Elasticity and Curvature of Discrete Choice Demand Models [⋆]

Eugenio J. Miravete[†] Katja Seim[‡] Jeff Thurk[§]

November 2024

Abstract

We study the determinants of cost pass-through in differentiated product markets. Random utility models of demand, such as mixed logit, are attractive because they place limited restrictions on customer substitution patterns. We show the shape of the distributions of customer preferences determines cost pass-through. Common functional form assumptions for these distributions 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) up to *CES* demand. Instruments and estimation are straightforward, and Monte Carlo analysis validates our ability to recover the underlying demand curvature. Using automobile data, we find the bias from ex-ante shape restrictions is large. Results show that flexibly estimating cost pass-through has important implications for evaluating trade and subsidy policies.

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, through its impact on cost pass-through, drives the conclusions to many substantive questions in industrial organization (IO), including the ability of digital platforms such as Amazon.com to affect the division of surplus between third-party sellers and consumers (Gutierrez, 2022), the welfare implications of uniform pricing observed in settings ranging from consumer packaged goods (DellaVigna and Gentzkow, 2019) to consumer financial products (Cuesta and Sepúlveda, 2021), or the predicted price effects of horizontal mergers that generate cost efficiencies.[1] 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).

These examples highlight the value to a flexible demand specification to prevent functional form assumptions from impacting curvature predictions. Bulow and Pfleiderer (1983) underscore this point by illustrating how the demand specification can skew statistical tests assessing the presence of market power in the context of the tobacco industry. Likewise, Froeb, Tschantz and Werden (2005) find that the predicted pass-through rates of cost efficiencies in the WorldCom–Sprint merger are seven times greater when employing a constant elasticity of substitution (CES) demand system compared to a linear demand system.

We focus on discrete-choice demand models and examine the connection between preference specification and the set of feasible substitution and pass-through combinations accommodated by the demand model. The mixed-logit (ML) model, in particular, can capture realistic substitution patterns across heterogeneous consumers. 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:

[1] Such price effects depend on the concavity of the profit function and thus demand curvature. Jaffe and Weyl (2013) suggest that for small merger-induced price increases, observed pass-through rates allow inference of the concavity of profit. For large price changes, Miller, Remer, Ryan and Sheu (2015) suggest conducting a merger simulation with a demand system constrained to mimic observed pass-through.

> ...[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.

We aim to highlight the implications of modeling choices for representing consumer preference heterogeneity in answering questions such as: When do assumptions on preference heterogeneity restrict feasible curvature estimates and pass-through? How can we model preference heterogeneity flexibly to simultaneously allow for the estimation of realistic market power and substitution (elasticity) and pass-through (curvature)?

Motivating Examples. We begin by illustrating the consequences of preference heterogeneity on elasticity and demand curvature estimates 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). We focus on elasticity and curvature pairs as descriptive statistics for the shape of demand. While using (own-price) elasticity as a simple measure of market power will likely be familiar to the reader, demand curvature as a simple measure of cost pass-through is likely new. In Section 2, we formally define demand curvature and its connection to cost pass-through. For now, the reader should take substitution estimates "as given" and to simplify, think of curvature and pass-through as equivalent; i.e., if a product's estimated demand curvature is 0.9, the model predicts the firm will increase price 90 cents when marginal cost increases by one dollar. Of course, the actual pass-through rate will also depend on the substitution patterns of the estimated demand, something that we explicitly addres later in the paper.

Figure 1 Panel A summarizes estimation results for a simple multinomial logit (*MNL*) model of cereal demand that represents preferences as a common linear function of cereal attributes and price. Each dot represents the estimated own-price elasticity (ε) and curvature (ρ) pair of a single product evaluated at observed prices. We find that estimated product demands are elastic and that estimated curvature is less than – and truncated at – one.

In Panel B, we present the results for the same data using Nevo's (2000) mixed-logit (*ML*) specification which adds normally distributed heterogeneity in both price sensitivity

Figure 1: Example Elasticity and Curvature Estimates

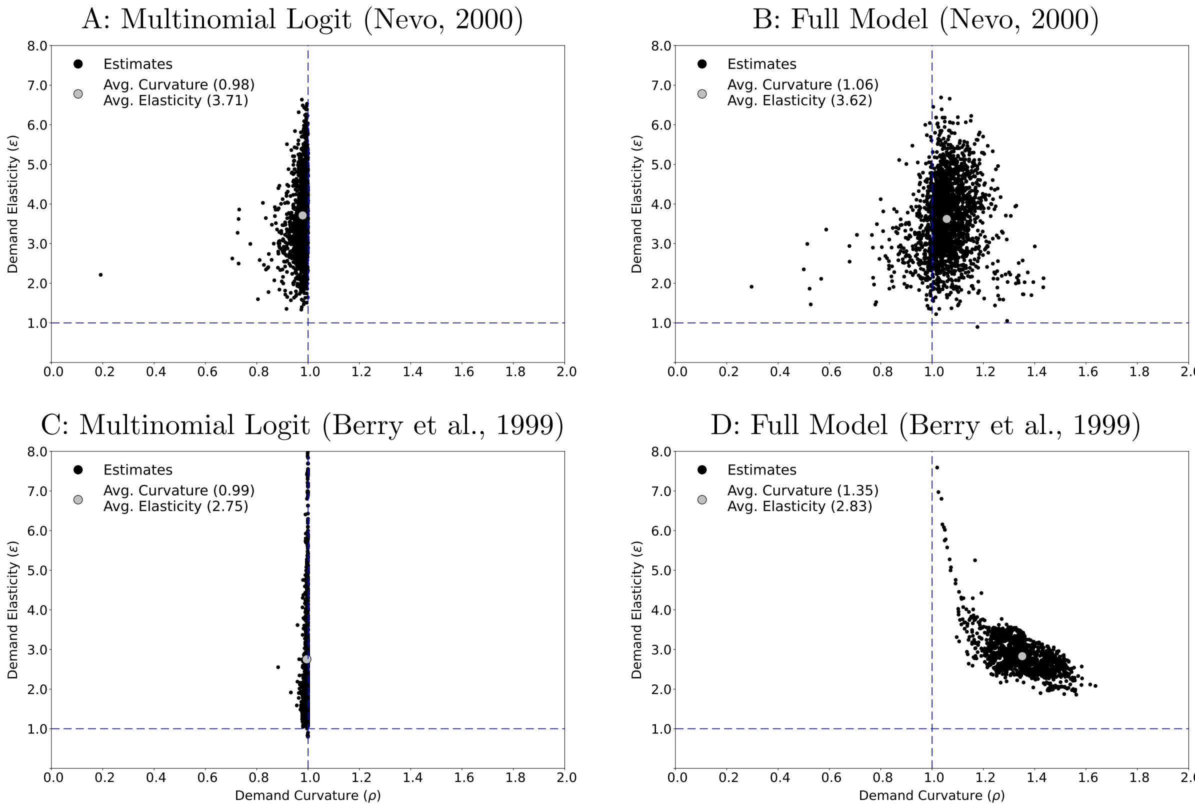

Figure Notes: Black dots represent the estimated own-price elasticity and curvature for a sample product. The gray dot corresponds to the average elasticity and curvature. Demand estimates in the top panels are based on the simulated ready-to-eat cereal data from Nevo (2000), while the bottom panels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

and valuation of product attributes to the *MNL* model.[2] We observe that estimated demand curvatures in the *ML* model exceed one for the majority of the products, which indicates that demand is "overshifted": a one-dollar increase in cost results in more than a one-dollar increase in price for these products. While the two specifications deliver similar average elasticities – a statistic often reported by researchers in the literature – the *MNL* model predicts near complete pass-through for the modal product, as in perfectly competitive settings. As the *ML* model nests the *MNL* model, this points to the importance of preference heterogeneity in generating these discrepancies.

[2] We estimate each model using Nevo's (2000) original set of Hausman-style price instruments. We also considered other demand specifications that rely on only an idiosyncratic price random coefficient or on only price-demographic interactions to represent heterogeneity in price sensitivity. The full description of the different specifications and the relevant estimates are reported in Online Appendix B.

We repeat the exercise for new automobile purchases in panels (C) and (D). We observe that *ML* demand again implies greater diversity of demand curvature and, therefore, also a greater diversity of pass-through. As a foreshadowing of later results, the shape of the dots in panel (D) reflects the underlying Cobb-Douglas utility specification, which connects the distribution of price sensitivity with the distribution of income via the budget constraint. 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 and the shape of demand is an important ingredient in building a model that delivers robust empirical predictions.

Contributions. Our goal in this paper is to identify the sources for the differences between the left-hand and right-hand side panels of Figure 1. We empirically assess the substantive implications of these differences and highlight when and how to include flexibility in such models. As the discrete choice model is a workhorse in empirical work, adding such flexibility in a tractable way is useful in addressing a wide variety of important questions in academic and policy settings. We make several theoretical and empirical contributions toward achieving this objective.

First, we identify how different components of customer preferences influence the shape of mixed-logit demand. We do so by adopting the "demand manifold" approach of Mrázová and Neary (2017) and focusing on the set of achievable demand elasticity and curvature pairs as sufficient statistics for the shape of demand.[3] We show that the shape of the mixing distribution in *ML* – how the researcher models preference heterogeneity – determines the set of achievable elasticity-curvature pairs and, therefore, the shape of demand. Using a simple single-product monopoly model, we show that curvature of demand is the outcome of a tug-of-war between heterogeneity of consumer preference over product attributes and heterogeneity in price sensitivity.

When consumers have heterogeneous tastes over product attributes, we show that demand curvature and pass-through decrease at all prices relative to the simple *MNL* model. Thus, by incorporating heterogeneity in tastes for product attributes, pass-through remains at most complete. This result builds on Caplin and Nalebuff (1991b), who show that a

[3] 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.

MNL model with heterogeneous valuations of product attributes preserves the curvature properties of the *MNL* model. Intuitively, while consumers have heterogeneous tastes over characteristics, their demand response to a change in price is uniform, which leads a firm with market power to absorb some changes in marginal cost.

Incorporating heterogeneity in price sensitivity increases demand curvature and pass-through. Here, the shape of the price mixing distribution plays a vital role. Idiosyncratic price responsiveness thus accommodates but does not impose more than complete pass-through. We consider three ways of 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 depicting estimated product-level demand in elasticity-curvature space for empirical applications – as in the motivating example above – is useful in visualizing this tug-of-war. It also amounts to a useful presentation of heterogeneous estimated demands by providing insight into possible restrictions imposed by the researcher's preference specification on the shape of demand (e.g., Figure 1, Panels A and C).

Our second contribution is connecting *ML* demand with *CES* demand – the dominant framework in the macro and international trade literature. Here, we demonstrate that discrete choice not only nests *CES* in the case of continuous demand (Anderson, de Palma and Thisse, 1992) but also in the case of the unit-demand discrete choice models commonly used in the Industrial Organization literature. However, where *CES* and *ML* differ is in equilibrium pass-through under oligopoly. Since its curvature exceeds one, *CES* demand implies a pass-through that is always over-shifted, and this pass-through does not vary with the number of competitors. In contrast, we show that competition reduces pass-through relative to the single-product monopoly case in *ML* models for such curvatures. Hence, while the *ML* framework provides a natural and intuitive foundation for *CES* demand and may generate identical estimates of demand elasticity and curvature, the *CES* model predicts larger counterfactual price responses to cost changes.

As a third contribution, we highlight the key modeling choices of the mixing distributions governing preference for price and product characteristics. As theory provides limited guidance on sources of preference heterogeneity, particularly for quasi-linear utility specifications, researchers often impose shape restrictions on these mixing distributions ex-ante. We offer an easy and parsimonious way of modulating how correlates of demand heterogeneity, such as consumer demographics, interact with product characteristics and price to increase the range of feasible elasticity and demand curvature pairs. An appealing feature of our approach is that it nests a *ML* model with standard distributional assumptions and that parameters can be recovered using a generalized method of moments estimator. Identification derives from data moments that trace price responses and consumption patterns across distributions of customer demographics.

We suggest an instrumentation strategy that recovers the nonlinear distribution of price sensitivity by connecting the price responses to cost shifts at different price points to the empirical distribution of demographics through a simple and parsimonious formulation. We demonstrate our identification strategy and the consequences of misspecification via Monte Carlo simulation of two discrete choice demand models: one with income effects and one with quasilinear demand. Consistent with our motivating examples in Figure 1, we show that elasticity and pass-through misspecification bias can be large even when demand elasticities are identical. Moreover, we show that common specifications of price sensitivity are unlikely to correctly estimate the distribution of price sensitivity among customers. Thus, a researcher interested in understanding the distributional consequences of a particular government policy (e.g., , tariffs, subsidies, health insurance) or an exogenous cost shock, will grossly under- or over-estimate the welfare effects if they impose income or log-income ex ante.

Using the data from Berry et al. (1999), we estimate automobile demand under alternative assumptions on the strength of income effects that introduce heterogeneity into price sensitivities. We evaluate a hypothetical $1,000 subsidy for consumers who buy a new US vehicle. The average estimated cost pass-through ranges from 0.99 to 1.79 at the vehicle level, depending on how income effects modulate the price sensitivity distribution. When price sensitivity is inversely proportional to consumer income, subsidy pass-through and hence predicted price declines of domestic cars are most pronounced, likely overstating the true effectiveness of the trade policy. While our empirical exercise explicitly evaluates trade policy, the insight extends naturally to the case of electric vehicle subsidies introduced in the *Inflation Reduction Act of 2022*. This example focuses on the empirical specification of income effects, which is likely important for large purchases of 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, be these connected to observable demographics or unobservable random taste variation. How to best specify these in any particular application is a function of the identifying variation available in the data.

Alternative Approaches. Our focus in this paper is to address the role of distributional assumptions in determining the shape of discrete choice demand. Our work complements Compiani (2022) who also focuses on estimating demand flexibly but uses a non-parametric approach. This attractive solution places even fewer restrictions on the shape of demand than our environment but it easily suffers from a curse of dimensionality. Our work suggests a tractable approach to modeling demand flexibly for a broad range of empirical settings.

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) or fractional units of the same product (Anderson and de Palma, 2020; Birchall and Verboven, 2022).

Finally, our focus on demand specification ignores the effect of supply-side frictions on cost pass-through. We do not consider supply moments as supply-side misspecification may bias estimation results (Conlon and Gortmaker, 2020). Such considerations, for example, the presence of 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). However, 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 via 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 before turning to our empirical application of evaluating a trade policy in automobiles. 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 with a discussion of the demand manifold for a single-product monopolist, as we rely on this setup in Sections 3-4 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 equilibria and comparative statics. Demand sub-convexity weakly limits the feasible elasticity and curvature combinations by ensuring demand becomes more elastic at higher prices.

2.1 Single-Product Monopoly

For simplicity, 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:

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

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

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

letting ρ denote the curvature of inverse demand. While demand can be concave ($\rho < 0$), linear ($\rho = 0$), or convex ($\rho > 0$), concavity of the profit function rules out excessively convex demands.

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:

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

The slope of demand plays a central role in the profit maximization necessary condition (1); in equilibrium, demand must be elastic whenever firms have market power. Economists frequently rewrite the necessary profit maximization condition regarding markups or the Lerner Index.

The sufficient condition for profit maximization further requires that at the equilibrium price, the monopolist's marginal revenue function is non-increasing, which we rewrite in turn in Equation (2) as a constraint on the equilibrium curvature of demand. Cournot (1838) first established the connection between demand curvature and pass-through for a monopolist with constant marginal costs:

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

Hence, 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.

Oligopoly. The monopoly case is useful 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 of discrete choice models affect the relationship between own-elasticity and curvature by plotting demand manifolds corresponding to a single product monopoly case. Obviously in practical applications oligopolists sell multiple products and this graphical representation should be understood as a tool, representing, at best, the relationship between elasticity and curvature over the residual demand of a particular product taking all other substitution estimates as given.

In oligopoly markets, however, the pass-through rate depends not only on demand curvature, but also on substitution between products affected by a common cost shock. We take into account all these substitution effects when plotting the clouds of Figure 1 for each preference specification. To convey the intuition of how substitution affects the pass-through rate of 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 perfectly competitive to $\theta = 1$ for monopoly. The case of $\theta = 1/n$ corresponds to the Cournot solution. Thus, the larger the number of firms, the importance of demand curvature diminishes and pass-through gets closer to complete.

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

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

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

The "monopolist" pass-through of equation (4) that ignores substitution effects could thus be understood as a rough upper-bound estimate of pass-through of any firm in oligopoly. Of course, in practice oligopoly equilibrium is not symmetric and different products are sold at different prices but statements of how the effective pass-through responds to general substitution patterns are difficult to characterize analytically. We evaluate the quantitative relationship between (4) and (5) in the context of our Monte Carlo study in Appendix E.

2.2 Demand sub-convexity

Demand is sub-convex (super-convex) if $\log[q(p)]$ is concave (convex) in $\log(p)$. In the monopoly manifold examples we consider in Sections 3 and4, we focus 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 elasticity price, this cutoff is the curvature of the Constant Elasticity of Substitution (*CES*) inverse demand, $p(p) = \beta p^{-1/\sigma}$. *CES* demand is the only demand system where a single parameter determines both elasticity and curvature: $\varepsilon^{CES} = \sigma$ and $\rho^{CES} = (\sigma + 1)\sigma^{-1} > 1$. 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.[4] More importantly, it is key to the equilibrium existence results for oligopoly models with differentiated products in Caplin and Nalebuff (1991a) for single-product firms and in Nocke and Schutz (2018) for multi-product aggregative games. 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

In this section, we rely on demand manifolds to explore the relationship between curvature and elasticity in the context of the discrete choice demand model, which forms the backbone of much empirical work in IO: mixed logit demand. We characterize the demand manifold, in general, for arbitrary specifications of preference heterogeneity, which we refine in the following sections.

[4] 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 in IO (Miravete, Seim and Thurk, 2018).

We 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 observed and unobserved characteristics of product j, respectively, p_j its price, and y_i consumer i's income. 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 heterogeneity over product characteristics. Lastly, we normalize the value of the outside option to zero.

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

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

Thus, 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"). The purchase decision is the outcome of a Bernoulli random process with a probability of success $\mathbb{P}_{ij}$, which 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 of $sk_{ij} = \sigma^2_{ij}(1 - 2\mathbb{P}_{ij})$. Aggregating over the measure of heterogeneous individuals summarized by $G(i)$, total demand for product j becomes:

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

We can now write the elements defining the demand manifold, elasticity, and curvature of product j, relegating the detailed derivations to Appendix C. 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:

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

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

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

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

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

From (14) we observe that the link between elasticity and curvature is driven by the mixing distributions (right-most terms) in the underlying distribution of customers. While these are best understood as primitives of customer demand, they are most often chosen by the researcher ex-ante. Our objective is to show how these choices impact the connection between not only estimated elasticity (market power) and curvature (pass-through) but also the counterfactual implications of common assumptions.

How the researcher defines the distribution of taste heterogeneity G and the sub-functions $f(\cdot)$ and $h(\cdot)$ plays a fundamental role in determining demand elasticity and curvature. Non-price tastes through $h(\cdot)$ are almost always assumed linear in customer demographics, and non-observed heterogeneity is captured via a standard normal distribution. These choices implicitly restrict $\{\sigma^2_{ij},\ sk_{ij}\}$, restricting the relationship between elasticity and curvature.

A common empirical pricing sub-function is simply the linear function of outside good spending, i.e., $f_{ij}(y_i, p_j) = \alpha_i^\star(y_i - p_j)$, resulting in quasi-linear demand. The curvature is now driven by heterogeneity in the idiosyncratic price sensitivity $\alpha_i^\star$ for a given elasticity. Alternatively, the researcher could impose a non-linear sub-function (Griffith, Nesheim and O'Connell, 2018) with different implications for curvature and pass-through which will become apparent later. Our setup, however, consist of discrete choice unit demands consistent with utility maximization, i.e., where Roy's identity hold for $q_{ij}1$. We illustrate graphically how choices of taste heterogeneity G and the pricing sub-function $f(\cdot)$ impact the shape of the demand manifold using this monopoly model as a simple example. This is two dimensional representation of a single-product monopoly discrete choice model used exclusively to convey the intuition of the relationship between demand elasticity and curvature estimates. Mrázová and Neary (2017, §3) show how demand manifolds could be used in a monopolistic competition oligopoly. After our Monte Carlo evaluation of the model we show in Section E how the insights of this graphical analysis extend to multi-product oligopoly settings including both direct own-price effects and indirect cross-price effects through the dependence of choice probabilities in Equation (10) on the vector of all prices p. Equation (14) should thus be better understood as the manifold of residual demand for product j.

4 Demand Manifolds of Common Demand Specifications

This section begins by considering demand manifolds for discrete choice demand models with quasi-linear preferences, which researchers commonly rely on for inexpensive products like the cereal varieties considered in Nevo (2000), before moving to the case of income effects.

4.1 Quasi-linear Preferences

In the case of quasi-linear preferences, the pricing subfunction f 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 $\beta_i = \beta + \sigma_x \nu_i$. These functional form decisions imply the distribution of price sensitivity has a mean of α with deviations driven by the shape of the mean-zero distribution Φ of ϕ_i, scaled by σ_p. Similarly, β denotes the mean valuation while ν_i captures the idiosyncratic heterogeneity in the valuation of the observed product characteristic, 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:

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

Curvature and elasticity are thus inversely related for any price-quantity pair as long as $sk_{ij} = \sigma_{ij}^2(1 - 2\mathbb{P}_{ij}) > 0$ when the probability $\mathbb{P}_{ij}$ of choosing product j is sufficiently small.

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 $\sigma_p = \sigma_x = 0$ and $\alpha_i^\star = \alpha$ 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. Irrespective of market shares, *MNL* restricts demand to be log-concave and $\rho_j(p) < 1$ for all possible prices. Thus, 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 is atomistic.

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

Figure 2: Multinomial and Mixed Logit Manifolds

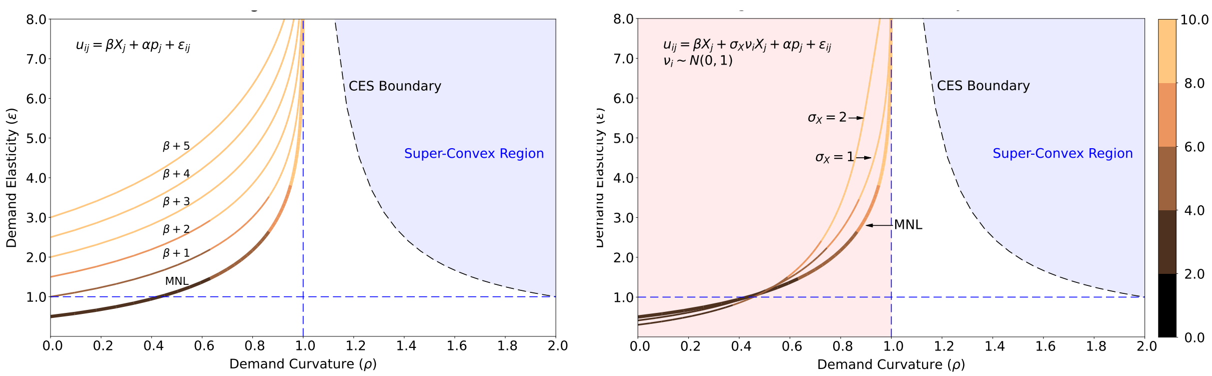

Figure Notes: The left panel shows six alternative *MNL* demand manifolds with one inside good assuming $\alpha = 0.5$, $X = 1$, and $\beta \in \{1, ..., 6\}$. The right panel shows manifolds for a *ML* model with a random coefficient on the product characteristic under alternative standard deviations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

each attribute valuation β, the manifold of the *MNL* is increasing as proven by Mrázová and Neary (2017, Appendix B.5). We set the price response coefficient $\alpha = 0.5$ and consider elasticity-curvature combinations at different price levels. Each manifold is color-coded by level of price, ranging from $p_j = 0$ (darkest) to $p_j = 10$ (lightest). Note that higher prices always correspond to more elastic demands and lower equilibrium markups.

Increasing the average valuation of the product attribute, β, to $\beta + 1$, $\beta + 2, \ldots$, shifts demand manifolds upwards from the base *MNL* manifold in Figure 2. Increasing mean demand for a product thus decreases both curvature and price elasticity for a given price.

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, de Palma and Thisse (1987, 1992) were the first to 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 (Dixit and Stiglitz, 1977). Thus, *CES* arises naturally in the context of discrete-continuous models (Hanemann, 1984), while *MNL* is most appropriate when consumers have unit demand. However, like the *MNL* model, *CES* choice probabilities suffer from the *IIA* property in producing unrealistic substitution patterns. Figure 2 also illustrates that for the same elasticity, the *CES* and *MNL* models imply different demand curvatures (and pass-through). The researcher's choice of one of these two demand models thus restricts pass-through in stark ways, accommodating only over- or under-shifted pass-through, which may not be consistent with the underlying 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 thus consider whether introducing individual heterogeneity in the valuation of the product attribute while continuing to assume that all consumers have the same price responsiveness similarly addresses the limitations of *MNL* in ex-ante restricting curvature.

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 other log-concave distribution of idiosyncratic preferences, comprising the vast majority of distributions used in economics (Bagnoli and Bergstrom, 2005). Further, this result extends naturally to the nested logit – a demand system commonly employed in antitrust economics – because it provides for more reasonable substitution patterns with a small computational burden.[5] Mathematically, Equation 14 demonstrates that curvature can only come through the shape of the choice probability distribution ($\mathbb{P}_{ij}$), particularly the skew.

It is evident that this version of a *ML* model has inherent limitations when used to study pass-through in non-competitive environments empirically: 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

[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.

[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.

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 thus over-shifting of pass-through? The only element of preferences that remains to be considered is idiosyncratic price responsiveness. Substituting $\alpha_i^\star = \alpha + \sigma_p \phi_i$ into the demand manifold for quasi-linear preferences (15) results in:

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

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). Thus, 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_i. 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 can generate. We begin with two price mixing distributions common in empirical work: 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

Figure 3: Normal and Log-Normal Price Mixing Distributions

(a) Standard Normal (b) Log-normal

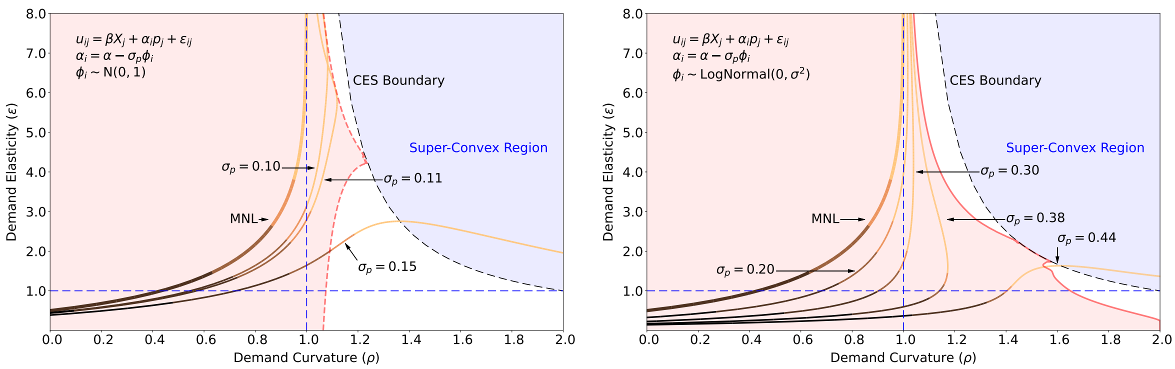

Figure Notes: Panels present demand manifolds in the (ε, ρ) plane under standard normal and log-normal price mixing distribution, respectively. Light-shaded regions represent all feasible (ε, ρ) pairs conditional on the price-mixing distribution. We define "feasible" as manifolds for demand that is sub-convex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

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 prices, elasticity decreases in price, violating Marshall's Second Law. For demand to be sub-convex for all price levels, we thus require less heterogeneity in price sensitivity among consumers, and the feasible elasticity-curvature pairs do not cover the full space up to the *CES* boundary: achieving pass-through greater than one requires relatively elastic demand.

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 larger values of σ_p continue to generate sub-convex demand. This results in a much larger range of feasible curvatures for a given demand elasticity, particularly for less elastic demands where firms enjoy more 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.

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

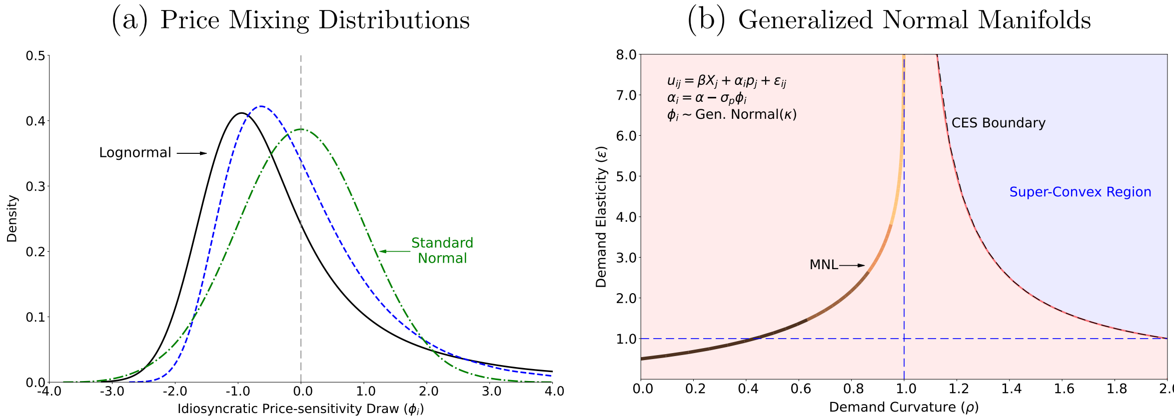

Figure Notes: The left panel shows three specifications of the price random coefficient distribution for different values of shape parameter κ with the Asymmetric Generalized Normal distribution. The right panel shows the combinations of all structural parameters generating well-behaved solutions for (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.

Incorporating a 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 different 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 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. In empirical applications, researchers frequently rely on the fact that idiosyncratic price responsiveness is correlated with demo-

[7] See Appendix for technical details.

graphics. 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.

4.2 Preferences with Income Effects

In contrast to the quasi-linear case where outside good spending enters consumers' indirect utility linearly, *BLP* specifies the preferences in Equation (8) with 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, which is consistent with utility maximization in discrete choice contexts for any value of the power parameter $\lambda \in \mathbb{R}$ driving the convexity or concavity of the transformation. We, therefore, specify the generalized price sub-function:[8]

$$f_i(y_i, p_j) = \alpha_i^\star \left(y_i - p_j\right)^{(\lambda)} = \begin{cases} \alpha_i^\star \dfrac{\left(y_i - p_j\right)^\lambda - 1}{\lambda}, & \text{if } \lambda \neq 0, \\[2ex] \alpha_i^\star \ln\left(y_i - p_j\right), & \text{if } \lambda = 0, \end{cases} \tag{19}$$

and 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). 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 $\lambda = 0$ thus yields the *BLP* model, while a power parameter of $\lambda = 1$ results in a *MNL* model. This

[8] Using a multi-unit demand model, Birchall and Verboven (2022) rely on a price sub-function with a different Box-Cox transformation, $f(y_i, p_j) = \gamma^{\lambda-1}\left(y_i^\lambda - 1\right)\lambda - \left(p_j^\lambda - 1\right)\lambda$ that depends on the share of income, γ, spent on a chosen product and is 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 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 instead to allow for curvature flexibility through an unconstrained Box-Cox power parameter that modulates income effects within the confines of a unit-demand setup consistent with utility maximization (e.g., Roy's identity holds for $q_{ij} = 1$).

means that the income distribution captures any idiosyncratic price responsiveness across individuals, modulated byλ. As in 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 presents 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 the motivating elasticity-curvature pairs in Figure 1 for Berry et al. (1999) are driven by the demand specification.

4.3 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 attributes and (b) flexible distributions of price sensitivity. The former basically expands the set of estimable elasticity-curvature combinations within the logconcave region of demand while the latter

[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 C.9 for the case of the Maclauring approximation of the Box-Cox price sub-function (19).

[11] While we consider a power parameter $\lambda \in [0, 1]$, in line with the empirical literature, Box and Cox (1964) consider $\lambda \in [-5, 5]$, which would expand the range of feasible curvature elasticity pairs beyond the ones depicted in Figure 5. Note that using the convenient Maclauring approximation leads to some estimation bias and thus *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 $\ln(y_i - p_j)$.

Figure 5: Demand Manifolds & Income Effects

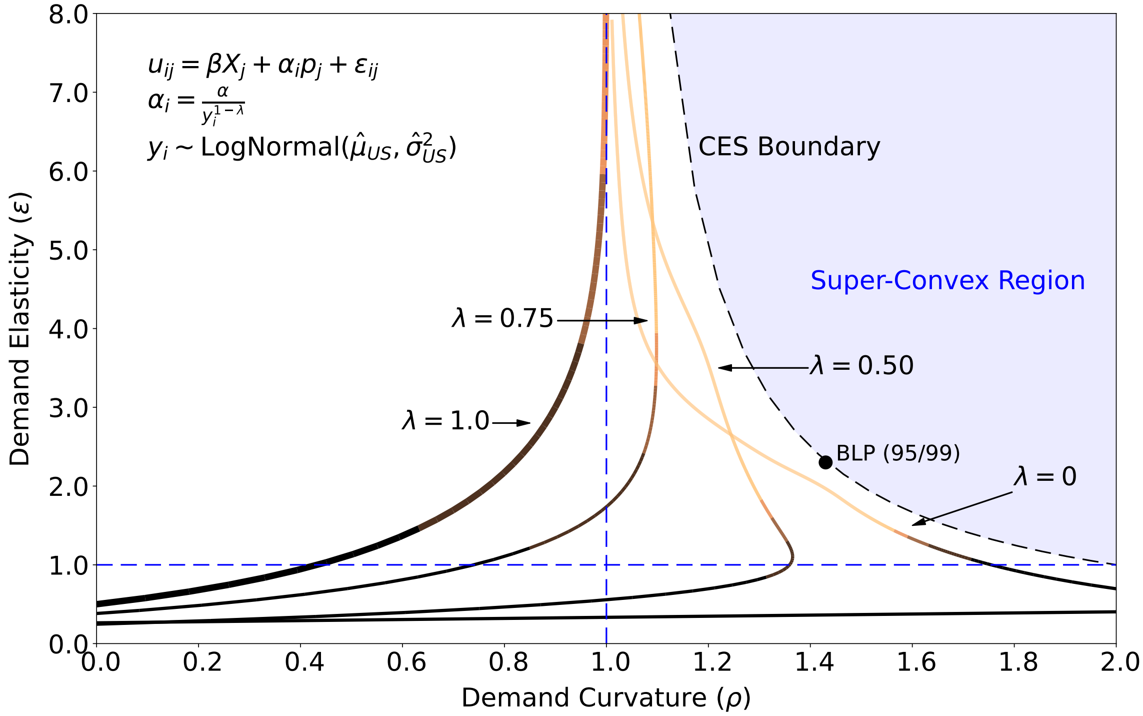

Figure Notes: Demand manifolds for different values of the Box-Cox transform parameter λ using the U.S. income distribution and the rest of the model specification of Berry et al. (1999). The dot identified as "BLP (95/99)" corresponds to the average estimated curvature and elasticity value using the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

extends this set to the logconvex region. 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 entails over-shifted pass-through because 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.

To make this intuition concrete, we present a simple example of pricing by a monopolist who caters to two consumers with linear demands of different slopes in Figure 6, 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 perfect price discrimination.[12] In Panel B, we show that in setting a uniform price, the monopolist now faces a kinked

[12] A subtle but important point is that the shape of demand depends on the markets firms choose to compete in because such decisions imply consumer preference heterogeneity.

Figure 6: Heterogeneous Customers and the Shape of Demand

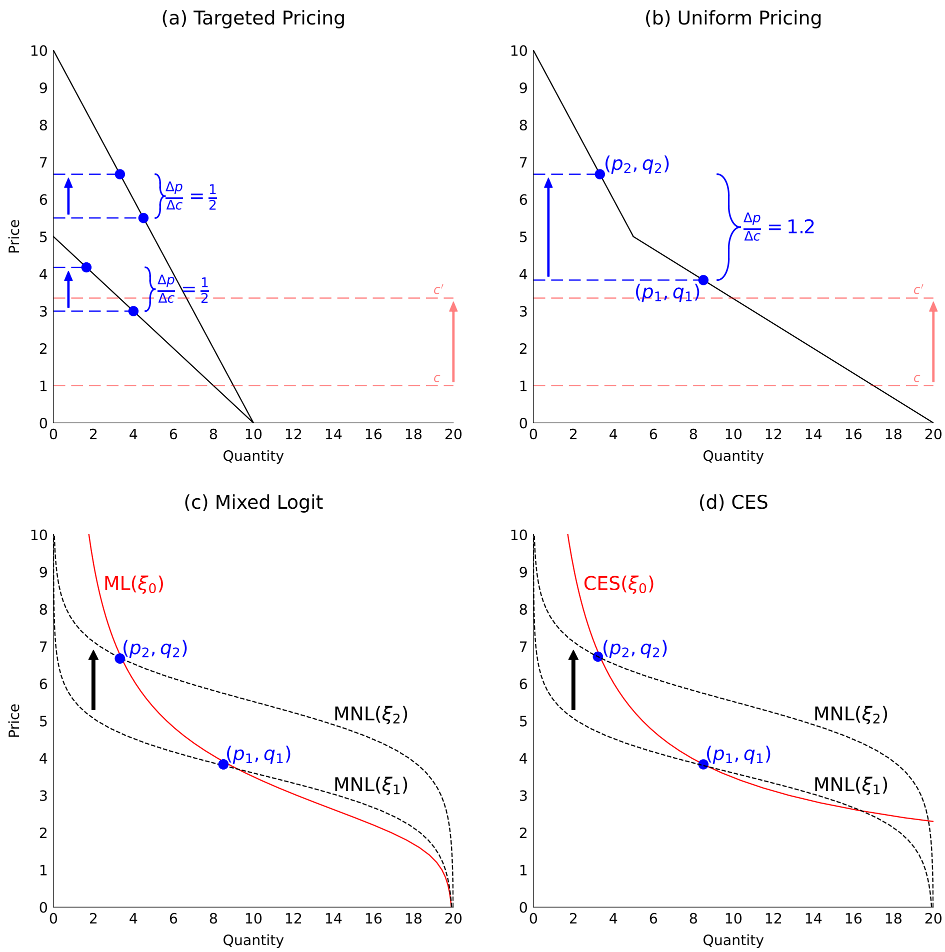

Figure Notes: Top panels present pass-through for linear demand under targeted and uniform pricing given data points (p_1, q_1) and (p_2, q_2). Bottom panels present discrete choice and *CES* rationalizations of the same data points.

demand.[13] 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 to a level that serves only the price-insensitive customer. Pass-through is 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.

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

[13] Kimball (1995) first suggests a smooth, differentiable version of this kinked demand to ensure subconvexity and markups that increase in the production scale.

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. Thus, while *CES* is a useful simplification of *ML* for estimation, its pass-through predictions in oligopoly settings are restrictive.

5 Guidance for Empirical Work

Our work thus far addresses how to model the shape of demand via non-price characteristics and price sensitivity flexibly to limit the impact of the utility specification on pass-through. Now, we turn to how to identify and estimate the shape of demand in an empirical setting. We focus on a model where demand has income effects as in Berry et al. (1999) and provide an extension to quasilinear demand (e.g., Nevo, 2001) in the appendix. Both models generate similar results.

We detail our identification strategy in Section 5.1. In Section 5.2, we use Monte Carlo evidence to demonstrate that our identification strategy is effective and illustrate the consequences of misspecification. In Section E, we provide evidence on how competition affects pass-through relative to the earlier single-product monopolist. We then address the empirical implications of imposing different price mixing distribution in Section 5.3 by simulating a counterfactual trade policy using alternative estimated new automobile demands.

5.1 Instruments to Identify Demand Manifolds

Here, we consider taking the preference specification in Equation (8) to data when allowing for non-linear taste heterogeneity. As an example, we focus on consumers' sensitivities to price; similar considerations arise for taste heterogeneity for product characteristics. Empirically, one could represent heterogeneous price sensitivities with a flexible function of consumers' demographics, a flexibly distributed price random coefficient, or nonlinear income effects through e.g., a Box-Cox power transform, as in Section 4.2. How can the data pin down such price mixing distributions?

We focus on the case of preference heterogeneity tied to observable consumer demographics – income –, which empirically frequently correlate well with purchase behavior. We model the distribution of heterogeneous price sensitivities using the Box-Cox power transform of income (Box and Cox, 1964), hence considering nonlinear income effects. While only one possible way of representing consumer heterogeneity, the Box-Cox transform is a simple, one-parameter mechanism to transform data into mixing distributions consistent with identifying data moments.

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 under the assumption that product characteristics are chosen before demand unobservables, ξ, are realized. A comparison between instances where there are 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:[14]

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

We run the above regression and use the results to construct the vector of predicted (exogenous) 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 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, 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 thus trace the demand manifolds using cost shocks, holding constant exogenous demand shifters at different price levels. In Section 5.2, we explore the instrument's performance in Monte Carlo simulations. Lastly, extending the argument to quasi-linear demand is straightforward and we cover identification in the Appendix, Section A.

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

5.2 Flexible Manifold Estimation: Monte Carlo Analysis

We conduct a Monte Carlo analysis to demonstrate the validity of our identification strategy and evaluate the potential for misspecified 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 prices 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, as above, 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.[15] Consumers, therefore, have preference heterogeneity over the J inside goods as a category via the constant random coefficient and over variation in the observable product characteristic across the J products and T time periods. 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 Equation 20's approximation to the Box-Cox transformation 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 $p^\star_t$, satisfy the set of J first-order conditions for the firms:

[15] 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.

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

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

$$\log c_{jt} = \gamma_0 + \gamma_1 \log x^1_{jt} + \gamma_2 \log x^2_{jt} + \omega_{jt} + \zeta_{jt} \tag{27}$$

We set all γ parameters equal to 1 and draw cost shocks $\{\omega_t, \zeta_t\}$ from standard normal distributions. The researcher observes ω_t, which identifies the price sensitivity distribution. We generate pricing equilibria 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 the objective of a researcher who estimates consumer demand given observed prices, quantities, and ω cost shocks following the best practices outlined in Conlon and Gortmaker (2020). The researcher also specifies the supply side as in Berry et al. (1999) and correctly specifies the outside option and the distribution generating the random coefficients for product characteristics ν_i. The researcher, however, *may incorrectly model* income effects and, hence, the distribution of price-sensitivities. This Monte Carlo analysis aims to investigate the success of an empirical demand model with a flexible Box-Cox power transformation of outside good spending at identifying and recovering the true demand curvature underlying the data-generating process, relative to simpler alternatives.

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'.

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.

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 ($\mathrm{corr}(\varepsilon_j, \rho_j)$).

Discussion of Results. We present the parameter estimates in Table 1 for seven distinct 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), the researcher models 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. The researcher, therefore, would overestimate both market power and pass-through. Moreover, specifying log preferences ex-ante 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 the researcher 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 Manifold Estimation: Empirical Implications

In this section we demonstrate the sizable quantitative and, perhaps, qualitative implications 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

Figure 7: Income Effects and Demand Manifolds

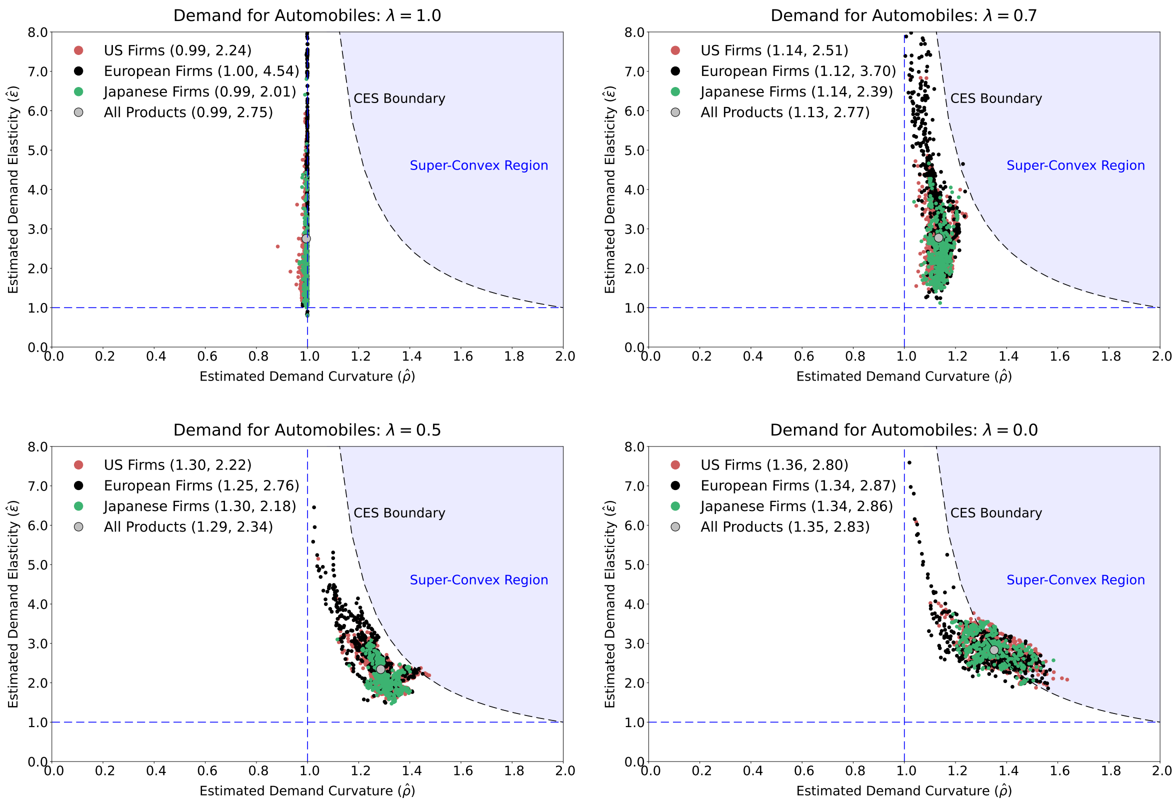

Figure Notes: Each dot represents the point elasticity and curvature estimates for each observation in the *BLP* automobile data, while the gray dot corresponds to the average elasticity and curvature estimates. Point estimates are colored according to vehicle origin and we see significant overlap in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

models, we incorporate income draws as in Andrews et al. (2017) and follow best practices outlined in Conlon and Gortmaker (2020).

Figure 7 shows the scatter plots of $(\hat{\varepsilon}, \hat{\rho})$ for each automobile model in the *BLP* data under these four alternative specifications.[16] 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 7. Reducing λ increases the importance of income effects through smaller price responses by higher-income households. Moving from quasi-linear demand to demand with income effects does not significantly

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

Figure 8: Demand Manifolds and Trade Policy

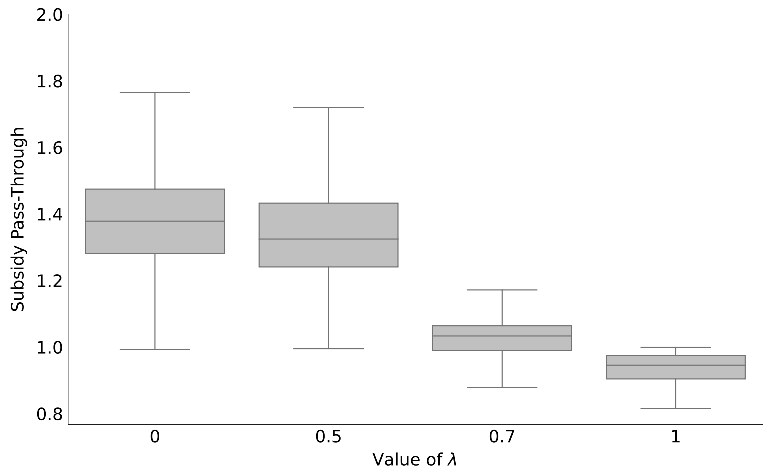

Figure Notes: Figure present pass-through of US vehicles for a hypothetical \$1,000 vehicle subsidy paid to US consumers for the purchase of domestic vehicle.

change the average estimated elasticity, reaching $\hat{\varepsilon}_{\text{BLP}} = 2.83$ 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 motivating RTE cereal case.

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 thus 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 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 8.

As expected, given the Monte Carlo results, the shape of the mixing distribution of first-order quantitative importance in evaluating this policy. Under the original *BLP* specification, median pass-through is 1.39. Hence, the subsidy elicits domestic automakers to reduce their vehicle prices by more than the subsidy to capture price-sensitive customers. As the mixing distribution becomes less skewed, firms pass through the subsidy less until it becomes under-shifted in the case of $\lambda = 1$. We also observe a compression of the distribution of subsidy pass-through across vehicles. For a policy maker choosing the generosity of 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 income effects by including the income distribution in the price mixing distribution will overestimate the effectiveness of subsidization.

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. This is useful as it both aids in the identification of the mixed-logit model and demystifies the mixed-logit model by enabling the researcher to articulate the path from data to model to empirical results. In particular, 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 we think there are a variety of empirical settings where our work will be valuable. First, cost pass-through in the international trade and macroeconomic literature is pinned-down ex ante via *CES* demand. Our results indicate this assumption leads to over-estimates of 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*.[17] Our results indicate that robust estimates of how effective this policy has been at generating incremental EV purchases requires modelling the distribution of customer preferences flexibly.

Finally, we know from Aguirre, Cowan and Vickers (2010) that estimating the welfare effects of uniform pricing requires estimates of both demand curvature and demand elasticity. Recent work by Adams and Williams (2019), DellaVigna and Gentzkow (2019), and Hitsch, Hortaçsu and Lin (2021) have demonstrated that such pricing policies are common in industry. Meanwhile, uniform pricing is often mandated by government in health care markets. For example, 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.[18,19] Our work demonstrates the importance of the preferences specification in recovering robust estimates of the distributional consequences of these policies. Adding supply-side considerations to our framework – such as menu costs or allowing marginal cost to be correlated with demand, as in the case of healthcare – is another exciting area for further research.

[17] Specifically, consumers are given a subsidy between $2,500 and $7,500 when they purchased a electric vehicle that meets certain domestic sourcing requirements.

[18] Within each rating area, insurers can only charge different prices based on age, the number of people on the plan, and the health plan tier. Geddes (2022) demonstrates that while insurers are constrained to offer plans based on actuarial values, insurers mitigate these constraints by selectively entering rating areas. Moreover, these entry decisions correlate with the people's characteristics in those areas.

[19] Recent studies of insurance demand under the ACA, such as Saltzman (2019) and Tebaldi (2022), incorporate customer heterogeneity in the underlying demand model, but both use shape restrictions that are common in the literature.

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.

Adams, Brian and Kevin Williams (2019): "Zone Pricing in Retail Oligopoly," *American Economic Journal: Microeconomics*, Vol. 11, pp. 124–56.

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

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 (1987): "The CES Is a Discrete Choice Model?" *Economics Letters*, Vol. 24, pp. 139–140.

——— (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 and Frank Verboven (2022): "Estimating Substitution Patterns and Demand Curvature in Discrete-Choice Models of Product Differentiation," Discussion Paper 16981, CEPR.

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.

Dixit, Avinash and Joseph E. Stiglitz (1977): "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, Vol. 67, pp. 297–308.

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.

Geddes, Eilidh (2022): "The Effects of Price Regulation in Markets with Strategic Entry: Evidence from Health Insurance Markets," working paper, Northwestern University.

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.

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.

Hitsch, Günter, Ali Hortaçsu, and Xiliang Lin (2021): "Prices and Promotions in U.S. Retail Markets," *Quantitative Marketing and Economics*, Vol. 19, 12.

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 (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.

MacGillivray, Helen L. (1986): "Skewness and Asymmetry: Measures and Orderings," *The Annals of Statistics*, Vol. 14, pp. 994–1011.

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.

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.

Saltzman, Evan (2019): "Demand for health insurance: Evidence from the California and Washington ACA exchanges," *Journal of Health Economics*, Vol. 63, pp. 197–222.

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

Tebaldi, Pietro (2022): "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA," Working Paper 29869, National Bureau of Economic Research.

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.

Appendix

A Flexibly Estimating Quasi-linear Demand

The main text demonstrated the importance of modeling the distribution of heterogenous preferences flexibly using preferences which generated income effects. In this appendix we demonstrate the same ideas – including identification and monte carlo results – extend to customer preferences which generate quasilinear demand. Quasi-linearity is a popular empirical choice of researchers even in settings with high-dollar goods where income effects are likely important (e.g., Grieco, Murry and Yurukoglu, 2021). This is because quasi-linearity provides the researcher added flexibility; i.e., it does not impose how they include income in the specification, nor does it require accounting for income effects in counterfactual welfare calculations. Quasi-linear preferences also imply simpler curvature derivations since $f''_{ij} = 0$.

As should be apparent by now, the choice of price mixing distribution nonetheless may materially impact 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):

$$u_{ij} = x_j \beta_i^\star + \alpha_i^\star \big(y_i - p_j\big) + \xi_j + \epsilon_{ij}\,, \quad i \in \mathcal{I},\ j \in \mathcal{J},\ \epsilon_{ij} \sim \mathrm{EV1}\,, \tag{A.1}$$

where we include $\alpha_i^\star$ to capture consumers' heterogeneous price sensitivity.

A researcher could allow demographics to 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. It is important to note, however, that in a quasi-linear model, 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 $\alpha_i^\star$. The monte carlo analysis of Section 5.2 demonstrates, however, that leveraging the Box-Cox transformation provides greater flexibility with minimal computational burden. One might allow price sensitivity to differing 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. We therefore model price-sensitivity as follows:

$$\alpha_i = -\exp\left(\alpha + \pi y_i^{(\lambda)}\right), \text{ where } y_i^{(\lambda)} \equiv \begin{cases} \dfrac{y_i^{\lambda} - 1}{\lambda}, & \text{if } \lambda > 0, \\ \ln\left(y_i\right), & \text{if } \lambda = 0 \end{cases} \tag{A.2}$$

This specification of α_i has the nice feature of guaranteeing that all customers have downward sloping demand.

A nice feature of the 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 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.

How exactly does the transform influence customer behavior in the model, and what are the implications for the distribution of α_i? Figure A.1 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 regulates 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

Figure A.1: Visualizing the Box-Cox Transform

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

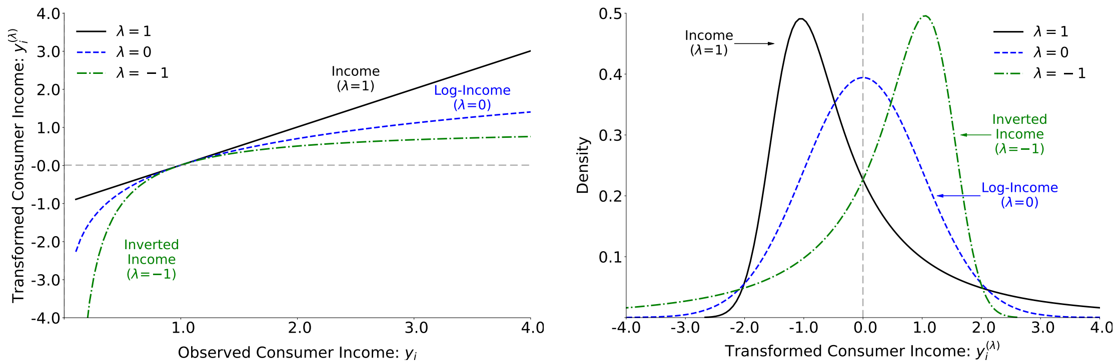

are uniform. Hence, the predicted average price of the chosen product changes uniformly across income groups, all-else-equal. When $\lambda = 0$, we observe that small differences in income will yield very different consumption sensitivities to price. We would, therefore, observe in the data that the 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 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 as the Monte Carlo simulation as described in Section 5.2 but now applied to (indirect) utility specification (A.1). We specify $\pi = -0.2$ so high-income agents 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. In the results section we discuss what happens when we choose, for example, $\lambda = 2$.

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 heterogenous 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.

Table A.1: Quaslinear 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.

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) this these indicator variables. The identifying assumption then follows the logic of Figure A.1; 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 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 A.1. 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. Bad things happen, however, when the researcher imposes the relationship between income and price-sensitivity but the true data-generating process is different.

The implications of imposing the relationship between income and price-sensitivity is demonstrated in Table A.2, particularly scenarios 6 and 7. The first two columns demonstrate that imposing log-income or income as the price interaction leads to estimated distributions of price sensitivity (summarized by their respective coefficients of variation) which are very

Table A.2: 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 ($\text{corr}(\varepsilon_j, \rho_j)$).

different from the true data. This indicates that a researcher interested in understanding distributional consequences of a government policy, cost shock, etc. will grossly underestimate the welfare effects if they impose income or log-income ex ante. We observe that in both 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 DGP is the flexible Box-Cox specification with $\lambda = -1$, they under-estimate demand curvature. We find the mis-specification bias when the researcher assume log-income on estimated demand curvature is small, however. This stems from the fact that our experiments assumed $\lambda = -1$. When we switch and set $\lambda = 2$, we get the opposite result where assuming log-income distorts the curvature estimate more than assuming income. 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 using commonly-used instruments.

Online Appendix

B 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):

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

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

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 B.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.

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

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

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

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

C 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:

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

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

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

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

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

from where we obtain standardized moment or *skewness* (MacGillivray, 1986) as:

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

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:

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

The derivative of the variance with respect to price is:

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

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

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

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

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

Similarly, the inverse demand curvature of product j is:

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

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

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

D 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):

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

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:

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

so that:

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

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

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

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

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

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

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

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

Figure E.1: Competition and Pass-Through Rates

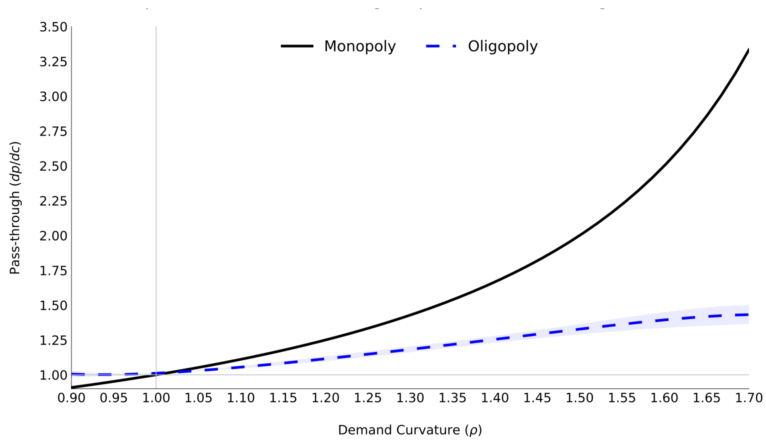

Figure Notes: Figure presents Monte Carlo results across equilibria of median demand curvature. We generate each equilibrium following the environment discussed in Section 5.2 for the Box-Cox utility specification where $\lambda \in [0, 1]$. For each market t in each equilibrium, we solve for the median (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.

other products similarly affected by the cost increase. The net effect thus depends on "how far" a particular product is from its closest substitutes in product space.

Many empirical questions depend critically on the relative importance of the "direct effect" and "net effect." The literature has thus far 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.