Discussion Paper Series – CRC TR 224

Discussion Paper No. 799
Project B 04

Recommendation Design, Pricing, and Regulation

Laurenz Marstaller[1]

November 2026

[1]University of Bonn, Adenauerallee 24- 42, 53113, Bonn, Germany. Email: laurenz.marstaller@uni-bonn.de

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Recommendation Design, Pricing, and Regulation*

Laurenz Marstaller†

September 8, 2026

Abstract

This paper studies how platforms jointly choose fees and recommendations and their implications for fee regulation. A platform charges sellers a commission rate and ranks products based on price and match-value. The analysis shows that price-sensitive rankings intensify seller competition, allowing the platform to extract more surplus. Commission-rate caps constrain fees, but platforms may respond by making recommendations less price-sensitive, attenuating consumer-surplus gains. By contrast, capped nominal fees can be more effective because they shift the platform's incentives toward transaction volume rather than transaction value. Effective fee regulation must therefore account for how platforms adjust their recommendation policies in response.

Algorithm Design, DMA, Platform Regulation, Platforms, Recommendations, Search
D43, D83, L13, L51, L86

*I am deeply grateful to my advisors, Stephan Lauermann, Andreas Kleiner, and Jidong Zhou, for their guidance, encouragement, and support. I owe special thanks to Jidong Zhou for detailed and insightful feedback, and to Carl-Christian Groh for generous support and many valuable discussions. I thank Yale University and the University of Mannheim for their hospitality during research visits, as well as participants at the Yale Breakfast Seminar, the BCFM PhD Conference, and the CRC Young Researchers workshop for helpful comments. This research was supported by the Bonn Graduate School of Economics (BGSE) and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/2–390838866 and CRC TR 224 (Project B04).

†University of Bonn, Adenauerallee 24 - 42, 53113, Bonn, Germany — laurenz.marstaller@uni-bonn.de

1 Introduction

Digital platforms such as Apple's App Store, Amazon Marketplace, and Booking.com not only facilitate transactions but also mediate consumer search. Consumers often search through large product spaces, and platforms shape this search process by deciding which products are made especially visible. The ordering of listings and badges such as "Amazon's Choice" may help consumers navigate these choices. At the same time, they influence which sellers receive attention and demand, thereby affecting the allocation of demand across sellers, the intensity of on-platform competition, and market prices. Platforms therefore face a joint design problem: how to monetize transactions and how to steer demand through recommendations.

A prominent feature of platform monetization is the use of commissions proportional to transaction value. These fees have become one focus of broader efforts to regulate digital platforms. In food-delivery markets, caps on commission rates have been a central policy instrument, with many U.S. jurisdictions adopting explicit limits on the percentage charged by delivery platforms. Related concerns about platform commissions are also part of the debate over the regulation of Apple's App Store and other digital platforms.[1] Yet commission regulation is complicated by the platform's ability to steer demand. A cap affects outcomes directly by constraining the fee and indirectly by changing the platform's incentives to adjust recommendation design.

To study how a platform sets fees and designs recommendations, the paper develops a stylized model of platform-mediated search. A monopolistic platform intermediates between horizontally differentiated sellers and consumers who face search frictions. The platform chooses a commission rate and a recommendation rule that trades off consumer-product fit against price. Sellers then set prices, anticipating how the recommendation rule affects their demand. The framework is used to study commission caps and to compare them with caps on nominal per-transaction fees.

The baseline model features a standard double-marginalization force. The platform charges sellers for access to consumers, while sellers set final consumer prices. Specifically, the commission raises sellers' effective marginal cost of making a sale. If sellers retain market power on the platform, final prices therefore reflect both the platform fee and a downstream seller markup. Recommendation design affects this downstream component: more price-sensitive recommendations make seller demand more elastic and compress seller markups.

This yields a simple characterization of platform incentives. In the absence of regulation, the platform chooses a maximally price-sensitive recommendation rule to intensify

[1] For examples of food-delivery commission caps, see Seattle Municipal Code (2022); Administrative Code of the City of New York (2025). In the App Store context, restrictions on commission rates typically arise from broader regulatory and litigation debates over platform conduct, rather than from isolated commission-cap policies. See, for example, Crémer et al. (2024); Competition Appeal Tribunal (2025).

seller competition and increase the surplus it can extract through its commission. At the optimum, downstream markups are competed away, and the platform adjusts its commission to replicate the vertically integrated monopoly outcome. Thus, strong on-platform seller competition can coexist with limited consumer gains and substantial platform surplus extraction.

Commission caps change these incentives. A binding cap constrains the platform's ability to monetize transactions through the commission rate but leaves recommendation design under its control. The platform may therefore respond by reducing the weight placed on price, allowing sellers to retain more market power. Commission caps can consequently have more modest effects on prices and consumer surplus than an analysis with a fixed recommendation rule would suggest. By contrast, when a nominal per-transaction fee is capped, the platform's incentives shift toward transaction volume rather than transaction value, aligning the platform's recommendation incentives more closely with consumer welfare.

Extensions in Sections 4.3 and 5 examine restrictions on algorithm design and forces that limit the optimal price sensitivity of recommendations. When algorithm design is coarse, commission caps can leave consumers worse off. Endogenous quality investment and vertically differentiated products, stylized by the presence of an inferior but cheap product, likewise reduce the platform's incentive to place maximal weight on price. The proofs of Sections 3 and 4 are in the Appendix. Additional results and the proofs of the extensions in Section 5 are in the Online Appendix.

Literature. This paper studies platform-mediated search and recommendation design and their interaction with platform monetization. The consumer side builds on search models of horizontally differentiated markets (Wolinsky, 1986; Anderson and Renault, 1999), but search is mediated by a platform that observes match values and chooses how prices enter recommendations. This differs from canonical models of two-sided platforms, where buyer–seller interactions are captured through reduced-form cross-group effects (Armstrong, 2006; Rochet and Tirole, 2003, 2006).

Hagiu and Jullien (2011) show that intermediaries may strategically divert search and thereby affect both consumer search and firms' pricing incentives. Such distortions in platform-mediated search can arise along several design margins. One is how much weight rankings place on price relative to product fit, which in this paper is chosen jointly with the platform's commission rate. Another is recommendation precision, that is, how accurately platforms direct consumers to their best-matching products (Nocke and Rey, 2024; Peitz and Sobolev, 2026; Ng, 2025).[2]

Two papers are particularly close. The closest paper on ranking design is Rhodes and

[2]A related question is whether commissions rather than prices should enter the recommendation algorithm; see Inderst and Ottaviani (2012) and Teh and Wright (2022).

Zhou (2026). They consider a fixed ranking rule combining price and product fit and compare it with a ranking based only on product fit. As in related work (Karle et al., 2026; Heidhues et al., 2023), the price sensitivity of the ranking is taken as given. Here, by contrast, the platform chooses how much weight to place on price jointly with its commission rate. On the broader design problem, Teh (2022) is especially close. He develops a general framework in which platform governance affects both buyer-side value creation and seller competition, and shows how the direction of governance distortions depends on whether platform revenues are more aligned with seller surplus or transaction volume. Recommendation design is one application. The present paper focuses specifically on recommendation rankings, deriving how their price sensitivity affects seller demand and markups, and uses this structure to study how fee regulation changes the platform's recommendation policy.

Empirical work also highlights the importance of price-sensitive rankings for prices and purchasing patterns (Dinerstein et al., 2018; Kaye, 2026).

More broadly, several strands of the theoretical literature study how platforms shape demand and seller competition through design choices beyond the price sensitivity of recommendations. These include sponsored placement, advertising, and targeting (Bergemann and Bonatti, 2011; de Cornière, 2016), the design and coarseness of search results (Eliaz and Spiegler, 2016), and seller discoverability (Hagiu and Wright, 2024). A related but distinct strand studies self-preferencing and biased intermediation by platforms (de Cornière and Taylor, 2019; Hagiu et al., 2022; Lee and Musolff, 2025), which this paper abstracts from.

A complementary literature studies platform-fee regulation when consumers can substitute between the platform and a direct channel. In Gomes and Mantovani (2025), price parity suppresses price-based channel substitution, and regulation trades off limiting platform extraction against maintaining incentives for socially valuable platform provision. In Wang and Wright (2025), prices differ across channels, so fees distort consumers' choice between the platform and the direct channel; efficient regulation balances efficient platform use against the lower seller margins generated by on-platform competition. This paper instead allows the platform to adjust recommendation design in response to fee regulation.

2 Model

There is a unit mass of consumers, a finite number of firms $n \geq 2$, and a single platform P that intermediates between them. Each consumer has unit demand and interacts with firms only through the platform. Firms are indexed by $i \in \{1, \ldots, n\}$.

Each firm offers a single horizontally differentiated product. For each consumer j, let v_{ij} denote the match value from firm i's product. The random variables $\{v_{ij}\}_{i,j}$ are

i.i.d. across consumers and firms with common CDF F. The distribution F has compact support $[\underline{v}, \bar{v}] \subset \mathbb{R}$, with $\underline{v} < \bar{v}$, and admits a strictly positive, log-concave density f on $(\underline{v}, \bar{v})$. Each firm produces at constant marginal cost $c \in (0, \bar{v})$.

Consumers learn prices and match values through sequential inspections. The first product inspected is selected by the platform's recommendation rule, described below. This initial inspection is free and reveals the recommended product's price and match value. The consumer may then inspect further products sequentially. Each additional inspection costs $s > 0$ and reveals the price and match value of one previously unseen firm, drawn uniformly at random from the remaining firms. Recall is free: after any history, the consumer may buy any inspected product or stop without buying.

Consumer j's utility from purchasing product i at price p_i is $v_{ij} - p_i$, while the outside option of not purchasing yields utility 0. If she has conducted k additional inspections beyond the initial recommendation, she incurs search cost ks. Her payoff is therefore $v_{ij} - p_i - ks$ if she purchases product i, and $-ks$ if she stops without purchasing.

The platform observes the entire vector of match values $v_j = (v_{1j}, \ldots, v_{nj})$ for each consumer and prices $p_1, \ldots, p_n$, while consumers observe (v_{ij}, p_i) only for products they inspect. Firms do not observe individual consumers' match values or search histories.

The platform chooses two policy instruments: a commission rate $\tau \in [0, 1)$ and a recommendation algorithm indexed by $\alpha \in [0, \infty]$. A sale at price p_i gives the platform τp_i and firm i retains $(1 - \tau)p_i$ before production costs. Section 3.5 considers nominal per-transaction fees.

For $\alpha < \infty$, given the match-value vector v_j and the price vector p, the platform assigns product i the score

$$\sigma_{ij} = v_{ij} - \alpha p_i$$

and recommends a score-maximizing product, breaking ties by the highest v_{ij}. The parameter α governs the price sensitivity of recommendations: $\alpha = 0$ recommends the product with the highest match value, and $\alpha = 1$ the product that maximizes consumer utility. When $\alpha = \infty$, the platform recommends the cheapest product, breaking ties by match value. This is the natural limit of the finite-α rule.

Given a price profile $p = (p_1, \ldots, p_n)$, recommendation parameter α, and consumer strategy, let $D_i(p; \alpha)$ denote the measure of consumers who purchase from firm i. Since the consumer population has unit mass, $D_i(p; \alpha)$ is equivalently the probability that a representative consumer purchases from firm i. Firm i's profit is

$$\pi_i(p_i, p_{-i}; \alpha, \tau) = \big((1 - \tau)p_i - c\big)\, D_i(p; \alpha). \tag{1}$$

The platform has zero marginal cost and earns commission revenue

$$\Pi_P(\alpha, \tau; p) = \tau \sum_{i=1}^{n} p_i \, D_i(p; \alpha). \tag{2}$$

Total welfare is defined as the sum of consumer surplus (net of search costs), firms' profits, and the platform's profit. Because the commission is a transfer between firms and the platform, total welfare equals expected gross consumption value minus production costs and realized search costs.

The game proceeds as follows.

1. **Platform commitment.** The platform publicly commits to a commission rate τ and a recommendation parameter α.

2. **Firms' pricing.** After observing (α, τ), firms simultaneously choose prices, with firm i choosing p_i.

3. **Search and purchase.** For each consumer, match values $(v_{1j}, \ldots, v_{nj})$ are realized. The platform observes v_j and p, applies the scoring rule, and recommends one product. The consumer observes the recommended product and (v_{ij}, p_i) and then decides whether to (i) buy one of the products she has inspected so far, (ii) stop without buying, or (iii) pay the search cost s to inspect another, randomly selected, unseen product. After each additional inspection, the consumer can stop without buying, buy any previously inspected product, or continue searching.

We study symmetric pure-strategy perfect Bayesian equilibria. Anticipating subsequent pricing and consumer behavior, the platform chooses (α, τ) to maximize its profit. Given the platform policy, rivals' strategies, and consumer behavior, each firm chooses its price to maximize profit, while consumers choose search and purchase actions sequentially optimally given their beliefs. On-path beliefs are Bayes-consistent. Following Zhou (2014), after observing an off-equilibrium price, consumers continue to believe that unseen firms charge the equilibrium price; beliefs about unseen match values are determined by F, the recommendation rule, and the observed history. Symmetry requires all firms to use the same pricing strategy and hence to charge the same price on the equilibrium path. As Section 3.1 shows, this restriction is consequential because it eliminates on-path search.

2.1 Discussion

Linear scoring. The linear score $v_{ij} - \alpha p_i$ provides a tractable reduced-form measure of the price sensitivity of recommendations. This interpretation is not limited to deterministic linear rules. For a deterministic but nonlinear rule, α reflects the average marginal

price sensitivity in cases where a small price change can alter which product is recommended. For a probabilistic rule, it reflects the expected local response of a product's recommendation probability to its price. The linear specification thus captures the local pricing incentives generated by a broader class of recommendation algorithms, even if it does not reproduce their behavior following large price changes.

The platform's commitment power. The platform commits to its recommendation rule and fee, (α, τ), before firms set prices, and this policy is common knowledge in the pricing subgame. For the fee τ, this assumption is natural in markets where platforms publicly post their fee schedules. By contrast, recommendation algorithms are more complex and proprietary, and are typically not publicly observable. The assumption that market participants know the relevant price sensitivity of recommendations can therefore be viewed as a reduced-form representation of environments in which consumers and sellers repeatedly interact with the platform and learn about its recommendation policy.

Search on the platform. Random sequential search is one way to model consumers' ability to inspect products beyond the platform's recommendation. For equilibrium pricing, however, what matters is only that some friction prevents further inspection on the symmetric equilibrium path and after sufficiently small observed price deviations. Local pricing incentives therefore depend on how prices affect recommendation and purchase probabilities, not on whether the friction arises from direct search costs, behavioral inattention, or costly information acquisition.

3 Platform and Firm Behavior without Regulation

3.1 Consumer Search

In the final stage of the game, consumers observe the platform's recommendation and decide whether to purchase, search further, or take the outside option.

On the symmetric-equilibrium path, all firms charge the same price. Scores, therefore, differ only by match values, so the platform recommends the product with the highest match value. Conditional on observing this recommendation, a consumer cannot improve upon it by searching, and thus does not search in equilibrium.

Lemma 1 (No Search on Path)
In any symmetric equilibrium, consumers do not search on the symmetric-equilibrium path and purchase the recommended product if and only if its match value exceeds the price.

Moreover, for every finite α, sufficiently small unilateral price deviations do not induce search. Consider a deviation by firm i from p^* to p_i. If the firm is no longer recommended,

the deviation is initially unobserved. If it remains recommended, then $v_{ij} - \alpha p_i \geq v_{\ell j} - \alpha p^*$ for every unseen rival ℓ, so $v_{\ell j} \leq v_{ij} - \alpha(p_i - p^*)$. Hence the utility gain from inspecting an unseen rival is at most $(1-\alpha)(p_i - p^*)$, and therefore no more than $|1-\alpha|\,|p_i - p^*|$. Since $s > 0$, this gain is smaller than the search cost for all p_i sufficiently close to p^*. Thus, search is inactive on and locally around the symmetric-equilibrium path, although sufficiently large observed deviations may induce search.

3.2 Firms' Pricing

Benchmark: Prohibitive Search Costs. We first consider the benchmark $s = \infty$, in which consumers never inspect a product beyond the platform's recommendation. Besides being analytically tractable, this benchmark turns out to characterize the symmetric equilibrium pricing for every strictly positive search cost.

In this benchmark, firm i's demand, when its rivals charge p^*, is

$$D_i^{\infty}(p_i, p^*; \alpha) = \int_{p_i}^{\bar{v}} F\big(v - \alpha p_i + \alpha p^*\big)^{n-1} f(v)\, dv. \tag{3}$$

Demand equals the probability that (i) the consumer's value for product i exceeds its price, $v_{ij} \geq p_i$, and (ii) product i attains the highest recommendation score:

$$v_{ij} - \alpha p_i \geq v_{\ell j} - \alpha p^* \qquad \text{for all } \ell \neq i.$$

Given (α, τ), firm i therefore solves

$$\max_{p_i} \big((1-\tau)p_i - c\big) D_i^{\infty}(p_i, p^*; \alpha).$$

Lemma 2 (Pricing with Prohibitive Search Costs)
Fix $\alpha \in [0, \infty)$ and $\tau \in [0, 1)$, and suppose $(1-\tau)\bar{v} > c$. Then:

(i) the pricing game with $s = \infty$ admits a symmetric equilibrium;

(ii) a price p^{∞} is a symmetric equilibrium price if and only if

$$p^{\infty} - \frac{c}{1-\tau} = \frac{1 - F(p^{\infty})^n}{n} \frac{1}{F(p^{\infty})^{n-1} f(p^{\infty}) + \alpha(n-1) \displaystyle\int_{p^{\infty}}^{\bar{v}} F(v)^{n-2} f(v)^2\, dv}. \tag{4}$$

The $s = \infty$ benchmark is tractable because consumers never inspect another product after observing the recommendation. Under log-concavity, each firm has a unique best response to every conjectured rival price. A price is therefore a symmetric equilibrium

price if and only if it satisfies the symmetric first-order condition (4).[3] At a symmetric price p, the recommended product has the highest of the n match values. Aggregate demand is therefore $1 - F(p)^n$, the survival probability of the highest order statistic, and each firm's demand is $\left[1 - F(p)^n\right]/n$. The first-order condition shows that the markup $p - \frac{c}{1-\tau}$ equals firm demand divided by the marginal demand loss from a unilateral price increase. The term $F(p)^{n-1}f(p)$ captures the purchase margin: it is the density that firm i's match value equals p while all rival match values lie below p. The integral captures the recommendation margin. Raising p_i lowers firm i's score and causes it to lose consumers for whom its match value is close to the highest rival match value, whose density is $(n-1)F(v)^{n-2}f(v)$. Thus, the recommendation effect is governed by the gap between the highest and second-highest match values. A larger α strengthens this effect, makes demand more elastic, and compresses the equilibrium markup.

Finite search costs. Returning to $s \in (0, \infty)$, the pricing equilibria turn out to coincide with those of the $s = \infty$ benchmark.

Proposition 1 (Pricing with Finite Search Costs)
Fix $\alpha \in [0, \infty)$, $s \in (0, \infty)$, and $\tau \in [0, 1)$, and suppose $(1-\tau)\bar{v} > c$. Then:

(i) a price p is a symmetric equilibrium price if and only if it is a symmetric equilibrium price in the $s = \infty$ benchmark;

(ii) equivalently, p is a symmetric equilibrium price if and only if it satisfies (4);

(iii) a symmetric pricing equilibrium exists.

To see the first direction of the equivalence, consider a symmetric equilibrium price of the $s = \infty$ benchmark. Local deviations yield the same payoff when $s > 0$, while larger deviations yield a weakly lower payoff because consumers may inspect and purchase another product. Hence, the price remains an equilibrium for every $s > 0$.

Conversely, every symmetric equilibrium price for $s > 0$ is an interior local maximizer of the $s = \infty$ profit function, because sufficiently small deviations do not induce search. Strict log-concavity of the benchmark profit then implies that it is a global maximizer. Thus, the two environments have the same set of symmetric equilibrium prices.

Although equation (4) does not guarantee uniqueness in general, uniqueness can be shown for the power distribution on $[\underline{v}, \bar{v}]$ with shape parameter $\theta \geq 1$, whose CDF and density are

$$F(v) = \left(\frac{v - \underline{v}}{\bar{v} - \underline{v}}\right)^{\theta}, \qquad f(v) = \frac{\theta}{\bar{v} - \underline{v}} \left(\frac{v - \underline{v}}{\bar{v} - \underline{v}}\right)^{\theta - 1}.$$

[3] Although each firm's best response to a given rival's price is unique, (4) may admit multiple solutions, so the pricing game need not have a unique symmetric equilibrium price.

Remark 1 (Uniqueness under the Power Distribution)
Suppose

$$\frac{c}{1-\tau} \geq \underline{v}.$$

For the power distribution with shape parameter $\theta \geq 1$, the $s = \infty$ pricing game has a unique symmetric equilibrium. Consequently, the pricing game has a unique symmetric equilibrium price for every $s \in (0, \infty)$.

The pricing characterization (4) applies to finite α. When $\alpha = \infty$, the platform recommends the lowest-priced product following an asymmetric price deviation rather than the maximizer of a finite-α score. The local argument above therefore does not apply directly. Pricing, however, becomes Bertrand-like: the cheapest firm receives all recommendations and thus has an incentive to undercut.

Lemma 3 (Pricing at $\alpha = \infty$)
Fix $\tau \in [0, 1)$ such that

$$\frac{c}{1-\tau} < \bar{v},$$

and let $\alpha = \infty$. Then:

(i) for every $s \in (0, \infty]$, there is a unique symmetric equilibrium price,

$$p^* = \frac{c}{1-\tau};$$

(ii) any sequence of finite-α symmetric equilibrium prices p_α converges to p^* as $\alpha \to \infty$.

3.3 Platform's Choice of α and τ

We first fix α and characterize the platform's optimal commission before turning to the joint decision of (α, τ).

Fixed recommendation rule. By Proposition 1, for finite α the equilibrium price satisfies

$$p^* - \frac{c}{1-\tau} = m(\alpha, p^*), \tag{5}$$

with $m(\alpha, p)$ defined by (4); at $\alpha = \infty$, Lemma 3 implies zero downstream markup. The decomposition makes the double-marginalization mechanism transparent. The commission raises firms' effective marginal cost above production cost, while firms add a downstream markup. Recommendation design affects the second wedge: increasing α makes demand more price-sensitive and compresses $m(\alpha, p)$. For the remainder of this subsection, fix $\alpha < \infty$ and write $m(p) := m(\alpha, p)$. Given the recommendation price sensitivity α and commission rate τ, any symmetric equilibrium price p is pinned down

by (5). Conversely, for any candidate price p with $p - m(p) \geq c$, the markup condition implies

$$p \;=\; \frac{c}{1-\tau} + m(p) \quad \Longleftrightarrow \quad \tau(p) \;=\; 1 - \frac{c}{p - m(p)}.$$

Therefore, the platform's maximization problem, from (2), given a fixed α and the symmetric pricing subgame, can be rewritten. Substituting $\tau(p)$, then for any price p with $p - m(p) \geq c$, the platform's payoff can be written as

$$\Pi_P(p) = [p - \tilde{c}(p)]\,(1 - F(p)^n)\,, \qquad \tilde{c}(p) := c + \frac{c\,m(p)}{p - m(p)}. \tag{6}$$

Lemma 4
For every $\alpha \geq 0$, an optimal commission exists.

Equation (6) is reminiscent of the standard monopoly problem: conditional on a cost level x, a vertically integrated monopolist chooses p to maximize

$$(p - x)\big(1 - F(p)^n\big).$$

The difference is that here the relevant cost term $\tilde{c}(p)$ is itself endogenous and varies with the induced price through the equilibrium markup function $m(\alpha, p)$. Accordingly, for finite α the platform does not face a standard monopoly problem. Nevertheless, the monopoly comparison can be useful.

Definition 1 (Vertically Integrated Monopoly Price)
For any $x \in [0, \overline{v})$, define

$$p_m(x) \in \arg\max_{p \in [\underline{v}, \bar{v}]} (p - x)\big(1 - F(p)^n\big).$$

That is, $p_m(x)$ is the common price chosen by a vertically integrated monopolist with marginal cost x.

Under the maintained log-concavity assumptions, the monopoly problem has the following standard properties. For every $x \in [0, \bar{v})$, $p_m(x)$ is uniquely defined, satisfies $p_m(x) \in (x, \bar{v})$, and is strictly increasing in x. Moreover, $p \mapsto (p - x)(1 - F(p)^n)$ is strictly increasing on $(x, p_m(x))$ and strictly decreasing on $(p_m(x), \bar{v})$. These properties follow from standard log-concavity and are proven in Online Appendix OA.1.

Endogenous choice of α. When the platform can choose α, it internalizes the effect of recommendation design on downstream competition and therefore on the size of the surplus available for fee extraction. The next proposition shows that the platform pushes the design of recommendations to the extreme.

Proposition 2 (Unregulated Platform Choice)
The game admits a unique symmetric equilibrium outcome. In that equilibrium,

$$\alpha^* = \infty, \qquad p^* = p_m(c) = \frac{c}{1-\tau^*}, \qquad \tau^* = 1 - \frac{c}{p_m(c)}.$$

Recommendation design affects the extent of double marginalization by shaping the price sensitivity of demand. When α is low, sellers retain market power and charge positive markups, so the final price reflects both downstream market power and the platform's commission. Increasing α makes demand more price-sensitive and compresses those markups. Hence, any finite α leaves firms with positive markups that the platform can profitably eliminate. Proposition 2 shows that the platform therefore pushes recommendation design to the extreme: it chooses $\alpha^* = \infty$, eliminates downstream markups, and sets τ^* to replicate the vertically integrated monopoly price $p_m(c)$. Thus, double marginalization is eliminated.

3.4 Comparative Statics in n

This subsection studies how entry affects markups and equilibrium prices. An increase in the number of firms has two opposing effects. On the one hand, each firm faces more rivals, which tends to intensify price competition and compress markups. On the other hand, a larger product space allows the platform to recommend a better match on average, raising consumers' willingness to pay. This selection effect gives firms an incentive to raise prices conditional on being recommended and, when the platform adjusts its fee, increases the surplus the platform can extract through higher fees. To isolate these forces, the analysis begins with a fixed platform fee, then allows the platform to adjust its fee.

Lemma 5 (Vanishing Markup)
Fix $\alpha > 0$. For each $n \geq 2$, let $\tau_n \in [0, 1)$ and let $p_n \in [c, \bar{v})$ satisfy

$$p_n - \frac{c}{1-\tau_n} = m_n(\alpha, p_n).$$

Then

$$m_n(\alpha, p_n) \to 0 \qquad \text{as } n \to \infty.$$

Lemma 5 characterizes the large-n behavior of equilibrium markups. For any fixed fee τ and any price-sensitive recommendation rule $\alpha > 0$, the equilibrium markup converges to zero as $n \to \infty$. Intuitively, when rankings respond to prices, the gain from undercutting becomes stronger as the number of competitors grows, so pricing becomes increasingly competitive. By contrast, when firms do not compete on price through the recommendation rule ($\alpha = 0$), this effect is naturally absent. In that case, conditional on

being selected, a firm prices as a monopolist. Since $p_m(\frac{c}{1-\tau}) \to \bar{v}$ as $n \to \infty$, the price charged upon selection approaches $\bar{v}$.

Lemma 5 is asymptotic and does not by itself determine how prices or markups move for finite n. For moderate n, entry has two opposing effects: it intensifies price competition, but it also improves the expected match quality of the recommended product and thereby raises willingness to pay. The next proposition shows that when $\alpha \geq 1$, the first force dominates for the markup function under an additional condition on F.

Proposition 3

Fix $\alpha \geq 1$. Suppose, in addition to the assumptions in Section 2, that f is absolutely continuous and that $f'(v) \geq 0$ for a.e. $v \in (\underline{v}, \bar{v})$. For each $n \geq 2$, let $m_n(p) := m(\alpha, p)$ denote the markup in (4) when there are n firms. Then, for every $p < \bar{v}$,

$$m_{n+1}(p) < m_n(p).$$

Proposition 3 shows that entry lowers $m_n(p)$ pointwise. Thus, at any fixed final price, the stronger competition generated by an additional rival dominates the increase in willingness to pay resulting from better expected match quality.

The case $\alpha = 1$ provides a useful special case. The platform then recommends the product that maximizes consumers' net utility, so the resulting purchase allocation is the same as under full information despite the search friction. The proposition shows that entry nevertheless compresses firms' markups at any given price whenever recommendations are at least as price-sensitive as under full information. For $\alpha > 1$, the recommendation rule places still greater weight on price, strengthening competition for being recommended. By contrast, when $\alpha < 1$, recommendations respond less strongly to prices, and the improvement in expected match quality may offset the competitive effect of entry.

The result remains pointwise: it does not by itself determine the induced equilibrium price, nor does it incorporate the platform's fee response. Proposition 4 therefore turns to equilibrium outcomes and studies the asymptotic behavior of prices and consumer surplus under fixed and endogenous platform fees.

Proposition 4 (Comparative Statics in n)

Fix $\alpha > 0$. For each $n \geq 2$, let p_n be a symmetric equilibrium price and let CS_n be the corresponding consumer surplus.

(i) ***Fixed platform fee.*** *Fix $\tau \in [0, 1)$ and assume $(1-\tau)\bar{v} > c$. Then, as $n \to \infty$,*

$$p_n \longrightarrow \frac{c}{1-\tau} \qquad \text{and} \qquad CS_n \longrightarrow \bar{v} - \frac{c}{1-\tau}.$$

(ii) ***Endogenous platform fee.*** *Suppose the platform chooses τ optimally given α.*

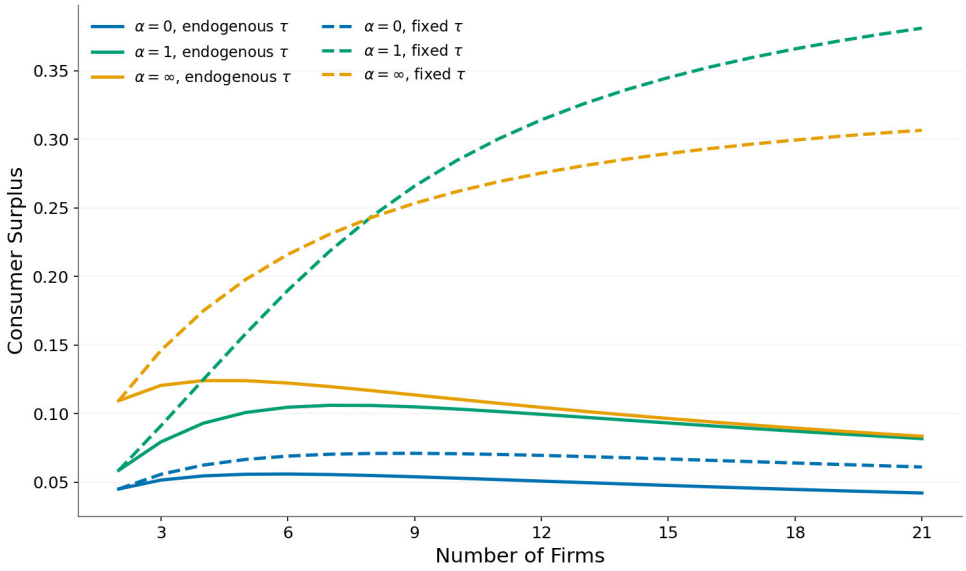
Then, as $n \to \infty$,

$$p_n \longrightarrow \bar{v} \qquad \textit{and} \qquad CS_n \longrightarrow 0.$$

Proposition 4 shows that the role of entry depends on whether the platform adjusts its fee. If the platform keeps τ fixed, increased competition eventually drives seller markups to zero, so the equilibrium price converges to the effective marginal-cost floor $c/(1-\tau)$ and consumer surplus converges to $\bar{v} - c/(1-\tau)$. The selection effect discussed above still operates for finite n: with more products, the platform can recommend a better match on average, which raises willingness to pay, but not enough to overturn this limiting result.

By contrast, when the platform chooses τ optimally, the improvement in match quality becomes a source of additional surplus extraction. As n grows, consumers' willingness to pay rises because the platform can select from a larger set of products; the platform responds by increasing its fee, so the induced price converges to $\bar{v}$ and consumer surplus converges to zero.

Figure 1: Consumer Surplus as a Function of the Number of Firms



$F \sim \mathrm{U}[0,1],\ c = 0.1$

Figure 1 illustrates these forces for $F \sim U[0,1]$ and $c = 0.1$. When τ is fixed at the platform-optimal τ for $n = 2$ and $\alpha > 0$, consumer surplus rises with n in this example as improved matching and markup compression dominate. With $\alpha = 0$, entry does not intensify competition for recommendations, so the gains are smaller and eventually reverse. When the platform instead adjusts τ optimally, the initial consumer-surplus gains are eventually offset by the endogenous fee response. For every $\alpha > 0$, consumer surplus consequently converges to zero.[4]

[4]This result is driven primarily by the platform's fee-setting power as an intermediary. As more firms participate, consumers can find better matches, but their surplus also becomes easier to predict and extract through higher prices. A similar force is present even without recommendation design, when

3.5 Alternative Monetization: Nominal Fees

Before turning to regulation, suppose that the platform charges a nominal per-transaction fee $t \geq 0$ instead of a commission. Firm i and the platform then earn

$$\pi_i^N = (p_i - c - t)D_i(\mathbf{p}; \alpha), \qquad \Pi_P^N = t \sum_{i=1}^{n} D_i(\mathbf{p}; \alpha).$$

For finite α, Proposition 1 applies with effective marginal cost $c+t$ in place of $c/(1-\tau)$. Hence every symmetric equilibrium price satisfies

$$p - c - t = m(\alpha, p). \tag{7}$$

When a platform policy supports multiple symmetric pricing equilibria, we select the one that maximizes platform profit. This avoids making the fee-form comparison depend on arbitrary equilibrium selection.[5]

Lemma 6 (Fee-Form Comparison for Finite α)
Fix $\alpha \in [0, \infty)$.

(i) The platform's maximal profit under commission pricing strictly exceeds its maximal profit under nominal-fee pricing.

(ii) Suppose that

$$p \longmapsto \frac{m(\alpha, p)}{p - m(\alpha, p)}$$

is weakly decreasing. Then, among platform-preferred equilibria, the highest equilibrium price under commission pricing is weakly lower than the highest equilibrium price under nominal-fee pricing.

The condition in part (ii) is satisfied, for example, under a constant absolute markup $m(\alpha, p)$ or a constant proportional markup $m(\alpha, p)/p$. It holds in particular if $0 \leq \alpha \leq 1$, f is absolutely continuous and $f' \geq 0$ a.e., or if F is the power distribution of Remark 1 with $\underline{v} = 0$.

For finite α, firms retain a positive downstream markup. A commission allows the platform to capture a share of this markup and therefore weakly raises its profit at any price implementable under both fee instruments, echoing Shy and Wang (2011). By partially internalizing the downstream markup, commissions also tend to mitigate double marginalization and induce lower prices.

the platform only charges a commission and consumers have to search through products; see Online Appendix OA.6.

[5] Equations (4) and (7) characterize the symmetric equilibrium prices under commissions and nominal fees but need not determine a unique price for every platform policy. If several platform-preferred equilibria remain, price comparisons use the highest equilibrium price. The selection rule is irrelevant whenever the pricing subgame has a unique symmetric equilibrium.

The monotonicity condition in part (ii) extends this intuition to endogenous markups: it requires the relative markup $\frac{m(\alpha,p)}{p-m(\alpha,p)}$ to decrease with the induced price. The platform's gain from sharing in the downstream markup is then greatest at lower prices, reinforcing the price-reducing effect of commissions. If the condition fails, the price ranking need not hold.

When recommendation design is endogenous, the platform eliminates the downstream markup under either fee instrument.

Proposition 5 (Optimal Nominal-Fee Policy)
If the platform chooses an unrestricted nominal fee $t \geq 0$ and recommendation parameter α, its unique optimal policy satisfies

$$\alpha^* = \infty, \qquad t^* = p_m(c) - c, \qquad p^* = p_m(c).$$

For a given t, equation (7) implies

$$p = c + t + m(\alpha, p) \geq c + t,$$

with equality at $\alpha = \infty$. Since nominal-fee profit is $t[1 - F(p)^n]$, the platform chooses $\alpha = \infty$ to maximize transaction volume. It then solves

$$\max_{t \geq 0} t\big[1 - F(c+t)^n\big] = \max_{p \geq c} (p - c)\big[1 - F(p)^n\big],$$

whose unique solution is $p_m(c)$.

Hence, fee form is irrelevant when both the fee and recommendation policy are unrestricted.

4 Regulation

Regulation breaks this equivalence because it constrains how the platform extracts surplus while leaving recommendation design under its control. Motivated by platform-fee caps that have been implemented or contemplated in several jurisdictions and markets, this section compares caps on commission rates with caps on nominal per-transaction fees. It also draws on Lemma 6 to evaluate a mandate to replace commissions with unrestricted nominal fees.

4.1 Caps on Commissions

Suppose the platform faces a commission cap $\bar{\tau} \in (0, 1)$, so that $\tau \leq \bar{\tau}$. Holding α fixed, a tighter cap lowers firms' effective marginal cost and therefore tends to reduce prices. Once recommendation design is endogenous, however, the platform can respond by lowering α,

thereby restoring seller markups. A commission cap therefore affects prices both directly through the fee and indirectly through the platform's recommendation policy.

Let

$$\tau^* = 1 - \frac{c}{p_m(c)}$$

denote the unconstrained optimal commission from Proposition 2. The next result characterizes the platform's response when the cap binds.

Lemma 7 (Platform Pricing under a Commission Cap)
Suppose the platform chooses (α, τ) *and faces a commission cap* $\tau \leq \bar{\tau}$.

If $\bar{\tau} < \tau^*$, *then the equilibrium price induced by an optimal platform choice is*

$$p^* = \max\left\{p_m(0), \frac{c}{1-\bar{\tau}}\right\}.$$

When the cap binds, the platform can no longer raise profits by increasing τ. Since τ is then fixed at $\bar{\tau}$, the platform chooses the price that maximizes platform revenue

$$\bar{\tau} p\big(1 - F(p)^n\big),$$

or equivalently $p(1-F(p)^n)$, namely $p_m(0)$. If the cap is only moderately tight, however, the platform may be unable to induce such a low price, because no equilibrium price below $c/(1-\bar{\tau})$ is feasible. Define

$$\tau^{**} = 1 - \frac{c}{p_m(0)},$$

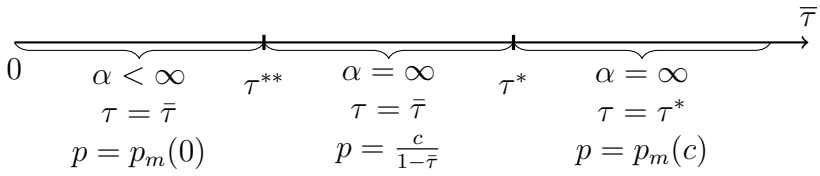
that is, the commission rate at which the effective marginal cost $c/(1-\tau)$ equals the monopoly revenue-maximizing price $p_m(0)$.

There are therefore two regions. For $\tau^{**} < \bar{\tau} < \tau^*$, tightening the cap lowers the induced price from the unregulated level $p_m(c)$ to $c/(1-\bar{\tau})$. Once $\bar{\tau} \leq \tau^{**}$, the lower bound $c/(1-\bar{\tau})$ has reached $p_m(0)$, so the platform can already sustain its preferred capped-regime price $p_m(0)$. Further tightening then leaves the induced price unchanged and therefore has no additional effect on consumer surplus or total welfare; it only reallocates profits from the platform to firms. Figure 2 shows how the induced price and recommendation policy vary with the commission cap.

Since a binding cap requires $\bar{\tau} < \tau^*$, Lemma 7 implies $p^* \leq p_m(c)$. This price reduction directly yields the following welfare ranking.

Corollary 1
Relative to the unregulated outcome, when the constraint $\tau \leq \bar{\tau}$ *binds, total welfare, consumer surplus, and firm profits weakly increase, while platform profit weakly decreases.*

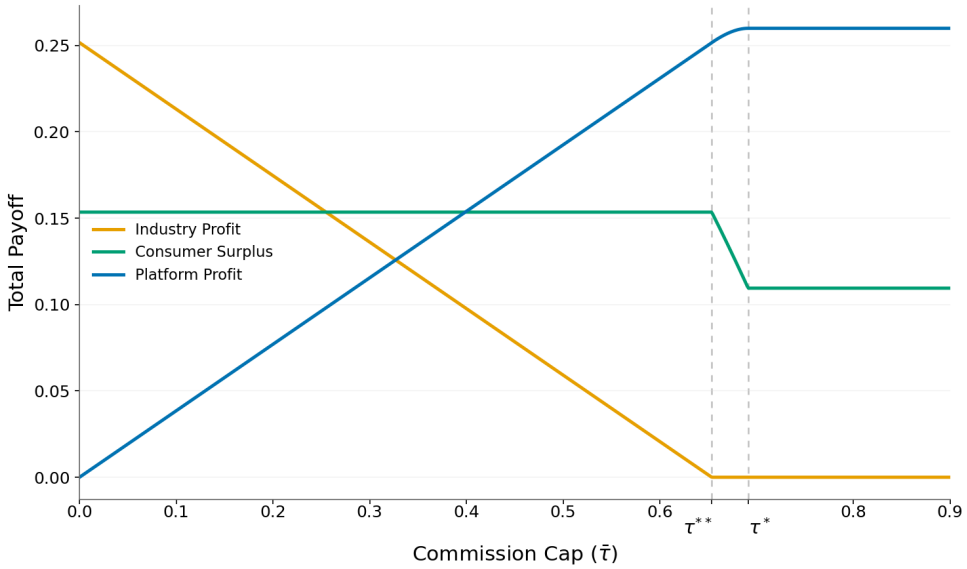
Figure 2: Pricing and Recommendation Regimes under a Commission Cap



The welfare gains from commission caps are limited by the endogenous adjustment of recommendation design. Even under a tight cap, the lowest price the platform chooses to induce is $p_m(0)$, the revenue-maximizing price under zero marginal cost. Figure 3 illustrates the corresponding effects on platform profit, firm profit, and consumer surplus.

For sufficiently high costs, $c \geq p_m(0)$, the threshold τ^{**} does not belong to $(0, 1)$.

Figure 3: Profits and Consumer Surplus as a Function of the Commission Cap $\bar{\tau}$



$F \sim U[0, 1]$, $n = 5$, and $c = 0.2$

4.2 Caps on Nominal Fees

Section 3.5 shows that replacing commissions with unrestricted nominal fees does not change the unregulated outcome when recommendation design is endogenous. Now suppose that the nominal per-transaction fee is capped:

$$t \leq \bar{t}.$$

For a given nominal fee t, the platform earns

$$t\big[1 - F(p)^n\big].$$

It therefore prefers the lowest price implementable at that fee. From (7),

$$p = c + t + m(\alpha, p) \geq c + t,$$

with equality at $\alpha = \infty$. The platform therefore chooses $\alpha = \infty$, induces $p = c + t$, and solves

$$\max_{t \leq \bar{t}} t\big[1 - F(c+t)^n\big].$$

Unlike a capped commission, a capped nominal fee aligns the platform's objective with transaction volume rather than transaction value.

Lemma 8 (Nominal-Fee Caps Dominate Commission Caps)
If

$$c < p_m(0) \qquad \textit{and} \qquad 0 < \bar{t} < p_m(0) - c,$$

then the equilibrium price under the nominal-fee cap is strictly lower than under any commission cap $\bar{\tau}$, and consumer surplus and total welfare are strictly higher. Moreover, when $p_m(0) - c > \bar{t} > \underline{v} - c$, lowering $\bar{t}$ further strictly lowers the equilibrium price and strictly raises consumer surplus and total welfare.

The condition $\bar{t} < p_m(0) - c$ implies that the nominal-fee cap binds and induces

$$p = c + \bar{t} < p_m(0).$$

By contrast, Lemma 7 implies that the platform never induces a price below $p_m(0)$ under a commission cap. Hence, a sufficiently tight nominal-fee cap yields strictly higher consumer surplus and total welfare. The following proposition summarizes the welfare implications of Lemmas 6, 7, and 8.

Proposition 6 (Comparative Welfare of Regulatory Regimes)

(i) Relative to the unregulated commission outcome, a binding commission cap weakly increases consumer surplus and total welfare;

(ii) for a fixed finite α, replacing commissions with unrestricted nominal fees can reduce consumer surplus and total welfare;

(iii) if $\bar{t} < p_m(0) - c$, a binding nominal-fee cap yields strictly higher consumer surplus and total welfare than any commission cap.

The difference between the two caps reflects the platform's residual objective. Under a commission cap, platform revenue remains proportional to transaction value, so the platform continues to benefit from high prices. Under a binding nominal-fee cap, revenue per transaction is fixed, so the platform instead maximizes transaction volume. In the

terminology of Teh (2022), proportional fees are seller-aligned, whereas capped nominal fees are volume-aligned.

4.3 Restricted Algorithm Design under Fee Regulation

The analysis so far has assumed that the platform can fine-tune the price sensitivity of its recommendation algorithm through a continuous choice of α. In practice, however, recommendation systems may offer only a small number of coarse design options, for example, ranking primarily by price or primarily by relevance rather than choosing an arbitrary intermediate weight between the two. To capture this possibility, suppose the platform can choose only between two regimes, $\alpha = 0$ and $\alpha = \infty$. In the absence of regulation, Proposition 2 continues to apply: the platform chooses the price-only rule and implements

$$(\alpha, \tau) = (\infty, \tau^*).$$

Proposition 7
Suppose the platform is restricted to $\alpha \in \{0, \infty\}$. For every $\tau \in [0, 1)$ satisfying $c/(1-\tau) < \bar{v}$, the pricing subgame admits a unique symmetric equilibrium for each $\alpha \in \{0, \infty\}$. Moreover, in the platform stage, the unique optimal choice is

$$(\alpha, \tau) = (\infty, \tau^*).$$

Under a binding cap on the commission rate, however, the inability to adjust α smoothly makes the platform's response discontinuous. The next result shows that, when algorithm design is restricted, a tighter cap can induce a discrete switch from the price-only rule to the match-value rule, raising prices and reducing consumer surplus and welfare. Recall from Proposition 2 that

$$\frac{c}{1-\tau^*} = p_m(c).$$

Proposition 8 (Binary Algorithm Design)
Suppose $\alpha \in \{0, \infty\}$ and $\bar{\tau} > 0$. Then:

(i) If $\bar{\tau} > \tau^*$, the cap is slack and $(\alpha, \tau) = (\infty, \tau^*)$.

(ii) If $c < p_m(0)$ and $\bar{\tau} \in [\tau^{**}, \tau^*]$, then $(\alpha, \tau) = (\infty, \bar{\tau})$ and $p = c/(1-\bar{\tau})$.

(iii) If

$$p_m(c)[1 - F(p_m(c))^n] > c[1 - F(c)^n],$$

there exists $\tau^{***} \in (0, \tau^{**})$ such that $\alpha = 0$ for every $\bar{\tau} \in (0, \tau^{***})$.

Whenever $\alpha = 0$ in equilibrium, $p > p_m(c)$ and both consumer surplus and total welfare are strictly lower than in the unregulated outcome.

A commission cap still creates an incentive to move toward less price-sensitive recommendations. When the platform can adjust α continuously, this response is smooth. Under restricted algorithm design, however, the optimal intermediate adjustment may be infeasible. The platform may then switch discretely to the match-value rule, leaving firms with more market power and thereby increasing prices. In that case, fee regulation can reduce both consumer surplus and total welfare.

5 Discussion and Extensions

The baseline model provides a sharp benchmark. When unconstrained, the platform chooses a maximally price-sensitive recommendation rule ($\alpha^* = \infty$) to eliminate seller markups and mitigate double marginalization, and then sets the commission rate to implement the vertically integrated monopoly price. Binding commission caps on τ raise consumer surplus and welfare in the benchmark, but the gains are attenuated because the platform can adjust α in response.

The benchmark also clarifies what drives the platform's strong incentive to make recommendations as price-sensitive as possible: in the baseline environment, stronger price sensitivity only helps by compressing downstream markups. In the extensions below, however, that force is offset by two countervailing considerations. Greater price sensitivity can weaken firms' incentives to invest in quality and increase the risk that the platform recommends an inferior product with a low purchase probability. Additional extensions in the Online Appendix study endogenous firm entry, positive platform marginal costs, and a benchmark without recommendation design; endogenous entry provides another force limiting recommendation price sensitivity because firms must retain rents to cover entry costs. Throughout this section, for every platform policy under consideration, the relevant subgame played by firms is assumed to admit a unique interior symmetric equilibrium. Moreover, this equilibrium is assumed to be the unique interior solution to the corresponding symmetric first-order conditions.

5.1 Endogenous quality

Quality investment is one factor that limits the platform's incentive to make highly price-sensitive recommendations. If price sensitivity is too high and equilibrium markups become too low, firms may no longer find it profitable to invest in product quality. The platform, therefore, faces a trade-off between intensifying price competition and preserving investment incentives.

Set-up. Consider the baseline setting, but now each firm chooses both its price and its product's quality, (p_i, q_i). Consumer j's valuation for firm i is $v_{ij} + q_i$, and marginal

cost is an increasing, convex function of quality, $c(q_i)$. The platform's recommendation algorithm incorporates quality through consumers' valuations, so the scoring rule remains a function of the quality-adjusted match value and price:

$$\sigma_{ij} = v_{ij} + q_i - \alpha p_i.$$

Assume that the market is not fully covered, i.e., $p^* - q^* \geq \underline{v}$. Given that all other firms choose the symmetric pair (p^*, q^*), firm i's profit is

$$\Pi_i = \big((1-\tau)p_i - c(q_i)\big) \int_{p_i - q_i}^{\bar{v}} \big[F(v + q_i - q^* - \alpha(p_i - p^*))\big]^{n-1} f(v)\, dv. \tag{8}$$

In a symmetric equilibrium $(p_i, q_i) = (p^*, q^*)$, the first-order conditions for price and quality imply that the common quality level q^* must satisfy

$$c'(q^*) = (1-\tau)\, \frac{F(p^*-q^*)^{n-1} f(p^*-q^*) + \displaystyle\int_{p^*-q^*}^{\bar{v}} (n-1) F(v)^{n-2} f(v)^2\, dv}{F(p^*-q^*)^{n-1} f(p^*-q^*) + \alpha \displaystyle\int_{p^*-q^*}^{\bar{v}} (n-1)\, F(v)^{n-2} f(v)^2\, dv}. \tag{9}$$

Equation (9) shows that two forces govern the firm's choice of quality. If τ is large, incentives to invest in quality are weak, because the firm captures only the smaller fraction $1-\tau$ of any price increase needed to recoup the higher cost. Furthermore, a higher α reduces the incentive to invest, because any price increase associated with higher quality is penalized more strongly in the recommendation score.

To illustrate this trade-off in a tractable finite-n environment, consider the case in which $v_{ij} \sim \mathrm{U}[0,1]$ and quality costs are quadratic. In that environment, equilibrium can be parameterized explicitly, and when quality investment is sufficiently cheap, the platform strictly prefers not to make recommendations too price-sensitive: any profit-maximizing policy satisfies $\alpha^* \in [0, 1)$ and hence never chooses $\alpha > 1$.

Proposition 9
Assume $v_{ij} \overset{iid}{\sim} \mathrm{U}[0,1]$ and costs are quadratic, $c(q) = \frac{k}{2}q^2$ with $k > 0$. If $k < \frac{n}{3}$, then any profit-maximizing platform policy satisfies $\alpha^* \in [0, 1)$.

The standard uniform-distribution example shows that when quality investment is sufficiently cheap, the platform prefers to limit α so that firms retain sufficiently strong incentives to invest in quality.[6]

[6] A cap on the commission rate can distort the platform's joint choice of (α, τ). Simulations with a uniform distribution for match values and polynomial quality costs suggest welfare need not be monotone in the cap $\overline{\tau}$, while consumer surplus increases as the cap tightens.

5.2 Vertical differentiation: an inferior product

A second force that can limit the platform's incentive to make rankings highly price-sensitive arises under vertical differentiation. Suppose the platform's candidate set contains a clearly inferior product. If the recommendation rule places too much weight on price, the platform may rank this product highly simply because it is cheap, even though no consumer would ever buy it. The platform, therefore, trades off markup compression against the risk of misrecommendations.

Set-up. In addition to the n regular products, suppose the platform's candidate set contains one inferior product, denoted U (for "useless"), with fixed price p_U and match value $v_U \leq p_U$. We assume that consumers choose the outside option when indifferent between purchasing and not purchasing. Hence U is never purchased. Economically, product U captures a product that is irrelevant to the consumer, for example, because it is misclassified and not actually relevant to the consumer's search, but that can nevertheless be ranked highly by the algorithm when it is sufficiently cheap, crowding out relevant products.

The scoring rule is unchanged: regular product i receives score $\sigma_{ij} = v_{ij} - \alpha p_i$, while U has score $\sigma_U = v_U - \alpha p_U$. Fix a symmetric regular-product price p and define the associated U-cutoff

$$t_U(\alpha, p_U, v_U; p) := v_U + \alpha(p - p_U).$$

At price p, U outranks a regular product exactly when the regular product's match value is at most $t_U(\alpha, p_U, v_U; p)$.

Call a recommendation a misrecommendation if U is recommended even though at least one regular product would yield strictly positive utility at price p and hence would be purchased in the baseline model without U.

Lemma 9 (Misrecommendations)
Fix $\alpha \geq 0$ and (p_U, v_U) with $v_U \leq p_U$. At a symmetric regular-product price p, misrecommendations occur with positive probability if and only if

$$F(t_U(\alpha, p_U, v_U; p)) > F(p),$$

where

$$t_U(\alpha, p_U, v_U; p) = v_U + \alpha(p - p_U).$$

In particular, when $p \in [\underline{v}, \bar{v})$, this condition is equivalent to

$$t_U(\alpha, p_U, v_U; p) > p.$$

Lemma 9 identifies when the inferior product can distort the ranking. For prices

$p > p_U$, define the threshold

$$\bar{\alpha}(p) := \frac{p - v_U}{p - p_U}. \tag{10}$$

Then misrecommendations are impossible whenever $\alpha \leq \bar{\alpha}(p)$. If, in addition, $p \in [\underline{v}, \bar{v})$, misrecommendations occur with positive probability exactly when $\alpha > \bar{\alpha}(p)$. Since $v_U \leq p_U$, the map $\bar{\alpha}(\cdot)$ is weakly decreasing on (p_U, ∞).

The next lemma records two implications of this threshold for an optimal policy.

Lemma 10 (Endogenous Threshold and High-Priced Inferior Product)
Let (α^*, τ^*) be an optimal design, and let p^* denote the induced symmetric regular-product price.

(i) If $p^* > p_U$, then

$$\alpha^* \geq \bar{\alpha}(p^*) = \frac{p^* - v_U}{p^* - p_U}.$$

(ii) If $p_U > p_m(c)$, then

$$\alpha^* = \infty.$$

As long as $\alpha \leq \bar{\alpha}(p^*)$, the inferior product generates no misrecommendations and is therefore irrelevant for equilibrium trade and pricing. In that region, the logic is the same as in the baseline model: a higher α compresses seller markups and allows the platform to raise its fee. Hence, the platform benefits from increasing α.

A genuine trade-off arises only once α exceeds $\bar{\alpha}(p^*)$, so that misrecommendations may occur. When $p_U < p_m(c)$, the platform's optimal choice of α therefore depends on what happens after a misrecommendation of U. If search costs are high, consumers do not search after a recommendation of U, so a misrecommendation destroys trade. If search costs are low, a recommendation of U instead triggers search, and trade may still occur. The platform is therefore more reluctant to induce misrecommendations when search costs are high than when they are low. Technical details on demand, pricing, and search are provided in Online Appendix OA.3.

High search costs. Suppose consumers do not search after a recommendation of U. Then a misrecommendation destroys a sale that would otherwise occur under a regular recommendation. The online appendix OA.3.3 derives the corresponding pricing condition and shows that, relative to the baseline model without U, the markup term is weakly lower, with strict inequality whenever the inferior product can crowd out a purchasable regular product. This is because product U still competes with other firms for recommendations, thereby exerting competitive pressure on the seller's margin.

The possibility of a lost sale when equilibrium prices and α are high disciplines both the induced price and the platform's choice of α. The next proposition shows that the risk

of lost sales can overturn the platform's incentive to make recommendations maximally price-sensitive.

Proposition 10 (High Search Costs)
Suppose consumers do not search after a recommendation of U.

(i) If $p_U < c$, then any optimal design (α^*, τ^*) satisfies

$$\alpha^* < \infty.$$

(ii) If $p_U = 0$ and $v_U = 0$, then in any optimal design (α^*, τ^*) we have

$$\alpha^* = 1.$$

High search costs make recommending U particularly costly for the platform: if U is recommended, trade is lost rather than merely delayed by consumer search. The platform, therefore, faces a trade-off. A higher α compresses downstream markups but also increases the risk that an inferior product is recommended and crowds out a relevant one.

Part (i) identifies a parameter region in which this force is strong enough to imply a finite optimal recommendation weight. If product U is so cheap that $p_U < c$, then under $\alpha = \infty$ it is always the cheapest available option whenever a regular product could be sold, so a recommendation of U destroys trade. In that region, the platform strictly prefers some finite α.

More generally, a low p_U makes high values of α costly because they increase the risk that U crowds out relevant products. Outside $p_U < c$, however, α also changes equilibrium prices and hence the region in which misrecommendations occur, preventing a simple general characterization.

Part (ii) pins down the optimum in the benchmark case $p_U = v_U = 0$. Then $\bar{\alpha}(p^*) = 1$, so increasing α beyond 1 can only raise the risk that U crowds out relevant products. Since a recommendation of U destroys the sale in the high-search-cost regime, the platform optimally chooses the smallest α that eliminates the underlying ranking distortion.

Online Appendix OA.3.3.2 provides an asymptotic refinement for general (p_U, v_U): as $n \to \infty$, the platform's optimal price sensitivity converges to the limiting threshold $\bar{\alpha}_\infty(p_U, v_U)$. Intuitively, once downstream markups vanish, there is no longer any gain from making recommendations more price-sensitive if doing so increases the risk of a wrong recommendation and a lost sale.

Low search costs. Now suppose search costs are low, so that a recommendation of U triggers consumer search with positive probability. Then misrecommendations need not destroy trade: they induce search, and the consumer may still buy a regular product.

If a regular product is recommended, the consumer follows the recommendation, and search cannot improve on the top-ranked regular product. If instead U is recommended, the consumer faces a standard sequential search problem as in Wolinsky (1986), characterized by a reservation value $\hat{x}(\alpha, s, p^*)$ under the posterior induced by observing U.[7]

When search is nearly costless, exhausting the search pool becomes nearly optimal, so the consumer's final choice approaches the full-information rule $\arg\max_i\{v_{ij} - p_i\}$. Thus, for small s, a misrecommendation triggers search rather than eliminating trade. The next proposition makes this point transparent by considering $p_U = v_U = 0$.

Proposition 11

Let $v_U = p_U = 0$. Then there exists $\bar{s} > 0$ such that for all $s \in (0, \bar{s})$, every platform-optimal choice $\alpha^*(s)$ satisfies

$$\alpha^*(s) > 1.$$

When search is cheap, consumers endogenously split into two groups. A follower receives a regular recommendation and buys without searching; conditional on being a follower, a higher α strengthens competition and compresses markups. A searcher receives a recommendation of U, rejects it, and starts searching. As $s \downarrow 0$, searchers behave as if they compare products using $v - p$, that is, as if the relevant price weight were 1 regardless of the platform's α.

Increasing α therefore has two opposing effects. On the one hand, it raises the price sensitivity of demand within the follower group. On the other hand, it increases the probability of a misrecommendation and shifts consumers from followers to searchers. This composition effect can make an interior α optimal: a larger α intensifies competition among followers, but simultaneously creates more searchers whose behavior is effectively pinned down by the $v - p$ comparison. Thus, when search is cheap, the platform may optimally tolerate some misrecommendations because they trigger search rather than destroy trade.

To summarize, the vertical-differentiation extension highlights two forces shaping recommendation design: greater price sensitivity compresses seller markups, but it also raises the risk that cheap inferior products crowd out relevant alternatives. How the platform balances these forces depends on search costs. Richer models of vertical differentiation may introduce additional forces. In particular, limiting α can allow a more expensive product to remain competitive for recommendations rather than being mechanically displaced by a cheaper alternative. This may itself strengthen competition and compress

[7] Conditional on U being recommended, Bayes' rule pins down the posterior distribution of unseen regular-product values, F_U; see Online Appendix Section OA.3.2.2. By consistency, unseen regular products are still conjectured to be priced at p^*. Hence, the search problem after observing U is well defined. Whenever U is recommended with positive probability, sequential rationality requires the consumer's search behavior after observing U, and after each subsequent search outcome that can arise from it, to coincide with the optimal search rule. Outcomes not generated under that behavior are irrelevant to the characterization below.

markups, while also giving the platform greater scope to steer consumers toward higher-priced products. The overall effect of recommendation price sensitivity on competition and platform incentives is therefore presumably ambiguous once vertical differentiation is sufficiently rich.

6 Conclusion

This paper studies how platform fees and recommendation design interact. In the model, the platform uses its recommendation rule to shape the intensity of competition among sellers and then monetizes the resulting transactions through fees. In the unconstrained benchmark, the platform makes recommendations that are maximally price-sensitive to compress seller markups and mitigate double marginalization. Consumers, however, do not fully benefit from this stronger downstream competition because the platform adjusts its commission rate so that final prices coincide with those of an integrated monopolist.

This interaction also shapes the effects of regulation. A cap on the commission rate directly limits one channel of surplus extraction by the platform, but it also changes the platform's incentive to use recommendation design to discipline seller markups. Its effect, therefore, depends not only on the direct reduction in the fee but also on the platform's adjustment of the recommendation rule. In this sense, the analysis echoes a standard second-best logic: when regulation constrains only one margin of a multi-instrument environment, its effects depend on how the remaining margins adjust in equilibrium.

The form of platform monetization is therefore important. At the unconstrained optimum, commission rates and nominal per-transaction fees can yield the same outcome when the platform is free to choose fees and recommendation design. Under regulation, however, the two fee instruments create different incentives. When the commission rate is capped, the platform continues to benefit from high transaction prices. It can therefore respond by making recommendations less price-sensitive, relaxing competition among sellers, and attenuating the cap's gains in consumer surplus. By contrast, when a nominal per-transaction fee is capped, the platform's objective shifts toward expanding transaction volume rather than sustaining high transaction prices. In the model, this aligns the design of recommendations more closely with consumer welfare than a cap on the commission rate alone.

The model focuses on platform-mediated search among horizontally differentiated sellers in order to isolate how fee regulation interacts with recommendation design. This leaves several aspects for future work. The analysis abstracts from off-platform channels and data sharing. The vertical-differentiation extension studies a stylized environment in which recommendation design trades off markup compression against the risk that a cheap inferior product crowds out relevant alternatives and, under high search costs, prevents trade altogether. Richer models of vertical differentiation may introduce additional forces,

including the platform's incentive to steer consumers toward more expensive products and changes in which products effectively compete for recommendations. Moreover, the focus on ex-ante symmetric sellers and symmetric equilibria abstracts from richer forms of seller heterogeneity and the search behavior they may induce.

References

Administrative Code of the City of New York (2025). Section 20-563.3: Fee Caps. New York, NY. As amended by Local Law No. 79 of 2025, effective June 30, 2025.

Anderson, S. P. and Renault, R. (1999). Pricing, product diversity, and search costs: A Bertrand-Chamberlin-Diamond model. *The RAND Journal of Economics*, 30(4):719–735.

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

Bergemann, D. and Bonatti, A. (2011). Targeting in advertising markets: implications for offline versus online media. *The RAND Journal of Economics*, 42(3):417–443.

Competition Appeal Tribunal (2025). Dr. rachael kent v apple inc. and apple distribution international ltd. [2025] CAT 67, judgment dated 23 October 2025.

Crémer, J., Heidhues, P., Schnitzer, M., and Scott Morton, F. (2024). Apple's exclusionary App Store scheme: An existential moment for the Digital Markets Act. VoxEU Column, CEPR.

de Cornière, A. (2016). Search advertising. *American Economic Journal: Microeconomics*, 8(3):156–188.

de Cornière, A. and Taylor, G. (2019). A model of biased intermediation. *The RAND Journal of Economics*, 50(4):854–882.

Dinerstein, M., Einav, L., Levin, J., and Sundaresan, N. (2018). Consumer price search and platform design in internet commerce. *American Economic Review*, 108(7):1820–1859.

Eliaz, K. and Spiegler, R. (2016). Search design and broad matching. *American Economic Review*, 106(3):563–586.

Gomes, R. and Mantovani, A. (2025). Regulating platform fees under price parity. *Journal of the European Economic Association*, 23(1):190–235.

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

Hagiu, A., Teh, T.-H., and Wright, J. (2022). Should platforms be allowed to sell on their own marketplaces? *The RAND Journal of Economics*, 53(2):297–327.

Hagiu, A. and Wright, J. (2024). Optimal discoverability on platforms. *Management Science*, 70(11):7770–7790.

Heidhues, P., Köster, M., and Kőszegi, B. (2023). Steering fallible consumers. *The Economic Journal*, 133(652):1430–1465.

Inderst, R. and Ottaviani, M. (2012). Competition through commissions and kickbacks. *American Economic Review*, 102(2):780–809.

Karle, H., Preuss, M., and Reisinger, M. (2026). Selling on recommender platforms: Demand boost versus customer migration. CESifo Working Paper 12465, CESifo.

Kaye, A. P. (2026). The personalization paradox: Welfare effects of personalized recommendations in two-sided digital markets. Job market paper, Boston University and Massachusetts Institute of Technology.

Lee, K.-H. and Musolff, L. (2025). Two-sided markets shaped by platform-guided search. Accepted for publication in *Econometrica*.

Ng, R. (2025). Competition through recommendations. CRC TR 224 Discussion Paper No. 718.

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

Peitz, M. and Sobolev, A. (2026). Inflated recommendations. *The RAND Journal of Economics*, 57(2):368–385.

Rhodes, A. and Zhou, J. (2026). Personalization and privacy choice. *The RAND Journal of Economics*. Forthcoming.

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

Rochet, J.-C. and Tirole, J. (2006). Two-sided markets: a progress report. *The RAND Journal of Economics*, 37(3):645–667.

Seattle Municipal Code (2022). § 7.30.032: Fees and charges. Seattle, WA.

Shy, O. and Wang, Z. (2011). Why do payment card networks charge proportional fees? *American Economic Review*, 101(4):1575–1590.

Teh, T.-H. (2022). Platform governance. *American Economic Journal: Microeconomics*, 14(3):213–254.

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

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

Wolinsky, A. (1986). True monopolistic competition as a result of imperfect information. *The Quarterly Journal of Economics*, 101(3):493–511.

Zhou, J. (2014). Multiproduct search and the joint search effect. *American Economic Review*, 104(9):2918–2939.

A.1 Deferred Proofs of Sections 3 and 4

A.1.1 Proofs of Section 3

Proof of Lemma 2. The proof proceeds in three steps. First, for every conjectured rivals' price, the pricing problem under $s = \infty$ has a unique interior best response. Second, continuity of the best-response function and Brouwer's fixed-point theorem establish the existence of a symmetric equilibrium. Third, interiority and strict log-concavity yield the equilibrium characterization.

Let

$$\underline{p} := \frac{c}{1-\tau}, \qquad \mathcal{P} := [\underline{p}, \bar{v}].$$

A price below $\underline{p}$ yields a negative margin whenever a sale occurs, whereas $\underline{p}$ yields zero profit. Prices above $\bar{v}$ yield zero demand. It is therefore without loss to restrict prices to $\mathcal{P}$.

For a given rivals' price $p^* \in \mathcal{P}$, firm i chooses p to maximize

$$\Pi_i^{\infty}(p, p^*) = \big((1-\tau)p - c\big) q(p \mid p^*),$$

where

$$q(p \mid p^*) := \int_p^{\bar{v}} F(v - \alpha p + \alpha p^*)^{n-1} f(v)\, dv.$$

Step 1: Best responses. For fixed p^*, define

$$\phi(v, p) := \mathbf{1}_{\{v \geq p\}} F(v - \alpha p + \alpha p^*)^{n-1} f(v).$$

Because f is log-concave, its CDF F is log-concave on its support. Consequently,

$$F(v - \alpha p + \alpha p^*)^{n-1}$$

is log-concave in (v, p). The indicator of the convex set $\{(v, p) : v \geq p\}$ is log-concave in the extended-value sense. Hence ϕ is log-concave in (v, p), and Prékopa's theorem implies that $q(\cdot \mid p^*)$ is log-concave.

On the region where the margin and demand are strictly positive,

$$\log \Pi_i^{\infty}(p, p^*) = \log\big((1-\tau)p - c\big) + \log q(p \mid p^*).$$

The first term is strictly concave in p, while the second is concave. Therefore, $\Pi_i^{\infty}(\cdot, p^*)$ has at most one maximizer.

At the boundaries,

$$\Pi_i^{\infty}(\underline{p}, p^*) = \Pi_i^{\infty}(\bar{v}, p^*) = 0.$$

Since $\underline{p} < \bar{v}$, one can choose $p \in (\underline{p}, \bar{v})$ sufficiently close to $\underline{p}$ that the margin and demand are both strictly positive. Thus, the unique maximizer lies in the interior of $\mathcal{P}$. Denote it by $BR(p^*)$.

The objective is continuous in (p, p^*), and $\mathcal{P}$ is compact, so a maximizer exists. Because the maximizer is unique, Berge's maximum theorem implies that $BR(\cdot)$ is continuous.

Step 2: Existence. The function

$$BR : \mathcal{P} \longrightarrow \mathcal{P}$$

is continuous, and $\mathcal{P}$ is nonempty, compact, and convex. Brouwer's fixed-point theorem therefore gives a price $p^\infty \in \mathcal{P}$ such that

$$BR(p^\infty) = p^\infty.$$

Hence, the pricing game under $s = \infty$ admits a symmetric equilibrium.

Step 3: Characterization. An equilibrium price is interior and therefore satisfies

$$\begin{aligned} 0 &= \frac{\partial \Pi_i^\infty}{\partial p_i}(p_i, p^*)\bigg|_{p_i = p^* = p} \\ &= (1-\tau) q(p_i \mid p^*) + \big((1-\tau)p_i - c\big) \frac{\partial q}{\partial p_i}(p_i \mid p^*)\bigg|_{p_i = p^* = p}. \end{aligned}$$

By Leibniz's rule,

$$\begin{aligned} \frac{\partial q}{\partial p_i}(p_i \mid p^*) = &- F\big((1-\alpha)p_i + \alpha p^*\big)^{n-1} f(p_i) \\ &- \alpha(n-1) \int_{p_i}^{\bar{v}} F(v - \alpha p_i + \alpha p^*)^{n-2} f(v - \alpha p_i + \alpha p^*) f(v)\, dv. \end{aligned}$$

At symmetry,

$$q(p \mid p) = \int_p^{\bar{v}} F(v)^{n-1} f(v)\, dv = \frac{1 - F(p)^n}{n},$$

and

$$\frac{\partial q}{\partial p_i}(p \mid p) = -F(p)^{n-1} f(p) - \alpha(n-1) \int_p^{\bar{v}} F(v)^{n-2} f(v)^2\, dv.$$

Substitution into the first-order condition gives

$$(1-\tau)\frac{1 - F(p)^n}{n} = \big((1-\tau)p - c\big) \left[F(p)^{n-1} f(p) + \alpha(n-1) \int_p^{\bar{v}} F(v)^{n-2} f(v)^2\, dv \right],$$

which is equivalent to (4).

Conversely, suppose that p satisfies (4). It then satisfies the symmetric first-order condition. By the strict log-concavity established in Step 1, this condition is sufficient for p to be the unique best response to itself. Hence $BR(p) = p$, and p is a symmetric equilibrium price. □

Proof of Proposition 1. Let p^∞ be a symmetric equilibrium price under $s = \infty$. By Lemma 2, p^∞ is interior and yields strictly positive profit. Hence, a deviation to

$$p_i < \frac{c}{1-\tau}$$

cannot be profitable, since it yields nonpositive profit.

It remains to consider a unilateral deviation to

$$p_i \geq \frac{c}{1-\tau}$$

when all rivals charge p^∞. Firm i's demand with a finite search cost is weakly lower than its demand under $s = \infty$:

$$D_i(p_i, p^\infty) \leq D_i^\infty(p_i, p^\infty).$$

To see this, first suppose that another firm is recommended. The consumer then observes a product charging the conjectured equilibrium price and, under her equilibrium beliefs, does not initiate search. Firm i therefore makes no sale. If firm i is recommended, subsequent search can only cause the consumer to purchase another product or not to purchase; it cannot increase firm i's purchase probability relative to the case in which no further product can be inspected.

At the symmetric price, consumers do not search, and hence

$$D_i(p^\infty, p^\infty) = D_i^\infty(p^\infty, p^\infty).$$

Because p^∞ is optimal under $s = \infty$,

$$\begin{aligned}
&\big((1-\tau)p^\infty - c\big) D_i(p^\infty, p^\infty) \\
&\quad = \big((1-\tau)p^\infty - c\big) D_i^\infty(p^\infty, p^\infty) \\
&\quad \geq \big((1-\tau)p_i - c\big) D_i^\infty(p_i, p^\infty) \\
&\quad \geq \big((1-\tau)p_i - c\big) D_i(p_i, p^\infty).
\end{aligned}$$

Thus, every symmetric equilibrium price under $s = \infty$ remains a symmetric equilibrium price for every $s > 0$. Together with Lemma 2(i), this also proves existence.

Conversely, let p^* be a symmetric equilibrium price for some $s > 0$. By Lemma 1, there exists $\bar{\Delta} > 0$ such that consumers do not search following any observed deviation

satisfying

$$|p_i - p^*| \leq \bar{\Delta}.$$

Consequently, in a neighborhood of p^*, firm i's profit coincides with its profit under $s = \infty$. Since p^* is an equilibrium price for $s > 0$, it is therefore a local maximizer of the $s = \infty$ profit function.

The equilibrium price is interior. If $p^* < \frac{c}{1-\tau}$, on-path demand is positive and profit is negative, whereas a price above $\bar{v}$ yields zero profit. If $p^* = \frac{c}{1-\tau}$, a sufficiently small upward deviation does not induce search and yields positive profit. Finally, if $p^* \geq \bar{v}$, profit is zero, whereas an interior price yields positive profit: firm i is recommended and purchased on the positive-measure event that its match value is sufficiently high and its rivals' match values are sufficiently low. Hence,

$$p^* \in \left(\frac{c}{1-\tau}, \bar{v} \right).$$

The first-order condition therefore holds at p^*, so p^* satisfies (4). By Lemma 2(ii), it is a symmetric equilibrium price under $s = \infty$.

Hence, the finite-s pricing game and the $s = \infty$ benchmark have the same set of symmetric equilibrium prices, and these prices are precisely the solutions to (4). □

Proof of Remark 1. For $\theta \geq 1$, the power density is log-concave. Hence Lemma 2 and Proposition 1 imply that both the $s = \infty$ benchmark and original pricing games admit a symmetric equilibrium, and that every symmetric equilibrium price satisfies (4). Thus it suffices to show that, under the power distribution, (4) has at most one solution. Let

$$\underline{p} := \frac{c}{1-\tau}, \qquad L := \bar{v} - \underline{v}, \qquad x := \frac{p - \underline{v}}{L}.$$

Since every equilibrium price satisfies

$$p > \underline{p} = \frac{c}{1-\tau} \geq \underline{v}$$

and $p < \bar{v}$, we have $x \in (0, 1)$.

Under the power distribution,

$$F(p) = x^\theta \qquad \text{and} \qquad f(p) = \frac{\theta}{L} x^{\theta-1}.$$

Also let

$$\gamma := \theta n - 1, \qquad A := \alpha \frac{(n-1)\theta^2}{\gamma}, \qquad B := \theta - A.$$

Then $F(p)^n = x^{\gamma+1}$, $F(p)^{n-1} f(p) = \frac{\theta}{L} x^\gamma$, and, with $y = (v - \underline{v})/L$,

$$\int_p^{\bar{v}} F(v)^{n-2} f(v)^2 \, dv = \frac{\theta^2}{L} \int_x^1 y^{\theta n - 2} \, dy = \frac{\theta^2}{L\gamma}(1 - x^\gamma).$$

Hence the denominator in (4) equals $\frac{1}{L}(A + Bx^\gamma)$, so (4) becomes

$$p - \underline{p} = L \, \frac{1 - x^{\gamma+1}}{n(A + Bx^\gamma)}.$$

Dividing by $p - \underline{v} = Lx$ gives

$$H(x) := \frac{1 - x^{\gamma+1}}{nx(A + Bx^\gamma)} = 1 - \frac{\underline{p} - \underline{v}}{p - \underline{v}}.$$

The right-hand side is weakly increasing in p, hence in x. Therefore it remains to show that H is strictly decreasing on $(0, 1]$.

A direct calculation yields

$$H'(x) = \frac{U(x)}{x^2(A + Bx^\gamma)^2}, \qquad n\,U(x) = -A\big(1 + \gamma x^{\gamma+1}\big) - (\gamma + 1)Bx^\gamma.$$

If $B \geq 0$, then $U(x) < 0$ for all $x \in (0, 1]$, so $H'(x) < 0$.

Now consider $B < 0$. Since $B = \theta - A$, this is equivalent to $A > \theta$, and hence

$$0 < -\frac{B}{A} = \frac{A - \theta}{A} < 1.$$

Define $t := -B/A \in (0, 1)$. The point of this definition is simply that it rewrites the B-term in units of A: since $B = -At$, the expression can be written as

$$n\,U(x) = -A\,h(x), \qquad h(x) := 1 - (\gamma + 1)tx^\gamma + \gamma x^{\gamma+1}.$$

Thus it suffices to show that $h(x) > 0$ on $(0, 1]$. But

$$h'(x) = \gamma(\gamma + 1)x^{\gamma-1}(x - t),$$

so h is strictly decreasing on $(0, t)$ and strictly increasing on $(t, 1]$. Hence its unique minimum is attained at $x = t$, where

$$h(t) = 1 - t^{\gamma+1} > 0.$$

Therefore $h(x) > 0$ for all $x \in (0, 1]$, so $U(x) < 0$ and again $H'(x) < 0$.

Thus H is strictly decreasing on $(0, 1]$, while the right-hand side above is weakly increasing. Therefore (4) has at most one solution. Since a symmetric equilibrium ex-

ists and (4) has at most one solution, the $s = \infty$ benchmark has a unique symmetric equilibrium price. Proposition 1 extends uniqueness to every $s > 0$. □

Proof of Lemma 3. The price $\frac{c}{1-\tau}$ is a symmetric equilibrium: a deviation to some $p_i < \frac{c}{1-\tau}$ yields a negative margin on every sale, while a deviation to some $p_i > \frac{c}{1-\tau}$ yields zero demand, since under $\alpha = \infty$ the deviating product is never recommended, so the deviation is unobserved and consumers follow the equilibrium strategy.

It remains to rule out any other symmetric equilibrium price. If $p < \frac{c}{1-\tau}$, then every firm makes losses on every sale, so p cannot be an equilibrium. If $p \geq \bar{v}$, a deviation to any $q \in (\frac{c}{1-\tau}, \bar{v})$ yields positive demand at a positive margin. Hence consider $p \in (\frac{c}{1-\tau}, \bar{v})$.

Because search is costly and utilities have bounded support, there exists a cutoff

$$\widetilde{u}(p, s) \in [0, \bar{v} - p)$$

such that a consumer whose current utility is strictly above $\widetilde{u}(p, s)$ buys immediately rather than searching. Choose $\gamma > 0$ such that $\tilde{u}(p, s) + \gamma < \bar{v} - p$, and define

$$B_i(\gamma) := \Big\{ v_i - p > \tilde{u}(p, s) + \gamma,\ \max_{j \neq i}(v_j - p, 0) > v_i - p \Big\}.$$

By full support, $\Pr(B_i(\gamma)) =: \kappa > 0$.

Under the symmetric price p, consumers in $B_i(\gamma)$ buy a rival, not firm i: all firms charge the same price, some rival yields strictly higher utility, and on the symmetric path there is no search. If firm i deviates to $p - \varepsilon$, where $0 < \varepsilon < \min\{\gamma, p - \frac{c}{1-\tau}\}$, then firm i becomes uniquely cheapest and is recommended to every consumer. For every consumer in $B_i(\gamma)$, $v_i - p + \varepsilon > \tilde{u}(p, s)$, so the consumer buys immediately and does not search. Thus the deviation creates an additional demand mass of at least κ, uniformly for all sufficiently small $\varepsilon > 0$. Furthermore, every consumer who bought from firm i at the symmetric price p still buys from firm i after the deviation, since firm i's utility strictly increases while all rivals remain at price p. The margin loss is of order ε from undercutting, while the demand increase is bounded away from zero, so for ε sufficiently small the deviation is profitable. Therefore no symmetric price $p > \frac{c}{1-\tau}$ can be an equilibrium.

This proves part (i). For part (ii), let $\alpha_k \to \infty$ and let p_k be any sequence of finite-α_k symmetric equilibrium prices. By Proposition 1, each p_k satisfies (4), so

$$p_k - \frac{c}{1-\tau} = \frac{1 - F(p_k)^n}{n\Big(F(p_k)^{n-1} f(p_k) + \alpha_k (n-1) \int_{p_k}^{\bar{v}} F(v)^{n-2} f(v)^2 \, dv\Big)}.$$

Since the right-hand side is nonnegative, $p_k \geq \frac{c}{1-\tau}$. Moreover, by the interiority of equilibrium prices established in the proof of Proposition 1, $p_k < \bar{v}$. Thus $p_k \in \left[\frac{c}{1-\tau}, \bar{v}\right)$.

Since $[\frac{c}{1-\tau}, \bar{v}]$ is compact, it suffices to show that every convergent subsequence of (p_k)

has limit $\frac{c}{1-\tau}$. Let $p_k \to \hat{p} \in [\frac{c}{1-\tau}, \bar{v}]$ along a subsequence.

If $\hat{p} < \bar{v}$, then $\int_{\hat{p}}^{\bar{v}} F(v)^{n-2} f(v)^2 \, dv > 0$, so by continuity there is $\eta > 0$ such that $\int_{p_k}^{\bar{v}} F(v)^{n-2} f(v)^2 \, dv \geq \eta$ for all large k. Hence

$$0 \leq p_k - \frac{c}{1-\tau} \leq \frac{1}{n\alpha_k(n-1)\eta} \to 0,$$

so $\hat{p} = \frac{c}{1-\tau}$.

If instead $\hat{p} = \bar{v}$, then dropping the α_k-term in the denominator gives

$$0 \leq p_k - \frac{c}{1-\tau} \leq \frac{1 - F(p_k)^n}{nF(p_k)^{n-1} f(p_k)} = \frac{1 + F(p_k) + \cdots + F(p_k)^{n-1}}{nF(p_k)^{n-1}} \cdot \frac{1 - F(p_k)}{f(p_k)}.$$

The first factor tends to 1 as $p_k \uparrow \bar{v}$. It is claimed that

$$\frac{1 - F(v)}{f(v)} \to 0 \qquad \text{as } v \uparrow \bar{v}.$$

Suppose not. Since $(1 - F)/f$ is decreasing and finite at every interior point, it cannot converge to $+\infty$. Hence there exists $L > 0$ such that

$$\frac{1 - F(v)}{f(v)} \to L \qquad \text{as } v \uparrow \bar{v}.$$

Then, for all v sufficiently close to $\bar{v}$, $\frac{f(v)}{1-F(v)} \leq \frac{2}{L}$.

Using $\frac{d}{dv} \log(1 - F(v)) = -\frac{f(v)}{1-F(v)}$, it follows that $\frac{d}{dv} \log(1 - F(v)) \geq -\frac{2}{L}$

for all such v. Fix $v_0 < \bar{v}$ sufficiently close to $\bar{v}$. Integrating from v_0 to $v < \bar{v}$ gives

$$\log(1 - F(v)) - \log(1 - F(v_0)) \geq -\frac{2}{L}(v - v_0),$$

so

$$1 - F(v) \geq (1 - F(v_0)) e^{-2(v - v_0)/L}.$$

Letting $v \uparrow \bar{v}$ yields

$$1 - F(\bar{v}-) \geq (1 - F(v_0)) e^{-2(\bar{v} - v_0)/L} > 0,$$

contradicting $1 - F(\bar{v}) = 0$. Thus $(1 - F(v))/f(v) \to 0$ as $v \uparrow \bar{v}$.

Hence again $p_k - \frac{c}{1-\tau} \to 0$, so $\hat{p} = \frac{c}{1-\tau}$.

Thus every convergent subsequence of (p_k) converges to $\frac{c}{1-\tau}$, and therefore $p_k \to \frac{c}{1-\tau}$. □

Proof of Lemma 4. By Proposition 1 for $\alpha < \infty$ and Lemma 3 for $\alpha = \infty$, for every positive admissible commission $\tau \in (0, 1)$ satisfying $(1 - \tau)\bar{v} > c$, the pricing subgame

admits a symmetric equilibrium price $p \in \left[\frac{c}{1-\tau}, \bar{v}\right)$ satisfying $p - \frac{c}{1-\tau} = m(\alpha, p)$, where the markup at $\alpha = \infty$ is understood to be zero.

Hence, the set of prices induced by admissible commissions is nonempty. Its closure in $[c, \bar{v}]$ is compact. By (6), the platform objective

$$\Pi^P(p) = \left[p - c - \frac{c\, m(\alpha, p)}{p - m(\alpha, p)}\right]\left(1 - F(p)^n\right)$$

extends continuously to this closure and therefore attains a maximum there. By continuity of $m(\alpha, \cdot)$, any price added by taking the closure corresponds to either $\tau = 0$ or the no-trade boundary $(1-\tau)\bar{v} = c$, and yields zero platform profit. Since the positive admissible commission yields strictly positive profit, the maximum is attained at a price induced by an admissible commission. □

Proof of Proposition 2. The proof proceeds in two steps. First, any symmetric equilibrium is shown to satisfy $\alpha^* = \infty$. Second, the corresponding equilibrium price and commission are characterized.

Suppose, toward a contradiction, that there is a symmetric equilibrium with some finite $\alpha_1 < \infty$ and commission rate $\tau_1 \in [0, 1)$, inducing the symmetric price p_1. Necessarily, $\frac{c}{1-\tau_1} < \bar{v}$. Otherwise there is no trade and platform profit is zero, whereas a deviation to $\alpha = \infty$ and any $\widehat{\tau} \in (0, 1 - c/\bar{v})$ yields strictly positive profit by Lemma 3. Proposition 1 therefore applies, giving

$$p_1 - \frac{c}{1-\tau_1} = m(\alpha_1, p_1).$$

Since $p_1 < \bar{v}$, both the numerator and denominator in the definition of m are strictly positive, so $m(\alpha_1, p_1) > 0$ and hence $p_1 > \frac{c}{1-\tau_1}$.

Now choose $\tau_2 \in (0, 1)$ so that

$$\frac{c}{1-\tau_2} = p_1. \tag{A.11}$$

Because $p_1 > \frac{c}{1-\tau_1}$, it follows that $\tau_2 > \tau_1$. Consider the platform deviation to $(\alpha_2, \tau_2) = (\infty, \tau_2)$.

If $\alpha = \infty$, the unique symmetric equilibrium price is $p = \frac{c}{1-\tau}$ by Lemma 3. At a symmetric price vector, rankings depend only on match values, so the recommended product and total quantity sold are unchanged. Thus demand remains $1 - F(p_1)^n$, while platform revenue per sale rises from $\tau_1 p_1$ to $\tau_2 p_1$. Since $\tau_2 > \tau_1$, platform profit strictly increases, contradicting the equilibrium.

Therefore, every symmetric equilibrium must have $\alpha^* = \infty$. Given $\alpha = \infty$, the unique symmetric equilibrium price in the pricing subgame is $p = \frac{c}{1-\tau}$. The platform therefore chooses τ to maximize $\tau p(1 - F(p)^n)$ subject to $p = \frac{c}{1-\tau}$, or equivalently chooses p to

maximize $(p-c)(1-F(p)^n)$. By Definition 1 and the properties of p_m, this problem has the unique solution $p_m(c)$. The corresponding commission is uniquely determined by $p_m(c) = \frac{c}{1-\tau^*}$, that is,

$$\tau^* = 1 - \frac{c}{p_m(c)}.$$

Hence, in the unique symmetric equilibrium outcome,

$$\alpha^* = \infty, \qquad p^* = p_m(c), \qquad \tau^* = 1 - \frac{c}{p_m(c)}.$$

Existence follows from this construction, and uniqueness follows from the uniqueness of the platform optimum and of the symmetric pricing equilibrium for each τ. □

Proof of Lemma 5. Suppose, toward a contradiction, that $m_n(\alpha, p_n) \not\to 0$. Then there exist $\eta > 0$ and a subsequence, still indexed by n, such that $m_n(\alpha, p_n) \geq \eta$ for all n. Passing to a further subsequence if necessary, assume that

$$1 - F(p_n)^n \to q \in [0,1].$$

By the pricing condition and (4),

$$m_n(\alpha, p_n) = \frac{1 - F(p_n)^n}{n\left[F(p_n)^{n-1} f(p_n) + \alpha(n-1)\int_{p_n}^{\bar{v}} F(v)^{n-2} f(v)^2\, dv\right]}.$$

Since f is log-concave, the proof of Lemma 3 shows that $f/(1-F)$ is increasing and $\frac{1-F(p)}{f(p)} \to 0$ as $p \uparrow \bar{v}$. Two cases are considered.

Case 1: $q < 1$. Then $F(p_n)^n \to 1 - q > 0$. In particular, $F(p_n) \to 1$, and therefore $p_n \to \bar{v}$. Moreover,

$$m_n(\alpha, p_n) \leq \frac{1 - F(p_n)^n}{nF(p_n)^{n-1} f(p_n)} = \frac{1 + F(p_n) + \cdots + F(p_n)^{n-1}}{nF(p_n)^{n-1}} \cdot \frac{1 - F(p_n)}{f(p_n)}.$$

Since $F(p_n)^n \to 1 - q > 0$ and $F(p_n) \to 1$, also $F(p_n)^{n-1} = \frac{F(p_n)^n}{F(p_n)} \to 1 - q > 0$, so the first factor is bounded. The second factor converges to 0 because $p_n \to \bar{v}$ and $(1 - F(p_n))/f(p_n) \to 0$. Hence $m_n(\alpha, p_n) \to 0$, a contradiction.

Case 2: $q = 1$. Then $F(p_n)^n \to 0$, so the numerator converges to 1. It therefore suffices to show that the denominator diverges.

For each n, let $y_n \in [\underline{v}, \bar{v})$ satisfy $F(y_n) = 1 - \frac{1}{n}$.

Then $y_n \uparrow \bar{v}$. Since $(1 - 1/n)^n \to e^{-1} > 0$ while $F(p_n)^n \to 0$, it follows that for all large n, $F(p_n) < 1 - \frac{1}{n}$, and hence $p_n \leq y_n$.

Because $f/(1-F)$ is increasing, for every $v \geq y_n$,

$$f(v) \geq \frac{f(y_n)}{1-F(y_n)}(1-F(v)).$$

Therefore

$$\int_{p_n}^{\bar{v}} F(v)^{n-2} f(v)^2 \, dv \geq \frac{f(y_n)}{1-F(y_n)} \int_{y_n}^{\bar{v}} F(v)^{n-2}(1-F(v)) f(v) \, dv.$$

Multiplying by $n(n-1)$ and using the change of variable $u = F(v)$ gives

$$n(n-1) \int_{p_n}^{\bar{v}} F(v)^{n-2} f(v)^2 \, dv \geq \frac{f(y_n)}{1-F(y_n)} \, c_n,$$

where $c_n := n(n-1) \int_{1-1/n}^{1} u^{n-2}(1-u) \, du$. Notice that

$$c_n = 1 - n\Big(1-\frac{1}{n}\Big)^{n-1} + (n-1)\Big(1-\frac{1}{n}\Big)^{n} \longrightarrow 1 - \frac{2}{e} > 0.$$

Hence c_n is eventually bounded below by some constant $\kappa > 0$, and thus

$$n\alpha(n-1) \int_{p_n}^{\bar{v}} F(v)^{n-2} f(v)^2 \, dv \geq \alpha\kappa \, \frac{f(y_n)}{1-F(y_n)}.$$

Since $y_n \uparrow \bar{v}$ and $(1-F(y_n))/f(y_n) \to 0$, the right-hand side diverges to $+\infty$. Therefore the denominator in the expression for $m_n(\alpha, p_n)$ tends to $+\infty$, while the numerator tends to 1. Hence $m_n(\alpha, p_n) \to 0$, again a contradiction.

Both cases lead to a contradiction, establishing $m_n(\alpha, p_n) \to 0$ as $n \to \infty$.

□

Proof of Proposition 3. For fixed $\alpha \geq 1$, the markup can be written as

$$m_n(p) = \frac{1-F(p)^n}{n \, d_n(p)},$$

where

$$d_n(p) := F(p)^{n-1} f(p) + \alpha(n-1) \int_{p}^{\bar{v}} F(v)^{n-2} f(v)^2 \, dv.$$

First suppose that $0 < F(p) < 1$. After integration by parts,

$$d_n(p) = \alpha b_n(p) + (1-\alpha) f(p) F(p)^{n-1}, \qquad b_n(p) := f(\bar{v}) - \int_{p}^{\bar{v}} f'(v) F(v)^{n-1} \, dv.$$

It will be shown that d_n is increasing and that $(1-F(p)^n)/n$ is decreasing in n.

Fix $x := F(p) \in (0,1)$ and define $a_n := \frac{1-x^n}{n}$. Consider the extension to the real

numbers $a(t) := (1 - x^t)/t$ for $t > 0$. Its derivative is

$$a'(t) = \frac{x^t - 1 - tx^t \log x}{t^2}.$$

Let $y := -\log x > 0$, so $x^t = e^{-ty}$. Then

$$e^{-ty} - 1 + ty\, e^{-ty} = e^{-ty}(1 + ty) - 1 < 0,$$

because $e^{ty} > 1 + ty$ for all $ty > 0$. Hence $a'(t) < 0$ for all $t > 0$, and therefore $a_{n+1} < a_n$ for all integers $n \geq 1$.

Further,

$$\begin{aligned} b_{n+1}(p) - b_n(p) &= -\int_p^{\bar{v}} f'(v)\big(F(v)^n - F(v)^{n-1}\big)\, dv \\ &= \int_p^{\bar{v}} f'(v)F(v)^{n-1}(1 - F(v))\, dv \;\geq\; 0, \end{aligned}$$

where the inequality uses $f'(v) \geq 0$ a.e. and $F(v)^{n-1}(1 - F(v)) \geq 0$. Moreover,

$$(1 - \alpha)\big(F(p)^n - F(p)^{n-1}\big) = (\alpha - 1)F(p)^{n-1}(1 - F(p)) \geq 0.$$

Hence

$$d_{n+1}(p) - d_n(p) = \alpha\big(b_{n+1}(p) - b_n(p)\big) + (1 - \alpha)f(p)\big(F(p)^n - F(p)^{n-1}\big) \geq 0.$$

Moreover, $d_n(p) > 0$.

Since $a_{n+1} < a_n$ and $d_{n+1}(p) \geq d_n(p) > 0$,

$$\frac{m_{n+1}(p)}{m_n(p)} = \frac{a_{n+1}}{a_n} \cdot \frac{d_n(p)}{d_{n+1}(p)} < 1 \cdot 1 = 1,$$

so $m_{n+1}(p) < m_n(p)$ for all $n \geq 2$.

If $F(p) = 0$, then $a_n = 1/n$ is strictly decreasing in n, while the same integration-by-parts argument, applied from $\underline{v}$ to $\bar{v}$, gives $d_{n+1}(p) \geq d_n(p) > 0$; hence $m_{n+1}(p) < m_n(p)$. □

Proof of Proposition 4. By Lemma 1, there is no search in a symmetric equilibrium. Hence, if the common price is p_n, the platform recommends the product with the highest match value, and the consumer buys iff that value exceeds p_n. Therefore

$$CS_n = \int_{p_n}^{\bar{v}} (1 - F(v)^n)\, dv.$$

For part (ii), fix $\alpha > 0$. For each $n \geq 2$, let τ_n be an optimal platform fee, let p_n be an induced symmetric equilibrium price, and define $x_n := \frac{c}{1-\tau_n}$. Then

$$p_n - x_n = m_n(\alpha, p_n),$$

where $m_n(\alpha, \cdot)$ denotes the markup function with n firms. The proof of Lemma 5 uses only that the sequence of prices satisfies the symmetric pricing condition, so the same argument applies here. Hence $p_n - x_n \to 0$.

Fix any $\delta \in (0, \bar{v} - c)$ and define $\tau^\delta := 1 - \frac{c}{\bar{v}-\delta}$. Let $\tilde{p}_n^\delta$ denote a symmetric equilibrium price induced by (α, τ^δ). Then

$$\tilde{p}_n^\delta - \frac{c}{1-\tau^\delta} = m_n(\alpha, \tilde{p}_n^\delta),$$

and therefore, again by Lemma 5,

$$\tilde{p}_n^\delta \to \frac{c}{1-\tau^\delta} = \bar{v} - \delta.$$

Since $\tilde{p}_n^\delta \to \bar{v} - \delta < \bar{v}$, for all large n one has $\tilde{p}_n^\delta \leq \bar{v} - \delta/2$, and thus

$$F(\tilde{p}_n^\delta)^n \leq F(\bar{v} - \delta/2)^n \to 0.$$

It follows that the platform profit under the feasible fee τ^δ satisfies

$$\tilde{\Pi}_n^\delta = \tau^\delta \tilde{p}_n^\delta \big(1 - F(\tilde{p}_n^\delta)^n\big) \to \bar{v} - \delta - c.$$

Since τ_n is optimal, $\Pi_n \geq \tilde{\Pi}_n^\delta$. On the other hand,

$$\Pi_n = \tau_n p_n \big(1 - F(p_n)^n\big) \leq \tau_n p_n \leq p_n - c,$$

because $(1 - \tau_n)p_n \geq c$. Hence $p_n - c \geq \Pi_n \geq \tilde{\Pi}_n^\delta$, so taking limits yields

$$\liminf_{n\to\infty} p_n \geq \bar{v} - \delta.$$

Since $\delta > 0$ is arbitrary and $p_n \leq \bar{v}$, it follows that $p_n \to \bar{v}$. Using $0 \leq 1 - F(v)^n \leq 1$, we obtain

$$0 \leq CS_n \leq \int_{p_n}^{\bar{v}} 1 \, dv = \bar{v} - p_n,$$

and therefore $CS_n \to 0$.

For part (i), let $p_0 := \frac{c}{1-\tau}$. By Lemma 5, $p_n - p_0 \to 0$, and thus $p_n \to \frac{c}{1-\tau}$.

Since $F(v) = 0$ for $v < \underline{v}$, we can rewrite

$$CS_n = \big(\underline{v} - p_n\big)_+ + \int_{\underline{v}}^{\bar{v}} \mathbf{1}\{v \geq p_n\}\big(1 - F(v)^n\big)\, dv.$$

Define

$$f_n(v) := \mathbf{1}\{v \geq p_n\}\big(1 - F(v)^n\big), \qquad f(v) := \mathbf{1}\{v \geq p_0\}, \qquad g(v) := 1.$$

Then $0 \leq f_n(v) \leq g(v)$ for all n and all v, and g is integrable on $[\underline{v}, \bar{v}]$. Moreover, since $p_n \to p_0$, $\mathbf{1}\{v \geq p_n\} \to \mathbf{1}\{v \geq p_0\}$ for every $v \neq p_0$. For every fixed $v < \bar{v}$, $F(v)^n \to 0$, so $1 - F(v)^n \to 1$. Hence $f_n(v) \to f(v)$ for almost every $v \in [\underline{v}, \bar{v}]$. By the dominated convergence theorem,

$$\int_{\underline{v}}^{\bar{v}} f_n(v)\, dv \longrightarrow \int_{\underline{v}}^{\bar{v}} f(v)\, dv = \bar{v} - \max\left\{\underline{v}, \frac{c}{1-\tau}\right\}.$$

Also, $\big(\underline{v} - p_n\big)_+ \to \big(\underline{v} - p_0\big)_+$. Therefore

$$CS_n \longrightarrow \big(\underline{v} - p_0\big)_+ + \bar{v} - \max\{\underline{v}, p_0\} = \bar{v} - p_0 = \bar{v} - \frac{c}{1-\tau}.$$

□

Proof of Lemma 6. Suppose $\alpha < \infty$ and write $m(p) := m(\alpha, p)$. Under nominal-fee pricing firm i solves $\max_{p_i} (p_i - c - t)\, D_i(p; \alpha)$, so the pricing game in the $s = \infty$ benchmark is exactly the same as in Lemma 2 with effective cost $c + t$ in place of $c/(1-\tau)$. Hence, whenever $c + t < \bar{v}$, the same arguments as in Lemma 2 and Proposition 1 imply that the original pricing game admits a positive-trade symmetric equilibrium, and every such equilibrium price satisfies $p - (c + t) = m(p)$, that is, $p = c + t + m(p)$. For finite α, fix any commission rate $\tau \in [0, 1)$ and define $t(\tau) := \frac{c\tau}{1-\tau}$. Then, for every price profile p and every firm i,

$$\big((1-\tau)p_i - c\big)D_i(p; \alpha) = (1-\tau)\big(p_i - c - t(\tau)\big)D_i(p; \alpha).$$

Hence the pricing subgame under commission pricing with fee τ is strategically equivalent to the pricing subgame under nominal-fee pricing with fee $t(\tau)$: firms' payoffs differ only by the positive multiplicative constant $(1-\tau)$. Therefore, the two subgames have the same best-response correspondences and, in particular, the same set of symmetric equilibrium prices. Since the map $\tau \mapsto t(\tau) = \frac{c\tau}{1-\tau}$ is a bijection from $[0, 1)$ to $[0, \infty)$, the two regimes generate the same family of pricing subgames up to this reparametrization.

Let $S(\alpha)$ denote the common set of supportable positive-trade symmetric equilibrium prices under the two regimes. By the argument above, this set is common to both

regimes and is nonempty. Let $\overline{S}(\alpha)$ denote its closure in $[c, \bar{v}]$. Then $\overline{S}(\alpha)$ is nonempty and compact.

Accordingly, any $p \in S(\alpha)$ satisfies, under nominal-fee pricing,

$$p = c + t + m(p),$$

whereas under commission pricing any symmetric equilibrium price satisfies

$$p = \frac{c}{1-\tau} + m(p).$$

Hence, writing platform profit as a function of the induced equilibrium price p,

$$\Pi^N(p) = \big(p - c - m(p)\big)\big(1 - F(p)^n\big)$$

under nominal fees, and

$$\Pi^C(p) = \left[p - c - \frac{c\, m(p)}{p - m(p)}\right]\big(1 - F(p)^n\big)$$

under commissions.

Thus, at the same induced price p,

$$\Pi^C(p) - \Pi^N(p) = \frac{m(p)\big(p - c - m(p)\big)}{p - m(p)}\big(1 - F(p)^n\big) \geq 0.$$

This proves weak dominance of commission pricing. The functions Π^N and Π^C extend continuously from $S(\alpha)$ to $\overline{S}(\alpha)$. By the same boundary argument as in Lemma 4, every point in $\overline{S}(\alpha) \setminus S(\alpha)$ corresponds to the no-trade boundary and yields zero profit under both fee regimes. On the other hand, choosing some $t > 0$ with $c + t < \bar{v}$, and the corresponding commission $\tau = t/(c+t) > 0$, yields strictly positive profit under both regimes. Therefore, both objectives attain their maxima over the compact set $\overline{S}(\alpha)$ at prices in $S(\alpha)$. Hence both objectives also attain their maxima over $S(\alpha)$. Let

$$p_N \in \arg\max_{p \in S(\alpha)} \Pi^N(p)$$

be any induced equilibrium price that attains the platform's maximal profit under nominal-fee pricing. Under nominal-fee pricing the platform can induce strictly positive profit by choosing some $t > 0$ with $c + t < \bar{v}$, so profits are positive. In particular, $1 - F(p_N)^n > 0$ and $p_N - c - m(p_N) > 0$. It follows that:

$$\Pi^C(p_N) - \Pi^N(p_N) > 0.$$

This proves part (i).

It remains to prove part (ii). Let

$$\mathcal{E}^N(\alpha) := \arg\max_{p\in S(\alpha)} \Pi^N(p), \qquad \mathcal{E}^C(\alpha) := \arg\max_{p\in S(\alpha)} \Pi^C(p).$$

By the preceding boundary argument, these sets are nonempty and coincide with the corresponding argmax sets over $\overline{S}(\alpha)$. Since the extended profit functions are continuous and $\overline{S}(\alpha)$ is compact, $\mathcal{E}^N(\alpha)$ and $\mathcal{E}^C(\alpha)$ are compact. In particular,

$$p_N^{\max} := \max \mathcal{E}^N(\alpha), \qquad p_C^{\max} := \max \mathcal{E}^C(\alpha)$$

are well defined.

Suppose that $p \mapsto \frac{m(p)}{p-m(p)}$ is weakly decreasing. Take any $p \in S(\alpha)$ with $p > p_N^{\max}$. Since

$$\mathcal{E}^N(\alpha) = \arg\max_{q\in S(\alpha)} \Pi^N(q)$$

and $p_N^{\max} = \max \mathcal{E}^N(\alpha)$, such a p does not belong to $\mathcal{E}^N(\alpha)$. Hence p does not maximize Π^N on $S(\alpha)$, whereas $p_N^{\max}$ does. Therefore

$$\Pi^N(p) < \Pi^N(p_N^{\max}).$$

Moreover,

$$1 + \frac{m(p)}{p-m(p)} \le 1 + \frac{m(p_N^{\max})}{p_N^{\max} - m(p_N^{\max})}.$$

Since

$$\Pi^C(p) = \left(1 + \frac{m(p)}{p-m(p)}\right) \Pi^N(p),$$

it follows that

$$\Pi^C(p) < \left(1 + \frac{m(p_N^{\max})}{p_N^{\max} - m(p_N^{\max})}\right) \Pi^N(p_N^{\max}) = \Pi^C(p_N^{\max}).$$

Hence no price in $S(\alpha)$ strictly above $p_N^{\max}$ can maximize Π^C on $S(\alpha)$. Therefore no platform-preferred equilibrium price strictly above $p_N^{\max}$ can arise under commission pricing, and

$$p_C^{\max} \le p_N^{\max}.$$

It remains to verify the two sufficient conditions. First, suppose $0 \le \alpha \le 1$, f is absolutely continuous on $[\underline{v}, \bar{v}]$, and $f'(v) \ge 0$ for a.e. $v \in (\underline{v}, \bar{v})$. Writing

$$m(p) = \frac{N(p)}{D(p)}, \qquad N(p) := \frac{1-F(p)^n}{n}, \qquad D(p) := F(p)^{n-1} f(p) + \alpha(n-1) \int_p^{\bar{v}} F(v)^{n-2} f(v)^2 \, dv,$$

one has, for a.e. $p \in (\underline{v}, \bar{v})$, $N'(p) = -F(p)^{n-1} f(p) < 0$ and

$$D'(p) = F(p)^{n-1} f'(p) + (1-\alpha)(n-1)F(p)^{n-2} f(p)^2 \geq 0.$$

Hence, for a.e. $p \in (\underline{v}, \bar{v})$,

$$m'(p) = \frac{N'(p)D(p) - N(p)D'(p)}{D(p)^2} \leq 0.$$

On the relevant domain, where $p - m(p) > 0$, the function $p \mapsto m(p)/(p - m(p))$ is locally absolutely continuous and, for a.e. p,

$$\left(\frac{m(p)}{p - m(p)}\right)' = \frac{p\, m'(p) - m(p)}{(p - m(p))^2} < 0,$$

because $m(p) > 0$ for every finite α and $p < \bar{v}$. The same conclusion is immediate below $\underline{v}$, where $m(p)$ is constant for $\alpha > 0$, while no such price is supportable for $\alpha = 0$. It follows that $p \mapsto m(p)/(p - m(p))$ is strictly decreasing.

Second, consider the power distribution of Remark 1 with $\underline{v} = 0$. In the proof of Remark 1, after the change of variables $x = p/\bar{v}$, the function

$$H(x) = \frac{1 - x^{\gamma+1}}{nx(A + Bx^{\gamma})}$$

is shown to be strictly decreasing for every $\alpha \geq 0$. Under $\underline{v} = 0$, this is exactly $m(p)/p$. Since

$$\frac{m(p)}{p - m(p)} = \frac{m(p)/p}{1 - m(p)/p},$$

and $x \mapsto x/(1-x)$ is strictly increasing on $[0, 1)$, it follows that $p \mapsto \frac{m(p)}{p - m(p)}$ is also strictly decreasing. This proves part (ii). □

Proof of Proposition 5. The argument follows the same double-deviation logic as Proposition 2. Consider a policy with finite α and nominal fee t that induces price $p = c + t + m(\alpha, p)$. The platform can instead choose $\alpha' = \infty$ and $t' = t + m(\alpha, p) = p - c$. This leaves the equilibrium price, and hence transaction volume, unchanged, while strictly increasing the platform's fee per transaction. Thus, no optimal policy can have finite α.

At $\alpha = \infty$, the equilibrium price is $p = c + t$, so the platform chooses t to maximize $t[1 - F(c + t)^n]$, equivalently choosing p to maximize $(p - c)[1 - F(p)^n]$. The unique solution is $p^* = p_m(c)$ and $t^* = p_m(c) - c$. □

A.1.2 Deferred Proofs for Section 4

Proof of Lemma 7. Suppose $\bar{\tau} < \tau^*$. Fix a feasible $\tau \leq \bar{\tau}$. A price p is supportable in a symmetric equilibrium for some $\alpha \geq 0$ if and only if $p - \frac{c}{1-\tau} = m(\alpha, p)$. Since, for fixed $p < \bar{v}$, the map $\alpha \mapsto m(\alpha, p)$ is continuous and strictly decreasing, with $\lim_{\alpha\to\infty} m(\alpha, p) = 0$ and $m(0, p) = \frac{1-F(p)^n}{nF(p)^{n-1}f(p)}$, the supportable prices are exactly those satisfying

$$0 \leq p - \frac{c}{1-\tau} \leq \frac{1-F(p)^n}{nF(p)^{n-1}f(p)}.$$

The right-hand inequality is equivalent to the derivative of $x \mapsto \left(x - \frac{c}{1-\tau}\right)(1-F(x)^n)$ being nonnegative at $x = p$. By the properties of p_m, this objective is single-peaked with unique maximizer $p_m\left(\frac{c}{1-\tau}\right)$. Hence, the supportable prices are exactly

$$\left[\frac{c}{1-\tau},\, p_m\left(\frac{c}{1-\tau}\right)\right].$$

The cap must bind. Suppose instead that an optimal policy satisfies $\tau < \bar{\tau}$ and induces price p.

If $p > \frac{c}{1-\tau}$, choose $\tau' \in (\tau, \bar{\tau}]$ sufficiently close to τ. Then $\frac{c}{1-\tau'} < p$, and since $p \leq p_m\left(\frac{c}{1-\tau}\right)$ and p_m is strictly increasing, also $p < p_m\left(\frac{c}{1-\tau'}\right)$. Thus p remains supportable under τ'. By the platform-preferred selection rule, the induced profit under some such policy is at least $\tau' p(1-F(p)^n) > \tau p(1-F(p)^n)$, a contradiction.

If instead $p = \frac{c}{1-\tau}$, then necessarily $\alpha = \infty$ by Lemma 3. Since $\bar{\tau} < \tau^* = 1 - \frac{c}{p_m(c)}$, one has $\frac{c}{1-\tau} < p_m(c)$. Hence $q \mapsto (q-c)(1-F(q)^n)$ is strictly increasing at $q = \frac{c}{1-\tau}$, so increasing τ slightly while keeping $\alpha = \infty$ strictly raises platform profit, again a contradiction.

Therefore every optimum satisfies $\tau = \bar{\tau}$.

Fix now $\tau = \bar{\tau}$. The supportable prices are

$$\left[\frac{c}{1-\bar{\tau}},\, p_m\left(\frac{c}{1-\bar{\tau}}\right)\right].$$

Since the platform-preferred equilibrium is selected, the platform solves

$$\max_{p \in \left[\frac{c}{1-\bar{\tau}},\, p_m\left(\frac{c}{1-\bar{\tau}}\right)\right]} \bar{\tau}\, p(1-F(p)^n).$$

Equivalently, it maximizes $p(1-F(p)^n)$ over that interval. By the standard monopoly price properties stated after Definition 1, this objective is single-peaked with unique maximizer $p_m(0)$. Hence the optimal price is $p_m(0)$ if $p_m(0) \geq \frac{c}{1-\bar{\tau}}$, and otherwise the lower boundary $\frac{c}{1-\bar{\tau}}$.

Thus

$$p^* = \max\left\{p_m(0), \frac{c}{1-\bar{\tau}}\right\}.$$

□

Proof of Lemma 8. For any commission cap $\bar{\tau}$, let p^τ denote the platform-preferred equilibrium price under that regime. Then $p^\tau \geq p_m(0)$. If $\bar{\tau} \geq \tau^*$, this follows from $p^\tau = p_m(c) > p_m(0)$ by Proposition 2 and the standard monopoly price properties stated after Definition 1. If $\bar{\tau} < \tau^*$, then by Lemma 7,

$$p^\tau = \max\left\{p_m(0), \frac{c}{1-\bar{\tau}}\right\} \geq p_m(0).$$

Under a nominal-fee cap $\bar{t}$, platform profit is

$$\Pi^P = t[1 - F(p)^n],$$

so for a given t the platform weakly prefers the lowest induced equilibrium price. Under nominal-fee pricing, every symmetric equilibrium price satisfies

$$p = c + t + m(\alpha, p) \geq c + t.$$

Since transaction volume is weakly decreasing in p, platform profit cannot exceed

$$t[1 - F(c+t)^n],$$

and this bound is attained at $\alpha = \infty$, which induces the unique pricing equilibrium $p = c + t$. If $c + t < \underline{v}$, other values of α may also be optimal, because any price at or below $\underline{v}$ generates unit transaction volume.

The platform therefore solves

$$\max_{t \leq \bar{t}} t[1 - F(c+t)^n],$$

or equivalently

$$\max_{p - c \leq \bar{t}} (p - c)[1 - F(p)^n].$$

By Definition 1, the unconstrained maximizer is $p_m(c)$, so the unconstrained optimal nominal fee is $p_m(c) - c$. Since $p_m(c) > p_m(0)$ by the standard monopoly price properties stated after Definition 1, the assumption $\bar{t} < p_m(0) - c$ implies that the cap binds. Thus every optimal policy sets $t = \bar{t}$. If $c + \bar{t} \geq \underline{v}$, then $\alpha = \infty$ is uniquely optimal and $p^t = c + \bar{t}$. If $c + \bar{t} < \underline{v}$, any optimal policy must induce $p^t \leq \underline{v}$. Because $n \geq 2$, $p_m(0) > \underline{v}$,

and thus, in either case

$$p^t < p_m(0) \leq p^\tau.$$

Hence every equilibrium price induced by an optimal policy under the nominal-fee cap is strictly lower than the platform-preferred equilibrium price under any commission cap $\bar{\tau}$.

Finally, by Lemma 1, there is no search under either regime, and at a common price the recommended product has the highest match value. Because $p^t < p^\tau$, $p^t > c$, and $p^\tau \geq p_m(0) > \underline{v}$, the lower nominal-fee price serves a strictly larger set of consumers whose match values exceed production cost. Consumer surplus and total welfare are therefore strictly higher under the nominal-fee cap than under the platform-preferred equilibrium under any commission cap $\bar{\tau}$.

Moreover, when $c + \bar{t} > \underline{v}$, the preceding argument gives $p^t = c + \bar{t}$. Hence lowering $\bar{t}$, while remaining in this region, strictly lowers the equilibrium price. Since consumer surplus and total welfare are strictly decreasing in p for $p > \max\{c, \underline{v}\}$, it strictly raises both. □

Proof of Proposition 7. First, fix $\tau \in [0, 1)$ such that $c/(1-\tau) < \bar{v}$, and characterize the symmetric equilibrium of the pricing subgame for each $\alpha \in \{0, \infty\}$.

If $\alpha = 0$, then by Proposition 1 any symmetric equilibrium price satisfies (4) with $\alpha = 0$, i.e.

$$p - \frac{c}{1-\tau} = \frac{1 - F(p)^n}{nF(p)^{n-1}f(p)}.$$

This is exactly the first-order condition for maximizing $(p - \frac{c}{1-\tau})(1 - F(p)^n)$. By Definition 1 and the properties of p_m, the unique solution is $p = p_m\big(\frac{c}{1-\tau}\big)$.

If $\alpha = \infty$, the unique symmetric equilibrium price is $p = \frac{c}{1-\tau}$ by Lemma 3.

Thus, for each $\alpha \in \{0, \infty\}$, the pricing subgame admits a unique symmetric equilibrium.

Turn now to the platform stage. If $c/(1-\tau) \geq \bar{v}$, every symmetric pricing equilibrium has zero trade and hence gives the platform zero profit. Such a policy cannot be optimal, since the platform obtains strictly positive profit by choosing $\alpha = \infty$ and a sufficiently small positive commission. Take any policy $(\alpha, \tau) = (0, \tau)$ with $c/(1-\tau) < \bar{v}$. The induced symmetric equilibrium price is then $p_0 = p_m(c/(1-\tau))$. By the standard monopoly price properties stated after Definition 1, $p_m(x) > x$, so $p_0 > \frac{c}{1-\tau}$. Define $\tilde{\tau} := 1 - \frac{c}{p_0} > \tau$. If the platform instead chooses $(\alpha, \tau) = (\infty, \tilde{\tau})$, the induced symmetric equilibrium price is again p_0, since under $\alpha = \infty$ it follows that $p = \frac{c}{1-\tilde{\tau}} = p_0$. Quantity sold is therefore unchanged, while the commission rate is strictly higher. Hence, platform profit is strictly higher. Therefore, no policy with $\alpha = 0$ can be optimal.

It follows that any optimal platform policy must have $\alpha = \infty$. Given $\alpha = \infty$, the induced price is $p = \frac{c}{1-\tau}$, so the platform chooses p to maximize $(p - c)(1 - F(p)^n)$. By Definition 1 and the properties of p_m, the unique maximizer is $p_m(c)$, which implies

the unique optimal fee $\tau^* = 1 - \frac{c}{p_m(c)}$. Hence the unique optimal platform choice is $(\alpha, \tau) = (\infty, \tau^*)$. □

Proof of Proposition 8. Part (i) follows from Proposition 7. If $\bar{\tau} > \tau^*$, the cap is slack, so the platform implements the unregulated outcome $(\alpha, \tau) = (\infty, \tau^*)$.

Part (ii) follows from Proposition 7 and Lemma 7. At $\bar{\tau} = \tau^*$, the claim follows from Proposition 7. If $c < p_m(0)$ and $\bar{\tau} \in [\tau^{**}, \tau^*)$, Lemma 7 implies that the cap binds and that the optimal price is $p = \frac{c}{1-\bar{\tau}} \geq p_m(0)$, with equality at $\bar{\tau} = \tau^{**}$. This lower-bound price requires zero markup and hence $\alpha = \infty$.

Now suppose

$$p_m(c)\big(1 - F(p_m(c))^n\big) > c\big(1 - F(c)^n\big).$$

By the standard monopoly price properties stated after Definition 1 the function $p \mapsto p(1 - F(p)^n)$ has the unique maximizer $p_m(0)$, and $p_m(c) > c$. Hence this condition implies $c < p_m(0)$: if instead $c \geq p_m(0)$, then single-peakedness together with $p_m(c) > c$ would imply

$$p_m(c)\big(1 - F(p_m(c))^n\big) \leq c\big(1 - F(c)^n\big),$$

a contradiction. In particular, τ^{**} exists.

Now consider $\bar{\tau} \leq \tau^{**}$. Under $\alpha = \infty$ and fee τ, the induced price is $p = c/(1-\tau)$ by Lemma 3, so platform profit is

$$\Pi_\infty(\tau) = \tau \frac{c}{1-\tau}\left(1 - F\left(\frac{c}{1-\tau}\right)^n\right) = \left(\frac{c}{1-\tau} - c\right)\left(1 - F\left(\frac{c}{1-\tau}\right)^n\right).$$

By the standard monopoly price properties stated after Definition 1, the function $p \mapsto (p-c)(1 - F(p)^n)$ has the unique maximizer $p_m(c)$, and $p_m(c) > p_m(0)$. Since $\tau \leq \tau^{**}$ implies $c/(1-\tau) \leq p_m(0) < p_m(c)$, it follows that $\Pi_\infty(\tau)$ is strictly increasing on $[0, \tau^{**}]$. Therefore, conditional on $\alpha = \infty$, the platform optimally sets $\tau = \bar{\tau}$.

Under $\alpha = 0$ and fee τ, firms set $p = p_m\big(\frac{c}{1-\tau}\big)$, so platform profit is

$$\Pi_0(\tau) = \tau\, p_m\left(\frac{c}{1-\tau}\right)\left(1 - F\left(p_m\left(\frac{c}{1-\tau}\right)\right)^n\right).$$

Hence, a sufficient condition for the platform to prefer $\alpha = 0$ is that $\Pi_0(\bar{\tau}) \geq \Pi_\infty(\bar{\tau})$, i.e.

$$p_m\left(\frac{c}{1-\bar{\tau}}\right)\left(1 - F\left(p_m\left(\frac{c}{1-\bar{\tau}}\right)\right)^n\right) \geq \frac{c}{1-\bar{\tau}}\left(1 - F\left(\frac{c}{1-\bar{\tau}}\right)^n\right).$$

At $\bar{\tau} = 0$, the difference between the two sides is strictly positive by assumption. Since p_m is continuous, the difference is continuous in $\bar{\tau}$ and therefore remains positive for all sufficiently small positive caps. Hence, shrinking this neighborhood if necessary, there exists $\tau^{***} \in (0, \tau^{**})$ such that for every $\bar{\tau} \in (0, \tau^{***})$ the platform prefers $\alpha = 0$.

Finally, whenever the platform chooses $\alpha = 0$, its chosen fee satisfies $\tau > 0$ (otherwise the platform profit is zero). The induced price is therefore $p = p_m\left(\frac{c}{1-\tau}\right) > p_m(c)$ by the properties of p_m. By Proposition 7, the unregulated outcome is $(\alpha, \tau) = (\infty, \tau^*)$, and by Proposition 2 it induces the price $p_m(c)$. In the symmetric equilibrium, the recommended product is still the highest-match product, so relative to the unregulated outcome the only change is the higher purchase threshold. Therefore, both consumer surplus and total welfare are strictly lower than in the unregulated outcome. □

Online Appendix

This online appendix contains auxiliary results and omitted proofs. Section OA.1 verifies the standard properties of the monopoly price used throughout the paper. Section OA.2 contains the proof and an additional result for the endogenous-quality extension in Section 5.1. Section OA.3 provides the omitted proofs, auxiliary derivations, and an additional large-market result for the inferior-product extension in Section 5.2. The remaining sections study additional aspects: endogenous firm entry, positive platform marginal costs, and a benchmark without recommendation design.

OA.1 Properties of the Monopoly Price

Lemma 11 (Properties of the Monopoly Price)
Suppose f is log-concave on $(\underline{v}, \bar{v})$. For every $x < \bar{v}$, define

$$p_m(x) \in \arg\max_{p \in [\underline{v}, \bar{v}]} (p - x)\big(1 - F(p)^n\big).$$

Then:

(i) $p_m(x)$ is uniquely defined and satisfies $p_m(x) \in (x, \bar{v})$.

(ii) The function $p \mapsto (p-x)\big(1-F(p)^n\big)$ is strictly increasing on $(x, p_m(x))$ and strictly decreasing on $(p_m(x), \bar{v})$.

(iii) $p_m(x)$ is strictly increasing in x.

Proof. Fix $x < \bar{v}$. Since $(p - x)\big(1 - F(p)^n\big) \leq 0$ for $p \leq x$, equals 0 at $p = \bar{v}$, and is strictly positive for $p > x$ sufficiently close to x, any maximizer lies in $(x, \bar{v})$.

Next, f is log-concave by assumption, and it is standard that F is log-concave on $(\underline{v}, \bar{v})$. Hence

$$v \mapsto nF(v)^{n-1} f(v)$$

is log-concave on $(\underline{v}, \bar{v})$, as the product of log-concave functions. Therefore

$$(v, p) \mapsto \mathbf{1}\{v \geq p\}\, nF(v)^{n-1} f(v)$$

is log-concave on the convex set $\{(v, p) : v \geq p\}$. By Prékopa's theorem,

$$p \mapsto \int_p^{\bar{v}} nF(v)^{n-1} f(v)\, dv = 1 - F(p)^n$$

is log-concave on $(\underline{v}, \bar{v})$.

It follows that on $(x, \bar{v})$,

$$\log\Big((p-x)\big(1-F(p)^n\big)\Big) = \log(p-x) + \log\big(1-F(p)^n\big)$$

is strictly concave, since $\log(p-x)$ is strictly concave and $\log(1-F(p)^n)$ is concave. Hence

$$p \mapsto (p-x)\big(1-F(p)^n\big)$$

is strictly log-concave on $(x, \bar{v})$ and therefore has a unique maximizer. This proves part (i).

Because the objective is strictly log-concave on $(x, \bar{v})$ and has a unique maximizer, it is single-peaked. Hence it is strictly increasing on $(x, p_m(x))$ and strictly decreasing on $(p_m(x), \bar{v})$. This proves part (ii).

For part (iii), let $x_2 > x_1$. Then

$$(p-x_2)\big(1-F(p)^n\big) - (p-x_1)\big(1-F(p)^n\big) = -(x_2-x_1)\big(1-F(p)^n\big),$$

which is strictly increasing in p, because $1-F(p)^n$ is strictly decreasing. Thus the objective has strictly increasing differences in (p, x), so the unique maximizer is weakly increasing in x.

To show strictness, suppose instead that $p_m(x_2) = p_m(x_1) =: p$. By part (i), this common maximizer lies in $(x_2, \bar{v})$ and therefore satisfies the first-order condition

$$1 - F(p)^n - (p-x_j)nF(p)^{n-1}f(p) = 0, \qquad j = 1, 2.$$

Subtracting the two equations gives

$$(x_2 - x_1)nF(p)^{n-1}f(p) = 0,$$

a contradiction, since $x_2 > x_1$ and $f(p) > 0$ on the interior of the support. Therefore $p_m(x_2) > p_m(x_1)$. □

OA.2 Endogenous Quality: Proofs and Additional Results

OA.2.1 Deferred Proofs

Proof of Proposition 9. First, for fixed net price x, platform profit is shown to be strictly decreasing in $\psi := x^{n-1} + \alpha(1-x^{n-1})$ on $[1, \infty)$. This fact is then used to show that any policy with $\alpha \geq 1$ admits a profitable deviation, and hence cannot be optimal.

Let (α, τ) be a profit-maximizing platform policy and let (p, q) be an associated symmetric equilibrium. Set $x := p - q$. Using the notation

$$I = \int_x^1 F(v)^{n-1} f(v)\, dv, \qquad A = F(x)^{n-1} f(x), \qquad B = \int_x^1 (n-1) F(v)^{n-2} f(v)^2\, dv,$$

and $M = (1-\tau)p - c(q)$, the symmetric first-order conditions are

$$(1-\tau)I = M(A + \alpha B), \qquad c'(q)I = M(A + B).$$

Under $F(v) = v$ on $[0,1]$, so $f \equiv 1$, $I = \Delta_n(x) := \int_x^1 v^{n-1}\, dv = (1-x^n)/n$, $A = x^{n-1}$, and $B = 1 - x^{n-1}$. Hence $A + B = 1$ and

$$A + \alpha B = x^{n-1} + \alpha(1 - x^{n-1}) =: \psi.$$

The corner cases $x = 1$ and $q = 0$ are ruled out first. If $x = 1$, then $\Delta_n(1) = 0$, so there is no trade and platform profit is zero. This cannot be optimal, since the platform can attain strictly positive profit under a feasible policy with positive trade.

Next suppose $x < 1$ but $q = 0$. Since $q \geq 0$ is a constraint, the firm's quality choice must satisfy the KKT condition $\partial \Pi_i / \partial q_i \big|_{q=0} \leq 0$, where the derivative is computed from the unconstrained profit function. In the uniform case,

$$\left.\frac{\partial \Pi_i}{\partial q_i}\right|_{q=0} = -c'(0)\Delta_n(x) + M(A + B) = M,$$

because $c'(0) = 0$ and $A + B = 1$. Thus $M \leq 0$. But $c(0) = 0$ and $\tau < 1$ imply $M = (1-\tau)p \geq 0$, hence $M = 0$ and therefore $p = 0$. Platform profit is then $\Pi^P = \tau p(1 - x^n) = 0$, again a contradiction. Therefore any profit-maximizing policy induces an equilibrium with $x < 1$ and $q > 0$.

Dividing the two first-order conditions gives

$$1 - \tau = c'(q)\psi. \tag{OA.12}$$

With quadratic costs $c(q) = \frac{k}{2}q^2$, $c'(q) = kq$ and $c(q)/c'(q) = q/2$. Using $p = x + q$ and $M = (1-\tau)(x+q) - c(q)$ together with (OA.12), one obtains

$$\psi(x + q) - \frac{q}{2} = \Delta_n(x). \tag{OA.13}$$

For $\psi > \frac{1}{2}$, rearranging (OA.13) gives

$$q\Big(\psi - \frac{1}{2}\Big) = \Delta_n(x) - \psi x,$$

and hence

$$q(\psi, x) = \frac{\Delta_n(x) - \psi x}{\psi - \frac{1}{2}}. \tag{OA.14}$$

Moreover, (OA.12) implies $1 - \tau(\psi, x) = k\, q(\psi, x)\psi$, so platform profit per consumer can be written as

$$\Pi^P(\psi, x) = \tau(\psi, x)\big(x + q(\psi, x)\big)(1 - x^n), \qquad \tau(\psi, x) = 1 - k\, q(\psi, x)\psi. \tag{OA.15}$$

A direct differentiation of (OA.15) using (OA.14) yields, for $\psi > \frac{1}{2}$,

$$\frac{\partial \Pi^P}{\partial \psi}(\psi, x) = -\frac{2\,(x^n - 1)\,\big(nx + 2x^n - 2\big)}{n^2(2\psi - 1)^3}\, G(\psi, x), \tag{OA.16}$$

where $G(\psi, x) := 2kn\psi x + 2k\psi x^n - 2k\psi + kx^n - k + 2n\psi - n$.

Now consider $\psi \geq 1$, which is equivalent to $\alpha \geq 1$. Since $x < 1$, $x^n - 1 < 0$ and $2\psi - 1 > 0$. If $x = 0$, then $nx + 2x^n - 2 = -2 < 0$. If instead $x \in (0, 1)$, then $q(\psi, x) > 0$ and (OA.14) imply $\Delta_n(x) > \psi x \geq x$, or equivalently

$$H_1(x) := nx + x^n - 1 < 0.$$

Since $nx + 2x^n - 2 = H_1(x) + (x^n - 1)$ and both terms on the right-hand side are strictly negative, one obtains $nx + 2x^n - 2 < 0$. Therefore the prefactor in (OA.16) is strictly negative for every $x \in [0, 1)$ and every $\psi \geq 1$ satisfying $q(\psi, x) > 0$, so

$$\operatorname{sign}\left(\frac{\partial \Pi^P}{\partial \psi}(\psi, x)\right) = -\operatorname{sign} G(\psi, x).$$

It remains to show that $G(\psi, x) > 0$ on this region. Write

$$G(\psi, x) = G(1, x) + 2\big(n + kH_1(x)\big)(\psi - 1),$$

where $G(1, x) = n + k\, s_n(x)$ and $s_n(x) := 2nx + 3x^n - 3$. Since $s_n'(x) = 2n + 3nx^{n-1} > 0$ and $s_n(0) = -3$, one has $s_n(x) \geq -3$ for all $x \in [0, 1]$. Therefore $G(1, x) = n + k\, s_n(x) \geq n - 3k > 0$ whenever $k < \frac{n}{3}$. Also, $H_1(0) = -1$ and H_1 is strictly increasing, so $H_1(x) \in [-1, 0)$ on this region. Hence $n + kH_1(x) \geq n - k > 0$. Thus $G(\psi, x)$ is strictly increasing in ψ, and so $G(\psi, x) \geq G(1, x) > 0$ on this region. It follows from (OA.16) that

$$\frac{\partial \Pi^P}{\partial \psi}(\psi, x) < 0 \qquad \text{whenever } \psi \geq 1,\ x \in [0, 1),\ \text{and } q(\psi, x) > 0.$$

Suppose now, toward a contradiction, that some profit-maximizing policy has $\alpha \geq 1$.

Let (p, q) be an associated equilibrium and write $x := p - q$. Since $x < 1$, one has

$$\psi = x^{n-1} + \alpha(1 - x^{n-1}) \geq 1.$$

By the previous display,

$$\frac{\partial \Pi^P}{\partial \psi}(\psi, x) < 0.$$

Hence, for sufficiently small $\varepsilon > 0$,

$$\Pi^P(\psi - \varepsilon, x) > \Pi^P(\psi, x).$$

Choose $\varepsilon > 0$ small enough that simultaneously

$$\Pi^P(\psi - \varepsilon, x) > \Pi^P(\psi, x), \qquad q(\psi - \varepsilon, x) > 0, \qquad \tau(\psi - \varepsilon, x) \in (0, 1),$$

and

$$\psi - \varepsilon \geq x^{n-1}.$$

This is possible because the original equilibrium has $q(\psi, x) > 0$, the original policy yields strictly positive platform profit and hence $\tau \in (0, 1)$, the functions $q(\cdot, x)$ and $\tau(\cdot, x)$ are continuous on $(\frac{1}{2}, \infty)$, and

$$\psi = x^{n-1} + \alpha(1 - x^{n-1}) > x^{n-1}$$

when $x < 1$ and $\alpha \geq 1$.

Now define

$$\alpha' := \frac{\psi - \varepsilon - x^{n-1}}{1 - x^{n-1}} \qquad \text{if } x > 0,$$

and $\alpha' := \psi - \varepsilon$ if $x = 0$. Then $\alpha' \geq 0$ and $\alpha' < \alpha$, and by construction

$$x^{n-1} + \alpha'(1 - x^{n-1}) = \psi - \varepsilon.$$

Thus the pair

$$(\alpha', \tau'), \qquad \tau' := \tau(\psi - \varepsilon, x),$$

implements the same x and the corresponding quality level $q(\psi - \varepsilon, x)$, but yields strictly higher platform profit. This contradicts optimality of (α, τ).

Therefore no profit-maximizing platform policy can have $\alpha \geq 1$. Since $\alpha \geq 0$ by construction, it follows that

$$\alpha^* \in [0, 1).$$

□

OA.2.2 Additional Results

The next result complements the finite-n quality analysis in Section 5.1. Condition (9) shows that recommendation price-sensitivity and commissions jointly determine firms' incentives to invest in quality. Proposition 9 shows, in a tractable finite-n example, that this force can lead the platform to choose a recommendation rule with limited price sensitivity. The following result shows that in large markets the optimal design aligns the recommendation weight with the commission wedge: along any convergent sequence of optimal policies, α converges to $1-\tau$.

Proposition 12 (Asymptotic Alignment of α and τ with Endogenous Quality)
Assume that there exist $\varepsilon_0 > 0$ and a constant $0 < \underline{f}$ such that $f(v) \geq \underline{f}$ for all $v \in [\bar{v} - \varepsilon_0, \bar{v}]$. Let $c : \mathbb{R}_+ \to \mathbb{R}_+$ be continuously differentiable and strictly convex with $c(0) = 0$, $c'(0) = 0$, and $\lim_{q\to\infty} c'(q) = \infty$.

For each $n \geq 2$, let (α_n, τ_n) be a platform-optimal design and let (p_n, q_n) be an associated symmetric equilibrium. Then every convergent subsequence of (α_n, τ_n) satisfies

$$\alpha_\infty = 1 - \tau_\infty.$$

Proof of Proposition 12. The proof proceeds in four steps. First, the symmetric first-order conditions and the associated wedge and markup identities are recorded. Second, it is shown that for any fixed $\alpha > 0$ the markup term vanishes along any sequence of net prices. Third, feasible platform policies that asymptotically attain the first-best total surplus are constructed and used to pin down the limiting trade probability and quality under the optimal sequence. Fourth, limits are taken in the wedge identity to obtain $\alpha_\infty = 1 - \tau_\infty$.

Step 0: No $\alpha = \infty$ First note that every optimal design has $\alpha_n < \infty$. Indeed, at $\alpha = \infty$, the Bertrand argument of Lemma 3 gives zero seller margin conditional on quality; any $q > 0$ is then defeated by choosing $q' = 0$ and a slightly lower price, so the symmetric outcome has $q = p = 0$ and gives the platform zero profit, whereas a finite-α policy with positive trade and $\tau > 0$ gives strictly positive profit.

Step 1: FOCs. Fix n and let (α_n, τ_n) be a platform-optimal design with associated symmetric equilibrium (p_n, q_n). Write

$$x_n := p_n - q_n, \qquad D_n := 1 - F(x_n)^n,$$

and define

$$I_n := \int_{x_n}^{\bar{v}} F(v)^{n-1} f(v)\, dv = \frac{1 - F(x_n)^n}{n} = \frac{D_n}{n},$$

$$M_n := (1-\tau_n)p_n - c(q_n),$$

$$A_n := F(x_n)^{n-1} f(x_n), \qquad B_n := \int_{x_n}^{\bar{v}} (n-1)F(v)^{n-2} f(v)^2 \, dv.$$

The firm's symmetric first-order conditions are

$$0 = (1-\tau_n)I_n - M_n\big(A_n + \alpha_n B_n\big), \tag{OA.17}$$

$$0 = -c'(q_n)I_n + M_n\big(A_n + B_n\big). \tag{OA.18}$$

Dividing (OA.18) by (OA.17) yields

$$c'(q_n) = (1-\tau_n)\frac{A_n + B_n}{A_n + \alpha_n B_n}. \tag{OA.19}$$

Moreover, (OA.17) implies

$$p_n - \frac{c(q_n)}{1-\tau_n} = \frac{\int_{x_n}^{\bar{v}} F(v)^{n-1} f(v)\, dv}{A_n + \alpha_n B_n} = \frac{(1 - F(x_n)^n)/n}{A_n + \alpha_n B_n}. \tag{OA.20}$$

Step 2: Vanishing markups. Fix $\alpha > 0$ and, for $x \in [\underline{v}, \bar{v}]$, define

$$m_n(x;\alpha) := \frac{\int_x^{\bar{v}} F(v)^{n-1} f(v)\, dv}{F(x)^{n-1} f(x) + \alpha \int_x^{\bar{v}} (n-1)F(v)^{n-2} f(v)^2 \, dv}.$$

It is claimed that for any sequence $x_n \in [\underline{v}, \bar{v}]$, one has $m_n(x_n;\alpha) \to 0$ as $n \to \infty$.

The argument is the same as in Lemma 5. Passing to a subsequence if necessary, suppose $1 - F(x_n)^n \to q \in [0,1]$.

If $q = 1$, then $F(x_n)^n \to 0$. Let $a_n := \max\{x_n, \bar{v} - \varepsilon_0\}$. Since $f(v) \geq \underline{f}$ on $[\bar{v} - \varepsilon_0, \bar{v}]$,

$$(n-1)\int_{x_n}^{\bar{v}} F(v)^{n-2} f(v)^2 \, dv \geq \underline{f} \int_{a_n}^{\bar{v}} (n-1)F(v)^{n-2} f(v)\, dv = \underline{f}\big(1 - F(a_n)^{n-1}\big).$$

If $a_n = \bar{v} - \varepsilon_0$, then $F(a_n) < 1$, so $F(a_n)^{n-1} \to 0$. If $a_n = x_n$, then $F(a_n)^{n-1} = F(x_n)^{n-1} \to 0$ because $F(x_n)^n \to 0$. Hence the denominator of $m_n(x_n;\alpha)$ is eventually bounded below by a positive constant, while the numerator is $(1 - F(x_n)^n)/n \to 0$. Therefore $m_n(x_n;\alpha) \to 0$.

If $q < 1$, then $F(x_n)^n \to 1 - q > 0$. This implies $F(x_n) \to 1$, for otherwise some subsequence would satisfy $F(x_n) \leq \rho < 1$, implying $F(x_n)^n \to 0$, a contradiction. Since $F(x) < 1$ for every $x < \bar{v}$, it follows that $x_n \to \bar{v}$. Hence, for all large n, $x_n \in [\bar{v} - \varepsilon_0, \bar{v}]$, so

$$F(x_n)^{n-1} f(x_n) \geq \underline{f}\, F(x_n)^{n-1}.$$

Moreover, $F(x_n)^{n-1} = F(x_n)^n / F(x_n) \to 1 - q > 0$. Thus the denominator of $m_n(x_n;\alpha)$ is eventually bounded below by a positive constant, while the numerator is again $(1 -$

$F(x_n)^n)/n \to 0$. Therefore $m_n(x_n;\alpha) \to 0$.

So in either case, $m_n(x_n;\alpha) \to 0$.

Step 3: Feasible near-first-best policies. Define $J(q) := \bar{v} + q - c(q)$. Since c is strictly convex, J is strictly concave and has a unique maximizer $q^\star$ satisfying $c'(q^\star) = 1$.

Fix $\varepsilon > 0$ small enough that $\bar{v} - \varepsilon > \underline{v}$ and define

$$\alpha^\varepsilon := 1 - \tau^\varepsilon := \frac{c(q^\star)}{\bar{v} - \varepsilon + q^\star} \in (0,1),$$

Since $c(0) = 0$, strict convexity, and $c'(q^\star) = 1$, one has $c(q^\star) = \int_0^{q^\star} c'(u)\,du < q^\star < \bar{v} - \varepsilon + q^\star$, so indeed $\alpha^\varepsilon \in (0,1)$. Equivalently,

$$\frac{c(q^\star)}{1-\tau^\varepsilon} - q^\star = \bar{v} - \varepsilon. \tag{OA.21}$$

Let $(p_n^\varepsilon, q_n^\varepsilon)$ be a symmetric equilibrium under this policy, and set

$$x_n^\varepsilon := p_n^\varepsilon - q_n^\varepsilon, \qquad D_n^\varepsilon := 1 - F(x_n^\varepsilon)^n.$$

Using (OA.19) with $\alpha^\varepsilon = 1 - \tau^\varepsilon$, one obtains

$$c'(q_n^\varepsilon) = (1-\tau^\varepsilon)\frac{A_n^\varepsilon + B_n^\varepsilon}{A_n^\varepsilon + (1-\tau^\varepsilon)B_n^\varepsilon} \le 1.$$

Since c' is strictly increasing and $c'(q^\star) = 1$, it follows that $q_n^\varepsilon \le q^\star$ for every n.

Next, (OA.20) gives

$$p_n^\varepsilon - \frac{c(q_n^\varepsilon)}{1-\tau^\varepsilon} = m_n(x_n^\varepsilon;\alpha^\varepsilon).$$

Subtracting q_n^ε yields

$$x_n^\varepsilon = \left(\frac{c(q_n^\varepsilon)}{1-\tau^\varepsilon} - q_n^\varepsilon\right) + m_n(x_n^\varepsilon;\alpha^\varepsilon).$$

Define $g(q) := \frac{c(q)}{1-\tau^\varepsilon} - q$. Then g is convex, $g(0) = 0$, and (OA.21) gives $g(q^\star) = \bar{v} - \varepsilon$. Since $0 \le q_n^\varepsilon \le q^\star$, convexity implies

$$g(q_n^\varepsilon) \le \max\{g(0), g(q^\star)\} = g(q^\star) = \bar{v} - \varepsilon.$$

By Step 2, $m_n(x_n^\varepsilon;\alpha^\varepsilon) \to 0$, and therefore

$$\limsup_{n\to\infty} x_n^\varepsilon \le \bar{v} - \varepsilon.$$

Hence, for all large n, $x_n^\varepsilon \le \bar{v} - \varepsilon/2$.

Let $\bar{f} := \sup_{[\underline{v},\bar{v}]} f < \infty$. Then, for all large n,

$$A_n^\varepsilon = F(x_n^\varepsilon)^{n-1} f(x_n^\varepsilon) \le \bar{f}\, F(\bar{v} - \varepsilon/2)^{n-1} \to 0.$$

Let $\delta_\varepsilon := \min\{\varepsilon/2, \varepsilon_0\}$. Since $x_n^\varepsilon \le \bar{v} - \delta_\varepsilon$ for all large n,

$$B_n^\varepsilon \ge \underline{f} \int_{\bar{v}-\delta_\varepsilon}^{\bar{v}} (n-1) F(v)^{n-2} f(v)\, dv = \underline{f}\big[1 - F(\bar{v} - \delta_\varepsilon)^{n-1}\big],$$

so B_n^ε is eventually bounded away from zero. Taking limits in (OA.19) yields $c'(q_n^\varepsilon) \to 1$. Since c' is strictly increasing and $c'(q^\star) = 1$, it follows that $q_n^\varepsilon \to q^\star$. Substituting back gives $x_n^\varepsilon \to \bar{v} - \varepsilon$ and $D_n^\varepsilon \to 1$.

Total profit per consumer under this policy is

$$\big(p_n^\varepsilon - c(q_n^\varepsilon)\big) D_n^\varepsilon = \big(x_n^\varepsilon + q_n^\varepsilon - c(q_n^\varepsilon)\big) D_n^\varepsilon \to \bar{v} - \varepsilon + q^\star - c(q^\star) = J(q^\star) - \varepsilon.$$

Aggregate firm profit per consumer equals $M_n^\varepsilon D_n^\varepsilon$, where

$$M_n^\varepsilon = \frac{(1 - \tau^\varepsilon) I_n^\varepsilon}{A_n^\varepsilon + \alpha^\varepsilon B_n^\varepsilon}.$$

Since $I_n^\varepsilon = D_n^\varepsilon / n \le 1/n$ and the denominator is eventually bounded away from zero, it follows that $M_n^\varepsilon D_n^\varepsilon \to 0$. Therefore the platform profit under this feasible policy satisfies $\Pi_n^{P,\varepsilon} \to J(q^\star) - \varepsilon$.

For every sufficiently small, fixed $\varepsilon > 0$, Step 3 provides a feasible policy with asymptotic platform profit $J(q^\star) - \varepsilon$, so optimality implies

$$\liminf_{n\to\infty} \Pi_n^P \ge J(q^\star) - \varepsilon.$$

Letting $\varepsilon \downarrow 0$ gives $\liminf_{n\to\infty} \Pi_n^P \ge J(q^\star)$.

On the other hand, a firm can always choose a price at which it makes no sale and earns zero, so $M_n \ge 0$, hence $p_n - c(q_n) = M_n + \tau_n p_n \ge 0$. If there is trade, then $x_n \le \bar{v}$, and therefore

$$\Pi_n^P = \tau_n p_n D_n \le \big(p_n - c(q_n)\big) D_n \le p_n - c(q_n) = x_n + q_n - c(q_n) \le \bar{v} + q_n - c(q_n) \le J(q^\star).$$

Thus $\Pi_n^P \to J(q^\star)$.

Also, $\Pi_n^P \le J(q^\star) D_n$, so necessarily $D_n \to 1$, equivalently $F(x_n)^n \to 0$. Finally, $\Pi_n^P \le \bar{v} + q_n - c(q_n) = J(q_n)$, and J is strictly concave with unique maximizer $q^\star$. Since $\Pi_n^P \to J(q^\star)$, it follows that $q_n \to q^\star$ and hence $c'(q_n) \to 1$.

Step 4: Limit of the wedge identity. Consider any convergent subsequence $(\alpha_n, \tau_n) \to (\alpha_\infty, \tau_\infty)$, allowing $\alpha_\infty = \infty$ a priori. By Step 3, $q_n \to q^\star$, $c'(q_n) \to 1$, and $F(x_n)^n \to 0$.

Let $\bar{f} := \sup_{[\underline{v},\bar{v}]} f < \infty$. Then

$$A_n = F(x_n)^{n-1} f(x_n) \le \bar{f}\, F(x_n)^{n-1} \to 0.$$

Moreover,

$$B_n \ge \underline{f} \int_{\max\{x_n, \bar{v}-\varepsilon_0\}}^{\bar{v}} (n-1)F(v)^{n-2} f(v)\, dv = \underline{f}\Big[1 - F\big(\max\{x_n, \bar{v}-\varepsilon_0\}\big)^{n-1}\Big].$$

Since $F(\bar{v}-\varepsilon_0) < 1$ and $F(x_n)^n \to 0$ implies $F(x_n)^{n-1} \to 0$, B_n is eventually bounded away from zero. Hence $A_n/B_n \to 0$.

If $\alpha_n \to \infty$ along this subsequence, then (OA.19) would imply

$$c'(q_n) = (1-\tau_n)\frac{A_n/B_n + 1}{A_n/B_n + \alpha_n} \to 0,$$

contradicting $c'(q_n) \to 1$. Therefore (α_n) is bounded along the subsequence.

Now rewrite (OA.19) as

$$c'(q_n)\left(\frac{A_n}{B_n} + \alpha_n\right) = (1-\tau_n)\left(\frac{A_n}{B_n} + 1\right).$$

Thus

$$c'(q_n)\alpha_n - (1-\tau_n) = \big((1-\tau_n) - c'(q_n)\big)\frac{A_n}{B_n} \to 0.$$

Since (α_n) is bounded along the subsequence and $c'(q_n) \to 1$, it follows that

$$\alpha_n - (1-\tau_n) \to 0.$$

Passing to the limit along the convergent subsequence yields

$$\alpha_\infty = 1 - \tau_\infty.$$

This completes the proof. □

When n is large, competition for being recommended drives markups to zero, so firms choose quality mainly to improve their score at the top of the match distribution. A marginal quality increase dq raises the score one-for-one, while a marginal price increase dp lowers the score by $\alpha\, dp$. Thus, to keep its score (and recommendation probability) roughly unchanged, a firm that raises quality by dq can raise its price by about $dp \approx dq/\alpha$, of which it retains only the fraction $1-\tau$. The resulting marginal return to quality is therefore approximately $(1-\tau)/\alpha$. Since the first-best satisfies $c'(q^\star) = 1$, asymptotic

optimality requires $(1-\tau)/\alpha \to 1$, i.e. $\alpha_\infty = 1 - \tau_\infty$.

OA.3 Inferior Products: Proofs, Auxiliary Derivations, and Additional Results

OA.3.1 Deferred Proofs

Proof of Lemma 9. Let $V_j := \max_{i \le n} v_{ij}$ denote the highest regular-product match value. At the symmetric price p, a regular product would be purchased in the baseline model if and only if $V_j \ge p$. Moreover, U has a weakly higher recommendation score than every regular product if and only if

$$v_U - \alpha p_U \ge V_j - \alpha p \quad \Longleftrightarrow \quad V_j \le v_U + \alpha(p - p_U) = t_U(\alpha, p_U, v_U; p).$$

Thus, almost surely, a misrecommendation occurs on the event $\{p \le V_j \le t_U(\alpha, p_U, v_U; p)\}$. Because V_j has CDF F^n, this event has positive probability if and only if $F\big(t_U(\alpha, p_U, v_U; p)\big)^n > F(p)^n$, which, since $n \ge 2$, is equivalent to $F\big(t_U(\alpha, p_U, v_U; p)\big) > F(p)$.

Finally, suppose $p \in [\underline{v}, \bar{v})$. Since f is strictly positive on $(\underline{v}, \bar{v})$,

$$F\big(t_U(\alpha, p_U, v_U; p)\big) > F(p) \quad \Longleftrightarrow \quad t_U(\alpha, p_U, v_U; p) > p.$$

□

Proof of Lemma 10. For $p > p_U$, Lemma 9 implies that misrecommendations are ruled out whenever $\alpha < \bar{\alpha}(p)$, where $\bar{\alpha}(\cdot)$ is defined in (10). Since $v_U < p_U$, $\bar{\alpha}'(p) = \frac{v_U - p_U}{(p-p_U)^2} < 0$, so $\bar{\alpha}(\cdot)$ is weakly decreasing on (p_U, ∞).

For part (i), let (α^*, τ^*) be an optimal design and let p^* be the induced symmetric regular-product price. Suppose $p^* > p_U$ and, toward a contradiction, that $\alpha^* < \bar{\alpha}(p^*)$. Choose any $\alpha' \in (\alpha^*, \bar{\alpha}(p^*))$. Then the inferior product U is irrelevant at price p^* under both α^* and α'. By the demand and markup derivation in Section OA.3.3.1, the symmetric pricing condition at p^* therefore coincides with the baseline one, namely $p^* - \frac{c}{1-\tau} = m(\alpha, p^*)$.

Since an optimal design yields positive profit, necessarily $p^* < \bar{v}$. For fixed $p^* < \bar{v}$, write

$$m(\alpha, p^*) = \frac{N(p^*)}{A(p^*) + \alpha B(p^*)},$$

where $N(p^*) := \frac{1 - F(p^*)^n}{n}$, $A(p^*) := F(p^*)^{n-1} f(p^*)$, and

$$B(p^*) := (n-1) \int_{p^*}^{\bar{v}} F(v)^{n-2} f(v)^2 \, dv.$$

Since $N(p^*) > 0$ and $B(p^*) > 0$, the markup $m(\alpha, p^*)$ is strictly decreasing in α. Hence $m(\alpha', p^*) < m(\alpha^*, p^*)$. Therefore there exists a unique $\tau' > \tau^*$ such that $p^* - \frac{c}{1-\tau'} = m(\alpha', p^*)$.

Equivalently, under the policy (α', τ'), the price p^* satisfies the corresponding symmetric first-order condition. By the maintained assumption for Section 5, the pricing subgame under (α', τ') admits a unique interior symmetric pure-strategy equilibrium, and its equilibrium price is the unique interior solution to that first-order condition. Hence, p^* is the symmetric equilibrium regular-product price under (α', τ'). Because the inferior product U is irrelevant at p^*, total quantity sold remains $1 - F(p^*)^n$, while platform revenue per sale rises from $\tau^* p^*$ to $\tau' p^*$. This strictly increases platform profit, contradicting the optimality of (α^*, τ^*). Therefore $\alpha^* \geq \bar{\alpha}(p^*)$.

For part (ii), suppose $p_U > p_m(c)$.

Let $Q(p) := 1 - F(p)^n$ and $\pi(p) := (p - c)Q(p)$. If a design induces price p and trade probability D, platform profit is $\Pi = \tau p\, D$. In any equilibrium with positive sales, firms must have nonnegative operating margins, so $(1-\tau)p \geq c$ and therefore $\tau p \leq p - c$. Also, trade can occur only if at least one regular product yields nonnegative utility at price p, so $D \leq Q(p)$. Hence $\Pi \leq (p - c)D \leq (p - c)Q(p) = \pi(p)$. By Definition 1 and the standard monopoly price properties stated after Definition 1, π is uniquely maximized at $p_m(c)$, so $\Pi \leq \pi\big(p_m(c)\big)$ under any design.

Now choose $\tau^\star$ so that $p_m(c) = \frac{c}{1-\tau^\star}$. Under $\alpha = \infty$, downstream markups vanish, and the induced price is $p_m(c)$. Since $p_U > p_m(c)$, product U is strictly more expensive than the regular products and is never recommended. Thus, trade probability is $1 - F\big(p_m(c)\big)^n$. Hence the platform attains profit $\Pi(\infty, \tau^\star) = \pi\big(p_m(c)\big)$, which reaches the upper bound.

Moreover, the upper bound is strict under every finite α. Indeed, in any finite-α equilibrium with positive sales, firms must have strictly positive operating margins; otherwise, a sufficiently small price increase would preserve positive demand while yielding a positive margin. Hence $(1 - \tau)p > c$, so $\tau p < p - c$ and therefore

$$\Pi < (p - c)D \leq \pi(p) \leq \pi\big(p_m(c)\big).$$

A design with zero sales yields zero platform profit. Therefore $\alpha^* = \infty$. □

Proof of Proposition 10. Part (i) is proven first. Suppose $p_U < c$ and, toward a contradiction, that some optimal design has $\alpha^* = \infty$. Under $\alpha = \infty$, the platform recommends the cheapest product. Since every regular product that is sold must yield a weakly nonnegative operating margin, any sale by a regular firm at price p_i requires $(1 - \tau)p_i - c \geq 0$, hence $p_i \geq c/(1 - \tau) \geq c > p_U$. Thus, U is strictly cheapest whenever any regular product could be sold, so U is always recommended. Because $v_U < p_U$, product U is never purchased; and since consumers do not search after a recommendation of U, no regular product is ever purchased either. Therefore, every design with $\alpha = \infty$ yields zero

platform profit.

Now consider instead $\alpha = 1$ and choose any $\tau \in (0,1)$ with $(1-\tau)\bar{v} > c$. For any price profile p, if a regular product yields nonnegative utility for consumer j, i.e. $v_{ij} - p_i \geq 0 > v_U - p_U$ for some i, then U cannot be recommended. If all regular products yield negative utility, then there is no trade whether or not U is recommended. Hence, under $\alpha = 1$, the pricing game with U coincides with the baseline pricing game without U. By Proposition 1, a symmetric equilibrium exists, and its price p satisfies

$$p - \frac{c}{1-\tau} = m(1,p) > 0.$$

Thus $p < \bar{v}$, so $1 - F(p)^n > 0$, and platform profit is

$$\tau p\big(1 - F(p)^n\big) > 0.$$

This contradicts optimality of $\alpha^* = \infty$. Therefore $\alpha^* < \infty$.

Part (ii) is proven next. Suppose $p_U = 0$ and $v_U = 0$. By Lemma 10, any optimal design satisfies

$$\alpha^* \geq \bar{\alpha}(p^*) = 1,$$

where the equality follows from $p_U = v_U = 0$.

It remains to show that no $\alpha > 1$ can be optimal. Fix a candidate optimal design (α, τ) with $\alpha > 1$, and let p be its symmetric equilibrium price. Write $t := t_U(\alpha, 0, 0; p) = \alpha p$. Since a design yielding strictly positive platform profit is feasible, optimality requires $p > 0$ and $t < \bar{v}$.

For any deviation $p_i \geq 0$, firm i's demand under high search costs is

$$D_i^U(p_i, p; 0, 0) = \int_{\alpha p_i}^{\bar{v}} F\big(v - \alpha p_i + \alpha p\big)^{n-1} f(v)\, dv,$$

because when $\alpha > 1$ the relevant cutoff is $t_U(\alpha, 0, 0; p_i) = \alpha p_i$. Now define an alternative design (α', τ') by $\alpha' = 1$ and

$$1 - \tau' = \frac{1-\tau}{\alpha},$$

and let the candidate symmetric price be $p' := \alpha p = t$. For any deviation $p_i \geq 0$, set $p_i' := \alpha p_i$. Under $(\alpha', \tau') = (1, \tau')$, firm i's demand at deviation p_i' against rivals' price p' is

$$D_i^U(p_i', p'; 0, 0) = \int_{p_i'}^{\bar{v}} F\big(v - p_i' + p'\big)^{n-1} f(v)\, dv = \int_{\alpha p_i}^{\bar{v}} F\big(v - \alpha p_i + \alpha p\big)^{n-1} f(v)\, dv = D_i^U(p_i, p; 0, 0).$$

Moreover,

$$(1-\tau')p_i' - c = (1-\tau')\alpha p_i - c = (1-\tau)p_i - c.$$

Thus the deviation payoff under $(1, \tau')$ at $p_i' = \alpha p_i$ is exactly the same as the deviation payoff under (α, τ) at p_i. Since p is a best response to p under (α, τ), it follows that $p' = \alpha p$ is a best response to itself under $(1, \tau')$. Therefore $(1, \tau')$ induces a symmetric equilibrium with cutoff

$$t_U(1, 0, 0; p') = p' = \alpha p = t_U(\alpha, 0, 0; p),$$

so the purchasing set, and hence the trade probability, is unchanged.

Platform profit under the original design is

$$\Pi(\alpha, \tau) = \tau p\big(1 - F(t)^n\big) = \tau \frac{t}{\alpha}\big(1 - F(t)^n\big),$$

whereas under the transformed design it is

$$\Pi(1, \tau') = \tau' t\big(1 - F(t)^n\big).$$

Finally,

$$\alpha\tau' = \alpha\left(1 - \frac{1-\tau}{\alpha}\right) = \alpha - 1 + \tau > \tau,$$

because $\alpha > 1$. Hence $\Pi(1, \tau') > \Pi(\alpha, \tau)$, contradicting the optimality of (α, τ). Therefore $\alpha^* \leq 1$. Combined with $\alpha^* \geq 1$, this yields $\alpha^* = 1$. □

Proof of Proposition 11. Since $p_U = v_U = 0$, Lemma 10 implies $\alpha^*(s) \geq 1$ for every s. It therefore suffices to rule out $\alpha^*(s) = 1$ for all sufficiently small s.

Suppose not. Then there exists a sequence $s_k \downarrow 0$ and, for each k, a platform-optimal policy with $\alpha_k^* = 1$. Let p_k be the induced symmetric equilibrium price and τ_k the associated commission.

Choose a benchmark policy $(1, \tau_0)$ with $\tau_0 \in (0, 1)$ and $(1-\tau_0)\bar{v} > c$, and let $p_0 \in (0, \bar{v})$ be an induced symmetric equilibrium price. Its platform profit, $\pi_0 := \tau_0 p_0(1 - F(p_0)^n)$, is strictly positive. Since (α_k^*, τ_k) is optimal at s_k,

$$\tau_k p_k\big(1 - F(p_k)^n\big) \geq \pi_0 \qquad \text{for all } k.$$

Because $\tau_k \leq 1$, it follows that $p_k(1 - F(p_k)^n) \geq \pi_0$, so $p_k \not\to \bar{v}$; otherwise $F(p_k) \to 1$ and the left-hand side would converge to 0. Passing to a subsequence if necessary, assume $p_k \to \bar{p} < \bar{v}$.

Fix $\alpha > 1$ and a symmetric candidate price p^*. Let $\hat{x}(\alpha, s, p^*)$ denote the reservation value after a recommendation of U, as characterized in Section OA.3.2.2. Also let

$$t_U(\alpha, p_U, v_U; q) := v_U + \alpha(q - p_U)$$

denote the highest regular-product match value consistent with an initial recommendation of U when the regular product charges q. In the benchmark considered here, $p_U = v_U = 0$, so $t_U(\alpha, 0, 0; q) = \alpha q$.

Suppose $\alpha > 1$ and s is sufficiently small that search after U is active. Then, locally around $p_i = p^*$, firm i's demand following Wolinsky (1986) is

$$\begin{aligned} D_i(p_i, p^*) = &\int_{t_U(\alpha,0,0;p_i)}^{\bar{v}} F(v - \alpha p_i + \alpha p^*)^{n-1}\, f(v)\, dv \\ &+ \int_{p_i}^{\hat{x}(\alpha,s,p^*)-p^*+p_i} F(v - p_i + p^*)^{n-1}\, f(v)\, dv \\ &+ \int_{\hat{x}(\alpha,s,p^*)-p^*+p_i}^{t_U(\alpha,0,0;p_i)} \frac{F(t_U(\alpha,0,0;p^*))^n - F(\hat{x}(\alpha,s,p^*))^n}{n\big(F(t_U(\alpha,0,0;p^*)) - F(\hat{x}(\alpha,s,p^*))\big)}\, f(v)\, dv. \end{aligned} \tag{OA.22}$$

The three terms correspond to three regions of firm i's match value. First, if $v \geq t_U(\alpha, 0, 0; p_i)$, firm i beats U in the recommendation score and is purchased directly whenever it is recommended. Second, if $p_i \leq v < \hat{x}(\alpha, s, p^*) - p^* + p_i$, firm i's product is purchasable but below the reservation threshold, so it is bought only after all regular products were searched. Third, if $\hat{x}(\alpha, s, p^*) - p^* + p_i \leq v < t_U(\alpha, 0, 0; p_i)$, firm i is reached through search after U is recommended, but it is good enough to stop search once encountered.

By (OA.22), at a symmetric price vector

$$D(p) := D_i(p, p; s, \alpha) = \frac{1 - F(p)^n}{n}.$$

Define

$$\Lambda(\alpha, p, s) := -\partial_{p_i} D_i(p_i, p; s, \alpha)\big|_{p_i=p}, \qquad \hat{m}(\alpha, p, s) := \frac{D(p)}{\Lambda(\alpha, p, s)}.$$

The symmetric first-order condition for firms' profit maximization in the pricing subgame can therefore be written as

$$p - \frac{c}{1-\tau} = \hat{m}(\alpha, p, s).$$

By the Leibniz rule,

$$\begin{aligned} \Lambda(\alpha, p, s) = &F(p)^{n-1} f(p) + (n-1) \int_p^{\hat{x}(\alpha,s,p)} F(v)^{n-2} f(v)^2\, dv \\ &+ \alpha(n-1) \int_{\alpha p}^{\bar{v}} F(v)^{n-2} f(v)^2\, dv \\ &+ \alpha \left[F(\alpha p)^{n-1} - \frac{F(\alpha p)^n - F(\hat{x}(\alpha,s,p))^n}{n\big(F(\alpha p) - F(\hat{x}(\alpha,s,p))\big)} \right] f(\alpha p) \\ &- \left[F(\hat{x}(\alpha,s,p))^{n-1} - \frac{F(\alpha p)^n - F(\hat{x}(\alpha,s,p))^n}{n\big(F(\alpha p) - F(\hat{x}(\alpha,s,p))\big)} \right] f(\hat{x}(\alpha,s,p)). \end{aligned}$$

For $\alpha > 1$, as $s \downarrow 0$, $\hat{x}(\alpha, s, p) \uparrow \alpha p$. The resulting convergence below is locally uniform in p. Hence, by continuity of F and f,

$$\frac{F(\alpha p)^n - F(\hat{x}(\alpha, s, p))^n}{n\big(F(\alpha p) - F(\hat{x}(\alpha, s, p))\big)} \to F(\alpha p)^{n-1},$$

so $\Lambda(\alpha, p, s) \to \Lambda^{\lim}(\alpha, p)$, where

$$\Lambda^{\lim}(\alpha, p) = F(p)^{n-1} f(p) + (n-1) \int_p^{\alpha p} F(v)^{n-2} f(v)^2 \, dv + \alpha(n-1) \int_{\alpha p}^{\bar{v}} F(v)^{n-2} f(v)^2 \, dv.$$

At $\alpha = 1$, search after U never starts, and direct differentiation gives $\Lambda(1, p, s) = \Lambda^{\lim}(1, p)$ and $\hat{m}(1, p, s) = m(1, p)$. For $\alpha > 1$,

$$\Lambda^{\lim}(\alpha, p) - \Lambda^{\lim}(1, p) = (\alpha - 1)(n-1) \int_{\alpha p}^{\bar{v}} F(v)^{n-2} f(v)^2 \, dv,$$

so $\Lambda^{\lim}(\alpha, p) > \Lambda^{\lim}(1, p)$ whenever $\alpha p < \bar{v}$.

Choose $\varepsilon > 0$ such that $\underline{v} < (1+\varepsilon)\bar{p} < \bar{v}$. Then $\Lambda^{\lim}(1+\varepsilon, \bar{p}) > \Lambda^{\lim}(1, \bar{p})$, and therefore $\Lambda(1+\varepsilon, p_k, s_k) > \Lambda(1, p_k, s_k)$ for all sufficiently large k. Since $D(p_k) = \frac{1-F(p_k)^n}{n}$ does not depend on α, it follows that

$$\hat{m}(1+\varepsilon, p_k, s_k) < \hat{m}(1, p_k, s_k)$$

for all sufficiently large k.

Fix such a k and set $\alpha' := 1 + \varepsilon$.

Since $(\alpha_k^*, \tau_k) = (1, \tau_k)$ induces the symmetric equilibrium price p_k, the assumption for the low-search-cost analysis implies that p_k satisfies the symmetric pricing first-order condition. Hence

$$p_k - \frac{c}{1 - \tau_k} = \hat{m}(1, p_k, s_k).$$

Now define $\tilde{\tau}_k > \tau_k$ by

$$p_k - \frac{c}{1 - \tilde{\tau}_k} = \hat{m}(1+\varepsilon, p_k, s_k).$$

Under the policy $(\alpha', \tilde{\tau}_k)$, the same price p_k satisfies the symmetric pricing first-order condition. By the maintained assumption for Section 5, this first-order condition characterizes the symmetric equilibrium price, so p_k is the symmetric equilibrium price under $(\alpha', \tilde{\tau}_k)$.

At that price, total quantity sold remains $1 - F(p_k)^n$, while per-sale platform revenue rises from $\tau_k p_k$ to $\tilde{\tau}_k p_k$. Hence, platform profit is strictly higher, contradicting the optimality of the policy with $\alpha_k^* = 1$.

Therefore, there exists $\bar{s} > 0$ such that every platform-optimal choice satisfies

$$\alpha^*(s) > 1 \qquad \text{for all } s \in (0, \bar{s}).$$

□

OA.3.2 Recommendation of the Inferior Product

OA.3.2.1 Equivalent Condition

A useful restatement of Lemma 9 is the following corollary, which makes transparent how the direction of the distortion depends on whether p_U is low or high.

Corollary 2 (Equivalent Forms)
Fix (p_U, v_U) with $v_U < p_U$ and a symmetric regular-product price $p \in [\underline{v}, \bar{v})$.

1. *If $p > p_U$, then misrecommendation at p occurs with positive probability if and only if*
$$\alpha \;>\; \frac{p - v_U}{p - p_U} \;>\; 1.$$

2. *If $p < p_U$, then misrecommendation at p occurs with positive probability if and only if*
$$0 \le \alpha \;<\; \frac{p - v_U}{p - p_U}.$$
In particular, if $v_U \le p < p_U$, then the right-hand side is ≤ 0, so misrecommendation is impossible for any $\alpha \ge 0$.

3. *If $p = p_U$, then misrecommendation is impossible, since it would require $v_U \ge p_U$.*

If U is relatively cheap ($p_U < p$), misrecommendations arise only for sufficiently large α, because the ranking becomes price-driven and U can outrank regular products despite its low match value. If instead U is expensive ($p_U > p$), misrecommendations can occur only when U's match value is high enough (necessarily $v_U > p$) and the ranking is sufficiently insensitive to price (small α).

OA.3.2.2 Search After Observing U

Fix a conjectured symmetric regular-product price p^* and algorithm parameter α, and define

$$t := t_U(\alpha, p_U, v_U; p^*), \qquad t_U(\alpha, p_U, v_U; p) := v_U + \alpha(p - p_U).$$

If $F(t) = 0$, then U is recommended with probability zero, so no posterior following a recommendation of U is required. Henceforth suppose $F(t) > 0$. Conditional on a

recommendation of U, the regular match values are independent and distributed as F truncated above at t. Denote this posterior cdf by F_U, i.e.

$$F_U(x) := \begin{cases} 0, & x < \underline{v}, \\ \dfrac{F(x)}{F(t)}, & \underline{v} \le x \le t, \\ 1, & x > t. \end{cases}$$

If $p^* > p_U$, then $t_U(\alpha, p_U, v_U; p^*)$ is increasing in α, so the posterior F_U shifts upward in the sense of first-order stochastic dominance, making search more attractive.

The key object is the posterior probability that an unseen regular product is purchasable at price p^*.

Corollary 3 (Posterior Support above the Purchase Threshold)
Conditional on a recommendation of U, the posterior assigns positive probability to a regular product with non-negative utility if and only if

$$F(t) > F(p^*).$$

When $p^* \in [\underline{v}, \bar{v})$, this is equivalent to $t > p^*$. More precisely, for each regular product,

$$\Pr(v_{ij} \ge p^* \mid U \text{ recommended}) = \begin{cases} 0, & F(t) \le F(p^*), \\ \dfrac{F(t) - F(p^*)}{F(t)}, & F(t) > F(p^*). \end{cases}$$

Hence, if $F(t) > F(p^*)$, the probability that at least one of the n regular products is purchasable under the posterior is

$$1 - \left(\frac{F(p^*)}{F(t)}\right)^n > 0.$$

In particular, when $p^* \in [\underline{v}, \bar{v})$, $p_U < p^*$, and

$$\alpha > \frac{p^* - v_U}{p^* - p_U},$$

one has $t_U(\alpha, p_U, v_U; p^*) > p^*$, so conditional on observing U there is a positive probability that some unseen regular product would be purchased. Taking $p^* = p_m(c)$, this condition is equivalent to $\alpha > \bar{\alpha}(p_m(c)) = \frac{p_m(c) - v_U}{p_m(c) - p_U}$.

If instead $F(t) \le F(p^*)$, then under the posterior, no regular product yields non-negative utility with positive probability, so search is never optimal after a recommendation of U.

Conditional on a recommendation of U, the regular-product values are independent,

and each has marginal distribution F_U. By consistency, observing off-equilibrium prices of inspected products does not change beliefs about the prices charged by unseen firms. Hence, after observing U, each unseen regular product is conjectured to be priced at p^*, and its match value is governed by F_U. The continuation problem after observing U is therefore the sequential-search problem over the remaining regular products under (p^*, F_U). Whenever search is started after observing U, sequential rationality requires the consumer's search behavior, after each subsequent search outcome that can arise in this continuation problem, to coincide with the optimal search rule. Outcomes not generated under that behavior are irrelevant to the characterization below.

Suppose therefore that $F(t) > F(p^*)$. Let $\hat{x}(\alpha, s, p^*)$ denote the reservation value defined by

$$\int_{\hat{x}(\alpha,s,p^*)}^{t} \big(v - \hat{x}(\alpha, s, p^*)\big)\, dF_U(v) = s.$$

By the characterization of this continuation problem following Wolinsky (1986), the optimal search rule after a recommendation of U is the reservation rule governed by $\hat{x}(\alpha, s, p^*)$: after each subsequent search outcome, the consumer continues searching if and only if the best currently inspected regular option yields net utility $v - p$ below

$$\hat{x}(\alpha, s, p^*) - p^*.$$

In particular, immediately after observing U, the consumer starts searching if and only if

$$\hat{x}(\alpha, s, p^*) > p^*.$$

Equivalently, define

$$\bar{s}(\alpha, p^*) := \int_{p^*}^{t} (v - p^*)\, dF_U(v).$$

Then the consumer starts searching after observing U if and only if

$$s < \bar{s}(\alpha, p^*).$$

Indeed, by the defining equation for $\hat{x}(\alpha, s, p^*)$, one has $\hat{x}(\alpha, s, p^*) > p^*$ if and only if

$$s < \int_{p^*}^{t} (v - p^*)\, dF_U(v) = \bar{s}(\alpha, p^*).$$

Whenever $F(t) > F(p^*)$, Corollary 3 implies that $F_U\big((p^*, t]\big) = \frac{F(t)-F(p^*)}{F(t)} > 0$, and therefore $\bar{s}(\alpha, p^*) = \int_{p^*}^{t} (v - p^*)\, dF_U(v) > 0$. Hence, whenever $F(t) > F(p^*)$, a recommendation of U triggers search for all sufficiently small search costs.

This also yields the high-search-cost region used in the text. If $s > \bar{s}(\alpha, p^*)$, then starting to search is not optimal after a recommendation of U. Since a recommendation

of U is an on-path history whenever it occurs with positive probability, no symmetric pure-strategy PBE can prescribe search after U in that case.

Finally, as $s \downarrow 0$, one has $\hat{x}(\alpha, s, p^*) \uparrow \min\{t, \bar{v}\}$, so the consumer nearly exhausts the search pool. In particular, her final choice approaches the full-information rule

$$\arg\max_i \{v_{ij} - p_i\},$$

that is, for sufficiently small search costs, the consumer behaves as if the relevant price weight were 1 regardless of the platform's α.

OA.3.3 High Search Costs

OA.3.3.1 Demand and Markups

Conditional on consumers never starting to search, the demand of firm i when it sets price p_i and all other regular products set price p is

$$D_i^U(p_i, p; p_U, v_U) \;=\; \int_{t(p_i)}^{\bar{v}} F\big(v - \alpha p_i + \alpha p\big)^{n-1} f(v)\, dv, \qquad t(p_i) := \max\{p_i,\; v_U + \alpha(p_i - p_U)\}. \tag{OA.23}$$

Whenever $t(p) \in (\underline{v}, \bar{v})$ and $t(\cdot)$ is differentiable at p, the first-order condition for firm i in a symmetric equilibrium can be written as

$$p - \frac{c}{1-\tau} = \tilde{m}(\alpha, p; p_U, v_U), \tag{OA.24}$$

where

$$\tilde{m}(\alpha, p; p_U, v_U) := \frac{1 - F\big(t(p)\big)^n}{n} \frac{1}{t'(p)\, F\big(t(p)\big)^{n-1} f\big(t(p)\big) + \alpha(n-1)\int_{t(p)}^{\bar{v}} F(v)^{n-2} f(v)^2\, dv}.$$

If $v_U + \alpha(p - p_U) < p$, then $t(p) = p$ and $t'(p) = 1$, so

$$\tilde{m}(\alpha, p; p_U, v_U) = m(\alpha, p) :$$

the inferior product U is irrelevant for pricing at p.

If instead $v_U + \alpha(p - p_U) > p$, then $t(p) = v_U + \alpha(p - p_U)$ and $t'(p) = \alpha$, so the cutoff a regular product must clear to be recommended is strictly higher than in the baseline. This affects demand with positive probability if and only if $F\big(t(p)\big) > F(p)$. When $p \in [\underline{v}, \bar{v})$, this is equivalent to $t(p) > p$. In the empirically relevant case $p \in [\underline{v}, \bar{v})$ and $p > p_U$, this occurs when α is large enough, as characterized in Corollary 2; in that case, the markup is pointwise lower than the baseline markup.

OA.3.3.2 Asymptotic Price Sensitivity

The main text studies the inferior-product extension under high search costs through the threshold logic in Lemma 10, the finiteness result and the sharp benchmark $p_U = v_U = 0$ in Proposition 10. One can also sharpen the large-market conclusion for general (p_U, v_U) as $n \to \infty$.

Recall the threshold $\bar{\alpha}(p)$ defined in (10). At a symmetric regular-product price $p > p_U$, misrecommendations are impossible whenever

$$\alpha \leq \bar{\alpha}(p).$$

Thus $\bar{\alpha}(p)$ is the threshold at which $t_U(\alpha, p_U, v_U; p) = p$. When $p \in [\underline{v}, \bar{v})$, it is also the threshold above which misrecommendations occur with positive probability. Since $v_U < p_U$, the map $\bar{\alpha}(\cdot)$ is strictly decreasing on (p_U, ∞).

Since $\bar{\alpha}(\cdot)$ is continuous on (p_U, ∞), define the limiting threshold

$$\bar{\alpha}_\infty(p_U, v_U) := \lim_{p \uparrow \bar{v}} \bar{\alpha}(p) = \frac{\bar{v} - v_U}{\bar{v} - p_U}.$$

The next result shows that, as $n \to \infty$, the platform's optimal choice of α converges to this limiting threshold.

Proposition 13 (Asymptotic Behavior of $\alpha^*(n, p_U, v_U)$)
Fix $0 \leq v_U < p_U < \bar{v}$. Then

$$\alpha^*(n, p_U, v_U) \to \bar{\alpha}_\infty(p_U, v_U).$$

Proof. Write $\bar{\alpha}_\infty := \bar{\alpha}_\infty(p_U, v_U)$. First, it is shown that $\limsup_{n\to\infty} \alpha^*(n, p_U, v_U) \leq \bar{\alpha}_\infty$. Suppose not. Then there exist $\tilde{\alpha} > \bar{\alpha}_\infty$ and a subsequence $n_k \to \infty$ such that $\alpha^*(n_k, p_U, v_U) \geq \tilde{\alpha}$ for all k.

Consider any $\alpha \in [\tilde{\alpha}, \infty]$ and a symmetric equilibrium under (α, τ) with positive trade. If $\alpha = \infty$, positive regular-product trade requires $p \leq p_U < p_{\max}(\tilde{\alpha})$. If $\alpha < \infty$, some regular product must beat U at some match value $v \leq \bar{v}$, so $v - \alpha p \geq v_U - \alpha p_U$, and hence

$$p \leq p_{\max}(\alpha) := p_U + \frac{\bar{v} - v_U}{\alpha} \leq p_{\max}(\tilde{\alpha}) < \bar{v},$$

where the last inequality uses $\tilde{\alpha} > \bar{\alpha}_\infty$. Thus, in either case, $p \leq p_{\max}(\tilde{\alpha})$. Also, positive trade and nonnegative equilibrium firm profit imply $(1 - \tau)p \geq c$, and hence $\tau p \leq p - c$. Since purchase probability is at most one,

$$\Pi_n(\alpha, \tau; p_U, v_U) \leq p_{\max}(\tilde{\alpha}) - c.$$

Hence

$$\Pi_{n_k}\big(\alpha^*(n_k, p_U, v_U), \tau^*(n_k, p_U, v_U); p_U, v_U\big) \le p_{\max}(\tilde{\alpha}) - c \qquad \text{for all } k.$$

Now choose

$$p' \in \big(\max\{p_{\max}(\tilde{\alpha}), c, \underline{v}\}, \bar{v}\big) \quad \text{and set} \quad \alpha' := \bar{\alpha}_\infty.$$

Since $\bar{\alpha}(\cdot)$ is strictly decreasing and $\bar{\alpha}(\bar{v}) = \bar{\alpha}_\infty$, one has $\alpha' < \bar{\alpha}(p')$, so the inferior product U is slack at p' under α'. Let τ_n solve

$$p' - \frac{c}{1-\tau_n} = m_n(\alpha', p').$$

Because $p' > c$ and $m_n(\alpha', p') \to 0$, $\tau_n \in [0, 1)$ for all large n and $\tau_n p' \to p' - c$. Since U is strictly slack at p', this is also the full pricing game's symmetric first-order condition; by the maintained uniqueness assumption, p' is the induced equilibrium price, and purchase occurs whenever at least one regular product has value above p'. So,

$$\Pi_n(\alpha', \tau_n; p_U, v_U) = \tau_n p'\big(1 - F(p')^n\big) \to p' - c.$$

As $p' > p_{\max}(\tilde{\alpha})$, one has $p' - c > p_{\max}(\tilde{\alpha}) - c$, so for all large k,

$$\Pi_{n_k}(\alpha', \tau_{n_k}; p_U, v_U) > p_{\max}(\tilde{\alpha}) - c,$$

contradicting the optimality of $\alpha^*(n_k, p_U, v_U)$. Therefore

$$\limsup_{n\to\infty} \alpha^*(n, p_U, v_U) \le \bar{\alpha}_\infty(p_U, v_U).$$

For the reverse inequality, let p_n^* be the symmetric regular-product price induced by an optimal design (α_n^*, τ_n^*), and let Π_n^* denote the corresponding platform profit. The same feasible-policy construction can be carried out for every fixed $p' \in (\max\{c, \underline{v}, p_U\}, \bar{v})$ such that $\alpha' := \bar{\alpha}_\infty < \bar{\alpha}(p')$. For each such p', choosing τ_n' from $p' - \frac{c}{1-\tau_n'} = m_n(\alpha', p')$

yields a feasible policy with asymptotic platform profit $p' - c$. Hence, for every sufficiently small $\varepsilon > 0$, taking $p' = \bar{v} - \varepsilon$ gives

$$\liminf_{n\to\infty} \Pi_n^* \ge \bar{v} - \varepsilon - c.$$

Since $\varepsilon > 0$ is arbitrary and always $\Pi_n^* \le \bar{v} - c$, it follows that $\Pi_n^* \to \bar{v} - c$.

On the other hand, if D_n^* is the associated trade probability, then $\Pi_n^* = \tau_n^* p_n^* D_n^* \le (p_n^* - c) D_n^* \le p_n^* - c \le \bar{v} - c$, where the first inequality uses $(1 - \tau_n^*) p_n^* \ge c$. Hence $p_n^* \to \bar{v}$.

Therefore $p_n^* > p_U$ for all sufficiently large n, and Lemma 10 yields $\alpha^*(n, p_U, v_U) \ge$

$\bar{\alpha}(p_n^*)$. Since $\bar{\alpha}(\cdot)$ is continuous on (p_U, ∞) and $p_n^* \to \bar{v}$, it follows that $\bar{\alpha}(p_n^*) \to \bar{\alpha}(\bar{v}) = \bar{\alpha}_\infty(p_U, v_U)$. Thus

$$\liminf_{n\to\infty} \alpha^*(n, p_U, v_U) \geq \bar{\alpha}_\infty(p_U, v_U).$$

Combining this with the lim sup bound yields

$$\alpha^*(n, p_U, v_U) \to \bar{\alpha}_\infty(p_U, v_U).$$

□

Thus, in large markets, the platform chooses the threshold level that just keeps the inferior product U from binding at the limiting benchmark price. The gain from raising α further vanishes with n because markup compression disappears, while the distortion from misrecommendations does not.

OA.4 Free entry

This appendix studies how endogenous entry changes the platform's joint choice of recommendation design and monetization. Because entry is costly, the platform must leave participating firms positive rents. Since greater recommendation price-sensitivity intensifies price competition and compresses those rents, endogenous entry can discipline the platform's choice of α. Through the free-entry condition, the platform's policy also determines the equilibrium mass of active firms. Relative to the baseline model, entry is now endogenous, and for tractability, I extend the symmetric pricing condition derived for integer $n \geq 2$ to continuous $n \geq 2$. Under commission caps, the induced entry response can further alter the welfare effects of regulation.

Set-up. Extend the baseline environment by allowing endogenous entry. There is a continuum of ex ante identical potential firms. An entering firm lists one product on the platform and incurs a fixed entry cost $k > 0$.

Let $n \in [2, \infty)$ denote the mass of entrants, treated as continuous for tractability. For integer $n \geq 2$, the symmetric pricing condition is the one from the baseline model. For real $n \geq 2$, I define the symmetric price by continuous interpolation of that symmetric pricing condition.

The timing is:

1. The platform chooses (α, τ) (possibly subject to a cap $\tau \leq \bar{\tau}$).

2. Firms enter. In a symmetric free-entry equilibrium, entry continues until expected operating profit equals k, determining n.

3. Given (α, τ, n), the symmetric price $p(\alpha, \tau, n)$ is determined by the continuous interpolation of the baseline symmetric pricing condition.

4. Consumers observe the recommendation and make their search and purchase decisions as in the baseline model.

For $n \in [2, \infty)$, let $p(\alpha, \tau, n)$ denote the symmetric price induced by this interpolated pricing condition. Given a commission rate τ, a price p, and mass n, a firm's profit is

$$\pi(p, n; \tau) = \big[(1-\tau)p - c\big]\frac{1 - F(p)^n}{n}.$$

A symmetric free-entry equilibrium is a tuple $(\alpha^*, \tau^*, p^*, n^*)$ such that $p^* = p(\alpha^*, \tau^*, n^*)$ and the free-entry condition holds:

$$\pi(p^*, n^*; \tau^*) = k.$$

Thus entry is pinned down by the zero-profit free-entry condition. We focus on symmetric equilibria and maintain the assumptions on F used in the baseline analysis.

Endogenous entry adds an additional constraint to the platform's problem: for any policy (α, τ), the induced equilibrium must leave firms enough operating profit to cover the entry cost k. This may force the platform to limit recommendation price-sensitivity in order to sustain entry.

When entry costs are very large, the free-entry condition may be inconsistent with $n \geq 2$ (or may require such large rents that the platform prefers an outcome with a single active firm). We therefore focus on sufficiently small entry costs under which the equilibrium features at least two active firms.

Proposition 14 (Free entry with unconstrained commission rate)
Suppose that for every finite $\alpha \geq 0$, every $\tau \in [0, 1)$, and every $n \in [2, \infty)$, the interpolated symmetric pricing condition admits a unique symmetric price $p(\alpha, \tau, n)$.

Then there exists $k_0 > 0$ such that for all $k \in (0, k_0]$, any symmetric free-entry equilibrium $(\alpha^*(k), \tau^*(k), p^*(k), n^*(k))$ satisfies

$$n^*(k) \geq 2 \qquad \textit{and} \qquad \alpha^*(k) < \infty.$$

Moreover, $n^*(k) \to \infty$ as $k \downarrow 0$.

Proof. Write $D(p, n) := 1 - F(p)^n$, $\pi(p, n; \tau) := \big[(1-\tau)p - c\big]D(p, n)/n$, and $\Pi(\tau, p, n) := \tau p\, D(p, n)$.

The proof has two parts. First, we construct a feasible policy with finite price sensitivity that, for small entry costs, supports a large mass of entrants and yields platform profit arbitrarily close to $\bar{v} - c$. Second, we show that any sequence of free-entry equilibria

with entry bounded by some finite N yields platform profit uniformly bounded away from $\bar{v} - c$. Hence, equilibrium entry must diverge as $k \downarrow 0$. Once $n^*(k) \geq 2$, the extreme policy $\alpha = \infty$ is impossible, because it induces Bertrand pricing and zero firm rents.

Fix $\varepsilon > 0$ sufficiently small that $\bar{v} - \varepsilon > \underline{v}$, choose τ^ε so that $c/(1-\tau^\varepsilon) = \bar{v} - \varepsilon$, and fix any finite $\alpha^\varepsilon > 0$. For each $n \geq 2$, let $p^\varepsilon(n) := p(\alpha^\varepsilon, \tau^\varepsilon, n)$ denote the unique symmetric equilibrium price under this fixed policy, and let $m_n^\varepsilon := p^\varepsilon(n) - c/(1-\tau^\varepsilon)$. By the symmetric pricing condition,

$$m_n^\varepsilon = \frac{1 - F(p^\varepsilon(n))^n}{n\Big[F(p^\varepsilon(n))^{n-1} f(p^\varepsilon(n)) + \alpha^\varepsilon(n-1)\int_{p^\varepsilon(n)}^{\bar{v}} F(v)^{n-2} f(v)^2\, dv\Big]}.$$

Under the continuous interpolation in n, the right-hand side is continuous in (p, n). Hence, since the equilibrium price is unique for each n, the map $n \mapsto p^\varepsilon(n)$ is continuous.

Now define

$$\phi^\varepsilon(n) := \pi(p^\varepsilon(n), n; \tau^\varepsilon),$$

so $\phi^\varepsilon(n)$ is simply the per-firm operating profit under the fixed policy $(\alpha^\varepsilon, \tau^\varepsilon)$ when the entrant mass is n. Since $\alpha^\varepsilon > 0$ is fixed, the proof of Lemma 5 extends directly to the interpolated domain $n \in [2, \infty)$, so $m_n^\varepsilon \to 0$ as $n \to \infty$. Therefore $p^\varepsilon(n) \to \bar{v} - \varepsilon$ and $D(p^\varepsilon(n), n) \to 1$. It follows that

$$\phi^\varepsilon(n) = (1-\tau^\varepsilon) m_n^\varepsilon \frac{D(p^\varepsilon(n), n)}{n} \to 0,$$

while platform profit under the same fixed policy satisfies

$$\Pi^\varepsilon(n) := \Pi(\tau^\varepsilon, p^\varepsilon(n), n) = \tau^\varepsilon p^\varepsilon(n) D(p^\varepsilon(n), n) \to \tau^\varepsilon(\bar{v} - \varepsilon) = \bar{v} - \varepsilon - c.$$

Moreover, $\phi^\varepsilon(n) > 0$ for every finite $n \geq 2$, because $p^\varepsilon(n) = \bar{v}$ would make the pricing equation read $\varepsilon = 0$, while $p^\varepsilon(n) < \bar{v}$ makes both the numerator and denominator in the markup formula strictly positive. Since ϕ^ε is continuous, $\phi^\varepsilon(2) > 0$, and $\phi^\varepsilon(n) \to 0$, the intermediate value theorem implies that, for every sufficiently small $k > 0$, the set of $n \geq 2$ satisfying

$$\phi^\varepsilon(n) = k$$

is nonempty. Let $n^\varepsilon(k)$ denote the entrant mass chosen from this set by an arbitrary rule. Then $n^\varepsilon(k) \to \infty$ as $k \downarrow 0$. Otherwise, there would exist $N < \infty$ and a sequence $k_j \downarrow 0$ such that $n^\varepsilon(k_j) \in [2, N]$ for all j. But continuity and strict positivity of ϕ^ε on $[2, N]$ imply

$$\min_{n \in [2,N]} \phi^\varepsilon(n) > 0,$$

contradicting

$$\phi^{\varepsilon}(n^{\varepsilon}(k_j)) = k_j \to 0.$$

Hence, under every continuation rule, the platform profit induced by the fixed policy $(\alpha^{\varepsilon}, \tau^{\varepsilon})$ satisfies

$$\Pi^{\varepsilon}(n^{\varepsilon}(k)) \to \bar{v} - \varepsilon - c.$$

Now suppose, toward a contradiction, that there exist $k_j \downarrow 0$ and $N < \infty$ such that the corresponding symmetric free-entry equilibria $(\alpha_j, \tau_j, p_j, n_j) := (\alpha^*(k_j), \tau^*(k_j), p^*(k_j), n^*(k_j))$ satisfy $n_j \leq N$ for all j. By free entry, $\big[(1-\tau_j)p_j - c\big]D(p_j, n_j)/n_j = k_j$, so

$$\Pi(\tau_j, p_j, n_j) = \tau_j p_j D(p_j, n_j) = (p_j - c)D(p_j, n_j) - n_j k_j.$$

Since $n_j \leq N$, we have $D(p_j, n_j) \leq 1 - F(p_j)^N$, hence

$$\Pi(\tau_j, p_j, n_j) \leq (p_j - c)\big(1 - F(p_j)^N\big) + Nk_j \leq \overline{\Pi}_N + Nk_j,$$

where $\overline{\Pi}_N := \max_{p \in [c, \bar{v}]}(p - c)(1 - F(p)^N)$. Because $N < \infty$, $\overline{\Pi}_N < \bar{v} - c$: if $p < \bar{v}$, then

$$(p - c)\big(1 - F(p)^N\big) \leq p - c < \bar{v} - c,$$

while at $p = \bar{v}$ the expression equals 0.

Choose $\varepsilon > 0$ sufficiently small that $\bar{v} - \varepsilon > \underline{v}$ and $\bar{v} - \varepsilon - c > \overline{\Pi}_N$. Since $\Pi^{\varepsilon}(n^{\varepsilon}(k_j)) \to \bar{v} - \varepsilon - c$ and $Nk_j \to 0$, for all large j,

$$\Pi^*(k_j) \geq \Pi^{\varepsilon}(n^{\varepsilon}(k_j)) > \overline{\Pi}_N + Nk_j \geq \Pi(\tau_j, p_j, n_j),$$

contradicting optimality. Hence no such finite N can exist, and therefore $n^*(k) \to \infty$ as $k \downarrow 0$.

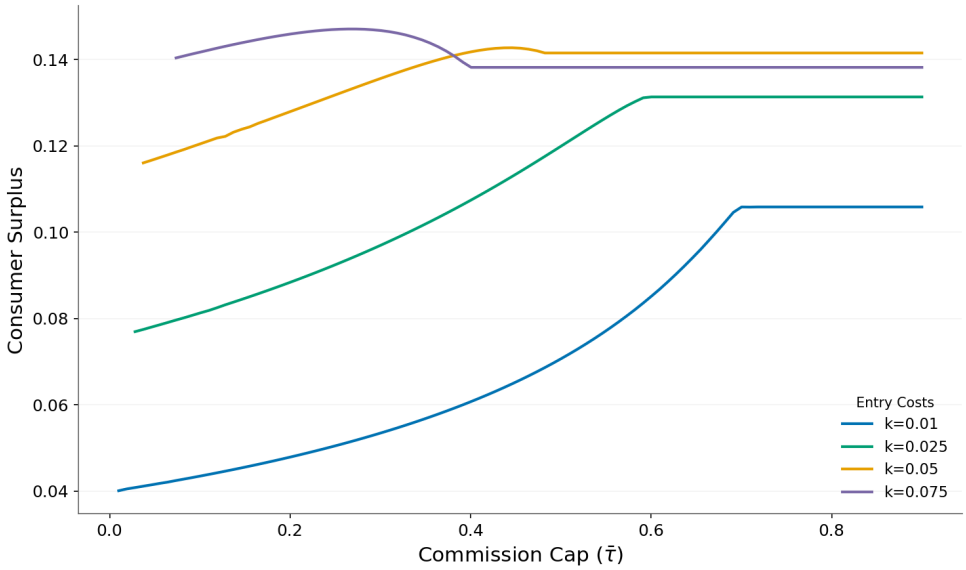
In particular, there exists $k_0 > 0$ such that $n^*(k) \geq 2$ for all $k \in (0, k_0]$. Now suppose $\alpha^*(k) = \infty$ for some such k. Then the induced pricing outcome is Bertrand. If $c/(1 - \tau^*(k)) < \bar{v}$, Lemma 3 implies $p^*(k) = c/(1 - \tau^*(k))$, so each firm's operating profit is zero; if instead $c/(1 - \tau^*(k)) \geq \bar{v}$, there is no trade, so each firm's operating profit is again zero. In either case $\pi(p^*(k), n^*(k); \tau^*(k)) = 0$, contradicting the free-entry condition $\pi(p^*(k), n^*(k); \tau^*(k)) = k > 0$. Hence $\alpha^*(k) < \infty$ for all $k \in (0, k_0]$. □

Proposition 14 formalizes how endogenous entry disciplines algorithm design. Holding (p, τ) fixed, per-firm operating profits fall with entry because demand is shared among more firms. When k is small, equilibrium entry becomes large, and indeed diverges as $k \downarrow 0$. But because entry requires positive rents, the platform cannot implement the extreme price-driven rule $\alpha = \infty$: with $n \geq 2$ this would eliminate operating profits and violate free entry for any $k > 0$.

Commission caps. With free entry, firms enter until expected operating profit equals the entry cost k. A cap on the commission rate τ shifts surplus from the platform to firms by limiting commission revenue per sale. Under endogenous entry, this rent shift can induce additional entry and thereby increase the equilibrium mass of active firms n.

This entry response matters because consumer surplus can be non-monotone in n in the baseline environment; see Subsection 3.4. In the simulation below, once the cap $\bar{\tau}$ binds and entry is sufficiently cheap, tightening $\bar{\tau}$ reduces consumer surplus even though, holding n fixed, it would improve consumer outcomes. Figure OA.4 illustrates this mechanism.

Figure OA.4: Consumer Surplus under Free Entry as a Function of $\overline{\tau}$



$F \sim U[0,1]$, $c = 0.3$, Lower $\bar{\tau}$ corresponds to a tighter cap.

OA.5 Positive Platform Marginal Cost

In this subsection, I reexamine the results on platform pricing and commission caps when the platform incurs a strictly positive marginal cost $c_p > 0$ per successful transaction and has an outside option of exiting with value zero.

Set-up. The firms' and consumers' problems are unchanged relative to the baseline. Hence, by Corollary 1, for any (α, τ), any symmetric equilibrium price in the pricing subgame satisfies

$$p - \frac{c}{1-\tau} = m(\alpha, p). \tag{5}$$

If the platform induces the symmetric price p, its profit is

$$(\tau p - c_p)\big(1 - F(p)^n\big).$$

Using (5), this can be rewritten as

$$\left[p - (c + c_p) - \frac{c\, m(\alpha, p)}{p - m(\alpha, p)}\right] \big(1 - F(p)^n\big).$$

Relative to equation (6), the cost term is only shifted by c_p. For finite α, this is only a convenient representation: the effective cost term depends on p through the markup function, so the platform still does not face a standard monopoly problem. A clean monopoly representation reappears in the two cases studied below: first, in the unregulated problem once $\alpha = \infty$ is shown to be optimal; and second, under a binding commission cap, when τ is fixed and the platform chooses among feasible induced prices.

Unconstrained commission.

Proposition 15 (Unregulated Platform Choice with Positive Platform Marginal Cost) *Suppose f is log-concave on $[\underline{v}, \bar{v}]$.*

If $c + c_p < \bar{v}$, then the game admits a unique symmetric equilibrium outcome. In that equilibrium,

$$\alpha^* = \infty, \qquad p^* = p_m(c + c_p), \qquad \tau^* = 1 - \frac{c}{p_m(c + c_p)}.$$

If $c + c_p \geq \bar{v}$, then there is no trade.

Proof Sketch. Existence of the pricing subgame follows from Proposition 1. The argument that $\alpha^* = \infty$ is the same double-deviation argument as in Proposition 2: any finite α leaves a strictly positive downstream markup, and the platform can profitably deviate to $\alpha = \infty$ while adjusting τ so as to keep the consumer price fixed. Quantity sold is unchanged, while revenue per sale strictly increases. Hence no operating optimum can have finite α.

Given $\alpha = \infty$, the unique symmetric equilibrium price in the pricing subgame is

$$p = \frac{c}{1 - \tau}.$$

The platform therefore chooses p to maximize

$$(p - c - c_p)\big(1 - F(p)^n\big).$$

By Definition 1 and the standard monopoly-price properties, if $c + c_p < \bar{v}$ this problem has the unique maximizer $p_m(c + c_p)$, which yields

$$\tau^* = 1 - \frac{c}{p_m(c + c_p)}.$$

If instead $c + c_p \geq \bar{v}$, there is no trade and the payoff is zero. □

Proposition 15 shows that in the unregulated benchmark, the platform still implements the outcome of a vertically integrated industry, except that the relevant marginal cost is now $c + c_p$.

Caps on commissions. Now suppose the regulator imposes a cap $\tau \leq \bar{\tau}$. Once τ is fixed, the platform uses α only to choose which symmetric equilibrium price is induced in the pricing subgame. Thus, under a binding cap, the relevant problem is to maximize platform profit over the set of prices that can be implemented at the capped commission.

Proposition 16 (Commission Caps with Positive Platform Marginal Cost)
Suppose $c + c_p < \bar{v}$, and for every $\tau \in [0, 1)$ with $(1-\tau)\bar{v} > c$ and every $\alpha \in [0, \infty]$, the pricing subgame admits a unique symmetric equilibrium price. Let

$$\tau^* := 1 - \frac{c}{p_m(c + c_p)}, \qquad \tau^{***} := \frac{c_p}{c + c_p},$$

let $\tau^{**} \in (\tau^{***}, \tau^*)$ be the unique solution to

$$p_m\left(\frac{c_p}{\tau^{**}}\right) = \frac{c}{1 - \tau^{**}},$$

and let $\tau^{****} \in (0, \tau^{***})$ be the unique solution to

$$\tau^{****} p_m\left(\frac{c}{1 - \tau^{****}}\right) = c_p.$$

Then

$$\tau^* > \tau^{**} > \tau^{***} > \tau^{****} > 0,$$

and under a cap $\tau \leq \bar{\tau}$,

$$(\alpha^*, \tau, p^*) = \begin{cases} (\infty, \tau^*, p_m(c + c_p)), & \bar{\tau} \geq \tau^*, \\ \left(\infty, \bar{\tau}, \dfrac{c}{1 - \bar{\tau}}\right), & \bar{\tau} \in [\tau^{**}, \tau^*), \\ \left(\alpha(\bar{\tau}), \bar{\tau}, p_m\left(\dfrac{c_p}{\bar{\tau}}\right)\right), \ \alpha(\bar{\tau}) \in (0, \infty), & \bar{\tau} \in (\tau^{***}, \tau^{**}), \\ \left(0, \bar{\tau}, p_m\left(\dfrac{c}{1 - \bar{\tau}}\right)\right), & \bar{\tau} \in (\tau^{****}, \tau^{***}], \\ \textit{either } \left(0, \bar{\tau}, p_m\left(\dfrac{c}{1 - \bar{\tau}}\right)\right) \textit{ or exit}, & \bar{\tau} = \tau^{****}, \\ \textit{exit (no trade)}, & \bar{\tau} < \tau^{****}. \end{cases}$$

In particular, the cap binds whenever $\bar{\tau} < \tau^*$.

Proof. Part (i) follows from Proposition 15, so fix $\bar{\tau} < \tau^*$. For any $\tau < 1$ with $(1-\tau)\bar{v} > c$,

let $p(\alpha,\tau)$ denote the unique symmetric equilibrium price in the pricing subgame. By Proposition 1, $p(\alpha,\tau)$ satisfies $p(\alpha,\tau)-\frac{c}{1-\tau}=m(\alpha,p(\alpha,\tau))$. Moreover, $p(0,\tau)=p_m\big(\frac{c}{1-\tau}\big)$ and $p(\infty,\tau)=\frac{c}{1-\tau}$. By continuity of the pricing equation, together with the tie-breaking convention at $\alpha=\infty$, the map $\alpha\mapsto p(\alpha,\tau)$ is continuous on $[0,\infty]$. Hence, for fixed τ, the set of implementable prices is exactly $\big[\frac{c}{1-\tau},\, p_m\big(\frac{c}{1-\tau}\big)\big]$.

Next, whenever $\bar\tau<\tau^*$, the cap binds. Suppose instead that the platform operates with some $\tau<\bar\tau$ and induces price p. If $p>\frac{c}{1-\tau}$, then for some $\tau'>\tau$ close enough to τ the same price remains feasible under τ', so the platform can choose some α' that still induces p. Quantity is unchanged, while per-sale revenue rises from $\tau p-c_p$ to $\tau' p-c_p$, a contradiction. If instead $p=\frac{c}{1-\tau}$, then necessarily $\alpha=\infty$, and profit is

$$(p-c-c_p)\big(1-F(p)^n\big).$$

By Definition 1 and the standard monopoly-price properties, this is strictly increasing for $p<p_m(c+c_p)$. Since $\bar\tau<\tau^*$, $\frac{c}{1-\bar\tau}<p_m(c+c_p)$, so deviating to $(\alpha,\tau)=(\infty,\bar\tau)$ raises the induced price from $\frac{c}{1-\tau}$ to $\frac{c}{1-\bar\tau}$ while keeping it below $p_m(c+c_p)$, and therefore strictly raises profit. Thus any optimum must satisfy $\tau=\bar\tau$.

Fix therefore $\tau=\bar\tau<\tau^*$. The platform chooses $p\in\big[\frac{c}{1-\bar\tau},\, p_m\big(\frac{c}{1-\bar\tau}\big)\big]$ to maximize $(\bar\tau p-c_p)(1-F(p)^n)$. Up to the positive factor $\bar\tau$, this is the monopoly problem with marginal cost $c_p/\bar\tau$, so the unconstrained maximizer is $\hat p(\bar\tau):=p_m\big(\frac{c_p}{\bar\tau}\big)$.

Since $p_m(\cdot)$ is strictly increasing, $\hat p(\tau)$ is strictly decreasing in τ, while $\frac{c}{1-\tau}$ is strictly increasing. At $\tau=\tau^*$,

$$\hat p(\tau^*)=p_m\Big(\frac{c_p}{\tau^*}\Big)<p_m(c+c_p)=\frac{c}{1-\tau^*},$$

because $\tau^*>\frac{c_p}{c+c_p}$. At $\tau^{***}:=\frac{c_p}{c+c_p}$,

$$\hat p(\tau^{***})=p_m(c+c_p)>c+c_p=\frac{c}{1-\tau^{***}}.$$

Hence there exists a unique $\tau^{**}\in(\tau^{***},\tau^*)$ such that

$$p_m\Big(\frac{c_p}{\tau^{**}}\Big)=\frac{c}{1-\tau^{**}}.$$

Therefore $\hat p(\bar\tau)\le\frac{c}{1-\bar\tau}$ iff $\bar\tau\in[\tau^{**},\tau^*)$, and in this region the optimum is the lower endpoint,

$$p^*(\bar\tau)=\frac{c}{1-\bar\tau},$$

implemented by $\alpha^*=\infty$. This proves part (ii).

Next, by strict monotonicity of p_m and $\hat{p}$,

$$\hat{p}(\bar{\tau}) \geq p_m\left(\frac{c}{1-\bar{\tau}}\right) \quad \Longleftrightarrow \quad \frac{c_p}{\bar{\tau}} \geq \frac{c}{1-\bar{\tau}} \quad \Longleftrightarrow \quad \bar{\tau} \leq \frac{c_p}{c+c_p} = \tau^{***}.$$

Hence, if $\bar{\tau} \in (\tau^{***}, \tau^{**})$, then

$$\frac{c}{1-\bar{\tau}} < \hat{p}(\bar{\tau}) < p_m\left(\frac{c}{1-\bar{\tau}}\right),$$

so the optimum is interior:

$$p^*(\bar{\tau}) = p_m\left(\frac{c_p}{\bar{\tau}}\right),$$

implemented by some $\alpha^*(\bar{\tau}) \in (0, \infty)$. This proves part (iii).

If instead $\bar{\tau} \leq \tau^{***}$, the optimum is the upper endpoint,

$$p^*(\bar{\tau}) = p_m\left(\frac{c}{1-\bar{\tau}}\right),$$

implemented by $\alpha^* = 0$, provided the platform operates. Operation requires

$$\bar{\tau}\, p_m\left(\frac{c}{1-\bar{\tau}}\right) - c_p > 0.$$

Now $\tau \mapsto \tau\, p_m\left(\frac{c}{1-\tau}\right)$ is continuous and strictly increasing, tends to 0 as $\tau \downarrow 0$, and at τ^{***} equals

$$\tau^{***} p_m(c + c_p) > \tau^{***}(c + c_p) = c_p.$$

Hence there exists a unique $\tau^{****} \in (0, \tau^{***})$ such that

$$\tau^{****} p_m\left(\frac{c}{1-\tau^{****}}\right) = c_p.$$

Therefore $\bar{\tau} \in (\tau^{****}, \tau^{***}]$ yields part (iv). If $\bar{\tau} < \tau^{****}$, then even the highest feasible price yields strictly negative profit, so the platform exits. If $\bar{\tau} = \tau^{****}$, then operating at the highest feasible price yields zero profit, so both exit and $\left(0, \bar{\tau}, p_m\left(\frac{c}{1-\bar{\tau}}\right)\right)$ are optimal. This proves parts (iv) and (v).

The threshold comparisons established above also imply the ordering

$$\tau^* > \tau^{**} > \tau^{***} > \tau^{****} > 0.$$

□

Relative to the benchmark without platform marginal cost, a binding cap on the commission now affects not only the division of surplus between firms and the platform, but also whether the platform can cover its own per-transaction cost c_p. As in Subsection 4.1,

once the cap binds the platform chooses α to implement the most profitable feasible price. The difference is that the relevant objective is now

$$(\bar{\tau}p - c_p)\big(1 - F(p)^n\big),$$

so the platform's desired price is $p_m(c_p/\bar{\tau})$ rather than the constant target $p_m(0)$ from the benchmark.

As the cap is tightened, this desired price rises, while the feasible price interval shifts down. When the cap first binds, the desired price lies below the feasible set, so the platform chooses the lowest feasible price. As the cap is tightened further, the desired price becomes feasible and is implemented by an interior recommendation rule, implying that tighter caps raise equilibrium prices in this region. For tighter caps still, the desired price exceeds the feasible set, so the platform chooses the highest feasible price. If the cap becomes too tight, the platform exits altogether.

Thus, unlike in the benchmark, tighter commission caps need not lower prices or raise consumer surplus monotonically. Prices first fall below the unregulated benchmark level $p_m(c + c_p)$, then rise as the cap tightens further, return exactly to $p_m(c + c_p)$ at $\bar{\tau} = \tau^{***} = c_p/(c + c_p)$, and then fall again before trade ceases altogether. Consumer surplus is therefore non-monotone as well: a tighter cap can make consumers worse off than a looser one, even though every capped equilibrium price remains weakly below the unregulated benchmark price.

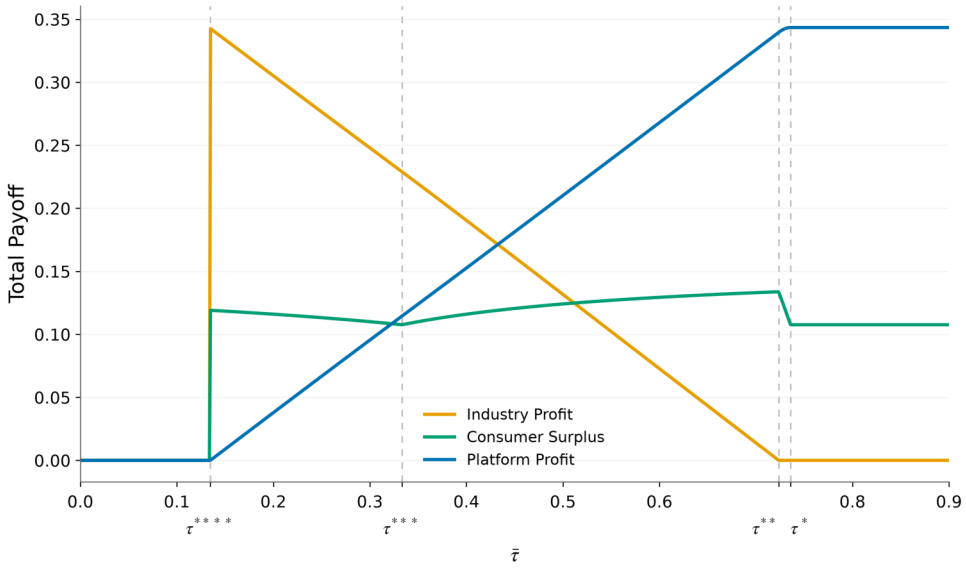
Figure OA.5 illustrates these comparative statics in the uniform example. When the cap first binds, platform profits fall and consumer surplus rises. In the interior-price region, tighter caps raise the equilibrium price and reduce consumer surplus. For tighter caps still, the platform switches to the highest feasible price; eventually the cap becomes so tight that the platform exits and all surplus falls to zero.

OA.6 Wolinsky benchmark without recommendation design

Consider the standard Wolinsky search model with a central platform that sets only a commission $\tau \in [0, 1)$. The platform does not rank products and does not otherwise steer demand. This benchmark isolates fee setting from the recommendation mechanism in the main text.

Set-up. There are $n \geq 2$ firms, each producing one product at marginal cost $c > 0$. For a representative consumer, match values v_i are i.i.d. according to F on $[\underline{v}, \bar{v}]$, with density f, where f is continuously differentiable on $[\underline{v}, \bar{v}]$. The timing is as follows. First, the platform chooses τ. Second, firms set prices. Third, consumers sequentially search

Figure OA.5: Commission caps: Payoffs with positive marginal costs of platforms



$F \sim U[0,1]$, $n = 5$, $c = 0.2$, $c_p = 0.1$. Lower $\bar{\tau}$ corresponds to a tighter cap.

across products and learn prices and match values. Each search costs s. Attention is restricted to symmetric equilibria. Consumers maximize utility, firms maximize profits, and the platform maximizes its commission revenue.

Consumer behavior. Assume that $s > 0$ is such that the equation below has a solution $\hat{x} \in (\underline{v}, \bar{v})$ with $f(\hat{x}) > 0$, and define

$$\int_{\hat{x}}^{\bar{v}} (v - \hat{x})\, dF(v) = s. \tag{OA.25}$$

Following Wolinsky (1986), if a consumer has already inspected product i, then she continues to search if and only if

$$\hat{x} - p^* > v_i - p_i,$$

where p^* denotes the expected equilibrium price. Moreover, a consumer starts searching if and only if

$$\hat{x} \geq p^*.$$

Pricing subgame for a given commission. Fix τ. Firm i solves

$$\max_{p_i} \big((1-\tau)p_i - c\big) D_i \iff \max_{p_i} \left(p_i - \frac{c}{1-\tau}\right) D_i.$$

Hence, a commission τ induces a standard Wolinsky pricing game with effective marginal cost

$$\tilde{c}(\tau) := \frac{c}{1-\tau}$$

instead of c. Wolinsky (1986) shows that any symmetric equilibrium price p^* must satisfy

$$p^* - \frac{c}{1-\tau} = \mu_n(p^*; \hat{x}), \tag{OA.26}$$

where

$$\mu_n(p; \hat{x}) := \frac{1 - F(p)^n}{\frac{f(\hat{x})\big(1-F(\hat{x})^n\big)}{1-F(\hat{x})} - n\int_p^{\hat{x}} F(v)^{n-1} f'(v)\, dv}. \tag{OA.27}$$

Throughout this subsection, we assume, for simplicity, that the FOC (OA.26) is necessary and sufficient for a symmetric pricing equilibrium and that the platform-preferred equilibrium is implemented.

Platform: commission choice. Suppose the platform wants to implement a symmetric price $p \le \hat{x}$. Then (OA.26) requires

$$\tau(p) = 1 - \frac{c}{p - \mu_n(p; \hat{x})}. \tag{OA.28}$$

Conditional on consumer participation, the platform therefore solves

$$\max_{p \le \hat{x}:\, p - \mu_n(p;\hat{x}) \ge c} \left(1 - \frac{c}{p - \mu_n(p; \hat{x})}\right) p\big[1 - F(p)^n\big]. \tag{OA.29}$$

Thus, the platform chooses the commission indirectly by choosing which symmetric price $p \le \hat{x}$ to implement.

The next result shows that, for sufficiently large n, the platform pushes the induced price to the participation boundary.

Proposition 17 (Comparative Statics in n)
Suppose

$$c < \hat{x} - \frac{1 - F(\hat{x})}{f(\hat{x})}.$$

Then there exists $\bar{N}$ such that, for all $n \ge \bar{N}$,

$$p^* = \hat{x} \qquad \text{and} \qquad CS_n = 0.$$

Proof. Since $\sup_{p \le \hat{x}} F(p)^n = F(\hat{x})^n \to 0$, we have $1 - F(p)^n \to 1$ uniformly in $p \le \hat{x}$.

Moreover, since $F(v) \leq F(\hat{x}) < 1$ for all $v \in [\underline{v}, \hat{x}]$,

$$\sup_{p \leq \hat{x}} n \int_p^{\hat{x}} F(v)^{n-1} |f'(v)| \, dv \leq nF(\hat{x})^{n-1} \int_{\underline{v}}^{\hat{x}} |f'(v)| \, dv \to 0.$$

Hence, uniformly in $p \leq \hat{x}$,

$$\mu_n(p; \hat{x}) \to \frac{1 - F(\hat{x})}{f(\hat{x})}.$$

Also,

$$\sup_{p \leq \hat{x}} nF(p)^{n-1} |f(p)| \to 0 \qquad \text{and} \qquad \sup_{p \leq \hat{x}} nF(p)^{n-1} |f'(p)| \to 0.$$

Together with the preceding convergence, these bounds imply that $\partial_p \mu_n(p; \hat{x}) \to 0$ uniformly. Hence, differentiating the objective in (OA.29) shows that its derivative converges uniformly on the feasible set to

$$1 + \frac{c \dfrac{1 - F(\hat{x})}{f(\hat{x})}}{\left(p - \dfrac{1 - F(\hat{x})}{f(\hat{x})} \right)^2},$$

which is strictly positive. Moreover, $p - \mu_n(p; \hat{x})$ is strictly increasing in p for all sufficiently large n (recall that $\partial_p \mu_n(p; \hat{x}) \to 0$), so the feasible set is an interval. Therefore, for all sufficiently large n, the objective in (OA.29) is strictly increasing on the feasible set.

The maintained condition

$$c < \hat{x} - \frac{1 - F(\hat{x})}{f(\hat{x})}$$

implies that $p = \hat{x}$ is feasible for all sufficiently large n, since $\mu_n(\hat{x}; \hat{x}) \to (1 - F(\hat{x}))/f(\hat{x})$. Hence, for all n large enough, the platform chooses the largest feasible participating price, namely $p^* = \hat{x}$.

At $p^* = \hat{x}$, expected consumer surplus from entering is zero. Hence $CS_n = 0$. □

The result shows that an increase in the number of firms can raise prices and reduce consumer surplus, even when the platform only sets fees and does not manipulate recommendations. For large n, the purchase probability becomes close to one at any participating price. Once the boundary price $p = \hat{x}$ is implementable, the platform therefore prefers the highest price consistent with consumer participation and drives consumer surplus to zero. Thus, the forces in Proposition 4 are not driven solely by recommendation design. They also reflect the platform's role as an intermediary: as more firms participate, access to the platform becomes more valuable, which allows the platform to extract a larger share of consumer surplus.