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? We decompose the problem into two layers. In the competition layer, consumers' consideration sets determine equilibrium distributions of normalized margins. In the curvature layer, demand elasticity maps these margins into prices and pass-through rates. We prove the pricing game is strategically equivalent to a game over normalized margins, so equilibrium margin distributions are invariant to demand and costs. This separation yields closed-form pass-through formulas at each quantile of the price distribution, robust bounds across demand specifications, and sharp comparative statics linking market structure to incidence.

1. Introduction

Empirical studies consistently document substantial price variation for homogeneous products.[1] 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 margins translate into prices and pass-through rates. Our key result, Theorem 3.3, establishes that **the equilibrium margin distribution depends only on market**

BA: International Center for Law & Economics. Email: mail@briancalbrecht.com. MW: Arizona State University. Email: mark.whitmeyer@gmail.com. We thank Rani Spiegler and Joseph Whitmeyer for their feedback. refine.ink helped us improve this paper.

[1]For evidence on dispersion in relatively homogeneous retail markets, see Sorensen (2000). For evidence from online price-comparison settings, see Baye et al. (2004).

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 fundamentals matter, but they interact with a distribution of markups rather than a single markup.[2] 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, including

[2]For search or limited-information models where equilibrium markup distributions are central (even with homogeneous goods), see, e.g., Menzio (2024a); Kaplan et al. (2019); Menzio and Trachter (2015); Burdett and Menzio (2018).

search costs, informational frictions, platform algorithms, and behavioral inattention.[3] Our results apply regardless of the underlying source.

The isomorphism we identify reformulates 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 bijection 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.

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, so incidence can be estimated 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 in this case. We derive bounds on pass-through that hold across all demand specifications, useful for policy analysis when demand is only partially known.

[3]For models where platforms/intermediaries shape consumer consideration and can steer or limit access, see Nocke and Rey (2024); Dinerstein et al. (2018); Hagiu and Jullien (2011); Teh and Wright (2022); Teh et al. (2024). For search-engine monetization or position effects, see Eliaz and Spiegler (2011c); Athey and Ellison (2011).

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 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 results make the framework applicable to structural estimation and policy counterfactuals.

1.1. **Road map.** The rest of the paper proceeds as follows. §2 sets up the paper; §3 develops the central μ-isomorphism result; and §4 characterizes the resulting margin distributions for symmetric markets, asymmetric duopoly, and independent consideration. §5 then uses the isomorphism to derive quantile pass-through formulas, while §6 converts these posted-price results into transaction-weighted pass-through statistics. §7 derives robust pass-through bounds when demand is only partially specified, and §8 develops comparative statics linking margin orderings to price and pass-through orderings. §9 discusses applications to gasoline markets and online retail. §10 extends the framework to endogenous consideration. All proofs omitted from the main text lie in the appendices.

1.2. **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.[4] Recent work extends these foundations in several directions: Guthmann (2024, 2025); Guthmann and Albrecht (2025)

[4]Other classic foundations include Salop and Stiglitz (1977); Butters (1977); Stahl (1989); Janssen and Moraga-González (2004). Baye and Morgan (2001); Baye et al. (2004) diagnose dispersion with information gatekeepers and price-comparison sites.

analyze dynamic settings, 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 adopt their consideration-set framework and equilibrium concept. They note that their equilibrium analysis is unaffected by downward-sloping demand, provided revenue is increasing up to the monopoly price (their footnote 7). McAfee (1994) makes a similar observation (his Remark 1). Neither paper formalizes this invariance. We provide the first formal statement and proof via the μ-isomorphism, and show that the resulting decomposition yields tractable pass-through formulas. Whereas Armstrong and Vickers focus on welfare and market structure with unit demand, we characterize pass-through with standard downward-sloping demand, deriving formulas that apply to any consideration structure.

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.[5] Their approach has been extended to vertical markets (Adachi and Ebina, 2014), welfare analysis (Adachi and Fabinger, 2022), multi-product firms (Armstrong and Vickers, 2023), and platform fees (Weyl, 2010).[6] 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 consideration structure rather than 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

[5] Bulow and Pfleiderer (1983); Seade (1985); Delipalla and Keen (1992); Anderson et al. (2001) provide earlier analyses of incidence under imperfect competition.

[6] Other canonical two-sided platform pricing frameworks are Rochet and Tirole (2003, 2006); Armstrong (2006). Rysman (2009) surveys this literature. Wang and Wright (2020, 2025) explore search-platform environments with marketplace commissions with a view toward regulation.

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.[7]

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). We characterize pass-through at each quantile of the price distribution, recovering heterogeneity that single-price models collapse.

2. Model

There are n firms, indexed by $N = \{1, \ldots, n\}$. The demand side consists of a unit mass of consumers, partitioned by their consideration sets.[8]

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 $\sum_{S \subseteq N} \alpha_S = 1$.[9]

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.[10]

Example 2.1 (Random Search)**.** If each consumer samples each firm independently with probability $\lambda \in (0,1)$, then $\alpha_S = \lambda^{|S|}(1-\lambda)^{n-|S|}$, so consideration is generated by i.i.d. Bernoulli sampling and $|S|$ is binomially distributed.

[7] Recently, Menzio (2023, 2024b) presents compelling explorations of search-theoretic imperfect competition and markup distributions.

[8] 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.

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

[10] 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).

Example 2.2 (Spatial Markets)**.** Consider n equally-spaced firms located on a circle. If consumers are distributed uniformly around the circle and observe only their $1 \leq k < n$ nearest neighbors, then $\alpha_S = 1/n$ if S consists of k consecutive firms and 0 otherwise.

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.[11] 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:

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,

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

[11] Throughout, we restrict attention without loss of generality to firms with strictly positive reach.

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

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

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.

Firm i's profit from posting price p is

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

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 $(p - c)x(p)$, has strictly positive derivative on $[c, 1]$: $x(p) + (p - c)x'(p) > 0$ 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. For the standard demand classes considered below, this assumption is equivalent to the effective margin $(p - c)x(p)$ being 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 $(p-c)p^{-\eta}$, with derivative $p^{-\eta-1}[(1-\eta)p + \eta c]$. Consequently, Assumption 2.4 holds automatically for $\eta \in [0,1)$, holds for $\eta = 1$ whenever $c > 0$ (while $c = 0$ makes $p \mapsto (p-c)x(p)$ constant on $(0,1]$), and for $\eta > 1$ holds on $[c,1]$ if and only if $c > 1 - 1/\eta$. When $\eta > 1$ and $c \in (0, 1-1/\eta]$, the effective margin is maximized at $\hat{p}(c) = \eta c/(\eta-1) \in (c,1]$. The isomorphism can be extended in this case by normalizing margins by $\bar{m}(c) \equiv \max_{p\in[c,1]}(p-c)x(p)$ rather than $(1-c)x(1)$, and restricting attention to prices $p \le \hat{p}(c)$, since any price above $\hat{p}(c)$ yields the same effective margin as a lower price and a weakly lower demand share. The competition layer is unchanged—what changes is the curvature layer, since $\bar{m}'(c) = -x(\hat{p}(c))$ 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 standard demand is isomorphic to a unit-demand game after a change of variables. Prior work has noted this invariance informally. Armstrong and Vickers (2022) observe in their footnote 7 that their analysis "is not affected if each consumer has a downward-sloping demand function $x(p)$, provided revenue $px(p)$ is an increasing function up to the monopoly price," and McAfee (1994) notes that "there is no difference in profits between different demand curves with the same monopoly profits" (Remark 1). Neither paper formalizes this observation. We do so here via a change of variables that makes the invariance precise and yields a decomposition useful for pass-through analysis. Firms compete over effective margin per customer served—how much profit they extract from each transaction. Once we

normalize these margins appropriately, the strategic problem becomes identical regardless of the underlying demand function.

3.1. **Normalized Effective Margins.** To see the isomorphism, start with what firms 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 strategic trade-off is the same. Higher margins mean higher profit per customer but lower probability of winning customers.

Definition 3. The *normalized effective margin* at price p is:

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

The numerator is firm profit per served consumer. The denominator is the profit per served consumer when pricing at the maximum price. 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 via the maximum price. Under Assumption 2.4, $\mu(p;c)$ is strictly increasing in p on the equilibrium support, ensuring a bijection between prices and normalized margins.

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

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

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 inverse map $\phi(\mu,c) \in [c,1]$ satisfying* (Inverse Map).

3.2. **The Main Isomorphism Result.** Once we reformulate the game in terms of normalized effective margins μ, the demand-curvature dependence vanishes. Equilibrium depends only on the consideration structure. Demand curvature and costs

matter only for translating these margin distributions back into price distributions.

We define the margin game firms take part in as follows.

Definition 4. 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 verifies the hypotheses of Dasgupta and Maskin (1986, Theorem 5).

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 $\Phi(F_i)(\mu) := F_i(\phi(\mu, c))$, so that $F_i^\mu = \Phi(F_i)$, is a bijection between equilibria of the pricing game and equilibria of the margin game.[12]

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

The normalized effective margin $\mu(p;c)$ gives a bijection between prices and margins, and payoffs in the pricing game are proportional to payoffs in the margin

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

game with a common scaling factor $(1-c)x(1) > 0$, *which does not affect best-response comparisons*.

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 clarifies pass-through analysis. Pass-through decomposes 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 simplifies the analysis. Rather than solving different equilibria for each $(c, x(\cdot))$ combination, we solve once in μ-space and apply different transformations. 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.

3.3. **Economic Intuition.** The normalized effective margin μ captures the tradeoff 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 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 yields a decomposition for merger analysis. Any merger changes the consideration structure α_S, inducing new equilibria 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 $\phi(\mu^{\text{post}}(u), c) - \phi(\mu^{\text{pre}}(u), c)$. 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.[13]

4. Equilibrium Characterization

Having established the isomorphism, we now characterize equilibrium in μ-space. The cases we consider—symmetric firms, asymmetric duopoly, and independent consideration—all fall within the "symmetric interactions" class of Armstrong and Vickers (2022) (which requires symmetry of the consideration structure's interaction parameters, not of firms' reaches), who prove existence, uniqueness, and the structural properties of equilibrium (nested supports, common minimum price, profits proportional to reach). We reformulate their equilibria in terms of normalized margins μ and derive the explicit quantile functions used below.

4.1. **Symmetric Firms.** When the consideration structure treats all firms symmetrically (for instance, when consumers randomly sample firms or when firms are

[13]For a tractable multiproduct oligopoly approach via aggregative games (with merger applications), see, in particular, Nocke and Schütz (2018, 2025). For complementary merger-approximation tools, see Jaffe and Weyl (2013).

arranged symmetrically in geographic or product space), equilibrium has a simple form: symmetric competition leads to a common distribution of market power.

Definition 5. 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 measures the degree of market power arising from limited consideration. When ρ is high, many consumers are captive to individual firms; 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 $i \in S$.*

$H(s)$ 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.

In this case, firm i's reach is $\sigma_i = \lambda$ and captive mass is $\alpha_{\{i\}} = \lambda(1-\lambda)^{n-1}$. The captive-to-reach ratio is $\rho = \alpha_{\{i\}}/\sigma_i = (1-\lambda)^{n-1}$, and the PGF of $|S|-1$ conditional on $i \in S$ is $H(s) = (\lambda s + (1-\lambda))^{n-1}$. The equilibrium μ-quantile function is:

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

4.3. **Asymmetric Duopoly.** Any duopoly has symmetric interactions, so Proposition 1 of Armstrong and Vickers (2022) guarantees a unique equilibrium with common lower support bound and profits proportional to reach. We derive the explicit μ-space CDFs.

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:

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

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 4.2. *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*[14]

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

The corresponding asymmetric duopoly quantile functions are

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

[14] 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 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 $u \leq 1 - \Delta_1$, 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 first by Ireland (1993) and McAfee (1994), yields closed-form equilibrium characterizations for arbitrary n.

Assumption 4.3. 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:

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

This separability is what enables closed-form solutions. In particular, we state the μ-space equilibrium and verify it directly. For the remainder of this section we posit Assumption 4.3 and order firms so that $\lambda_1 > \lambda_2 > \lambda_3 \geq \cdots \geq \lambda_n$, where we impose the first two strict inequalities to simplify the equilibrium construction. With this arrangement in hand, we define the common lower bound $\underline{\mu} := \rho_1 = \prod_{h=2}^n (1 - \lambda_h)$ and the upper support bounds $\bar{\mu}_1 = \bar{\mu}_2 := 1$, and for each $k \in \{3, \ldots, n\}$,

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

For each $m \in \{2, \ldots, n\}$, we also 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 $\Gamma(\mu) := 1 - \left(\underline{\mu}/(\mu C_m)\right)^{1/(m-1)}$.

Proposition 4.4. *The unique equilibrium in μ-space has the following structure:*

1. *Firm 1 mixes on $\left[\underline{\mu}, 1\right]$ and has an atom at $\mu = 1$. Firm 2 mixes continuously on $\left[\underline{\mu}, 1\right]$ 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 $m \in \{2, \ldots, n\}$, for each $\mu \in [\bar{\mu}_{m+1}, \bar{\mu}_m]$,*

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

On the common overlap $[\underline{\mu}, \bar{\mu}_n]$ (where $m = n$),[15]

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

[15] Firm 1 has a mass point at $\mu = 1$ of size $\Delta_1 = 1 - F_1^\mu(1^-) = 1 - \frac{\lambda_2}{\lambda_1}$. No other firm has an atom there.

The quantile functions follow by inverting the CDFs. We note the following stochastic dominance of margins, namely, that higher-reach firms price higher at each quantile.

Corollary 4.5. *The margin CDFs satisfy* $F_1^\mu(\mu) \le F_2^\mu(\mu) \le \cdots \le F_n^\mu(\mu)$ *for all* μ. *Equivalently,* $\mu_1(u) \ge \mu_2(u) \ge \cdots \ge \mu_n(u)$ *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 active firms on each support interval. Hence, whenever firm j is active, $F_j^\mu(\mu) = \Gamma(\mu)/\lambda_j$. Higher-reach firms ($\lambda_j$ large), therefore, have flatter CDFs: they place less mass on low margins and, equivalently, maintain higher margins at each quantile. The leader's atom at $\mu = 1$ reflects this hierarchy. Firm 1 has the largest equilibrium profit target, $\pi_1^* = \lambda_1 \underline{\mu}$, and the flattest CDF, so its continuous part does not exhaust all mass below 1; indeed, $F_1^\mu(1^-) = \lambda_2/\lambda_1 < 1$, and the remaining mass is placed at the monopoly margin.

Lower-reach firms have less captive power under independence. Within a given market,

$$\rho_j = \prod_{k \ne j}(1 - \lambda_k) = \frac{\alpha_\emptyset}{1 - \lambda_j},$$

so ρ_j is increasing in λ_j. Instead, what lower-reach firms have are smaller equilibrium profit targets and steeper CDFs, which allow them to achieve those profits without an atom at $\mu = 1$. Firm 2 reaches $\mu = 1$ only continuously, while firms $k \ge 3$ have supports ending at $\bar{\mu}_k < 1$. The common lower bound $\underline{\mu} = \prod_{h \ne 1}(1 - \lambda_h) = \rho_1$ depends on the reach of all rival firms: when rivals have higher reach, the competitive floor on margins falls. Pass-through heterogeneity across firms then follows from their positions in the margin distribution. Higher-reach firms post higher margins at each quantile and, when ϕ_c is decreasing in μ (as under 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 (Demand_I)

factors multiplicatively, implying that $\lambda_j F_j^{\mu}(\mu)$ must be equal across active firms on each support interval. (1) solves this equal-Γ structure explicitly on the lowest interval where all n firms compete. For general (correlated) consideration structures, this factorization fails. Nevertheless, for duopoly, the qualitative properties (common support, stochastic dominance ordering, atoms for high-reach firms, pass-through ranking) hold for general consideration structures.

5. Quantile Pass-Through

We now derive our main pass-through results. The μ-isomorphism lets us compute pass-through rates by differentiating the mapping from margins to prices. All the action comes from how ϕ responds to costs. Throughout, fix a firm i with equilibrium μ-quantile function $\mu_i(u) \coloneqq \inf\left\{\mu:\ F_i^{\mu}(\mu) \geq u\right\}$.

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 6. 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:*

$$\tau_i^Q(u;c) = \phi_c(\mu_i(u), c) = \frac{x(p_i(u;c))(1-p_i(u;c))}{(1-c)[x(p_i(u;c)) + (p_i(u;c)-c)x'(p_i(u;c))]} \tag{2}$$

5.2. Economic Interpretation.

The pass-through formula (2) factors as

$$\tau_i^Q(u;c) = \frac{1-p}{(1-c)(1-\varepsilon(p))}, \qquad \text{and} \qquad \varepsilon(p) \equiv -\frac{(p-c)x'(p)}{x(p)} \in [0,1),$$

since $x(p) + (p-c)x'(p) = x(p)(1-\varepsilon)$ is the slope of the effective-margin function $(p-c)x(p)$. Note that ε is a property of per-customer demand $x(\cdot)$, not of the firm's equilibrium residual demand $x(p)q_i(p)$, which is endogenous to the equilibrium.

The numerator $(1-p)/(1-c)$ is the firm's normalized headroom below the reservation value. The denominator $(1-\varepsilon)$ measures the slope of the effective margin. When ε is close to 1, the effective margin is nearly flat and the firm must adjust price substantially to preserve $(p-c)x(p) = \mu(1-c)x(1)$ after a cost shock. Hence, more elastic $x(\cdot)$ at a fixed quantile raises τ.

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. Consumers do not randomly draw from the price distribution but rather systematically buy more at lower prices.[16] This selection effect fundamentally changes how we think about pass-through and incidence.

A key simplification comes from the equilibrium indifference condition: the mass of transactions at price p satisfies $T_i(p;c) = \pi_i(c)/(p-c)$, so transaction volume is proportional to $1/(p-c)$, independent of the demand function $x(\cdot)$. Firms selling at lower prices must compensate with proportionally higher volume to maintain the same profit. Define the *harmonic integral* $B_i(c) \equiv \int (p-c)^{-1} dF_i(p;c)$, in which case the mean transaction-weighted price is then $\bar{p}_i^{\text{trans}}(c) = c + 1/B_i(c)$.

Proposition 6.1. *Posit Assumptions 2.3 and 2.4. If a firm i's posted-price support is bounded away from c, the transaction-weighted pass-through rate is*

$$\tau_i^{trans}(c) = 1 + \frac{\int_0^1 \frac{\phi_c(\mu_i(u),c)-1}{(\phi(\mu_i(u),c)-c)^2} du}{\left[\int_0^1 \frac{1}{\phi(\mu_i(u),c)-c} du\right]^2}. \tag{3}$$

Under unit demand, this reduces to a closed-form expression in terms of the consideration structure alone. For each firm i, with $\mu_i(u)$ being firm i's equilibrium quantile function in the margin game, we define $K_i := \int_0^1 \frac{1}{\mu_i(u)} du$. Under unit demand, $B_i(c) = \frac{K_i}{1-c}$ for each firm i. Thus,

[16] Empirically, consumers disproportionately purchase at low prices, so dispersion in posted prices need not mirror dispersion in transaction prices. Reference, e.g., Sorensen (2000); Kaplan et al. (2019).

Corollary 6.2. *Under unit demand ($x(p) = 1$), firm i's mean transaction price and transaction-weighted pass-through satisfy*

$$\bar{p}_i^{trans}(c) = c + \frac{1-c}{K_i}, \qquad \text{and} \qquad \tau_i^{trans} \equiv \frac{d\bar{p}_i^{trans}(c)}{dc} = 1 - \frac{1}{K_i}. \tag{4}$$

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 $\mu(u) = \rho/H(1-u)$, so $K = \bar{H}/\rho$, where $\bar{H} \equiv \int_0^1 H(s)ds$, and the general formula (4) reduces to $\tau^{\text{trans}} = 1 - \rho/\bar{H}$.

Under independent consideration, the per-firm formulas aggregate nicely. Let $\Lambda \equiv \sum_{k=1}^n \lambda_k$ denote total reach and $\alpha_\emptyset \equiv \prod_{k=1}^n (1-\lambda_k)$ the no-purchase mass.

Proposition 6.3. *Posit Assumption 4.3 and order firms so that $\lambda_1 > \lambda_2 \geq \cdots \geq \lambda_n$. The aggregate expected transaction margin is*

$$\bar{\mu}^{trans} = \frac{\underline{\mu}\,\Lambda}{1-(1-\lambda_1)\underline{\mu}}. \tag{5}$$

Corollary 6.4. *Under unit demand and independent consideration, the industry pass-through rate is*

$$\tau^{trans,agg} = 1 - \frac{\underline{\mu}\,\Lambda}{1-(1-\lambda_1)\underline{\mu}}. \tag{6}$$

Industry pass-through is one minus the average market power among transacting consumers. The aggregate margin is determined by three scalars: the support floor $\underline{\mu}$, the total reach Λ, and the leader's reach λ_1.

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. Under unit demand, the new $K_i = \int \mu_i(u)^{-1} du$ determines $\tau_i^{\text{trans}} = 1 - 1/K_i$. Under independent consideration, if the merged entity's consideration probability satisfies $1 - \lambda_{jk} = (1 - \lambda_j)(1 - \lambda_k)$—i.e., the merged firm is considered whenever either predecessor would have been—then 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.

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, weakly decreasing, and continuously differentiable, have strictly positive demand at the upper bound ($x(1) > 0$), and satisfy Assumption 2.4. We normalize $x(1) = 1$.

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

$$(p-c)x(p) = \mu(1-c)x(1) \tag{7}$$

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 $x(p) \ge 1$ for $p \in [c,1]$. From (7),

$$\mu(1-c) = (p-c)x(p) \ge p-c \quad \Longrightarrow \quad p \le c+\mu(1-c).$$

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

For the pass-through bound, from Theorem 5.1, define $\varepsilon \equiv -(p-c)x'(p)/x(p) \ge 0$. From Assumption 2.4, $\frac{d}{dp}[(p-c)x(p)] = x(p)(1-\varepsilon) > 0$, so $\varepsilon < 1$. Then

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

as $\varepsilon \ge 0$ and $\phi(\mu,c) \le c+\mu(1-c)$. ∎

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. *Let $\mu \in [0,1)$. For common demand families (with $d := 1-c$), in addition to the universal lower bound $1-\mu \le \phi_c(\mu,c)$, the following hold:*

(1) ***Linear demand.*** *$x(p) = 1 + b(1-p)$ with $b \in [0, 1/d)$: $\phi_c(\mu,c) \le \frac{1+\sqrt{1-\mu}}{2}$. The lower bound is attained at $b = 0$ (unit demand), and the upper envelope is tight as $b \uparrow 1/d$.*

(2) ***Constant semi-elasticity.*** *$x(p) = \mathrm{e}^{\beta(1-p)}$ with $\beta \in [0, 1/d)$: $\phi_c(\mu,c) \le 1$. The lower bound is attained at $\beta = 0$ (unit demand), and the upper envelope is tight as $\beta \uparrow 1/d$.*

(3) ***Constant elasticity*** *$x(p) = p^{-\eta}$ with $0 \le \eta < 1/d$ and $c > 0$: for $\mu \in (0,1)$,*

$$\phi_c(\mu,c) = \frac{1-\phi(\mu,c)}{d\left(1-\eta\frac{\phi(\mu,c)-c}{\phi(\mu,c)}\right)},$$

where $\phi(\mu,c) \in (c,1)$ is the unique solution to

$$(\phi(\mu,c)-c)\,\phi(\mu,c)^{-\eta} = \mu d.$$

For constant-elasticity demand, writing $p := \phi(\mu,c)$, we have

$$\phi_c(\mu,c) > 1 \iff p < \eta d, \quad \phi_c(\mu,c) = 1 \iff p = \eta d, \quad \text{and} \quad \phi_c(\mu,c) < 1 \iff p > \eta d.$$

In particular, when $\eta = 1$,

$$\phi(\mu,c) = \frac{c}{1-\mu d}, \qquad \text{and} \qquad \phi_c(\mu,c) = \frac{1-\mu}{(1-\mu d)^2}.$$

Figure 1 illustrates the linear envelope and a collection of CES curves. In the linear case, pass-through lies between the universal lower bound $\tau = 1-\mu$ (unit demand, $b = 0$) and the upper envelope $\tau = \left(1+\sqrt{1-\mu}\right)/2$, which is tight as $b \uparrow$

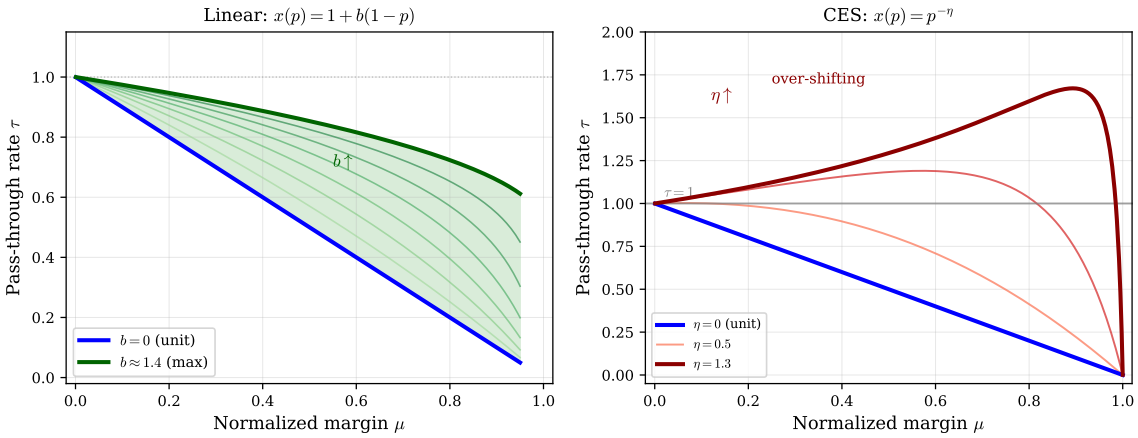


FIGURE 1. **Robust pass-through bounds by demand family.**

$1/(1-c)$. Thus, linear demand never generates over-shifting, and for $\mu \in (0,1)$ pass-through is strictly below one. For CES demand with $c > 0$, there is no analogous family-wide upper envelope. Instead, pass-through is governed by the expression in Theorem 7.3: under the maintained admissible range $\eta < 1/(1-c)$, pass-through exceeds one exactly when the implied price satisfies $p < \eta(1-c)$. When the CES pass-through curve crosses the line $\tau = 1$, therefore, depends jointly on η, c, and the induced price level, rather than on a universal "critical elasticity."

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. 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; 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 $\phi(\cdot, c)$ being 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 sufficiently high elasticity, ϕ_c need not be decreasing in μ—it can increase over parts of the support, so the pass-through ordering is ambiguous.

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. Under independent consideration, Corollary 4.5 establishes that higher-reach firms maintain higher margins at each quantile. Likewise, in asymmetric duopoly, Proposition 4.2 orders firms by their captive-to-reach ratios. From Proposition 8.1, these firms also charge higher prices at each quantile; and, 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.

Of course, for welfare analysis, transaction-weighted prices matter. The following corollary shows that price dominance carries through to mean paid prices.

Corollary 8.2. *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 $\underline{m}_i^M(c) := \inf_{u\in[0,1]}\left(p_i^M(u;c) - c\right) > 0$.*[17] *If $p_i^A(u;c) \geq p_i^B(u;c)$ for all $u \in [0,1]$, then $B_i^A(c) \leq B_i^B(c)$.*

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

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 8.2 extends price orderings to transaction-weighted objects relevant for welfare.

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: within a single market, informed and uninformed consumers face a common posted-margin distribution but transact at different quantiles of it, and quantile pass-through varies across the distribution (§5).

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.[18] 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.[19]

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. For 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

[18]Related obfuscation mechanisms include add-on pricing (Ellison, 2005) and consumer myopia (Gabaix and Laibson, 2006).

[19]Reference Dinerstein et al. (2018) on platform design and consumer search; Nocke and Rey (2024) on steering and choice overload; and Teh et al. (2024) on platforms strategically limiting accessibility. Hagiu and Jullien (2011) and Teh and Wright (2022) formulate related models of intermediation and steering.

$\mu(u) = [(1-\lambda)/(1-\lambda u)]^{n-1}$. 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. Endogenous consideration and general-equilibrium pass-through

Our baseline analysis treats the consideration structure $\{\alpha_S\}_{S\subseteq N}$ as exogenous and invariant to the cost level c. In many environments, however, consideration itself is an equilibrium outcome of marketing, placement, and product-design choices. A leading example is Eliaz and Spiegler (2011a,b), in which firms choose marketing strategies jointly with payoff-relevant product characteristics, and these choices determine which products enter consumers' consideration sets. In such settings, a cost shock can shift the consideration structure and thereby shift the equilibrium distribution of normalized effective margins. This section records the resulting modification to our pass-through formulas in reduced form.

Let $\Delta\left(2^N\right)$ denote the simplex over subsets of N. We represent endogenous consideration by a reduced-form map

$$\begin{aligned} \alpha\colon [0,1) &\to \Delta\left(2^N\right) \\ c &\mapsto (\alpha_S(c))_{S\subseteq N}. \end{aligned}$$

For each c, the pricing game is played with marginal cost c and consideration structure $\alpha(c)$. In the margin game, Theorem 3.3 continues to imply that any equilibrium profile of μ-CDFs $\left(F_i^{\mu}(\cdot;c)\right)_{i\in N}$ depends on the demand system only through $\alpha(c)$. Consequently, shifts in the consideration structure mediate any c-dependence of equilibrium μ-quantile functions $\mu_i(\cdot;c)$.

To separate the two channels, we define the *conditional* (*viz.*, competition-layer fixed) quantile pass-through rate as

$$\tau_i^{Q,\mathrm{cond}}(u;c) := \phi_c(\mu_i(u;c),c),$$

which is exactly the object characterized in Theorem 5.1, evaluated at the equilibrium μ-quantile $\mu_i(u;c)$.

When consideration shifts with costs, the *general-equilibrium* (total-derivative) quantile pass-through rate is

$$\tau_i^{Q,\mathrm{GE}}(u;c) \coloneqq \frac{d}{dc} p_i(u;c).$$

Assume $c \mapsto \mu_i(u;c)$ is differentiable at the cost level of interest. Via the chain rule, we have the decomposition

$$\tau_i^{Q,\mathrm{GE}}(u;c) = \tau_i^{Q,\mathrm{cond}}(u;c) + \phi_\mu(\mu_i(u;c),c)\,\mu_{i,c}(u;c), \quad \text{where} \quad \mu_{i,c}(u;c) \coloneqq \frac{d}{dc}\mu_i(u;c).$$

In the baseline model with fixed α, we have $\mu_{i,c}(u;c) = 0$, so $\tau_i^{Q,\mathrm{GE}}(u;c) = \tau_i^{Q,\mathrm{cond}}(u;c)$ and Definition 6 coincides with the total derivative.

The new term depends on ϕ_μ, which we can compute in closed form. Differentiating (Inverse Map) with respect to μ (holding c fixed), we obtain

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

By Assumption 2.4, the denominator is strictly positive, so $\phi_\mu(\mu,c) > 0$. Evaluating at the equilibrium quantile, and using $p_i(u;c) = \phi(\mu_i(u;c),c)$, we can write

$$\phi_\mu(\mu_i(u;c),c) = \frac{(1-c)\,x(1)}{x(p_i(u;c)) + (p_i(u;c)-c)\,x'(p_i(u;c))}.$$

Similarly, Theorem 5.1 implies

$$\tau_i^{Q,\mathrm{cond}}(u;c) = \frac{x(p_i(u;c))\,(1-p_i(u;c))}{(1-c)\,[x(p_i(u;c)) + (p_i(u;c)-c)\,x'(p_i(u;c))]}.$$

Combining these expressions yields an explicit formula:

$$\begin{aligned}\tau_i^{Q,\mathrm{GE}}(u;c) = &\frac{x(p_i(u;c))\,(1-p_i(u;c))}{(1-c)\,[x(p_i(u;c)) + (p_i(u;c)-c)\,x'(p_i(u;c))]} \\ &+ \frac{(1-c)\,x(1)}{x(p_i(u;c)) + (p_i(u;c)-c)\,x'(p_i(u;c))}\,\mu_{i,c}(u;c).\end{aligned}$$

Thus, relative to the conditional formula in Theorem 5.1, endogenous consideration adds a term proportional to the cost-induced shift in the equilibrium μ-quantile function. Because $\phi_\mu(\mu, c) > 0$, the sign of the additional term is governed by $\mu_{i,c}(u; c)$, i.e., by whether a cost increase shifts the equilibrium distribution of normalized effective margins upward or downward.

In models where consideration is shaped by equilibrium marketing and product-design choices, such as Eliaz and Spiegler (2011a,b), a cost change can shift $\alpha(c)$ through changes in advertising intensity, platform placement, or product characteristics that affect consumers' attention and sampling. Our decomposition implies that any such mechanism affects pass-through only through the induced response $\mu_{i,c}(u; c)$ of the competition layer, while the curvature layer ϕ continues to govern the mapping from (μ, c) to prices.

Example 10.1 (Unit Demand)**.** When $x(p) \equiv 1$, we have $\phi(\mu, c) = c + (1 - c)\mu$, so

$$p_i(u; c) = c + (1 - c)\mu_i(u; c)$$

and the GE quantile pass-through simplifies to

$$\tau_i^{Q,\mathrm{GE}}(u; c) = 1 - \mu_i(u; c) + (1 - c)\mu_{i,c}(u; c).$$

In this benchmark, the full wedge between conditional and general-equilibrium pass-through is $(1 - c)\mu_{i,c}(u; c)$.

All of the demand-robust envelope results in §7 apply directly to the conditional component $\tau_i^{Q,\mathrm{cond}}(u; c)$, evaluated at the equilibrium price quantile $p_i(u; c)$. The additional general-equilibrium term depends on the endogenous competition-layer response $\mu_{i,c}(u; c)$, which demand curvature alone does not discipline.

Appendix A. §3 Proofs

A.1. **Proof of Lemma 3.1.**

Proof of Lemma 3.1. 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 $g(p^*) = 0$. Uniqueness follows because $g'(p) = x(p) + (p-c)x'(p) > 0$ on $[c,1]$ by Assumption 2.4, so g is strictly increasing. ■

A.2. **Proof of Lemma 3.2.** Let $U_i(\mu) \coloneqq \mu_i q_i^{\mu}(\mu)$ denote firm i's payoff.

Proof of Lemma 3.2. 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 $A_i \coloneqq [0,1]$, a closed interval. Moreover, for all $\mu \in [0,1]^n$,

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

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 function $f_{ij}^1(a_i) \coloneqq a_i$. 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

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

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

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

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 $\liminf_{t \nearrow t_0} U_i(t, \mu_{-i}) \geq U_i(t_0, \mu_{-i})$. This left-limit inequality establishes Dasgupta and Maskin (1986, Definition 6) with $\lambda = 1$ (full weight on the left-hand liminf) for all $t_0 \in (0,1]$.

At the left endpoint, $t_0 = 0$, we have $U_i(0, \mu_{-i}) = 0$ and $U_i \geq 0$, so $\liminf_{t \searrow 0} U_i(t, \mu_{-i}) \geq 0 = U_i(0, \mu_{-i})$, 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.3. **Proof of Theorem 3.3.**

Proof of Theorem 3.3. 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:

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

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

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

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:

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

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 3.2 combined with the bijection established here. ∎

Appendix B. §4 and §5 Proofs

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

Proof of Proposition 4.1. By the μ-isomorphism (Theorem 3.3), the margin game is strategically equivalent to the unit-demand pricing game of Armstrong and Vickers (2022). Their Proposition 1 establishes that with symmetric consideration structure, the unique symmetric equilibrium has interval support $\left[\underline{\mu}, 1\right]$ with $\underline{\mu} = \rho$, equilibrium profit $\pi^* = \rho\sigma$, and no atoms on $\left(\underline{\mu}, 1\right)$.

It remains to derive the quantile function. Let $G(\mu) \equiv 1 - F^{\mu}(\mu)$. Since the distribution is atomless on the interior, the demand share at μ is

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

Indifference on $[\rho, 1]$ requires $\mu \cdot q^{\mu}(\mu) = \rho\sigma$. Dividing by σ and defining

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

we obtain $\mu \cdot H(G(\mu)) = \rho$, so $\mu = \rho / H(G(\mu))$. With the quantile transformation $u = F^{\mu}(\mu) = 1 - G(\mu)$, this yields $\mu(u) = \frac{\rho}{H(1-u)}$. The support endpoints are confirmed by $H(1) = 1$ and $H(0) = \rho$, delivering $\mu(0) = \rho$ and $\mu(1) = 1$. ∎

B.2. **Proof of Proposition 4.2.**

Proof of Proposition 4.2. Any duopoly has symmetric interactions, so uniqueness follows from Armstrong and Vickers (2022, Proposition 1). We derive the CDFs from the indifference conditions.

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 4.4.**

Proof of Proposition 4.4. Independent consideration satisfies symmetric interactions (Armstrong and Vickers, 2022, Section 3), so their Proposition 1 guarantees a unique equilibrium with common lower bound $\underline{\mu} = \rho_1$, nested interval supports, profits $\pi_j^* = \lambda_j \underline{\mu}$, and the common-CDF property $\lambda_j F_j^{\mu}(\mu) = \lambda_k F_k^{\mu}(\mu) =: \Gamma(\mu)$ for all firms active at μ. We verify the claimed μ-space CDFs.

When firms $\{1,\dots,m\}$ are active at μ (i.e., $\mu \in [\bar{\mu}_{m+1}, \bar{\mu}_m]$), each active firm $j \le m$ has $\lambda_j F_j^{\mu}(\mu) = \Gamma(\mu)$ and each inactive firm $i > m$ has $F_i^{\mu}(\mu) = 1$. Firm j's profit is

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

where $C_m = \prod_{h=m+1}^{n}(1-\lambda_h)$. Setting $\Pi_j(\mu) = \lambda_j \underline{\mu}$ and solving yields

$$(1-\Gamma(\mu))^{m-1} = \frac{\underline{\mu}}{\mu\, C_m}, \qquad \text{so} \qquad \Gamma(\mu) = 1 - \left(\frac{\underline{\mu}}{\mu\, C_m}\right)^{1/(m-1)}.$$

The cutoff $\bar{\mu}_k$ where firm k's CDF reaches 1 satisfies $\Gamma(\bar{\mu}_k) = \lambda_k$, yielding

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

By taking $\mu \uparrow 1$ on the top interval (where $m = 2$) we obtain $\Gamma(1^-) = \lambda_2$, so $F_1^{\mu}(1^-) = \lambda_2/\lambda_1 < 1$, confirming firm 1's atom $\Delta_1 = 1 - \lambda_2/\lambda_1$.

For the no-deviation check, if $\mu < \underline{\mu}$ then $\Pi_j(\mu) = \mu\lambda_j < \underline{\mu}\lambda_j = \pi_j^*$. If $j > m$ at some $\mu \in [\bar{\mu}_{m+1}, \bar{\mu}_m]$, then $\Pi_j(\mu) = \lambda_j \underline{\mu}(1-\Gamma(\mu))/(1-\lambda_j) \le \pi_j^*$ since $\Gamma(\mu) \ge \lambda_{m+1} \ge \lambda_j$. ∎

B.4. **Proof of Theorem 5.1.**

Proof of Theorem 5.1. Recall the implicit definition (Inverse Map):

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

Differentiating both sides with respect to c,[20] holding μ fixed, and rearranging yields

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

From (Inverse Map), 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,

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

Since $\mu_i(u)$ does not depend on c (Theorem 3.3(b)), setting $\mu = \mu_i(u)$ and noting that $\phi(\mu_i(u),c) = p_i(u;c)$ yields (2). ∎

Appendix C. §6 Transaction-Weighted Proofs and Extensions

C.1. **Supporting Lemmas.**

Lemma C.1. *At equilibrium, for all $p \in \text{supp}(F_i)$ with $p > c$, $T_i(p;c) = x(p)q_i(p) = \frac{\pi_i(c)}{p-c}$.*

Proof. Immediate from the equilibrium indifference condition. ∎

Lemma C.2. *The mean transaction-weighted price is $\bar{p}_i^{trans}(c) = c + \frac{1}{B_i(c)}$, where $B_i(c) \equiv \int (p-c)^{-1} dF_i(p;c)$.*

Proof. The *transaction-weighted CDF* is:

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

so the mean paid price is

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

[20] Permitted by the implicit function theorem, by the smoothness from Assumptions 2.3 and 2.4.

C.2. **Proof of Proposition 6.1.**

Proof of Proposition 6.1. From Lemma C.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 $B_i'(c)$. By the μ-isomorphism, $\mu_i(\cdot)$ is fixed as c varies, so

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

Differentiating under the integral sign is justified by the support assumption: locally around c, $p_i(u;\tilde{c}) - \tilde{c}$ is bounded below uniformly in u, and $\phi_c(\mu_i(u), \tilde{c})$ is bounded. Consequently,

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

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 (3). ∎

C.3. **Proof of Corollary 6.2.**

Proof. We have

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

Thus, using Lemma C.2, $\bar{p}_i^{\text{trans}}(c) = c + (1-c)/K_i$. Differentiating, $\tau_i^{\text{trans}} = 1 - 1/K_i$. ∎

C.4. **Proofs of Proposition 6.3 and Corollary 6.4.** Denote

$$\bar{\mu}^{\text{trans}} := \mathbb{E}\left[\frac{(p-c)x(p)}{(1-c)x(1)} \,\middle|\, \text{a purchase occurs}\right],$$

where p is the price paid by a randomly selected purchasing consumer. Thus, $\bar{\mu}^{\text{trans}}$ is the average normalized effective margin among purchasing consumers. Also denote $\Lambda := \sum_{i=1}^n \lambda_i$ and recall from Proposition 4.4 that under independent consideration the common lower support bound is $\underline{\mu} = \prod_{h=2}^n (1 - \lambda_h)$.

Proof of Proposition 6.3. If some non-leaders have equal reaches, the payoff formula used below follows by continuity from the (generic) strict-order case. Thus it suffices to write the argument for the strict-order case covered by Proposition 4.4.

Let π_j^{μ} denote firm j's equilibrium payoff in the margin game. From Proposition 4.4(2), $\pi_j^{\mu} = \lambda_j \underline{\mu}$. Moreover, equilibrium payoffs in the pricing game are $(1-c)x(1)$ times equilibrium payoffs in the margin game (Theorem 3.3). Consequently, firm j's equilibrium profit in the pricing game is $\pi_j(c) = (1-c)x(1)\lambda_j\underline{\mu}$; and by summing across firms, total expected industry profit is $\sum_{j=1}^{n} \pi_j(c) = (1-c)x(1)\underline{\mu}\Lambda$.

The same total expected industry profit can be computed from purchases. A purchase occurs with probability $1-\alpha_\emptyset$, and conditional on a purchase, the average normalized effective margin is $\bar{\mu}^{\text{trans}}$. Therefore,

$$\sum_{j=1}^{n} \pi_j(c) = (1-c)x(1)\,(1-\alpha_\emptyset)\,\bar{\mu}^{\text{trans}}.$$

Equating the two expressions, we get

$$(1-\alpha_\emptyset)\,\bar{\mu}^{\text{trans}} = \underline{\mu}\Lambda.$$

Under independent consideration,

$$\alpha_\emptyset = \prod_{j=1}^{n} \left(1-\lambda_j\right) = (1-\lambda_1)\,\underline{\mu}.$$

Substituting gives

$$\bar{\mu}^{\text{trans}} = \frac{\underline{\mu}\Lambda}{1-\alpha_\emptyset} = \frac{\underline{\mu}\Lambda}{1-(1-\lambda_1)\,\underline{\mu}},$$

which is exactly (5). ∎

The aggregate margin formula immediately yields Corollary 6.4.

Proof of Corollary 6.4. Under unit demand, the normalized effective margin of a purchase at price p is $\frac{p-c}{1-c}$. Thus, since $\bar{p}^{\text{trans}}(c)$ is the average price paid conditional

on purchase, we use the definition of $\bar{\mu}^{\text{trans}}$ to obtain

$$\bar{\mu}^{\text{trans}} = \frac{\bar{p}^{\text{trans}}(c) - c}{1 - c} \quad \Longleftrightarrow \quad \bar{p}^{\text{trans}}(c) = c + (1 - c)\,\bar{\mu}^{\text{trans}}.$$

(5) tells us that $\bar{\mu}^{\text{trans}}$ depends only on the reach vector and not on c, so

$$\tau^{\text{trans,agg}} = \frac{d\bar{p}^{\text{trans}}(c)}{dc} = 1 - \bar{\mu}^{\text{trans}}.$$

Substituting (5) yields (6). ∎

Appendix D. §7 Pass-Through Envelope Proofs

D.1. **Proof of Theorem 7.3.**

Proof. Let $d := 1 - c$. Fix $\mu \in (0,1)$, write $p := \phi(\mu, c)$, $m := p - c$, and $q := 1 - p$, so $m, q \in (0, d)$ and $m + q = d$. From $\phi(0, c) = c$ and $\phi(1, c) = 1$, we have $\phi_c(0, c) = 1$ and $\phi_c(1, c) = 0$. The universal lower bound $1 - \mu \le \phi_c(\mu, c)$ is Theorem 7.2, so we focus on the family-specific upper envelopes and formulas.

Case 1: Linear demand. Let $x(p) = 1 + b(1 - p)$ with $b \in [0, 1/d)$. When $b = 0$, we recover unit demand, so $\phi_c(\mu, c) = 1 - \mu$. For $b > 0$, the effective-margin equation

$$m\,(1 + b(d - m)) = \mu d, \tag{D1}$$

is solved to get

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

where we take the smaller root to keep $p \le 1$.

Since $x'(p) = -b$, Theorem 5.1 delivers

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

Using the equilibrium relation $b = (\mu d - m)/(mq)$, this simplifies to

$$\phi_c(\mu, c) = \frac{\mu q^2}{(m - \mu d)^2 + \mu(1 - \mu)d^2}.$$

Differentiating with respect to m yields

$$\frac{\partial \phi_c}{\partial m} = -\frac{2\mu(1-\mu)dm(d-m)}{\left[(m-\mu d)^2 + \mu(1-\mu)d^2\right]^2} < 0.$$

Thus, ϕ_c is strictly decreasing in m. Implicit differentiation of ($D1$) produces

$$\frac{\partial m}{\partial b} = -\frac{m(d-m)}{1+bd-2bm} = -\frac{mq}{1+bd-2bm} < 0,$$

where the denominator is positive by Assumption 2.4. Therefore, ϕ_c is increasing in b. The supremum over admissible $b < 1/d$ is the limit as $b \uparrow 1/d$. In that weak-boundary case, ($D1$) becomes $m\left(2 - \frac{m}{d}\right) = \mu d$, so $m = d\left(1 - \sqrt{1-\mu}\right)$.

Substituting into the pass-through formula, (2) in Theorem 5.1, yields

$$\sup_{b<1/d} \phi_c(\mu, c) = \frac{1+\sqrt{1-\mu}}{2}.$$

Case 2: Constant semi-elasticity. Let $x(p) = e^{\beta(1-p)}$ with $\beta \in [0, 1/d)$. When $\beta = 0$, we again recover unit demand, so $\phi_c(\mu, c) = 1 - \mu$. For $\beta > 0$, the inverse equation is

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

Because $x'(p) = -\beta x(p)$, Theorem 5.1 produces

$$\phi_c(\mu, c) = \frac{x(p)(1-p)}{d\left[x(p) + (p-c)x'(p)\right]} = \frac{1-p}{d\left(1-\beta(p-c)\right)} = \frac{q}{d(1-\beta m)}.$$

Hence,

$$\phi_c(\mu, c) \le 1 \iff q \le d(1-\beta m) \iff d - m \le d - d\beta m \iff m(1-d\beta) \ge 0,$$

which holds for every $\beta < 1/d$. Thus, $\phi_c(\mu, c) \le 1$. The bound is tight because in the weak-boundary case $\beta = 1/d$,

$$\phi_c(\mu, c) = \frac{q}{d(1-m/d)} = 1$$

for every $\mu \in (0, 1)$, so $\sup_{\beta<1/d} \phi_c(\mu, c) = 1$.

Case 3: Constant elasticity. Let $x(p) = p^{-\eta}$ with $0 \le \eta < 1/d$ and $c > 0$. The effective-margin equation is

$$(p-c)p^{-\eta} = \mu d.$$

Moreover,

$$\frac{d}{dp}\left[(p-c)p^{-\eta}\right] = p^{-\eta-1}\left[p - \eta(p-c)\right],$$

so admissibility is equivalent to $p - \eta(p-c) > 0$ on the relevant support.

Applying Theorem 5.1 produces

$$\phi_c(\mu, c) = \frac{p^{-\eta}(1-p)}{d\left[p^{-\eta} - \eta(p-c)p^{-\eta-1}\right]} = \frac{1-p}{d\left(1 - \eta\frac{p-c}{p}\right)}.$$

Subtracting one and simplifying,

$$\phi_c(\mu, c) - 1 = \frac{(p-c)(\eta d - p)}{d\left[p - \eta(p-c)\right]}.$$

Since the denominator is positive on the interior support, for $p \in (c, 1)$ we have

$$\phi_c(\mu, c) > 1 \iff p < \eta d, \qquad \phi_c(\mu, c) = 1 \iff p = \eta d, \qquad \phi_c(\mu, c) < 1 \iff p > \eta d.$$

In the special case $\eta = 1$,

$$(p-c)p^{-1} = \mu d \Longrightarrow p = \frac{c}{1-\mu d}.$$

Substituting into the pass-through formula, (2) in Theorem 5.1, yields

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

At $\eta = 0$, the CES family collapses to unit demand and recovers the lower bound $\phi_c(\mu, c) = 1 - \mu$. ■

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.

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.

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

Brian C Albrecht. Price competition and the use of consumer data. Mimeo, 2020.

Simon P. Anderson, André de Palma, and Brent Kreider. Tax incidence in differentiated product oligopoly. *Journal of Public Economics*, 81(2):173–192, 2001.

Mark Armstrong. Competition in two-sided markets. *RAND Journal of Economics*, 37(3):668–691, 2006.

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

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

Mark Armstrong and John Vickers. Multiproduct cost pass-through: Edgeworth's paradox revisited. *Journal of Political Economy*, 131(10):2645–2665, 2023.

Susan Athey and Glenn Ellison. Position auctions with consumer search. *The Quarterly Journal of Economics*, 126(3):1213–1270, 2011.

Michael R. Baye and John Morgan. Information gatekeepers on the internet and the competitiveness of homogeneous product markets. *American Economic Review*, 91(3):454–474, 2001.

Michael R. Baye, John Morgan, and Patrick Scholten. Price dispersion in the small and in the large: Evidence from an internet price comparison site. *Journal of Industrial Economics*, 52(4):463–496, 2004.

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

Jeremy I. Bulow and Paul Pfleiderer. A note on the effect of cost changes on prices. *Journal of Political Economy*, 91(1):182–185, 1983.

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

Kenneth Burdett and Guido Menzio. The (Q,S,s) pricing rule. *The Review of Economic Studies*, 85(2):892–928, 2018.

Gerard R. Butters. Equilibrium distributions of sales and advertising prices. *The Review of Economic Studies*, 44(3):465–491, 1977.

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

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

Sofia Delipalla and Michael Keen. The comparison between ad valorem and specific taxation under imperfect competition. *Journal of Public Economics*, 49(3): 351–367, 1992.

Michael Dinerstein, Liran Einav, Jonathan Levin, and Neel Sundaresan. Consumer price search and platform design in internet commerce. *American Economic Review*, 108(7):1820–1859, 2018.

Kfir Eliaz and Ran Spiegler. Consideration sets and competitive marketing. *The Review of Economic Studies*, 78(1):235–262, 2011a.

Kfir Eliaz and Ran Spiegler. On the strategic use of attention grabbers. *Theoretical Economics*, 6(1):127–155, 2011b.

Kfir Eliaz and Ran Spiegler. A simple model of search engine pricing. *The Economic Journal*, 121(556):F329–F339, 2011c.

Glenn Ellison. A model of add-on pricing. *The Quarterly Journal of Economics*, 120 (2):585–637, 2005.

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

Kai Fischer, Simon Martin, and Philipp Schmidt-Dengler. Indirect taxation in consumer search markets: The case of retail fuel. Mimeo, 2024.

Xavier Gabaix and David Laibson. Shrouded attributes, consumer myopia, and information suppression in competitive markets. *The Quarterly Journal of Economics*, 121(2):505–540, 2006.

Luke Garrod, Ruochen Li, Antonio Russo, and Chris M. Wilson. Understanding cost pass-through when prices are dispersed. Mimeo, 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.

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

Rafael R Guthmann. Price dispersion in dynamic competition. *Economic Theory*, 78:1203–1232, 2024.

Rafael R Guthmann. Imperfect competition as a result of unawareness. *Review of Industrial Organization*, 67:35–53, 2025.

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

Andrei Hagiu and Bruno Jullien. Why do intermediaries divert search? *RAND Journal of Economics*, 42(2):337–362, 2011.

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

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.

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.

Norman J Ireland. The provision of information in a Bertrand oligopoly. *Journal of Industrial Economics*, 41(1):61–76, 1993.

Sonia Jaffe and E. Glen Weyl. The first-order approach to merger analysis. *American Economic Journal: Microeconomics*, 5(4):188–218, 2013.

Maarten C. W. Janssen and José Luis Moraga-González. Strategic pricing, consumer search and the number of firms. *The Review of Economic Studies*, 71(4): 1089–1118, 2004.

Zi Yang Kang and Shoshana Vasserman. Robustness measures for welfare analysis. *American Economic Review*, 115(8):2449–2487, 2025.

Greg Kaplan, Guido Menzio, Leena Rudanko, and Nicholas Trachter. Relative price dispersion: Evidence and theory. *American Economic Journal: Microeconomics*, 11(3):68–124, 2019.

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.

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

Guido Menzio. Search theory of imperfect competition with decreasing returns to scale. Mimeo, 2023.

Guido Menzio. Markups: A search-theoretic perspective. Mimeo, 2024a.

Guido Menzio. Efficient imperfect competition with an application to international trade. Mimeo, 2024b.

Guido Menzio and Nicholas Trachter. Equilibrium price dispersion with sequential search. *Journal of Economic Theory*, 160:188–215, 2015.

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.

Felix Montag, Robin Mamrak, Alina Sagimuldina, and Monika Schnitzer. Imperfect price information, market power, and tax pass-through. Mimeo, 2023.

Volker Nocke and Patrick Rey. Consumer search, steering, and choice overload. *Journal of Political Economy*, 132(5):1684–1739, 2024.

Volker Nocke and Nicolas Schütz. Multiproduct-firm oligopoly: An aggregative games approach. *Econometrica*, 86(2):523–557, 2018.

Volker Nocke and Nicolas Schütz. An aggregative games approach to merger analysis in multiproduct-firm oligopoly. *RAND Journal of Economics*, 2025.

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

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

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

Jean-Charles Rochet and Jean Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029, 2003.

Jean-Charles Rochet and Jean Tirole. Two-sided markets: A progress report. *RAND Journal of Economics*, 37(3):645–667, 2006.

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.

Marc Rysman. The economics of two-sided markets. *Journal of Economic Perspectives*, 23(3):125–143, 2009.

Steven C. Salop and Joseph E. Stiglitz. Bargains and ripoffs: A model of monopolistically competitive price dispersion. *The Review of Economic Studies*, 44(3): 493–510, 1977.

Jesús Seade. Profitable cost increases and the shifting of taxation: Equilibrium response of markets in oligopoly. Mimeo, 1985.

Alan T Sorensen. Equilibrium price dispersion in retail markets for prescription drugs. *Journal of Political Economy*, 108(4):833–850, 2000.

Dale O. Stahl. Oligopolistic pricing with sequential consumer search. *American Economic Review*, 79(4):700–712, 1989.

Samuel Stolper. Who bears the burden of energy taxes? the role of local passthrough. Mimeo, 2017.

Tat-How Teh and Julian Wright. Intermediation and steering: Competition in prices and commissions. *American Economic Journal: Microeconomics*, 14(2):281–321, 2022.

Tat-How Teh, Chengsi Wang, and Makoto Watanabe. Strategic limitation of market accessibility: Search platform design and welfare. *Journal of Economic Theory*, 216:105798, 2024.

Hal R Varian. A model of sales. *American Economic Review*, 70(4):651–659, 1980.

Chengsi Wang and Julian Wright. Search platforms: Showrooming and price parity clauses. *RAND Journal of Economics*, 51(1):32–58, 2020.

Chengsi Wang and Julian Wright. Regulating platform fees. *Journal of the European Economic Association*, 23(2):746–783, 2025.

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

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.

Appendix E. Supplementary Appendix

E.1. **§6 Further Discussion.** Three features deserve emphasis.

Remark E.1. Proposition 6.3 is just an accounting identity for expected industry profit:

$$\sum_{j=1}^{n} \pi_j(c) = (1-c)\,x(1)\underline{\mu}\Lambda = (1-c)\,x(1)\,(1-\alpha_\emptyset)\,\bar{\mu}^{\text{trans}}.$$

The left-hand side computes expected industry profit by summing firm-level equilibrium payoffs. The right-hand side computes the same object by averaging realized transaction margins. In particular, the proof of Proposition 6.3 uses only the margin-game payoff identity. The unit-demand assumption enters in Corollary 6.4, when normalized margins are translated back into transaction prices.

Remark E.2. The conditional aggregate margin is determined by three market-level numbers: the support floor $\underline{\mu}$, the total reach Λ, and the no-purchase mass $\alpha_\emptyset$ (equivalently, λ_1). This is still a substantial dimension reduction relative to the full reach vector $(\lambda_1,\ldots,\lambda_n)$, but unlike the industry-profit identity above, the conditioned object cannot be recovered from $\left(\underline{\mu},\Lambda\right)$ alone.

Remark E.3. Proposition 6.3 yields sharp signs for the conditional aggregate margin. The leader's reach raises the average margin paid by transacting consumers:

$$\frac{\partial \bar{\mu}^{\text{trans}}}{\partial \lambda_1} = \frac{\underline{\mu}\left[1-\underline{\mu}(1+\Lambda-\lambda_1)\right]}{\left[1-(1-\lambda_1)\,\underline{\mu}\right]^2} > 0,$$

where the sign follows from

$$\frac{1}{\underline{\mu}} = \prod_{h=2}^{n}\left(1+\frac{\lambda_h}{1-\lambda_h}\right) \geq 1+\sum_{h=2}^{n}\frac{\lambda_h}{1-\lambda_h} > 1+\sum_{h=2}^{n}\lambda_h = 1+\Lambda-\lambda_1.$$

Each non-leader's reach lowers it:

$$\frac{\partial \bar{\mu}^{\text{trans}}}{\partial \lambda_k} = \frac{\underline{\mu}}{\left[1-(1-\lambda_1)\,\underline{\mu}\right]^2}\left(1-(1-\lambda_1)\,\underline{\mu}-\frac{\Lambda}{1-\lambda_k}\right) < 0$$

for $k \neq 1$, and the sign follows because

$$1 - (1 - \lambda_1)\underline{\mu} = 1 - \alpha_\emptyset \leq \Lambda < \frac{\Lambda}{1 - \lambda_k}.$$

Conditioning on transaction therefore yields a clean asymmetry: greater leader reach pushes the average paid margin up, whereas greater non-leader reach pushes it down.