Self-Preferencing and Welfare in Hybrid Platforms*

Bo Hu† Hongzhe Huang‡ Zonglai Kou§ Makoto Watanabe¶

March 31, 2026

Abstract

Regulatory efforts to ban *self-preferencing*—the practice of hybrid platforms favoring their own products—are gaining traction. This paper develops a model demonstrating that self-preferencing serves a strategic function beyond merely disadvantaging rivals: it prevents marketplace saturation by curbing excessive seller entry. We show that an outright ban yields ambiguous welfare outcomes. By forcing platforms to rely on a "blunter instrument"—higher commission fees—a ban may over-deter entry, ultimately reducing social welfare and harming consumers. Conversely, a ban may enhance welfare by inducing greater seller participation. Our results suggest that a ban's welfare effect is typically positive when first-party presence is low, but becomes negative as that presence increases.

Keywords: *Platforms*, *Hybrid Marketplace*, *Self-Preferencing*, *Seller Entry*

*Hu acknowledges the financial supports from the Natural Science Foundation of China (72003041), the Shanghai Pujiang Program (21PJC011), and Shanghai Institute of International Finance and Economics; Kou acknowledges the financial support from the National Social Science Fund of China (22AZD031).

†Institute of World Economy, Fudan University; hu_bo@fudan.edu.cn.

‡School of Economics, Fudan University; huang_hongzhe@outlook.com

§School of Economics, Research Institute of Innovation and Digital Economy, Fudan University; zlkou@fudan.edu.cn.

¶KIER, Kyoto University; makoto.wtnb@gmail.com.

1 Introduction

Hybrid e-commerce platforms—Amazon and JD.com are leading examples—sit at the center of a global regulatory debate. In Europe, the Digital Markets Act (DMA) explicitly bars dominant "gatekeeper" platforms from engaging in *self-preferencing*, by which is meant favoring their own first-party products over those of third-party sellers. This regulatory scrutiny is not confined to Europe: the E.C.'s case against Google and the FTC's suit against Amazon target the same phenomenon. What motivates these actions is the sharp conflict of interest inherent in the platform's dual role as both marketplace operator and first-party retailer.

Yet the platform's incentive to self-preference is constrained by a fundamental economic tension. On one hand, third-party sellers are critical: they generate variety, attract consumers, and drive marketplace revenue. On the other, these same sellers introduce cannibalization risk, competing directly with the platform's own first-party offerings. This tension motivates our core questions: When is it profit-maximizing for a hybrid platform to self-preference? What are the welfare implications of policies—like the DMA—that ban this practice?

To answer these questions, we develop a framework of hybrid intermediation where the platform acts as a gatekeeper, allocating visibility across its first-party and third-party listings. Visibility is the key battleground: the platform recommends products from its assortment, and a listing's market share depends directly on how prominently it is featured. This is not just a modeling assumption—on Amazon, sellers compete for the algorithmically managed "Buy Box," and on Booking.com, hotel rankings are a primary driver of revenue.

The platform manages its marketplace through two levers. First, it sets a commission fee on third-party sales, which introduces the classic double marginalization problem. Second, it controls listing visibility. By self-preferencing we mean the platform's choice to lower the relative visibility of third-party listings compared to its own first-party products. These two instruments jointly determine the profitability of third-party participation, which is governed by a free-entry condition. Self-preferencing thus serves as a strategic instrument: it allows the platform to manage the entry margin—to curb "excessive" entry that cannibalizes first-party sales, or to fine-tune marketplace composition in response to its commission fee.[1]

We first show that self-preferencing is not an inherent feature of hybrid intermediation. To

[1]Our modeling approach allows us to cleanly isolate the platform's curation power from direct price rivalry. We thus set aside a pricing channel where a platform might raise commission fees to induce higher third-party prices, thereby diverting demand to its own first-party offerings. While both mechanisms may operate in practice, our focus is highly relevant to the contemporary debate. Widespread complaints from third-party sellers—who report being pushed out of critical visibility spots like Amazon's "Buy Box" by the platform's own identical products—suggest that the direct, algorithmic competition for visibility is a primary and contentious front in the battle against self-preferencing.

see why, consider the platform's problem in the absence of third-party entry constraint—suppose it could freely choose the number of third-party listings. The platform would select a profit-maximizing number, trading off the benefits of additional transactions against the cost of shifting the sales mix away from its first-party offerings. This trade-off works as follows: more third-party sellers generate more transactions, which is good; but a larger third-party presence also cannibalizes first-party sales, which is bad.

Now reintroduce third-party entry constraint, where sellers pay a cost to join the platform and trade. We say that "excessive entry" occurs when the equilibrium number of sellers under free entry condition exceeds the profit-maximizing number. It is precisely in this scenario that the platform actively self-preferences its first-party listings. Absent excessive entry, the platform treats third-party and first-party listings equally. We show that excessive entry is likely to arise when the platform's first-party presence is large, total market demand is low, or seller entry costs are small.

Then, why does the platform not simply deter entry by raising its commission rate? It could, but doing so would exacerbate double marginalization on all remaining sales. Self-preferencing thus emerges as the preferred tool: it manages entry without distorting retail prices. More precisely, it allows the platform to decouple two distinct objectives—setting the commission rate to optimize the fee revenue of each trade while separately using visibility control to fine-tune seller quantity. Therefore, banning self-preferencing removes this decoupling instrument, forcing the platform onto a single, blunter tool: high commission fees. But this lever is less efficient, and may trigger a surprising consequence we call *over-deterrence,* i.e., the fees chosen by the platform to contain entry can be so high that the equilibrium number of sellers actually falls below what the platform would have under self-preferencing—the cure, in other words, leaves fewer sellers than the disease.

We evaluate the welfare consequences of banning self-preferencing and find it to be ambiguous. To illustrate, consider a scenario where the platform is prone to listing too few sellers relative to the social optimum.[2] A ban in this case forces the platform to raise its commission fee to manage seller entry. Two effects follow. The *fee effect* is unambiguously negative: higher commissions intensify double marginalization, raising retail prices and harming consumers. But there is also a third-party participation effect or a *trade volume effect*: the fee increase alters how many sellers find entry worthwhile, and hence the volume of trade. When first-party presence is low, the plat-

[2] This scenario arises when the revenue-to-consumer-surplus ratio is lower for third-party sales than for first-party sales, so that the platform internalizes a smaller share of the gains from third-party trade.

form raises fees only modestly, so the trade volume effect runs in the favorable direction—more sellers enter, expanding market trade volume, and overall welfare can improve. But when first-party presence is high, the platform has more at stake and raises fees aggressively. This triggers over-deterrence: third-party participation actually falls relative to the self-preferencing equilibrium. Both effects now cut in the same direction, and the ban unambiguously reduces welfare. The policy implication is clear: whether to ban self-preferencing depends on the trade-off between the fee effect—higher commissions after a ban—and the third-party participation, which may either increase or decrease market trade.

Finally, we also consider whether a lump-sum participation fee could substitute for self-preferencing. We find that the lump sum fee is not only a perfect substitute for achieving a target number of sellers but also a superior instrument from both a profit and welfare perspective. Self-preferencing is wasteful because some sellers incur entry costs but are ultimately never displayed to buyers. A participation fee, in contrast, is an efficient transfer that ensures only the desired number of sellers enter in the first place. The superiority of fees suggests that the real-world prevalence of self-preferencing may stem from practical or contractual barriers that prevent platforms from implementing direct entry pricing (see discussions, e.g., Wang and Wright 2017).

Related Literature

Our paper contributes to the growing literature on platform self-preferencing (e.g., de Cornière and Taylor 2019; Hagiu, Teh and Wright 2022; Zennyo 2022; Gautier, Hu and Watanabe 2023; Anderson and Defolie 2024; Dendorfer 2024; Sato and Kittaka 2024; Chi, Choi, Hahn and Kim 2025; Hervas-Drane and Shelegia 2025). Much of this work treats the population of third-party sellers as fixed. Under that assumption, banning self-preferencing triggers a single, well-understood effect: the platform raises its commission fee, intensifying double marginalization and harming consumers. The welfare prescription is then, self-preferencing, by keeping fees low, is the lesser evil.

We take seller participation under a free-entry condition, and its welfare consequences, seriously. We concur with prior work that a ban on self-preferencing raises commission fees—but we continue one step further and ask: what does this fee increase do to the volume of trade? Higher fees reduce seller profits, which in turn reduces entry. The answer, we show, depends on the platform's first-party presence: when it is large, the fee response is aggressive and over-deterrence harms welfare; when it is modest, fees rise only slightly and welfare can improve. This is what makes the welfare consequences genuinely ambiguous.

This insight connects our work to a broader tradition in industrial organization on the efficiency of free entry. Mankiw and Whinston (1986) and Anderson, De Palma and Nesterov (1995) show that market entry can be socially excessive when entrants steal business from incumbents without creating commensurate value. Our framework extends this logic to platforms: self-preferencing emerges as a curation tool that manages the entry margin, curbing excessive participation when variety gains are outweighed by cannibalization of first-party sales.

Finally, it is worth noting that our analysis is commitment-proof. In several related models (e.g., Dendorfer 2024), the platform faces a time-inconsistency problem: it would like to promise no self-preferencing ex ante to attract sellers, only to deviate ex post once they have sunk their entry costs. Our framework sidesteps this issue. The platform's profit depends on the number of displayed sellers—not on the total pool of entrants beyond those recommended. There is no ex-post temptation to deviate, which allows us to isolate the strategic trade-offs of self-preferencing cleanly.

The paper proceeds as follows. Section 2 lays out the baseline model of hybrid intermediation without self-preferencing. Section 3 introduces self-preferencing and derives our key findings. We then evaluate the welfare implications of banning this practice in Section 4. Section 5 explores several extensions, and Section 6 concludes. All proofs are provided in the Appendix.

2 A Basic Model of Hybrid Intermediation

2.1 Set-ups

Consider a large economy with a measure b of buyers and a large pool of potential third-party sellers. The measure of sellers participating in trade, s, is determined by a free-entry condition where each entrant incurs a cost $k \in (0, \bar{k})$.[3] Trade revolves around products that are potentially horizontally differentiated. Each buyer has unit demand, and each seller can produce their variant at a constant marginal cost of c.

A monopolistic platform serves as the exclusive venue for trade. The platform hosts third-party sellers and charges an ad valorem commission fee $\gamma \in (0, 1)$ on their sales revenue. The analysis of other fee structures is deferred to later sections. The platform also acts as a retailer, operating a measure $h > 0$ of its own first-party sellers. We treat this first-party capacity, h, as an

[3] See the proof of Proposition 3 for the definition of $\bar{k}$.

exogenous parameter. These first-party sellers also produce at a constant marginal cost c. Thus, the platform operates as a hybrid marketplace.

The platform leverages a recommendation algorithm to match sellers with buyers based on partial information regarding consumer preferences. The total measure of successful recommendations is governed by a function, $M(b, h+s)$, given b buyers and a total of $h+s$ sellers. This, in turn, defines the measure of successful recommendations for each seller as $\mu^s(b, h+s) \equiv \frac{M(b,h+s)}{h+s}$. We avoid making functional form assumptions, and the modeling of $M(\cdot)$ allows for several independent interpretations.[4]

Example 1: $M(\cdot)$ is a general matching function. *It is common that $M(\cdot)$ is a general matching function in the search and matching literature. In this context, $M(\cdot)$ serves as an abstraction of the underlying search process, capturing its aggregate outcomes while remaining agnostic to the specific mechanism. Standard regularity conditions include that $M(\cdot)$ is strictly increasing, concave in arguments, and exhibits constant returns to scale. This general specification is analytically convenient and nests several widely-used functional forms, such as urn-ball and CES. A key economic implication of this approach is an environment of ex-ante product homogeneity, and price discovery occurs only after a match is established.*

Example 2: $M(\cdot)$ is micro-founded as the platform recommendation function. *Consider a space of N potential product variants where the platform hosts H first-party and S third-party sellers, with each offering a unique variant. Each buyer has a consideration set containing ω variants from which they derive positive utility. The platform's recommendation algorithm has an efficiency of $\zeta \in (0,1]$. A buyer is successfully matched with probability ζ if their consideration set contains at least one variant offered by any of the $H+S$ sellers; otherwise, the match probability is zero. The probability that a single, randomly chosen variant is offered by a seller is $(H+S)/N$. The probability of a buyer finding at least one match is therefore $\zeta\left(1-(1-(H+S)/N)^{\omega}\right)$. Taking the large market limit where $N, H, S \to \infty$ such that the ratios of sellers to variants converge to constant measures, $H/N \to h$ and $S/N \to s$. Each buy's matching probability becomes $\mu^b(h+s, \omega, \zeta) = \zeta\left(1-e^{-(h+s)\omega}\right)$. The total measure of matches is then $M(b, h+s) = b \cdot \mu^b$, which is increasing in both b and $h+s$ and is strictly concave in $h+s$. Moreover, M is higher when the platform algorithm is more efficient (ζ is higher) and consumers have a larger scope of interest (ω is higher). The matching probability for sellers are $\mu^s(b, h+s) = \frac{b\zeta\left(1-e^{-(h+s)\omega}\right)}{h+s}$.*

Example 3: $M(\cdot)$ is derived from the Van Zandt (2004) and Johnson (2013). *There are b buyers and $h+s$ listings. For any given buyer-listing pair, match probability q is an i.i.d. random variable drawn*

[4] In our model, the platform's recommendation objective is to maximize the number of successful matches between buyers and sellers. Therefore, price is not a factor in the recommendation process. Moreover, in the analysis that follows, $M(\cdot)$ will be considered exogenous.

from a cdf $F(q)$ on $[0,1]$. The platform observes q realizations and recommends the listing with highest match probability for each consumer: $q_{max} = \max_i\{q_i\}$. The ex-ante match probability for a buyer is $\mathbb{E}[q_{max}] = 1 - \int_0^1 [F(q)]^{h+s} dq$ [5] The total volume of successful recommendations is $M(b, h+s) = b \cdot \left(1 - \int_0^1 [F(q)]^{h+s} dq\right)$, which is increasing in the number of listings $(h+s)$. The matching probability for sellers, defined as $\mu^s(b, h+s) = \frac{b \cdot \left(1 - \int_0^1 [F(q)]^{h+s} dq\right)}{h+s}$, is decreasing in $h+s$.

The timing of the game is as follows: First, the platform announces the commission rate γ. Observing γ, potential third-party sellers simultaneously decide whether to enter the market. Each sellers set their prices p, and the platform simultaneously use its algorithm to to recommend goods from sellers to buyers. Finally, Conditional on a successful recommendation, a buyer generates a demand of $D(p)$. The solution concept is Subgame Perfect Nash Equilibrium.[6]

For the rest of the paper, we normalize $b = 1$ for simplicity. This allows us to express both $M(\cdot)$ and $\mu^s(\cdot)$ solely as functions of the measure of sellers.

2.2 Equilibrium

We solve the model by working backward from the final stage, beginning with the sellers' post-matching pricing decisions. A first-party seller acts as a monopolist for a given match and chooses a price to maximize its profit:

$$\pi_m = \max_p (p - c)D(p). \tag{1}$$

A third-party seller, taking commission γ as given, maximizes its profits:

$$\pi_{ss}(\gamma) = \max_p [p(1-\gamma) - c]D(p). \tag{2}$$

By the Envelope Theorem, $\pi_{ss}(\gamma)$ strictly decreases in γ. There exists $\bar{\gamma} \leq 1$ such that $\pi_{ss}(\gamma) = 0$ for $\gamma \geq \bar{\gamma}$.[7] Denote the profit-maximizing price by $p_s(\gamma)$. Intuitively, $p_s(\gamma)$ decreases in γ.

From each successful third-party transaction, the platform extracts a revenue of

$$\pi_{sm}(\gamma) = \gamma p_s(\gamma) D(p_s(\gamma)). \tag{3}$$

[5] Since $F(1) = 1$ and $F(0) = 0$, we obtain $\mathbb{E}[q_{max}] = \int_0^1 q\, dF(q)^{h+s} = [qF(q)^{h+s}]_0^1 - \int_0^1 [F(q)]^{h+s}\, dq = 1 - \int_0^1 [F(q)]^{s+h}\, dq$.

[6] Consider an alternative timing where third-party sellers set their price immediately upon entering the market, rather than after matching. In such a setting, we must impose that the platform's recommendation function is price-independent. This assumption establishes our benchmark case (no self-preferencing), in which the platform treats all products equally and assigns them the same recommendation weight. This benchmark stands in contrast to the self-preferencing case analyzed in Section 3, where the platform strategically introduces bias to maximize its own profits.

[7] A seller can make a positive profit as long as they can find a price such that the effective margin is positive, $p > c/(1-\gamma)$.

There are two observations. First, the effect of the commission rate γ on the platform's revenue is determined by two opposing forces. On one hand, as γ increases, the platform captures a larger share of sales revenue. Concurrently, this increase in γ raises sellers' marginal cost, causing them to raise their price. This, in turn, dampens buyer demand, ultimately reducing sales revenue $p_s(\gamma)D(p_s(\gamma))$. In the following, we assume $\gamma p_s(\gamma)D(p_s(\gamma))$ is single-peaked in γ. We then define $\hat{\gamma}$ as the rate that maximizes the platform's per-match fee revenue:

$$\hat{\gamma} \equiv \arg\max_{\gamma} \gamma p_s(\gamma) D(p_s(\gamma)) \in (0, \bar{\gamma}). \tag{4}$$

The entry of sellers. Given first-party capacity h and commission rate γ, the equilibrium measure of third-party sellers, denoted by $s_0(\gamma)$, is determined by the free-entry condition

$$\mu^s\,(b, s+h)\,\pi_{ss}(\gamma) = k, \tag{5}$$

where the left-hand side is the expected profit of a third-party seller participating in the platform, and the right-hand side is the entry cost. Since the sellers' profits $\pi_{ss}(\gamma)$ decrease as γ increases, it follows that $s_0(\gamma)$ also decreases in γ:[8]

$$\frac{\partial s_0(\gamma)}{\partial \gamma} < 0.$$

Platform profit maximization. The platform chooses its commission fee γ to maximize total profits subject to the free-entry condition:

$$\max_{\gamma \in [0,1]} \mu^s\,(b, h+s)\,\Big(h\pi_m + s\pi_{sm}(\gamma)\Big), \text{ s.t. } s = s_0(\gamma),$$

The objective is the sum of revenues from the platform's first-party and third-party sales. The choice of γ affects this profit in two ways: it *directly* controls the revenue extracted from each third-party seller, $\pi_{sm}(\gamma)$, and it *indirectly* determines the equilibrium number of sellers entering the platform, $s_0(\gamma)$. This creates a fundamental tension. To preserve a large seller base, the platform must choose a γ that deviates from $\hat{\gamma}$-the rate that would maximize revenue from an individual seller match.

To build intuition, we proceed in two steps. First, we consider a hypothetical, unconstrained scenario where the platform can choose both the commission γ and the number of sellers s directly, thus bypassing the trade-off between extraction and participation. Second, we solve the

[8] There exists $\bar{\bar{\gamma}} \in (0, \bar{\gamma}]$ such that $s_0(\gamma)$ is strictly decreasing in γ with $s_0(0) > 0$ and $s_0(\gamma) = 0$ for $\gamma \in [\bar{\bar{\gamma}}, 1]$.

actual constrained problem, where s is endogenously set by seller entry. By comparing the two results, we can clearly see how the platform must adjust its commission rate to account for equilibrium entry.

2.2.1 The unconstrained platform problem

In the unconstrained problem, since γ only appears in the third-party profit term $\pi_{sm}(\gamma)$, the optimal commission is simply $\hat{\gamma}$ that maximizes $\pi_{sm}(\gamma)$. For any given γ (including the special case $\gamma = \hat{\gamma}$), the platform's problem over s is then to maximize the total number of matches multiplied by its average profit per match:

$$\Pi(\gamma, s) \equiv M(b, h+s)\left(\frac{h}{h+s}\pi_m + \frac{s}{h+s}\pi_{sm}(\gamma)\right)$$

An increase in s generates a positive matching effect, as the total number of matches rises $\partial M/\partial s > 0$, while also introducing a negative weighting effect by shifting the sales mix toward less profitable third-party sellers. Since the third-party sales are less profitable than the first-party with $\pi_{sm} < \pi_m$ holds for all γ, this shifts down the average profit per match. This trade-off implies a unique, interior optimum if the profit function is well-behaved. Specifically, we assume the objective is single-peaked in s. A sufficient condition is the elasticity of the matching function, $\varepsilon_M(s) \equiv \frac{\partial M}{\partial(h+s)}\frac{h+s}{M}$, is sufficiently decreasing in s.[9]

Lemma 1 *Assume that the platform's profits are single-peaked with respect to s[10]. Let $\hat{s}(\gamma)$ denote the unique profit-maximizing measure of third-party sellers. Then, $\hat{s}(\gamma)$ is single-peaked in γ. When $\hat{s}(\gamma)$ is interior, it is increasing for $\gamma < \hat{\gamma}$ and decreasing for $\gamma > \hat{\gamma}$.*

The intuition for why $\hat{s}(\gamma)$ is single-peaked in γ is as follows. The profit-maximizing number of sellers, $\hat{s}(\gamma)$, increases monotonically with the platform's profit from each third-party seller, $\pi_{sm}(\gamma)$. This is because a higher $\pi_{sm}(\gamma)$ narrows the margin gap between first-party and third-party sales, diminishing the negative weighting effect. When $\gamma < \hat{\gamma}$, increasing γ raises $\pi_{sm}(\gamma)$: as the negative weighting effect weakens, the platform has an incentive to add more third-party sellers, causing $\hat{s}$ to rise. In contrast, when $\gamma > \hat{\gamma}$, increasing γ lowers $\pi_{sm}(\gamma)$: as the negative weighting effect strengthens, the platform reduces the number of third-party sellers, leading $\hat{s}$ to fall. In short, the shape of $\hat{s}(\gamma)$ with respect to γ directly mirrors that of $\pi_{sm}(\gamma)$—which is

[9] If $\varepsilon'_M(s) < -\frac{h\left(\frac{\pi_m}{\pi_{sm}(\hat{\gamma})}-1\right)}{\left(h\frac{\pi_m}{\pi_{sm}(\hat{\gamma})}+s\right)^2} < 0$, then the platform profit is single-peaked with respect to s.

[10] For the matching probability specified in Example 2, this is a derived result, not an assumption. See Appendix for the formal proof.

assumed to be single-peaked, with its maximum at $\hat{\gamma}$.[11]

Finally, note that the platform is unconstrained in choosing s and γ separately. The solution to the unconstrained platform problem is therefore:

$$(\gamma^*, s^*) = (\hat{\gamma}, \hat{s}(\hat{\gamma})).$$

2.2.2 The constrained problem with free entry of sellers

We reintroduce the free-entry constraint, $s = s_0(\gamma)$. Now the number of sellers becomes endogenously determined by the platform's commission fee. The platform's problem thus reduces to choosing γ to maximize its profit:

$$\Pi_{nsp}(\gamma) \equiv \Pi(\gamma, s_0(\gamma)) = M\left(b,\ h + s_0(\gamma)\right)\left(\frac{h}{h + s_0(\gamma)}\pi_m + \frac{s_0(\gamma)}{h + s_0(\gamma)}\pi_{sm}(\gamma)\right)$$

Unlike the unconstrained problem, the platform can no longer choose the number of sellers s directly. Instead, it uses γ as an instrument to steer seller entry toward its preferred level, which may require setting a fee γ_{nsp} that deviates from the per-seller fee-maximizing fee $\hat{\gamma}$.

The trade-off is characterized by the first-order condition, which decomposes the marginal impact of a fee change into two effects:

$$\left.\frac{d\Pi_{nsp}}{d\gamma}\right|_{\gamma=\gamma_{nsp}} = \underbrace{\left.\frac{\partial \Pi_{nsp}}{\partial \pi_{sm}}\frac{\partial \pi_{sm}(\gamma)}{\partial \gamma}\right|_{\gamma=\gamma_{nsp}}}_{\text{Revenue Effect}} + \underbrace{\left.\frac{\partial \Pi_{nsp}}{\partial s_0}\frac{\partial s_0(\gamma)}{\partial \gamma}\right|_{\gamma=\gamma_{nsp}}}_{\text{Entry Effect}} = 0 \tag{6}$$

The *revenue effect* captures how a change in γ affects the per-seller fee revenue, $\pi_{sm}(\gamma)$. As established previously, this effect is positive when $\gamma < \hat{\gamma}$ and negative when $\gamma > \hat{\gamma}$. The *entry effect* captures the indirect impact on total profit via the change in the number of sellers, $s_0(\gamma)$. This effect is a product of two terms: how the fee affects entry ($\frac{\partial s_0}{\partial \gamma} < 0$, as higher fees deter entry), and whether the platform benefits from more or fewer sellers at the margin ($\frac{\partial \Pi_{nsp}}{\partial s_0}$). The sign of $\frac{\partial \Pi_{nsp}}{\partial s_0}$ is ambiguous, depending on whether the positive matching effect dominates the negative weighting effect or not.

Figure 1 illustrates this trade-off. In each panel, the downward-sloping gray curve represents the number of sellers who enter, $s_0(\gamma)$, while the single-peaked red curve is the platform's preferred number of sellers, $\hat{s}(\gamma)$. The platform's decision hinges on the condition at the unconstrained optimum $\gamma = \hat{\gamma}$, where the revenue effect is zero. So the platform's incentive to adjust its fee is driven entirely by the entry effect.

[11] Since $\pi_{sm}(0) = \pi_{sm}(\bar{\gamma}) = 0$, we must have $\hat{s}(0) = \hat{s}(\bar{\gamma}) = 0$, and $\hat{s}(\gamma) = 0$ for $\gamma \in [\bar{\gamma}, 1]$.

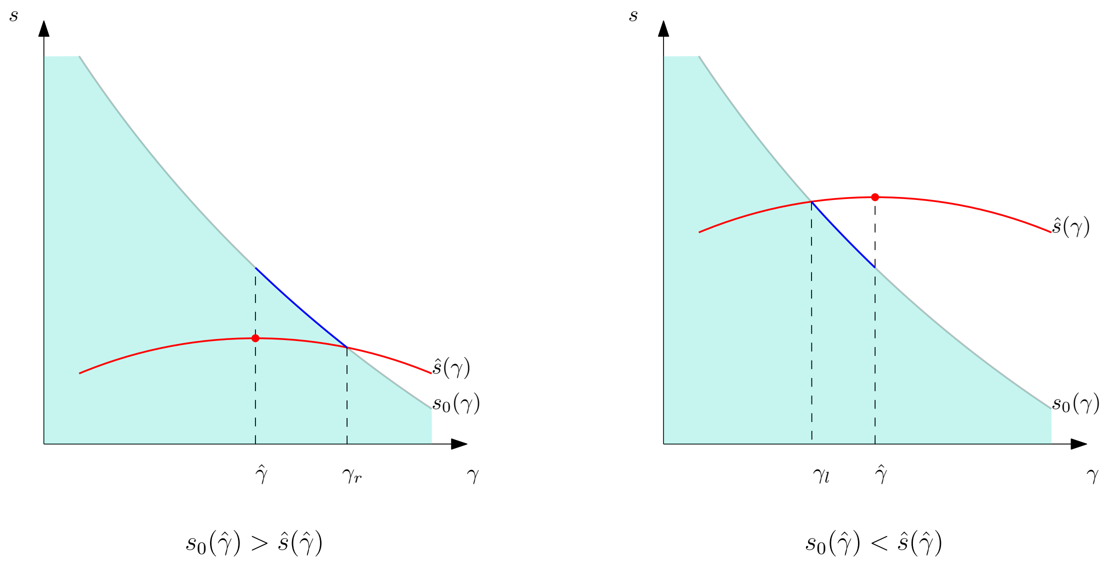

Figure 1: Optimal constrained fee under excessive entry (left) and insufficient entry (right).

If, at $\gamma = \hat{\gamma}$, the platform faces *excessive entry* at $\hat{\gamma}$, i.e., $s_0(\hat{\gamma}) > \hat{s}(\hat{\gamma})$, the negative weighting effect dominates ($\frac{\partial \Pi_{nsp}}{\partial s_0} < 0$). This incentivizes the platform to raise its fee above $\hat{\gamma}$, sacrificing per-seller fee revenue to drive out excessive sellers and improve the average profit of matches. As shown in the left panel, the optimal fee $\gamma_{nsp} > \hat{\gamma}$ lies on the blue section of the $s_0(\gamma)$ curve. Note that γ_{nsp} is bounded above by γ_r, where γ_r is the rate that turns excessive entry into insufficient entry.[12] If the platform overshoots and sets a fee that is even higher than γ_r, it creates a new problem of scarcity of sellers and unnecessarily hurts revenue.

If, at $\gamma = \hat{\gamma}$, the platform faces *insufficient entry*, i.e., $s_0(\hat{\gamma}) < \hat{s}(\hat{\gamma})$, the positive matching effect dominates ($\frac{\partial \Pi_{nsp}}{\partial s_0} > 0$). The platform will then be better off lowering its fee below $\hat{\gamma}$. As depicted in the right panel, the platform accepts a lower per-seller fee revenue to attract more sellers and increase total transaction volume. Again, the platform will not set a fee that is lower than γ_l, which would then create a new problem of excessive seller entry.[13]

The results are summarized in the following proposition.

Proposition 1 (Optimal Fee Under Free Entry) *The platform's optimal fee γ_{nsp} under free entry of sellers is characterized by (6) and is bounded by γ_l and γ_r. If $s_0(\hat{\gamma}) > \hat{s}(\hat{\gamma})$ (excessive entry), $\gamma_{nsp} \in (\hat{\gamma}, \gamma_r]$. If $s_0(\hat{\gamma}) < \hat{s}(\hat{\gamma})$ (insufficient entry), $\gamma_{nsp} \in [\gamma_l, \hat{\gamma})$.*

[12] Since $\hat{s}(\bar{\bar{\gamma}}) \geq 0 = s_0(\bar{\bar{\gamma}})$(where $\bar{\bar{\gamma}}$ is defined in footnotes 8), $s_0(\gamma)$ and $\hat{s}(\gamma)$ must intersect at least once on the right of $\hat{\gamma}$. We let $\gamma_r = \min(\{\gamma \in (\hat{\gamma}, \bar{\bar{\gamma}}] \mid s_0(\gamma) = \hat{s}(\gamma)\}$.

[13] γ_l, is defined as the solution to $s_0(\gamma) = \hat{s}(\gamma)$ for $\gamma \in [0, \hat{\gamma})$. Since $s_0(0) > 0 = \hat{s}(0)$, γ_l must exist.

3 Self-Preferencing

In this section, we allow the platform to engage in self-preferencing. Previously, the commission rate γ was a blunt tool used to influence both per-seller revenue and the equilibrium number of entrants. Now, the platform can use self-preferencing to directly control the number of third-party sellers to be displayed.

3.1 Self-preferencing

We model self-preferencing as the platform's choice of a recommendation weight $\alpha \in [0,1]$ applied to third-party sellers[14]. In the first stage, the platform chooses the commission rate γ and the weight α simultaneously. Let s_E be the measure of third-party sellers who enter the market. The weight α implies an effective matching pool of $h + \alpha s_E$, which defines the matching function as $M(b, h + \alpha s_E)$. The free entry condition is

$$\mu^s\left(b, h + \alpha s_E\right) \pi_{ss}(\gamma) \cdot \alpha = k.$$

This weighting mechanism is equivalent to that only an effective mass (or proportion) of sellers, $s_{sp} = \alpha s_E$, actually participates in the recommendation process. For a potential entrant, the expected profit is the profit earned if visible, $\mu^s\left(b, h + s_{sp}\right) \pi_{ss}(\gamma)$, multiplied by the probability of being visible, $\frac{s_{sp}}{s_E}$. This yields the new free entry condition for the mass of entering sellers, s_E:

$$\mu^s\left(b, h + s_{sp}\right) \pi_{ss}(\gamma) \cdot \frac{s_{sp}}{s_E} = k. \tag{7}$$

This implies that the weighting self-preferencing strategy can be reframed as the platform determining the mass of sellers s_{sp} to recommend to buyers. The following lemma identifies the feasible set for s_{sp} and establishes the relationship between the number of displayed sellers (s_{sp}), entered sellers (s_E), and the entry level from our benchmark model ($s_0(\gamma)$).

Lemma 2 *For any given commission rate γ, the feasible set for the platform's choice of displayed sellers is $s_{sp} \in [0, s_0(\gamma)]$ and s_E satisfies $s_E \leq s_0(\gamma)$. Furthermore, $s_E = s_0(\gamma)$ if and only if $s_{sp} = s_0(\gamma)$.*

A key takeaway from the lemma is that the ability to curate the marketplace comes at no cost to the platform's choice set. A reader might expect that chilled entry ($s_E \leq s_0(\gamma)$) would restrict the platform's options, but this is not the case. The platform retains the capability to achieve any outcome up to the benchmark level, as it can always trigger the full entry of $s_0(\gamma)$ sellers by

[14] When $\alpha = 1$, the platform does not engage in self-preferencing, as all third-party sellers are given full recommendation weight. This case corresponds to our no self-preferencing benchmark.

simply committing to display them all. Thus, self-preferencing gives the platform a superior tool to select its ideal number of displayed sellers, s_{sp}, without forfeiting any of its original feasible outcomes.

Discussion. Our model focuses on platform-controlled visibility. We thus set aside an important pricing channel where a higher commission γ might be passed through to consumers as a higher third-party price, potentially diverting demand to the platform's own first-party offerings, see e.g., Anderson and Defolie (2024). The analytical merit of this simplification is that it allows a clean isolation of the search and curation channel, where the platform directly manages which (set of) sellers buyers discover. This focus is reasonable and grounded in the operations of many dominant platforms where competition for visibility precedes and often outweighs direct price wars.

Consider the example of Amazon's Buy Box. On a typical Amazon product page, over 80% of sales go through the "Buy Box" — the default "Add to Cart" option. Winning the Buy Box is therefore the primary battle for sellers of the same product. While price is a factor, Amazon's proprietary algorithm awards the Buy Box based on a host of non-price variables, including fulfillment method (Fulfillment by Amazon is heavily favored), shipping speed, seller rating, and inventory levels. A seller with a higher price but superior fulfillment and ratings will often win the Buy Box over a cheaper rival. This demonstrates a market where the key competition is for algorithmic approval to gain visibility, justifying our focus on platform curation.

Another example is Booking.com's Hotel Rankings. When a user searches for a hotel, the platform presents a ranked list that is not sorted by price by default. A hotel's position in this ranking is a critical determinant of its sales and is influenced by a complex algorithm. Key factors include the commission rate the hotel pays Booking.com, its historical conversion rate on the platform, guest review scores, and participation in programs like "Preferred Partner." A hotel's strategic priority is often to improve its algorithmic ranking to appear on the first page, even if it means paying higher fees or offering amenities, rather than simply being the cheapest option. This again highlights an environment where the battle for platform-curated visibility is the central competitive dynamic. In both examples, the platform acts as a powerful gatekeeper, shaping competition through its control of visibility. Our framework is therefore well-suited to analyzing the strategic implications of this curation power, a central feature of the modern digital economy.

3.2 Platform profit maximization

We now analyze the platform's optimal strategy when self-preferencing is an available, but not mandatory, choice. The game proceeds sequentially: the platform sets a commission rate γ, third-party sellers enter, and finally, the platform chooses how many entrants to display, s_{sp}. Backward induction simplifies this to the platform choosing γ and s_{sp} to maximize profit, subject to the entry constraint established in Lemma 2. The platform's optimization problem is therefore:

$$\max_{s\in[0,\bar{s}],\gamma\in[0,1]} \Pi(s,\gamma) \quad \text{s.t. } s \leq s_0(\gamma)$$

The readers shall recognize that this optimization problem is the same as the unconstrained problem of Section 2, except that the platform is subject to the free-entry condition. The solution is characterized in the following proposition.

Proposition 2 (Optimal Self-Preferencing Strategy) *The platform's decision to self-preference depends on the level of entry at $\hat{\gamma}$.*

1. *Excessive Entry ($s_0(\hat{\gamma}) > \hat{s}(\hat{\gamma})$): It is strictly optimal for the platform to self-preference. The platform achieves its unconstrained maximum profit by setting $\gamma_{sp} = \hat{\gamma}$ and displaying $s_{sp} = \hat{s}(\hat{\gamma})$.*
2. *Insufficient Entry ($s_0(\hat{\gamma}) < \hat{s}(\hat{\gamma})$): It is optimal for the platform to **not** actively self-preference. The outcome is identical to the no-self-preferencing benchmark.*

This proposition highlights that self-preferencing is a tool the platform uses only when facing an excess of third-party sellers. In the *excessive entry* case, self-preferencing becomes a powerful tool that allows the platform to solve two problems with two instruments. It sets the commission rate $\gamma = \hat{\gamma}$ to maximize revenue from each sale, and it then uses self-preferencing to reduce the number of displayed sellers from the excessive level $s_0(\hat{\gamma})$ down to its ideal level, $\hat{s}(\hat{\gamma})$. By separating these objectives, the platform achieves its unconstrained profit.

In the *insufficient entry* case, the platform's primary problem is attracting more sellers, not fewer. Using self-preferencing to further reduce the number of displayed sellers would only lead to lower profits. The platform, therefore, forgoes this option, and the outcome is identical to the no-self-preferencing benchmark.

It is worth mentioning that a key methodological feature of our model is its time-consistency. The platform's profit function, $\Pi(\gamma, s_{sp})$, depends on the commission rate (γ) and the mass of display sellers (s_{sp}), not the total mass of entrants (s_E). The mass of entrants s_E is an equilibrium outcome that adjusts to satisfy the sellers' free-entry condition, but it does not directly enter the platform's profit calculation. In the relevant "excessive entry" case, the platform's unconstrained optimal number of displayed sellers, $\hat{s}(\hat{\gamma})$, is always feasible, as the market will provide a sufficient mass of entrants ($s_E \geq \hat{s}(\hat{\gamma})$). Because the platform directly chooses the variable (s_{sp}) that maximizes its profit, its ex-ante optimal plan is identical to its ex-post action, and it has no incentive to deviate.

3.3 Determinants of self-preferencing

Proposition 2 shows that self-preferencing is triggered by excessive seller entry. We now investigate the underlying conditions that creates this entry gap $s_0(\hat{\gamma}) - \hat{s}(\hat{\gamma})$.

First-party capacity. The primary driver of excessive entry is the platform's own retail presence. The following proposition demonstrates that as the platform's first-party business grows, a conflict of interest that incentivizes self-preferencing becomes inevitable.

Proposition 3 *The entry gap, $s_0(\hat{\gamma}) - \hat{s}(\hat{\gamma})$, which triggers self-preferencing when it is positive, increases with the mass of first-party sellers h, provided that both $s_0(\hat{\gamma})$ and $\hat{s}(\hat{\gamma})$ are positive. Consequently, with k sufficiently small, there exists a threshold $\tilde{h}$, such that for $h \in (\tilde{h}, \bar{h})$, self-preferencing is profit-maximizing.*

The intuition stems from a "race" between the number of sellers the market supplies (s_0) and the number the platform desires ($\hat{s}$). As the platform's first-party size (h) increases, it becomes more protective of its own sales. While more first-party competition naturally reduces third-party entrants (s_0 decreases), the platform's ideal number of third-party sellers ($\hat{s}$) shrinks even faster because the negative *weighting effect* becomes more severe. This widening gap means that as a platform's own presence grows, it transitions from a neutral facilitator into a direct competitor with its own sellers. The platform's first-party market share thus becomes a simple yet powerful heuristic for monitoring anticompetitive risk.

Other market conditions. In a similar vein, self-preferencing becomes more likely when the mass of buyers is small or when seller entry costs are low.

Corollary 1 *The entry gap, $s_0(\hat{\gamma}) - \hat{s}(\hat{\gamma})$, which triggers self-preferencing, decreases with the the mass of buyers, b, and seller entry cost, k, provided that both $s_0(\hat{\gamma})$ and $\hat{s}(\hat{\gamma})$ are positive. Consequently, self-preferencing is a strictly profit-maximizing strategy when b and k are sufficiently small.*

The intuition is as follows. An increase in the number of buyers b makes the market more attractive, causing a flood of third-party entrants (s_0 increases). However, the platform's demand for sellers to serve this larger market increases even more rapidly. This narrows the entry gap, making curation less necessary.

A higher entry cost k acts as a natural barrier to entry, reducing the number of third-party sellers, s_0. The platform's ideal number of sellers, $\hat{s}$, is unaffected by this external cost. A higher entry cost, therefore, shrinks the gap by solving the excessive entry problem on the platform's behalf, making self-preferencing a less necessary strategy.

4 Welfare Analysis

The platform's pursuit of profits is not necessarily aligned with the interests of society. To pinpoint the sources of inefficiency, In this section, we derive the social optimal allocation and pinpoint the sources of inefficiency. We then use this framework to evaluate the welfare consequences of a ban on self-preferencing. We conduct this evaluation using both social welfare and consumer surplus.

4.1 Social Welfare Optimality and Distortions

Consider a social planner who can set the commission fee γ and select the number of displayed third-party sellers to maximize welfare, subject to the seller's free-entry condition. A first-party match generates total surplus $w_m \equiv \pi_m + cs_m$, where cs_m is the consumer surplus at the first-party price p_m. A third-party match generates total surplus $w_s(\gamma) \equiv \pi_{ss}(\gamma) + \pi_{sm}(\gamma) + cs_s(\gamma)$, where $cs_s(\gamma)$ is the consumer surplus at third-party price $p_s(\gamma)$. The planner's objective is to maximize the social welfare:

$$W(\gamma, s) = \mu^s\left(b, h+s\right)\left\{hw_m + s\Big(\pi_{ss}(\gamma) + \pi_{sm}(\gamma) + cs_s(\gamma)\Big)\right\} - s_E k,$$

which is the total trade surplus minus the entry costs of sellers $s_E k$ with s_E the measure of sellers that enter the platform.

The free-entry condition dictates that third-party sellers enter the market until their expected profit is exactly offset by their entry costs. This drives the net surplus of third-party sellers to zero. The objective then reduces to the combined surplus of consumers and the platform (i.e., first-party sales and commission revenue):

$$\max_{\gamma,s} \quad \mu^s\left(b, h+s\right)\left\{hw_m + s\Big(\pi_{sm}(\gamma) + cs_s(\gamma)\Big)\right\} \quad \text{s.t.} \quad s \leq s_0(\gamma).$$

The planner's problem can be broken down into two parts: setting the optimal commission and choosing the optimal number of sellers.

First, consider the commission. The socially optimal commission rate, denoted γ_w, must be lower than the platform's private optimum $\hat{\gamma}$. This is because the social planner maximizes the joint surplus $\pi_{sm}(\gamma) + cs_s(\gamma)$, whereas the platform maximizes only its own profit $\pi_{sm}(\gamma)$. The platform's choice ignores the negative impact of a higher commission on consumer surplus (as fees are passed through as higher prices), whereas the planner internalizes this negative externality.

With the commission optimally set at γ_w, the planner's problem reduces to choosing the number of displayed third-party sellers. This choice involves a trade-off analogous to the one the platform faces. Increasing s generates a positive matching effect by increasing the total trade volume, while also yileds a negative weighting effect as the surplus from a first-party sale is higher than the joint surplus from a third-party sale, i.e., $w_m > \pi_{sm}(\gamma_w) + cs_s(\gamma_w)$. The social planner's optimal number of sellers should balance these two effects. Despite that both the planner and the platform actively manage the number of sellers, the platform's curation decision is weighted by its own profits, whereas the planner's is weighted by total trade surplus. Consequently, the platform may feature either too many or too few third-party sellers compared to the social optimum.

In conclusion, there are two distortions, a *fee distortion* where the platform sets a higher commission fee which exacerbates the double marginalization problem and leads to inefficiently high consumer prices; and a *quantity distortion* where the platform curates its marketplace to maximize profits, neglecting the influence on consumer surplus.

Understanding the nature of these distortions is crucial for designing effective regulation. Since a common regulatory approach is to control prices (e.g., by capping the commission fee γ) rather than dictating marketplace composition, it is particularly useful to analyze the quantity distortion in isolation. Therefore, we now fix γ to determine the conditions under which the platform lists too many or too few sellers.

The sign of the quantity distortion. The platform's profit-maximizing choice of sellers, $\hat{s}(\gamma)$, is driven by the ratio of its own revenue streams: $\frac{\pi_{sm}(\gamma)}{\pi_m}$, which is the commission revenue from a third-party sale relative to the profit from a first-party sale. In contrast, the planner's choice, $\hat{s}_w(\gamma)$, is driven by the ratio of total social surplus generated by each type of sale: $\frac{cs_s(\gamma)+\pi_{sm}(\gamma)}{w_m}$. The following lemma formalizes this distortion.

Lemma 3 *Fixing $\gamma \in (0, \bar{\gamma})$, the platform lists too few third-party sellers, $\hat{s}(\gamma) < \hat{s}_w(\gamma)$, if and only if:*

$$\frac{\pi_{sm}(\gamma)}{cs_s(\gamma)} < \frac{\pi_m}{cs_m}.$$

The lemma shows that the platform lists too few third-party sellers whenever its revenue-to-consumer-surplus ratio for third-party sales is less than the same ratio for its first-party sales. This is a classic case where the profit motive leads to an under-provision of goods that generate high social value. Conversely, when third-party sellers generate high commission revenue for the platform but offer little surplus to consumers, the platform has an incentive to list too many of them.

4.2 Banning Self-Preferencing

We not turn to the welfare implications of a regulatory ban on self-preferencing. Such a policy requires the platform to display all sellers who enter the market. The impact of this ban depends critically on the platform's equilibrium strategy. Since a ban only has an effect when the market is characterized by *excessive entry* ($s_0(\hat{\gamma}) > \hat{s}(\hat{\gamma})$), our subsequent analysis focuses on this scenario.

When self-preferencing is allowed, the platform sets the optimal commission rate $\hat{\gamma}$, and actively curates the marketplace by displaying a smaller mass of third-party sellers $\hat{s}(\hat{\gamma})$. The resulting social welfare is:

$$W_{sp} \equiv W(\hat{\gamma}, \hat{s}(\hat{\gamma})) = \mu^s\left(b, h + \hat{s}(\hat{\gamma})\right)\{hw_m + \hat{s}(\hat{\gamma})[\pi_{sm}(\hat{\gamma}) + cs_s(\hat{\gamma})]\}.$$

If self-preferencing is banned, the platform loses its ability to curate and must display all entering sellers. Its only strategic tool is the commission rate, which it sets at γ_{nsp}. Social welfare is then:

$$W_{nsp} \equiv W(\gamma_{nsp},\ s_0(\gamma_{nsp})) = \mu^s\left(b, h + s_0(\gamma_{nsp})\right)\{hw_m + s_0(\gamma_{nsp})[\pi_{sm}(\gamma_{nsp}) + cs_s(\gamma_{nsp})]\}.$$

To understand the welfare trade-offs, we can decompose the total welfare difference into two distinct components:

$$W_{sp} - W_{nsp} = \underbrace{W(\hat{\gamma},\ \hat{s}(\hat{\gamma})) - W(\hat{\gamma},\ s_0(\gamma_{nsp}))}_{\text{Quantity Effect}} + \underbrace{W(\hat{\gamma},\ s_0(\gamma_{nsp})) - W(\gamma_{nsp},\ s_0(\gamma_{nsp}))}_{\text{Fee Effect}}.$$

First, the *pricing effect* is always positive. A ban on self-preferencing incentivizes the platform to raise its commission rate to control entry: $\gamma_{nsp} > \hat{\gamma}$. This harms both consumer surplus and the platform's commission revenue, thereby reducing social welfare.

Second, the *quantity effect* captures the welfare change from moving the number of displayed sellers from $\hat{s}(\hat{\gamma})$, to the level under the ban, $s_0(\gamma_{nsp})$. This change can move the number of sellers either closer to or further from the social optimum $s_w(\gamma)$. Thus, the sign of this effect is ambiguous. Assuming the welfare function $W(\hat{\gamma}, s)$ is single-peaked in s, the quantity effect is positive if $\hat{s}(\hat{\gamma})$ brings the number of third-party sellers closer to the social optimum than $s_0(\gamma_{nsp})$ would.

The relationship between the curated number of sellers, $\hat{s}(\hat{\gamma})$, and the uncurated number, $s_0(\gamma_{nsp})$, is crucially determined by the platform's first-party capacity, h. We define two thresholds: $\tilde{h}$, below which the platform never self-preferences, and $\bar{h}$, above which the platform operates as a pure first-party monopolist by deterring all third-party entry. The following lemma describes scenario for intermediate values of $h \in (\tilde{h}, \bar{h})$.

Lemma 4 *There exist h_0 and h_1, satisfying $\tilde{h} < h_0 \leq h_1 < \bar{h}$, such that $\hat{s}(\hat{\gamma}) \leq s_0(\gamma_{nsp})$ when $h \leq h_0$, and $\hat{s}(\hat{\gamma}) \geq s_0(\gamma_{nsp})$ when $h \geq h_1$.*

The lemma reveals a non-monotonic change in the number of third-party sellers when a ban on self-preferencing is imposed. When the platform's first-party presence is small ($h \leq h_0$), a ban on self-preferencing leads to more third-party sellers being displayed. In contrast, when the platform's first-party presence is large ($h \geq h_1$), it responds to a ban by aggressively raising its commission fee to defend its first-party sales. This "over-deterrence" can be so strong that the number of sellers under a ban is actually *lower* than what the platform would have curated: $s_0(\gamma_{nsp}) \leq \hat{s}(\hat{\gamma})$.

Figure 2 demonstrates the case of over-deterrence. In the absence of a ban, the platform would select the optimal fee $\gamma = \hat{\gamma}$ at the peak of the red curve, yielding its preferred seller count of $\hat{s}(\hat{\gamma})$. In response to a ban on self-preferencing, however, the platform sets a much higher commission fee, γ_{nsp}. As illustrated by the blue point, this new fee corresponds to a sharply reduced number of sellers, $s_0(\gamma_{nsp})$. Crucially, this shows that the number of sellers under the ban is lower than

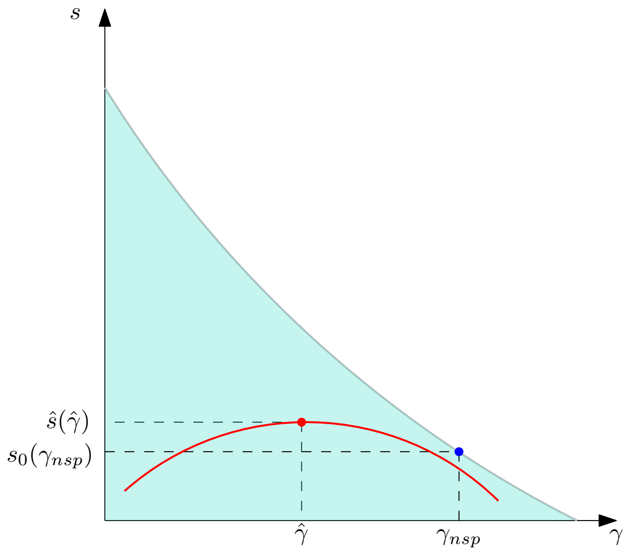

Figure 2: Over-deterrence.

what the platform would have curated.

Combining Lemma 3 and Lemma 4, we have our main result on the welfare effects of self-preferencing.

Proposition 4 *Self-preferencing yields higher social welfare ($W_{sp} > W_{nsp}$) under two distinct sets of conditions: (1) $\frac{\pi_{sm}(\hat{\gamma})}{cs_s(\hat{\gamma})} > \frac{\pi_m}{cs_m}$ and $h \in (\tilde{h}, h_0)$; (2) $\frac{\pi_{sm}(\hat{\gamma})}{cs_s(\hat{\gamma})} < \frac{\pi_m}{cs_m}$ and $h \in (h_1, \bar{h})$. If $h \leq \tilde{h}$ or $h \geq \overline{h}$, a ban has no effect, and $W_{sp} = W_{nsp}$.*

The proposition outlines sufficient conditions under which banning self-preferencing can harm welfare. It examines how the welfare outcome is influenced by the interaction of two factors: the type of self-preferencing distortion and the platform's first-party capacity.

In the first case, the platform is prone to listing too many sellers ($\hat{s}(\hat{\gamma}) > s_w(\hat{\gamma})$). When its first-party capacity is small ($h \in (\tilde{h}, h_0)$), a ban exacerbates this problem by having even more sellers onto the platform ($s_0(\gamma_{nsp}) \geq \hat{s}(\hat{\gamma})$), moving the outcome further from the social optimum. This creates a positive quantity effect. This, combined with the positive fee effect, makes self-preferencing welfare-superior.

In the second case, the platform is prone to listing too few sellers ($\hat{s}(\hat{\gamma}) < s_w(\hat{\gamma})$). When its first-party capacity is large ($h \in (h_1, \bar{h})$), a ban triggers the over-deterrence response, resulting in even fewer sellers ($s_0(\gamma_{nsp}) \leq \hat{s}(\hat{\gamma})$). Once again, the ban pushes the market further from the social optimum, yielding a positive quantitative effect that reinforces the positive pricing effect.

Finally, at the boundaries, a ban is ineffective. For a very small first-party presence ($h \leq \tilde{h}$), the platform does not self-preference. For a very large first-party presence ($h \geq \bar{h}$), it acts as a monopolist regardless of the policy. In both scenarios, the equilibrium outcome is identical, and thus $W_{sp} = W_{nsp}$.

4.3 Evaluating a Ban from the Consumer Welfare Perspective

We now evaluate the impact of a self-preferencing ban from the perspective of consumer welfare. Consumer surplus can also be expressed as a function of the commission rate, γ, and the mass of displayed third-party sellers, s:

$$CS(\gamma, s) = \mu^s(b, h+s)(h \cdot cs_m + s \cdot cs_s(\gamma)).$$

Assume $CS(\gamma, s)$ is single-peaked in s, and let $\hat{s}_c(\gamma)$ be the number of third-party sellers that maximizes consumer surplus (taking γ as given):

$$\hat{s}_c(\gamma) = \arg\max_s CS(\gamma, s).$$

The first-order condition of this maximization problem reveals that $\hat{s}_c(\gamma)$ is determined by the ratio of consumer surplus from a third-party match to that from a first-party match, $\frac{cs_s(\gamma)}{cs_m}$. This ratio governs the "weighting effect" from a purely consumer-centric viewpoint. The following lemma compares the platform's profit-maximizing choice of sellers to the consumer-optimal choice, which parallels Lemma 3.

Lemma 5 *Fixing $\gamma \in (0, \bar{\gamma})$, the platform displays too few third-party sellers from a consumer surplus perspective, $\hat{s}(\gamma) < \hat{s}_c(\gamma)$, if and only if: $\frac{\pi_{sm}(\gamma)}{cs_s(\gamma)} < \frac{\pi_m}{cs_m}$.*

We now state the main result for consumer surplus, denoting CS_{sp} as the consumer surplus with self-preferencing and CS_{nsp} as the surplus without it.

Proposition 5 *Self-preferencing yields higher consumer surplus ($CS_{sp} > CS_{nsp}$) under two distinct sets of conditions: (1) $\frac{\pi_{sm}(\hat{\gamma})}{cs_s(\hat{\gamma})} > \frac{\pi_m}{cs_m}$ and $h \in (\tilde{h}, h_0)$; (2) $\frac{\pi_{sm}(\hat{\gamma})}{cs_s(\hat{\gamma})} < \frac{\pi_m}{cs_m}$ and $h \in (h_1, \bar{h})$. If $h \leq \tilde{h}$ or $h \geq \bar{h}$, a ban has no effect, and $CS_{sp} = CS_{nsp}$.*

The reasoning behind this proposition mirrors that of Proposition 4, hinging on the combined impact of a fee effect and a quantity effect on consumers. First, a positive fee effect exists because a ban induces the platform to set a higher commission fee, which harms consumers. Second, the quantity effect is positive under the two sets of conditions, because a ban consistently pushes the number of displayed sellers further away from the consumer-optimal level.

5 Extensions on Fee Structures

5.1 Alternative Fee Structure: Per-Transaction Fee

As a robustness check, we now analyze an alternative fee structure: a per-transaction fee, $\tau \geq 0$, instead of an ad valorem commission fee. We show in this subsection that this change in fee structure does not alter the main conclusions of our baseline analysis.

Similarly in the final stage, a third-party seller, taking the per-transaction fee τ as given, solves the following profit-maximization problem:

$$\pi_{ss}(\tau) = \max_{p} [p - \tau - c] D(p). \tag{8}$$

By the Envelope Theorem, the seller's maximized profit, $\pi_{ss}(\tau)$, is strictly decreasing in τ. There exists a prohibitive fee $\bar{\tau}$ such that $\pi_{ss}(\tau) = 0$ for all $\tau \geq \bar{\tau}$.[15] From each successful third-party transaction, now the platform extracts a revenue of

$$\pi_{sm}(\tau) = \tau D(p_s(\tau)). \tag{9}$$

Following the baseline model, we assume that the platform's per-match revenue, $\pi_{sm}(\tau)$, is single-peaked in τ. We can thus define $\hat{\tau}$ as the fee that maximizes this revenue:

$$\hat{\tau} \equiv \arg\max_{\tau} \tau D(p_s(\tau)) \in (0, \bar{\tau}). \tag{10}$$

Seller Entry and Platform Profit Maximization. As before, given the first-party capacity h and the fee τ, the equilibrium measure of third-party sellers, $s_0(\tau)$, is determined by the standard free-entry condition:

$$\mu^s(b, h+s)\, \pi_{ss}(\tau) = k. \tag{11}$$

Since a higher fee τ reduces seller profits $\pi_{ss}(\tau)$, it follows that $s_0(\tau)$ is a decreasing function of τ: [16] $\frac{\partial s_0(\tau)}{\partial \tau} < 0$.

The platform's problem is to choose its fee, τ, to maximize total profits, subject to the sellers' entry decisions. The unconstrained problem, as in the baseline analysis, is a hypothetical scenario

[15] A seller earns a positive profit as long as it can set a price such that the effective margin is positive, i.e., $p > c + \tau$.

[16] There exists $\bar{\bar{\tau}} \in (0, \bar{\tau}]$ such that $s_0(\tau)$ is strictly decreasing in τ with $s_0(0) > 0$ and $s_0(\tau) = 0$ for $\tau \in [\bar{\bar{\tau}}, \infty)$.

where the platform can choose both s and τ independently. This problem is:

$$\Pi(\tau, s) \equiv M(b, h+s)\left(\frac{h}{h+s}\pi_m + \frac{s}{h+s}\pi_{sm}(\tau)\right).$$

The trade-off between the positive matching effect and the negative weighting effect is structurally identical to our baseline model. This implies that the logic from Lemma 1 extends directly, leading to a parallel result.

Lemma 6 *Assume that the platform's profit $\Pi(\tau, s)$ is single-peaked with respect to s. Let $\hat{s}(\tau)$ denote the unique profit-maximizing measure of third-party sellers for a given fee τ. Then, $\hat{s}(\tau)$ is single-peaked in τ. When interior, $\hat{s}(\tau)$ is increasing for $\tau < \hat{\tau}$ and decreasing for $\tau > \hat{\tau}$.*

With the free-entry constraint, the platform's problem is to choose τ to maximize $\Pi_{nsp}(\tau) \equiv \Pi(\tau, s_0(\tau))$. The first-order condition preserves the decomposition into a revenue effect and an entry effect. This leads to a proposition analogous to Proposition 1.

Proposition 6 (Optimal Transaction Fee Under Free Entry) *The platform's optimal fee τ_{nsp} is characterized by the first-order condition $\frac{d\Pi_{nsp}}{d\tau} = 0$. If $s_0(\hat{\tau}) > \hat{s}(\hat{\tau})$ (excessive entry), then $\tau_{nsp} > \hat{\tau}$. If $s_0(\hat{\tau}) < \hat{s}(\hat{\tau})$ (insufficient entry), then $\tau_{nsp} < \hat{\tau}$.*

Robustness to Fee Structure. The analysis above confirms that switching from the ad valorem commission fee to the per-transaction fee does not change the fundamental mechanics of our model. All subsequent results in the paper—including the self-preferencing strategy, the determinants of self-preferencing, and the welfare analysis—are derived from the key functions $\pi_{ss}(\cdot)$, $\pi_{sm}(\cdot)$, $s_0(\cdot)$, and $\hat{s}(\cdot)$. Since these functions retain the same essential properties (e.g., monotonicity or single-peakedness) under the per transaction fee, all our results hold. The core implications of the baseline model are therefore robust to this alternative specification of the platform's fee instrument.

5.2 Participation Fee as a Substitute for Self-Preferencing

In our baseline model, the platform uses self-preferencing to remedy excessive seller entry. This raises a natural question: could a participation fee serve as a substitute for self-preferencing? In this section, we explore this possibility and compare the platform's optimal strategy when using these two instruments.

For any given commission rate γ,[17] suppose the platform wishes to have a target number of third-party sellers, s'_{sp}, where $s'_{sp} < s_0(\gamma)$. The platform can achieve this in two ways. Using *self-preferencing*, more sellers (s'_E) enter than are displayed (s'_{sp}), with the entry condition being

$$\mu^s\left(b, h + s'_{sp}\right)\pi_{ss}(\gamma)\frac{s'_{sp}}{s'_E} = k.$$

Alternatively, the platform can charge a non-negative *participation fee*), $F \geq 0$. This fee effectively raises the entry cost, reducing the number of entrants to the target level, s'_{sp}, without needing to curate. The entry condition becomes

$$\mu^s\left(b, h + s'_{sp}\right)\pi_{ss}(\gamma) = k + F.$$

By setting $F = k\left(\frac{s'_E}{s'_{sp}} - 1\right)$, the platform can achieve the exact same number of displayed sellers. This shows that for reaching a specific seller count, fees and self-preferencing are perfect substitutes. However, as the following lemma states, they are not strategically equivalent.

Lemma 7 *When the platform can use both self-preferencing and a participation fee, its profit-maximizing strategy is to use the participation fee and not engage in self-preferencing.*

The intuition behind this result lies in efficiency. Self-preferencing is inherently wasteful because more sellers enter the market (s'_E) than are ultimately displayed (s'_{sp}). These undisplayed sellers incur the entry cost k but do not contribute to the matching process. A participation fee, by contrast, is a direct transfer. It allows the platform to ensure that only the desired number of sellers find it profitable to enter, and all who pay the cost are subsequently displayed. Given that the fee is a more efficient instrument for controlling entry, the platform will always prefer it when both tools are available.

Furthermore, if the platform were to use a participation fee to replicate the seller quantity from the self-preferencing outcome, $\hat{s}(\hat{\gamma})$, the resulting social welfare would be

$$\begin{aligned} W_F &= \mu^s\left(b, h + \hat{s}(\hat{\gamma})\right)\{hw_m + \hat{s}(\hat{\gamma})[\pi_{sm}(\hat{\gamma}) + \pi_{ss}(\hat{\gamma}) + cs_s(\hat{\gamma})]\} - \hat{s}(\hat{\gamma})k \\ &= \underbrace{\mu^s\left(b, h + \hat{s}(\hat{\gamma})\right)\{hw_m + \hat{s}(\hat{\gamma})[\pi_{sm}(\hat{\gamma}) + cs_s(\hat{\gamma})]\}}_{=W_{sp}} + \hat{s}(\hat{\gamma})F > W_{sp}. \end{aligned}$$

Therefore, the participation fee is a superior instrument to self-preferencing from a social welfare perspective. By directly pricing market access, it ensures all sellers who pay to enter are displayed, obviating the need for the wasteful curation inherent in self-preferencing. The superior

[17] The platform's profit is $\Pi = \mu^s\left(b, h + s'_{sp}\right)\left[h\pi_m + s'_{sp}(\pi_{sm}(\gamma) + \pi_{ss}(\gamma))\right] - s'_E k$. For any target outcome—a pair of entered and displayed sellers, (s'_E, s'_{sp})—that can be implemented with a commission rate $\gamma > 0$, the same outcome can also be implemented with $\gamma = 0$. This is because setting $\gamma = 0$ maximizes the potential pool of entrants ($s_0(0) > s_0(\gamma)$). Since setting $\gamma = 0$ is always feasible and yields a higher profit for any given quantity of sellers (as $\max_\gamma\left(\pi_{sm}(\gamma) + \pi_{ss}(\gamma)\right) = \pi_{sm}(0) + \pi_{ss}(0)$), the optimal commission rate must be $\gamma = 0$.

efficiency of participation fees points toward a more nuanced policy direction. A constructive regulatory approach could therefore shift its focus from banning specific practices to promoting transparent mechanisms like participation fees.

5.2.1 Comparing Participation Fee and Self-Preferencing

Having established that a participation fee is the superior instrument, we now compare the market structures emerging from the two distinct regimes: one reliant on self-preferencing and the other on a participation fee. While a platform with access to a powerful surplus-extraction tool like a participation fee would ideally set its commission to zero, our analysis proceeds for a fixed commission rate. This approach allows us to cleanly isolate how each instrument—a lump sum fee versus curation—incentivizes the platform's choice of seller quantity. Furthermore, this approach is directly relevant to policy contexts, such as a commission floor, where a platform's commission may be constrained to be positive.

First, consider the participation fee regime. When the platform uses a lump sum fee to control entry, it can precisely tune the number of entrants. The fee allows the platform to act as a residual claimant on the total surplus generated by the supply side. By substituting the free-entry condition ($F = \mu^s(\cdot)\pi_{ss}(\gamma) - k$), the platform's problem becomes equivalent to choosing the number of sellers that maximizes the total profits of the platform and third-party sellers:

$$\max_{s\in[0,\bar{s}]} \quad \Pi_F(s;\gamma) = \mu^s\left(b, h+s\right)\Big[h\pi_m + s\big(\pi_{sm}(\gamma) + \pi_{ss}(\gamma)\big)\Big] - sk$$

Assume the objective is single-peaked in s and let $\hat{s}_F(\gamma)$ be the solution to this problem. This is the platform's ideal number of sellers when it uses F, as it perfectly internalizes the sellers' entry cost. $\hat{s}_F(\gamma)$ decreases in γ, as a higher γ exacerbates double marginalization and reduces the total surplus that can be extracted.

Next, we compare this outcome, $\hat{s}_F(\gamma)$, to the optimal number of sellers in the unconstrained optimal number of sellers, $\hat{s}(\gamma)$. Recall that $\hat{s}(\gamma)$ is the number of sellers that maximizes the platform's gross revenue from sales, *ignoring* seller entry costs. In contrast, $\hat{s}_F(\gamma)$ maximizes the platform's net profit by *internalizing* seller entry costs. This difference drives the result in the following proposition.

Proposition 7 *Suppose the platform has a concave objective under both regimes, then $\hat{s}_F(\gamma) \leq \hat{s}(\gamma)$ if* *and only if*

$$\mu^s \left(b, h + \hat{s}(\gamma)\right) \pi_{ss}(\gamma) \cdot \frac{h\pi_m}{h\pi_m + \hat{s}(\gamma)\pi_{sm}(\gamma)} \leq k. \tag{12}$$

Condition (12) highlights the difference between the two regimes. A platform using a participation fee *internalizes* the full surplus generated by a marginal seller, since it can be extracted via the fee. However, adding sellers also dilutes its profitable first-party sales, creating a crowding-out effect. The condition compares the seller's surplus, adjusted for this first-party composition effect, with the seller's entry cost, k. The direction of this comparison depends critically on market parameters. A higher *entry cost* k, for instance, directly reduces the net surplus the platform can extract, thus lowering its optimal seller count, $\hat{s}_F(\gamma)$, and making it more likely that $\hat{s}_F(\gamma) \leq \hat{s}(\gamma)$. The effect of the platform's *first-party capacity*, h, is more complex. As h increases, the negative crowding-out effect on the platform's own sales becomes more severe. In response, the curation-only platform sharply reduces its target number of sellers, $\hat{s}(\gamma)$. This reduction in total competition makes each remaining seller slot more valuable by increasing the matching probability. A platform using a fee can capture this additional surplus, creating a stronger incentive to attract sellers. Therefore, for larger values of h, we may see $\hat{s}_F(\gamma) > \hat{s}(\gamma)$.

6 Conclusion

This paper provides a hybrid platform framework which shows that platform self-preferencing is not an inherently anti-competitive practice, but rather a response to the excessive entry of third-party sellers. The strategic value of self-preferencing lies in that it allows the platform to decouple its objectives of managing per-transaction revenue and fine-tuning seller quantity. A regulatory ban removes this tool, forcing the platform to rely on the blunt instrument of high commission fees to deter entry. This can trigger an inefficient over-deterrence, where the number of sellers is pushed even lower than what the platform would have chosen if self-preferencing were allowed. We further highlight the critical role of the platform's first-party capacity and the ambiguous nature of welfare outcomes of banning self-preferencing. These findings caution against one-size-fits-all bans, which can trigger unintended consequences and ultimately harm consumers and social welfare.

References

Anderson, Simon P. and Özlem Bedre Defolie, "Hybrid platform model: monopolistic competition and a dominant firm," *RAND Journal of Economics*, 2024, *55* (4), 684–718.

__ , Andre De Palma, and Yurii Nesterov, "Oligopolistic Competition and the Optimal Provision of Products," *Econometrica*, 1995, *63* (6), 1281–1301.

Chi, Chang-Koo, Jay Pil Choi, Jong-Hee Hahn, and Seongkyun Kim, "Platform Incentives and the Limits of Self-Preferencing Regulation," 2025. Working paper.

de Cornière, Alexandre and Greg Taylor, "A model of biased intermediation," *RAND Journal of Economics*, 2019, *50* (4), 854–882.

Dendorfer, Florian, "First-party selling and self-preferencing," *International Journal of Industrial Organization*, 2024, *97* (4), 103098.

Gautier, Pieter, Bo Hu, and Makoto Watanabe, "Marketmaking Middlemen," *The RAND Journal of Economics*, 2023, *54* (1), 83–103.

Hagiu, Andrei, Tat-How Teh, and Julian Wright, "Should platforms be allowed to sell on their own marketplaces?," *The RAND Journal of Economics*, 2022, *53* (2), 297–327.

Hervas-Drane, Andres and Sandro Shelegia, "Retailer-Led Marketplaces," *Management Science*, 2025.

Mankiw, N. Gregory and Michael D. Whinston, "Free Entry and Social Inefficiency," *The RAND Journal of Economics*, 1986, *17* (1), 48–58.

Sato, Susumu and Yuta Kittaka, "Search-Order Design by Dual-Role Platforms," 2024. Working paper.

Wang, Zhu and Julian Wright, "Ad valorem platform fees, indirect taxes, and efficient price discrimination," *RAND Journal of Economics*, 2017, *48* (2), 467–484.

Zennyo, Yusuke, "Platform Encroachment and Own-content Bias," *The Journal of Industrial Economics*, 2022, *70* (3), 684–710.

Appendix Proofs

Proof of Proposition 1

Case 1: Excessive Entry $s_0(\hat{\gamma}) > \hat{s}(\hat{\gamma})$. Condition (6) shows the platform is incentivized to raise its fee to deter sellers. Thus, $\gamma_{nsp} > \hat{\gamma}$. Next, we show $\gamma_{nsp} \leq \gamma_r$, where γ_r is defined in footnote 12. Assume that the platform overshoots by setting $\gamma_{nsp} > \gamma_r$. Since $s_0(\gamma)$ is a decreasing function of the fee, $s_0(\gamma_{nsp}) < s_0(\gamma_r)$. Consider the profit at this supposed optimum. We can show a profitable deviation exists:

$$\Pi_{nsp}(\gamma_{nsp}) = \Pi(\gamma_{nsp}, s_0(\gamma_{nsp})) < \Pi(\gamma_r, s_0(\gamma_{nsp})) \leq \Pi(\gamma_r, s_0(\gamma_r)) = \Pi_{nsp}(\gamma_r)$$

The first inequality holds because, for a fixed number of sellers, the platform's profit function is strictly decreasing in the fee when $\gamma > \hat{\gamma}$. Since $\gamma_{nsp} > \gamma_r > \hat{\gamma}$, the lower fee γ_r yields higher profit. The second inequality holds because, for a fixed fee γ_r, the platform's profit is maximized when the number of third-party sellers is $s_0(\gamma_r)$. By definition of γ_r, $s_0(\gamma_r) = \hat{s}(\gamma_r)$. Thus, the platform must set $\gamma_{nsp} \leq \gamma_r$.

Case 2: Insufficient Entry $s_0(\hat{\gamma}) < \hat{s}(\hat{\gamma})$. The first-order condition implies that with insufficient entry, the platform is incentivized to lower its fee to attract sellers. Thus, $\gamma_{nsp} < \hat{\gamma}$. Next, assume the contrary that $\gamma_{nsp} < \gamma_l$, where γ_l is defined in footnote 13. Then, $s_0(\gamma_{nsp}) > s_0(\gamma_l) = \hat{s}(\gamma_l)$. A profitable deviation to γ_l exists:

$$\Pi_{nsp}(\gamma_{nsp}) = \Pi(\gamma_{nsp}, s_0(\gamma_{nsp})) < \Pi(\gamma_l, s_0(\gamma_{nsp})) \leq \Pi(\gamma_l, s_0(\gamma_l)) = \Pi_{nsp}(\gamma_l).$$

The first inequality holds because profit is strictly increasing in the fee when $\gamma < \hat{\gamma}$. The second inequality holds because, for a fixed fee γ_l, profit is maximized at $\hat{s}(\gamma_l)$. Thus, we must have $\gamma_{nsp} \geq \gamma_l$. ■

Proof of Lemma 2

First, we prove that $s_{sp} > s_0(\gamma)$ is infeasible.

Suppose $s_{sp} > s_0(\gamma)$. Since $s_E \geq s_{sp}$, we have $\frac{s_{sp}}{s_E} \leq 1$. This leads to:

$$\mu^s\left(b, h + s_{sp}\right) \pi_{ss}(\gamma) \cdot \frac{s_{sp}}{s_E} \leq \mu^s\left(b, h + s_{sp}\right) \pi_{ss}(\gamma) < \mu^s\left(b, h + s_0(\gamma)\right) \pi_{ss}(\gamma) = k$$

This implies the expected profit is strictly less than k, which contradicts the free entry condition.

Next, we prove that the set $[0, s_0(\gamma)]$ is feasible. Specifically, we show that for any $s_{sp} \in [0, s_0(\gamma)]$, there exists a corresponding mass of entrants $s_E \geq s_{sp}$ that satisfies the free entry condition.

By the definition of $s_0(\gamma)$, we have the following condition hold:

$$\mu^s\left(b, h + s_0(\gamma)\right) \pi_{ss}(\gamma) = k \Leftrightarrow \mu^s\left(b, h + s_0(\gamma)\right) \pi_{ss}(\gamma) s_0(\gamma) = s_0(\gamma) k.$$

Suppose the platform chooses a smaller number of sellers to display, $s_{sp} < s_0(\gamma)$, we have

$$\mu^s\left(b, h + s_{sp}\right) \pi_{ss}(\gamma) > k \Leftrightarrow \mu^s\left(b, h + s_{sp}\right) \pi_{ss}(\gamma) s_{sp} > s_{sp} k.$$

Since $\mu^s(b, h+s) \cdot s$ is increasing in s, we have $\mu^s\left(b, h + s_{sp}\right) s_{sp} < \mu^s\left(b, h + s_0(\gamma)\right) s_0(\gamma)$. Taking all these conditions together, we have

$$s_{sp} k < \mu^s\left(b, h + s_{sp}\right) \pi_{ss}(\gamma) s_{sp} < \mu^s\left(b, h + s_0(\gamma)\right) \pi_{ss}(\gamma) s_0(\gamma) = s_0(\gamma) k.$$

By the Intermediate Value Theorem, there must exist a unique $s_E \in (s_{sp}, s_0(\gamma))$ such that

$$\mu^s\left(b, h + s_{sp}\right) \pi_{ss}(\gamma) s_{sp} = s_E k.$$

Finally, it is easy to see that $s_E = s_{sp} = s_0(\gamma)$ if and only if $s_{sp} = s_0(\gamma)$. ■

Proof of Proposition 2

With excessive entry at $\hat{\gamma}$, i.e., $s_0(\hat{\gamma}) > \hat{s}(\hat{\gamma})$, the platform can attain the unconstrained optimum $(\gamma_{sp},\, s_{sp}) = (\hat{\gamma},\, \hat{s}(\hat{\gamma}))$ by engaging in self-preferencing.

With insufficient entry at $\hat{\gamma}$, i.e., $s_0(\hat{\gamma}) < \hat{s}(\hat{\gamma})$, suppose the platform engages in self-preferencing by choosing some $\hat{s}(\gamma_{sp}) < s_0(\gamma_{sp})$. Then, by the same argument as in Proposition 1, there exists a $\gamma_l \in (\gamma_{sp}, \hat{\gamma})$ such that:

$$\Pi(\gamma_l, \hat{s}(\gamma_l)) > \Pi(\gamma_l, \hat{s}(\gamma_{sp})) \geq \Pi(\gamma_{sp}, \hat{s}(\gamma_{sp})),$$

or a $\gamma_r \in (\hat{\gamma}, \gamma_{sp})$ such that:

$$\Pi(\gamma_r, \hat{s}(\gamma_r)) > \Pi(\gamma_r, \hat{s}(\gamma_{sp})) \geq \Pi(\gamma_{sp}, \hat{s}(\gamma_{sp})).$$

This contradicts that (γ_{sp}, s_{sp}) is optimal. Therefore, any optimal choice must satisfy $s_0(\gamma_{sp}) \leq \hat{s}(\gamma_{sp})$. Accordingly, the free-entry condition must be binding, and the outcome is identical to the no-self-preferencing benchmark. ■

Proof of Proposition 3

We first show the gap $s_0(\hat{\gamma}) - \hat{s}(\hat{\gamma})$ is strictly increasing in h. For this, we examine the derivatives of $s_0(\cdot)$ and $\hat{s}(\cdot)$ with respect to h. $s_0(\cdot)$ is determined by (5). Differentiating with respect to h yields $\frac{\partial s_0(\hat{\gamma})}{\partial h} = -1$. The platform's optimal third-party presence $\hat{s}(\gamma)$ is defined by the first-order condition $\frac{\partial \Pi(\gamma,s)}{\partial s} = 0$. Applying the implicit function theorem gives

$$\frac{\partial \hat{s}(\hat{\gamma})}{\partial h} = -\frac{\frac{\partial^2 \Pi(\gamma,s)}{\partial h \partial s}}{\frac{\partial^2 \Pi(\gamma,s)}{\partial^2 s}} = -\frac{\frac{\partial^2 \Pi(\gamma,s)}{\partial^2 s} - \frac{\partial \mu^s(b,h+s)}{\partial s}[\pi_m - \pi_{sm}(\gamma)]}{\frac{\partial^2 \Pi(\gamma,s)}{\partial^2 s}} < -1.$$

Therefore, $s_0 - \hat{s}$ is strictly increasing in h, provided that both s_0 and $\hat{s}$ are positive.

Next, we establish the sign of the gap at low and high values of h. As $h \to 0$, $\hat{s}(\hat{\gamma}) \to \infty$ while $s_0(\hat{\gamma})$ remains finite. Consequenly, for a sufficiently small h, $s_0(\hat{\gamma}) - \hat{s}(\hat{\gamma}) < 0$. To find a high value of h such that $s_0(\hat{\gamma}) - \hat{s}(\hat{\gamma}) > 0$, and with some abuse of notations, let $\bar{h}$ be the level of first-party presence at which the platform's optimal choice is to deter all third-party entry, i.e., $\hat{s}(\bar{h}, \hat{\gamma}) = 0$. Now, let $\bar{k}$ be the entry cost that would make equilibrium third-party entry zero at this same level of first-party presence, i.e., $s_0(\bar{k}, \bar{h}, \hat{\gamma}) = 0$. By assumption, the actual entry cost k is less than this prohibitive level ($k < \bar{k}$). Therefore, at $h = \bar{h}$, third-party entry is still profitable, meaning $s_0(\hat{\gamma}) > 0$. Since $\hat{s}(\hat{\gamma}) = 0$ (and $\hat{s}$ is independent of k), the gap at $h = \bar{h}$ is positive: $s_0(\hat{\gamma}) - \hat{s}(\hat{\gamma}) > 0$.

The arguments above ensure that there exists a unique threshold $\tilde{h} \in (0, \bar{h})$ such that

$$s_0(\hat{\gamma})\big|_{h=\tilde{h}} = \hat{s}(\hat{\gamma})\big|_{h=\tilde{h}} > 0.$$

Because the gap is strictly increasing, it follows that for all $h < \tilde{h}$, the gap is negative, and for all $h > \tilde{h}$, the gap is positive. ■

Proof of Lemma 4

⊙ Step 1: Prove that $\hat{s}(\hat{\gamma}) \geq s_0(\gamma_{nsp})$ when $h \geq h_1$. Let $\hat{s}(\hat{\gamma}, h) \equiv \hat{s}(\hat{\gamma})$ be the optimal seller measure with self-preferencing. It is the solution to the first-order condition:

$$\frac{\partial[\mu^s(b, h+s)\, s]}{\partial s} \pi_{sm}(\hat{\gamma}) + \frac{\partial \mu^s(b, s+h)}{\partial s} h \pi_m = 0. \tag{13}$$

Let $s_{nsp}(h) \equiv s_0(\gamma_{nsp})$ be the optimal seller measure without self-preferencing. In this case, the platform effectively chooses s, which in turn determines the commission rate $\gamma_0(s)$ via the free-entry condition. The corresponding first-order condition is:

$$\frac{\partial[\mu^s(b, h+s)\, s]}{\partial s}\pi_{sm}(\gamma_0(s)) + \frac{\partial \mu^s(b, h+s)}{\partial s}h\pi_m + \frac{\partial \pi_{sm}(\gamma_0(s))}{\partial s}\mu^s(b, h+s)\, s = 0. \tag{14}$$

Both $\hat{s}(\hat{\gamma}, h)$ and $s_{nsp}(h)$ are decreasing functions of h. We define their respective "shutdown" thresholds, $\bar{h}$ and $\bar{h}_{nsp}$, as the levels of first-party presence that drive optimal third-party presence to zero:

$$\hat{s}(\hat{\gamma}, \bar{h}) = 0 \quad \text{and} \quad s_{nsp}(\bar{h}_{nsp}) = 0.$$

To prove the proposition, we will show that $\bar{h} > \bar{h}_{nsp}$. By definition, $\bar{h}$ is equivalently given by (13) evaluated at $s = 0$:

$$\frac{\partial[\mu^s(b, h+s)\, s]}{\partial s}\pi_{sm}(\hat{\gamma}) + \frac{\partial \mu^s(b, h+s)}{\partial s}h\pi_m\Big|_{h=\bar{h}, s=0} = 0.$$

Condition (14) evaluated at the same point $h = \bar{h}, s = 0$ satisfies

$$\begin{aligned}
&\frac{\partial[\mu^s(b, h+s)\, s]}{\partial s}\pi_{sm}(\gamma_0^{-1}(s)) + \frac{\partial \mu^s(b, h+s)}{\partial s}h\pi_m + \frac{\partial \pi_{sm}(\gamma_0^{-1}(s))}{\partial s}\mu^s(b, h+s)\, s\Big|_{h=\bar{h}, s=0} \\
&= \frac{\partial[\mu^s(b, h+s)\, s]}{\partial s}\pi_{sm}(\gamma_0^{-1}(s)) + \frac{\partial \mu^s(b, h+s)}{\partial s}h\pi_m\Big|_{h=\bar{h}, s=0} \\
&< \frac{\partial[\mu^s(b, h+s)\, s]}{\partial s}\pi_{sm}(\hat{\gamma}) + \frac{\partial \mu^s(b, h+s)}{\partial s}h\pi_m\Big|_{h=\bar{h}, s=0} = 0.
\end{aligned}$$

In words, for $h \geq \bar{h}$, we have $s_{nsp}(h) = 0$. Then it must be that $\bar{h}_{nsp} < \bar{h}$. Consequently, for $h \in [\bar{h}_{nsp}, \bar{h}]$, $\hat{s}(\hat{\gamma}, h) > s_{nsp}(h) = 0$. By continuity, there exists a threshold $h_1 < \bar{h}_{nsp}$ such that for $h \in (h_1, \bar{h})$, we have $\hat{s}(\hat{\gamma}, h) > s_{nsp}(h)$.

⊙ Step 2: Prove that $\hat{s}(\hat{\gamma}) \leq s_0(\gamma_{nsp})$ when $h \leq h_0$. Given the maximized profit under no self-preferencing

$$\Pi_{nsp} = \mu^s\left(b, h + s_0(\gamma_{nsp})\right)\left(h\pi_m + s_0(\gamma_{nsp})\pi_{sm}(\gamma_{nsp})\right),$$

and that under self-preferencing

$$\Pi_{sp} = \mu^s\left(b, h + \hat{s}(\hat{\gamma})\right)\left(h\pi_m + \hat{s}(\hat{\gamma})\pi_{sm}(\hat{\gamma})\right),$$

using the Envelope theorem, we can derive the marginal revenue for the platform as h marginally increases

$$\begin{aligned}
MR_{nsp} &\equiv \frac{\mathrm{d}\Pi_{nsp}}{\mathrm{d}h} = [\pi_m - \pi_{sm}(\gamma_{nsp})]\mu^s\left(b, h + s_0(\gamma_{nsp})\right), \\
MR_{sp} &\equiv \frac{\mathrm{d}\Pi_{sp}}{\mathrm{d}h} = [\pi_m - \pi_{sm}(\hat{\gamma})]\mu^s\left(b, h + \hat{s}(\hat{\gamma})\right).
\end{aligned}$$

We assert that there exists a threshold $h_0' > \tilde{h}$, such that for all $h \in [\tilde{h},\ h_0')$, we have $MR_{sp} > MR_{nsp}$. To prove this by contradiction, suppose such a threshold h_0' does not exist. Then there must exist another $h_0'' > \tilde{h}$, such that for all $h \in [\tilde{h},\ h_0'']$, we would have $MR_{sp} \leq MR_{nsp}$. Since at $h = \tilde{h}$, $\Pi_{sp} = \Pi_{nsp}$, it follows that $\Pi_{sp} < \Pi_{nsp}$ for $h \in [\tilde{h},\ h_0'']$. This contradicts that $\Pi_{sp} > \Pi_{nsp}$ when $h > \tilde{h}$ (implying $\hat{s}(\hat{\gamma}) < s_0(\hat{\gamma})$, see the proof of Proposition 3). Therefore, such a threshold h_0' must exist.

Since $MR_{sp} > MR_{nsp}$ for $h \in [\tilde{h},\ h_0')$, and we also have $[\pi_m - \pi_{sm}(\hat{\gamma})] < [\pi_m - \pi_{sm}(\gamma_{nsp})]$, this implies that $\hat{s}(\hat{\gamma}) < s_0(\gamma_{nsp})$ within this interval. By continuity, we know that there exists a threshold h_0 ($h_0 \geq h_0'$), such that when $\tilde{h} < h < h_0$, we have $\hat{s}(\hat{\gamma}) < s_0(\gamma_{nsp})$. ■

Proof of Lemma 7

Let s be the platform's target number of displayed third-party sellers. Suppose the platform uses a combination of a participation fee, F', and self-preferencing to achieve this. Let s' be the total mass of third-party sellers who choose to enter, where by definition $s' \geq s$. The display probability is therefore s/s', with $s/s' \leq 1$.

The entry condition is now:

$$\mu^s(b, h+s)\,\pi_{ss}(\gamma)\frac{s}{s'} = k + F'.$$

The platform's profit is the sum of its profit from matching and the revenue from participation fees:

$$\begin{aligned}\Pi &= \mu^s(b, h+s)\,[h\pi_m + s\pi_{sm}(\gamma)] + s'F' \\ &= \mu^s(b, h+s)\,[h\pi_m + s\big(\pi_{sm}(\gamma) + \pi_{ss}(\gamma)\big)] - s'k.\end{aligned}$$

From the expression, we can see that the platform's profit, Π, is strictly decreasing in the total mass of third-party sellers who choose to enter, s'. Therefore, to maximize its profit for a given target s, the platform has an incentive to minimize s'. The minimum feasible value for s' is s, which occurs when the display probability $s/s' = 1$.

Thus, the optimal strategy is to set a participation fee, F, that induces exactly s sellers to enter and then displays all of them. This fee is determined by the simplified entry condition where $s' = s$:

$$\mu^s(b, h+s)\,\pi_{ss}(\gamma) = k + F.$$

■

Proof of Proposition 7

The first-order condition for the optimal number of sellers, $\hat{s}_F(\gamma)$, is

$$\frac{\partial\mu^s(b, s+h)}{\partial s}\Big[h\pi_m + s\big(\pi_{sm}(\gamma) + \pi_{ss}(\gamma)\big)\Big] + \pi_{ss}(\gamma))\mu^s(b, h+s)\Big|_{s=\hat{s}_F(\gamma)} = k.$$

Substitute $\hat{s}(\gamma)$, into the left-hand side (LHS) of the first-order condition for $\hat{s}_F(\gamma)$, which yields:

$$LHS|_{s=\hat{s}(\gamma)} = \frac{\partial\mu^s(b, s+h)}{\partial s}s\pi_{ss}(\gamma) + \pi_{ss}(\gamma)\mu^s(b, h+s)\Bigg|_{s=\hat{s}(\gamma)}.$$

From the first-order condition for $\hat{s}(\gamma)$, we can further obtain:

$$\begin{aligned}&-\frac{\partial\mu^s(b, s+h)}{\partial s}s\pi_{ss}(\gamma) \\ &= \frac{s\pi_{sm}(\gamma)}{h\pi_m + s\pi_{sm}(\gamma)}\pi_{ss}(\gamma)\mu^s(b, h+s)\Bigg|_{s=\hat{s}(\gamma)}.\end{aligned}$$

This is equivalent to stating that:

If and only if $\mu^s(b, h+\hat{s}(\gamma))\,\pi_{ss}(\gamma)\cdot\frac{h\pi_m}{h\pi_m + \hat{s}(\gamma)\pi_{sm}(\gamma)} \leq k$, then $\hat{s}_F(\gamma) \leq \hat{s}(\gamma)$.■

Appendix: Proof of Single-peakedness for Example 2

To prove that the platform's profit function, $\Pi(\gamma, s)$, is single-peaked with respect to s, we analyze its partial derivative, $\frac{\partial \Pi(\gamma,s)}{\partial s}$. By substituting the specified matching probability, We show that the sign of the derivative hinges on a simpler function, denoted $G(s)$, which is derived from the numerator of the derivative. Therefore, the single-peaked property of $\Pi(\gamma, s)$ is established if we can show that $G(s) = 0$ has at most one positive root.

The sign of the partial derivative $\frac{\partial \Pi(\gamma,s)}{\partial s} \gtrless 0$ is equivalent to

$$\frac{\partial \mu^s(b, h+s)}{\partial s}\left(h\pi_m + s\pi_{sm}(\gamma)\right) + \mu^s(b, h+s)\pi_{sm}(\gamma) \gtrless 0 \Leftrightarrow -\frac{\partial \mu^s(b, h+s)}{\partial s}\frac{1}{\mu^s(b, h+s)} \lessgtr \frac{\pi_{sm}(\gamma)}{h\pi_m + s\pi_{sm}(\gamma)}.$$

After rearranging and substituting the matching probability $\mu^s(b, h+s) = \frac{b\zeta(1-e^{-(h+s)w})}{h+s}$ from Example 2, and defining $\beta = \frac{\pi_m}{\pi_{sm}(\gamma)} > 1$, the condition reduces to

$$\frac{e^{-w(h+s)}(w(h+s)+1)-1}{1-e^{-w(h+s)}} \lessgtr \frac{h+s}{h\beta+s} \Leftrightarrow e^{w(h+s)}h(\beta-1) - s^2w - swh(\beta+1) \lessgtr h(\beta-1) + h^2w\beta.$$

To facilitate the analysis of this inequality, we define the function $G(s)$ such that:

$$G(s) = e^{w(h+s)}h(\beta-1) - s^2w - swh(\beta+1) - h(\beta-1) - h^2w\beta.$$

Thus, the profit function $\Pi(\gamma, s)$, is single-peaked with respect to s if and only if the function $G(s)$ has at most one root for $s \geq 0$.

To determine the number of roots of $G(s)$, we now analyze its essential characteristics. Since the parameters w, h, and $(\beta - 1)$ are all positive, the exponential term $e^{w(h+s)}$ dominates the polynomial terms as $s \to +\infty$. This leads to

$$\lim_{s\to\infty} G(s) = +\infty.$$

Next, we consider the function's first and second derivatives:

$$G'(s) = we^{w(h+s)}h(\beta-1) - 2sw - wh(\beta+1),$$
$$G''(s) = w^2e^{w(h+s)}h(\beta-1) - 2w.$$

Similarly, the limit of the first derivative is $\lim_{s\to\infty} G'(s) = +\infty$. Furthermore, $G''(s)$ is a strictly increasing function, implying that $G''(s) = 0$ can have at most one root.

Having established the analytical properties of $G(s)$, we now prove that it has at most one root for $s \geq 0$. The proof proceeds by analyzing two cases based on the sign of $G(0)$, the function's value at the boundary.

⊙ Case 1: $G(0) < 0$.

We proceed by contradiction that $G(s)$ has more than one root for $s \geq 0$. The conditions $G(0) < 0$ and $\lim_{s\to\infty} G(s) = +\infty$, combined with the continuity of $G(s)$, guarantee the existence of at least one positive root by the Intermediate Value Theorem. For more than one root to exist under these conditions, the function must change direction multiple times. Thus, it implies that $G(s)$ must have at least three positive roots.

By a repeated application of Rolle's Theorem, if $G(s)$ has at least three positive roots, then its derivative, $G'(s)$, must have at least two distinct positive roots. Applying the theorem again to $G'(s)$ implies that its second derivative, $G''(s)$, must have one positive root. To show that this conclusion is impossible, we now proceed with a case analysis based on the sign of $G'(0)$.

Case 1: $G'(0) < 0$. The function $G'(s)$ is continuous, starts at a negative value ($G'(0) < 0$), and approaches infinity ($\lim_{s\to\infty} G'(s) = +\infty$). As established previously, since $G''(s)$ has at most one root, $G'(s)$ can have at most one local extremum. A continuous function with these properties—starting negative, ending at positive infinity, and having at most one extremum—must cross the s-axis exactly once. Therefore, $G'(s)$ has exactly one positive root. This conclusion directly contradicts the necessary condition derived from Rolle's Theorem, which requires $G'(s)$ to have at least two positive roots. Thus, our initial assumption of three roots for $G(s)$ is false in this case.

Case 2: $G'(0) \geq 0$. $G'(0) \geq 0$ is equivalent to the inequality $we^{hw}h(\beta-1) \geq wh(\beta+1)$. ***Case 2:*** $G'(0) \geq 0$. This case is defined by the inequality $we^{hw}h(\beta-1) \geq wh(\beta+1)$. When combined with the initial assumption that $G(0) < 0$, further simplification yields a crucial property $hw\beta > 2$.

Since $G''(s)$ is a strictly increasing function, its minimum value occurs at $s = 0$:

$$G''(s) \geq G''(0) = w^2 e^{hw} h(\beta-1) - 2w$$

From the condition for this case, $we^{hw}h(\beta-1) \geq wh(\beta+1)$, we can substitute this into the inequality to establish a lower bound:

$$G''(s) \geq w\big(wh(\beta+1)\big) - 2w = w(wh\beta + wh - 2)$$

Because we derived that $hw\beta > 2$, and we know $wh > 0$, the term $(wh\beta + wh - 2)$ is strictly positive. Therefore, we conclude that $G''(s) > 0$ for all $s \geq 0$. This implies that $G''(s)$ has zero positive roots, which directly contradicts the necessary condition from Rolle's Theorem that $G''(s)$ must have one positive root.

⊙ Case 2: $G(0) \geq 0$.

In this case, we will show that $G(s) > 0$ for $s > 0$. First, the condition $G(0) \geq 0$ is equivalent to $e^{hw}h(\beta-1) \geq h(\beta-1) + h^2 w\beta$. Rearranging this inequality reveals a condition on β:

$$\beta \geq \frac{e^{hw}-1}{e^{hw}-1-hw}.$$

It is a known property that for any $x > 0$, the function $f(x) = x\frac{e^x-1}{e^x-1-x}$ is always greater than or equal to 2. Applying this property with $x = hw$, we establish the key inequality for this case: $hw\beta \geq 2$.

To prove that $G(s)$ is positive, we construct a simpler lower-bound function, $G_2(s)$, by substituting the expression for $G(0) \geq 0$ back into the definition of $G(s)$. This ensures that $G(s) \geq G_2(s)$ for all $s \geq 0$, where:

$$G_2(s) = e^{sw}[h^2 w\beta + h(\beta-1)] - s^2 w - swh(\beta+1) - h(\beta-1) - h^2 w\beta.$$

This function has three key properties. First, by substitution, $G_2(0) = 0$. Second, its second derivative is strictly positive for $s \geq 0$, since $G_2''(s) = w^2 e^{sw}[h^2 w\beta + h(\beta-1)] - 2w > 0$, which follows from our derived inequality $hw\beta \geq 2$. This implies $G_2(s)$ is a convex function. Third, its first derivative at the origin is non-negative, as $G_2'(0) = wh(hw\beta - 2) \geq 0$.

The three properties of $G_2(s)$—starting at zero ($G_2(0) = 0$), having a non-negative initial slope ($G_2'(0) \geq 0$), and being convex ($G_2''(s) > 0$)—collectively ensure that $G_2(s) > 0$ for all $s > 0$. Since we know $G(s) \geq G_2(s)$, it follows that $G(s) > 0$ for all $s > 0$, which equals to $G(s)$ has at most one root for $s \geq 0$.