

연세대학교 상경대학

경제연구소

Economic Research Institute

Yonsei University



서울시 서대문구 연세로 50

50 Yonsei-ro, Seodaemun-gu, Seoul, Korea

TEL: (+82-2) 2123-4065 FAX: (+82-2) 364-9149

E-mail: yeri4065@yonsei.ac.kr

<http://yeri.yonsei.ac.kr/new>

Regulating Platform Bias: Self-preferencing and Transaction Fees as Strategic Substitutes

Chang-Koo Chi

Yonsei University

Jay Pil Choi

Michigan State University

Jong-Hee Hahn

Yonsei University

Seongkyun Kim

Yonsei University

June 2026

2026RWP-294

Economic Research Institute

Yonsei University

REGULATING PLATFORM BIAS: SELF-PREFERENCING AND TRANSACTION FEES AS STRATEGIC SUBSTITUTES

Chang-Koo Chi¹, Jay Pil Choi², Jong-Hee Hahn³, and Seongkyun Kim^{*4}

^{1,2,3,4}School of Economics, Yonsei University

²Department of Economics, Michigan State University

June 15, 2026

Abstract

This paper studies self-preferencing incentives by vertically integrated platforms that operate both marketplaces and affiliated retail businesses. We show that self-preferencing and transaction fees are substitute instruments for profit extraction, implying that restrictions on self-preferencing may induce offsetting increases in transaction fees and thereby generate unintended consequences for consumer welfare. We characterize the platform's optimal choice of self-preferencing and transaction fees and evaluate the welfare effects of behavioral and structural remedies. We also extend the analysis to settings with platform competition and consumer search, examining how market forces shape self-preferencing incentives and evaluating the robustness of our main results.

JEL CLASSIFICATION: L2, L5, D2, D8

KEYWORDS: self-preferencing, vertically integrated platforms, transaction fees, regulation, hierarchical Hotelling model, search

1. Introduction

Vertically integrated platforms increasingly serve as both marketplace intermediaries and direct competitors to the third-party sellers that rely on their platforms. This dual role has raised concerns that platforms may engage in self-preferencing by steering consumers toward their own

^{*}We are grateful to the seminar participants at Korea University, Michigan State University, Seoul National University, Southern Methodist University, University of Seoul, the Annual Winter Meeting of the Korean Economic Association, the Asia-Pacific IO Conference, and the Midwest Theory Meeting for their valuable feedback and suggestions.
Email addresses: ck.changkoo@yonsei.ac.kr (Chi); choijay@msu.edu (Choi); hahnjh@yonsei.ac.kr (Hahn); ksk7757@yonsei.ac.kr (Kim)

products. This paper studies the incentives for self-preferencing in a model where a platform can influence consumer choices through algorithmic recommendations while simultaneously charging transaction fees to third-party sellers. We show that self-preferencing and transaction fees act as substitute instruments for profit extraction. Consequently, regulatory interventions aimed at curbing self-preferencing may induce platforms to adjust alternative instruments under their control, such as transaction fees, thereby partially offsetting the intended effects of regulation and generating unintended consequences. We characterize the platform’s optimal choice of self-preferencing and transaction fees, evaluate the welfare effects of behavioral and structural remedies, and extend the analysis to settings with platform competition and consumer search.

Self-preferencing has become a central concern in competition policy and platform regulation across many jurisdictions, including the United States, Europe, and Asia.¹ Recent high-profile cases illustrate these concerns. The European Union found that Google unlawfully favored its own comparison-shopping service in search results, while the European Commission has investigated Amazon’s Buy Box practices and their treatment under the DMA. In South Korea, the major e-commerce platform Naver faced fines and remedial orders for allegedly favoring its own SmartStore merchants, although the Korean Supreme Court overturned the sanctions on October 16, 2025, citing insufficient evidence of anticompetitive effects or exclusionary intent despite the firm’s dominant position. Kakao Mobility, a leading mobility platform, similarly faced sanctions for allegedly favoring its affiliated taxi services.² With the DMA now in force and similar policy debates emerging worldwide, understanding the welfare consequences of self-preferencing and the effectiveness of alternative regulatory approaches has become increasingly important.

To provide a framework for evaluating regulatory interventions targeting self-preferencing, we first analyze a baseline model of a monopoly platform that intermediates transactions between consumers and two sellers. The sellers are horizontally differentiated in the spirit of Hotelling: one is the platform’s vertically integrated subsidiary, and the other is an independent third-party seller. The platform possesses superior information about consumers’ locations on the Hotelling line, as inferred from their observed shopping histories and underlying product preferences. Consumers, by contrast, do not know their locations and therefore rely on the platform’s recommendation to make purchase. The platform presents each consumer with only one of the two products, and consumers purchase the recommended option. Although consumers are aware of the availability of alternative products, they do not purchase the non-recommended product, for instance, because doing so would require incurring search or transaction costs. Later, we relax this assumption by allowing consumers to search for the alternative product.

Following [de Cornière and Taylor \(2014\)](#), we model self-preferencing by allowing the platform’s recommendation algorithm to distort the allocation rule, shifting the effective indifference

¹The EU’s Digital Markets Act (DMA) defines self-preferencing as giving preferential treatment to a platform’s own products or services in search results, rankings, or other platform features relative to those of third parties.

²Another important antitrust concern involving hybrid platforms is the use of proprietary data from third-party sellers to develop competing private-label products, which may weaken innovation incentives. See [Hagiu, Teh and Wright \(2022\)](#), [Madsen and Vellodi \(2025\)](#), and [Choi, Kim and Mukherjee \(2025\)](#) for analyses of this issue and related regulatory proposals.

point along the Hotelling line in favor of its subsidiary's product. The magnitude of this shift captures the degree of self-preferencing.

In our model, we assume that once consumers join the platform, the market is fully covered in the sense that each participating consumer purchases one of the two products, as in the standard Hotelling model, even though platform participation itself is elastic. This combination of full market coverage and elastic participation generates a quasi-linear structure in both user surplus and the platform's per-user profit with respect to transaction fees. This quasi-linear structure is the key feature of the model, as it yields a particularly transparent characterization of equilibrium and its welfare implications.

In particular, the quasi-linear structure implies that transaction fees and self-preferencing play fundamentally different roles in the platform's profit-maximization problem. The transaction fee is competitively neutral: under full market coverage, it is fully passed through into retail prices and therefore operates as a one-for-one transfer from consumers to the platform. Self-preferencing, by contrast, has nonlinear effects on both user surplus and platform profit. Starting from an unbiased recommendation algorithm, a marginal increase in self-preferencing raises platform profit while having only a second-order effect on user surplus. Consequently, at low levels of bias, self-preferencing is a more effective instrument for surplus extraction than the transaction fee. As the degree of self-preferencing increases, however, the marginal profit gain relative to the marginal utility loss declines monotonically. The platform therefore increases self-preferencing up to the point at which this marginal benefit-to-cost ratio equals unity—the constant ratio associated with the transaction fee. This condition uniquely determines the optimal level of self-preferencing. Beyond this threshold, transaction fees become the more effective instrument for extracting surplus and are primarily used to govern consumer participation.

More importantly, the quasi-linear structure also implies that the platform's choice of self-preferencing maximizes the joint surplus of the platform and its users at the expense of the third-party seller. The resulting joint surplus is then divided between consumers and the platform through the transaction fee. This characterization has immediate implications for regulation. Because self-preferencing and transaction fees are substitute instruments, any restriction on self-preferencing induces the platform to rely more heavily on transaction fees. As a result, interventions that reduce self-preferencing lower joint surplus and can reduce consumer welfare even though they improve allocative efficiency. We examine two widely discussed antitrust interventions: a behavioral remedy that bans self-preferencing and a structural remedy in the form of vertical separation. We show that a complete ban on self-preferencing can reduce consumer welfare because the platform responds by increasing transaction fees sufficiently to offset the gains from a less distorted recommendation algorithm. We further show that structural separation can be even more harmful to consumers. By eliminating the platform's downstream retail profits, separation makes transaction fees the platform's sole source of revenue, leading to even higher fees, lower consumer surplus, and reduced platform participation. Thus, a seemingly stricter structural remedy may generate worse consumer outcomes than a behavioral ban.

Having analyzed the baseline monopoly setting, in which consumers follow the platform's

recommendation, we next examine the robustness of our main insights under platform competition and consumer search. In Section 3, we introduce competition between two horizontally differentiated platforms within a hierarchical Hotelling framework. The equilibrium outcome depends on the degree of platform differentiation relative to product differentiation. When platforms are relatively homogeneous, implying intense platform competition, the equilibrium level of self-preferencing is lower than under monopoly and transaction fees are driven down to their lower bound. By contrast, when platforms are sufficiently differentiated, the equilibrium degree of self-preferencing coincides with the monopoly outcome, and competition affects only transaction fees. Thus, platform competition disciplines self-preferencing incentives only when competition is sufficiently intense. In this environment, a ban on self-preferencing can increase consumer welfare because competitive pressure limits the platforms’ ability to offset the ban through higher transaction fees. Otherwise, the ban reduces consumer welfare, as in the monopoly benchmark. This finding challenges the conventional wisdom that platform competition alone is sufficient to discipline self-preferencing and render regulation unnecessary. Paradoxically, the pro-consumer effect of a ban on self-preferencing arises precisely when platform competition is intense rather than weak. By contrast, structural separation unambiguously reduces consumer welfare. As in the monopoly setting, eliminating the platform’s downstream profit source induces greater reliance on transaction fees, resulting in higher prices and lower consumer surplus.

In Section 4, we incorporate consumer search into the model. This extension is of independent interest, as it develops a framework for search in the presence of algorithmic recommendations, where the platform’s optimal level of steering and consumers’ search incentives are jointly determined. The possibility of search disciplines the platform’s incentives for both self-preferencing and pricing, as consumer search leads to a discontinuous reduction in platform profits. In this environment, we characterize the equilibrium in which the platform sets self-preferencing at a level that deters search. This mechanism offers a rational explanation for the widespread tendency of consumers to purchase the product highlighted by the platform—for instance, Amazon’s Buy Box—without actively considering alternatives. Rather than reflecting behavioral biases such as naivete, inattention, or inertia, such outcomes may arise endogenously from platform incentives and rational consumer search decisions.³

In Section 5, we discuss the generality of our approach and several extensions of the baseline model. We adopt the Hotelling framework primarily for analytical tractability and closed-form characterization. However, our approach and main insights are considerably more general. In particular, we show that our results extend to a broader discrete-choice framework that accommodates multiple third-party sellers and imposes no restrictions on the functional forms governing users’ matching utility or the platform’s retail profits. The Hotelling model analyzed in the preceding sections is a tractable special case of this more general framework. We also discuss two

³Amazon’s Buy Box features a default seller, and according to Juul (2021), more than 80 percent of Amazon sales occur through this channel. Raval (2022) and Chen and Tsai (2024) document that Amazon’s algorithms systematically favor its private-label products in determining Buy Box placements and “Frequently Bought Together” recommendations, respectively.

additional extensions. First, we examine how the analysis can be adapted to evaluate the effects of self-preferencing regulation on total surplus, inclusive of third-party seller profits. Second, we analyze the implications of cost asymmetries between the platform’s affiliated seller and the independent seller.

Our paper contributes to the recent literature on self-preferencing by vertically integrated platforms. [de Cornière and Taylor \(2014\)](#) develop a search-bias model in which a search engine allocates user traffic across publishers. Their setting features ad-financed firms, no transaction fees, and consumers who bear nuisance costs from advertising; in that model, the level of ad load serves as the relevant price instrument. By contrast, our analysis focuses on a marketplace platform’s joint choice of transaction fees and bias. We show that these two instruments are strategic substitutes, which is central to assessing the effects of regulation. Because our model is designed for a marketplace, whereas theirs studies an ad-financed search engine, the two papers address different dimensions of self-preferencing and are best viewed as complementary.⁴

[Inderst and Ottaviani \(2012\)](#) analyze a market in which consumers rely on intermediaries for advice to select suitable specialized offerings. They highlight the role of commissions in making the advisor responsive to supply-side incentives, while also showing that regulatory interventions such as mandatory disclosure or commission caps may have unintended welfare effects by inefficiently reducing the responsiveness of advice to supply-side differences. Similarly, [Teh and Wright \(2022\)](#) examine steering by an informed intermediary facing many firms, showing that consumers benefit from intermediation only when search costs are sufficiently high and that equilibrium prices increase with the number of competing firms. In contrast, we consider a marketplace in which the intermediary endogenously sets its fee structure and analyze how vertical integration shapes the extent of recommendation bias.

Our model of self-preferencing in the presence of consumer search relates to [Armstrong, Vickers and Zhou \(2009\)](#), who study the implications of “prominence” for price equilibrium under a logit demand specification. In their sequential search model, the prominent firm’s product is sampled first by all consumers; however, their analysis is not set in a platform environment, since prominence is exogenous and transaction fees are absent.⁵ While [Kittaka and Sato \(2022\)](#) and [Kittaka and Sato \(2024\)](#) study the welfare effects of self-preferencing by hybrid platforms that make their own products more prominent, their framework assumes an *exogenously* given transaction fee. They therefore do not examine the interaction between self-preferencing and fee setting. By contrast, our analysis focuses on the joint determination of self-preferencing and transaction fees, and on the strategic interaction between these two instruments.

[Zenny \(2022\)](#) extend the [Armstrong et al. \(2009\)](#) framework to model a hybrid platform that

⁴Meanwhile, [de Cornière and Taylor \(2019\)](#) show that, for platforms that earn profits only from sales of their own products and do not rely on advertising or fee revenue, the welfare effects of search bias depend on whether the induced changes in sellers’ decisions increase or decrease consumer utility.

⁵[Armstrong and Zhou \(2011\)](#) develop a model in which firms compete to become “prominent,” thereby influencing the order of consumer consideration. See also [Varian \(2007\)](#), [Athey and Ellison \(2011\)](#), and [Chen and He \(2011\)](#) for auction-based models of competition for favorable search positions. These studies, however, focus on competition among third parties for prominence and do not address self-preferencing by a vertically integrated platform that controls the ranking mechanism.

exercises self-preferencing by systematically including its own product in search results, while third-party products enter randomly. This increases the probability of first-party sales, raises expected profit per consumer, and induces the platform to lower transaction fees to attract participation; the resulting fee reduction lowers consumer prices and thus increases consumer surplus and social welfare. Because logit demand makes products horizontally symmetric, self-preferencing entails no taste-mismatch cost in their models, so welfare depends only on the transaction fee. By contrast, our framework endogenizes both the degree of self-preferencing and the transaction fee, while allowing for taste mismatch and potential cost disadvantages of the platform’s product. Our analysis of self-preferencing with consumer search (Section 4) captures the joint determination of platform bias and consumer search behavior. Since the platform’s ranking is based on relative match quality, it acts as a signal that updates consumers’ beliefs about the value of further search, yielding an equilibrium that is much richer than in settings with exogenous prominence.

[Hervas-Drane and Shelegia \(2025\)](#) analyze a “retailer-led marketplace” in which a monopolistic retailer relies on third-party sellers to relax its capacity constraints. Although they also show that transaction fees and steering are strategic substitutes, steering in their model affects only “inattentive” consumers, whose prevalence is exogenously fixed. By contrast, our framework makes the platform’s bias endogenous and centers on the joint determination of self-preferencing and transaction fees.

[Anderson and Bedre-Defolie \(2024\)](#) study a hybrid platform as a dominant firm competing against a fringe of sellers under monopolistic competition with logit demand. In their model, the platform engages in “insidious steering” by raising transaction fees to favor its own products, rather than using explicit self-preferencing.⁶ By contrast, our framework treats transaction fees and self-preferencing as separate policy instruments.

The remainder of the paper is organized as follows. Section 2 introduces the baseline model of a monopolistic hybrid platform and characterizes the optimal transaction fee and degree of self-preferencing. It also evaluates the welfare implications of standard regulatory interventions, including bans on self-preferencing and structural separation. Section 3 extends the analysis to a duopoly setting using a hierarchical Hotelling framework to examine whether platform competition disciplines self-preferencing incentives. Section 4 incorporates consumer search and obedience constraints into the platform’s recommendation problem, thereby analyzing steering through recommendation algorithms and testing the robustness of the baseline results. Section 5 discusses the generality of our approach and several extensions, including a more general discrete-choice framework, the effects of self-preferencing regulation on total surplus, and the implications of cost asymmetries between the platform’s affiliated seller and independent sellers. The paper concludes in Section 6. All technical proofs are relegated to the Appendix.

⁶See also [Etro \(2023\)](#) for an analysis of pricing by hybrid platforms in a setting with free entry of sellers.

2. Monopoly Platform

This section presents our baseline model, where a monopoly platform intermediates transactions between consumers and sellers and attempts to steer consumers toward a particular product. The framework builds on [de Cornière and Taylor \(2014\)](#). After characterizing the market equilibrium, we analyze the welfare implications of alternative regulatory policies designed to address the underlying market distortion. The analysis in this section serves as a benchmark for the subsequent investigation of how platform competition and consumer search incentives affect the degree of self-preferencing in Sections 3 and 4, respectively.

2.1. Model and Equilibrium

The model consists of two sellers, indexed by $i = 1, 2$, offering horizontally differentiated products on a platform. The sellers are located at opposite endpoints of a Hotelling line of unit length. Seller 1, located at zero, operates as a subsidiary of the platform, whereas seller 2, located at one, functions as an independent third-party seller. For expositional simplicity, we assume that the platform incurs no operational overhead and that both sellers share an identical unit production cost c . We explore the implications of downstream cost asymmetry in Section 5.⁷ The sellers compete in prices by simultaneously setting $\mathbf{p} = (p_1, p_2)$ for their respective products.

A unit mass of consumers is uniformly distributed along the Hotelling line, with each consumer having a unit demand for the product. A consumer located at $x \in [0, 1]$ obtains utility $v - tx - p_1$ from purchasing product 1 and $v - t(1 - x) - p_2$ from purchasing product 2, where v denotes the intrinsic value of the products and t is the transportation cost parameter. Under perfect information about product characteristics and consumers' own locations, consumers choose the product yielding higher utility. In this standard Hotelling framework, the demand for product 1 is determined by the location of the indifferent consumer, given by $\hat{x}(\mathbf{p}) \equiv (t - p_1 + p_2)/2t$.

In contrast, our model assumes that consumers have imperfect information about the match between their preferences and product characteristics—specifically, they cannot identify which product is better suited to their tastes. We capture this informational friction by assuming that consumers do not observe their locations on the Hotelling line. This lack of information induces consumers to rely on the platform's recommendations when making purchasing decisions. The platform, by contrast, is better informed: using data on past behavior and product characteristics, it can infer consumers' locations and provide personalized recommendations. In our baseline model, consumers purchase the product recommended by the platform.⁸ We relax this assumption in Section 4 by introducing a search model where consumers can acquire information about their preferences and alternative products by incurring search costs.⁹

⁷A comprehensive analysis of the asymmetric cost case is provided in the Online Appendix.

⁸In this two-product setting, the non-recommended product can be interpreted as the best alternative among many third-party options that consumers cannot easily identify. The assumption that consumers do not purchase the non-recommended product reflects sufficiently high search or transaction costs.

⁹The baseline model thus corresponds to a case of prohibitively high search costs.

Within this framework, we consider a simple recommendation algorithm. The platform recommends product 1 to a consumer located at x if and only if

$$v - tx - p_1 + b \geq v - t(1 - x) - p_2, \quad (1)$$

where $b \geq 0$ measures the degree of algorithmic bias, or self-preferencing. In practice, recommendation algorithms operate as scoring mechanisms that evaluate and rank products based on user characteristics and product attributes. In our framework, the platform assigns products scores reflecting expected consumer-product match quality. The parameter b can therefore be interpreted as a “bonus” score awarded exclusively to the platform’s subsidiary, thereby biasing recommendations against the third-party seller.

Under the recommendation rule in (1), the mass of consumers directed to the platform’s subsidiary is given by

$$\bar{x}_b(\mathbf{p}) \equiv \frac{t + b - p_1 + p_2}{2t}, \quad (2)$$

which constitutes the demand for product 1. Note that $\bar{x}_b$ mirrors the standard Hotelling demand $\hat{x}$, but augmented by the bias parameter b . The effect of self-preferencing is therefore to shift the location of the indifferent consumer toward the independent seller, allowing the platform’s subsidiary to capture a larger market share for any given price pair.

The game unfolds in three stages. First, the platform chooses the degree of self-preferencing b and sets a transaction fee f levied on the third-party seller for each completed transaction.¹⁰ Upon observing the platform’s choice (b, f) , consumers decide whether to join the platform. Non-participating consumers receive an outside option z , which is drawn from a log-concave distribution G with a continuous density g , independently of the consumer’s location x .¹¹ Finally, the two sellers simultaneously set the prices of their products. We assume that the intrinsic value v is sufficiently large so that all participating users purchase one of the two products, implying full market coverage within the platform.

We characterize the subgame perfect equilibrium using backward induction. In the final stage, seller 1—vertically integrated with the platform—sets the price p_1 to maximize its *per-user* profit, π_1 . This objective function accounts for both the retail margin from the subsidiary and the fee revenue collected from the third-party seller:

$$\pi_1 = \bar{x}_b(\mathbf{p})(p_1 - c) + [1 - \bar{x}_b(\mathbf{p})]f = f + \bar{x}_b(\mathbf{p})(p_1 - c - f).$$

In this formulation, $c + f$ represents the effective marginal cost of the subsidiary’s product. The transaction fee f therefore acts as an opportunity cost: each unit sold by the subsidiary displaces a

¹⁰Under full market coverage, the assumption that the fee is paid solely by the seller entails no loss of generality, as any fee burden is fully passed on to consumers, as in [Edelman and Wright \(2015\)](#).

¹¹To distinguish between different stages of the model, we refer to agents as *consumers* when considering the participation decision and as *users* once they have joined the platform.

sale by the independent seller and thus forgoes the fee revenue that the platform would otherwise collect. Simultaneously, seller 2 sets p_2 to maximize her own per-user profit, $\pi_2 = [1 - \bar{x}_b(\mathbf{p})] (p_2 - c - f)$.

Solving the system of first-order conditions yields the equilibrium prices:

$$p_1^* = t + c + f + \frac{b}{3} \quad \text{and} \quad p_2^* = t + c + f - \frac{b}{3}. \quad (3)$$

The transaction fee f induces a parallel shift in retail prices; under full market coverage, it is fully passed through to consumers without affecting the price differential between the two sellers. In contrast, the sellers respond in opposite directions to the bias parameter b . As b increases, the platform leverages its steering capability to raise p_1 , while the third-party seller is compelled to lower p_2 in order to defend market share. Consequently, the equilibrium price differential, $p_1^* - p_2^*$, is increasing in the degree of self-preferencing. Substituting the equilibrium prices in (3) into (2), the equilibrium demand for the platform's subsidiary is given by $\bar{x}_b^* = 1/2 + b/6t$.

To analyze consumers' participation decisions in the second stage, we now derive the expected surplus from joining the platform. Since users rely on the platform's recommendation, a user is directed to product 1 with probability $\bar{x}_b^*$ and to product 2 with probability $1 - \bar{x}_b^*$. Expected user surplus is therefore given by

$$U = \int_0^{\bar{x}_b^*} (v - tx - p_1^*) dx + \int_{\bar{x}_b^*}^1 (v - t(1-x) - p_2^*) dx = \phi(b) - f,$$

where $\phi(b) = v - c - 5t/4 - 5b^2/36t$ captures the gross utility generated by the recommendation algorithm, net of expected transportation costs and mismatch distortions induced by algorithmic bias. A consumer joins the platform if and only if the reservation utility z from the outside option does not exceed the surplus offered by the platform, i.e., $z \leq U$. Observe that $\phi(b)$ is monotonically decreasing in b . Self-preferencing distorts the allocation rule away from the match-efficient benchmark $\hat{x}$, thereby unambiguously reducing user surplus.¹²

In the first stage, the platform anticipates equilibrium pricing and participation decisions when choosing the degree of self-preferencing b and transaction fee f . These strategic choices affects expected user surplus U , which in turn determines the platform's user base. Given $z \sim G$, the mass of participating consumers is $G(U)$. The platform's total expected profit is therefore

$$\Pi \equiv G(U)\pi_1,$$

where per-user profit π_1 can be expressed as

$$\pi_1 = f + \bar{x}_b^* [p_1^* - c - f] = f + \sigma(b),$$

¹²When the platform subsidiary operates at a cost disadvantage, however, $\phi(b)$ can become hump-shaped. In that case, a moderate degree of self-preferencing in favor of the less efficient subsidiary may increase user welfare by intensifying price competition and improving expected match outcomes. See the Online Appendix for details.

with $\sigma(b) = (3t + b)^2/18t$. The per-user profit of a hybrid platform thus consists of two distinct revenue sources: transaction-fee revenue f collected from the third-party seller and the retail profit $\sigma(b)$ earned by the platform's subsidiary. In contrast, a pure marketplace lacks the latter source and relies solely on fee revenue. This distinction plays a central role in our subsequent analysis of structural separation as a regulatory remedy for self-preferencing incentives.

Observe that the per-user profit π_1 is increasing in $b \geq 0$. Thus, absent user participation constraints, the platform would increase b until $\bar{x}_b^* = 1$, fully foreclosing the third-party seller. To understand this incentive, note from (3) that the platform earns a margin of $t + f + b/3$ from a sale by its subsidiary, whereas it earns only f from a sale by the third-party seller. Even when $b = 0$, the platform gains an additional margin of t by steering a transaction toward its own subsidiary. Increasing b not only expands the subsidiary's market share but also widens the margin differential between the two types of sales. Consequently, the platform has a strong incentive to engage in self-preferencing as aggressively as possible, which would ultimately lead to complete foreclosure of the third-party seller unless constrained by consumer participation decisions.

Taking consumer participation into account, the platform's problem can be written as

$$\max_{b, f} \Pi(b, f) = G(\phi(b) - f)[f + \sigma(b)].$$

A noteworthy feature of this problem is the quasi-linear structure of both user surplus U and per-user profit π_1 with respect to the transaction fee f . This property arises because, under full market coverage, the fee is competitively neutral and fully passed through into retail prices, as shown in (3). Consumers therefore bear the entire burden of the fee, while the platform captures the full amount regardless of whether the transaction occurs through its subsidiary or the independent seller. In this sense, the transaction fee operates as a pure transfer from consumers to the platform. As formalized in the following lemma, this neutrality property yields a particularly clean characterization of the equilibrium: the platform's optimal choice of algorithmic bias is determined independently of both the pricing instrument f and the outside option distribution G .

Although the matching utility $\phi(b)$ and retail profit $\sigma(b)$ admit explicit closed-form expressions in our stylized framework, the equilibrium characterization does not rely on these specific functional forms. Since most of our main results extend to alternative specifications provided that the quasi-linear structure is preserved, we present our subsequent analysis using this generalized formulation.

LEMMA 1. *The platform's optimal level of self-preferencing, b^* , is characterized by the maximization of the joint per-user surplus:*

$$b^* = \operatorname{argmax}_{b \geq 0} S(b) \equiv U + \pi_1 = \phi(b) + \sigma(b).$$

PROOF OF LEMMA 1: Let (b^*, f^*) denote the platform's optimal strategy that maximizes total expected profits Π . Suppose, to the contrary, that there exists some $b^\circ \geq 0$ such that $S(b^\circ) > S(b^*)$.

Consider an alternative transaction fee f° satisfying

$$\phi(b^\circ) - f^\circ = \phi(b^*) - f^* \iff U(b^\circ, f^\circ) = U(b^*, f^*). \quad (4)$$

By construction, expected user surplus—and hence the mass of participating consumers—is identical under (b^*, f^*) and (b°, f°) . However, the platform’s per-user profit under the alternative strategy is

$$\pi_1(b^\circ, f^\circ) = S(b^\circ) - U(b^*, f^*) > \pi_1(b^*, f^*),$$

where the equality follows from (4), and the inequality follows from the assumption that $S(b^\circ) > S(b^*)$. Since the mass of participating consumers is unchanged while per-user profit is strictly higher, the platform would earn strictly greater total profit under (b°, f°) , contradicting the optimality of (b^*, f^*) . $\square$

Lemma 1 shows that the platform chooses b to maximize the joint surplus of itself and its users. This allows the platform to extract surplus from the third-party seller, which is forced to lower its price to defend market share against algorithmic bias. In this respect, the role of b is reminiscent of [Aghion and Bolton \(1987\)](#), where exclusive dealing is used to extract surplus from a potential entrant. Lemma 1 is central to our regulatory analysis because interventions that restrict b —such as a ban on self-preferencing—limit the platform’s ability to maximize this joint surplus. As we show formally in the following subsection, the resulting reduction in the total surplus available to the platform and its users ultimately leaves users worse off once the platform optimally adjusts its alternative instrument, f , to preserve profitability.

PROPOSITION 1 (Monopoly Platform). *The platform’s optimal degree of self-preferencing is $b^* = 2t$, and the optimal transaction fee f^* is uniquely determined by*

$$\frac{G}{g}(\phi(b^*) - f) = f + \sigma(b^*), \quad \text{where } b^* = 2t. \quad (5)$$

Furthermore, the aggregate profit function Π is log-submodular in (b, f) , implying that the two instruments b and f act as strategic substitutes.

PROOF OF PROPOSITION 1: See Appendix A.1. $\square$

The equilibrium is characterized by the *relative* effectiveness of self-preferencing and transaction fees as instruments for increasing platform profitability. Both instruments raise the platform’s per-user profit π_1 at the expense of user surplus U , thereby reducing platform participation. However, their marginal effects differ fundamentally. Because U and π_1 are quasi-linear in the transaction fee f , an increase in f generates a one-for-one transfer from users to the platform, implying a constant marginal benefit-to-cost ratio of unity. By contrast, the effects of self-preferencing b are nonlinear. In particular, U is decreasing and concave in b , with slope zero at $b = 0$, whereas π_1

is increasing and convex in b , with a strictly positive slope at $b = 0$. Thus, starting from $b = 0$, an increase in self-preferencing raises platform profit while having only a second-order effect on user surplus. Consequently, at low levels of bias, self-preferencing is a more effective instrument for surplus extraction than the transaction fee.

As b increases, however, the ratio $(d\pi_1/db)/(-dU/db)$, which measures the marginal profit gain relative to the marginal utility loss, declines monotonically. The platform therefore increases b up to the point at which this marginal benefit-to-cost ratio equals unity—the constant marginal benefit-to-cost ratio associated with the transaction fee. This occurs at $b = 2t$. Beyond this threshold, the transaction fee becomes the more effective instrument for extracting surplus.¹³

Condition (5) for the optimal transaction fee can also be expressed in Lerner index form. Define

$$\eta(f|b) \equiv f \frac{g}{G} (\phi(b) - f)$$

as the price elasticity of consumer participation for a given level of self-preferencing b . Then, the optimal transaction fee satisfies

$$\frac{f - (-\sigma(b))}{f} = \frac{1}{\eta(f|b)},$$

where $-\sigma(b)$ represents the platform's opportunity cost of losing a user to the outside option.

2.2. Regulation

Building on the preceding results, we now evaluate two policy remedies aimed at curbing self-preferencing by vertically integrated platforms: a behavioral remedy banning self-preferencing and a structural remedy separating the affiliated seller from the platform. We focus on their effects on user welfare U , leaving the analysis of total surplus—including platform and third-party seller profits—to Section 5.

The quasi-linear structure of U and π_1 implies a monotonic relationship between joint per-user surplus S and user surplus U . As a result, any policy that raises or lowers total surplus has a corresponding directional effect on consumer welfare.

PROPOSITION 2. *Let $f^*(b)$ denote the platform's profit-maximizing transaction fee for a given level of self-preferencing b . Under optimal fee adjustment, user surplus $U(b) = \phi(b) - f^*(b)$ is increasing in the joint per-user surplus $S(b) = \phi(b) + \sigma(b)$:*

$$\frac{dU}{dS} > 0.$$

PROOF OF PROPOSITION 2: Consider the platform's choice of f for a given level of self-preferencing b , which determines the joint per-user surplus $S(b) = \phi(b) + \sigma(b)$. The platform's

¹³Note that the optimal level of self-preferencing, $b^* = 2t$, is increasing in the degree of product differentiation t . This reflects the fact that the platform's incentive to steer users toward its subsidiary arises from capturing margins that would otherwise accrue to the rival seller. When product differentiation is weak, equilibrium margins are small, reducing the gain from steering. In the limiting case where the rival earns zero equilibrium margins, the platform has no incentive to engage in self-preferencing.

problem can then be written as

$$\max_f \Pi(f; b) = G(\phi(b) - f)[f + \sigma(b)].$$

Using the change of variables $U = \phi(b) - f$, the problem can equivalently be expressed as the platform choosing the level of user surplus:

$$\max_U \Pi(U; S) = G(U)(S - U).$$

The first-order condition is given by:¹⁴

$$\frac{G}{g}(U) = S - U. \quad (6)$$

Totally differentiating (6) yields

$$\frac{dU}{dS} = \left[1 + \left(\frac{G}{g}(U) \right)' \right]^{-1}.$$

This expression is positive because the inverse hazard rate G/g is increasing in U under the log-concavity of G .¹⁵ $\square$

REMARK 1. Equation (6) also sheds light on how the joint surplus $S(b)$ is divided between the platform and consumers through the platform's choice of f . To see this, define the elasticity of consumer participation with respect to user surplus U as $\varepsilon \equiv Ug(U)/G(U)$. Rearranging (6), it follows that consumers receive a share $\varepsilon/(1 + \varepsilon)$ of the joint surplus, while the platform captures the remaining share $1/(1 + \varepsilon)$. Moreover, if G exhibits constant elasticity ε , then $dU/dS = \varepsilon/(1 + \varepsilon)$, implying that consumers receive a fixed share of any increase in joint surplus.

Proposition 2 provides a useful shortcut for evaluating regulatory interventions: consumer welfare can be assessed by examining their effect on joint surplus. In the absence of regulation, the platform chooses the level of self-preferencing to maximize joint surplus, as established in Lemma 1. Consequently, as long as the platform retains the ability to adjust its transaction fee, any intervention that distorts the optimal level of self-preferencing necessarily reduces consumer welfare. We examine this implication below for behavioral and structural remedies, respectively.

¹⁴Since the objective function $\Pi(U; S)$ is the product of two log-concave functions, it is itself log-concave in U . Hence, the first-order condition is both necessary and sufficient for a global optimum.

¹⁵Alternatively, from $\Pi(U; S) = G(U)(S - U)$, we have $\partial^2 \Pi / \partial S \partial U = g(U) > 0$. Thus, Π exhibits increasing differences in (U, S) , which yields the desired property.

2.2.1. Ban on Self-Preferencing

A ban on self-preferencing, such as that imposed under the European Union’s Digital Markets Act (DMA), prohibits gatekeeper platforms from favoring their own products or services over those of third-party providers. In our framework, such a policy corresponds to setting $b = 0$, so that $\bar{x}_0(\mathbf{p}) = \hat{x}(\mathbf{p}) = (t - p_1 + p_2)/2t$, thereby restoring an unbiased recommendation algorithm. The following result is a direct corollary of Proposition 2.

COROLLARY 1. *A ban on self-preferencing reduces user surplus.*

PROOF OF COROLLARY 1: Under a ban on self-preferencing, the joint per-user surplus is $S(0) = \phi(0) + \sigma(0)$, which is lower than $S(b^*)$ by Lemma 1. Proposition 2 then implies $\phi(b^*) - f^* > \phi(0) - f_{\text{ban}}$, where f_{ban} denotes the platform’s optimal transaction fee under the ban. $\square$

The intuition is that self-preferencing and transaction fees are instruments of strategic substitutes. Although self-preferencing distorts match quality, prohibiting it induces the platform to raise transaction fees in order to preserve profitability. The resulting increase in fees more than offsets the allocative gains from eliminating bias, generating a waterbed effect that ultimately reduces user surplus.

2.2.2. Structural Separation

Behavioral remedies require ongoing monitoring because algorithmic fairness or neutrality is difficult to verify in practice. Structural separation avoids this problem by eliminating the platform’s incentive to bias recommendations altogether. However, it also entails the same unintended welfare consequences as a behavioral ban.

COROLLARY 2. *Structural separation reduces user surplus.*

PROOF OF COROLLARY 2: Under structural separation, the platform no longer owns the affiliated seller and therefore earns no subsidiary profit, implying $\sigma = 0$. The platform consequently chooses b to maximize user surplus $\phi(b)$. Since

$$\max_{b \geq 0} \{\phi(b) + \sigma(b)\} > \max_{b \geq 0} \phi(b),$$

structural separation reduces joint per-user surplus. Proposition 2 then implies that user surplus is also lower under structural separation. $\square$

Under structural separation, the platform sets $b = 0$, thereby maximizing ϕ . However, this gain from eliminating algorithmic bias is more than offset by the platform’s optimal increase in the transaction fee.

2.2.3. Ban on Self-Preferencing versus Structural Separation

Our framework also allows a simple comparison between the two policy options, assuming that one must be implemented despite both being inferior to laissez-faire from the standpoint of user welfare, perhaps for other policy objectives such as protecting third-party sellers.

COROLLARY 3. *A ban on self-preferencing is strictly preferred to structural separation from the users' perspective.*

PROOF OF COROLLARY 3: Under both policy regimes, the equilibrium level of self-preferencing is $b = 0$. However, the joint per-user surplus under a behavioral ban is $S_{\text{ban}} = \phi(0) + \sigma(0)$, whereas under structural separation it is $S_{\text{sep}} = \phi(0)$. Since $\sigma(0) > 0$, it follows that $S_{\text{ban}} > S_{\text{sep}}$. Proposition 2 then implies that user surplus is strictly higher under a ban on self-preferencing than under structural separation. $\square$

The intuition is straightforward: although both regimes eliminate algorithmic bias, structural separation additionally removes the subsidiary's profit margin. As a result, the platform responds by raising the transaction fee more aggressively in order to preserve profitability. This finding is consistent with Hagiu *et al.* (2022), who show that structural separation harms consumer surplus more than targeted behavioral remedies because it eliminates internal product competition. By contrast, our framework highlights a different mechanism: eliminating vertical profit margins intensifies the platform's incentive to extract surplus through transaction fees, ultimately harming consumers.

REMARK 2. *The equilibrium transaction fee may be negative, as is often the case in two-sided platform markets; see Armstrong (2006) and Rochet and Tirole (2006). In our framework, this can arise when product differentiation t is sufficiently large, so that the retail margin earned by the platform's subsidiary is high. Even if a non-negativity constraint were imposed on transaction fees, Proposition 2 would continue to hold, and the qualitative welfare implications of the monopoly model would remain unchanged. However, as we show in the next section, a non-negative fee constraint can fundamentally alter the welfare implications of self-preferencing under platform competition.*

3. Platform Competition

In this section, we extend the baseline model to a setting with two competing hybrid platforms. Each platform can steer users toward its affiliated product while charging transaction fees to third-party sellers. To assess whether platform competition disciplines self-preferencing incentives, we develop a *hierarchical* Hotelling model with a *nested* structure and characterize the resulting market equilibrium. We also analyze the implications of self-preferencing regulation for user welfare, paralleling the monopoly analysis.




Figure 1: The Hierarchical Hotelling Model with Platform and Product Layers

3.1. Model and Equilibrium

The competition model consists of two hierarchical layers (see Figure 1). In the upper layer, two symmetric platforms, indexed by $i = L, R$, are located at the endpoints of a unit Hotelling line, with consumers uniformly distributed along the line. Each consumer joins one platform and purchases the product recommended there.¹⁶ A consumer located at $z \in [0, 1]$ obtains utility $U_L - \tau z$ from joining platform L and utility $U_R - \tau(1 - z)$ from joining platform R , where $\tau > 0$ measures the degree of platform differentiation and U_i denotes the expected utility generated by platform i . The consumer indifferent between the two platforms is located at

$$\hat{z} \equiv \frac{1}{2} + \frac{1}{2\tau}(U_L - U_R),$$

which determines platform L 's market share.

After choosing a platform, consumers enter the lower layer, where the platform recommends either its affiliated product (Product 1) or the product of an independent third-party seller (Product 2). As in the monopoly setting, users rely on the platform's recommendation algorithm because they are imperfectly informed about their product-match preferences. Given prices $\mathbf{p}_i = (p_{i1}, p_{i2})$, platform i 's recommendation rule is characterized by $\bar{x}_i = (t + b_i - p_{i1} + p_{i2})/2t$, where b_i measures the degree of self-preferencing by platform i . We assume that a consumer's location z on the upper layer is independent of her location x on the lower layer, and that the intrinsic value v is identical across products and platforms.

The timing of the game is the same as in the monopoly model. First, the two platforms simultaneously choose their levels of self-preferencing b_i and transaction fees f_i . Given these choices, consumers decide whether to join platform L or R . Finally, the two sellers on each platform simultaneously set their product prices. The final two stages parallel those in the monopoly setting, except that the consumer's outside option is replaced by competitive pressure from the rival platform, which is determined endogenously in equilibrium.

Consequently, both user surplus U_i and per-user profit π_i for platform $i \in \{L, R\}$ retain the

¹⁶Consumers do not multi-home, and the market is fully covered.

same functional forms as in Section 2:

$$U_i = \phi(b_i) - f_i \quad \text{and} \quad \pi_i = f_i + \sigma(b_i).$$

The joint per-user surplus generated within platform i is therefore

$$S(b_i) = \phi(b_i) + \sigma(b_i).$$

Using this, platform i 's total profit can be written as

$$\Pi_i = \hat{z}(U_i, U_j) \pi_i = \hat{z}(U_i, U_j) [S(b_i) - U_i], \quad (7)$$

where $\hat{z}(U_i, U_j)$ denotes platform i 's market share in the upper-layer Hotelling competition. A key implication of (7) is that, for any given (U_i, U_j) , platform i maximizes profit by choosing b_i to maximize the joint per-user surplus $S(b_i)$. Hence, the logic of Lemma 1 carries over directly to the competitive setting: each platform chooses its level of self-preferencing to maximize the total surplus available for extraction. Accordingly, absent any restriction on transaction fees—such as a non-negativity constraint—both platforms choose $b_i = b^* = 2t$.

Let $S^* = S(b^*)$ denote the maximized joint per-user surplus. With self-preferencing fixed at $b^* = 2t$, the first stage reduces to a surplus-offer game in which each platform i chooses its transaction fee—or equivalently, the level of user surplus U_i —to maximize total profit $\hat{z}(U_i, U_j) [S^* - U_i]$. Since this objective function is concave in U_i , the unique Nash equilibrium is characterized by the first-order condition:

$$\frac{\partial \Pi_i}{\partial U_i} = \frac{1}{2\tau} [S^* - U_i] - \hat{z}(U_i, U_j) = 0, \quad i = L, R. \quad (8)$$

The equilibrium is symmetric, implying $U_L = U_R = U^*$ and $\hat{z} = 1/2$. Substituting these conditions into (8) yields the equilibrium transaction fee:

$$f^* = \underbrace{-\sigma(b^*)}_{\text{opportunity cost}} + \underbrace{\tau}_{\text{market power}}. \quad (9)$$

Expression (9) admits a standard Hotelling interpretation. Because the transaction fee is fully passed through to consumers, it effectively serves as the price of platform participation. Meanwhile, $\sigma(b^*)$ captures the platform's opportunity cost of losing a user and therefore lowers its effective marginal cost of attracting consumers. The transportation cost parameter τ then determines the equilibrium price-cost margin, reflecting the degree of market power arising from platform differentiation.

As in the monopoly case, each platform's objective function Π_i exhibits increasing differences in (U_i, S^*) . This implies that the equilibrium level of user surplus U^* is increasing in the joint per-user surplus S^* . Hence, any regulatory intervention that reduces the joint per-user surplus necessarily lowers consumer welfare. This finding establishes that, absent any constraint on trans-

action fees, the unintended welfare consequences of self-preferencing regulation identified in the monopoly setting remain robust under platform competition.

3.2. Non-Negative Transaction Fee Constraint

The preceding analysis allows for negative transaction fees, which may arise naturally in two-sided platform markets. In practice, however, such subsidies may be infeasible and can generate moral hazard or adverse selection concerns. In this subsection, we impose the non-negativity constraint $f_i \geq 0$ and show that, when platform competition is sufficiently intense, this constraint breaks the parallel between the monopoly and competition cases.¹⁷

As shown in (9), the non-negativity constraint binds when platform competition is sufficiently intense, namely when $\tau \leq \sigma(b^*) = 25t/18$. In this region, the platforms can no longer offset downstream retail profits through negative transaction fees, and the logic underlying Lemma 1 no longer applies. By contrast, when platform differentiation is sufficiently high so that $\tau > \sigma(b^*)$, the constraint is slack and the analysis from the previous subsection continues to hold.

When $\tau \leq \sigma(b^*)$, the platforms are constrained to set $f_i = 0$. The next result characterizes the equilibrium level of self-preferencing in this constrained environment.

PROPOSITION 3 (Platform Competition with Non-Negative Transaction Fees). *The hierarchical Hotelling model admits a unique symmetric equilibrium characterized as follows.*

- If $\tau > \sigma(b^*)$, the platforms charge a positive transaction fee $f^* = \tau - \sigma(b^*)$. In this case, the equilibrium level of self-preferencing is $b^* = 2t$, which maximizes the joint per-user surplus $S(b)$ and satisfies

$$\phi'(b^*) + \sigma'(b^*) = 0.$$

- If $\tau \leq \sigma(b^*)$, the platforms are constrained to set $f^+ = 0$, so that no transaction fee is charged to the third-party seller. The equilibrium level of self-preferencing is then $b^+ (\leq b^*)$, uniquely characterized by

$$\sigma(b^+)\phi'(b^+) + \tau\sigma'(b^+) = 0.$$

PROOF OF PROPOSITION 3: See Appendix A.2. $\square$

With a non-negative transaction fee constraint, the model yields two qualitatively distinct equilibria. When platform competition is relatively mild, i.e., $\tau > \sigma(b^*)$, the equilibrium level of self-preferencing coincides with that in the monopoly benchmark. In this regime, self-preferencing is used to maximize joint per-user surplus, while the transaction fee serves as the primary competitive instrument for attracting users and expanding at the platform level. As competition in-

¹⁷For related analyses of non-negative price constraints in other contexts, see [Choi and Jeon \(2023\)](#) on platform design and [Bisceglia and Tirole \(2024\)](#) on access-fee regulation for dominant gatekeepers.

tensifies and τ declines, platforms respond by lowering transaction fees while maintaining self-preferencing at the surplus-maximizing level b^* .

Once competition becomes sufficiently intense so that $\tau \leq \sigma(b^*)$, the transaction fee reaches its lower bound of zero and the equilibrium changes qualitatively. Platforms then compete by moderating self-preferencing below b^* , with the equilibrium level b^+ uniquely characterized by

$$\mathfrak{J}(b) \equiv \sigma(b) \left[\phi'(b) + \tau \frac{\sigma'}{\sigma}(b) \right] = 0. \quad (10)$$

Uniqueness follows because ϕ' is decreasing and σ is log-concave, implying that the defined function $\mathfrak{J}$ changes sign only once from positive to negative. Since $S'(b^*) = \phi'(b^*) + \sigma'(b^*) = 0$, the condition $\tau < \sigma(b^*)$ leads to $\mathfrak{J}(b^*) < 0$, which in turn leads to $b^+ < b^*$.

Moreover, standard comparative statics imply that b^+ is increasing in τ , so the equilibrium level of self-preferencing declines as platform competition intensifies.¹⁸ In the limiting case of perfect platform competition ($\tau \rightarrow 0$), the equilibrium condition (10) reduces to $\sigma(b)\phi'(b) = 0$, which implies $b^+ = 0$. Thus, under sufficiently intense competition, platforms abandon self-preferencing altogether in order to maximize consumer surplus and attract consumers.

3.3. Self-Preferencing Regulation with Non-Negative Transaction Fee Constraint

Under intense platform competition, the regulatory implications differ markedly from those in the monopoly setting. When the non-negativity constraint on transaction fees binds, the cost of regulation may be borne by the platform rather than passed on to users through higher fees. As a result, regulation can potentially increase user surplus. To examine these effects, we consider both a ban on self-preferencing and structural separation.

From Proposition 3, equilibrium user surplus under platform competition in the absence of regulation is given by

$$U^* = \begin{cases} \phi(b^*) - [\tau - \sigma(b^*)] & \text{if } \tau > \sigma(b^*) \\ \phi(b^+) & \text{if } \tau \leq \sigma(b^*), \end{cases} \quad (11)$$

where $b^+ \leq b^*$ is uniquely determined by condition (10). Since b^+ is increasing in τ and ϕ is decreasing in b , it follows that U^* is decreasing in τ . Hence, user surplus rises as platform competition intensifies.

¹⁸By the implicit function theorem applied to (10),

$$\frac{db^+(\tau)}{d\tau} = -\frac{\mathfrak{J}_\tau(b^+; \tau)}{\mathfrak{J}_b(b^+; \tau)} = -\frac{\sigma'}{\sigma}(b^+) \left[\phi''(b^+) + \tau \frac{d}{db} \left(\frac{\sigma'}{\sigma}(b) \right) \Big|_{b=b^+} \right]^{-1}.$$

Since $\phi(b)$ is strictly concave and σ is log-concave, the bracketed expression is negative, implying that $\frac{db^+(\tau)}{d\tau} > 0$.

3.3.1. Ban on Self-Preferencing

Under a ban on self-preferencing ($b_L = b_R = 0$), platform i 's profit in (7) becomes $\Pi_i = \hat{z}(U_i, U_j)[S(0) - U_i]$. Proceeding as in the derivation of the unregulated equilibrium, the equilibrium transaction fee under the ban and the nonnegative fee constraint is given by

$$f_{\text{ban}} = \max\{0, \tau - \sigma(0)\}.$$

Accordingly, equilibrium user surplus is

$$U_{\text{ban}} = \begin{cases} \phi(0) - [\tau - \sigma(0)] & \text{if } \tau > \sigma(0) \\ \phi(0) & \text{if } \tau \leq \sigma(0). \end{cases} \quad (12)$$

Comparing U_{ban} in (12) with the unregulated user surplus U^* in (11) yields the following result.

PROPOSITION 4. *The effect of a ban on self-preferencing on user surplus depends on the intensity of platform competition. There exists a threshold $\hat{\tau} \in (\sigma(0), \sigma(b^*))$ such that a ban increases user surplus if $\tau < \hat{\tau}$ and decreases it otherwise.*

PROOF OF PROPOSITION 4: See Appendix A.3. $\square$

In contrast to the monopoly case, where a ban on self-preferencing unambiguously reduces user surplus, platform competition creates a parameter region in which such a ban can improve user welfare. The key difference is that competition limits the platforms' ability to offset the ban through higher transaction fees. As in the monopoly setting, transaction fees and self-preferencing remain substitute instruments, so platforms respond to the ban by attempting to raise fees. However, when competition is sufficiently intense ($\tau < \hat{\tau}$), this adjustment is constrained by competitive pressure. In that case, the reduction in algorithmic bias dominates, leading to higher user surplus. By contrast, when competition is weak ($\tau > \hat{\tau}$), platforms retain sufficient pricing power to offset the welfare gains from eliminating bias, and the ban reduces user surplus as in the monopoly benchmark.

These results underscore that the consumer-welfare effects of banning self-preferencing depend critically on market conditions, particularly the intensity of platform competition. Contrary to the view that competition alone disciplines self-preferencing and therefore renders regulation unnecessary, our analysis shows that a ban may improve user surplus precisely when competition is intense rather than weak.

3.3.2. Structural Separation

Under structural separation, $\sigma(b_i) = 0$ for all b_i , so platform i 's profit becomes

$$\Pi_i = \hat{z}(U_i, U_j)f_i = \hat{z}(U_i, U_j)[\phi(b_i) - U_i].$$

Hence, for any given (U_i, U_j) , the platform optimally sets $b_i = 0$. The equilibrium transaction fee is then equal to τ , yielding the standard Hotelling outcome in which platforms compete in utility space solely through transaction fees. User surplus under structural separation is therefore

$$U_{\text{sep}} = \phi(0) - \tau.$$

Comparing U_{sep} with the unregulated user surplus U^* in (11) yields the following result.

PROPOSITION 5. *Structural separation reduces user surplus under platform competition for all $\tau > 0$.*

PROOF OF PROPOSITION 5 : See Appendix A.4. $\square$

Structural separation eliminates the platform's incentive to favor its affiliated seller, thereby making the recommendation algorithm consumer-optimal. However, this gain is outweighed by a countervailing price effect: once the platform no longer earns secondary margins from vertical integration, it has a stronger incentive to extract surplus through transaction fees. The resulting increase in fees dominates the benefit from eliminating recommendation bias, leading to lower user surplus.

3.3.3. Ban on Self-Preferencing versus Structural Separation

The preceding analysis demonstrates that behavioral and structural remedies have distinct implications for consumer welfare. The comparison is summarized in the following proposition.

PROPOSITION 6. *A ban on self-preferencing is strictly preferred to structural separation from the consumer's perspective for all $\tau > 0$.*

PROOF OF PROPOSITION 6: Under a behavioral ban, equilibrium user surplus is $U_{\text{ban}} = \phi(0) - \max\{0, \tau - \sigma(0)\}$. Under structural separation, equilibrium user surplus is $U_{\text{sep}} = \phi(0) - \tau$ for every $\tau > 0$. Therefore,

$$U_{\text{ban}} - U_{\text{sep}} = \min\{\tau, \sigma(0)\} > 0,$$

implying that a ban on self-preferencing strictly dominates structural separation in terms of consumer welfare. $\square$

Both remedies eliminate algorithmic distortion. However, unlike structural separation, a behavioral ban preserves the platform's vertically integrated business model. By retaining the subsidiary margin $\sigma(0) > 0$, the platform continues to internalize gains from downstream sales, which weakens its incentive to raise transaction fees aggressively. Consequently, user surplus is always higher under a ban on self-preferencing than under structural separation.

4. Self-Preferencing and Consumer Search

So far, we have assumed passive users who purchase the platform’s recommended product without further consideration. We now relax this assumption by incorporating *consumer search* into the monopoly model in Section 2. In particular, we introduce an additional stage in which users may inspect the alternative product before making a purchase decision. The analysis in this section is of independent theoretical interest, as it develops a model of consumer search under biased recommendations. In addition, it illustrates how consumer search can discipline and shape a platform’s algorithmic bias, with search behavior and self-preferencing jointly and endogenously determined in equilibrium.

Suppose the platform recommends its affiliated product to consumers with $x \leq \bar{x}_b(\mathbf{p})$, as defined in equation (2). After observing the recommendation and the recommended product’s price, a consumer may incur a search cost $s > 0$ to inspect the alternative product.¹⁹ Search enables the consumer to make a fully informed purchase decision by revealing both the consumer’s true relative preference between the two products and the prices of both products. With costless recall, the post-search purchase decision therefore follows the standard Hotelling allocation rule: a consumer with $x \leq \hat{x}(\mathbf{p})$ purchases product 1, where $\hat{x}(\mathbf{p}) \equiv (t - p_1 + p_2)/2t$ denotes the indifferent consumer.

We solve for the subgame-perfect equilibrium by backward induction and analyze how consumer search affects the welfare implications of regulation. Consider first the user’s search decision. Because $\hat{x} < \bar{x}_b$ under a biased recommendation algorithm, consumers who are recommended product 2 would continue to purchase product 2 even after search and therefore never have an incentive to search. By contrast, consumers who are recommended product 1 may search if the expected benefit from discovering the alternative product exceeds the search cost. More precisely, search is optimal only when the observed price of the subsidiary product, p_1 , exceeds a threshold level $\bar{p}_1$. As in standard search models such as [Armstrong et al. \(2009\)](#), this threshold depends on consumers’ beliefs about the third-party price p_2 . Hence, the independent seller affects search behavior only indirectly through equilibrium beliefs.

Next, we analyze how consumer search affects product pricing. Search introduces a discontinuity in the platform’s per-user profit π_1 , which is the platform’s objective function conditional on consumer participation. When $p_1 \leq \bar{p}_1$, consumers do not search, and product 1’s demand is given by $\bar{x}_b$.²⁰ When $p_1 > \bar{p}_1$, consumers recommended product 1 engage in search, so demand

¹⁹The search cost is exogenously given. Although assuming a homogeneous search cost is restrictive, it yields a tractable model with closed-form solutions.

²⁰When $p_1 = \bar{p}_1$, consumers are indifferent between searching and not searching. We assume that indifferent consumers do not search and instead follow the platform’s recommendation. This tie-breaking rule ensures upper semi-continuity of π_1 , thereby facilitating the analysis.




Figure 2: Consumer Search and Pricing

falls to the standard unbiased Hotelling demand $\hat{x}$. Accordingly, the per-user profit function,

$$\pi_1(\mathbf{p}) = \begin{cases} \pi_1^{ns}(\mathbf{p}) \equiv f + \bar{x}_b(\mathbf{p})(p_1 - c - f) & \text{for } p_1 \leq \bar{p}_1 \\ \pi_1^s(\mathbf{p}) \equiv f + \hat{x}(\mathbf{p})(p_1 - c - f) & \text{for } p_1 > \bar{p}_1, \end{cases}$$

exhibits a discrete downward jump at $\bar{p}_1$, as illustrated in Figure 2. Due to this discontinuity, the platform's optimal pricing strategy depends on the third-party price p_2 , with two key threshold values, $\underline{p}_2$ and $\bar{p}_2$.

When p_2 is sufficiently high, namely $p_2 \geq \bar{p}_2$, consumer search is not binding. The platform therefore sets its standard no-search best response,

$$p_1^{ns}(p_2) \equiv \operatorname{argmax}_{p_1} \pi_1^{ns}(p_1, p_2) = \frac{t + b + c + f + p_2}{2},$$

as in the baseline model.²¹ Because consumers expect a sufficiently high third-party price ($p_2 > \bar{p}_2$), the expected benefit from search is lower than the search cost, so no consumer searches in equilibrium. The platform can therefore exploit self-preferencing without distorting its price.

When p_2 lies in the intermediate range $[\underline{p}_2, \bar{p}_2]$, consumer search becomes binding, and the platform sets the subsidiary price at the threshold $\bar{p}_1$. At this price, consumers recommended product 1 are indifferent between searching and not searching. The price $\bar{p}_1$ therefore acts as a limit price that deters consumer search, so the outcome remains a no-search equilibrium. Intuitively, the platform responds to the threat of search by lowering p_1 just enough to make search unattractive, thereby dissipating part of the rent generated by self-preferencing.

Finally, when p_2 is sufficiently low, i.e., $p_2 < \underline{p}_2$, deterring search becomes too costly because the required price reduction outweighs the associated profit gain. The platform then accommo-

²¹As p_2 falls toward $\bar{p}_2$, the search threshold $\bar{p}_1$ converges to p_1^{ns} (see Figure 2-(a)). Hence $\bar{p}_2$ is uniquely determined by the condition $p_1^{ns} = \bar{p}_1$.




Figure 3: Optimal Pricing under Consumer Search

dates search and sets the standard Hotelling best response,

$$p_1^s(p_2) \equiv \operatorname{argmax}_{p_1} \pi_1^s(p_1, p_2) = \frac{t + c + f + p_2}{2}.$$

Because active search fully eliminates the visibility advantage from self-preferencing, the biased recommendation algorithm no longer generates informational rents for the platform. Figure 3 summarizes the platform's optimal pricing rule $p_1^*(p_2)$ across the three regions.

A notable property emerges at $p_2 = \underline{p}_2$. At this lower threshold, the platform earns the same expected per-user profit from deterring search with $\bar{p}_1$ as from accommodating search with $p_1^s > \bar{p}_1$, as illustrated in Figure 2-(b). This payoff equivalence creates the possibility of a mixed-strategy equilibrium in which the platform randomizes between search deterrence and accommodation while consumers hold consistent beliefs.

The next lemma fully characterizes equilibrium prices as a function of the degree of self-preferencing b . For simplicity, we focus on the case of relatively low search costs, ($s \leq 3t/4$). The qualitative structure of the equilibrium remains unchanged for higher search costs, although the threshold values and equilibrium pricing formulas differ slightly.²² To characterize the pricing equilibrium, we adopt the terminology of Bain (1956), originally developed to describe an incumbent firm's response to the threat of entry.

LEMMA 2 (Price Equilibrium under Consumer Search). *Suppose $s \leq 3t/4$. Equilibrium prices vary with the degree of self-preferencing b as follows:*

²²The complete characterization for all parameter ranges of s is available from the authors upon request.

- (i) **Blockaded Search:** If $b \leq \underline{b}$, equilibrium prices coincide with those in the baseline model and no consumer searches.
- (ii) **Deterred Search:** If $b \in (\underline{b}, \hat{b}]$, the platform sets $p_1 = \bar{p}_1$ to deter search. The third-party seller then chooses

$$p_2 = c + 2t + f - \frac{b^2}{2s},$$

and the platform charges $\bar{p}_1(p_2) = c + 3t + f + b - b^2/2s$.

- (iii) **Mixed Strategy with Partially Deterred Search:** If $b \in (\hat{b}, \bar{b})$, the third-party seller sets $p_2 = \underline{p}_2$, and the platform randomizes between $\bar{p}_1$ (detering search) and p_1^s (accommodating search).
- (iv) **Accommodated Search:** If $b \geq \bar{b}$, the platform accommodates search, and the equilibrium prices are $p_1 = p_2 = c + t + f$ as in the standard Hotelling model.

The thresholds $\underline{b}$, $\hat{b}$, and $\bar{b}$ are uniquely determined by the conditions

$$p_2^{ns} = \bar{p}_2, \quad p_2 = \underline{p}_2, \quad p_2^s = p_1^s,$$

respectively, where $p_2 = c + 2t + f - b^2/2s$. Substituting the equilibrium expressions yields the following equivalent characterizations in terms of the model parameters:

$$\frac{3b^2}{2s} - b = 3t, \quad \frac{3b^2}{2s} - b\sqrt{\frac{2b}{s}} = 3t, \quad \frac{3b^2}{2s} - \frac{3b}{2}\sqrt{\frac{2b}{s}} = 3t,$$

which determine $\underline{b}$, $\hat{b}$, and $\bar{b}$, respectively.

PROOF OF LEMMA 2: See Appendix A.5. $\square$

The proof is straightforward except in case (iii), where the platform randomizes between two prices because its best response is discontinuous at $\underline{p}_2$. Since the third-party seller cannot directly affect consumer search, its best response is simple. If the third-party seller expects $p_1 \leq \bar{p}_1$, so that no one searches, then its best response is $p_2^{ns}(p_1) = (t + c + f - b + p_1)/2$. If instead the third-party seller expects $p_1 > \bar{p}_1$, so that consumers search, then its best response becomes $p_2^s(p_1) = (t + c + f + p_1)/2$, which is independent of b . A mixed strategy equilibrium arises when these search-consistent best responses do not intersect the platform's best response. As Figure 3 shows, although $p_2^s(p_1)$ intersects the platform's best response $p_1^*(p_2)$, the intersection occurs in the intermediate region of p_2 where the platform prefers to set a defensive limit price to deter search. Hence, it cannot constitute an equilibrium.

To illustrate the role of randomization, suppose the platform mixes between $\bar{p}_1$ and p_1^s with probabilities μ_1 and $1 - \mu_1$, respectively. The expected platform price is then $\mathbb{E}_{\mu_1}[p_1] = \mu_1\bar{p}_1 + (1 - \mu_1)p_1^s$. Under this strategy, a fraction $1 - \mu_1$ of consumers initially recommended product 1 engage in search. Anticipating this, the third-party seller's best response is the corresponding weighted




Figure 4: The effect of self-preferencing on the joint surplus[#]

[#] The vertical axis represents the components of the joint surplus $S = \phi + \sigma$ that vary with b ; constant terms independent of b are omitted.

average of its two conditional best responses, $\mu_1 p_2^{\text{ns}} + (1 - \mu_1) p_2^{\text{s}}$, as a function of $\mathbb{E}_{\mu_1}[p_1]$. There exists a unique $\mu_1 \in (0, 1)$ such that $(\mathbb{E}_{\mu_1}[p_1], \underline{p}_2)$ lies on this best response function, implying that $\underline{p}_2$ is optimal for the third-party seller. Consequently, both $\bar{p}_1$ and p_1^{s} maximize the platform's expected per-user profit, sustaining randomization in equilibrium.

Based on Lemma 2, we can now derive the platform's optimal choice of self-preferencing b in the first stage.

PROPOSITION 7. *Suppose that $s \leq 3t/4$. The platform's optimal degree of self-preferencing under consumer search is given by*

$$b^* = \min \left\{ \sqrt{\frac{10st}{3}}, \hat{b} \right\},$$

where $\hat{b}$ is the search-deterrence threshold defined in Lemma 2.

PROOF OF PROPOSITION 7: See Appendix A.6. $\square$

By Lemma 1, the platform's optimal degree of self-preferencing b^* maximizes the joint surplus, $S(b) = \phi(b) + \sigma(b)$. When the search cost is relatively small ($s \leq 3t/4$), the surplus function $S(b)$ is increasing for $b \leq \underline{b}$ and decreasing for $b \geq \hat{b}$. Consequently, the optimal level b^* must lie in the intermediate region $(\underline{b}, \hat{b}]$, where the platform sets the subsidiary's price exactly equal to the search-deterrence threshold $\bar{p}_1$, as characterized in Case (ii) of Lemma 2. As a result, *no* consumers engage in costly search in equilibrium.

In the search-deterrence regime, both p_1 and p_2 decrease in b . These defensive price adjustments, necessary to deter consumer search, initially dominate the allocative distortion created by self-preferencing, so the joint surplus rises up to a unique peak at $b^\dagger \equiv \sqrt{10st/3}$, and declines thereafter. Whether this peak lies in the interior of $(\underline{b}, \hat{b}]$ depends on the magnitude of the search

cost s . If $s > 27t/250$, then the equilibrium level of self-preferencing is given by $b^* = b^\dagger$, as illustrated in Figure 4-(a). By contrast, if $s \leq 27t/250$, then the equilibrium is $b^* = \hat{b}$, as shown in Figure 4-(b). Thus, when search costs are sufficiently low, the platform increases self-preferencing up to the maximal level at which consumer search can still be deterred. Beyond $\hat{b}$, consumers begin to search and switch away from the subsidiary product, making $\hat{b}$ a natural upper bound on the platform's steering incentives.

Proposition 7 shows that, as long as consumer search remains costly ($s > 0$), the platform retains an incentive to steer users toward its affiliated product. By Lemma 1, the equilibrium level of self-preferencing maximizes joint surplus in the absence of regulation, and Proposition 2 then implies that any policy prohibiting self-preferencing lowers user surplus. The underlying mechanism is the same as in the monopoly benchmark without search. Although a ban on self-preferencing reduces mismatch distortions, the resulting user welfare gain is more than offset by the platform's endogenous increase in transaction fees. Consumer search may mitigate the initial distortion by disciplining equilibrium self-preferencing, but it does not overturn the qualitative result. Because restricting self-preferencing reduces joint per-user surplus, consumers are ultimately worse off under self-preferencing regulation.

5. Extensions and Discussion

As noted earlier, our approach and analytical results are considerably more general than the Hotelling framework adopted for analytical tractability and closed-form characterization. In this section, we first outline a more general discrete-choice framework that allows for multiple third-party sellers and places no restrictions on the functional forms describing users' matching utility or the platform's retail profits. The Hotelling model analyzed in the previous sections is a special case of this broader framework. We then discuss two additional issues within the baseline monopoly model: the effects of regulatory policies on total surplus, including third-party seller profits, and the role of cost asymmetries between the platform's affiliated seller and independent sellers.

5.1. Beyond the Hotelling Model

Our analysis in Section 2 and 3 regarding the welfare effects of self-preferencing regulation hinges on the quasilinear form of both user surplus, $U = \phi(b) - f$, and the platform's per-user profit, $\pi_1 = \sigma(b) + f$. Importantly, our results do not depend on the specific functional forms of the user's matching utility ϕ or the platform's retail profit σ . Thus, the insights extend generally to an environment where U and π_1 maintain quasilinearity with respect to the platform's two instruments, b and f , provided that ϕ is decreasing and σ is increasing in the level of self-preferencing b . To illustrate the scope of possible extensions, we first formalize the micro-foundations under which U and π_1 exhibit these properties.

Denote by u_1 and u_2 a consumer's match values for the platform's subsidiary (Product 1) and

the third-party seller (Product 2), respectively. These values are observed by the platform, but not by consumers. Given a price vector $\mathbf{p} = (p_1, p_2)$ and an algorithmic bias level $b \geq 0$, the platform steers the consumer by recommending its subsidiary if and only if

$$u_1 - p_1 + b \geq u_2 - p_2. \quad (13)$$

Let H denote the cumulative distribution function of the idiosyncratic match value differential $u_2 - u_1$, which we assume to be log-concave with a continuous density h . The market share to the platform's subsidiary is then given by $\bar{x}_b(\mathbf{p}) = H(b - p_1 + p_2)$, a log-concave distribution. As a result, the equilibrium prices $\mathbf{p}^*$ are characterized by the first-order conditions:

$$p_1^* = c + f + \frac{H(v^*)}{h(v^*)} \quad \text{and} \quad p_2^* = c + f + \frac{1 - H(v^*)}{h(v^*)}, \quad (14)$$

where $v^* = b - p_1^* + p_2^*$.

From these conditions, it follows that the transaction fee f is passed through to the equilibrium prices exactly on a one-to-one basis. Furthermore, the equilibrium price gap $p_1^* - p_2^*$ depends solely on b , implying that v^* is a function of b alone and increasing in b . Using these properties and (14), we can write the platform's per-user profit as

$$\pi_1 = \bar{x}_b(\mathbf{p}^*) [p_1^* - c - f] + f = f + \underbrace{\frac{H^2(v^*)}{h(v^*)}}_{\equiv \sigma(b)}.$$

Consequently, the function π_1 retains the quasilinear structure as in Section 2. Furthermore, because log-concavity implies that the hazard ratio $H/h(v^*)$ is increasing in its argument, the retail profit function $\sigma(b)$ is increasing in the level of self-preferencing b .

The expected user surplus generated within the marketplace can be expressed as

$$\begin{aligned} U &= \mathbb{E} [(u_1 - p_1^*)\mathbb{1}_{\{u_2 - u_1 \leq v^*\}} + (u_2 - p_2^*)\mathbb{1}_{\{u_2 - u_1 > v^*\}}] \\ &= \mathbb{E}[u_2] - p_2^* + (p_2^* - p_1^*)H(v^*) - \int_{-\infty}^{v^*} v h(v) dv. \end{aligned}$$

Since the transaction fee f enters U exclusively via the one-to-one pass-through in p_2^* , while the other components are functions of b , the two instruments are additively separable and the surplus function retains the same form $U = \phi(b) - f$, where ϕ is decreasing in b . Therefore, the two functions U and π_1 take a quasilinear form in a general environment, provided that the platform's recommendation algorithm operates according to (13) and demand for the platform's subsidiary is represented by a log-concave distribution.

This environment encompasses a generalized Hotelling framework. To illustrate, suppose consumers are distributed according to an absolutely continuous distribution F along a line of unit length, facing an increasing transportation cost function $\kappa : [0, 1] \rightarrow \mathbb{R}$. The standalone matching utilities are $u_1 = v - \kappa(x)$ and $u_2 = v - \kappa(1 - x)$, respectively. Given the recommendation

algorithm (13), a consumer located at $x \sim F$ is recommended the subsidiary with probability

$$\Pr(\kappa(x) - \kappa(1-x) \leq b - p_1 + p_2) = F\left(\Delta_\kappa^{-1}(b - p_1 + p_2)\right),$$

where Δ_κ^{-1} denotes the inverse function of $\Delta_\kappa(x) \equiv \kappa(x) - \kappa(1-x)$. Let $H(v) \equiv F(\Delta_\kappa^{-1}(v))$, which serves as the distribution of match differentials $u_2 - u_1$ in a generalized Hotelling framework, mapping $v = b - p_1 + p_2$ to the subsidiary's demand. The distribution H is log-concave—preserving our welfare analysis—under mild conditions: namely, that the distribution F is log-concave and the function Δ_κ is weakly convex.²³

Our welfare analysis also extends naturally to a multi-product environment in which consumers derive random match values u_i from each product $i \in \{1, \dots, n\}$, as in the classic random-utility framework of [Perloff and Salop \(1985\)](#). Suppose that the match value u_1 from the platform's subsidiary (Product 1) is drawn from distribution H_1 with a log-concave density h_1 , and that the values u_j for the third-party sellers $j \in \{2, \dots, n\}$ are independently drawn from a common log-concave distribution. Observing the full realization vector $\{u_i\}_{i=1}^n$, the platform recommends its own subsidiary if and only if $u_1 - p_1 + b \geq \max_{j \geq 2} \{u_j - p_j\}$. Under a symmetric downstream pricing strategy $p_j = p_*$ among the independent sellers, the demand directed to the subsidiary is

$$H(v) \equiv \Pr(u_* - u_1 \leq v) = \int H_*(v + u) h_1(u) du,$$

where $v = b - p_1 + p_*$ and H_* is the cumulative distribution function of the first-order statistic $u_* = \max\{u_2, \dots, u_n\}$. Whenever the underlying match values are independent and log-concave, the distribution of the maximum H_* is inherently log-concave ([Bagnoli and Bergstrom, 2005](#)). Moreover, since log-concavity is preserved under convolution, the demand function $H(v)$ is log-concave. Consequently, even in a multi-product framework á la the discrete choice model, the platform's optimization problem retains its quasilinear structure, ensuring the robustness of our regulatory evaluations.

5.2. Impacts of Regulatory Policies on Total Surplus

So far, we have evaluated regulatory policies primarily from the perspective of consumer welfare. We now extend the analysis to total surplus, including third-party seller profits.

To this end, define total per-user surplus W as the sum of user surplus U , platform profits π_1 , and third-party seller profits π_2 . Since the transaction fee f is fully passed through to consumers and recovered by the platform as revenue, it represents a pure transfer and therefore does not affect aggregate welfare. Consequently, total per-user surplus depends only on the degree of self-

²³Since the function $\Delta_\kappa(x)$ is increasing, the inverse function $\Delta_\kappa^{-1}(x)$ is weakly concave if and only if $\Delta_\kappa(x)$ is weakly convex, or equivalently, the second derivative $\kappa''(x)$ is symmetric about $x = 1/2$.

preferencing:

$$W(b) \equiv \underbrace{U(b) + \pi_1(b)}_{= S(b)} + \pi_2(b), \quad \text{where } \pi_2(b) = [1 - \bar{x}(b)](p_2^* - c - f) = \frac{(3t - b)^2}{18t}$$

Note that $W(b)$ is monotonically decreasing in b over the relevant domain $[0, 3t]$. This reflects the allocative distortion created by self-preferencing: greater bias reduces match quality and shifts demand away from the independent seller toward the platform's affiliated seller. Consequently, *conditional* on platform participation, total per-user surplus is maximized under complete neutrality, corresponding to $b^0 = 0$.

We next define the aggregate total surplus, incorporating consumers' platform participation decisions as

$$TS(b) \equiv \int_0^U [W(b) - z] dG(z).$$

Recall that a consumer participates if and only if user surplus U exceeds her reservation utility $z \sim G$. Since a participating consumer generates net surplus $W(b) - z$, integrating over all consumers who choose to participate yields the aggregate total surplus created within the platform ecosystem.

Let $b^{SB} = \operatorname{argmax}_{b \geq 0} TS(b)$ denote the second-best welfare-maximizing level of self-preferencing. This corresponds to a regulatory regime in which the authority can directly regulate the degree of self-preferencing but leaves the transaction fee to be optimally chosen by the platform. The following proposition characterizes the resulting second-best level of self-preferencing.

PROPOSITION 8. *The second-best welfare-maximizing level of self-preferencing lies strictly between the level that maximizes total per-user surplus and the platform's privately optimal level; that is,*

$$b^{SB} \in (b^0, b^*).$$

PROOF OF PROPOSITION 8: See Appendix A.7. $\square$

Proposition 8 implies that appropriately restricting self-preferencing can increase total surplus. Although such a policy reduces both consumer surplus and platform profits, the resulting increase in third-party seller profits may more than offset these losses. Implementing the second-best level b^{SB} , however, is unlikely to be feasible in practice. Even with continuous monitoring, the precise degree of algorithmic bias is difficult to observe and verify in complex recommendation systems. A more realistic policy would be therefore a complete ban on self-preferencing, though its effect on total surplus is ambiguous and depends on the location of b^{SB} relative to b^0 and b^* .

The effect of structural separation on total surplus is likewise ambiguous, even though it unambiguously reduces consumer welfare. Still, structural separation is inferior to a behavioral

ban as a policy instrument. Under both policies, steering bias is eliminated, so the per-user total surplus generated by platform participants, $W(0)$, is identical. However, consumer surplus is higher under a behavioral ban because transaction fees are lower. Under structural separation, the platform loses its downstream profit source and therefore relies more heavily on transaction fees, leading to higher prices, lower consumer surplus, and reduced consumer participation. Structural separation is therefore justified only when algorithmic steering cannot be monitored and a direct ban cannot be enforced.

5.3. Cost Asymmetry

Our preceding analysis has assumed cost symmetry between the platform’s subsidiary and the independent seller. We now relax this assumption by allowing for cost asymmetry, under which the subsidiary may be less efficient than the third-party seller, and examine the implications for equilibrium self-preferencing and regulatory outcomes. The main insights are summarized below, with the formal analysis relegated to the Online Appendix.

First, cost asymmetry can make a moderate degree of self-preferencing directly beneficial to consumers, even holding transaction fees fixed. When a more efficient third-party seller enjoys a substantial cost advantage and therefore serves a large share of the market, steering some consumers toward the less efficient subsidiary can partially offset the third party’s market power and induce a more aggressive pricing response. Although the favored subsidiary raises its own price, the reduction in the third-party seller’s price affects a larger mass of consumers. Consequently, the net effect on consumer prices may be positive despite the increase in the subsidiary’s price. As a result, a moderate degree of self-preferencing may increase consumer surplus despite the mismatch distortion created by biased recommendations.²⁴

This observation also has implications for the strategic relationship between self-preferencing and transaction fees. While our baseline analysis establishes that the two instruments are strategic substitutes, cost asymmetry can generate circumstances in which they become complements, particularly when $\phi(b^*) > \phi(0)$; see the Online Appendix for an illustrative example. However, such cases arise only under extreme parameter values involving a sufficiently large cost advantage for the third-party seller and highly elastic consumer participation. As a result, this possibility does not affect any of the paper’s main qualitative conclusions.

Cost asymmetry also interacts with platform competition in meaningful ways. Competition may either mitigate or amplify the effects of self-preferencing depending on the relative efficiency of sellers and the degree of platform differentiation. In particular, contrary to the conventional view that intense platform competition disciplines recommendation bias, self-preferencing may persist even in the limit of perfect platform competition ($\tau \rightarrow 0$) when third-party sellers enjoy a sufficiently large cost advantage. Moreover, the parameter region identified in Proposition 4,

²⁴Related issues arise in the literature on most-favored-nation (MFN) clauses, such as in models of input price discrimination (DeGraba, 1990) and discriminatory tariffs in international trade (Hwang and Mai, 1991). Our framework highlights analogous strategic interactions in the context of vertically integrated platforms through the joint use of algorithmic steering and fee setting.

in which a ban on self-preferencing improves user surplus, may disappear altogether when that advantage becomes sufficiently large.

Finally, cost asymmetry has important implications for dynamic investment incentives. A larger cost advantage reduces the platform's incentive to distort recommendations, potentially encouraging third-party sellers to invest in quality improvements or cost reductions. However, if the platform cannot commit to its recommendation policy *ex ante*, a hold-up problem emerges. Anticipating that future efficiency improvements will be met with more aggressive self-preferencing, third-party sellers may have weaker incentives to invest in the first place.

6. Concluding Remarks

Self-preferencing by vertically integrated platforms has become a major antitrust concern across jurisdictions, prompting policymakers to consider both behavioral and structural remedies. This article has examined the welfare consequences of such interventions in models of monopoly and platform competition.

Our analysis yields three main conclusions. First, self-preferencing and transaction fees are strategic substitutes: when platforms are constrained to reduce self-preferencing, they tend to respond by raising transaction fees. As a result, a ban on self-preferencing need not improve consumer welfare and may in fact reduce it by inducing higher retail prices. Second, structural remedies such as vertical separation can amplify this effect by depriving the platform of an important source of revenue from its affiliated seller. As a consequence, the platform may rely even more heavily on transaction fees, leading to lower consumer surplus than under a behavioral ban. Third, in a competitive platform environment, the interaction between competition and self-preferencing becomes more nuanced. Stronger platform competition may limit the ability of platforms to increase transaction fees in response to regulatory interventions, creating circumstances in which a ban on self-preferencing can raise consumer welfare even though structural separation remains harmful. Ironically, the disciplining role of platform competition does not render regulation unnecessary. Rather, competition and self-preferencing regulation can complement one another: a ban on self-preferencing may improve user welfare precisely when platform competition is most intense.

Taken together, our findings highlight the importance of accounting for the full set of instruments available to platforms when evaluating regulatory interventions. When a platform can adjust multiple margins of conduct, restricting one dimension of behavior may induce offsetting responses in another, potentially undermining the intended effects of regulation. More broadly, the welfare consequences of self-preferencing regulation depend critically on market conditions, particularly the intensity of platform competition and the platform's ability to substitute across instruments. These considerations suggest that effective platform regulation requires a holistic approach that anticipates strategic responses rather than evaluating individual practices in isolation.

References

- AGHION, P. and BOLTON, P. (1987). Contracts as a Barrier to Entry. *The American Economic Review*, **77** (3), 388–401.
- ANDERSON, S. and BEDRE-DEFOLIE, O. (2024). Hybrid Platform Model: Monopolistic Competition and a Dominant Firm. *The RAND Journal of Economics*, **55** (4), 684–718.
- ARMSTRONG, M. (2006). Competition in Two-Sided Markets. *The RAND Journal of Economics*, **37** (3), 668–691.
- , VICKERS, J. and ZHOU, J. (2009). Prominence and Consumer Search. *The RAND Journal of Economics*, **40** (2), 209–233.
- and ZHOU, J. (2011). Paying for Prominence. *The Economic Journal*, **121** (556), F368–F395.
- ATHEY, S. and ELLISON, G. (2011). Position Auctions with Consumer Search. *The Quarterly Journal of Economics*, **126** (3), 1213–1270.
- BAGNOLI, M. and BERGSTROM, T. (2005). Log-Concave Probability and Its Applications. *Economic Theory*, **26** (2), 445–469.
- BAIN, J. S. (1956). *Barriers to New Competition: Their Character and Consequences in Manufacturing Industries*. Cambridge, MA: Harvard University Press.
- BISCEGLIA, M. and TIROLE, J. (2024). Fair Gatekeeping in Digital Ecosystems. *TSE Working Paper* (1452).
- CHEN, N. and TSAI, H.-T. (2024). Steering via Algorithmic Recommendations. *The RAND Journal of Economics*, **55** (4), 501–518.
- CHEN, Y. and HE, C. (2011). Paid Placement: Advertising and Search on the Internet. *The Economic Journal*, **121** (556), F309–F328.
- CHOI, J. P. and JEON, D.-S. (2023). Platform Design Biases in Ad-Funded Two-Sided Markets. *The RAND Journal of Economics*, **54** (2), 240–267.
- , KIM, K. and MUKHERJEE, A. (2025). “Sherlocking” and Platform Information Policy. *Management Science*, **71** (11), 9812–9830.
- DE CORNIÈRE, A. and TAYLOR, G. (2014). Integration and Search Engine Bias. *The RAND Journal of Economics*, **45** (3), 576–597.
- and — (2019). A Model of Biased Intermediation. *The RAND Journal of Economics*, **50** (4), 854–882.

- DEGRABA, P. (1990). Input Market Price Discrimination and the Choice of Technology. *The American Economic Review*, **80** (5), 1246–1253.
- EDELMAN, B. and WRIGHT, J. (2015). Price Coherence and Excessive Intermediation. *The Quarterly Journal of Economics*, **130** (3), 1283–1328.
- ETRO, F. (2023). Platform Competition with Free Entry of Sellers. *International Journal of Industrial Organization*, **89** (102903).
- HAGIU, A., TEH, T. and WRIGHT, J. (2022). Should Platforms Be Allowed to Sell on Their Own Marketplaces? *The RAND Journal of Economics*, **53** (2), 297–327.
- HERVAS-DRANE, A. and SHELEGIA, S. (2025). Retailer-Led Marketplaces. *Management Science*, **71** (11), 9650–9669.
- HWANG, H. and MAI, C.-C. (1991). Optimum Discriminatory Tariffs under Oligopolistic Competition. *The Canadian Journal of Economics / Revue canadienne d'Economique*, **24** (3), 693–702.
- INDERST, R. and OTTAVIANI, M. (2012). Competition through Commissions and Kickbacks. *The American Economic Review*, **102** (2), 780–809.
- JUUL, M. (2021). The Amazon Buy Box Playbook for Sellers and Retailers. Feedvisor, available at <https://feedvisor.com/resources/e-commerce-strategies/the-amazon-buy-box-playbook-for-sellers-and-retailers/>.
- KITAKA, Y. and SATO, S. (2022). Dual-Role Platforms and Self-Preferencing: Sequential Search Approach. *SSRN Working Paper* (3736574).
- and — (2024). Search-Order Design by Dual-Role Platforms. *SSRN Working Paper* (4665476).
- MADSEN, E. and VELLODI, N. (2025). Insider Imitation. *The Journal of Political Economy*, **133** (2), 652–709.
- PERLOFF, J. M. and SALOP, S. C. (1985). Equilibrium with Product Differentiation. *The Review of Economic Studies*, **52** (1), 107–120.
- RAVAL, D. (2022). *Steering in One Click: Platform Self-Preferencing in the Amazon Buy Box*. Unpublished Manuscript.
- ROCHET, J.-C. and TIROLE, J. (2006). Two-Sided Markets: A Progress Report. *The RAND Journal of Economics*, **37** (3), 645–667.
- TEH, T.-H. and WRIGHT, J. (2022). Intermediation and Steering: Competition in Prices and Commissions. *American Economic Journal: Microeconomics*, **14** (2), 281–321.
- VARIAN, H. R. (2007). Position Auctions. *International Journal of Industrial Organization*, **25** (6), 1163–1178.

ZENNYO, Y. (2022). Platform Encroachment and Own-Content Bias. *Journal of Industrial Economics*, **70** (3), 684–710.

Appendix A Omitted Proofs

This Appendix provides the proofs omitted from the main text.

A.1 Proof of Proposition 1

In the monopoly model, the joint per-user surplus S is a twice-differentiable function of b . Hence it follows by Lemma 1 that the optimal level of self-preferencing b^* has to meet the first-order condition $S'(b) = 0$:

$$\phi'(b) + \sigma'(b) = -\frac{10b}{36t} + \frac{3t+b}{9t} = 0.$$

Since $S'' < 0$ for all $b \geq 0$, the condition is indeed sufficient for b^* . Solving for b yields $b^* = 2t$.

To derive the platform's optimal transaction fee, we substitute $b = b^*$ into the platform's aggregate profit function, to get $\Pi(b^*, f) = G(\phi(b^*) - f) [\sigma(b^*) + f]$. The objective is log-concave in f . To see this, we take the natural logarithm and then take the derivative with respect to f , to obtain

$$\frac{d \log \Pi(b^*, f)}{df} = -\frac{g}{G}(U) + \frac{1}{\sigma(b^*) + f}, \quad (\text{A.1})$$

where $U = \phi(b^*) - f$. Under the assumption that G is log-concave, the first term $-g(U)/G(U)$ is nondecreasing in U and thus non-increasing in f . Hence the log-profit is strictly concave in f , and the optimal fee f^* is uniquely determined by (A.1).

To establish that b and f are strategic substitutes, recall that the function ϕ is monotone decreasing and σ is monotone increasing in b . This ensures that

$$\frac{\partial \log \Pi}{\partial f} = -\frac{g}{G}(\phi(b) - f) + \frac{1}{\sigma(b) + f},$$

is decreasing in b , demonstrating Π is log-submodular in (b, f) . $\square$

A.2 Proof of Proposition 3

As the high differentiation regime ($\tau > \sigma(b^*)$) is analyzed in the main text, we restrict attention to the case of intense competition where $\tau \leq \sigma(b^*)$.

We first formally show that both platforms are forced to set $f = 0$ in any equilibrium. Suppose to the contrary that platform i imposes a positive transaction fee $f_i > 0$ and its rival platform j sets $f_j = 0$. We first note that any $b_i > b^*$ is strictly dominated regardless of the non-negativity constraint on f_i ; since $S'(b_i) < 0$ for every $b_i > b^*$, a platform could reduce b_i toward b^* and raise f_i to keep U_i constant, thereby increasing its total profit Π_i without affecting its market share. This tells us that both b_i and b_j are bounded above by b^* in any equilibrium. Given $b_i \leq b^*$ and U_j , we

examine the marginal incentive to increase user surplus U_i (by decreasing f_i):

$$\frac{\partial \Pi_i}{\partial U_i} = \frac{1}{2\tau} (S(b_i) - U_i) - z(U_i, U_j) = \frac{1}{2\tau} (2f_i + \sigma(b^*) - \tau - \phi(b_i) + \phi(b_j)).$$

Under the given conditions $\sigma(b^*) - \tau \geq 0$ and $f_i > 0$, $b_i \geq b_j$ results in the partial derivative above being strictly positive, implying that platform i prefers to lower its fee and attract more users whenever $b_i \geq b_j$. This leads to $b_i < b^*$. However, when $b_i < b^*$ and $f_i > 0$, the platform increases its total profit by increasing b_i toward the efficient level b^* and simultaneously decreasing f_i accordingly.²⁵ Therefore, intense platform competition drives the transaction fee to zero in any equilibrium.

It remains to determine the equilibrium level of self-preferencing b^+ when the constraint $f_i \geq 0$ is binding. In this regime, the platform no longer use the transaction fee to compete for market share and thus $U_i = \phi(b_i)$ and $\pi_i = \sigma(b_i)$ for each platform $i = L, R$. Hence platform i 's problem (7) reduces to

$$\max_{b_i \geq 0} \Pi_i = z(\phi(b_i), \phi(b_j)) \sigma(b_i),$$

namely, the platforms compete for users solely through the quality of their recommendation. The marginal profit to platform i with respect to b_i is given by

$$\frac{1}{2\tau} [\phi'(b_i) \sigma(b_i) + (\tau + \phi(b_i) - \phi(b_j)) \sigma'(b_i)] \quad (\text{A.2})$$

Because ϕ is a decreasing function and $\sigma' > 0$, this marginal profit is increasing in b_j , indicating that the objective function Π_i exhibits increasing differences in $(b_i; b_j)$. The problem of platform j is symmetric, implying that the self-preferencing strategies are complements. As each platform's best response function is non-decreasing, the existence of an equilibrium is guaranteed by Tarski's fixed point theorem. Substituting the fixed point $b_i = b_j = b^+$ into (A.2) and setting it equal to zero, we obtain the necessary condition for b^+ , which is $\phi'(b^+) \sigma(b^+) + \tau \sigma'(b^+) = 0$. The uniqueness of b^+ is established in the main text. The proof is now complete. $\square$

A.3 Proof of Proposition 4

To evaluate the welfare impact of the regulatory ban, we define the surplus difference as $\Delta_U(\tau) \equiv U^* - U_{\text{ban}}$. Because the equilibrium user surplus functions U^* and U_{ban} defined in (11) and (12) are piecewise functions with different cutoffs. To investigate the effect of the regulatory ban on the user surplus, we analyze the sign of Δ_U across three distinct regimes: (i) $\tau > \sigma(b^*)$ (*mild competition*), (ii) $\tau \leq \sigma(0)$ (*intense competition*), and (iii) $\tau \in (\sigma(0), \sigma(b^*))$ (*intermediate competition*).

Analysis of the first two cases is straightforward. In case (i), the surplus difference simplifies to $\Delta_U(\tau) = S(b^*) - S(0) > 0$, because the joint per-user surplus $S(b)$ is uniquely maximized at the

²⁵To be specific, consider a local deviation $db_i > 0$ and $df_i < 0$ such that $dU_i = \phi'(b_i)db_i - df_i = 0$. This deviation increases the joint per-user surplus $S(b_i)$ and, since the platform is the residual claimant of that surplus when $f_i > 0$, it increases Π_i without changing its market share.

efficient level b^* . Consequently, as in the monopoly benchmark, the platform fully passes through the efficiency loss of the ban to its users via higher transaction fees, causing user surplus to fall. In case (ii) where the transaction fees are zero in both regimes, the surplus difference reduces to $\Delta_U(\tau) = \phi(b^\dagger) - \phi(0)$. Since $b^\dagger > 0$ and the consumer's matching utility $\phi(b)$ is monotone decreasing in b , it follows that $\Delta_U(\tau) < 0$ for $\tau \leq \sigma(0)$. Thus, when platform competition is sufficiently intense, the regulatory ban on self-preferencing increases user surplus.

In the intermediate regime where $\tau \in (\sigma(0), \sigma(b^*)]$, the surplus difference takes form of $\Delta_U(\tau) = \phi(b^\dagger(\tau)) + \tau - \phi(0) - \sigma(0)$. Below we prove that $\Delta_U(\tau)$ is monotone increasing in τ . For this purpose, we first establish the following lemma:

LEMMA A.1. *For each $\tau \in (\sigma(0), \sigma(b^*)]$, denote by $b^\dagger(\tau)$ the equilibrium level of self-preferencing satisfying*

$$\mathfrak{J}(b^\dagger(\tau); \tau) = \phi'(b^\dagger(\tau)) + \frac{\tau}{\sigma(b^\dagger(\tau))} \sigma'(b^\dagger(\tau)) = 0. \quad (\text{A.3})$$

Then the ratio $\tau/\sigma(b^\dagger(\tau))$ is increasing in τ .

PROOF OF LEMMA A.1: Consider any $\tau_1 > \tau_2$ within the interval $(\sigma(0), \sigma(b^*)]$, and let $b_i^\dagger = b^\dagger(\tau_i)$ for $i = 1, 2$. Recall that $b^\dagger$ is increasing in τ with $b^\dagger < b^*$. Since the joint per-user surplus $S(b) = \phi(b) + \sigma(b)$ is concave and maximized at $b = b^*$, its derivative is positive and decreasing on this domain, yielding

$$0 < \phi'(b_1^\dagger) + \sigma'(b_1^\dagger) < \phi'(b_2^\dagger) + \sigma'(b_2^\dagger).$$

As ϕ is decreasing and concave, to be consistent with the condition (A.3), it must be the case that $\tau_1/\sigma(b_1^\dagger) > \tau_2/\sigma(b_2^\dagger)$, establishing that the ratio is increasing. $\square$

We are now ready to show that the surplus difference Δ_U is increasing in τ . To see this concisely, we change variables by defining $y \equiv \sigma(b^\dagger(\tau)) > 0$ and let $J(y) \equiv \phi(\sigma^{-1}(y))$. With them and the inverse function theorem, we can restate the equilibrium condition (A.3) for $b^\dagger$ as

$$\mathfrak{J}(b^\dagger; \tau) = \sigma'(b^\dagger) \left[\sigma(b^\dagger) \frac{\phi'}{\sigma'}(b^\dagger) + \tau \right] = 0 \implies yJ'(y) + \tau = 0.$$

Differentiating $\Delta_U(\tau)$ with respect to τ and using $J'(y) = -\frac{\tau}{y}$ gives

$$\frac{d}{d\tau} [\phi(b^\dagger(\tau)) + \tau] = \frac{d}{d\tau} [J(y(\tau)) + \tau] = -\frac{\tau y'(\tau)}{y(\tau)} + 1 = y(\tau) \frac{d}{d\tau} \left(\frac{\tau}{y(\tau)} \right) > 0,$$

where the inequality is due to Lemma A.1.

To complete the proof of the proposition, note that as $\tau \rightarrow \sigma(0)$, the surplus difference continuously matches case (ii), where $\Delta_U(\sigma(0)) = \phi(b^\dagger) - \phi(0) < 0$. On the other hand, as $\tau \rightarrow \sigma(b^*)$, the function continuously matches case (i), where $\Delta_U(\sigma(b^*)) = S(b^*) - S(0) > 0$. Therefore, by the

intermediate value theorem, there exists a unique threshold $\hat{\tau} \in (\sigma(0), \sigma(b^*))$ such that $\Delta_U(\tau) < 0$ for every $\tau < \hat{\tau}$. $\square$

A.4 Proof of Proposition 5

We evaluate the welfare impact of structural regulation by analyzing the sign of the user surplus difference, that is given by $\Delta_U(\tau) = U^* - U_{\text{sep}}$. Since the unregulated user surplus U^* is a piecewise function, we consider two regimes: (i) $\tau > \sigma(b^*)$ and (ii) $\tau \leq \sigma(b^*)$.

When $\tau > \sigma(b^*)$, the user surplus difference is

$$\Delta_U(\tau) = \phi(b^*) + \sigma(b^*) - \phi(0) > S(b^*) - S(0) > 0,$$

where the first inequality is due to $S(0) = \phi(0) + \sigma(0) > \phi(0)$ and the second inequality is due to the fact that the joint per-user surplus $S(b)$ is uniquely maximized at $b = b^*$.

When $\tau \leq \sigma(b^*)$, the unregulated platform sets its transaction fee to zero, and the user surplus difference is $\Delta_U(\tau) = \phi(b^+(\tau)) + \tau - \phi(0)$. As is shown in the proof of Proposition 4, the non-constant component $\phi(b^+(\tau)) + \tau$ is increasing in τ for $\tau \leq \sigma(b^*)$. In the limiting case of perfect competition where $\tau \rightarrow 0$, the equilibrium level of self-preferencing converges to zero ($b^+ \rightarrow 0$). Hence $\lim_{\tau \rightarrow 0} \Delta_U(\tau) = 0$. Since Δ_U is increasing in τ throughout the interval $(0, \sigma(b^*))$, we have $\Delta_U(\tau) > 0$ for all $\tau > 0$. Therefore, platform users are always worse off under structural separation. $\square$

A.5 Proof of Lemma 2

Consumer Search. We begin with the consumer's search decision. Consider a consumer who observes (b, p_1) and expects the third-party product to be priced at p_2^e . The consumer benefits from search only if, after learning his true location x , he prefers product 2 over product 1, namely when $x \in (\hat{x}, \bar{x}_b)$. The expected benefit from switching to product 2 is²⁶

$$v - t \left(1 - \frac{\bar{x}_b + \hat{x}}{2} \right) - p_2^e - \left[v - \frac{t(\bar{x}_b + \hat{x})}{2} - p_1 \right] = \begin{cases} b/2 & \text{if } \hat{x} > 0 \\ \frac{b-t+p_1-p_2^e}{2} & \text{otherwise.} \end{cases}$$

Since the consumer is willing to purchase product 2 only in the event of $x \in (\hat{x}, \bar{x}_b)$, the expected benefit from search, denoted Γ , equals the above payoff multiplied by the conditional probability of this event:

$$\Gamma(p_1, p_2^e, b) = \begin{cases} \frac{b^2}{2(t+b-p_1+p_2^e)} & \text{for } p_1 < p_2^e + t \\ \frac{b-t+p_1-p_2^e}{2} & \text{otherwise.} \end{cases}$$

²⁶If $\hat{x}(p_1, p_2^e) = 0$, or equivalently, if $p_2^e \leq p_1 - t$, the consumer would prefer to purchase product 2 after search. In this case the expected benefit from search is computed differently because switching occurs with probability one.

Let λ denote the fraction of consumers who follow the platform's recommendation and purchase product 1. Consumers recommended product 1 choose to search iff $\Gamma \geq s$, which is equivalent to $p_1 > \bar{p}_1$, where the threshold price is given by

$$\bar{p}_1 = \begin{cases} t + p_2^e - b + 2s & \text{for } b \in [0, 2s] \\ t + p_2^e + b - \frac{b^2}{2s} & \text{for } b > 2s. \end{cases} \quad (\text{A.4})$$

Hence the optimal search decision is $\lambda = 1$ iff $p_1 \leq \bar{p}_1$.

Optimal Pricing. Represent seller i 's pricing (possibly mixed) strategy by μ_i , where $\mu_i(p) = \Pr(p_i \leq p)$. Given (μ_1, λ) , the third-party seller's expected per-user profit is

$$\mathbb{E}\pi_2(p_2|\mu_1, \lambda) = \frac{t - \lambda b + \mathbb{E}_{\mu_1}[p_1] - p_2}{2t} (p_2 - c - f).$$

Since this function is strictly concave in p_2 , the first-order condition yields the unique optimal price:

$$p_2 = \frac{1}{2} (t + c + f - \lambda b + \mathbb{E}_{\mu_1}[p_1]). \quad (\text{A.5})$$

This demonstrates that the independent seller adopts a pure strategy in equilibrium, which we denote simply by p_2 .

We next characterize the platform's optimal pricing. Given (p_2, λ) and correct consumer beliefs $p_2 = p_2^e$, the platform's per-user profit is

$$\pi_1(p_1|p_2, \lambda) = \begin{cases} f + \frac{t+b-p_1+p_2}{2t} (p_1 - c - f) & \text{for } p_1 \leq \bar{p}_1 \\ f + \frac{t-p_1+p_2}{2t} (p_1 - c - f) & \text{for } p_1 > \bar{p}_1 \text{ and } p_1 < p_2 + t \\ f & \text{for } p_1 > \bar{p}_1 \text{ and } p_1 \geq p_2 + t. \end{cases}$$

When $p_1 > \bar{p}_1$, all consumers recommended product 1 engage in search ($\lambda = 0$). Moreover, if $p_1 \geq p_2 + t$, then $\hat{x} = 0$; put differently, consumers who search always switch to product 2. In this case demand for product 1 vanishes, and thus the platform collects only the transaction fee f .

In light of the expression for $\bar{p}_1$ in (A.4), we divide the analysis into two cases depending on whether $b \leq 2s$. When $b \leq 2s$, we have $\bar{p}_1 \geq p_2 + t$, implying that any $p_1 > \bar{p}_1$ is dominated. The platform's best response is therefore p_1^{ns} if $p_2 \geq \hat{p}_2$, and $\bar{p}_1$ otherwise, where $\hat{p}_2$ is determined by the condition $p_1^{\text{ns}} = \bar{p}_1$:

$$\frac{t + c + f + b + \hat{p}_2}{2} = t + \hat{p}_2 - b + 2s \implies \hat{p}_2 = c + f - t + 3b - 4s.$$

When $b > 2s$, the profit function π_1 achieves its maximum at p_1^{ns} whenever $p_1^{\text{ns}} \leq \bar{p}_1$, or

equivalently,

$$p_2 \geq \bar{p}_2 = c + f - t - b + \frac{b^2}{s}$$

As p_2 falls below $\bar{p}_2$, the optimal price becomes $\bar{p}_1$ until p_2 reaches the lower threshold $\underline{p}_2$, at which the platform earns the same profit from $\bar{p}_1$ and p_1^s . Solving the indifference condition yields the unique relevant root

$$\underline{p}_2 = c + f - t + \frac{b^2}{s} - \sqrt{\frac{2b^3}{s}}. \quad (\text{A.6})$$

Lastly, for $p_2 < \underline{p}_2$, the platform optimally sets $p_1 = p_1^s > \bar{p}_1$, thereby accommodating consumer search.

Equilibrium Prices. Consider a subgame initiated by the platform's choice of $b \in [0, 2s]$. When b is sufficiently small, consumers follow the platform's recommendation and do not search. Anticipating $\lambda = 1$, the two sellers set prices as in the baseline model, where the third-party seller charges $p_2 = c + t + f - b/3$. Observe that as b increases, p_2 declines but remains above the cutoff $\hat{p}_2$, provided that $s \leq 3t/4$. Hence, for $b \leq 2s$, the equilibrium prices of the baseline model are supported as an equilibrium of the search model. As b increases further, however, p_2 eventually hits the threshold $\bar{p}_2$, which occurs at $b = \underline{b} > 2s$. Consequently, consumer search does not affect equilibrium prices when $b \leq \underline{b}$.

As b increases beyond $\underline{b}$, the third-party seller charges a lower price than $\bar{p}_2$, provided consumers do not search. In response, the platform sets its subsidiary price at $\bar{p}_1$ to deter search. As a result, the equilibrium prices are determined by the intersection of the best response (A.5) with $\lambda = 1$ and $\bar{p}_1$. These prices constitute an equilibrium until b rises to $\hat{b}$, at which point the equilibrium price $p_2 = c + 2t + f - b^2/2s$ reaches the lower threshold $\underline{p}_2$ in (A.6).

As b increases beyond $\hat{b}$, the platform's best response no longer intersects the third-party seller's best response (A.5) for consumer search behavior λ consistent with p_1 . Hence no pure-strategy equilibrium exists. Since the third-party seller plays a pure pricing strategy, we consider the possibility that the platform adopts a mixed pricing strategy. However, when $b > \hat{b}$ and $p_2 \neq \underline{p}_2$, the objective function π_1 is maximized at a single point. Therefore, for the platform to mix, it must be the case that $p_2 = \underline{p}_2$, in which case the platform is indifferent between $\bar{p}_1$ and p_1^s , as displayed in Figure 2-(b). With a slight abuse of notation, let $\mu_1 = \Pr(p_1 = \bar{p}_1)$. Under this mixed strategy, only consumers who observe the high price p_1^s would engage in search, implying $\lambda = \mu_1$. Substituting $\mathbb{E}_{\mu_1}[p_1] = \mu_1 \bar{p}_1 + (1 - \mu_1)p_1^s$ into (A.5) and setting it equal to $\underline{p}_2$ yields

$$\mu_1 = \sqrt{\frac{s}{2b^3}} \left(6t - \frac{3b^2}{s} + 3\sqrt{\frac{2b^3}{s}} \right).$$

Notice that μ_1 equals one at $b = \hat{b}$ and decreases in b . The mixed-strategy equilibrium persists

until $\mu_1 = 0$; in other words, the third-party seller's best response under search, $p_2^s = (t + c + f + p_1)/2$, intersects p_1^s . This determines the threshold $\bar{b}$. Consequently, for $b \geq \bar{b}$, the equilibrium prices of the standard Hotelling model constitute an equilibrium. This completes the proof of the lemma. $\square$

A.6 Proof of Proposition 7

To characterize the optimal level of self-preferencing b^* under consumer search, we first note that by Lemma 2, the transaction fee f is completely passed through users regardless of the platform's choice of b . As a result, the expected user surplus U takes the form $\phi(b) - f$ and the platform's per-user profit π_1 is $f + \sigma(b)$. It then follows by Lemma 1 that b^* maximizes the joint surplus, $S(b) = U + \pi_1 = \phi(b) + \sigma(b)$. In contrast to the baseline model, however, the function $\phi(b) + \sigma(b)$ is not differentiable at the thresholds, $\underline{b}$, $\hat{b}$, and $\bar{b}$, which are defined in Lemma 2. For this reason, the optimal b^* cannot be characterized solely by the first-order condition. Nonetheless, since the equilibrium prices change continuously with respect to b , the objective $S(b)$ is continuous in b , which guarantees the existence of an optimal b . To identify the optimal b , therefore, we need to examine the behavior of $S(b)$ within each regime, as follows.

When $b \leq \underline{b}$, the two sellers set their product prices as in the baseline model, and as established in Proposition 1, the joint surplus $U(b)$ is increasing for all $b \leq 2t$. Recall that the threshold $\underline{b}$ is defined as the positive solution to the quadratic equation,

$$\frac{3b^2}{2s} - b - 3t = 0.$$

Under the assumption $s \leq 3t/4$, evaluating this quadratic at $b = 2t$ yields a positive value, leading to $\underline{b} < 2t$. Consequently, $S(b)$ is continuously increasing throughout the interval $[0, \underline{b}]$, and therefore, the platform's optimal choice b^* will never fall within this range.

When the recommendation algorithm is excessively biased ($b \geq \bar{b}$), all users who are recommended product 1 engage in search, and in response, the sellers set prices as in the standard Hotelling model. As a result, the platform's per-user profit π_1 remains invariant to further increases in b . By contrast, the user surplus U is decreasing in b , because a higher b increases $\bar{x}$, i.e., the mass of users who incur search costs to find their preferred product. Therefore, the joint surplus is continuously decreasing for all $b \geq \bar{b}$.

When $b \in (\hat{b}, \bar{b})$, the third-party seller sets $p_2 = p_2$, while the platform adopts a mixed strategy, randomizing between $\bar{p}_1$ and p_1^s . In this regime, the third-party price p_2 , defined in (A.6), increases with b .²⁷ Due to strategic complementarity in pricing in a Hotelling model, the subsidiary's prices, both $\bar{p}_1$ and p_1^s , also increase in b . Consequently, self-preferencing helps the platform increase its

²⁷Differentiating with respect to b yields

$$\frac{dp_2}{db} = \frac{b}{s} \left(2 - 3\sqrt{\frac{s}{2b}} \right),$$

which is positive for all $b > 2s$. Since the assumption $s \leq 3t/4$ implies $2s < \underline{b}$, the derivative is positive for all $b > \hat{b} > \underline{b}$.

per-user profit π_1 within this regime. However, the inflated prices cause the consumer surplus U to decrease at a rate that outweighs the gain in π_1 . Therefore, the objective $S(b)$ is decreasing in b . A formal and exhaustive proof of this monotone behavior can be found in the Online Appendix.

Our preceding discussion shows that the platform will choose neither a low level of self-preferencing ($b \leq \underline{b}$) nor a high level ($b > \hat{b}$). To see how b affects the objective within the remaining regime $b \in (\underline{b}, \hat{b}]$, recall that the platform sets $p_1 = \bar{p}_1$ to deter consumer search. Anticipating this inactive search, the third-party seller responds with a price $p_2 = c + 2t + f - b^2/2s$. As b increases within this regime, the third-party seller is compelled to lower her price, and thus the platform also has to reduce the subsidiary's price to effectively deter consumer search. This suggests that users benefit from lower product prices, whereas the platform's per-user profit necessarily declines. To determine the net effect on the joint surplus, we substitute the equilibrium prices to obtain

$$S(b) = v - c - \frac{5t}{2} + \frac{b^2}{4st} \left(5t - \frac{3b^2}{4s} \right),$$

which is unimodal, reaching a unique peak at $b^\dagger \equiv \sqrt{10st/3} > \underline{b}$. Therefore, the function $S(b)$ attains a global maximum at $b = b^\dagger$, provided that $b^\dagger < \hat{b}$, put differently, the search cost s exceeds the threshold $\hat{s}$. Conversely, if s is lower than $\hat{s}$, the function $S(b)$ is increasing throughout the interval $(\underline{b}, \hat{b}]$, so the global maximum is achieved at the boundary $b = \hat{b}$. $\square$

A.7 Proof of Proposition 8

Suppose the regulator mandates a behavioral bias $b \geq 0$. As in the proof of Proposition 2, the platform's problem of optimizing its transaction fee can be reformulated as choosing the level of user surplus: $\max_U G(U) [S(b) - U]$. The optimal choice of user surplus, denoted by $U(b)$, is uniquely characterized by the first-order condition

$$\frac{G}{g}(U(b)) = S(b) - U(b) = \pi_1(b). \quad (\text{A.7})$$

Let $\pi_1(b) = \sigma(b) + f^*(b)$ represents the platform's per-user profit from optimizing the fee. Since the joint per-user surplus $S(b)$ is increasing in b for all $b \in (0, b^*)$, it follows by Proposition 2 that $U(b)$ is strictly increasing over this region, namely $U'(b) > 0$ for all $b \in (0, b^*)$.

Given the platform's response $U(b)$, the regulator's problem of choosing b can be written as

$$\max_{b \geq 0} TS(b) = \int_0^{U(b)} [W(b) - z] dG(z).$$

Differentiating $TS(b)$ with respect to b yields

$$\begin{aligned} TS'(b) &= g(U(b)) [W(b) - U(b)] U'(b) + G(U(b)) W'(b) \\ &= g(U(b)) [(\pi_1 + \pi_2)(b) U'(b) + \pi_1(b) W'(b)], \end{aligned}$$

where the second equality follows by the definition of W and the condition (A.7).

To locate the optimum b^{SB} , we evaluate the sign of this derivative at the relevant boundaries. By definition, b^0 maximizes per-user total surplus, implying $W'(b^0) = 0$. On the other hand, $b^0 < b^*$ implies $U'(b^0) > 0$. Evaluating the derivative at $b = b^0$ yields $TS'(b) > 0$, demonstrating that the regulator always has an incentive to introduce a strictly positive level of bias starting from b^0 . Conversely, for all $b \geq b^*$, $S'(b) \leq 0$ implies $U'(b) \leq 0$. Combined with the fact that $W(b)$ is monotone decreasing everywhere, we have $TS'(b) < 0$ for $b \geq b^*$. Therefore, the second-best level of self-preferencing must lie between b^0 and b^* , as claimed. $\square$

Appendix B Supplementary Material

Supplementary material related to this article can be found in the Online Appendix (not for publication).