed in antitrust economics – provides for more reasonable substitution patterns with a limited computational burden.[5] As with *MNL*, nested logit demand is log-concave. An economic implication of log-concave demand is that demand curvature and pass-through are always less than one (Amir, Maret and Troege, 2004). This means that while nested logit provides a simple framework for approximating substitution patterns, this simplicity comes at the expense of restricting counterfactual pricing equilibria in merger analysis.

[4] Jaffe and Weyl (2010) show that random utility models are inconsistent with linear demand beyond the symmetric duopoly case, when $\mathbb{P}_j = s_j(p) = 1/2$, which follows immediately from (16b).

[5] McFadden and Train (2000) demonstrate that a *ML* specification with random coefficients on product characteristics can generate equivalent substitution patterns to the nested logit model.

Figure 2: Multinomial and Mixed Logit Manifolds

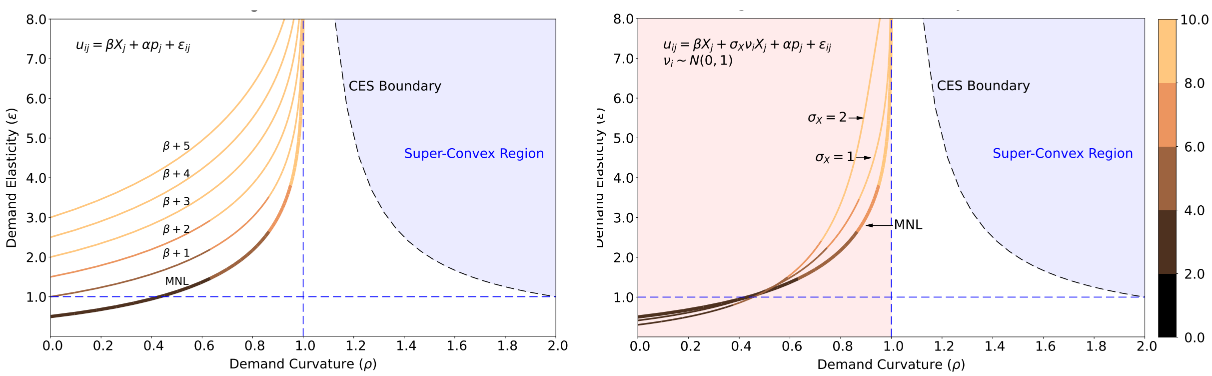

Figure Notes: The left panel shows six alternative *MNL* demand manifolds with one inside good assuming $\alpha = 0.5$, $X = 1$, and $\beta \in \{1, ..., 6\}$. The right panel shows manifolds for a *ML* model with a random coefficient on the product characteristic under alternative standard deviations σ_x and $\beta = 1$.

Constant Elasticity of Substitution Preferences. The decreasing and convex black dashed curve in Figure 2 represents the (ε, ρ) combinations for *CES* demand under alternative values for the elasticity of substitution. Anderson et al. (1992) show that a discrete choice model where individuals spend a fraction of their income on a continuous quantity of a single product can generate the *CES* utility function of the representative consumer model. *CES* therefore arises naturally in the context of discrete-continuous models (Hanemann, 1984), while *MNL* is most appropriate when consumers have unit demands. However, like the *MNL* model, *CES* choice probabilities suffer from the *IIA* property in producing unrealistic substitution patterns, e.g., Head and Mayer (2025). Figure 2 also illustrates that for the same elasticity, the *CES* and *MNL* models imply different demand curvatures (and pass-through). This means model choice starkly restricts pass-through, often inconsistently with the data.

ML with Characteristic Random Coefficients. The literature has highlighted that accounting for idiosyncratic preferences for product attributes can relax the restricted substitution patterns generated by *MNL* demand. We consider whether heterogeneity in the valuation of the product attribute addresses the limitations of *MNL* in restricting curvature, while still assuming that all consumers have the same price responsiveness.

The right panel of Figure 2 shows several demand manifolds for such a *ML* model, allowing the standard deviation of the random coefficient on the product attribute to increase from $\sigma_x = 1$ to $\sigma_x = 2$, while holding fixed the mean product valuation at $\beta = 1$. Adding individual preference heterogeneity "rotates" manifolds: for a given demand elasticity, preference heterogeneity reduces demand curvature and, hence, pass-through. The firm now faces a segment of consumers with high valuations for its attribute over whom it has market power locally, and it reduces its pass-through relative to the case of uniform preferences.

The light-red shaded area denotes the combinations of elasticity and curvature that a *ML* model with heterogeneity in the valuation of the product characteristic can generate for mean valuations of $\beta \geq 1$. The figure illustrates that the *ML* model with normally distributed attribute preferences continues to generate log-concave demand. Caplin and Nalebuff (1991b) show that *ML* demand remains log-concave under any log-concave distribution of idiosyncratic preferences, comprising the vast majority of distributions used in economics (Bagnoli and Bergstrom, 2005). Mathematically, Equation 14 demonstrates that curvature can only come through the shape of the choice probability distribution ($\mathbb{P}_{ij}$), particularly the skewness.

It is evident that this version of a *ML* model has inherent limitations when empirically studying pass-through in non-competitive environments: pass-through is necessarily restricted to be incomplete.[6] In empirical settings with log-convex demand, firms with market power aim to over-shift cost shocks. Employing a *MNL* or a *ML* model with idiosyncratic preferences over attributes in such instances would result in biased preference estimates that generate the closest demand curvature to the true data-generating process that these models can produce, a curvature of effectively one. Figure 2 illustrates that to exhibit such demand curvature, the estimated model would either understate the true degree of idiosyncratic product attribute preferences or overstate consumers' true price sensitivity, generating the appearance of a competitive environment with full pass-through.

ML with Price Random Coefficients. How can we expand the range of curvatures that the *ML* estimates can accommodate to allow for log-convex demand and over-shifting of pass-through? The only remaining element of preferences to consider is idiosyncratic price responsiveness. Substituting $\alpha_i^\star = \alpha + \sigma_p \phi_i$ into the demand manifold for quasi-linear preferences (15) results in:

$$\rho_j[\varepsilon_j(p)] = \frac{p_j^2}{\varepsilon_j^2(p) \cdot Q_j(p)} \cdot \int (\alpha + \sigma_p \phi)^2 \cdot sk_{ij} \, d\Phi(i) \tag{17}$$

[6] This is at odds with evidence of pass-through rates exceeding 100% in horizontally differentiated products industries such as groceries (Besley and Rosen, 1999); clothing and personal care items (Poterba, 1996); branded retail products (Besanko, Dubé and Gupta, 2005); gasoline and diesel fuel (Marion and Muehlegger, 2011); as well as beer, wine, and spirits (Kenkel, 2005) among others.

In the absence of idiosyncratic price heterogeneity, $\sigma_p = 0$, this demand manifold coincides with the manifold of the *MNL* in Equations (16a) and (16b). For any given demand elasticity and price-quantity pair, an increase in the spread of the distribution of idiosyncratic price heterogeneity via σ_p *expands* the range of demand curvatures that the model can generate. Indeed, the shift of each manifold to the right is proportional to the second-order moment of the distribution Φ. With a sufficiently large σ_p relative to the mean price coefficient α, the manifolds cross the unit curvature threshold, allowing discrete choice demand to accommodate pass-through rates above 100%. We illustrate this argument next by considering particular price mixing distributions.

Normal and Log-normal Price Random Coefficients. We now consider the choice of price mixing distribution, focusing on the range of feasible elasticity and curvature combinations up to the *CES* boundary that a candidate price mixing distribution could generate. We begin with two price-mixing distributions commonly used in empirical work: the normal and log-normal distributions. Figure 3 depicts the demand manifolds when price random coefficients are normally and log-normally distributed for alternative values of σ_p. The light-red shaded area identifies all combinations of (ε, ρ) within the sub-convex region of demand that are feasible under each model for any combination of the structural parameters $(\alpha, \sigma_p, \beta)$. Both panels show that increasing the variation in idiosyncratic price responsiveness σ_p increases the feasible curvatures the *ML* model can accommodate for a given elasticity value. Manifolds now cross into the log-convex region of demand with more than complete pass-through, a result that is consistent with many of the $(\hat{\varepsilon}, \hat{\rho})$ estimates in Figure 1.

In the left panel, we depict, among others, the demand manifold corresponding to the particular demand specification with a normally distributed price coefficient with $\sigma_p = 0.15$. The figure shows that for this value of σ_p, the maximum elasticity is reached precisely at the price level where the demand manifold intersects the *CES* locus. For higher price levels, elasticity decreases in price, violating Marshall's Second Law. For demand to be sub-convex for all price levels, we therefore require less heterogeneity in price sensitivity among consumers; i.e., low σ_p.

Utilizing a one-tailed log-normal distribution for price sensitivity introduces skewness (Equation 17). Beyond ensuring that the demand of all simulated consumers is downward sloping (Train, 2009), it expands the scope for more prominent differences in price sensitivity and curvature; the right panel in Figure 3 shows that larger values of σ_p continue to generate

Figure 3: Normal and Log-Normal Price Mixing Distributions

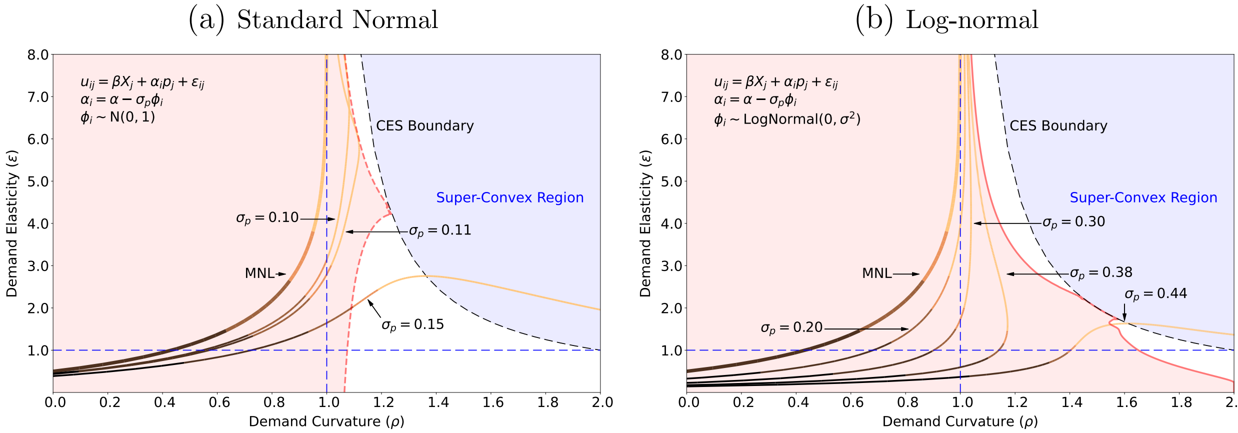

Figure Notes: Panels present demand manifolds in the (ε, ρ) plane under standard normal and log-normal price mixing distribution, respectively. Light-shaded regions represent all feasible (ε, ρ) pairs conditional on the price-mixing distribution. We define "feasible" as manifolds for demand that is sub-convex for all prices. We identify this region through a grid search across all parameters, conditional on the utility specification, including the price mixing distribution.

sub-convex demand. This results in a significantly larger range of feasible curvatures for a given demand elasticity, particularly for less elastic demands, where firms enjoy greater market power. Figure 3 hence shows that a model with a log-normal price random coefficient can admit the majority of curvature-elasticity pairs in the sub-convex region of demand.

A More Flexible Mixing Distribution. While log-normality increases the set of achievable elasticity-curvature pairs, there are still small gaps in coverage. Here, we consider a more flexible mixing distribution – the three-parameter Asymmetric Generalized Normal distribution (Nadarajah, 2005) which collapses to a one-parameter distribution in the case of mixed logit demand summarized by the shape parameter κ.[7] Figure 4 explores the implications of using this flexible mixing distribution for the price random coefficient.

In panel A we present three different variants of how the price mixing distribution may look using various values of κ: ranging from standard normal to log-normal. We also consider an intermediate case that might represent a particular mixture of these two distributions. Panel (a) therefore provides intuition of how the Asymmetric Generalized Normal modulates the shape of the price mixing distribution to cover (ε, ρ) space in panels A and B in Figures 3 as well as the space between. We confirm this intuition in panel B where we see the Asymmetric Generalized Normal does indeed cover (ε, ρ) space conditional on

[7] See Online Appendix C for technical details.

Figure 4: Covering the Space with a Flexible Price Mixing Distribution

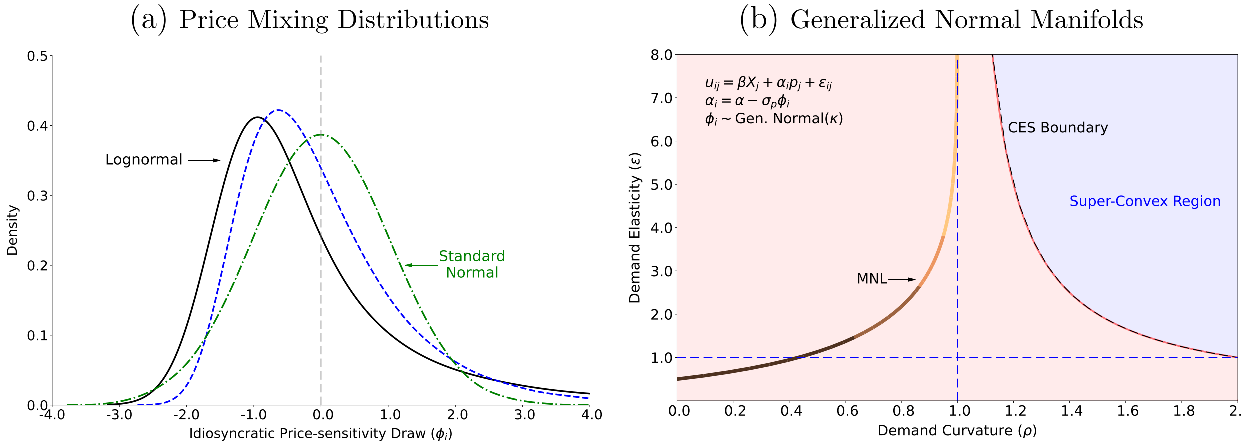

Figure Notes: The left panel shows three specifications of the price random coefficient distribution for different values of shape parameter κ with the Asymmetric Generalized Normal distribution. The right panel shows the combinations of all structural parameters generating well-behaved solutions for (ε, ρ) in the sub-convex region. The model is capable of covering the sub-convex region in Panel (b) because of the flexibility provided by the Asymmetric Generalized Normal as the pricing mixing distribution.

maintaining demand sub-convexity (light-shaded region). This result indicates that flexibility in the price mixing distribution can be achieved parsimoniously.

Demographics as Mixing Distributions. Empirical applications frequently exploit the fact that idiosyncratic price responsiveness is correlated with demographics. Rather than imposing a distribution on idiosyncratic price sensitivities, as we did above, one might therefore specify the idiosyncratic price sensitivity α_i as a function of an observable demographic d_i, i.e., $\alpha_i^\star = \alpha + \pi_d d_i$. The equivalence to the analysis of Section 3 is apparent: it is now the empirical distribution of demographic d_i that underlies measure $G(i)$ in the manifold expression (3) and that determines the feasible combinations of (ε, ρ) pairs that the demand system can accommodate. Since researchers choose demographics to use as mixing distributions, this result highlights that these choices restrict, ex ante, the set of achievable (ε, ρ) pairs. Equivalently, these choices restrict the estimates of market power and achievable pricing equilibria.

4.2 Preferences with Income Effects

We now extend the analysis to consumer preferences with income effects. In contrast to the quasi-linear case, where outside good spending enters consumers' indirect utility function linearly, it is natural to allow for income effects when studying expensive products, such as automobiles. For example, *BLP* specify Cobb-Douglas utility which generates indirect utility in Equation (8) via the following price sub-function:

$$f_i(y_i, p_j) = \alpha \ln \left(y_i - p_j\right) . \tag{18}$$

Both the quasi-linear price sub-function and *BLP*'s alternative are, however, special cases of a Box-Cox power transformation (Box and Cox, 1964) of outside good spending. We, therefore, find it useful to specify the following generalized price sub-function to vary the importance of income effects:

$$f_i(y_i, p_j) = \alpha_i^\star \left(y_i - p_j\right)^{(\lambda)} = \begin{cases} \alpha_i^\star \dfrac{\left(y_i - p_j\right)^\lambda - 1}{\lambda} , & \text{if } \lambda \neq 0 , \\ \alpha_i^\star \ln \left(y_i - p_j\right) , & \text{if } \lambda = 0 , \end{cases} \tag{19}$$

The power parameter $\lambda \in \mathbb{R}$ drives the convexity or concavity of the transformation, thereby varying the importance of income effects. This enables us to explore how the value of the power parameter λ affects demand elasticity (12), curvature (13), and the shape and position of the manifold (14) through its effect on f'_{ij} and f''_{ij} in Equation (9). This flexibility is a valuable feature that we utilize in our Monte Carlo simulations and empirical illustrations, where we estimate, rather than impose ex-ante, the importance of income effects. A subtle point in (19) is that the researcher still has the freedom to choose the consumer's disposable income. This could be the consumer's wage income, their wealth, or some other measure of discretionary income. What we have shown is that this choice will drive the shape of demand as well as the set of attainable elasticity-curvature pairs. The Box-Cox parameter modulates this choice to better fit the distribution of price sensitivities in the data.

In line with the *BLP* specification, we abstract from heterogeneity in price sensitivity and consider the case of $\alpha_i^\star = \alpha$. A power parameter of $\lambda = 0$ yields the *BLP* model, while a power parameter of $\lambda = 1$ results in a quasi-linear model such as Nevo (2001). This means that the income distribution captures any idiosyncratic price responsiveness across individuals, modulated by λ.[8]

[8] Using a multi-unit demand model, Birchall et al. (2025) rely on a price sub-function with a different Box-Cox transformation, $f(y_i, p_j) = \gamma^{\lambda-1} \left(y_i^\lambda - 1\right) \lambda - \left(p_j^\lambda - 1\right) \lambda$, which depends on the share of income, γ, spent on a chosen product and is only well behaved for $\lambda \in (0, 1)$ (Anderson and de Palma, 2020). Their transformation is an *h-function* bridging *MNL* and *CES* demands (Nocke and Schutz, 2018). The resulting

Following Berry et al. (1999), we adopt a first-order Maclaurin series approximation (at $p_j = 0$) of the Box-Cox transformation:[9]

$$f_i(y_i, p_j) = \alpha\big(y_i - p_j\big)^{(\lambda)} \simeq \alpha y_i^{(\lambda)} - \frac{\alpha p_j}{y_i^{1-\lambda}}\,. \tag{20}$$

yielding a demand manifold of:[10]

$$\rho_j[\varepsilon_j(p)] = \frac{p_j^2}{\varepsilon_j^2(p)\cdot Q_j(p)} \cdot \int\limits_{i\in\mathcal{I}} \frac{\alpha^2\cdot\big[(1-\lambda)y_i^{-\lambda}\cdot\sigma_{ij}^2 \;+ sk_{ij}\big]}{y_i^{2(1-\lambda)}}\, dG(i)\,. \tag{21}$$

Figure 5 depicts demand manifolds under various λ values in $[0, 1]$ when the characteristic sub-function is βX_j and income y_i is a log-normal approximation to the U.S. income distribution. The figure illustrates that accommodating income effects via the approximate Box-Cox transformation of outside good spending yields preferences that can accommodate curvatures close to those of the *CES* boundary as λ approaches zero; i.e., when the pricing sub-function is $f_{ij} = \log(y_i - p_j)$.[11] This suggests that the motivating patterns in the pairs of elasticity and curvature in Figure 1 for Berry et al. (1999) are driven by the shape of the income distribution.

Discussion. The preceding sections demonstrate that the *ML* model exhibits significant flexibility in capturing realistic substitution patterns and generating a wide range of cost pass-through when we allow for (a) heterogeneity in consumer valuations for product at-

---

curvature flexibility disappears for $\lambda = 1$, when the specification reduces to the quasi-linear unit-demand case. Our goal in specifying sub-function (19) is to expand the curvature range of a unit demand model through a Box-Cox power parameter that modulates income effects within the confines of a unit-demand setup consistent with utility maximization.

[9] Note that for $\lambda = 0$, the price sub-function becomes $\alpha \ln y_i - \alpha p_j / y_i$, which only coincides with equation (19) for $y_i = 1$. Hence, the preference specification based on equation (20) is only approximately consistent with utility maximization.

[10] This is the particular solution of the demand manifold derived in Online Appendix B for the case of the Maclaurin approximation of the Box-Cox price sub-function (19).

[11] We consider a power parameter $\lambda \in [0, 1]$, in line with the empirical literature, to ensure that marginal utility of income is increasing and concave (so that higher income households are always less responsive to price at different rates as modulated by λ). Note that using the convenient Maclaurin approximation leads to some estimation bias and *BLP* estimates sit on the *CES* boundary, which should correspond to $\lambda = 0$. Using the approximation will produce estimates of λ smaller than zero, corresponding to a more concave function than $\ln(y_i - p_j)$. Values $\lambda < 0$ still ensure that the marginal utility of income is increasing and concave, so that higher income individuals are less responsive to prices.

Figure 5: Demand Manifolds & Income Effects

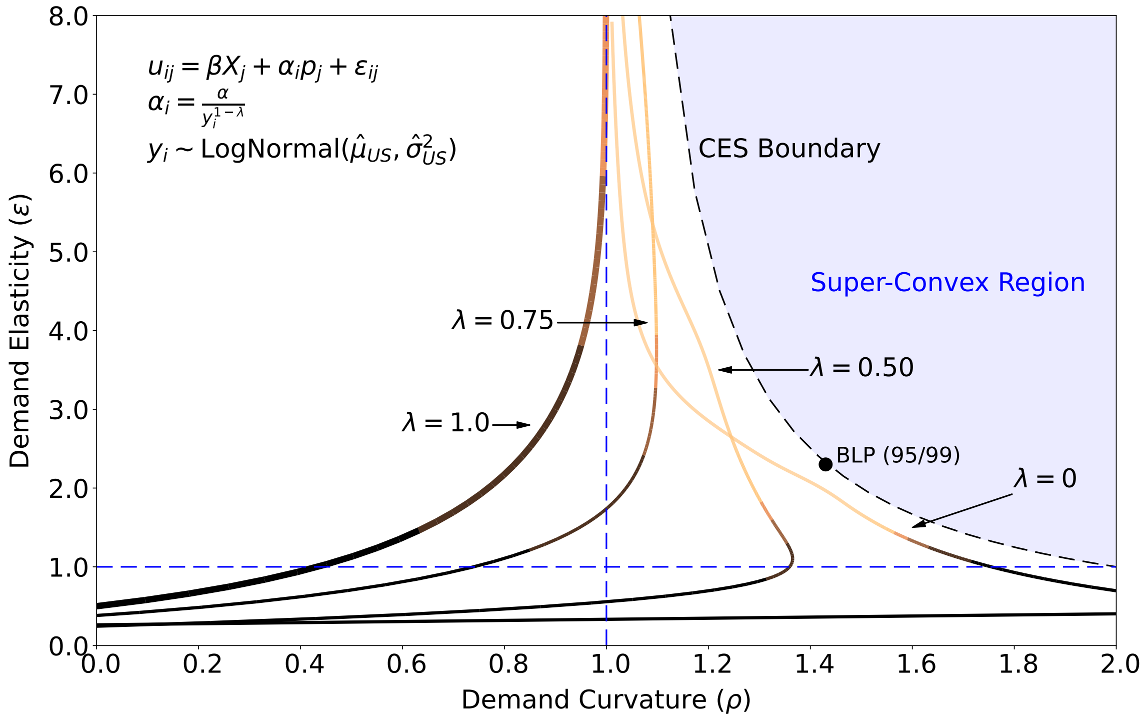

Figure Notes: Demand manifolds for different values of the Box-Cox transform parameter λ using the U.S. income distribution and the rest of the model specification of Berry et al. (1999). The dot identified as "BLP (95/99)" corresponds to the average estimated curvature and elasticity value using the *BLP* automobile data and estimation best practices as outlined in Conlon and Gortmaker (2020).

tributes and (b) flexible distributions of price sensitivity. The former expands the set of estimable elasticity-curvature combinations within the log-concave region of demand, while the latter extends this set to log-convex demand. Moreover, the interaction between (a) and (b) is non-trivial and is likely unknown before estimation. The intuition is that idiosyncratic attribute valuations give firms localized market power, leading to under-shifted pass-through. At the same time, consumer heterogeneity in price sensitivity results in over-shifted pass-through, as the firm focuses on different customer types in response to cost changes. The combined effect of these two forces drives a given product's pass-through. We provide a simple linear example of over-shifted pass-through (i.e., log-convex demand) via heterogeneous price sensitivities in Appendix D.

5 Guidance for Empirical Work

We suggest modeling preference heterogeneity using the Box-Cox power transform introduced in Section 4.2. While there are many ways to model consumer heterogeneity flexibly, the Box-Cox transform provides a simple, one-parameter mechanism for transforming data into mixing distributions that are consistent with identifying data moments. It also has the added advantage of being consistent with unit-demand models that rationalize consumer utility maximization.

In the following sections, we address the identification and estimation of the shape of demand in horizontally differentiated, multi-product oligopoly typical to empirical settings. We begin with a model where demand exhibits income effects, as in Berry et al. (1999), and then continue by evaluating the quasi-linear demand case presented in Nevo (2001). Both models generate results that are consistent with the predictions from the single-product monopoly case from Section 2 to connect mixing distributions to elasticity, curvature, and the shape of demand. These results also demonstrate that our approach provides an easy alternative for implementing empirical strategies to extend the range of estimable pass-through ratios that discrete choice models can accommodate in a wide variety of applied settings. While our focus in this section is on using the Box-Cox power transform to estimate the distribution of price sensitivities flexibly, modifying our identification strategy to non-price characteristics to accommodate lower demand curvature and pass-through is straightforward.

We detail our identification strategy in Section 5.1. In Section 5.2, we conduct a Monte Carlo analysis to demonstrate the effectiveness of our identification strategy and to illustrate the consequences of misspecification. In Section 5.3 we simulate a counterfactual trade policy using alternative estimated demand specifications for new automobiles. We conclude in Section 5.4 by addressing the empirical implications of imposing different price mixing distributions in a quasi-linear preference setup.

5.1 Instruments to Identify Demand Manifolds

We employ an identification strategy that exploits heterogeneous consumer responses to exogenous price changes by relying on a variant of the instruments proposed initially by *BLP* and refined as "Differentiation IVs" by Gandhi and Houde (2020): the distance of the focal product from rivals in product characteristic space. Changes in the focal product's isolation in characteristic space exogenously shift its demand, assuming that product characteristics are chosen before demand unobservables, ξ, are realized. A comparison between instances with many versus few similar products reveals the extent to which consumers substitute between similar products, akin to observing exogenous variation in choice sets. As the Box-Cox transformation allows for nonlinear heterogeneity in such substitution as a function of income or other observable consumer demographics, we interact the Differentiation IV with moments from the demographic distribution, e.g., income. This allows us to recover the shape of the distribution of consumers' price sensitivities and attribute valuations and, hence, the curvature of a unit demand function.

A challenge, of course, when employing this instrument to identify price sensitivity is the endogeneity of prices in an oligopoly equilibrium: unobserved demand shocks ξ may confound the response in price to a change in cost ω. We follow Gandhi and Houde (2020) and construct exogenous price predictions via a reduced-form hedonic price regression based on exogenous characteristics x_t and cost shocks ω_t:[12]

$$p_t = \gamma_0 + \gamma_1 x_t + \gamma_2 \omega_t + u_t \,. \tag{22}$$

We run the above regression and use the results to construct the vector of predicted prices $\hat{p}_t$. We then construct differences in price-space between product j and its competitors:

$$Z^p_{jt} = \sum_r \left(\hat{p}_{rt} - \hat{p}_{jt} \right)^2 . \tag{23}$$

Equation (22) enables us to construct exogenous prices by separating price effects due to changes in demand (via ξ) from changes in cost (via ω). It is also a simple pass-through regression. Cost pass-through $\hat{\gamma}_2$ informs the identification of demand primitives related to curvature via the demand shocks captured in equation (23). While Gandhi and Houde (2020) recommend relying on Z^p to identify the distribution of unobserved preference heterogeneity, interacting it with moments of the distribution of observable demographics serves to identify the case when price sensitivity is correlated with the same demographics. For example,

[12] Alternatively, we could construct prices non-linearly using firm first-order conditions as in Berry et al. (1999).

when estimating demand allowing for flexible income effects, we include the interactions of the above price differentiation instrument Z^p with moments of the income distribution:

$$Z^P_{jt} = \sum_r \left(\hat{p}_{rt} - \hat{p}_{jt}\right)^2 , \tag{24a}$$

$$Z^D_{jt} = Z^P_{jt} \otimes \left\{\text{inc}^{10\%}_t,\ \text{inc}^{50\%}_t,\ \text{inc}^{90\%}_t\right\} . \tag{24b}$$

We trace the demand manifolds using cost shocks, while holding exogenous demand shifters constant at different price levels. Section 5.2 explores the instrument's performance in Monte Carlo simulations. Lastly, we extend the argument to quasi-linear preferences in Section 5.4.

5.2 Flexible Discrete Choice Estimation with Income Effects

We conduct a Monte Carlo analysis to demonstrate the validity of our identification strategy and evaluate the potential for mis-specified demand systems to introduce biases in the economic outcomes of interest, namely elasticity and curvature. Consider a setting with $J = 20$ differentiated products sold by single-product firms competing in price for $T = 50$ periods. Consumer indirect utility takes the following form:

$$u_{jlt} = \underbrace{\beta_0 + \beta_1 x^1_{jt}}_{\substack{\text{Common Across} \\ \text{Consumers}}} + \underbrace{\sum_{k=1}^{K} \left(\beta_{2,k} + \sigma_{X,k}\nu_{ik}\right) x^2_{jt,k}}_{\substack{\text{Idiosyncratic} \\ \text{Characteristic Tastes}}} - \underbrace{\alpha \cdot p_{jt} \cdot y^{\lambda-1}_{it}}_{\substack{\text{Idiosyncratic} \\ \text{Price Sensitivities}}} + \xi_{jt} + \epsilon_{ijt} , \tag{25}$$

where income effects decrease as λ moves from zero to one. In this specification, some product characteristics are observed by the researcher (x^1_{jt}, x^2_{jt}) while others are only observed by consumers and firms (ξ_{jt}). Valuation of the product attribute x^1_{jt} is common across individuals, and we draw x^1 independently from a uniform distribution. We model consumer preference heterogeneity in product characteristics via x^2_{jt} with two elements $(K = 2)$, including a constant and a uniformly distributed product characteristic. As in Gandhi and Houde (2020), product attributes (other than the constant) vary across time.[13] Consumers, therefore, have preference heterogeneity over the J inside goods via the constant random coefficient and over variation in the observable product characteristic across the J products and T time periods.

[13] In empirical applications, such as automobiles, this is due to product remodels, which the researcher treats as exogenous to unobserved variation in demand via ξ. This is equivalent to allowing for exogenous product entry and exit – a common assumption in the empirical literature.

We set $\beta_2 = 1$ and $\sigma_X = 5$ for $k = 1, 2$. We assume that the unobservable characteristic ξ_{jt} is distributed standard normal. We model heterogeneous price sensitivity using the approximation to the Box-Cox transformation (20) of outside good spending modulated by parameter λ. We assume that consumer income y_{it} is drawn from a log-normal distribution and parameterize these draws following Andrews, Gentzkow and Shapiro (2017), generating market and time variation by allowing the variance of income to vary.

Single-product firms choose prices simultaneously each period, given their constant marginal costs c_{jt}. In the static oligopoly Bertrand-Nash equilibrium, period t equilibrium prices $p_t^\star$, satisfy the set of J first-order conditions for the firms:

$$p_{jt}^\star = c_{jt} - s_j(\delta_t, p_t^\star; \sigma_X, \sigma_p) \times \left[\frac{\partial s_j(\delta_t, p_t^\star; \sigma_X, \sigma_p)}{\partial p_{jt}^\star} \right]^{-1} . \tag{26}$$

Marginal costs are a function of product characteristics and cost shocks:

$$\log c_{jt} = \tau_0 + \tau_1 \log x_{jt}^1 + \tau_2 \log x_{jt}^2 + \omega_{jt} + \zeta_{jt} \tag{27}$$

We set all τ parameters equal to 1 and draw cost shocks $\{\omega_t, \zeta_t\}$ from standard normal distributions. We assume that ω_t, which identifies the distribution of price sensitivity, is observed. We generate pricing equilibria under Bertrand-Nash pricing as in Berry et al. (1999) in the true data-generating processes by selecting α and β_0 so that the average own-price elasticity is 2.5 with a 20% aggregate inside share for each simulation.

We consider alternative specifications of the role of outside good spending in demand and use the Monte Carlo analysis to investigate the success of an empirical demand model with a flexible Box-Cox power transformation of outside good spending at recovering the true demand curvature underlying these data-generating processes, relative to simpler alternatives. We employ the best practices outlined in Conlon and Gortmaker (2020) to estimate consumer demand given observed prices, quantities, and ω cost shocks.

We consider three data-generating processes: we simulate demand and cost data assuming that (1) $\lambda = 0$, as in the original *BLP* specification; (2) $\lambda = 1$, resulting in quasi-linear demand; and (3) $\lambda = 0.7$, an in-between case with weaker income effects than case (1): the distribution of α_i is compressed, with a coefficient of variation of only 0.56, relative to 3.57 for the case of $\lambda = 0$. In the following, we denote case (1) as 'log'; case (2) as 'linear'; and case (3) as 'box-cox' or 'bc'.

Table 1: Monte-Carlo: Parameter Estimates

Scenario	α (varies)		λ (varies)		$\sigma_x = 5$		$\sigma_0 = 5$		*Coeff.Var*		*MAB*		*Corr.*	
True-Specification	*A.Bias*	*RMSE*	*A.Bias*	*RMSE*	*A.Bias*	*RMSE*	*A.Bias*	*RMSE*	σ_α/α	$\hat{\sigma}_\alpha/\hat{\alpha}$	ε	ρ	(ε,ρ)	$(\hat{\varepsilon},\hat{\rho})$
1: log–log	0.003	0.161	0.000	0.000	-0.006	0.072	-0.012	0.231	-3.81	-3.79	0.00	0.00	0.66	0.66
2: linear–linear	0.001	0.011	-	-	0.015	0.090	-0.082	0.947	0.00	0.00	0.00	0.00	0.66	0.66
3: bc–bc	0.000	0.037	-0.001	0.024	0.006	0.079	-0.001	0.735	-0.57	-0.57	0.00	0.00	-0.47	-0.47
4: log–bc	0.331	0.379	0.005	0.006	-0.012	0.070	0.025	0.121	-3.81	-3.77	0.00	0.00	-0.47	-0.47
5: linear–bc	-0.031	0.048	-0.060	0.085	0.006	0.091	0.093	1.109	0.00	-0.11	0.00	-0.01	-0.44	-0.43
6: bc–log	-15.514	15.612	-	-	0.851	0.947	-2.211	2.218	-0.57	-3.77	0.55	-0.69	-0.44	0.63
7: bc-linear	0.248	0.248	-	-	0.015	0.091	-0.272	0.987	-0.57	0.00	-0.16	0.22	-0.44	0.43

Table Notes: The first column indicates the true data-generating process and the researcher's assumed specification of the price-income interactions. The next three (double) columns report the average bias (*A.Bias*) and root mean standard error (*RMSE*) of the income parameter λ and drivers of the idiosyncratic characteristics tastes using 1,000 replications for each scenario. The price coefficient, α, varies for each replication to ensure that $\varepsilon = 2.5$. The attribute random coefficients σ_x and σ_0 (constant) are both set to 5. Column "*Coeff.Var*" reports the coefficient of variation of the distribution of price responsiveness of the data-generating process and the estimated model. The remaining set of columns report the coefficient of variation for idiosyncratic prices-sensitivity parameters (α_i), the median average bias (*MAB*) for average product elasticity and curvature (ε, ρ), and the average correlation between product-level elasticity and curvature (corr(ε_j, ρ_j)).

With these three data sets, we then estimate seven specifications. In scenarios (1)-(3), we specify the demand model correctly and verify that we can recover the underlying preferences using the above instrumentation strategy. In scenarios (4) and (5), we specify general 'box-cox' preferences to recover the simpler 'log' and 'linear' preferences. Lastly, in scenarios (6) and (7) we investigate model misspecification by using either a 'log' or a 'linear' demand model in estimation to recover 'box-cox' preferences.

Discussion of Results. We present the parameter estimates in Table 1 for seven scenarios. In general, across curvature targets, the estimation succeeds at recovering the underlying parameters when the researcher's preference specification coincides with the true underlying data-generating process, i.e., Scenarios (1)-(3), consistent with Gandhi and Houde (2020) and Conlon and Gortmaker (2020). The estimates of elasticity (market power), curvature (pass-through), and their correlation are consistent with the true quantities in the data.

In Scenarios (4) and (5), we model consumer price-sensitivities flexibly using a Box-Cox transformation of outside expenditure and estimates the income parameter λ. The estimates of the Box-Cox model accurately identify the true λ and the random coefficients of product attributes when the underlying preferences include a logarithmic function of income. However, it overestimates the average price responsiveness α. We also observe that the the Box-Cox model accurately recovers the distribution of price sensitivity (columns labeled 'Coeff. Var') and the elasticity-curvature pairs.

Scenarios (6) and (7) address misspecification biases of imposing particular price-income interactions when the true data-generating process is Box-Cox. Scenario (6) assumes the logarithmic transformation of outside good spending, while Scenario (7) assumes quasi-linear preferences of Nevo (2001). The assumed logarithmic specification leads to a substantial misspecification bias in all estimated parameters. The large positive average bias for the random coefficients on the characteristic, σ_x, leads to greater substitution within inside products than the true data. In comparison, the average bias of -2.2 for the random coefficient on the constant indicates greater substitution to the outside option than the true data. Not surprisingly, the economic implications are significant as the average estimated elasticity is -1.95, or 0.55 points less elastic than the true data-generating process. In contrast, the average estimated curvature is 0.69 points above the true data-generating process. Misspecification of price-income interactions as logarithmic, therefore, results in an overestimate of both market power and pass-through. Moreover, specifying log preferences amounts to imposing a different rate of change of the demand elasticity with income from the true relationship under Box-Cox preferences, leading to much greater heterogeneity in price sensitivity than the underlying data. Such a bias has consequences for welfare calculations, especially since solving for changes in consumer surplus requires accounting for income effects. Suppose that we assumed that preferences are quasi-linear, instead, as in scenario (7). Then the estimated elasticity of -2.66 understates firms' true market power while the estimated curvature is 0.22 points below the true data, indicating the estimated model also under-predict the firm pass-through.

The final two columns of Table 1 demonstrate that misspecification impacts the distribution of estimated elasticity-curvature pairs among products. Looking across the different data-generating processes, we observe that the shape of the distribution of price sensitivities via the income distribution determines the demand manifold relationship between demand elasticities and curvature. Imposing specific price sensitivity distributions – Scenarios (6) and (7) – results in a flipped sign of the correlation between product-level elasticities and curvatures, or the slope of the manifold, leading to a mischaracterization of the relationship between market power and pass-through among the products. This could have large consequences for evaluating the economic effects of mergers, cost changes, taxation, or tariffs, particularly for different consumer and firm types.

5.3 Flexible Income Effects: Empirical Implications

In this section, we demonstrate the quantitative and qualitative implications of misspecifying income effects for empirical research. We rely on the automobile data from Berry, Levinsohn and Pakes (1995) to illustrate the elasticity and curvature properties of a *ML* model with income effects modulated by the power parameter λ. Using the same model specification and identification strategy as Berry et al. (1999), we estimate four sets of preferences holding λ fixed at $\lambda = 0$ (*BLP* preferences), $\lambda = 1$ (quasi-linear preferences), $\lambda = 0.5$ and $\lambda = 0.75$. For all estimated models, we incorporate income draws as in Andrews et al. (2017) and follow best practices outlined in Conlon and Gortmaker (2020).

Figure 6 shows the scatter plots of $(\hat{\varepsilon}, \hat{\rho})$ for each automobile model in the *BLP* data under these four alternative specifications.[14] The top left panel represents the quasi-linear case. The average estimated automobile demand elasticity is $\hat{\varepsilon} = 2.75$ with nearly full (single-product) pass-through, $\hat{\rho} = 0.99$, as any mixed *MNL* without idiosyncratic price sensitivity is necessarily log-concave, as shown in Section 4. Note also the sorting of automobiles by price: the estimated demand is substantially more elastic for the most expensive vehicles.

The demand estimates are log-convex for all automobile models whenever we allow for some income effects, as shown in the other three panels of Figure 6. Reducing λ increases the importance of income effects through smaller price responses by higher-income households. Moving from quasi-linear preferences to accounting for income effects does not significantly change the average estimated elasticity, e.g., $\hat{\varepsilon}_{\text{BLP}} = 2.83$ when $\lambda = 0$. Despite the similar average price elasticity, the curvature distribution (pass-through) varies substantially across specifications. This is similar to what we observed in the RTE cereal case of Figure 1.

Curvatures decrease monotonically with λ, with $\hat{\rho} = 0.99$ when $\lambda = 1$ to $\hat{\rho}_{\text{BLP}} = 1.35$ when $\lambda = 0$ (which, in this case, coincides with the curvature of the *CES* model evaluated at the average elasticity: $\hat{\rho}_{CES} = 1 + 1/2.83 = 1.35$). Average pass-through rates increase from 99% in the quasi-linear specification without income effects to 179% with the strong income effect specification of *BLP* demand – dramatically different predictions. The estimated demand for all vehicles is hence sub-convex in the quasi-linear case, but only 55.8% of estimated vehicle demands are sub-convex under the original *BLP* specification. The intermediate cases of $\lambda = 0.5$ and $\lambda = 0.7$ make clear that income effects broadly not only

[14]We report average elasticity, curvature, price markup, and pass-through rate estimates for each scenario in Table F.1 in Appendix F.

Figure 6: Income Effects and Demand Manifolds

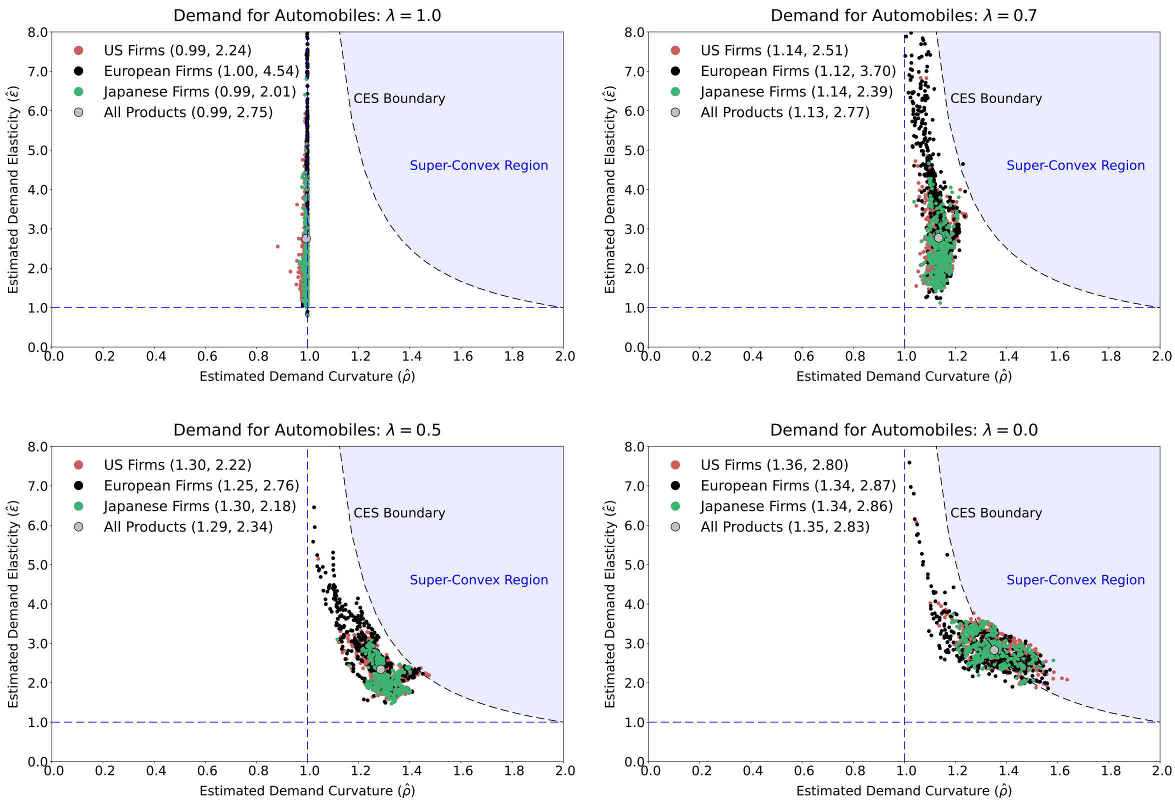

Figure Notes: Each dot represents the point elasticity and curvature estimates for each observation in the *BLP* automobile data, while the gray dot corresponds to the average elasticity and curvature estimates. Point estimates are colored according to vehicle origin and we see significant overlap in (ε, ρ) space.

restrict the range of demand elasticity (and markup) estimates but also expand the range of demand curvature (and pass-through rate) estimates that a discrete choice model of demand can deliver. Appendix F summarizes these results.

Do these differences matter for economic research and policy? We answer this question by giving consumers in the four estimated equilibria a \$1,000 subsidy for purchasing a new domestic vehicle and recompute the Bertrand-Nash pricing equilibrium. We are particularly interested in the degree to which firms adjust their prices to capture or amplify the subsidy across the different demand specifications. We present the equilibrium pass-through rates of the subsidy for domestic vehicles in Figure 7.

As expected, the shape of the mixing distribution is of first-order quantitative importance in evaluating this policy. Under the original *BLP* specification, median pass-through is 1.39. Hence, the subsidy induces domestic automakers to reduce their vehicle prices by more than the subsidy to capture price-sensitive customers. As the mixing distribution

Figure 7: Demand Manifolds and Trade Policy

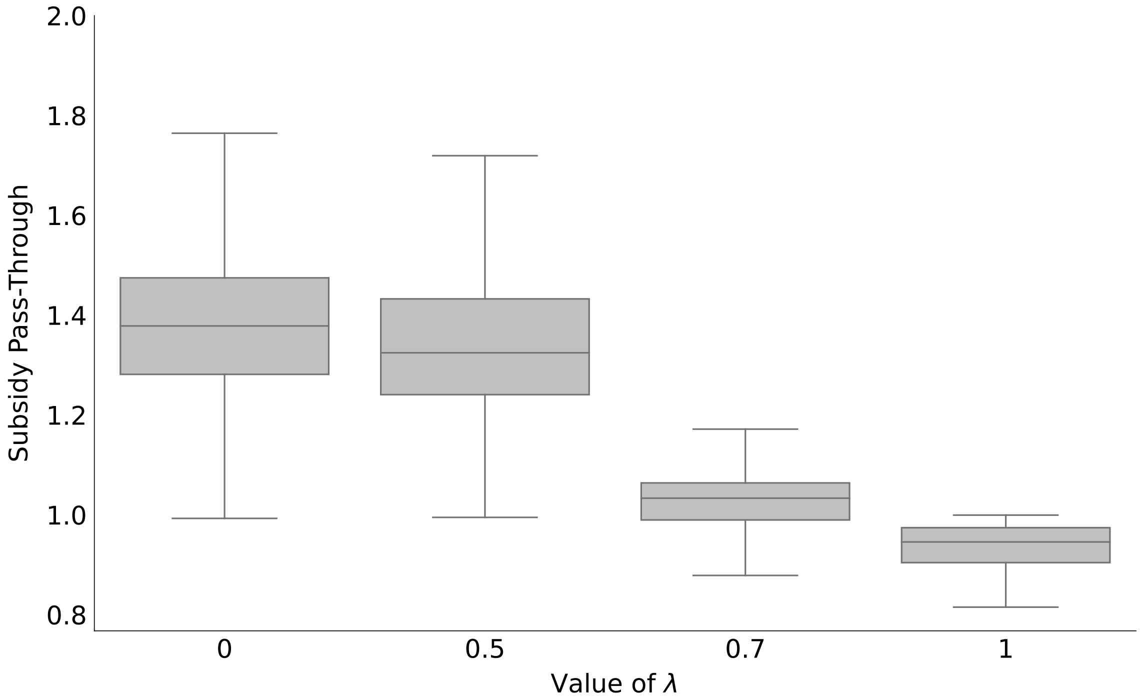

Figure Notes: Figure present pass-through of US vehicles for a hypothetical $1,000 vehicle subsidy paid to US consumers for the purchase of domestic vehicle.

becomes less skewed, pass-through decreases until it becomes under-shifted when $\lambda = 1$. We also observe a compression of the distribution of subsidy pass-through across vehicles. For a policy maker choosing the subsidy required to elicit a particular demand response, these differences are important, as are the welfare implications for consumers. For example, suppose consumers' true preferences are quasi-linear, so that $\lambda = 1$. A researcher who specifies preferences to account for income effects will overestimate the effectiveness of subsidization.

5.4 Flexible Discrete Choice Estimation with Quasi-Linear Preferences

The preceding analysis demonstrated the importance of modeling the distribution of heterogeneous preferences flexibly using preferences with income effects. We now show that identification and Monte Carlo results extend to quasi-linear preferences. Quasi-linearity is a common assumption, even in settings with expensive products where income effects are likely important (e.g., Grieco, Murry and Yurukoglu, 2021), because it provides added flexibility; i.e., it does not require specifying how to include income in the specification or accounting for income effects in counterfactual welfare calculations.

As should be apparent by now, the choice of price mixing distribution materially impacts answers to important empirical questions posed by researchers. Applying the Box-Cox transformation in the case of quasi-linear demand is straightforward and can be done to any mixing distribution, provided the researcher has access to identifying moments which connect changes in consumption across a distribution of individual (or household) characteristics. We consider the following variant of equation (8) where we include $\alpha_i^\star$ to capture consumers' heterogeneous price sensitivity:

$$u_{ij} = x_j \beta_i^\star + \alpha_i^\star \big(y_i - p_j\big) + \xi_j + \epsilon_{ij}\,, \quad i \in \mathcal{I},\ j \in \mathcal{J},\ \epsilon_{ij} \sim \text{EV1}\,. \tag{28}$$

Demographics can flexibly enter the price coefficient $\alpha_i^\star$ in a variety of ways. Nevo (2001) accommodates a non-linear effect of household income on price sensitivity, as prior work has found sizable differences in price elasticities across low- and high-income consumers in a wide variety of markets. However, with quasi-linear preferences, such patterns do not actually represent income effects; they simply capture differences in purchase behavior by consumers of different income levels. There are several ways of introducing such flexibility in $\alpha_i^\star$.[15] The Monte Carlo analysis of Section 5.2 demonstrates, however, that leveraging the Box-Cox transformation provides greater flexibility with minimal computational burden. We therefore model price-sensitivity as follows:

$$\alpha_i = -\exp\big(\alpha + \pi y_i^{(\lambda)}\big), \text{ where } y_i^{(\lambda)} \equiv \begin{cases} \dfrac{y_i^\lambda - 1}{\lambda}\,, & \text{if } \lambda > 0\,, \\ \ln\big(y_i\big)\,, & \text{if } \lambda = 0 \end{cases} \tag{29}$$

This specification of α_i guarantees that all customers have down-ward sloping demands.

A nice feature of this Box-Cox transformation is that it nests common empirical applications. A power parameter of $\lambda = 1$ corresponds to a linear effect of income on price sensitivity, and $\lambda = 0$ denotes the case of log income, but the transform can also accommodate a convex relationship between income and price sensitivity with $\lambda > 1$. While our focus above

[15] One might allow price sensitivity to differ by income bin or sieve estimation, e.g., Wang (2022). We found that both approaches implicitly introduced discrete customer types into the mixing distribution, thereby limiting the shape of the mixing distributions and leading to elasticity-curvature pairs, which deviated substantially from the true shape of demand. We therefore focus on Box-Cox transformations of continuous rather than discrete distribution, and use the power parameter λ to reflect differences in price sensitivity between low- and high-income consumers.

is on allowing flexibility in the price mixing distribution, we could also introduce flexibility via a Box-Cox transformation on demographics for non-price characteristics.

Figure 8: Visualizing the Box-Cox Transform

(a) Transform: $y_i^{(\lambda)}$ (b) Distribution: α_i

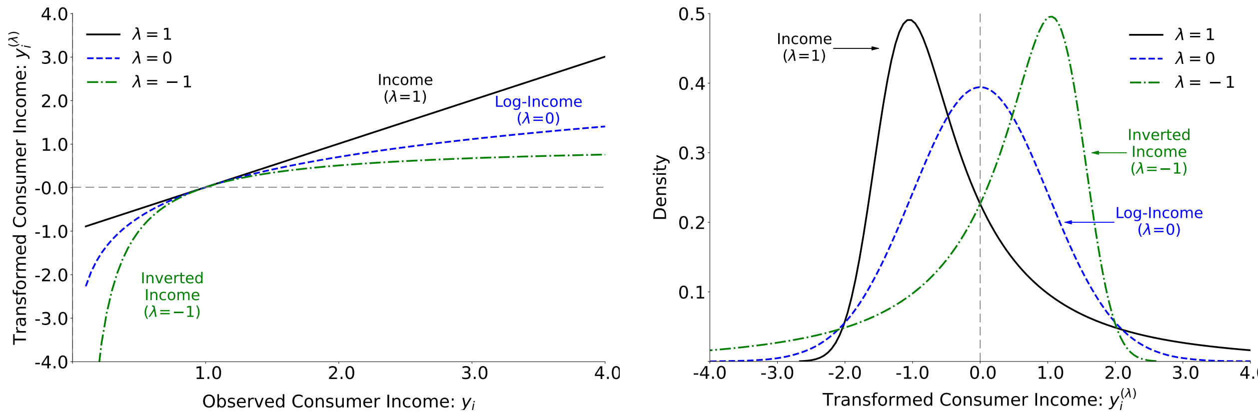

How exactly does the transform influence customer behavior in the model, and what are the implications for the distribution of α_i? Figure 8 depicts examples of the Box-Cox transformation for three values of λ. In Panel (a), we show the transformation for each value of y_i, while in Panel (b), we illustrate the distributional implications for the price sensitivity parameter α_i. For small values of λ, the transform generates most of the variation in price sensitivity among low-income consumers (Panel a), and these low-income consumers drive skewness in the distribution of price sensitivity. As λ increases, variation in price sensitivity increasingly shifts to higher-income consumers.

Identification. As λ modulates the distribution of price-sensitivity across consumers and, therefore, consumption patterns among low- and high-income consumers, identification comes from the likelihood that consumers buy inexpensive versus expensive varieties conditional on income. For example, when $\lambda = 1$, marginal differences in price sensitivity across income levels are uniform. Hence, the predicted average price of the chosen product changes uniformly across income groups, all-else-equal. When $\lambda = 0$, we observe that small differences in income will yield very different consumption sensitivities to price. We would, therefore, observe in the data that the average price paid between consumers across the lowest income groups would look very different while the average price paid among the highest income groups would change little. The opposite is true for the case when $\lambda > 1$ as the gradient in the average price paid across low-income consumers is flat while we observe a large gradient

Table 2: **Quasi-linear Monte-Carlo: Parameter Estimates**

Scenario	α (varies)		$\sigma_x = 5$		$\sigma_0 = 5$		$\pi = -0.2$		λ (varies)	
True-Specification	*A.Bias*	*RMSE*	*A.Bias*	*RMSE*	*A.Bias*	*RMSE*	*A.Bias*	*RMSE*	*A.Bias*	*RMSE*
1: log–log	-0.014	0.016	-0.003	0.056	-0.772	0.787	0.022	0.022	-0.010	0.010
2: linear–linear	-0.043	0.044	-0.003	0.055	-0.671	0.693	0.015	0.016	-	-
3: bc–bc	-0.024	0.025	-0.013	0.055	-0.276	0.293	0.012	0.013	-0.031	0.034
4: log–bc	0.057	0.132	-0.009	0.058	-0.695	0.725	0.040	0.057	0.267	0.556
5: linear–bc	-0.011	0.019	-0.005	0.055	-0.612	0.637	0.032	0.034	0.167	0.194
6: bc–log	-0.562	0.564	-0.189	0.204	-1.761	1.768	-0.292	0.292	-	-
7: bc-linear	-1.015	1.019	-0.402	0.405	0.095	0.286	-1.369	1.375	-	-

Table Notes: The first column indicates the true data-generating process and the researcher's assumed specification of the price-income interactions. The next four (double) columns report the average bias (*A.Bias*) and root mean standard error (*RMSE*) of different model parameters. Each scenario involves 1,000 pricing equilibrium and estimation simulations. The price coefficient, α, varies for each replication to ensure that $\varepsilon = 2.5$. The attribute random coefficients σ_x and σ_0 (constant) are both set to 5, the income price coefficient is set to $\pi = -0.2$, and the Box-Cox parameter is set to $\lambda = -1$ for all simulations, where applicable.

across high-income consumers. A similar argument holds for the Box-Cox transform of demographics related with non-price product characteristics.

Monte Carlo Simulation. We follow a similar process to the Monte Carlo simulation described in Section 5.2 but now applied to (indirect) utility specification (28). We specify $\pi = -0.2$ so high-income agents are less sensitive to changes in price and therefore their demand is less elastic than low-income consumers. We set $\lambda = -1$ to demonstrate an interesting case which aligns with the *BLP* specification we studied earlier.

All other parameters are specified – or solved for – as described in Section 5.2. Similarly, estimation and identification is no different with the exception of adding instruments to separately identify mean price sensitivity (α) from heterogeneous price sensitivity (π). We identify α by including $\hat{p}$ as an instrument (or equivalently cost shocks ω). Movement in these cost shocks drive exogenous shifts in prices which elicit a common demand response.

We identify π by via aggregate income shocks which vary by market t. Specifically, we solve for the average income for each market (or period) t and generate the income distribution across markets. We then construct three indicator variables which is equal to one if a market t is in the bottom ten percent, top ten percent, or between 40th and 60th percentiles. We create the identifying instruments for π by interacting $\hat{p}$ (or equivalently cost shocks ω) with these indicator variables. The identifying assumption then follows the logic of Figure 8, i.e., the instruments trace out how common cost shocks impact different income groups at different rates. For example, if prices increase exogenously and the demand

**Table 3: Quasi-linear Monte-Carlo:
Implications for the Estimated Shape of Demand**

True-Specification	σ_α/α	$\hat{\sigma}_\alpha/\hat{\alpha}$	ε	ρ	(ε,ρ)	$(\hat{\varepsilon},\hat{\rho})$
1: log–log	-0.256	-0.229	-0.143	0.009	-0.410	-0.371
2: linear–linear	-0.115	-0.109	-0.125	-0.002	-0.379	-0.362
3: bc–bc	-9.927	-10.631	-0.092	0.004	-0.251	-0.263
4: log–bc	-0.256	-0.190	-0.148	0.019	-0.410	-0.374
5: linear–bc	-0.115	-0.105	-0.124	-0.001	-0.379	-0.361
6: bc–log	-9.927	-0.713	-0.525	-0.005	-0.252	0.202
7: bc-linear	-9.927	-0.489	-0.692	0.172	-0.251	-0.311

Table Notes: The first column indicates the true data-generating process and the researcher's assumed specification of the price-income interactions. Column "*Coeff. Var*" reports the coefficient of variation of the distribution of price responsiveness of the data-generating process and the estimated model. The remaining set of columns report the coefficient of variation for idiosyncratic prices sensitivity parameters (α_i), the median average bias (MAB) for average product elasticity and curvature (ε,ρ), and the average correlation between product-level elasticity and curvature ($\text{corr}(\varepsilon_j,\rho_j)$).

response is largest among markets in bottom ten percent, the estimator will choose λ values closer to negative one (and vice-versa). We use these instruments for all specifications since these are commonly used in the existing literature and therefore provide a useful example.

Results. We present Monte Carlo estimation results in Table 2. As before, we find that we are able to recover the λ parameter when the true data are generated from the flexible Box-Cox model and we also allow for flexibility. We are also able to capture nested models popular in the literature which use either income or log-income when we allow for flexibility. Estimation bias increases substantially when the researcher imposes the relationship between income and price-sensitivity but the true data-generating process is different.

We demonstrate the implications of imposing a relationship between income and price sensitivity in Table 3, particularly scenarios 6 and 7. The first two columns demonstrate that imposing price sensitivity to be log-linear or linear in income leads to estimated distributions of price sensitivity (summarized by their respective coefficients of variation) that differ substantially from the true data, with important implications for estimated distributional welfare consequences of e.g., policy-induced cost changes. We observe that in both scenarios 6 and 7, the researcher over-estimates demand elasticity, or equivalently under-estimates the firms' market power. When the researcher imposes income interacted with price and the true data generating process is the flexible Box-Cox specification with $\lambda = -1$, they under-estimate demand curvature. We find that the misspecification bias when the researcher assume log-income on estimated demand curvature is small, however. This stems from the fact that our experiments assumed $\lambda = -1$. In all cases, our results indicate that we can recover the transform parameter (λ) and accurately estimate average elasticity (ε), demand curvature (ρ), and the distribution of price sensitivities using the Box-Cox model combined with commonly used instruments.

6 Concluding Remarks

We have shown that the unit-demand mixed-logit model accommodates a wide array of empirically relevant elasticity-curvature pairs, thereby providing further evidence of the power of the mixed-logit model as a demand framework and policy tool. We have also demonstrated how different components of the demand specification contribute to expanding the set of attainable elasticity-curvature pairs to better approximate the true shape of demand. Our theoretical and empirical results highlight the importance of modeling mixing distributions flexibly to keep a healthy distance between assumptions and results. As the Box-Cox transformation we rely on is simple to incorporate and the estimation can be done with standard econometric techniques, allowing for this flexibility has a high substantive return with only minor additional cost.

Our empirical setting demonstrated that modeling the distribution of customer preferences flexibly is important for designing and evaluating trade policy but the analyses here inform a range of other empirical contexts. First, cost pass-through in the international trade and macroeconomic literature is often driven by the assumption of *CES* demand. Our results indicate that this assumption may lead to over-estimated cost pass-through, including exchange rates. Second, the trade subsidy we offered customers in our empirical exercise is similar to subsidies given to consumers who buy an electric vehicle (EV) under the *Inflation Reduction Act of 2022*. Our results indicate that robust estimates of the effectiveness of this policy at generating incremental EV purchases requires modeling the distribution of customer preferences flexibly.

Another area where accurate measurement of demand curvature plays a key role in evaluating welfare is third-degree price discrimination (Aguirre, Cowan and Vickers, 2010). DellaVigna and Gentzkow (2019) document that retailers frequently choose to price uniformly despite the well-documented socioeconomic differences across local markets. Uniform pricing may also result from regulation, e.g., the 2010 Affordable Care Act (ACA) requires health insurers to set uniform prices within predefined "rating areas," covering a collection of counties or zip codes with a variety of customer types. In our companion paper, Miravete, Seim and Thurk (2025), we explore how misspecification of price responsiveness impact the sign and size of redistribution effects of uniform pricing in the breakfast cereal market.

References

Adão, Rodrigo, Arnaud Costinot, and Dave Donaldson (2017): "Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade," *American Economic Review*, Vol. 107, pp. 633–689.

Aguirre, Iñaki, Simon Cowan, and John Vickers (2010): "Monopoly Price Discrimination and Demand Curvature," *American Economic Review*, Vol. 100, pp. 1601–1615.

Amir, Rabah, Isabelle Maret, and Michael Troege (2004): "On Taxation Pass-Through for a Monopoly Firm," *Annals of Economics and Statistics*, Vol. 75-76, pp. 155–172.

Anderson, Simon P. and André de Palma (2020): "Decoupling the CES Distribution Circle with Quality and Beyond: Equilibrium Distributions and the CES-Logit Nexus," *Economic Journal*, Vol. 130, pp. 911–936.

Anderson, Simon P., André de Palma, and Jaques-François Thisse (1992): *Discrete Choice Theory of Product Differentiation*, Cambridge, MA: The MIT Press.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro (2017): "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *Quarterly Journal of Economics*, Vol. 132, pp. 1553–1592.

Bagnoli, Mark and Ted Bergstrom (2005): "Log-Concave Probability and Its Applications," *Economic Theory*, Vol. 26, pp. 445–469.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, Vol. 63, pp. 841–890.

——— (1999): "Voluntary Export Restraints on Automobiles: Evaluating Strategic Trade Policy," *American Economic Review*, Vol. 89, pp. 400–430.

Berry, Steven T. and Philip A. Haile (2021): "Foundations of Demand Estimation," in Kate Ho, Ali Hortaçsu, and Alessandro Lizzeri eds. *Handbook of Industrial Organization,* Vol. 4, New York, NY: North-Holland.

Besanko, David, Jean-Pierre Dubé, and Sachin Gupta (2005): "Own-Brand and Cross-Brand Retail Pass-Through," *Marketing Science*, Vol. 24, pp. 123–137.

Besley, Timothy and Harvey S. Rosen (1999): "Sales Taxes and Prices: An Empirical Analysis," *National Tax Journal*, Vol. 52, pp. 157–78.

Birchall, Cameron, Debashrita Mohapatra, and Frank Verboven (2025): "Estimating Substitution Patterns and Demand Curvature in Discrete-Choice Models of Product Differentiation," *Review of Economics and Statistics*, Vol. 107, p. Forthcoming.

Björnerstedt, Jonas and Frank Verboven (2016): "Does Merger Simulation Work? Evidence from the Swedish Analgesics Market," *American Economic Journal: Applied Economics*, Vol. 8, pp. 125–164.

Box, George E.P. and David R. Cox (1964): "An Analysis of Transformations," *Journal of the Royal Statistical Society, Series B*, Vol. 26, pp. 211–252.

Bulow, Jeremy I. and Paul Pfleiderer (1983): "A Note on the Effect of Cost Changes on Prices," *Journal of Political Economy*, Vol. 91, pp. 182–185.

Caplin, Andrew and Barry Nalebuff (1991a): "Aggregation and Imperfect Competition: On the Existence of Equilibrium," *Econometrica*, Vol. 59, pp. 25–59.

——— (1991b): "Aggregation and Social Choice: A Mean Voter Theorem," *Econometrica*, Vol. 59, pp. 1–23.

Compiani, Giovanni (2022): "Market Counterfactuals and the Specification of Multi-Product Demand: A Nonparametric Approach," *Quantitative Economics*, Vol. 13, pp. 545–591.

Conlon, Christopher T. and Jeff Gortmaker (2020): "Best Practices for Demand Estimation with pyBLP," *RAND Journal of Economics*, Vol. 51, pp. 1108–1161.

Conlon, Christopher T. and Nirupama L. Rao (2020): "Discrete Prices and the Incidence and Efficiency of Excise Taxes," *American Economic Journal: Economic Policy*, Vol. 12, pp. 111–43, November.

Cournot, Antoine A. (1838): *Recherches sur les principes mathématiques de la théorie des richesses*, Paris, France: Hachette.

Cuesta, José Ignacio and Alberto Sepúlveda (2021): "Price Regulation in Credit Markets: A Trade-off Between Consumer Protection and Credit Access," working paper, Stanford University.

De Loecker, Jan, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik (2016): "Prices, Markups and Trade Reform," *Econometrica*, Vol. 84, pp. 445–510.

DellaVigna, Stefano and Matthew Gentzkow (2019): "Uniform Pricing in U.S. Retail Chains," *The Quarterly Journal of Economics*, Vol. 134, pp. 2011–2084.

Fabra, Natalia and Mar Reguant (2014): "Pass-Through of Emissions Costs in Electricity Markets," *American Economic Review*, Vol. 104, pp. 2872–2899.

Froeb, Luke M., Steven Tschantz, and Gregory J. Werden (2005): "Pass-Through Rates and the Price Effects of Mergers," *International Journal of Industrial Organization*, Vol. 23, pp. 703–715.

Gandhi, Amit and Jean-François Houde (2020): "Measuring Substitution Patterns in Differentiated Products Industries," working paper, University of Wisconsin-Madison.

Genakos, Christos and Mario Pagliero (2022): "Competition and Pass-Through: Evidence from Isolated Markets," *American Economic Journal: Microeconomics*, Vol. 14, p. forthcoming.

Grieco, Paul L. E, Charles Murry, and Ali Yurukoglu (2021): "The Evolution of Market Power in the U.S. Auto Industry," Discussion Paper 29013, National Bureau of Economic Research.

Griffith, Rachel, Lars Nesheim, and Martin O'Connell (2018): "Income Effects and the Welfare Consequences of Tax in Differentiated Product Oligopoly," *Quantitative Economics*, Vol. 9, pp. 305–341.

Gutierrez, German (2022): "The Welfare Consequences of Regulating Amazon," working paper, University of Washington.

Hanemann, W. Michael (1984): "Discrete/Continuous Models of Consumer Demand," *Econometrica*, Vol. 52, pp. 541–562.

Head, Keith and Thierry Mayer (2025): "Poor Substitutes? Counterfactual Methods in IO and Trade Compared," *Review of Economics and Statistics*, Vol. 107, p. Forthcoming.

Hellerstein, Rebecca and Pinelopi Koujianou Goldberg (2013): "A Structural Approach to Identifying the Sources of Local Currency Price Stability," *Review of Economic Studies*, Vol. 80, pp. 175–210.

Hollenbeck, Brett and Kosuke Uetake (2021): "Taxation and Market Power in the Legal Marijuana Industry," *RAND Journal of Economics*, Vol. 19, pp. 417–455.

Jaffe, Sonia and E. Glen Weyl (2010): "Linear Demand Systems are Inconsistent with Discrete Choice," *B.E. Journal of Theoretical Economics (Advances)*, Vol. 10. `https://doi.org/10.2202/1935-1704.1729`.

——— (2013): "The First-Order Approach to Merger Analysis," *American Economic Journal: Microeconomics*, Vol. 5, pp. 188–218.

Kenkel, Don (2005): "Are Alcohol Tax Hikes Fully Passed Through to Prices? Evidence from Alaska," *American Economic Review*, Vol. 95, pp. 273–277.

Kimball, Miles S. (1995): "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit and Banking*, Vol. 27, pp. 1241–1277.

Marion, Justin and Erich Muehlegger (2011): "Fuel Tax Incidence and Supply Conditions," *Journal of Public Economics*, Vol. 95, pp. 1202–1212.

Marshall, Alfred (1920): *Principles of Economics*, London, UK: Macmillan, 8th edition.

McFadden, Daniel L. and Kenneth E. Train (2000): "Mixed MNL Models for Discrete Response," *Journal of Applied Econonometrics*, Vol. 15, pp. 447–470.

Miller, Nathan H., Marc Remer, Conor Ryan, and Gloria Sheu (2015): "Pass-Through Rates and the Prediction of Mergers Price Effects," *Journal of Industrial Economics*, Vol. 64, pp. 684–709.

Miller, Nathan H., Matthew Osborne, and Gloria Sheu (2017): "Pass-Through in a Concentrated Industry: Empirical Evidence and Regulatory Implications," *The RAND Journal of Economics*, Vol. 48, pp. 69–93.

Miravete, Eugenio J., Katja Seim, and Jeff Thurk (2018): "Market Power and the Laffer Curve," *Econometrica*, Vol. 85, pp. 1651–1687.

——— (2025): "Targeted Pricing and the Shape of Demand," working paper, University of Texas at Austin.

Mrázová, Monika, J. Peter Neary, and Mathieu Parenti (2021): "Sales and Markup Dispersion: Theory and Empirics," *Econometrica*, Vol. 89, pp. 1753–1788.

Mrázová, Monika and J. Peter Neary (2017): "Not So Demanding: Demand Structure and Firm Behavior," *American Economic Review*, Vol. 107, pp. 3835–3874.

Nadarajah, Saralees (2005): "A Generalized Normal Distribution," *Journal of Applied Statistics*, Vol. 32, pp. 685–694.

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

Nevo, Aviv (2000): "Mergers With Differentiated Products: The Case of the Ready-To-Eat Cereal Industry," *RAND Journal of Economics*, Vol. 31, pp. 395–421.

——— (2001): "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, Vol. 69, pp. 307–342.

Nocke, Volker and Nicolas Schutz (2018): "Multiproduct-Firm Oligopoly: An Aggregative Games Approach," *Econometrica*, Vol. 86, pp. 523–557.

Poterba, James M. (1996): "Retail Price Reactions to Changes in State and Local Sales Taxes," *National Tax Journal*, Vol. 49, pp. 165–76.

Shapiro, Carl (1996): "Mergers with Differentiated Products," *Antitrust*, Vol. 10, pp. 23–30.

Train, Kenneth E. (2009): *Discrete Choice Methods with Simulation*, New York, NY: Cambridge University Press, 2nd edition.

Wang, Ao (2022): "Sieve BLP: Sieve Blp: A Semi-nonparametric Model of Demand for Differentiated Products," *Journal of Econometrics*, p. forthcoming.

Weyl, E. Glen and Michal Fabinger (2013): "Pass-Through as an Economic Tool: Principles of Incidence Under Imperfect Competition," *Journal of Political Economy*, Vol. 121, pp. 528–583.

Online Appendix

A Elasticity and Curvature of Demand for Breakfast Cereal

This appendix describes the estimation of Figure 1 for ready-to-eat breakfast cereal. As we note in the main test, the specification for Berry et al. (1999) presented in Panel D follows Conlon and Gortmaker (2020) using income data as presente in . Nevo (2000) specifies preferences as follows (ignoring market location and time indices):

$$u_{ij} = x_j \beta_i^\star + \alpha_i^\star p_j + \xi_j + \epsilon_{ij}\,, \qquad i \in \mathcal{I},\ j \in \mathcal{J},\ \epsilon_{ij} \sim \text{EV1}\,, \tag{A.1a}$$

$$\begin{pmatrix} \alpha_i^\star \\ \beta_i^\star \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} + \Pi D_i + \Sigma \nu_i\,, \qquad \nu_i \sim N(0, I_{n+1})\,, \tag{A.1b}$$

where x_j is the $(n \times 1)$ vector of observed product characteristics and p_j is the price of (inside) product j available in each market, $\mathcal{J}$, with $J = |\mathcal{J}|$. Payoff of the outside good is $u_{i0} = \epsilon_{i0}$. There are random coefficients of product characteristics, β_i^* and price responsiveness, α_i^*. Preferences might be correlated with a d-vector of demographic traits D_i through the $(n{+}1) \times d$ matrix Π of interaction estimates that allow for cross-price elasticity to vary across markets with different demographic composition. To further account for individual preferences over unobservable product attributes, ν_i captures mean-zero, unobserved preference shifters with a diagonal variance-covariance matrix Σ. Lastly, the idiosyncratic unobserved preference by consumer i for product j, ϵ_{ij}, follows the Type-I extreme value distribution across all products in $\mathcal{J}$.

We consider four alternative specifications. The estimation results of Model A are represented graphically in Panel B of Figure 1. Each specification also includes the product characteristics product characteristics following Nevo (2001) but are not reported. Robust standard errors are in parentheses.

Where does curvature come from in this model? In Model B we removed the price-interactions with demographics, while in Model C we remove the normally-distributed price random coefficient. We observe that curvature is driven by the shape of the price mixing

Table A.1: Breakfast Cereal: Price Related Estimates

	Means	Std. Dev.	Demographic Interactions (π_p)			Manifold	
SPECIFICATION	(α)	(σ_p)	log(INCOME)	log(INCOME)^2	CHILD	ε	ρ
[A]	-62.7299	3.3125	588.3252	-30.1920	11.0546	3.62	1.06
	(14.8032)	(1.3402)	(270.4410)	(14.1012)	(4.1226)		
[B]	-30.9982	2.0216	—	—	—	3.74	0.96
	(0.9674)	(0.9367)	—	—	—		
[C]	-53.1367	—	444.7281	-22.3987	16.3664	3.60	1.08
	(12.1023)	—	(209.6548)	(10.7282)	(4.7824)		
[D]	-30.8902	—	—	—	—	3.74	0.96
	(0.9944)	—	—	—	—		

Table Notes: *GMM* estimates of parameters related to price sensitivity using simulated breakfast cereal data estimated via "best practices" described in Conlon and Gortmaker (2020). The remaining parameters for product characteristics follow Nevo (2001) and are included in each demand specification but are not reported. Robust standard errors are in parentheses.

distribution connect to demographics. This finding is supported in Model D where we observe no heterogeneity in price sensitivity leads to log-concave estimated demand.

$$\text{[A]} \quad \alpha_i^\star = \alpha + \sum_{k=1}^{d} \pi_{\alpha k} D_i + \sigma_\alpha \nu_i \,, \qquad \text{(Nevo – Full Model)} \tag{A.2a}$$

$$\text{[B]} \quad \alpha_i^\star = \alpha + \sigma_\alpha \nu_i \,, \qquad \text{(Only Price Random Coefficient)} \tag{A.2b}$$

$$\text{[C]} \quad \alpha_i^\star = \alpha + \sum_{k=1}^{d} \pi_{\alpha k} D_i \,, \qquad \text{(Only Demographic Price Interactions)} \tag{A.2c}$$

$$\text{[D]} \quad \alpha_i^\star = \alpha \,, \qquad \text{(No Price Interactions)} \tag{A.2d}$$

B Probability Distributions and Demand Manifolds

In this section we provide detail behind the derivations in the main text. Because of the additive i.i.d. type-I extreme value distribution of ϵ_{ij}, the individual i's choice probability of product j given by (10) is also the mean of an individual-specific Bernoulli distribution:

$$\mu_{ij} = \mathbb{P}_{ij} \,, \tag{B.1}$$

which are functions of the vector of prices p that we omit to reduce clutter. The variance is:

$$\sigma_{ij}^2 = \mathbb{P}_{ij}(1 - \mathbb{P}_{ij}) \,. \tag{B.2}$$

And finally, the third central moment or non-standardized skew is:

$$sk_{ij} = \mathbb{P}_{ij}(1-\mathbb{P}_{ij})^2 - \mathbb{P}_{ij}^2(1-\mathbb{P}_{ij}) = \mathbb{P}_{ij}(1-\mathbb{P}_{ij})(1-2\mathbb{P}_{ij})\,, \tag{B.3}$$

from where we obtain standardized moment or *skew* as:

$$\tilde{\mu}_{ij,3} = \frac{sk_{ij}}{\sigma_{ij}^3} = \frac{\mathbb{P}_{ij}(1-\mathbb{P}_{ij})(1-2\mathbb{P}_{ij})}{\sqrt{[\mathbb{P}_{ij}(1-\mathbb{P}_{ij})]^3}} = \frac{1-2\mathbb{P}_{ij}}{\sqrt{\mathbb{P}_{ij}(1-\mathbb{P}_{ij})}}\,, \tag{B.4}$$

where σ_{ij}^3 is the third raw moment of the individual choice probability distribution.

Moment Derivatives. We use the derivative of the choice probability (10) with respect to price repeatedly:

$$\mathbb{P}'_{ij} = \frac{\partial \mathbb{P}_{ij}}{\partial p_j} = f'_{ij} \cdot \mathbb{P}_{ij}(1-\mathbb{P}_{ij})\,. \tag{B.5}$$

The derivative of the variance with respect to price is:

$$\frac{\partial \sigma_{ij}^2}{\partial p_j} = \frac{\partial \mathbb{P}_{ij}(1-\mathbb{P}_{ij})}{\partial p_j} = \mathbb{P}'_{ij}(1-\mathbb{P}_{ij}) - \mathbb{P}_{ij}\mathbb{P}'_{ij} = f'_{ij} \cdot \mathbb{P}_{ij}(1-\mathbb{P}_{ij})(1-2\mathbb{P}_{ij}) = f'_{ij} \cdot sk_{ij}\,. \tag{B.6}$$

To conclude, we obtain the price derivative of skewness by differentiating the first equality in (B.3):

$$sk'_{ij} = \left[(1-\mathbb{P}_{ij})^2 - 4\mathbb{P}_{ij}(1-\mathbb{P}_{ij}) + \mathbb{P}_{ij}^2\right] \cdot \mathbb{P}'_{ij} = \left[(1-2\mathbb{P}_{ij})^2 - 2\mathbb{P}_{ij}(1-\mathbb{P}_{ij})\right] \cdot f'_{ij} \cdot \mathbb{P}_{ij}(1-\mathbb{P}_{ij})\,. \tag{B.7}$$

Demand Manifold. Price differentiate (11) and substitute (B.5) to obtain demand elasticity of product j with respect to p:

$$\varepsilon_j(p) \equiv -\frac{p_j}{Q_j(p)} \cdot \frac{\partial Q_j(p)}{\partial p_j} = -\frac{p_j}{Q_j(p)} \int_{i \in \mathcal{I}} f'_{ij} \cdot \mathbb{P}_{ij}\left(1-\mathbb{P}_{ij}\right)\, dG(i)\,. \tag{B.8}$$

Similarly, the inverse demand curvature of product j is:

$$\rho_j(p) \equiv Q_j(p) \cdot \frac{\partial^2 Q_j(p)/\partial p_j^2}{[\partial Q_j(p)/\partial p_j]^2} = \int_{i\in\mathcal{I}} \mathbb{P}_{ij} dG(i) \times \frac{\left[\begin{array}{c}\int f''_{ij} \cdot \mathbb{P}_{ij}\left(1-\mathbb{P}_{ij}\right) dG(i) + \\ \int \left(f'_{ij}\right)^2 \cdot \left[\mathbb{P}_{ij}\left(1-\mathbb{P}_{ij}\right)\left(1-2\mathbb{P}_{ij}\right)\right] dG(i)\end{array}\right]}{\left[\int f'_{ij} \cdot \mathbb{P}_{ij}\left(1-\mathbb{P}_{ij}\right) dG(i)\right]^2} . \tag{B.9}$$

Equations (12) and (13) follow after substituting (11), (B.2) and (B.3) into these expressions. Combining elasticity and curvature we obtain the expression for the demand manifold (14):

$$\rho_j[\varepsilon_j(p)] = \frac{p_j^2}{\varepsilon_j^2(p) \cdot Q_j(p)} \cdot \left[\int f''_{ij} \cdot \mathbb{P}_{ij}\left(1-\mathbb{P}_{ij}\right) dG(i) + \int \left(f'_{ij}\right)^2 \cdot \left[\mathbb{P}_{ij}\left(1-\mathbb{P}_{ij}\right)\left(1-2\mathbb{P}_{ij}\right)\right] dG(i)\right] . \tag{B.10}$$

C A General Mixing Distribution

Without loss of generality, suppose idiosyncratic demand sensitivity is modeled as $\alpha_i^\star = \alpha + \pi\phi_i$, where α is the mean slope of demand and π captures the effect on price heterogeneity of preferences across individuals. We model draws of individual types ϕ_i after the following three-parameter Asymmetric Generalized Normal distribution (Nadarajah, 2005):

$$\text{Prob}\big(\phi < x; \iota, \zeta, \kappa\big) = \Phi_N(y) \text{ where } = \begin{cases} \dfrac{-1}{\kappa} \log\left(1 - \dfrac{\kappa(x-\iota)}{\zeta}\right) , & \text{if } \kappa \neq 0 , \\ \dfrac{x-\iota}{\zeta} , & \text{if } \kappa = 0 , \end{cases} \tag{C.1}$$

and where $\Phi_N(\cdot)$ denotes the cumulative distribution function of a standard normal. To avoid an overparameterized model, we normalize the scale parameter $\zeta = 1$, and $\kappa < 0$ so that the support of the distribution is $(\iota + 1/\kappa, \infty)$. The distribution is right-skewed, mimicking a log-normal distribution for $\kappa = -1$ and converging to a normal distribution as $\kappa \longrightarrow 0$. Furthermore, we center the distribution around the mean slope:

$$E[\phi] = \iota - \frac{\zeta}{\kappa}\left(e^{\kappa^2/2} - 1\right) = 0 , \tag{C.2}$$

so that:

$$\iota = \frac{1}{\kappa}\left(e^{\kappa^2/2} - 1\right) . \tag{C.3}$$

The one-parameter (κ) Asymmetric Generalized Normal distribution can then be written as:

$$\text{Prob}\big(\phi < x; \kappa\big) = \Phi_N(y) \text{ where } = \begin{cases} -\dfrac{\log\left(e^{\kappa^2/2} - \kappa x\right)}{\kappa}, & \text{if } \kappa \neq 0, \\ \dfrac{x - \iota}{\zeta}, & \text{if } \kappa = 0, \end{cases} \tag{C.4}$$

with ι and ζ defined above. The mean, variance, and skewness are:

$$\mu[\phi; \kappa] = 0, \tag{C.5}$$

$$\sigma^2[\phi; \kappa] = \frac{e^{\kappa^2/2}\big(e^{\kappa^2/2} - 1\big)}{\kappa^2}, \tag{C.6}$$

$$\tilde{\mu}_3[\phi; \kappa] = \frac{3e^{\kappa^2/2} - e^{3\kappa^2/2} - 2}{\big(e^{\kappa^2/2} - 1\big)^{3/2}}. \tag{C.7}$$

D Over-shifted Pass-Through via Heterogenous Price Sensitivities

To make intuition connecting heterogeneous price sensitivity and over-shifted pass-through concrete, we present a simple example of pricing by a monopolist who caters to two consumers with linear demands of different slopes in Figure D.1, Panel A. The monopolist sets prices for each customer and responds to an increase in cost (red lines) by increasing equilibrium prices by half of the cost increase; i.e., pass-through is "under-shifted."

In many empirical settings, firms do not practice such first-degree price discrimination. Panel B shows that in setting a uniform price, the monopolist now faces a kinked demand.[16] At the initial marginal cost and implied optimal price, the monopolist serves both customer types. Once marginal cost increases, the firm maximizes profit by increasing the price and excluding the price-sensitive customer. Pass-through is now over-shifted. More generally, in responding to an increase in cost, a firm serving heterogeneous consumers with a uniform price trades off the standard incentive to remain on the elastic portion of demand and the benefits of catering to less price-sensitive customers only.

[16]Note that the shape of demand depends on the markets firms choose to compete in because such decisions imply consumer preference heterogeneity. Kimball (1995) first suggests a smooth, differentiable version of this kinked demand to ensure subconvexity and markups that increase in the production scale.

Figure D.1: Heterogeneous Customers and the Shape of Demand

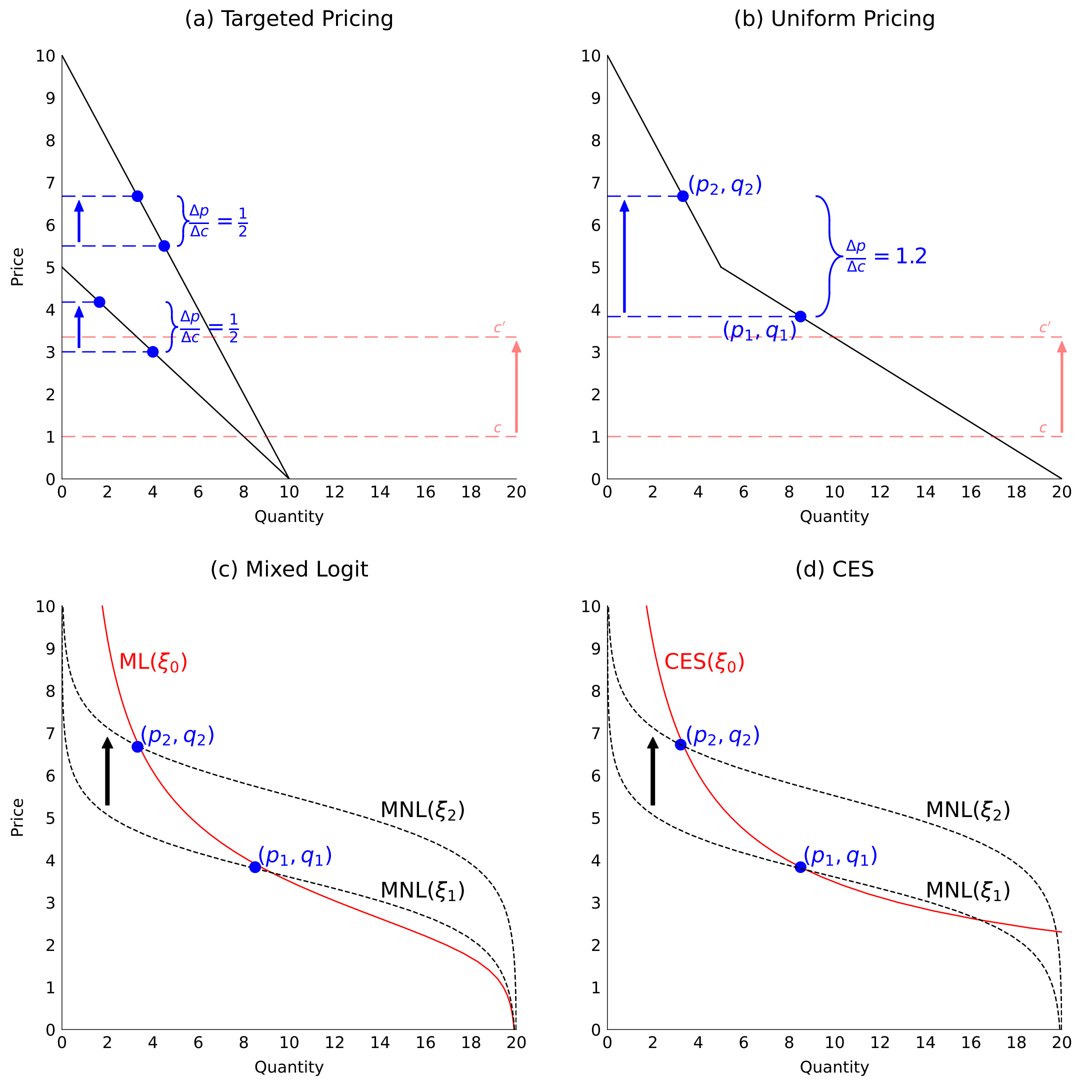

Figure Notes: Top panels present pass-through for linear demand under targeted and uniform pricing given data points (p_1, q_1) and (p_2, q_2). Bottom panels present discrete choice and *CES* rationalizations of the same data points.

How do these theoretical points translate to empirical work? In Panel C, we observe that the kinked demand intuition extends naturally to the *ML* framework since product demand is a function of the underlying mixing distributions.

If we restricted ourselves to *MNL* demand and only observed the two price-quantity pairs in the data, our estimator would infer incorrectly that a positive demand shock had also occurred ($\xi_2 > \xi_1$). This is because the shape of *MNL* demand is not sufficiently flexible to reconcile the first-order conditions at both points without adding a demand shift. When we add flexibility via heterogeneity in price sensitivity, we observe that the demand function (red line) can now contain both points on a single demand curve – just as in the kinked linear demand case.

Panel D shows that these two points are consistent with a single *CES* demand function. The *CES* demand entails two differences from *ML*. First, we observe in the figure that the difference between *ML* and *CES* demand becomes large as the price drops. Consumers purchase discrete quantities in *ML* but can choose arbitrarily small quantities in *CES*. Second, *CES* constrains curvature and pass-through to be constant. Hence, in a neighborhood where the *ML* and *CES* demand functions have similar demand curvatures, e.g., when price exceeds three, *ML* cost pass-through is far below what *CES* would predict. While *CES* is a useful simplification of *ML* for estimation, its pass-through predictions in oligopoly settings are restrictive.

E Competition, Demand Curvature, and Pass-through

Our analysis of pass-through in the main text focused on demand curvature (i.e., the shape of demand) and ignored the impact of competition (i.e., shifts of the demand curve). We address the interaction of curvature and competition in our Monte Carlo environment by varying λ to generate equilibria of varying degrees of demand curvature. We then shock each simulated equilibrium with a common 10% increase in marginal costs and consider two alternate counterfactual equilibria. First, we assume each firm operates as a single-product monopolist. We call this scenario "Monopoly." Second, we assume firms internalize the price choices of their competitors and we therefore solve for new Bertrand-Nash pricing equilibrium. We call this scenario "Oligopoly." We present the median Monopoly (solid line) and Oligopoly (dashed-line) pass-through rates for different levels of demand curvature in Figure E.1.

We find that competition pushes equilibrium pass-through towards one, thereby muting the upward pricing pressure generated by the change in marginal costs. The increase in the common cost leads to both direct and indirect pass-through effects. The price of a product always increases with its own cost. This is the direct effect captured by Monopoly pass-through. The indirect effect collects substitution effects induced by price changes of other products similarly affected by the cost increase. The net effect depends on "how far" a particular product is from its closest substitutes in product space.

Figure E.1: Competition and Pass-Through Rates

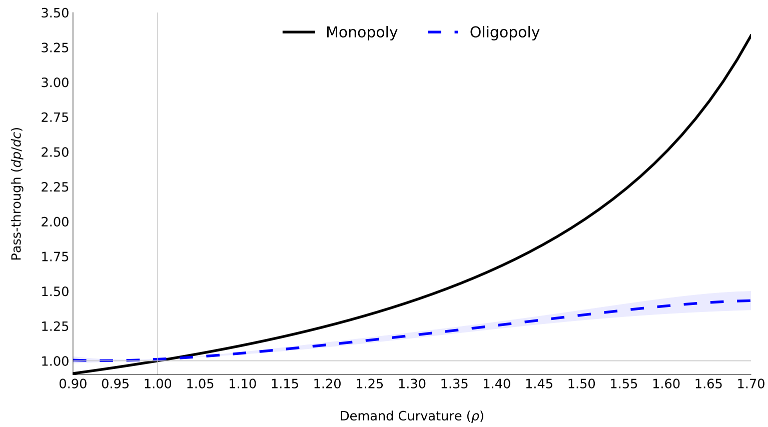

Figure Notes: Figure presents Monte Carlo results across equilibria of median demand curvature. We generate each equilibrium following the environment discussed in Section 5.2 for the Box-Cox utility specification where $\lambda \in [0, 1]$. For each market t in each equilibrium, we solve for the median (across 20 products) demand curvature. "Monopoly" represents the pass-through rate of a single-product monopolist, e.g., (4). "Oligopoly" is the median pass-through rate for each market t in each equilibrium generated by a 10% increase in marginal costs. The shaded region reflects the 95% confidence interval.

Many empirical questions depend critically on the relative importance of the "direct effect" and "net effect." The literature has focused on doing this by getting substitution patterns right. Our work highlights the importance of also getting the shape of demand right. We turn now to demonstrating that this focus on the shape of demand is important for designing and evaluating trade policy.

F Additional Results

Table F.1: Income Effects, Markups, and Pass-Through Rates

	$\lambda = 0$		$\lambda = 0.5$		$\lambda = 0.75$		$\lambda = 1$	
Elasticity (ε)	2.83	(0.26)	2.34	(0.48)	2.77	(1.01)	2.75	(2.05)
Curvature (ρ)	1.35	(0.08)	1.19	(0.07)	1.13	(0.05)	0.99	(0.01)
Markup (%)	44.41	(5.26)	46.25	(8.77)	44.48	(13.77)	48.12	(20.55)
Pass-Through (%)	178.99	(18.33)	145.91	(16.38)	117.90	(7.27)	99.41	(0.01)

Table Notes: Mean and standard deviations (in parentheses) of demand elasticity and curvature plus their implied price markup and pass-through rate.