PASS-THROUGH WITH PRICE DISPERSION

BRIAN C. ALBRECHT AND MARK WHITMEYER

Abstract. How do cost shocks pass through to prices in markets with price dispersion? Pass-through analysis typically assumes a single equilibrium price, but empirical studies consistently document substantial price variation, even for homogeneous products. This paper develops a tractable framework that decomposes the pass-through problem into two distinct tiers. The first is a competition layer where consumers' *consideration sets* determine equilibrium distributions of normalized margins. The second is a curvature layer where demand elasticity determines how these margins translate into prices and pass-through rates. The key theoretical innovation is showing that the strategic pricing game with arbitrary downward-sloping demand is order-isomorphic to a baseline unit-demand game once reformulated in terms of normalized effective margins. This decomposition yields closed-form pass-through formulas, robust bounds across demand specifications, and clear comparative statics linking market structure to incidence.

1. Introduction

Empirical studies consistently document substantial price variation for homogeneous products. The theoretical toolkit for analyzing pass-through presumes a unique price to differentiate. This paper develops a tractable framework for pass-through analysis when equilibrium features a price distribution rather than a price point.

The core contribution is a decomposition. We show that pass-through in dispersed-price markets separates into two tiers. The first is a *competition layer*, where market structure determines an equilibrium distribution of normalized effective margins. The second is a *curvature layer*, where demand elasticity determines how these

BA: International Center for Law & Economics. Email: mail@briancalbrecht.com. MW: Arizona State University. Email: mark.whitmeyer@gmail.com.

margins translate into prices and pass-through rates. Our key result, Theorem 3.3, establishes that **the equilibrium margin distribution depends only on market structure, not on demand curvature or cost levels**. Demand and costs matter only for the mechanical translation from margins to prices.

This decomposition has three immediate payoffs for applied work. First, it yields closed-form pass-through formulas at each quantile of the price distribution, computable without resolving the full equilibrium for each cost level. Second, the invariance of margin distributions to costs means pass-through follows from differentiating a function, rather than from comparative statics on a complex game. Third, the separation identifies exactly when demand specification matters. The competition layer is demand-free, while the curvature layer requires knowing the demand function.

The paper contributes to two literatures that have developed largely in parallel. The literature on pass-through, following Weyl and Fabinger (2013), derives simple formulas linking pass-through to demand curvature but requires a single equilibrium price, either under monopoly or symmetric oligopoly. We show how their curvature insights extend to equilibrium price distributions. The same demand primitives matter, but they interact with a distribution of markups rather than a single markup. The modern industrial organization literature on price dispersion, building on Varian (1980) and Burdett and Judd (1983), characterizes when and why prices are dispersed, but has not addressed how cost shocks transmit through these distributions. We provide a general framework of pass-through characterization for dispersed-price equilibria.

We adopt the consideration-set framework of Armstrong and Vickers (2022) as our model of price dispersion. Consumers observe prices only from firms in their consideration set and purchase from the cheapest option. This generates mixed-strategy equilibria where firms randomize over prices. The consideration structure (who considers whom) is reduced-form and captures many phenomena–search costs, informational frictions, platform algorithms, behavioral inattention, etc. Our results apply regardless of the underlying source.

The isomorphism works by reformulating the pricing game in terms of normalized effective margins rather than prices. The effective margin captures profit per customer served, normalized to lie in the unit interval. Under a natural monotonicity condition, there is a one-to-one mapping between prices and margins. The key observation is that equilibrium behavior depends on the trade-off between margin and market share, and this trade-off is determined entirely by the consideration structure. Demand curvature and costs affect only the translation from margins back to prices. This is why the equilibrium margin distribution is invariant. When costs change, the entire price distribution shifts, but the underlying distribution of competitive positions remains fixed.

To illustrate the decomposition's power, consider how it simplifies comparative statics. A standard question asks how pass-through varies with market structure. In our framework, this question separates cleanly. Changes in market structure shift the margin distribution; we characterize exactly how in terms of the consideration structure. Changes in demand curvature alter the translation from margins to prices; we provide closed-form expressions for common demand families. These effects combine in a tractable way, enabling sharp predictions about which structural changes matter most for incidence.

The framework yields several results that would be difficult to obtain without the decomposition. We derive pass-through formulas at each quantile of the price distribution, revealing heterogeneity invisible to single-price analysis. Consumers buying at low prices face different pass-through than those buying at high prices. Under unit demand, transaction-weighted pass-through (what matters for welfare) reduces to a sufficient statistic depending only on the consideration-set structure. One can then estimate incidence from consideration data (platform clicks, surveys, geographic proximity) rather than from a demand system, building on methods surveyed in Honka et al. (2019). The identification problem shifts from "how does quantity respond to price?" to "who considers whom?," which is often more tractable. We derive bounds on pass-through that hold across all demand specifications, useful for policy analysis when demand is only partially known. And we clarify when aggregate pass-through can exceed unity, showing that the same market structure yields qualitatively different incidence patterns depending on demand curvature.

We develop these results for arbitrary consideration structures before specializing to cases that admit closed-form solutions. With symmetric firms, equilibrium has a simple characterization in terms of the probability generating function for competitor counts. With asymmetric firms under independent consideration, we derive piecewise closed-form equilibria exhibiting a hierarchical structure. Higher-reach firms price higher at each quantile and, if demand is not too convex, have lower pass-through at each quantile. These closed-form results make the framework applicable to structural estimation and policy counterfactuals.

1.1. **Related Literature.** This paper connects two distinct literatures: the industrial organization literature on price dispersion and limited consideration and the literature on tax incidence and pass-through.

Price dispersion in homogeneous goods markets traces to Varian (1980) and Burdett and Judd (1983), who show that when consumers observe different subsets of prices, firms randomize in equilibrium. Recent work extends these foundations in several directions: Guthmann (2024, 2025); Guthmann and Albrecht (2025) analyze dynamic settings, Elliott et al. (2021) study platform design, Bergemann et al. (2021) and Albrecht (2020) derive robust equilibrium bounds that hold across information structures, while Armstrong (2017), Rhodes and Zhou (2019), and Rhodes et al. (2021) explore richer consideration structures including ordered search and multiproduct settings.

Most closely connected is Armstrong and Vickers (2022), who develop a general framework for competition with arbitrary consideration patterns. We build directly on their framework, adopting their consideration set and equilibrium concepts. But where they focus on welfare and market structure with unit demand, we characterize pass-through with general downward-sloping demand. Our main theorem, Theorem 3.3, shows that their unit-demand results extend to general demand functions once we work with normalized effective margins. More broadly, we extend this literature in two directions: we derive pass-through formulas that apply to any consideration structure, and we show that equilibrium margin distributions depend only on consideration patterns, not on demand curvature, via a separation principle that simplifies comparative statics dramatically.

A growing empirical literature estimates consideration sets and documents their importance for demand analysis. Different subfields use different terminology: ”awareness sets” when the friction is ignorance (Honka et al., 2017), ”choice sets” when it is institutional (Gaynor et al., 2016), and ”attention” when it is cognitive (Abaluck and Adams-Prassl, 2021). The core insight is the same: ignoring consideration sets biases elasticity estimates and distorts counterfactual predictions. Goeree (2008) shows that full-information models underestimate PC markups by a factor of four. The gap arises because high-share firms have larger consideration sets, not just better products. Abaluck and Adams-Prassl (2021) demonstrate that consideration-set frictions explain why Medicare defaults are sticky–a pattern full-information models cannot match.

The pass-through literature, following Weyl and Fabinger (2013), provides a unified framework showing that pass-through depends on the curvature of demand relative to its slope. Their approach has been extended to vertical markets (Adachi and Ebina, 2014), welfare analysis (Adachi and Fabinger, 2022), and platform fees (Weyl, 2010). Empirical work documents heterogeneity across markets (Marion and Muehlegger, 2011; Stolper, 2017; Miller et al., 2017). We extend the Weyl-Fabinger framework to markets with equilibrium price *dispersion*. The same demand-curvature fundamentals matter, but they interact with a distribution of markups rather than a single markup. Our separation into competition and curvature layers parallels their decomposition into conduct and curvature, with consideration structure playing the role of conduct, but microfounded through consumer information and search rather than through a reduced-form parameter.

Three recent papers have begun bridging these literatures by analyzing pass-through when prices are dispersed. Garrod et al. (2024) provide a theoretical analysis of the Varian model with general demand, showing that whether captive or non-captive consumers bear more of a cost increase depends on whether demand is log-concave or log-convex. Montag et al. (2023) develop a Varian-style model with informed and uninformed consumers and test it using tax changes in German and French fuel markets. They find that pass-through is higher for the minimum price paid by informed consumers than for the average price paid by uninformed consumers.

Fischer et al. (2024) estimate a structural search model using German fuel data, finding that informed consumers face higher effective pass-through rates and that excise tax reductions would benefit consumers more than equivalent VAT cuts. The first two papers work with binary consumer types–captive versus non-captive, or informed versus uninformed–while Fischer et al. (2024) estimate a distribution of search intensity. We build on their insights using the Armstrong-Vickers framework, which allows arbitrary consideration patterns, and we characterize pass-through at each quantile of the price distribution rather than by consumer type.

Beyond specific applications, our approach to transaction-weighting in mixed-strategy equilibria provides a general framework for computing consumer-relevant statistics when firms use mixed strategies. This paper also connects to the literature on order statistics and quantile methods in economics (Koenker, 2005; Chernozhukov et al., 2013). By characterizing pass-through at each quantile of the price distribution, we provide a complete picture of how costs transmit through the market, which is information that single-price models cannot capture.

2. Model

There are n firms, indexed by $N = \{1, \dots, n\}$. The demand side consists of a unit mass of consumers, partitioned by their consideration sets.[1]

Definition 1. A *consideration structure* is a probability distribution $\{\alpha_S\}_{S\subseteq N}$ where $\alpha_S \geq 0$ represents the mass of consumers who consider exactly the set S of firms, with $\sum_{S\subseteq N} \alpha_S = 1$.[2]

Each consumer observes prices only from firms in her consideration set and purchases from the lowest-priced firm, provided that price does not exceed 1.[3]

Example 2.1 (Random Search)**.** If each consumer samples each firm independently with probability $\lambda \in (0,1)$, then $\alpha_S = \lambda^{|S|}(1-\lambda)^{n-|S|}$. This yields the binomial consideration structure.

Example 2.2 (Spatial Markets)**.** Consider firms located on a circle. If consumers observe only their k nearest neighbors:

$$\alpha_S = \begin{cases} 1/n & \text{if } S \text{ consists of } k \text{ consecutive firms} \\ 0 & \text{otherwise} \end{cases}$$

For analytical convenience, we define:

Definition 2. Firm i's *reach* is $\sigma_i \equiv \sum_{S\ni i} \alpha_S$, the mass of consumers who consider i. Firm i's *captive share* is $\alpha_{\{i\}}$, the mass who consider only i. Firm i's *captive-to-reach ratio* is $\rho_i \equiv \alpha_{\{i\}}/\sigma_i \in [0,1]$.

Each consumer who purchases at price p demands quantity $x(p)$ where:

[1] Other studies, such as Perla et al. (2023), Guthmann (2024), McAfee (1994), Albrecht (2020), Guthmann and Albrecht (2025), and Armstrong and Vickers (2022), use terms such as "awareness," "availability rate," "choice set," "loyal customers," and "consideration set" to indicate the subset of firms that buyers have access to.

[2] We allow $\alpha_\emptyset \geq 0$, representing consumers who consider no firms.

[3] We interpret 1 as an upper bound on feasible prices rather than a choke price–consumers have positive demand even at $p = 1$ (see Assumption 2.3).

Assumption 2.3. The function $x\colon [0,1] \to \mathbb{R}_+$ is continuous, weakly decreasing, and continuously differentiable, with $x(1) > 0$.

Firms share a common marginal cost $c \in [0,1)$. The cost may represent production costs, taxes, or input prices whose changes we wish to study. Firms simultaneously choose price distributions. We break ties uniformly: whenever multiple firms offer the same lowest price within a consumer's consideration set, we assume that the consumer randomizes uniformly across the tied lowest-priced firms. Formally, for any nonempty consideration set $S \subseteq N$ and realized price profile $\mathbf{p} \in [c,1]^N$, define the set of minimizers

$$M_S(\mathbf{p}) := \arg\min_{j \in S} p_j.$$

A consumer with consideration set S purchases from firm $i \in S$ with probability

$$\frac{\mathbf{1}\{i \in M_S(\mathbf{p})\}}{|M_S(\mathbf{p})|}.$$

Consequently, given rival mixed strategies F_{-i}, the *demand* faced by firm i when it posts price p is, therefore,

$$q_i(p) := \sum_{S \ni i} \alpha_S \cdot \mathbb{E}\left[\frac{\mathbf{1}\{i \in M_S(p, p_{-i})\}}{|M_S(p, p_{-i})|}\right]. \quad \text{(Demand)}$$

If the mixed-strategy cumulative distribution functions (CDFs) are atomless on the interior of their support (which is typical in price competition with continuous payoffs on each side of any price), so ties occur with probability zero, the demand formula simplifies to

$$q_i(p) = \sum_{S \ni i} \alpha_S \prod_{j \in S \setminus \{i\}} (1 - F_j(p)), \quad \text{(Demand*)}$$

where we follow the convention that the empty product (when $S = \{i\}$) equals 1. (Demand*) is the mass of consumers who consider i and find i strictly cheaper than all rivals in their consideration set.

Now let

$$\Pi_i(p;c) = (p-c)x(p)q_i(p) \tag{Profit}$$

denote firm i's profit from posting price p. A profile of price CDFs $(F_i)_{i\in N}$ on $[c,1]$ constitutes an equilibrium if for each firm i two conditions hold:

(1) Each firm is indifferent over all prices on the support of its distributions: for all prices p in the support of F_i, $\Pi_i(p;c) = \pi_i(c)$.
(2) No firm has a profitable deviation: for all prices $p \in [c,1]$ outside the support of F_i, $\Pi_i(p;c) \leq \pi_i(c)$.

We also make the following monotonicity assumption on the effective margin per served consumer:

Assumption 2.4. The *effective margin per served consumer*, function $(p-c)x(p)$, is strictly increasing on $[c,1]$ for all $p \in [c,1]$.

This assumption ensures that higher prices correspond to higher effective margins, maintaining the trade-off between margin and market share. The assumption holds whenever the effective margin $(p-c)x(p)$ is maximized at the boundary $p=1$ rather than at an interior price. To ensure invertibility, we will always need some assumption like this. With unit demand, which is common in the search and price dispersion literature, the assumption holds automatically since $(p-c)\cdot 1$ is linear in p. When the effective margin is instead maximized at an interior price $\hat{p}(c) < 1$, the assumption fails on $[c,1]$. For example, with CES demand $x(p) = p^{-\eta}$, the effective margin is maximized at $\hat{p}(c) = \eta c/(\eta-1)$, so the assumption holds if and only if $c \geq 1-1/\eta$. The isomorphism can be extended to this case by normalizing margins by $\bar{m}(c) \equiv \max_{p\in[c,1]}(p-c)x(p)$ rather than $(1-c)x(1)$. The competition layer is unchanged; what changes is the curvature layer, since $\bar{m}'(c) = -x(\hat{p}(c))$ varies with c when the maximum is interior. We maintain Assumption 2.4 throughout to preserve the simpler structure.

3. The μ-Isomorphism

This section establishes our central theoretical result: the pricing game with general demand is isomorphic to a unit-demand game after a change of variables. The key insight is that firms fundamentally care about their effective margin per customer served–how much profit they extract from each transaction. Once we normalize these margins appropriately, the strategic considerations become identical regardless of the underlying demand function.

3.1. Normalized Effective Margins.

To understand the isomorphism, we need to think about what firms actually compete over. When a firm sets price p, it earns $(p-c)x(p)$ from each customer it serves–this is the effective margin. Different demand functions $x(\cdot)$ change the mapping from prices to effective margins, but the underlying strategic trade-off remains the same: higher margins mean higher profit per customer but lower probability of winning customers.

Definition 3. Given demand $x(\cdot)$ and cost c, the *normalized effective margin* at price p is:

$$\mu(p;c) \equiv \frac{(p-c)x(p)}{(1-c)x(1)} \in [0,1] \tag{1}$$

The numerator is firm profit per served consumer. The denominator is the maximum feasible profit when pricing at the reservation value. This normalization maps all possible effective margins to the unit interval $[0,1]$, with $\mu = 0$ corresponding to pricing at cost (zero margin) and $\mu = 1$ corresponding to extracting maximum profit at the reservation price. Under Assumption 2.4, $\mu(p;c)$ is strictly increasing in p on the equilibrium support, ensuring a one-to-one correspondence between prices and normalized margins.

Definition 4. The inverse map $\phi(\mu,c)$ solves:

$$(\phi(\mu,c)-c)x(\phi(\mu,c)) = \mu(1-c)x(1) \tag{2}$$

We record the existence and uniqueness of the inverse map ϕ:

Lemma 3.1. *Under Assumptions 2.3 and 2.4, for each $\mu \in [0,1]$ and $c \in [0,1)$, there exists a unique $\phi(\mu, c) \in [c,1]$ satisfying (2).*

Proof. Define $g(p) \equiv (p-c)x(p) - \mu(1-c)x(1)$ for $p \in [c,1]$. By the continuity of $x(\cdot)$, the function g is continuous. Moreover,

$$g(c) = -\mu(1-c)x(1) \leq 0 \leq (1-c)x(1)(1-\mu) = g(1),$$

so by the intermediate value theorem, there exists $p^* \in [c,1]$ with $g(p^*) = 0$. Finally, by Assumption 2.4, $g'(p) = x(p) + (p-c)x'(p) > 0$ on the relevant domain, so g is strictly increasing, yielding uniqueness. ■

3.2. **The Main Isomorphism Result.** We now state and prove our central theorem. The result shows that once we reformulate the game in terms of normalized effective margins μ, all the complexity from demand curvature disappears. The equilibrium depends only on the consideration set structure. Demand curvature and costs then only matter for translating these margin distributions back into price distributions.

We first define the transformed game that will be central to our analysis. We define the margin game as follows

Definition 5. In the *margin game*, firms simultaneously choose probability distributions over margins $\mu \in [0,1]$. When firm i posts margin μ and rivals use CDFs $(F_j^\mu)_{j\neq i}$, firm i's payoff is $\Pi_i^\mu(\mu) = \mu \cdot q_i^\mu(\mu)$, where the demand share $q_i^\mu(\mu)$ is defined analogously to (Demand):

$$q_i^\mu(\mu) = \sum_{S \ni i} \alpha_S \cdot \mathbb{E}\left[\frac{\mathbf{1}\{i \in M_S(\mu, \mu_{-i})\}}{\left|M_S(\mu, \mu_{-i})\right|}\right],$$

with $M_S(\mu, \mu_{-i})$ denoting the set of margin-minimizers in S and ties broken uniformly. When rival CDFs are atomless at μ, ties occur with probability zero and the demand share simplifies as in (Demand*):

$$q_i^\mu(\mu) = \sum_{S \ni i} \alpha_S \prod_{j \in S \setminus \{i\}} \left[1 - F_j^\mu(\mu)\right].$$

An equilibrium in the margin game exists:

Lemma 3.2. *The margin game admits a mixed-strategy Nash equilibrium.*

The proof, which verifies the hypotheses of Dasgupta and Maskin (1986, Theorem 5), appears in Appendix A.1.

Theorem 3.3. *Consider the pricing game with demand $x(\cdot)$ satisfying Assumptions 2.3 and 2.4 and cost c. Then:*

(a) The map Φ defined by $F_i^{\mu} = \Phi(F_i)$ is a bijection between equilibria of the pricing game and equilibria of the margin game.[4]

(b) The equilibrium μ-distributions depend only on the consideration structure $(\alpha_S)_{S \subseteq N}$, not on $x(\cdot)$.

The proof appears in Appendix A.2. The key insight is that the normalized effective margin $\mu(p;c) = (p-c)x(p)/[(1-c)x(1)]$ provides a bijection between prices and margins, and payoffs in the pricing game are proportional to payoffs in the margin game with a common scaling factor $(1-c)x(1)$, *which does not affect best-response comparisons*.

We pause to remark on the role of Assumption 2.4.

Remark 3.4. The invertibility condition that $p \mapsto (p-c)x(p)$ is strictly increasing ensures that the effective margin is strictly increasing in price. Without this, multiple prices could yield the same margin, breaking the bijection. Economically, the condition requires that higher prices always generate higher profit per customer served; *viz.*, firms never face a ”backward-bending” margin curve. This holds for unit demand (where $x(p) \equiv 1$) and for most standard demand specifications when markups are moderate.

The μ-isomorphism helps us in understanding pass-through. It reveals that cost pass-through in markets with price dispersion can be decomposed into two distinct layers:

[4]Thus, the pricing game admits a mixed-strategy equilibrium.

Corollary 3.5. *Cost pass-through analysis separates into two distinct layers. The competition layer solves for equilibrium μ-distributions using only the consideration structure $\{\alpha_S\}$, determining how market power is distributed across the price distribution. The curvature layer maps normalized margins to prices via $p = \phi(\mu, c)$ using the demand function $x(\cdot)$ and cost c, determining how a given level of market power translates into actual prices. Pass-through then follows by differentiating the mapping: $\tau = \phi_c(\mu, c)$. The pass-through rate depends on both layers but in a separable way.*

This separation dramatically simplifies analysis. Rather than solving different equilibria for each $(c, x(\cdot))$ combination, we solve once in μ-space and apply different transformations. To see the power of this approach, note that the equilibrium μ-distribution is invariant to cost changes—when costs rise, the entire price distribution shifts, but the underlying distribution of market power (captured by μ) remains fixed. When multiple equilibria exist, this invariance holds for the set, but the uniqueness results in Section 4 ensure a well-defined equilibrium path for differentiation. This is why pass-through can be computed as a simple derivative of the mapping function ϕ.

3.3. **Economic Intuition.**

The normalized effective margin μ captures the fundamental trade-off in price-setting: higher μ means higher profit per customer but lower probability of winning customers. Think of μ as the firm's "aggressiveness" in extracting surplus–choosing μ is like choosing a position on the competition spectrum from aggressive (low μ, low margins, high market share) to passive (high μ, high margins, low market share).

The consideration structure $\{\alpha_S\}$ determines how this trade-off resolves in equilibrium: markets with more overlapping consideration sets (more competition) push firms toward lower μ, while markets with many captive customers allow higher μ. Crucially, the demand curvature $x(\cdot)$ only affects the translation between these strategic positions and actual prices, not the positions themselves.

3.4. **Application to Merger Analysis.** The μ-isomorphism provides a clean decomposition for merger analysis. Any merger changes the consideration structure α_S, inducing a new equilibrium in the margin game. Since the margin game is invariant to demand, this step can be computed once without specifying demand. The price effects at each quantile then follow from $\Delta p(u) = \phi(\mu^{\text{post}}(u), c) - \phi(\mu^{\text{pre}}(u), c)$. This separates the problem: the competition layer determines how margins shift, whereas the curvature layer determines how those margin shifts translate to prices. The first depends only on consideration structure, the second only on demand.

4. Equilibrium Characterization

Having established the isomorphism, we now characterize equilibrium in two important cases. These results show how the consideration set structure–the pattern of consumer search and awareness–determines the equilibrium distribution of market power. We begin with the symmetric case where all firms face identical competitive environments, then discuss how heterogeneous market positions affect equilibrium.

4.1. **Symmetric Firms.** When the consideration structure treats all firms symmetrically, for instance, when consumers randomly sample firms or when firms are arranged symmetrically in geographic or product space–equilibrium takes a particularly simple form. The key insight is that symmetric competition leads to a common distribution of market power, though firms still mix over different prices in equilibrium.

Definition 6. The structure $\{\alpha_S\}$ is *symmetric* if α_S depends only on $|S|$, not on the identity of firms in S.

Under symmetry, all firms have identical reach σ (the mass of consumers who consider them) and captive-to-reach ratio $\rho = \alpha_{\{i\}}/\sigma$ (the fraction of their potential customers who consider no other firms). The ratio ρ is a crucial parameter, measuring the degree of market power arising from limited consideration. When ρ is high, many consumers are captive to individual firms, whereas when ρ is low, most consumers compare multiple options.

Proposition 4.1. *With symmetric consideration structure, the unique symmetric equilibrium in μ-space has quantile function*

$$\mu(u) = \frac{\rho}{H(1-u)},$$

where $H(s) = \frac{1}{\sigma}\sum_{S \ni i} \alpha_S s^{|S|-1}$ *is the probability generating function of* $|S|-1$ *conditional on* $i \in S$.

Intuitively, $H(s)$ encodes the competitive environment: it tells us the distribution of how many rival firms a consumer considers, conditional on considering firm i. The formula shows that firms mix over higher margins (higher μ) when they have more captive customers (higher ρ). Perhaps less obviously, margins are also higher when consumers consider more rivals (lower H). The logic, which goes back to Rosenthal (1980), is that with more rivals, competing aggressively for contested consumers becomes less valuable since they are spread across more firms. Firms respond by shifting their mixing toward higher margins, focusing on extracting surplus from captive customers rather than competing fiercely for contested ones.

4.2. **Example: Binomial Consideration.** To make these concepts concrete, consider a market where each consumer independently considers each firm with probability λ. This captures settings like online markets where consumers randomly encounter products, or markets where advertising reaches consumers stochastically.

Example 4.2 (Random Search Equilibrium)**.** With binomial consideration where each consumer considers each firm independ90

$$\mu(u) = \frac{\rho}{H(1-u)} = \frac{(1-\lambda)^{n-1}}{(\lambda(1-u)+(1-\lambda))^{n-1}} = \left(\frac{1-\lambda}{1-\lambda u}\right)^{n-1}.$$

4.3. **Asymmetric Duopoly.** We now characterize equilibrium when two firms have different captive shares. This case is tractable and reveals how asymmetric market positions affect equilibrium price dispersion and pass-through.

Consider $n = 2$ firms with captive shares $\alpha_1 \equiv \alpha_{\{1\}}$ and $\alpha_2 \equiv \alpha_{\{2\}}$, shared segment $\alpha_{12} \equiv \alpha_{\{1,2\}}$ (consumers who consider both), and outside option $\alpha_\emptyset$. Define the captive-to-reach ratios:

$$\rho_1 = \frac{\alpha_1}{\alpha_1 + \alpha_{12}}, \quad \rho_2 = \frac{\alpha_2}{\alpha_2 + \alpha_{12}}$$

The captive-to-reach ratio ρ_i measures the fraction of firm i's potential customers who have no alternative. Higher ρ_i means more market power from captive consumers.

Proposition 4.3. *Assume $1 > \rho_i > 0$. Let $\rho_1 > \rho_2$. The unique equilibrium in μ-space has both firms mixing on the common support $[\underline{\mu}, 1]$ with lower bound $\underline{\mu} = \rho_1$. The equilibrium CDFs on $\mu \in [\underline{\mu}, 1)$ are*

$$F_1^\mu(\mu) = 1 - \frac{1}{1-\rho_2}\left(\frac{\rho_1}{\mu} - \rho_2\right) \quad \text{and} \quad F_2^\mu(\mu) = 1 - \frac{\rho_1}{1-\rho_1}\left(\frac{1-\mu}{\mu}\right).$$

Firm 1 (with higher ρ_1) has a mass point at $\mu = 1$ of size $\Delta_1 = 1 - F_1^\mu(1) = (\rho_1 - \rho_2)/(1 - \rho_2)$, while firm 2 has no atom ($F_2^\mu(1) = 1$).

The corresponding asymmetric duopoly quantile functions are

$$\mu_1(u) = \begin{cases} \frac{\rho_1}{1-u(1-\rho_2)} & \text{if} \quad u \leq 1 - \Delta_1 \\ 1 & \text{if} \quad u > 1 - \Delta_1, \end{cases} \quad \text{and} \quad \mu_2(u) = \frac{\rho_1}{1 - u(1-\rho_1)}.$$

The asymmetric equilibrium reveals how differences in captive shares shape competitive behavior. The firm with more captives ($\rho_1 > \rho_2$) prices at the monopoly margin $\mu = 1$ with positive probability. The firm with fewer captives earns rents: it benefits from the "price umbrella" set by the stronger firm, earning expected profits strictly above its captive value ($\pi_2^* = \rho_1(\alpha_2 + \alpha_{12}) > \alpha_2$). Both firms share the same support-lower-bound $\underline{\mu} = \rho_1$, determined by the stronger firm's desire to exploit its large captive base.

We can also order the duopoly margins. For $u \le 1 - \Delta_1$, the quantile functions satisfy:

$$\mu_1(u) = \frac{\rho_1}{1 - u(1-\rho_2)} > \frac{\rho_1}{1 - u(1-\rho_1)} = \mu_2(u).$$

since $\rho_1 > \rho_2$ implies $1-\rho_1 < 1-\rho_2$. Firm 1 (with more captive customers) maintains higher margins and prices at every quantile.

4.4. **Extension to n Asymmetric Firms: Independent Consideration.** The duopoly analysis extends to n asymmetric firms when consideration sets exhibit a particular structure: *independent awareness*. Under this assumption, whether a consumer considers firm i is statistically independent of whether she considers firm j. This structure, used by Guthmann (2025), yields closed-form equilibrium characterizations for arbitrary n.

Assumption 4.4. Each consumer considers firm j independently with probability $\lambda_j \in (0,1)$. The consideration structure is

$$\alpha_S = \prod_{j \in S} \lambda_j \prod_{k \notin S} (1 - \lambda_k).$$

Under independence, firm j's reach is $\sigma_j = \lambda_j$ and its captive share is $\alpha_{\{j\}} = \lambda_j \prod_{k \neq j}(1 - \lambda_k)$. The captive-to-reach ratio becomes:

$$\rho_j = \frac{\alpha_{\{j\}}}{\sigma_j} = \prod_{k \neq j}(1 - \lambda_k).$$

A key property of independence is that firm j's demand share when posting μ takes a multiplicatively separable form:

$$q_j^{\mu}(\mu) = \sum_{S \ni j} \alpha_S \prod_{i \in S \setminus \{j\}} (1 - F_i^{\mu}(\mu)) = \lambda_j \prod_{i \neq j} \left[1 - \lambda_i F_i^{\mu}(\mu)\right]. \tag{3}$$

This separability is what enables closed-form solutions. We state and prove it for completeness.

Proposition 4.5. *Posit Assumption 4.4 and order firms so that* $\lambda_1 > \lambda_2 \geq \cdots \geq \lambda_n$. *Define the common lower bound* $\underline{\mu} := \rho_1 = \prod_{h=2}^{n}(1-\lambda_h)$ *and define the upper support bounds by* $\bar{\mu}_1 = \bar{\mu}_2 := 1$, *and for each* $k \in \{3, \ldots, n\}$,

$$\bar{\mu}_k := \frac{\prod_{h=2}^{k-1}(1-\lambda_h)}{(1-\lambda_k)^{k-2}} \in (0,1), \quad \textit{with the convention } \bar{\mu}_{n+1} := \underline{\mu}.$$

Then the unique equilibrium in μ*-space has the following structure:*

(1) *Firm 1 mixes on* $[\underline{\mu}, 1]$ *and has an atom at* $\mu = 1$. *Firm 2 mixes continuously on* $[\underline{\mu}, 1]$ *with no atom. Each firm* $k \geq 3$ *mixes continuously on* $[\underline{\mu}, \bar{\mu}_k]$ *with* $\bar{\mu}_k < 1$. *The supports are nested:* $1 = \bar{\mu}_1 = \bar{\mu}_2 > \bar{\mu}_3 \geq \cdots \geq \bar{\mu}_n > \underline{\mu}$.

(2) *Each firm* j *earns equilibrium profit* $\pi_j^* = \lambda_j \underline{\mu}$, *with* $\pi_1^* = \lambda_1 \underline{\mu} = \alpha_{\{1\}}$.

(3) *For each* $m \in \{2, \ldots, n\}$, *define* $C_m := \prod_{h=m+1}^{n}(1-\lambda_h)$ *(so* $C_n = 1$ *and* $C_2 = \prod_{h=3}^{n}(1-\lambda_h)$*), and for* $\mu \in [\bar{\mu}_{m+1}, \bar{\mu}_m]$ *define the common multiplier* $\Gamma(\mu) := 1 - \left(\underline{\mu}/(\mu C_m)\right)^{1/(m-1)}$. *Then for each* $\mu \in [\bar{\mu}_{m+1}, \bar{\mu}_m]$,

$$F_j^{\mu}(\mu) = \begin{cases} \Gamma(\mu)/\lambda_j, & j \leq m, \\ 1, & j > m, \end{cases} \qquad \textit{and} \qquad F_j^{\mu}(\mu) = 0 \ \textit{ for } \mu < \underline{\mu}.$$

On the common overlap $[\underline{\mu}, \bar{\mu}_n]$ *(where* $m = n$*),*[5]

$$\Gamma(\mu) = 1 - \left(\underline{\mu}/\mu\right)^{1/(n-1)} \quad \textit{and} \quad F_j^{\mu}(\mu) = \frac{\Gamma(\mu)}{\lambda_j}. \tag{4}$$

For firm $j > 1$ (no atom), on domain $[0,1]$ mapping to $[\underline{\mu}, \bar{\mu}_j]$:

$$\mu_j(u) = \underline{\mu} \cdot \left[1 - \lambda_j u\right]^{-(n-1)}$$

We note the following stochastic dominance of margins, namely, that higher-reach firms price higher at each quantile.

[5]Firm 1 has a mass point at $\mu = 1$ of size $\Delta_1 = 1 - F_1^{\mu}(1-) = 1 - \lambda_2/\lambda_1$. All firms $j \geq 2$ have no atom at $\mu = 1$.

Corollary 4.6. *Under independence with $\lambda_1 > \lambda_2 \geq \cdots \geq \lambda_n$, the margin CDFs satisfy $F_1^\mu(\mu) \leq F_2^\mu(\mu) \leq \cdots \leq F_n^\mu(\mu)$ for all μ in the common support. Equivalently, $\mu_1(u) \geq \mu_2(u) \geq \cdots \geq \mu_n(u)$ for all $u \in [0,1]$.*

The independent consideration model reveals a clean hierarchical structure. The key equilibrium object is $\Gamma(\mu) = \lambda_j F_j^\mu(\mu)$, which is *common across all firms* in the overlapping support region. Higher-reach firms (λ_j large) have flatter CDFs: they spread the same $\Gamma(\mu)$ over a larger mass of aware consumers, so each consumer is less likely to see low prices from them. The firm with highest reach faces the most competition and must sometimes charge the monopoly margin to earn its equilibrium profit, thus, the atom at $\mu = 1$.

Firms with lower reach have relatively more captive power and can achieve their profits at margins below monopoly, so their supports end before $\mu = 1$. The support lower bound $\underline{\mu} = \prod_{h \neq 1}(1 - \lambda_h) = \rho_1$ depends on the reach of *all other firms*; when rivals have high reach, the competitive floor on margins falls. Finally, pass-through heterogeneity across firms arises purely from their positions in the margin distribution; firms with higher reach post higher margins at each quantile and, if demand is not too convex (as for unit or linear demand), have lower pass-through at each quantile.

Of course, the independent consideration structure is a special case of general consideration sets. The key simplification is that the demand share (3) factors multiplicatively, implying that $\lambda_j F_j^\mu(\mu)$ must be equal across active firms on each support interval. Equation (4) solves this equal-Γ structure explicitly on the lowest interval where all n firms compete. For general (correlated) consideration structures, this factorization fails and closed-form solutions are unavailable for $n \geq 3$. Nevertheless, the qualitative properties (common support, stochastic dominance ordering, atoms for high-reach firms, pass-through ranking) extend to the general case.

5. Quantile Pass-Through

We now derive our main pass-through results. The key insight from the μ-isomorphism is that pass-through rates can be computed by simply differentiating the mapping from margins to prices. Since the equilibrium margin distribution is invariant to cost changes, all the action comes from how the mapping ϕ responds to costs. Throughout, fix a firm i with equilibrium μ-quantile function $\mu_i(u)$.

5.1. **The Pass-Through Formula.**

Quantile pass-through tells us how each price in the distribution responds to cost changes. Because firms randomize in equilibrium, different quantiles of the price distribution can have different pass-through rates. This heterogeneity in pass-through across the price distribution is a key feature of markets with price dispersion.

Definition 7. The *pass-through rate at quantile u* is $\tau_i^Q(u;c) \equiv \frac{\partial p_i(u;c)}{\partial c}$, where $p_i(u;c) = \phi(\mu_i(u), c)$.

Theorem 5.1. *Under Assumptions 2.3 and 2.4, the quantile pass-through rate is:*

$$\tau_i^Q(u;c) = \phi_c(\mu_i(u), c) = \frac{x(p_i(u;c))(1-p_i(u;c))}{(1-c)[x(p_i(u;c)) + (p_i(u;c)-c)x'(p_i(u;c))]} \tag{5}$$

Proof. Recall the implicit definition (2):

$$(\phi(\mu,c)-c)x(\phi(\mu,c)) = \mu(1-c)x(1).$$

Differentiating both sides with respect to c, holding μ fixed, and rearranging yields

$$\phi_c[x(\phi) + (\phi-c)x'(\phi)] = x(\phi) - \mu x(1).$$

From (2), we have $\mu x(1) = \frac{(\phi-c)x(\phi)}{1-c}$. Substituting this in produces

$$\phi_c[x(\phi) + (\phi-c)x'(\phi)] = x(\phi)\frac{1-\phi}{1-c},$$

so,

$$\phi_c(\mu,c) = \frac{x(\phi(\mu,c))(1-\phi(\mu,c))}{(1-c)[x(\phi(\mu,c)) + (\phi(\mu,c)-c)x'(\phi(\mu,c))]}.$$

Finally, setting $\mu = \mu_i(u)$ and noting that $\phi(\mu_i(u), c) = p_i(u; c)$ yields (5). ■

5.2. **Economic Interpretation.** The pass-through formula (5) reveals how market power and demand curvature jointly determine cost incidence. The numerator $x(p)(1-p)$ captures the direct effect of cost on profit: the factor $x(p)$ reflects that firms selling more quantity face a larger cost burden per customer served, while $(1-p)$ represents the "headroom" for price increases. Firms pricing closer to the reservation value have less room to raise prices. Together, these determine the pressure to pass costs through.

The denominator $(1-c)[x(p)+(p-c)x'(p)]$ is the slope of the effective margin function, measuring how responsive profit-per-customer is to price changes. When demand is highly elastic (large $|x'|$), small price increases cause large quantity losses, dampening pass-through. When demand is inelastic, firms can raise prices without losing much quantity, facilitating pass-through.

Limiting cases provide important benchmarks. Under *unit demand* ($x' = 0$), pass-through is $\tau = (1-p)/(1-c)$; firms pricing higher have lower pass-through because they are already extracting rents. Under *perfect competition* ($p \to c$), pass-through approaches 1 since firms with zero markups must fully pass through cost changes to break even. Under *monopoly pricing*, pass-through depends entirely on demand curvature, potentially exceeding or falling below 100%.

6. Transaction-Weighted Pass-Through

While quantile pass-through describes how each price in the distribution responds to costs, welfare analysis requires understanding what consumers actually pay. This distinction matters because consumers do not randomly draw from the price distribution; they systematically buy more at lower prices. This section derives pass-through for transaction-weighted prices, which captures the average pass-through experienced by consumers.

6.1. **The Transaction-Weighting Problem.** With price dispersion, the average posted price differs from the average price paid by consumers. Lower prices attract more buyers, creating a composition effect. To see why this matters, consider a market where half the firms charge \$10 and half charge \$20. The average posted price is \$15, but if 90% of the units sold are at the more attractive \$10 price, the average transaction price is only \$11. This selection effect fundamentally changes how we think about pass-through and incidence.

Lemma 6.1. *Under the equilibrium indifference condition, the mass of transactions at price p is:*

$$T_i(p;c) = x(p)q_i(p) = \frac{\pi_i(c)}{p-c} \tag{6}$$

Proof. From the equilibrium indifference condition, for all $p \in \text{supp}(F_i)$:

$$(p-c)x(p)q_i(p) = \pi_i(c)$$

Rearranging yields (6). ■

The key insight is that transaction volume is proportional to $1/(p-c)$, independent of the demand function $x(\cdot)$. Intuitively, the equilibrium indifference condition ensures that all prices yield the same profit, so firms selling at lower prices (with lower margins) must compensate with proportionally higher volume. This invariance property greatly simplifies the analysis of transaction-weighted pass-through.

Definition 8. The *transaction-weighted CDF* is:

$$F_i^{\text{trans}}(p;c) = \frac{\int_c^p \frac{1}{s-c} dF_i(s;c)}{\int_c^1 \frac{1}{s-c} dF_i(s;c)} \tag{7}$$

6.2. **Mean Paid Prices and Markups.** Next we note the harmonic mean of posted markups. For firm i's price distribution with support bounded away from c (which holds when $\rho > 0$), we define

$$B_i(c) \equiv \int \frac{1}{p-c} dF_i(p;c) \tag{8}$$

This integral is well-defined when the support of F_i is contained in $[c+\epsilon, 1]$ for some $\epsilon > 0$, which occurs whenever there are captive customers ($\rho > 0$). Then, the mean paid markup is the harmonic mean of posted markups:

Proposition 6.2. *The mean transaction-weighted price is* $\bar{p}_i^{trans}(c) = c + \frac{1}{B_i(c)}$.

Proof. From (7), the mean paid price is

$$\bar{p}_i^{\text{trans}}(c) = \int p\, dF_i^{\text{trans}}(p;c) = \frac{\int p \cdot \frac{1}{p-c} dF_i(p;c)}{\int \frac{1}{p-c} dF_i(p;c)} = \frac{1 + cB_i(c)}{B_i(c)} = c + \frac{1}{B_i(c)}.$$

■

6.3. **Transaction-Weighted Pass-Through.** We now derive how mean paid prices respond to cost changes. The analysis reveals a surprising result: transaction-weighted pass-through can differ substantially from the simple average of quantile pass-through rates. The reason is that cost changes affect not just prices but also the distribution of transactions across prices.

Theorem 6.3. *Under Assumptions 2.3 and 2.4, and when the support is bounded away from c (which holds when* $\rho > 0$*), the transaction-weighted pass-through rate is*

$$\tau_i^{trans}(c) = 1 + \frac{\int_0^1 \frac{\phi_c(\mu_i(u),c)-1}{(\phi(\mu_i(u),c)-c)^2} du}{\left[\int_0^1 \frac{1}{\phi(\mu_i(u),c)-c} du\right]^2} \tag{9}$$

Proof. From Proposition 6.2:

$$\tau_i^{\text{trans}}(c) = \frac{d\bar{p}_i^{\text{trans}}(c)}{dc} = 1 - \frac{B_i'(c)}{B_i(c)^2}$$

We need to compute $B_i'(c)$. We take the quantile representation

$$B_i(c) = \int_0^1 \frac{1}{p_i(u;c) - c} du,$$

and differentiate under the integral sign (valid by dominated convergence when $\rho > 0$, which ensures $p_i(u;c) - c \geq \epsilon > 0$ uniformly):

$$B_i'(c) = \int_0^1 \frac{\partial}{\partial c}\left(\frac{1}{p_i(u;c)-c}\right)du = \int_0^1 \frac{-(\tau_i^Q(u;c)-1)}{(p_i(u;c)-c)^2}du = -\int_0^1 \frac{\phi_c(\mu_i(u),c)-1}{(\phi(\mu_i(u),c)-c)^2}du,$$

using $p_i(u;c) = \phi(\mu_i(u),c)$ and $\tau_i^Q(u;c) = \phi_c(\mu_i(u),c)$. Substituting into the expression for $\tau_i^{\text{trans}}(c)$ yields (9). ■

6.4. **Pass-Through in Terms of Consideration Parameters.** The transaction-weighted pass-through formula (9) simplifies dramatically under unit demand, yielding a closed-form expression in terms of the consideration structure. This result provides a direct link between market structure (as captured by consideration sets) and cost incidence, without requiring knowledge of demand curvature.

Proposition 6.4. *Under unit demand ($x(p) = 1$), define for each firm i $K_i \equiv \int_0^1 \frac{1}{\mu_i(u)}du$, where $\mu_i(u)$ is firm i's equilibrium quantile function in the margin game. Then firm i's mean transaction price and transaction-weighted pass-through satisfy*

$$\bar{p}_i^{trans}(c) = c + \frac{1-c}{K_i}, \qquad \text{and} \qquad \tau_i^{trans} \equiv \frac{d\bar{p}_i^{trans}(c)}{dc} = 1 - \frac{1}{K_i}. \tag{10}$$

Proof. Under unit demand, $p = c+(1-c)\mu$, so $p_i(u;c)-c = (1-c)\mu_i(u)$. The transaction-weighting identity

$$\bar{p}_i^{\text{trans}}(c) = c + \frac{1}{B_i(c)}, \qquad \text{where} \qquad B_i(c) = \int_0^1 \frac{1}{p_i(u;c)-c}\,du$$

implies

$$B_i(c) = \int_0^1 \frac{1}{(1-c)\mu_i(u)}\,du = \frac{K_i}{1-c}.$$

Substituting yields $\bar{p}_i^{\text{trans}}(c) = c+(1-c)/K_i$. The μ-isomorphism implies $\mu_i(\cdot)$ is cost-invariant, so K_i does not depend on c. Differentiating delivers $\tau_i^{\text{trans}} = 1 - 1/K_i$. ■

The object K_i measures the intensity of competition facing firm i, aggregated across its transaction distribution. When the margin distribution places substantial weight on low μ (contested transactions), K_i is large and pass-through is high:

firms facing stiff competition pass cost shocks to consumers. When the distribution concentrates on high μ (captive transactions), K_i is small and pass-through is low: firms with market power absorb cost shocks.

Under unit demand and symmetric consideration, the equilibrium quantile function satisfies $\mu(u) = \rho/H(1-u)$, where $\rho = \alpha_{\{i\}}/\sigma$ is the captive-to-reach ratio and H is the probability generating function. Hence, $K = \frac{\bar{H}}{\rho}$, where $\bar{H} \equiv \int_0^1 H(s)d$, and the general formula (10) reduces to $\tau^{\text{trans}} = 1 - \frac{\rho}{\bar{H}}$.

There is a natural application to mergers. Standard merger analysis focuses on price levels. But mergers also affect how future cost shocks are transmitted to consumers. By the separation principle, a merger changes the consideration structure, which shifts each firm's equilibrium margin distribution. Pass-through at each quantile then follows from $\tau^Q(u) = \phi_c(\mu(u), c)$, and transaction-weighted pass-through from Theorem 6.3. Under unit demand, this simplifies: the new $K_i = \int \mu_i(u)^{-1} du$ determines $\tau_i^{\text{trans}} = 1 - 1/K_i$. Under independent consideration, a merger between non-leader firms leaves the margin floor $\underline{\mu}$ unchanged but shifts posted distributions toward lower margins, raising K_i for the remaining firms and, hence, the transaction-weighted pass-through. Even mergers with no immediate price effect can shift the incidence of future cost shocks toward consumers.

7. Pass-Through Envelopes

In many empirical settings, we may not know the exact demand function. Perhaps we observe that demand is downward-sloping but cannot pin down its precise curvature. Or we may know that demand belongs to a particular family (e.g., linear, constant elasticity) but not the exact parameters. This section shows that even with such partial knowledge, we can still derive robust bounds on pass-through. These bounds provide ”worst-case” scenarios for policy analysis, following the robust welfare approach of Kang and Vasserman (2025), and help identify when precise demand estimation is crucial versus when rough knowledge suffices.

7.1. **The Envelope Problem.** Now let us turn the pass-through problem on its head. Instead of starting with a known demand function and computing pass-through, we ask: given a particular margin level μ and cost c, what is the range of possible pass-through rates across all admissible demand functions?

Assumption 7.1. Let $\mathcal{X}$ denote the class of *admissible demand functions* $x\colon [0,1] \to \mathbb{R}_+$ that are continuous and weakly decreasing, have positive demand at the upper bound ($x(1) > 0$, which we can without loss of generality normalize to $x(1) = 1$ by rescaling quantity units), and satisfy the invertibility assumption (Theorem 2.4).

Formally, fix $\mu \in [0,1]$ and $c \in [0,1)$. The inverse problem is to find all prices p consistent with

$$(p-c)x(p) = \mu(1-c)x(1) \tag{11}$$

for some admissible demand function $x(\cdot)$.

7.2. **Universal Bounds.** We begin with the most general case: what can we say about pass-through knowing only that demand is downward-sloping? The answer provides universal bounds that apply regardless of the specific functional form.

Theorem 7.2. *For any admissible* $x \in \mathcal{X}$, $\mu \in [0,1]$, *and* $c \in [0,1)$, *we have the price bounds* $c \le \phi(\mu,c) \le c + \mu(1-c)$ *and the pass-through bounds* $1-\mu \le \phi_c(\mu,c)$. *The upper bound for prices and the lower bound for pass-through are attained by unit demand.*

These bounds have important economic implications. The pass-through bound $\tau \ge 1-\mu$ tells us that firms with lower market power (lower μ) have pass-through rates bounded further from zero. In the limit, firms pricing at cost ($\mu = 0$) must have pass-through of at least 1, recovering the perfect competition result. Conversely, firms extracting maximum margins (μ near 1) could have pass-through rates approaching zero.

Proof. For price bounds, we normalize $x(1) = 1$ (by rescaling quantity units). Since x is decreasing, we have $x(p) \geq 1$ for $p \in [c, 1]$. From (11),

$$\mu(1-c) = (p-c)x(p) \geq p - c \quad \Longrightarrow \quad p \leq c + \mu(1-c).$$

The lower bound $p \geq c$ is trivial.

For the pass-through bound, from Theorem 5.1, with $\varepsilon \equiv -(p-c)x'(p)/x(p) \in [0,1)$,

$$\phi_c(\mu, c) = \frac{1 - \phi(\mu, c)}{(1-c)(1-\varepsilon)} \geq \frac{1 - (c + \mu(1-c))}{1-c} = 1 - \mu,$$

as $\varepsilon \geq 0$ and $\phi(\mu, c) \leq c + \mu(1-c)$. ■

7.3. **Bounds for Specific Demand Families.** While universal bounds are useful, we can derive tighter bounds when we know more about demand. Different assumptions about demand curvature–whether demand is linear, exponential, or has constant elasticity–yield different pass-through bounds. These family-specific bounds help connect theoretical predictions to empirical demand estimation.

Theorem 7.3. *For common demand families (with $d \equiv 1 - c$):*

(1) ***Linear demand*** $x(p) = 1 + b(1-p)$ *with* $b \in [0, 1/d]$*:*

$$\phi_c(\mu, c) \leq \frac{1 + \sqrt{1-\mu}}{2}, \quad \textit{with equality at maximum slope } b = \frac{1}{d}.$$

(2) ***Constant semi-elasticity*** $x(p) = e^{\beta(1-p)}$ *with* $\beta \in [0, 1/d]$*:*

$$\phi_c(\mu, c) \leq 1, \quad \textit{with equality at } \beta = \frac{1}{d}.$$

(3) ***Constant elasticity*** $x(p) = p^{-\eta}$ *with* $\eta \geq 0$*:*

$$\phi_c(\mu, c) = \frac{1-\mu}{(1-\mu d)^2} \quad \textit{when } \eta = 1.$$

The proof, which provides complete derivations for each demand family, appears in Appendix B.4.

Figure 1 illustrates these bounds for the linear and CES demand families. The left panel shows the linear demand family $x(p) = 1 + b(1-p)$ as the slope parameter

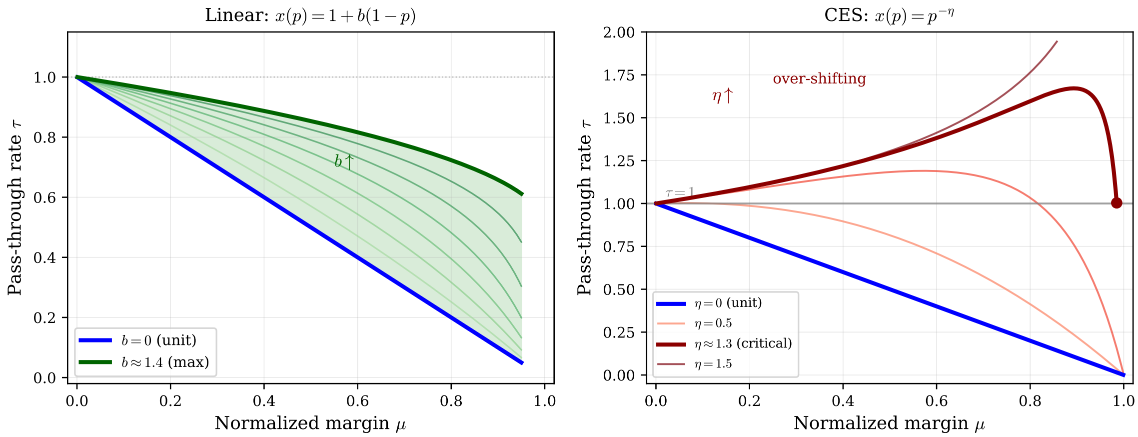

FIGURE 1. **Robust pass-through bounds by demand family.** Left panel: For linear demand $x(p) = 1+b(1-p)$ with slope $b \in [0, 1/d]$, pass-through lies between the unit demand lower bound $\tau = 1-\mu$ and the upper bound $\tau = (1 + \sqrt{1-\mu})/2$. All linear demands yield $\tau \leq 1$. Right panel: For CES demand $x(p) = p^{-\eta}$, higher elasticity η produces over-shifting ($\tau > 1$). Below the critical elasticity $\eta \approx 1.3$, pass-through eventually falls below one at high margins; above it, the invertibility condition binds before τ can fall to one.

b varies from zero (unit demand) to its maximum $b = 1/d$. Unit demand delivers the universal lower bound $\tau = 1 - \mu$, while steeper demand curves yield higher pass-through, reaching the upper bound $\tau = (1+\sqrt{1-\mu})/2$ at $b = 1/d$. Crucially, all linear demands produce incomplete pass-through ($\tau < 1$) throughout the margin support.

The right panel reveals strikingly different behavior for CES demand $x(p) = p^{-\eta}$. As elasticity η increases, the pass-through curve rotates upward, eventually crossing into the over-shifting region ($\tau > 1$). A critical elasticity $\eta \approx 1.3$ marks the transition. Below this threshold, pass-through is incomplete at high margins but can exceed one at low margins. Above it, the invertibility condition binds before pass-through falls to one, yielding pure over-shifting throughout the feasible range. The endpoints of the CES curves reflect this constraint. High-elasticity demands cannot support margins close to one because the required price would violate invertibility.

8. Comparative Statics

The separation principle yields a simple logic for comparative statics. Any comparison of prices or pass-through–whether across market structures, across firms, or across demand specifications–reduces to comparing margin distributions and applying the maps ϕ and ϕ_c. We develop this logic and its applications below.

The quantile pass-through formula (Theorem 5.1) expresses pass-through as a composition $\tau^Q(u;c) = \phi_c(\mu(u),c)$. Any ordering of margins, therefore, translates into an ordering of prices and pass-through, subject to the monotonicity properties of ϕ and ϕ_c.

Proposition 8.1. *Let $\mu^A(u)$ and $\mu^B(u)$ be two margin quantile functions (for the same firm across different markets, or for different firms in the same market). If $\mu^B(u) \geq \mu^A(u)$ for all $u \in [0,1]$, then: **(a)** prices inherit the ordering, $p^B(u;c) \geq p^A(u;c)$ for all u, since $\phi(\mu,c)$ is increasing in μ; and **(b)** pass-through inherits the ordering if ϕ_c is increasing in μ, or the reverse ordering if ϕ_c is decreasing in μ.*

Proof. Part (a) follows from $\phi(\cdot,c)$ strictly increasing (Lemma 3.1). Part (b) follows from the formula $\tau^Q(u;c) = \phi_c(\mu(u),c)$ and monotonicity of ϕ_c. ■

For unit demand, $\phi_c(\mu,c) = 1-\mu$, which is decreasing in μ. So higher margins imply lower pass-through. For linear demand, ϕ_c is also decreasing in μ. For CES demand with high elasticity, ϕ_c can be increasing in μ, reversing the pass-through ordering.

Consider first comparing the same firm across two consideration structures A and B. If structure B induces higher margins–say, because fewer consumers are informed or more are captive–then B has higher prices at each quantile. If demand is not too convex (so ϕ_c is decreasing), then B also has lower pass-through at each quantile: the less competitive market absorbs more of cost shocks.

The same logic applies to comparing different firms within a single market. We established in Section 4 that higher-reach firms maintain higher margins at each quantile in asymmetric equilibria. By the ordering principle, these firms charge higher prices at each quantile. If demand is not too convex, they also have lower pass-through at each quantile. Firms with more captive customers extract higher margins and absorb more of cost changes, partially insulating captive consumers.

In the symmetric benchmark (Definition 6), the equilibrium μ-quantile satisfies

$$\mu(u) = \frac{\rho}{H(1-u)}, \quad \text{where} \quad H(s) := \frac{1}{\sigma}\sum_{S \ni i} \alpha_S s^{|S|-1}.$$

Equivalently, if $K := |S|-1$ under the conditional law $\Pr(\cdot \mid i \in S)$, then $H(s) = \mathbb{E}\left[s^K\right]$ is the probability generating function of K. Letting $\beta_k := \Pr(K = k \mid i \in S)$, we have $H(s) = \sum_{k=0}^{n-1} \beta_k s^k$. When α_S depends only on $m = |S|$, let $A_m := \sum_{|S|=m} \alpha_S$ be the mass of consumers considering exactly m firms; then $\beta_{m-1} = mA_m / \sum_{j=1}^n jA_j$ and $H(s) = \sum_{m=1}^n mA_m s^{m-1} / \sum_{j=1}^n jA_j$.

Corollary 8.2. *Let K^A, K^B be the conditional rival-count random variables in two symmetric markets A, B. Then,*

$$H_A(s) \le H_B(s) \;\; \forall s \in [0,1] \qquad \Longleftrightarrow \qquad \mathbb{E}\left[\exp\left(-tK^A\right)\right] \le \mathbb{E}\left[\exp\left(-tK^B\right)\right] \;\; \forall t \ge 0,$$

using $s = \exp(-t)$ and $H(s) = \mathbb{E}\left[s^K\right]$.

This equivalence connects the PGF ordering to the familiar Laplace-transform ordering used in reliability theory and stochastic dominance.

Corollary 8.3. *Consider two symmetric markets A, B with the same ρ. Then*

$$H_A(s) \le H_B(s) \;\; \forall s \in [0,1] \qquad \Longleftrightarrow \qquad \mu_A(u) \ge \mu_B(u) \;\; \forall u \in [0,1].$$

Consequently, under Assumption 2.3 one also has $p_A(u;c) \ge p_B(u;c)$ for all $u \in [0,1]$ and all $c \in [0,1)$.

Proof. If $H_A(s) \le H_B(s)$, then for each $u \in [0,1]$ one has $H_A(1-u) \le H_B(1-u)$, hence $\mu_A(u) = \rho/H_A(1-u) \ge \rho/H_B(1-u) = \mu_B(u)$. Conversely, if $\mu_A(u) \ge \mu_B(u)$ for all u, then $\rho/H_A(1-u) \ge \rho/H_B(1-u)$, so $H_A(1-u) \le H_B(1-u)$ for all u, i.e. $H_A(s) \le H_B(s)$ for all $s \in [0,1]$. The price claim follows by the monotonicity of ϕ and composition. ∎

Corollary 8.4. *Fix two symmetric markets A, B with the same ρ. If there exists $s^* \in (0,1)$ such that $H_A(s^*) > H_B(s^*)$, then letting $u^* := 1 - s^*$,*

$$\mu_A(u^*) = \frac{\rho}{H_A(1-u^*)} < \frac{\rho}{H_B(1-u^*)} = \mu_B(u^*),$$

so global first-order stochastic dominance of μ and, hence, of posted prices under Assumption 2.3 fails.

In the symmetric benchmark, $\mu(u) = \rho/H(1-u)$ makes clear that global dominance can fail even if H-order holds, whenever the market-structure change also shifts ρ.

The preceding results concern posted-price quantiles. For welfare analysis, transaction-weighted prices matter. The following corollary shows that price dominance carries through to mean paid prices.

Corollary 8.5. *Fix $c \in [0,1)$ and a firm i. Assume that in each market $M \in \{A, B\}$ the posted-price quantile function $p_i^M(\cdot; c)$ satisfies $\underline{m}_i^M(c) := \inf_{u\in[0,1]}\left(p_i^M(u;c) - c\right) > 0$.*[6] *If $p_i^A(u;c) \geq p_i^B(u;c)$ for all $u \in [0,1]$, then $B_i^A(c) \leq B_i^B(c)$.*

Proof. Directly,

$$p_i^A(u;c) \geq p_i^B(u;c) \quad \Longrightarrow \quad \frac{1}{p_i^A(u;c) - c} \leq \frac{1}{p_i^B(u;c) - c} \ (\forall\ u) \quad \Longrightarrow \quad B_i^A(c) \leq B_i^B(c),$$

as $p \mapsto 1/(p-c)$ is strictly decreasing on (c, ∞). ■

Together, these results provide a toolkit for comparative statics: the ordering principle (Proposition 8.1) translates margin comparisons into price and pass-through comparisons, whether across markets or across firms. Corollary 8.5 extends price orderings to transaction-weighted objects relevant for welfare.

[6]This ensures that $B_i^M(c) := \int_0^1 \frac{1}{p_i^M(u;c)-c}\,du$ is finite.

9. Economic Applications

Incidence varies across the price distribution. The framework characterizes pass-through at each quantile, with patterns that depend on demand curvature. We illustrate with two markets featuring well-documented price dispersion.

9.1. **Gasoline Markets.** Retail gasoline exhibits substantial price dispersion even for a homogeneous product, with nearby stations often charging different prices. This dispersion reflects heterogeneous consumer search: some drivers compare prices across stations while others buy from the nearest option. Our framework applies directly.

The empirical pass-through literature has documented significant heterogeneity in gasoline markets. Marion and Muehlegger (2011) find that pass-through of state fuel taxes varies with supply conditions, while Stolper (2017) shows that station-level pass-through varies with local competition and spatial isolation. Montag et al. (2023) document heterogeneity at the consumer level, finding that informed consumers (who compare prices and buy cheap) face higher pass-through than uninformed consumers (who buy from one station). Our theory provides an explanation: differences in consideration patterns generate different equilibrium margins, leading to different pass-through rates even with identical demand.

Under the framework, stations in competitive locations (where many consumers compare prices) operate at low margins. If demand is not too convex, these stations exhibit high pass-through. Stations in captive locations (highway exits, isolated areas) charge higher prices and, under moderate demand curvature, absorb more of cost shocks. The framework predicts that markets with more price-comparing consumers have higher aggregate pass-through when demand is not too convex, consistent with the empirical finding that pass-through is higher in urban areas with more stations.

9.2. **Online Retail.** Online markets feature persistent price dispersion despite low search costs. Ellison and Ellison (2009) show that firms actively obfuscate to soften price competition. Platforms shape consideration sets through search rankings, algorithmic recommendations, and sponsored listings–even when consumers can easily compare prices, they often consider only a subset of sellers.

Heim (2021) provides evidence linking search behavior to pass-through. Using data from price comparison sites, he finds that pass-through of input costs depends on consumer search intensity: cost increases pass through less when consumers search more. Our framework captures this mechanism: search intensity determines which part of the margin distribution consumers transact at, and hence which pass-through rate they face. Under sufficiently convex demand, consumers transacting at lower margins face lower pass-through, consistent with Heim's finding.

The framework also applies to platform fees. Platforms charge sellers commissions. These fees enter as costs and pass through to consumers. When a platform shows each product to fraction λ of users, the symmetric equilibrium has $\mu(u) = [(1-\lambda)/(1-\lambda u)]^{n-1}$. Platforms that increase λ (showing more options) intensify competition, lowering margins and raising pass-through. Platforms that decrease λ (curating selections) create captive segments, raising margins and lowering pass-through. The same platform design choices that affect price levels also affect who bears the platform's fees.

10. Conclusion

This paper develops a framework for analyzing cost pass-through when equilibrium features a price distribution rather than a price point. The core contribution is a decomposition: pass-through separates into a competition layer, where the exogenous consideration-set structure determines the equilibrium distribution of normalized effective margins, and a curvature layer, where demand elasticity determines how these margins translate into prices. The μ-isomorphism shows that the equilibrium margin distribution depends only on the consideration structure, not on demand curvature or cost levels.

We characterize equilibrium for both symmetric and asymmetric cases. With symmetric firms, the equilibrium margin distribution has a simple representation in terms of the probability generating function for competitor counts. With asymmetric firms under independent consideration, equilibria exhibit a hierarchical structure: the two highest-reach firms share the full support $[\underline{\mu}, 1]$, with only the leader placing an atom at the monopoly margin, while lower-reach firms have truncated supports. Higher-reach firms price higher at each quantile. This pattern carries direct implications for pass-through under unit or linear demand.

The decomposition yields pass-through results that would be difficult to obtain otherwise. We derive closed-form formulas at each quantile of the price distribution, revealing heterogeneity invisible to single-price analysis. Under unit demand, transaction-weighted pass-through takes a particularly simple form: $\tau_i^{\text{trans}} = 1 - 1/K_i$, where $K_i = \int \mu_i(u)^{-1} du$ depends only on the firm's equilibrium margin distribution. With symmetric firms, this reduces to $\tau^{\text{trans}} = 1 - \rho/\bar{H}$, linking market structure directly to incidence without requiring demand estimation. We also characterize when aggregate pass-through can exceed unity: CES demand produces over-shifting at low prices while linear demand does not.

When demand is only partially known, we can still derive robust bounds on pass-through. The universal lower bound $\tau \geq 1 - \mu$ holds for any downward-sloping demand, with unit demand achieving this bound. For specific demand families, tighter bounds apply: linear demand yields $\tau \leq (1 + \sqrt{1 - \mu})/2$, while CES demand can exceed unity depending on the elasticity parameter. These bounds are useful for policy analysis when demand estimation is impractical or unreliable.

The framework suggests a natural empirical strategy: estimate consideration-set patterns from choice data, then use demand estimates to compute pass-through. With symmetric firms and unit demand, the sufficient statistic $\rho/\bar{H}$ simplifies this further, and the identification problem shifts from estimating demand elasticities to measuring consideration sets, which may be easier to recover in some settings.

Several extensions merit future research. Allowing firm-specific costs would capture heterogeneous responses to input price changes. Perhaps most important would be endogenizing the consideration structure. We take consideration sets as primitive, but in practice they emerge as a part of the competitive process: firms invest in advertising to expand their reach, platforms design algorithms that shape which products consumers see, and consumers decide how much to search based on expected gains. Each of these margins responds to cost shocks. When costs rise, firms may advertise more aggressively to steal competitors' captive customers, platforms may adjust ranking algorithms, and consumers may search harder as price dispersion widens. These responses feed back into the consideration structure, potentially amplifying or dampening the direct pass-through effects we characterize. A full analysis of incidence would trace these indirect effects alongside the direct price responses.

Appendix A. §3 Omitted Proofs

A.1. **Proof of Lemma 3.2.** Let $U_i(\mu) := \mu_i q_i^{\mu}(\mu)$ denote firm i's payoff.

Lemma A.1. *Fix a consideration structure $\{\alpha_S\}_{S\subseteq N}$ with $\alpha_S \geq 0$ and $\sum_{S\subseteq N}\alpha_S = 1$ and assume uniform tie-breaking within each S. Then the μ-game admits a mixed-strategy Nash equilibrium.*

Proof. We verify the hypotheses of Dasgupta and Maskin (1986, Theorem 5).

1. Compactness of the strategy sets and boundedness of the payoffs. For each i, the pure strategy set is $A_i := [0,1]$, a closed interval. Moreover, for all $\mu \in [0,1]^n$,

$$0 \leq q_i^{\mu}(\mu) \leq \sum_{S\ni i}\alpha_S \leq 1, \quad \implies \quad 0 \leq U_i(\mu) = \mu_i q_i^{\mu}(\mu) \leq 1,$$

so U_i is bounded.

2. Discontinuities occur only on diagonals. Fix i. If μ satisfies $\mu_i \neq \mu_j$ for all $j \neq i$, then in each set $S \ni i$ either $i \notin M_S(\mu)$ or else $M_S(\mu) = \{i\}$, and this classification is locally constant in a neighborhood of μ. Hence, $q_i^{\mu}(\cdot)$ is locally constant and $U_i(\cdot)$

is continuous at such μ. Therefore, U_i can be discontinuous only when $\mu_i = \mu_j$ for some $j \neq i$, i.e., only on a union of diagonal hyperplanes.

Equivalently, in the notation of Dasgupta and Maskin (1986, Equation 2), we may take $D(i) = n-1$ and, for each $j \neq i$, define the one-to-one continuous map $f_{ij}^1(a_i) := a_i$. Then

$$A^*(i) = \left\{\mu \in [0,1]^n : \exists j \neq i \text{ such that } \mu_j = f_{ij}^1(\mu_i) = \mu_i\right\},$$

and the set of discontinuities of U_i is a subset of $A^*(i)$.

3. The sum of payoffs is continuous (therefore, upper semi-continuous). For any $S \neq \emptyset$, uniform tie-breaking implies that the total payoff generated by mass α_S equals $\alpha_S m_S(\mu)$:

$$\sum_{i\in S} \mu_i \alpha_S \frac{\mathbf{1}\{i \in M_S(\mu)\}}{\left|M_S(\mu)\right|} = \alpha_S\, m_S(\mu) \sum_{i\in M_S(\mu)} \frac{1}{\left|M_S(\mu)\right|} = \alpha_S m_S(\mu).$$

Summing over all $S \subseteq N$ yields

$$\sum_{i\in N} U_i(\mu) = \sum_{S\subseteq N, S\neq\emptyset} \alpha_S m_S(\mu).$$

Because each $m_S(\mu) = \min_{j\in S} \mu_j$ is continuous and there are finitely many sets S, the sum $\sum_{i\in N} U_i(\mu)$ is continuous.

4. U_i is weakly lower semi-continuous in μ_i (Dasgupta and Maskin, 1986, Definition 6). Fix i and fix μ_{-i}. For each $S \ni i$, define $m_{S,-i} := \min_{j\in S\setminus\{i\}} \mu_j$, with the convention $m_{\{i\},-i} := 1$. Consider the one-variable function $t \mapsto U_i(t, \mu_{-i})$ on $[0,1]$.

We claim that for every $t_0 \in (0,1]$, $\liminf_{t\nearrow t_0} U_i(t, \mu_{-i}) \geq U_i(t_0, \mu_{-i})$. To see this, fix $S \ni i$ and examine the S-contribution to U_i. If $t_0 < m_{S,-i}$, then for all t sufficiently close to t_0 from below we still have $t < m_{S,-i}$, so i is the unique minimizer in S and the S-contribution equals $\alpha_S t$, which is continuous at t_0. If $t_0 > m_{S,-i}$, then for all t sufficiently close to t_0 from below we still have $t > m_{S,-i}$, so $i \notin M_S(\cdot)$ and the S-contribution is identically 0 near t_0. If $t_0 = m_{S,-i}$, then at t_0 firm i is tied for the minimum in S, so the S-contribution at t_0 equals

$$\alpha_S t_0 \frac{1}{\left|M_S\left(t_0, \mu_{-i}\right)\right|} \leq \alpha_S t_0.$$

For any $t < t_0$, however, i becomes the unique minimizer in S, so the S-contribution equals $\alpha_S t$. Consequently,

$$\liminf_{t \nearrow t_0} \alpha_S t = \alpha_S t_0 \geq \alpha_S t_0 \frac{1}{\left|M_S\left(t_0, \mu_{-i}\right)\right|}.$$

In all cases, the S-contribution satisfies the desired lower-semicontinuity inequality, *viz.*, $\liminf_{t \nearrow t_0} U_i(t, \mu_{-i}) \geq U_i(t_0, \mu_{-i})$ for all $t_0 \in (0,1]$; and summing over $S \ni i$ yields $\liminf_{t \nearrow t_0} U_i\left(t, \mu_{-i}\right) \geq U_i\left(t_0, \mu_{-i}\right)$. This left-limit inequality establishes Dasgupta and Maskin (1986, Definition 6) with $\lambda = 1$ (full weight on the left-hand lim inf) for all $t_0 \in (0,1]$.

At the left endpoint, $t_0 = 0$, we have $U_i\left(0, \mu_{-i}\right) = 0$ and $U_i \geq 0$, so $\liminf_{t \searrow 0} U_i\left(t, \mu_{-i}\right) \geq 0 = U_i\left(0, \mu_{-i}\right)$, so the condition holds trivially. Thus, U_i satisfies Dasgupta and Maskin (1986, Definition 6) for all $t_0 \in [0,1]$. .

5. Apply Dasgupta and Maskin (1986, Theorem 5). Steps 1-4 verify all hypotheses of Dasgupta and Maskin (1986, Theorem 5). Therefore, the μ-game possesses a mixed-strategy equilibrium. ∎

A.2. **Proof of Theorem 3.3.**

Proof. Recall that firm i's profit at price p is

$$\Pi_i(p;c) = (p-c)x(p)q_i(p) = (1-c)x(1) \cdot \mu(p;c) \cdot q_i(p),$$

with the factor $(1-c)x(1)$ being constant across all firms and prices. Recall also firm i's normalized effective margin at price p:

$$\mu(p;c) \equiv \frac{(p-c)x(p)}{(1-c)x(1)} \in [0,1],$$

and the inverse map $\phi(\mu, c)$, which solves

$$(\phi(\mu,c)-c)x(\phi(\mu,c)) = \mu(1-c)x(1). \tag{A1}$$

Setting, in turn, $\mu = 0$ and $\mu = 1$ in (A1) yield $\phi(0,c) = c$ and $\phi(1,c) = 1$. Moreover ϕ is increasing in μ. Hence, there is a bijection between price-CDFs F_i on $[c,1]$ and μ-CDFs F_i^μ on $[0,1]$ via $F_i^\mu(\mu) = F_i(\phi(\mu,c))$.

For a profile $\mu \in [0,1]^n$ and a nonempty set $S \subseteq N$, define

$$m_S(\mu) := \min_{j\in S} \mu_j, \quad \text{and} \quad M_S(\mu) := \left\{ j \in S : \mu_j = m_S(\mu) \right\}.$$

Under our uniform tie-breaking stipulation, consumers with consideration set S allocate their mass α_S equally across the minimizers $M_S(\mu)$.

Recall that firm i's demand when posting price p is

$$q_i(p) = \sum_{S\ni i} \alpha_S \cdot \mathbb{E}\left[\frac{\mathbf{1}\{i \in M_S(p, p_{-i})\}}{|M_S(p, p_{-i})|} \right].$$

When firm i posts μ and rivals use $(F_j^\mu)_{j\neq i}$, the demand share becomes

$$q_i^\mu(\mu) := \sum_{S\ni i} \alpha_S \cdot \mathbb{E}\left[\frac{\mathbf{1}\{i \in M_S(\mu, \mu_{-i})\}}{\left|M_S(\mu, \mu_{-i})\right|} \right]. \tag{A2}$$

The equilibrium indifference condition in the original game requires constant profit on the support:

$$(p-c)x(p)q_i(p) = \pi_i(c), \quad \forall p \in \operatorname{supp}(F_i).$$

Dividing by $(1-c)x(1)$:

$$\mu(p;c) \cdot q_i(p) = \frac{\pi_i(c)}{(1-c)x(1)} \equiv \tilde{\pi}_i \quad \forall p \in \operatorname{supp}(F_i).$$

Since ϕ is strictly increasing, the lowest-price firm in each consideration set is also the lowest-margin firm, so $q_i(p) = q_i^\mu(\mu(p;c))$. Under the bijection, the indifference condition becomes:

$$\mu \cdot q_i^\mu(\mu) = \tilde{\pi}_i \quad \forall \mu \in \operatorname{supp}(F_i^\mu),$$

which is precisely the equilibrium condition for a game with:

- Unit demand (quantity fixed at 1);

- zero marginal cost;
- ”price” $\mu \in [0,1]$; and
- demand shares computed via ($A2$) using μ-distributions.

Moreover, the transformed equilibrium conditions depend only on $\{\alpha_S\}$ through the demand-share formula. Neither c nor $x(\cdot)$ appears in the μ-space equilibrium, though they determine the mapping ϕ back to prices.

To complete the bijection, we verify that any μ-equilibrium induces an equilibrium in the original game. Suppose $(F_i^\mu)_{i\in N}$ is an equilibrium in the μ-game with constant profits $\tilde{\pi}_i$ on the support. Then, for any p in the support of F_i:

$$\begin{aligned}\Pi_i(p;c) &= (p-c)x(p)q_i(p)\\ &= (1-c)x(1)\cdot\mu(p;c)\cdot q_i(p)\\ &= (1-c)x(1)\cdot\mu(p;c)\cdot q_i^\mu(\mu(p;c))\\ &= (1-c)x(1)\cdot\tilde{\pi}_i\end{aligned}$$

where the third equality uses that $q_i(p) = q_i^\mu(\mu(p;c))$ by construction, and the fourth uses that $\mu(p;c)$ is in the support of F_i^μ when p is in the support of F_i.

For any $p \in [c,1]$ outside the support of F_i, the corresponding $\mu(p;c)$ is outside the support of F_i^μ. Since the map $p \mapsto \mu(p;c)$ is strictly increasing (by Assumption 2.4), each deviation in price space corresponds to a unique deviation in μ-space. Therefore, the ”no profitable deviation” inequality is preserved: if $\mu\cdot q_i^\mu(\mu) \le \tilde{\pi}_i$ for μ off support, then $(p-c)x(p)q_i(p) \le (1-c)x(1)\tilde{\pi}_i = \pi_i(c)$ for the corresponding p off support. Thus, the induced price distributions form an equilibrium with profits $\pi_i(c) = (1-c)x(1)\tilde{\pi}_i$.

Existence of an equilibrium in the pricing game follows from Lemma A.1 (proved above) combined with the bijection established here. ■

Appendix B. §4 Omitted Proofs

B.1. **Proof of Proposition 4.1.**

Proof of Proposition 4.1. In any symmetric equilibrium, all firms use the same μ-distribution F^μ. Let $\underline{\mu} \equiv \inf \operatorname{supp}(F^\mu)$ and $\bar{\mu} \equiv \sup \operatorname{supp}(F^\mu)$. A standard undercutting argument implies that, whenever $0 < \rho < 1$, the distribution F^μ has no atoms on $\left(\underline{\mu}, 1\right]$: if there were an atom of size $a > 0$ at some $\mu_0 \in \left(\underline{\mu}, 1\right]$, then deviating to $\mu_0 - \varepsilon$ would strictly increase the probability of winning any consideration set S with $|S| \geq 2$ in the event that all rivals in $S \setminus \{i\}$ draw μ_0, which occurs with probability $a^{|S|-1} > 0$. The resulting discrete gain in demand dominates the $O(\varepsilon)$ loss in margin for ε small, contradicting optimality. Consequently, ties occur with probability zero on the interior of the support, and the demand share simplifies to the atomless (Demand*) expression. Thus, when firm i posts μ and all rivals use F^μ, the demand share is

$$q^\mu(\mu) = \sum_{S \ni i} \alpha_S \prod_{j \in S \setminus \{i\}} [1 - F^\mu(\mu)] = \sum_{S \ni i} \alpha_S [1 - F^\mu(\mu)]^{|S|-1}.$$

Let $G(\mu) \equiv 1 - F^\mu(\mu)$ denote the complementary CDF. The equilibrium profit at μ is

$$\pi^\mu(\mu) = \mu \cdot q^\mu(\mu) = \mu \cdot \sum_{S \ni i} \alpha_S G(\mu)^{|S|-1}.$$

In any mixed-strategy equilibrium, profits must be constant on $\operatorname{supp}(F^\mu)$.

We first handle degenerate cases. If $\rho = 0$ (no captive customers), then $\alpha_{\{i\}} = 0$ and the unique equilibrium is pure with all firms setting $p = c$ (so $\mu = 0$), yielding zero profits à la standard Bertrand competition. If $\rho = 1$, then $\alpha_{\{i\}} = \sigma$ and consumers who consider firm i never consider any rival, so the unique symmetric equilibrium is pure with $\mu = 1$. Henceforth, assume $0 < \rho < 1$.

We next show the support is a connected interval. Suppose, for contradiction, $\operatorname{supp}(F^\mu)$ were not connected. There would exist a gap (μ_1, μ_2) such that rivals assign zero probability mass to (μ_1, μ_2). On that gap, $F^\mu(\mu)$ (so, $G(\mu)$ too) is constant, so $q^\mu(\mu)$ is constant, while $\pi^\mu(\mu) = \mu q^\mu(\mu)$ is strictly increasing in μ. Therefore, π^μ is strictly increasing on $[\mu_1, \mu_2]$, so $\pi^\mu(\mu_1) < \pi^\mu(\mu_2)$. But the gap endpoints μ_1, μ_2 lie in $\operatorname{supp}(F^\mu)$, so equilibrium indifference requires equal profits at both, giving the contradiction. Thus, $\operatorname{supp}(F^\mu)$ is an interval $\left[\underline{\mu}, \bar{\mu}\right]$.

We now show $\bar{\mu} = 1$. Suppose instead that $\bar{\mu} < 1$. Since F^μ is atomless on $\left(\underline{\mu}, 1\right]$, we have $G(\mu) \to 0$ as $\mu \uparrow \bar{\mu}$. It follows that

$$\lim_{\mu \uparrow \bar{\mu}} q^\mu(\mu) = \sum_{S \ni i} \alpha_S \lim_{\mu \uparrow \bar{\mu}} G(\mu)^{|S|-1} = \alpha_{\{i\}},$$

because $G(\mu)^{|S|-1} \to 0$ for $|S| \geq 2$ and equals 1 for $S = \{i\}$. Hence, the (constant) equilibrium profit satisfies

$$\pi^* = \lim_{\mu \uparrow \bar{\mu}} \pi^\mu(\mu) = \lim_{\mu \uparrow \bar{\mu}} \mu \, q^\mu(\mu) = \bar{\mu} \, \alpha_{\{i\}} < \alpha_{\{i\}}.$$

But deviating to $\mu = 1$ yields profit $1 \cdot \alpha_{\{i\}} = \alpha_{\{i\}}$, because the firm then sells only to captive consumers. This contradicts optimality, so $\bar{\mu} = 1$. Therefore $\operatorname{supp}(F^\mu) = \left[\underline{\mu}, 1\right]$, and equilibrium profit equals

$$\pi^* = 1 \cdot \alpha_{\{i\}} = \rho\sigma,$$

where $\sigma \equiv \sum_{S \ni i} \alpha_S$ and $\rho \equiv \alpha_{\{i\}}/\sigma$.

Indifference on $\left[\underline{\mu}, 1\right]$ implies that for every μ in the support,

$$\mu \cdot \sum_{S \ni i} \alpha_S G(\mu)^{|S|-1} = \rho\sigma.$$

Dividing by σ and defining

$$H(s) \equiv \frac{1}{\sigma} \sum_{S \ni i} \alpha_S s^{|S|-1},$$

we obtain the equilibrium condition

$$\mu \cdot H\left(G(\mu)\right) = \rho, \qquad \text{so} \qquad \mu = \frac{\rho}{H\left(G(\mu)\right)}.$$

With the quantile transformation $u = F^\mu(\mu) = 1 - G(\mu)$, this yields the quantile function

$$\mu(u) = \frac{\rho}{H\left(1 - u\right)}.$$

This formula also pins down the support endpoints. Since $H(1) = \frac{1}{\sigma}\sum_{S \ni i} \alpha_S = 1$ and $H(0) = \frac{1}{\sigma}\alpha_{\{i\}} = \rho$, we have

$$\mu(0) = \frac{\rho}{H(1)} = \rho, \qquad \text{and} \qquad \mu(1) = \frac{\rho}{H(0)} = 1,$$

so $\operatorname{supp}(F^\mu) = [\rho, 1]$.

To verify no profitable deviations exist, note first that profits equal $\pi^* = \rho\sigma$ on $[\rho, 1]$ by construction. If a firm deviates to $\mu < \rho$, then it undercuts rivals surely and wins all consumers who consider it, so $q^\mu(\mu) = \sigma$ and profit is $\mu\sigma < \rho\sigma = \pi^*$. Deviations to $\mu > 1$ are infeasible since the strategy space is $[0,1]$. We conclude that the constructed F^μ is an equilibrium.

Finally, the equilibrium is unique within symmetric strategies. When $0 < \rho < 1$, there is positive probability of facing at least one rival conditional on being considered, so H is strictly increasing on $[0,1]$ (equivalently, if $K \equiv |S| - 1$ under the conditional law $\mathbb{P}(\cdot \mid i \in S)$, then $H(s) = \mathbb{E}\left[s^K\right]$ and $H'(s) = \mathbb{E}\left[K s^{K-1}\right] > 0$ for $s \in (0,1)$). Thus, H is invertible on $[0,1]$, and the indifference condition $\mu H(G(\mu)) = \rho$ uniquely determines $G(\mu)$ (so F^μ too) on $[\rho, 1]$, delivering the unique quantile function $\mu(u) = \rho / H(1-u)$. ■

B.2. **Proof of Proposition 4.3.**

Proof of Proposition 4.3. The proof follows the standard construction for asymmetric price dispersion games (see Narasimhan, 1988).

First we pin down equilibrium profits. Firm 1 guarantees profit α_1 by pricing at $\mu = 1$ (serving only captives). Firm 2 benefits from Firm 1's high price floor ($\underline{\mu} = \rho_1$), guaranteeing rents above its captive share:

$$\pi_1^* = \alpha_1, \quad \text{and} \quad \pi_2^* = \underline{\mu}(\alpha_2 + \alpha_{12}) = \rho_1(\alpha_2 + \alpha_{12}).$$

Note that $\pi_2^* > \alpha_2$ since $\rho_1 > \rho_2$. Moreover, from $\rho_i = \alpha_i/(\alpha_i + \alpha_{12})$, we have $\frac{\alpha_i}{\alpha_{12}} = \frac{\rho_i}{1-\rho_i}$ and $\frac{\alpha_i + \alpha_{12}}{\alpha_{12}} = \frac{1}{1-\rho_i}$.

Next, we use the two firms' indifference conditions to back out the cdfs. When firm 1 posts μ and firm 2 uses CDF F_2^μ, firm 1's profit is:

$$\mu[\alpha_1 + \alpha_{12}(1 - F_2^\mu(\mu))] = \alpha_1 \quad \Longrightarrow \quad 1 - F_2^\mu(\mu) = \frac{\rho_1}{1-\rho_1}\left(\frac{1-\mu}{\mu}\right).$$

When firm 2 posts μ and firm 1 uses CDF F_1^μ, firm 2's profit is:

$$\mu[\alpha_2 + \alpha_{12}(1 - F_1^\mu(\mu))] = \pi_2^* = \rho_1(\alpha_2 + \alpha_{12}) \quad \Longrightarrow \quad 1 - F_1^\mu(\mu) = \frac{1}{1-\rho_2}\left(\frac{\rho_1}{\mu} - \rho_2\right).$$

Finally, we pin down the support. At the lower bound, $F_1^\mu(\underline{\mu}) = 0$ requires $\frac{\rho_1}{\underline{\mu}} - \rho_2 = 1 - \rho_2$, yielding $\underline{\mu} = \rho_1$. At the upper bound, $F_2^\mu(1^-) = 1$ confirms firm 2 mixes continuously to 1. However, $F_1^\mu(1) = 1 - \frac{\rho_1 - \rho_2}{1-\rho_2} < 1$, confirming the atom Δ_1. ■

B.3. **Proof of Proposition 4.5.**

Proof of Proposition 4.5. Fix $j \in \{1, \dots, n\}$ and fix rivals' μ-CDFs $\left(F_i^\mu\right)_{i \neq j}$. If the rivals' CDFs are atomless at μ (so ties occur with probability zero), firm j's demand share when it posts μ is

$$q_j^\mu(\mu) = \lambda_j \prod_{i \neq j} \Big((1 - \lambda_i) + \lambda_i \big(1 - F_i^\mu(\mu)\big)\Big) = \lambda_j \prod_{i \neq j} \big[1 - \lambda_i F_i^\mu(\mu)\big],$$

and profit is

$$\Pi_j(\mu) = \mu q_j^\mu(\mu) = \mu \lambda_j \prod_{i \neq j} \big[1 - \lambda_i F_i^\mu(\mu)\big].$$

In what follows we will only apply this formula at $\mu \in \left(\underline{\mu}, 1\right)$, where the candidate equilibrium CDFs are continuous. If a deviation ever selected a point at which some rival had an atom (here, the only such point is $\mu = 1$ for firm 1), the deviator could weakly improve by undercutting by an arbitrarily small $\varepsilon > 0$, because ties are split while a strict undercut wins all tied consumers.

Define $\underline{\mu}$, $(\bar{\mu}_k)_{k=1}^{n+1}$, C_m, and the piecewise function $\Gamma(\cdot)$ as in the statement. Define the candidate profile by

$$F_j^\mu(\mu) = 0 \quad \text{for } \mu < \underline{\mu},$$

and for $\mu \in [\bar{\mu}_{m+1}, \bar{\mu}_m]$,

$$F_j^\mu(\mu) := \begin{cases} \Gamma(\mu)/\lambda_j, & j \le m, \\ 1, & j > m. \end{cases}$$

Finally set $F_1^\mu(1) = 1$ and $F_j^\mu(1) = 1$ for each $j \ge 2$.

These are well-defined CDFs. On each interval $[\bar{\mu}_{m+1}, \bar{\mu}_m]$, $\Gamma(\mu) = 1 - \left(\underline{\mu}/(\mu C_m)\right)^{1/(m-1)}$ is weakly increasing in μ, hence, so is $F_j^\mu(\mu)$. Moreover, $\Gamma(\underline{\mu}) = 0$ (here $m = n$ and $C_n = 1$), so $F_j^\mu(\underline{\mu}) = 0$ for all j. At each cutoff $\bar{\mu}_{m+1}$ (for $m \in \{2, \ldots, n-1\}$), the left- and right-limit values coincide:

$$1 - \left(\frac{\underline{\mu}}{\bar{\mu}_{m+1} C_{m+1}}\right)^{1/m} = 1 - \left(\frac{\underline{\mu}}{\bar{\mu}_{m+1} C_m}\right)^{1/(m-1)} = \lambda_{m+1},$$

where we used $C_m = (1 - \lambda_{m+1}) C_{m+1}$ and the definition of $\bar{\mu}_{m+1}$. Thus, each F_j^μ is continuous on $\left(\underline{\mu}, 1\right)$. In particular, ties occur with probability zero for any $\mu \in \left(\underline{\mu}, 1\right)$, so the demand formula above applies on all interior support points.

Now fix $m \in \{2, \ldots, n\}$ and $\mu \in [\bar{\mu}_{m+1}, \bar{\mu}_m) \subseteq \left[\underline{\mu}, 1\right)$. For each active firm $i \le m$, $\lambda_i F_i^\mu(\mu) = \Gamma(\mu)$, while for each $i > m$, $F_i^\mu(\mu) = 1$. Accordingly, for any active firm $j \le m$,

$$\Pi_j(\mu) = \mu \lambda_j \left(\prod_{\substack{i \le m \\ i \ne j}} [1 - \Gamma(\mu)] \right) \left(\prod_{i > m} (1 - \lambda_i) \right) = \mu \lambda_j C_m (1 - \Gamma(\mu))^{m-1}.$$

By the definition of $\Gamma(\mu)$ on this interval,

$$(1 - \Gamma(\mu))^{m-1} = \frac{\underline{\mu}}{\mu C_m},$$

so $\Pi_j(\mu) = \lambda_j \underline{\mu}$, which is constant in μ and does not depend on m. This establishes indifference on every interval where firm j is active and yields the equilibrium profit claim $\pi_j^* = \lambda_j \underline{\mu}$.

The cutoff points $\bar{\mu}_k$ are pinned down by where each firm's CDF reaches 1. On the lowest interval $\left[\underline{\mu}, \bar{\mu}_n\right]$ we have $m = n$ and $C_n = 1$, so $\Gamma(\underline{\mu}) = 0$ and so $F_j^\mu(\underline{\mu}) = 0$ for all j, establishing the common lower bound. Take $k \in \{3, \ldots, n\}$. Firm k reaches

$F_k^\mu(\mu) = 1$ precisely when $\Gamma(\mu) = \lambda_k$. On the interval where the active set is $\{1, \ldots, k\}$, we have $m = k$ and $C_k = \prod_{h=k+1}^n (1-\lambda_h)$, so $\Gamma(\bar{\mu}_k) = \lambda_k$ implies

$$\lambda_k = 1 - \left(\frac{\underline{\mu}}{\bar{\mu}_k C_k}\right)^{1/(k-1)},$$

equivalently,

$$\bar{\mu}_k = \frac{\underline{\mu}}{C_k (1-\lambda_k)^{k-1}} = \frac{\prod_{h=2}^{k-1}(1-\lambda_h)}{(1-\lambda_k)^{k-2}},$$

as claimed (and $\bar{\mu}_2 = 1$ by definition).

On the top interval $[\bar{\mu}_3, 1]$ the active set is $\{1, 2\}$, so $m = 2$ and $C_2 = \prod_{h=3}^n (1-\lambda_h)$. Then

$$\Gamma(\mu) = 1 - \frac{\underline{\mu}}{\mu C_2}.$$

Taking $\mu \uparrow 1$ gives $\Gamma(1-) = 1 - \underline{\mu}/C_2 = 1 - (1-\lambda_2) = \lambda_2$, whence

$$F_2^\mu(1-) = \frac{\Gamma(1-)}{\lambda_2} = 1 \qquad \text{and} \qquad F_1^\mu(1-) = \frac{\Gamma(1-)}{\lambda_1} = \frac{\lambda_2}{\lambda_1} < 1.$$

Thus, firm 1 has an atom at $\mu = 1$ of size $\Delta_1 = 1 - \lambda_2/\lambda_1$, while firm 2 has no atom at $\mu = 1$. Since firms $k \geq 3$ have $\bar{\mu}_k < 1$, they also have no atom at $\mu = 1$.

It remains to show that no firm can profitably deviate. We already showed that any μ in firm j's support yields profit $\lambda_j \underline{\mu}$.

If $\mu < \underline{\mu}$, then $F_i^\mu(\mu) = 0$ for all rivals i, so

$$\Pi_j(\mu) = \mu \lambda_j < \underline{\mu} \lambda_j = \pi_j^*.$$

Now fix $\mu \in [\bar{\mu}_{m+1}, \bar{\mu}_m]$ and suppose $j > m$. Then for each active firm $i \leq m$, $\lambda_i F_i^\mu(\mu) = \Gamma(\mu)$, hence $\left[1 - \lambda_i F_i^\mu(\mu)\right] = 1 - \Gamma(\mu)$; and for each inactive firm $i > m$, $F_i^\mu(\mu) = 1$, hence, $\left[1 - \lambda_i F_i^\mu(\mu)\right] = 1 - \lambda_i$. Therefore,

$$\Pi_j(\mu) = \mu \lambda_j \prod_{i \neq j} \left[1 - \lambda_i F_i^\mu(\mu)\right] = \mu \lambda_j (1 - \Gamma(\mu))^m \prod_{\substack{i > m \\ i \neq j}} (1 - \lambda_i).$$

Using $C_m = \prod_{i>m}(1-\lambda_i)$, we have

$$\prod_{\substack{i>m \\ i\neq j}}(1-\lambda_i) = \frac{C_m}{1-\lambda_j},$$

so

$$\Pi_j(\mu) = \mu\,\lambda_j\,\frac{C_m}{1-\lambda_j}\,(1-\Gamma(\mu))^m.$$

On $[\bar{\mu}_{m+1}, \bar{\mu}_m]$, the definition of $\Gamma(\mu)$ implies

$$(1-\Gamma(\mu))^{m-1} = \frac{\underline{\mu}}{\mu\,C_m},$$

hence,

$$\Pi_j(\mu) = \mu\,\lambda_j\,\frac{C_m}{1-\lambda_j}\,(1-\Gamma(\mu))^{m-1}\,(1-\Gamma(\mu)) = \lambda_j\,\underline{\mu}\,\frac{1-\Gamma(\mu)}{1-\lambda_j}.$$

Moreover, $\Gamma(\mu)$ is increasing in μ and satisfies $\Gamma(\bar{\mu}_{m+1}) = \lambda_{m+1}$, so for all $\mu \in [\bar{\mu}_{m+1}, \bar{\mu}_m]$,

$$\Gamma(\mu) \geq \Gamma(\bar{\mu}_{m+1}) = \lambda_{m+1}.$$

Since $j > m$ and $\lambda_2 \geq \cdots \geq \lambda_n$, we have $\lambda_j \leq \lambda_{m+1}$, hence $\Gamma(\mu) \geq \lambda_j$, which implies

$$\frac{1-\Gamma(\mu)}{1-\lambda_j} \leq 1 \quad \Longrightarrow \quad \Pi_j(\mu) \leq \lambda_j\,\underline{\mu} = \pi_j^*.$$

Deviations to $\mu = 1$ cannot do better than taking $\mu \uparrow 1$ from below, because $\mu = 1$ creates a tie with firm 1's atom and ties are split. Finally, $\mu > 1$ is infeasible. This establishes that no firm has a profitable deviation, so the constructed profile is an equilibrium.

Uniqueness follows from the same logic that generates the construction. On the lowest interval $\left[\underline{\mu}, \bar{\mu}_n\right]$, all firms are active and each must be indifferent, so $\Pi_j(\mu) = \pi_j^*$ for all j and all μ in that interval. Evaluating at $\mu = \underline{\mu}$ (where $F_i^\mu(\underline{\mu}) = 0$ for all i since there are no atoms at the lower bound) implies $\pi_j^* = \lambda_j\underline{\mu}$. Comparing the indifference conditions for any two firms j,k on this same interval then forces $\lambda_j F_j^\mu(\mu) = \lambda_k F_k^\mu(\mu)$. Writing the common value as $\Gamma(\mu)$ pins down $\Gamma(\mu)$ on $\left[\underline{\mu}, \bar{\mu}_n\right]$ via $\mu(1-\Gamma(\mu))^{n-1} = \underline{\mu}$, hence, $\Gamma(\mu) = 1 - \left(\underline{\mu}/\mu\right)^{1/(n-1)}$. The cutoff $\bar{\mu}_n$ is then uniquely characterized by $\Gamma(\bar{\mu}_n) = \lambda_n$. Proceeding inductively, on each higher interval where exactly firms $\{1, \dots, m\}$ remain active, the boundary condition at $\bar{\mu}_{m+1}$ and indifference again force a common $\Gamma(\mu)$ and pin it down by $\mu C_m (1 - \Gamma(\mu))^{m-1} = \underline{\mu}$, and the next cutoff $\bar{\mu}_m$ is uniquely determined by $\Gamma(\bar{\mu}_m) = \lambda_m$. Finally, $\Gamma(1-) = \lambda_2$ pins down $F_1^{\mu}(1-) = \lambda_2/\lambda_1$, so the only possible atom at $\mu = 1$ is $\Delta_1 = 1 - \lambda_2/\lambda_1$, and no other atoms are compatible with indifference because they can be profitably undercut. We conclude that the equilibrium CDFs are uniquely determined. ■

B.4. **Proof of Theorem 7.3.**

Proof. We provide complete derivations for each case.

Case 1: Linear demand $x(p) = 1 + b(1-p)$, $b \in [0, 1/d]$.

The effective margin equation $(p-c)(1+b(1-p)) = \mu d$ becomes:

$$(p-c)(1+b(1-p)) = \mu d$$

Expanding: $(p-c) + b(p-c)(1-p) = \mu d$. Let $m = p - c$, so $p = c + m$ and $1 - p = 1 - c - m = d - m$. Then:

$$m + bm(d-m) = \mu d$$

$$m(1 + bd - bm) = \mu d$$

$$bm^2 - m(1+bd) + \mu d = 0$$

Using the quadratic formula:

$$m = \frac{(1+bd) \pm \sqrt{(1+bd)^2 - 4b\mu d}}{2b}$$

For the relevant root (taking the smaller value to ensure $p \le 1$):

$$p - c = \frac{(1+bd) - \sqrt{(1+bd)^2 - 4b\mu d}}{2b}$$

For pass-through, we need $\phi_c(\mu, c)$. The invertibility condition requires:

$$\frac{\partial}{\partial p}[(p-c)(1+b(1-p))] = 1 + b(1-p) - b(p-c) = 1 + b(1 - 2p + c) > 0$$

This gives $p < \frac{1+b(1+c)}{2b}$. At the boundary $b = 1/d$, the condition becomes $p < 1$.

At $b = 1/d$, the quadratic simplifies. Setting $b = 1/d$:

$$\frac{1}{d}m^2 - m(1+1) + \mu d = 0$$

$$m^2 - 2md + \mu d^2 = 0$$

$$m = d(1 - \sqrt{1-\mu})$$

Thus $p = c + d(1 - \sqrt{1-\mu}) = 1 - d\sqrt{1-\mu}$.

For pass-through at $b = 1/d$, we use implicit differentiation of $(p-c)x(p) = \mu d$ with b held fixed:

$$\phi_c(\mu, c) = \frac{x(p)(1-p)}{d[x(p) + (p-c)x'(p)]}$$

At $b = 1/d$, $x'(p) = -b = -1/d$, and substituting the equilibrium values:

$$\phi_c(\mu, c) = \frac{(1+\sqrt{1-\mu})d\sqrt{1-\mu}}{d[(1+\sqrt{1-\mu}) - d(1-\sqrt{1-\mu})/d]} = \frac{1+\sqrt{1-\mu}}{2}$$

Case 2: Constant semi-elasticity $x(p) = e^{\beta(1-p)}$, $\beta \in [0, 1/d]$.

The effective margin equation becomes:

$$(p-c)e^{\beta(1-p)} = \mu d e^{\beta}$$

$$(p-c) = \mu d e^{\beta p}$$

This is a transcendental equation. Taking logs:

$$\ln(p-c) = \ln(\mu d) + \beta p$$

Define W as the Lambert W function (satisfying $W(z)e^{W(z)} = z$). We can write:

$$p = \frac{1}{\beta}[W(\beta \mu d e^{\beta c}) + \beta c]$$

At the boundary $\beta = 1/d$, the invertibility condition becomes critical. The derivative:

$$\frac{\partial}{\partial p}[(p-c)e^{\beta(1-p)}] = e^{\beta(1-p)}[1 - \beta(p-c)]$$

At $\beta = 1/d$ and using the equilibrium condition $(p - c) = \mu d e^{\beta p}$:

$$1 - \beta(p - c) = 1 - \frac{\mu d e^{\beta p}}{d} = 1 - \mu e^{p/d}$$

For this to be positive requires $\mu < e^{-p/d}$.

For pass-through at $\beta = 1/d$, using the general formula:

$$\phi_c(\mu, c) = \frac{e^{(1-p)/d}(1 - p)}{d[e^{(1-p)/d}(1 - \mu e^{p/d})]}$$

In the limit as we approach the boundary where $1 - \mu e^{p/d} \to 0$, the pass-through approaches 1.

Case 3: Constant elasticity $x(p) = p^{-\eta}$.

For $\eta = 1$: $(p - c)p^{-1} = \mu d$ gives $p = c/(1 - \mu d)$, and:

$$\phi_c(\mu, c) = \frac{1 - p}{d(1 - \varepsilon)} = \frac{1 - \mu}{(1 - \mu d)^2}$$

Boundary behavior: At the boundaries ($b = 1/d$ for linear, $\beta = 1/d$ for exponential), the invertibility assumption (Theorem 2.4) becomes an equality at $p = 1$. These are knife-edge cases where the strict inequality becomes weak. The pass-through formulas remain valid by continuity, taking limits as we approach these boundaries from the interior of the parameter space. ■

References

Jason Abaluck and Abi Adams-Prassl. What do consumers consider before they choose? identification from asymmetric demand responses. *Quarterly Journal of Economics*, 136(3):1611–1663, 2021. doi: 10.1093/qje/qjab008.

Takanori Adachi and Takeshi Ebina. Cost pass-through and inverse demand curvature in vertical relationships with upstream and downstream competition. *Economics Letters*, 124(3):465–468, 2014. doi: 10.1016/j.econlet.2014.07.008.

Takanori Adachi and Michal Fabinger. Pass-through, welfare, and incidence under imperfect competition. *Journal of Public Economics*, 211:104659, 2022. doi: 10.1016/j.jpubeco.2021.104589.

Brian C Albrecht. Price competition and the use of consumer data. Working Paper, 2020. URL https://briancalbrecht.github.io/albrecht_price_competition_consumer_data.pdf.

Mark Armstrong. Ordered consumer search. *Journal of the European Economic Association*, 15(5):989–1024, 2017. doi: 10.1093/jeea/jvx017.

Mark Armstrong and John Vickers. Patterns of competitive interaction. *Econometrica*, 90(1):153–191, 2022. doi: 10.3982/ECTA18937.

Dirk Bergemann, Benjamin Brooks, and Stephen Morris. Search, information, and prices. *Journal of Political Economy*, 129(8):2275–2319, 2021. doi: 10.1086/714443.

Kenneth Burdett and Kenneth L Judd. Equilibrium price dispersion. *Econometrica*, 51(4):955–969, 1983. doi: 10.2307/1912045.

Victor Chernozhukov, Iván Fernández-Val, and Blaise Melly. Inference on counterfactual distributions. *Econometrica*, 81(6):2205–2268, 2013. doi: 10.3982/ECTA10582.

Partha Dasgupta and Eric Maskin. The existence of equilibrium in discontinuous economic games, i: Theory. *Review of Economic Studies*, 53(1):1–26, 1986. doi: 10.2307/2297588.

Matthew Elliott, Andrea Galeotti, Andrew Koh, and Wenhao Li. Market segmentation through information. Working Paper, 2021.

Glenn Ellison and Sara Fisher Ellison. Search, obfuscation, and price elasticities on the internet. *Econometrica*, 77(2):427–452, 2009. doi: 10.3982/ECTA5708.

Kai Fischer, Simon Martin, and Philipp Schmidt-Dengler. Indirect taxation in consumer search markets: The case of retail fuel. CEPR Discussion Paper No. 19095, 2024.

Luke Garrod, Ruochen Li, Antonio Russo, and Chris M. Wilson. Understanding cost pass-through when prices are dispersed. CESifo Working Paper No. 11635, 2024.

Martin Gaynor, Carol Propper, and Stephan Seiler. Free to choose? reform, choice, and consideration sets in the english national health service. *American Economic Review*, 106(11):3521–3557, 2016. doi: 10.1257/aer.20121532.

Michelle Sovinsky Goeree. Limited information and advertising in the u.s. personal computer industry. *Econometrica*, 76(5):1017–1074, 2008. doi: 10.3982/ECTA4158.

Rafael R Guthmann. Price dispersion in dynamic competition. *Economic Theory*, 78:1203–1232, 2024. doi: 10.1007/s00199-024-01601-9.

Rafael R Guthmann. Imperfect competition as a result of unawareness. *Review of Industrial Organization*, 67:35–53, 2025. doi: 10.1007/s11151-024-10006-w.

Rafael R. Guthmann and Brian C. Albrecht. Market microstructure and informational complexity. Working Paper, 2025.

Sven Heim. Asymmetric cost pass-through and consumer search: Empirical evidence from online platforms. *Quantitative Marketing and Economics*, 19(2):227–260, 2021. doi: 10.1007/s11129-021-09233-2.

Elisabeth Honka, Ali Hortaçsu, and Maria Ana Vitorino. Advertising, consumer awareness, and choice: Evidence from the u.s. banking industry. *RAND Journal of Economics*, 48(3):611–646, 2017. doi: 10.1111/1756-2171.12188.

Elisabeth Honka, Ali Hortaçsu, and Matthijs Wildenbeest. Empirical search and consideration sets. In *Handbook of the Economics of Marketing*, volume 1, pages 193–257. Elsevier, 2019. doi: 10.1016/bs.hem.2019.05.002.

Zi Yang Kang and Shoshana Vasserman. Robustness measures for welfare analysis. *American Economic Review*, 115(8):2449–2487, 2025. URL https://www.aeaweb.org/articles?id=10.1257/aer.20220673.

Roger Koenker. *Quantile Regression*. Cambridge University Press, 2005.

Justin Marion and Erich Muehlegger. Fuel tax incidence and supply conditions. *Journal of Public Economics*, 95(9-10):1202–1212, 2011. doi: 10.1016/j.jpubeco.2011.04.003.

R Preston McAfee. Endogenous availability, cartels, and merger in an equilibrium price dispersion. *Journal of Economic Theory*, 62(1):24–47, 1994. doi: 10.1006/jeth.1994.1002.

Nathan H Miller, Matthew Osborne, and Gloria Sheu. Pass-through in a concentrated industry: Empirical evidence and regulatory implications. *RAND Journal of Economics*, 48(1):69–93, 2017. doi: 10.1111/1756-2171.12168.

Felix Montag, Robin Mamrak, Alina Sagimuldina, and Monika Schnitzer. Imperfect price information, market power, and tax pass-through. Working Paper, 2023. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4622647.

Chakravarthi Narasimhan. Competitive promotional strategies. *Journal of Business*, 61(4):427–449, 1988. doi: 10.1086/296442.

Jesse Perla, Laurent Cavenaile, Murat Alp Celik, and Pau Roldan-Blanco. A model of product awareness and industry life cycles. Working Paper, 2023.

Andrew Rhodes and Jidong Zhou. Consumer search and retail market structure. *Management Science*, 65(6):2607–2623, 2019. doi: 10.1287/mnsc.2018.3058.

Andrew Rhodes, Makoto Watanabe, and Jidong Zhou. Multiproduct intermediaries. *Journal of Political Economy*, 129(2):421–464, 2021. doi: 10.1086/711917.

Robert W Rosenthal. A model in which an increase in the number of sellers leads to a higher price. *Econometrica*, 48(6):1575–1579, 1980. doi: 10.2307/1912828.

Samuel Stolper. Who bears the burden of energy taxes? the role of local pass-through. Working Paper, 2017.

Hal R Varian. A model of sales. *American Economic Review*, 70(4):651–659, 1980. URL https://www.jstor.org/stable/1803562.

E Glen Weyl. A price theory of multi-sided platforms. *American Economic Review*, 100(4):1642–1672, 2010. doi: 10.1257/aer.100.4.1642.

E Glen Weyl and Michal Fabinger. Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy*, 121(3):528–583, 2013. doi: 10.1086/670401.