Impressing the Algorithm*

Sales-Based Ranking, Learning, and Off-Platform Pricing

Heski Bar-Isaac† and Sandro Shelegia‡

May 6, 2026

Abstract

We study how platform ranking design shapes sellers' off-platform pricing when the platform learns product appeal from on-platform sales. Because those sales depend on cross-channel pricing, ranking creates a career-concerns force: by restraining off-platform discounts, the seller shifts demand on-platform and improves future exposure. Thus, even without conditioning directly on off-platform prices, a sales-based algorithm can reproduce the incentive effects of explicit price-parity restrictions. But because the platform also relies on on-platform sales to infer seller appeal, a sales-based algorithm may distort promotion decisions and reduce welfare. We compare sales-based algorithms with rules that condition directly on off-platform prices and show that a price-parity-clause algorithm can attain the first best.

*Thanks to the audiences at the 19th Berlin IO Day, the Economics of the DMA Workshop, USC Marshall, the Annual Conference of MaCCI, the University of Toronto, ESSET 2025, Mad-Bar 2025, the IO Winter Symposium - Bergamo 2024, the Pricing, Algorithms and Competition Conference, EARIE 2025, University of Barcelona, The 18th Toulouse Digital Economics conference, the Singapore Initiative on Digital Economics (SIDE) Workshop, Yi Chen, Martin Peitz, and Anton Sobolev. The authors are grateful for financial support from the Severo Ochoa Program for Centres of Excellence in R&D (Barcelona School of Economics CEX2019-000915-S) funded by MCIN/AEI/10.13039/501100011033, Spanish Ministry of Science and Universities (Grant MICIU/AEI/UE - PID2023-153318NB-I00) and the European Commission via ERC (grant FAPoD (101044072)).

†University of Toronto Distinguished Professor of Economics and Finance, Distinguished Fellow at the School of Business Administration at The Hebrew University of Jerusalem, CEPR Fellow, and CRESSE Associate. Email: heski.bar-isaac@rotman.utoronto.ca

‡Universitat Pompeu Fabra, BSE, and CEPR. Email: sandro.shelegia@upf.edu

1 Introduction

Ranking algorithms determine which sellers consumers see on digital platforms, and therefore whether those sellers make sales at all.[1] Because visibility is a prerequisite for sales, sellers devote substantial effort to understanding and influencing ranking rules, while platforms design those rules strategically.[2] Ranking design has also become a policy concern. Recent regulatory initiatives, including the European Union's Digital Markets Act and Digital Services Act, have focused attention on the information that platforms use in ranking and recommendation systems, particularly where those systems affect fairness, contestability, or sellers' ability to compete across channels.[3] This raises a natural question: how should a platform design its ranking rule when sales are informative about seller appeal, but sellers can also influence those sales through pricing on other channels?

We study this question in a setting where the platform learns from on-platform sales. Once a consumer discovers a seller on the platform, they may purchase either on-platform or directly from the seller. The seller can therefore influence the signal that the platform observes. By charging a higher off-platform price, the seller shifts demand on-platform, generates more observed sales, and appears more attractive to the algorithm than it otherwise would.[4] In that sense, the seller can *impress the algorithm*. The distortion is dynamic rather than static: a seller who would otherwise offer an off-platform discount instead restrains that discount in order to improve future exposure. Anticipating this response, the platform may design its ranking rule not only to allocate visibility, but also to shape off-platform competition.

[1] Ursu (2018) provides causal evidence that rankings have significant effects on consumer search and purchase decisions.

[2] Evidence of seller investment in ranking optimization is abundant. Numerous online forums are dedicated to Amazon SEO (search engine optimization), and a competitive digital marketing industry offers advice and consulting on strategies to rise in search rankings. As one illustration, Webfx states that "Amazon SEO is essential for businesses selling on Amazon. If your company doesn't optimize your listings for SEO, you will struggle to sell products and turn a profit," and invites firms to contact them for "our SEO services for Amazon." See `https://www.webfx.com/blog/marketing/what-is-amazon-seo/`, accessed March 5, 2025.

[3] The Digital Markets Act is Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector; the Digital Services Act is Regulation (EU) 2022/2065 of 19 October 2022 on a Single Market for Digital Services. Both are supported by the European Centre for Algorithmic Transparency, which has a mandate to audit algorithms used by digital gatekeepers.

[4] This is not lost on the search engine optimization industry, which emphasizes that channeling sales through the platform improves a seller's ranking. For example, BigCommerce lists as its first optimization strategy: "Sales are King... Your first option for generating sales is by driving both internal and external traffic to your Amazon listing" (`https://www.bigcommerce.com/blog/amazon-seo-strategy/`, accessed March 5, 2025) and advises sellers to "Point all product links to Amazon," noting that "Your goal should be to place on the first page of search results — preferably in the top three listings. This is where big money is made, not by a single sale on your site" (`https://www.bigcommerce.com/articles/omnichannel-retail/how-to-sell-on-amazon/`, accessed March 5, 2025).

We analyze this force in a two-period model of a monopoly platform that is the only channel through which consumers discover the seller. Once consumers are aware of the product, however, they may buy either on-platform or through a direct channel. The seller chooses the off-platform price in each period. The platform commits ex ante to a ranking rule and observes only on-platform sales. That informational restriction is central: sales are informative about seller appeal, but they are also manipulable through cross-channel pricing. To isolate this force, we hold the on-platform price fixed. In the baseline, consumers who find the product appealing have homogeneous unit demand on the platform, so the welfare analysis turns on two distortions: *sales-channel misallocation*, which arises when prices steer consumers toward the wrong channel, and *promotion inefficiency*, which arises when exposure is allocated inefficiently across sellers.

We compare three classes of algorithms. A *baseline* algorithm uses only information observable on the platform, so in period 2 it may condition on past on-platform sales but not on the seller's contemporaneous off-platform price. An *OPP* algorithm can condition directly on the seller's off-platform price. A *PPC* algorithm is more limited: it does not otherwise use off-platform information, but it excludes sellers that offer a positive off-platform discount. This distinction is useful both economically and institutionally. Baseline algorithms capture environments in which platforms do not observe all off-platform prices, while parity-type rules are often easier to monitor and enforce than fully general price-contingent ranking rules.

The analysis yields three main conclusions. First, sales-based ranking distorts seller pricing even when the algorithm never conditions directly on off-platform prices. Under a baseline rule, the seller reduces her off-platform discount in period 1 in order to increase on-platform sales and improve her chances of being featured in period 2. The platform exploits this force by sometimes setting a more demanding promotion rule than would be efficient from a period-2 perspective. When it does, the higher cutoff raises on-platform sales in period 1, but it also creates promotion inefficiency in period 2. In sufficiently strong cases, the same force effectively closes the direct channel altogether, even though the algorithm never conditions on off-platform prices. Restricting the information available to the platform is therefore not neutral: a sales-based rule can reproduce some of the incentive effects usually associated with explicit parity restrictions.

Second, allowing the platform to use off-platform information can improve welfare. An OPP algorithm gives the platform direct control over current pricing: by conditioning exposure on the seller offering a sufficiently large off-platform premium, it forces all interested consumers to buy on-platform. Because pricing is then disciplined directly rather than through continuation incentives, the platform no longer needs to distort the period-2 promotion cutoff. This removes promotion inefficiency, but it may worsen the sales-channel misallocation by

shutting down purchases that would efficiently occur off-platform. A PPC algorithm is more limited, but in some parameter regions it is also more effective. When it binds in both periods, it eliminates the platform's incentive to distort the promotion cutoff and equalizes prices across channels. In that region, it removes both promotion inefficiency and the sales-channel misallocation, and therefore attains the (welfare) first best.

Third, welfare and consumer surplus need not rank algorithms in the same way. Because the market is covered in the baseline, higher prices do not reduce total trade; instead, they affect which channel consumers use and which seller receives exposure. An algorithm that improves total surplus by reducing promotion inefficiency may nonetheless reduce consumer surplus by worsening the sales-channel misallocation, and vice versa. A regulator who focuses only on whether an algorithm uses off-platform information, or only on whether direct-channel prices rise or fall, may therefore miss the underlying trade-off and as we discuss in Section 5.4 policy design may be involved and require information off-platform.

1.1 Related literature

This paper sits at the intersection of three literatures: ranking algorithms and platform design, career concerns and dynamic incentives, and price parity restrictions.

Ranking algorithms and platform design. Most of the literature on ranking algorithms studies how rankings shape competition among products or sellers within a platform. Castellini et al. (2023) and Calvano et al. (2025) examine how ranking and recommendation systems affect within-platform competition, while Johnson, Rhodes and Wildenbeest (2023) studies how algorithmic design shapes competition among sellers who use pricing algorithms. Eisfeld (2023) provides empirical evidence on such effects, and Bar-Isaac and Shelegia (2025) contrasts algorithm-based and auction-based product ranking in a setting where the platform is the sole sales channel. Our contribution is to study ranking design when the platform competes with an off-platform channel and learns from sales that the seller can manipulate through cross-channel pricing.

Several recent papers consider the interaction between on- and off-platform pricing in static environments. Hagiu and Wright (2024) studies platform strategies to prevent off-platform leakage when the direct channel is less efficient.[5] Wang and Wright (2020) consider consumer search either on or off the platform and how PPCs might affect commission fees. Karle, Preuss and Reisinger (2026) examines how personalized platform recommendations and commission rates jointly shape sellers' dual-channel pricing through the composition of

[5]Our static model extends theirs by allowing some consumers to derive higher utility from the direct channel, which turns out to matter for the welfare analysis.

demand across channels. Bergemann and Bonatti (2024) and Bergemann, Bonatti and Wu (2025) analyze how platform mechanisms shape the interaction between on- and off-platform prices using consumer-level data. Peitz and Sobolev (2025) show that prohibiting PPCs can reduce welfare, as can arise in our model. In their model, the platform cannot commit to sellers in its reaction to on- and off-platform prices and their result relies on tailored recommendations to heterogeneous consumers. These papers analyze profit-maximizing designs in one-shot settings and therefore abstract from the dynamic incentive effects created by sales history and platform learning.[6] Relative to this literature, our contribution is to show that once a platform learns from on-platform sales, a seller's cross-channel pricing becomes a tool for influencing future exposure.

Career concerns and platforms with sellers of unknown type. The paper is also related to the literature on career concerns and dynamic incentives following Holmström (1999) and Bar-Isaac and Tadelis (2008) for a broad survey. In standard career-concerns models, current actions affect future rewards by shaping beliefs about an agent's type. Here, off-platform pricing plays that role: the seller restrains off-platform discounts not because doing so is statically profitable, but because stronger on-platform sales improve future promotion. The application differs from the standard career-concerns setup in two respects. The platform commits to an algorithmic promotion rule rather than relying on market-determined continuation payoffs, and the seller's strategic variable is a cross-channel price that changes both current allocation and future beliefs, rather than hidden effort that affects only beliefs.

A smaller literature studies platforms that interact with sellers of unknown type. Jiang, Jerath and Srinivasan (2011) and Madsen and Vellodi (2025) study environments in which the platform may enter the seller's market as a competitor, so sellers have an incentive to conceal information about high types in order to deter entry. We study the opposite incentive: sellers want to reveal, or even exaggerate, high types in order to secure continued exposure through the ranking algorithm.

Most-favored-nation and price-parity clauses. Our paper is also related to the literature on explicit contractual restrictions on sellers' prices across channels. Scott Morton (2012), Boik and Corts (2016), Johnson (2017), and Enache and Rhodes (2025*b*) study the welfare effects of most-favored-nation and price-parity clauses in a range of environments. Much of the existing literature emphasizes pricing distortions that create deadweight loss by raising prices above efficient levels. By contrast, our model studies how algorithms can generate PPC-like outcomes even in the absence of explicit contractual restrictions. To isolate

[6] Enache and Rhodes (2025*a*) is an exception in allowing dynamic interaction, focusing on how front-loaded platform pricing mitigates disintermediation though it does not consider learning.

this mechanism and highlight other effects (specifically, promotional inefficiency and sales-channel misallocation) we assume full market coverage, thereby abstracting from the more standard deadweight-loss channel.

Closest to our paper, Hunold et al. (2018) analyzes the Booking.com case, and Peitz (2022) and Franck and Peitz (2024) argue that algorithmic design may circumvent regulatory restrictions on PPCs. Our contribution is to make that concern precise: we show how and when a sales-based ranking rule can reproduce the incentive effects of an explicit PPC without any contractual restriction, and we characterize parameter regions in which the algorithmic "shadow PPC" is more distortive than an explicit one.

The limited empirical evidence on PPCs is broadly consistent with this mechanism. Hunold et al. (2018) and Ennis, Ivaldi and Lagos (2023) find that direct-channel prices fall relative to OTA prices after PPC bans in some countries and for certain hotel types, while Mantovani, Piga and Reggiani (2021) and Ma et al. (2024) find short-lived or small effects. Both patterns are consistent with our analysis: removing PPCs has no effect if the platform's ranking algorithm adapts to restore parity, but it can lower direct-channel prices if the induced algorithmic distortion is too weak to replicate PPC. In the same setting, Hunold, Kesler and Laitenberger (2020) shows that online travel agents rank hotels lower when their direct-channel prices are low, which is directly aligned with our mechanism.[7]

The rest of the paper is organized as follows. Section 2 presents the model. Section 3 develops a static benchmark. Section 4 studies dynamic equilibrium under baseline, OPP, and PPC algorithms. Section 5 compares welfare and consumer surplus across algorithm classes. Section 6 considers extensions and robustness.

2 Model

There are four sets of agents: a monopoly platform, the seller, a reserve seller, and consumers. All are risk neutral. Neither the platform nor the sellers incur costs. Consumers and each period's reserve seller are short-lived. The platform and the seller maximize expected profits summed over two periods, with no discounting between periods; equivalently, the parameter $\beta \geq 0$, described below, captures both the mass of period-2 consumers and any discounting.

Platform. The platform gives the seller access to consumers in exchange for a fee $f \in [0, 1)$ on each on-platform sale. The seller may also sell through a direct, off-platform channel, but that channel is accessible only to consumers who first become aware of the product through the platform. If the platform features the seller, it forgoes the revenue it could have earned

[7] Hunold, Laitenberger and Thébaudin (2022) documents a similar mechanism for Amazon's Buy Box.

by allocating the same exposure to the reserve seller. This forgone value is what generates the possibility of inefficient promotion in the analysis below. Before period 1, the platform commits to a ranking algorithm, that is, a rule that maps available information, including off-platform discounts when the algorithm class allows it but not off-platform sales, into a featuring decision in each period. The algorithm classes we study are defined below.

Seller. The seller is characterized by a type $\theta \sim U[0,1]$. We interpret type θ as the fraction of consumers who find the product appealing once they become aware of it. Initially, neither the platform nor the seller observes the realized type before period 1, so both share this prior. Section 6 considers alternative information structures that vary who knows θ.

In each period $t \in \{1,2\}$ in which it is active, the seller chooses an off-platform discount $\Delta_t \in \mathbb{R}$, where $\Delta_t > 0$ means that the off-platform price is below the on-platform price and $\Delta_t < 0$ means that the seller charges an off-platform premium. Since consumers who buy on-platform have valuation 1, the on-platform price is fixed at 1. The seller therefore earns $1 - f$ from an on-platform sale and $1 - \Delta_t$ from an off-platform sale.

Reserve seller(s). In each period, the platform has a reserve seller as an outside option. This seller is short-lived, so one can think of a fresh reserve seller arriving in each period; for notational ease, we refer to these period-specific outside options collectively as "the reserve seller." It has known type $\alpha \in [0, 1/2)$, with α interpreted analogously to θ. Apart from being short-lived and of known type, it faces the same demand and pricing environment as the seller. Throughout, unqualified references to "the seller" refer to the type-θ seller, while "reserve seller" refers to this outside option. Because the platform controls exposure, consumers see at most one of the two sellers in a given period. Thus the model abstracts from direct consumer-side competition between them, even though they may compete for the platform's exposure slot.[8]

Consumers. A unit mass of consumers arrives in period 1 and a mass β arrives in period 2. Consumers observe whichever seller the platform features, and the featured seller then faces the same demand system in that period.

If the featured seller has type $x \in \{\theta, \alpha\}$, then a fraction x of consumers who observe that seller finds the product appealing. Conditional on finding the featured seller's product appealing, a consumer may buy either on-platform or off-platform. Consumers differ in their disutility from buying off-platform. Let δ denote consumer-specific off-platform disutility,

[8] We therefore do not study seller competition here. In earlier work, Bar-Isaac and Shelegia (2025), we study seller competition and show that when competition is mediated by an algorithm that controls on-platform prices, the equilibrium outcome can coincide with the no-competition benchmark.

distributed uniformly on $[\underline{\delta}, 1]$, where $\underline{\delta} \leq 0$. Allowing $\underline{\delta} < 0$ means that some consumers prefer the off-platform channel even when prices are equal. Given the featured seller's discount Δ_t, a consumer who likes the product buys on-platform if and only if $\delta \geq \Delta_t$. Thus, conditional on discovery and on finding the product appealing, the market is covered: the consumer buys either on-platform or off-platform. Hence, among consumers who find the product appealing, the fraction who buy on-platform is

$$q(\Delta_t) \equiv \frac{1 - \Delta_t}{1 - \underline{\delta}}, \tag{1}$$

and the remaining fraction $1 - q(\Delta_t)$ buys off-platform.

Algorithm classes. In what follows, ranking is binary: in each period the platform features either the seller or the reserve seller. We now introduce the three classes of algorithms that we compare. Since $\alpha < 1/2$, it is easy to show that on path the seller is featured in period 1.[9] We therefore write algorithms as if they condition only on the seller's observables, suppressing arguments involving the reserve seller, and let ϕ_t denote the probability that the seller is featured in period t under the algorithm.[10]

- A **baseline algorithm** uses only information observable on the platform. In period 2, the seller's featuring probability is $\phi_2(S_1)$, where S_1 denotes period-1 on-platform sales.
- An **OPP algorithm** (off-platform price algorithm) can condition on the contemporaneous off-platform discount. In period 1, the seller's featuring probability is $\phi_1^{OPP}(\Delta_1)$, while in period 2 it is $\phi_2^{OPP}(\Delta_2, S_1)$.
- A **PPC algorithm** (price parity clause algorithm) is a baseline algorithm augmented with the constraint that any seller offering a positive off-platform discount is excluded from being featured. Thus the seller's featuring rule can be written as $\phi_1^{PPC}(\mathbf{1}_{\{\Delta_1>0\}})$ in period 1 and $\phi_2^{PPC}(\mathbf{1}_{\{\Delta_2>0\}}, S_1)$ in period 2, where $\mathbf{1}_{\{\Delta_t>0\}}$ is an indicator function that takes the value 1 if and only if the seller gives a positive off-platform discount (charges a lower off-platform price) in period t.

The first two classes are natural benchmarks. In practice, platforms may not have access to all off-platform prices — a hotel, for instance, can offer idiosyncratic quotes by telephone

[9] Ex ante, the seller's expected type is $E[\theta] = 1/2$, whereas the reserve seller has known type α.

[10] For example, writing Δ_t^r for the reserve seller's off-platform discount, a fully general OPP algorithm could in period 1 condition on (Δ_1, Δ_1^r), and in period 2 on $(\Delta_1, \Delta_1^r, \Delta_2, \Delta_2^r)$ together with the period-1 sales of whichever seller was featured in period 1. We suppress these arguments because on path the seller receives the period-1 slot, and because the reserve seller is short-lived, so one can think of a fresh reserve seller in each period and hence its sales do not generate dynamic incentive effects.

that would be costly to monitor at scale — which motivates the baseline algorithm. The OPP algorithm is an extreme whereby the platform can use all information. Price parity clauses, by contrast, are a focal point of regulatory concern and are typically easier to monitor and enforce than fully general OPP algorithms, which motivates the PPC class. Unless otherwise specified, we refer to the baseline algorithm simply as *the algorithm*.

Timing and equilibrium. The timing is as follows.

1. Before period 1, the platform commits to an algorithm and the seller observes it.
2. In each period $t = 1, 2$, the seller, or, when relevant, the reserve seller, chooses off-platform discounts.
3. The platform then features either the seller or the reserve seller according to the algorithm.
4. Consumers observe the featured seller's discount and decide whether and where to buy.

Pricing precedes the featuring decision because under OPP and PPC algorithms the platform conditions directly on Δ_t. Under the baseline algorithm, this within-period ordering is immaterial since Δ_t does not enter the rule directly.

Our solution concept is subgame-perfect equilibrium. The platform commits to its algorithm before any learning takes place, the seller does not observe its type in period 1, and the seller's period-2 pricing choice is myopic. One can think of the platform as using observed period-1 on-platform sales S_1, together with anticipated seller pricing, to infer θ; because the algorithm's behavior is fixed ex ante, however, these "beliefs" are only an interpretive device. As a result, the relevant strategic issues concern continuation play after observed histories rather than belief refinements, so subgame perfection is the natural equilibrium concept.

3 Static Benchmark

This section characterizes the seller's myopic pricing problem and the corresponding period-2 outcome under each algorithm class. By period 2, the seller has learned its realized type θ. Conditional on being featured, however, either the seller or the reserve seller faces the same static pricing problem: type only scales the mass of interested consumers, so the optimal discount and induced on-platform share are the same for both. In period 2 there are no dynamic incentives. Under the baseline algorithm, off-platform pricing affects only current channel allocation. Under OPP and PPC, it may also affect whether the seller satisfies

the platform's featuring condition, but conditional on being featured it again affects only current channel allocation. The section therefore isolates the sales-channel force that will later interact with learning and promotion incentives in period 1.

3.1 Baseline algorithm

Suppose the algorithm features a seller in period 2. For $\Delta \leq \underline{\delta}$, all interested consumers buy on-platform and the seller earns $1 - f$ per interested consumer regardless of Δ, so discounts below $\underline{\delta}$ are payoff-equivalent. For $\Delta > 1$, the seller subsidizes off-platform purchases, which is never optimal. We may therefore restrict attention to $\Delta \in [\underline{\delta}, 1]$ without loss.

Conditional on a consumer finding the product appealing, the seller's profit given an off-platform discount Δ is

$$\pi(\Delta) \equiv 1 - f + (1 - q(\Delta))(f - \Delta). \tag{2}$$

The first term, $1 - f$, is the profit the seller earns on any sale. The second term captures the incremental gain from inducing an off-platform purchase: the seller avoids the platform fee f, but gives up the discount Δ, on the fraction $1 - q(\Delta)$ of consumers who buy off-platform.

Proposition 1. *Under the algorithm, if a seller is featured in period 2, its optimal off-platform discount is*

$$\Delta_2^* = \frac{f + \underline{\delta}}{2}. \tag{3}$$

This discount is positive if $f > -\underline{\delta}$, and negative if $f < -\underline{\delta}$. The corresponding on-platform share, seller profit per interested consumer, and platform revenue per interested consumer are

$$q_2^* \equiv \frac{2 - \underline{\delta} - f}{2(1 - \underline{\delta})}, \qquad \pi_2^* \equiv \frac{(2 - \underline{\delta} - f)^2}{4(1 - \underline{\delta})}, \qquad r_2^* \equiv \frac{f(2 - \underline{\delta} - f)}{2(1 - \underline{\delta})}.$$

Proof. Differentiating (2) with respect to Δ and setting the derivative to zero yields $\Delta_2^* = (f + \underline{\delta})/2$. Since π is concave in Δ, this is the unique maximizer on $[\underline{\delta}, 1]$. The expressions for q_2^*, π_2^*, and r_2^* follow by substitution. □

The seller trades off the fee savings from diverting demand off-platform against the need to offer a lower off-platform price in order to attract that demand. Because some consumers may intrinsically prefer the direct channel ($\underline{\delta} < 0$), the optimal "discount" can be negative: when f is small relative to $-\underline{\delta}$, the seller charges a premium off-platform and still attracts some direct sales from consumers with strong channel preferences.

Since the market is covered, welfare is maximized when consumers buy through the channel they prefer, which requires no price differential across channels. Thus, when some consumers prefer buying off-platform (that is, $\underline{\delta} < 0$), a positive platform fee can raise welfare by inducing the seller to lower the off-platform price and thereby move the allocation of sales closer to the efficient benchmark.

3.2 OPP algorithm

Suppose the platform can condition its featuring decision on the seller's contemporaneous off-platform discount. Because this conditioning occurs *before* sales are realized—at the featuring stage rather than ex post—the platform has a direct control instrument over the seller's pricing. In particular, it can require $\Delta_2 \leq \underline{\delta}$ as a condition for exposure, forcing all interested consumers onto the platform. If the seller violates this condition, it is not featured and makes no sales through either channel. Clearly, this is what the platform will choose to do.

Proposition 2. *Under the optimal OPP algorithm, the seller sets $\Delta_2^{OPP} = \underline{\delta}$, so all consumers who find the product appealing buy on-platform. Relative to the baseline algorithm, consumer surplus is weakly lower. Welfare is higher when $f > -3\underline{\delta}$ and lower when $f < -3\underline{\delta}$.*

Proof. If featured, the seller complies with the condition $\Delta_2 \leq \underline{\delta}$ because the alternative is zero sales, so under OPP all interested consumers buy on-platform.

Consumer surplus under OPP is therefore zero. Under the algorithm, only consumers with $\delta < \Delta_2^*$ buy off-platform and obtain surplus $\Delta_2^* - \delta$, so expected consumer surplus per interested consumer is

$$CS_2^* \;=\; \int_{\underline{\delta}}^{\Delta_2^*} \frac{\Delta_2^* - \delta}{1 - \underline{\delta}}\, d\delta \;=\; \frac{(\Delta_2^* - \underline{\delta})^2}{2(1 - \underline{\delta})} \;=\; \frac{(f - \underline{\delta})^2}{8(1 - \underline{\delta})},$$

where the last equality uses $\Delta_2^* = (f + \underline{\delta})/2$. Hence consumer surplus is weakly lower under OPP.

Because the market is covered, welfare under OPP is simply $W_2^{OPP} = 1$. Under the algorithm,

$$W_2^* \;=\; 1 \;-\; \int_{\underline{\delta}}^{\Delta_2^*} \frac{\delta}{1 - \underline{\delta}}\, d\delta \;=\; 1 \;-\; \frac{(\Delta_2^*)^2 - \underline{\delta}^2}{2(1 - \underline{\delta})} \;=\; 1 \;-\; \frac{(f - \underline{\delta})(f + 3\underline{\delta})}{8(1 - \underline{\delta})}.$$

Therefore

$$W_2^{OPP} - W_2^* \;=\; \frac{(f - \underline{\delta})(f + 3\underline{\delta})}{8(1 - \underline{\delta})}.$$

Since $f - \underline{\delta} > 0$, welfare is higher under OPP when $f > -3\underline{\delta}$ and lower when $f < -3\underline{\delta}$. □

The welfare comparison reflects two offsetting forces. Under the baseline algorithm, the seller may set a high off-platform discount, drawing some consumers off-platform even though they would obtain greater utility on-platform. OPP eliminates this misallocation but also forecloses the direct channel entirely, preventing consumers who genuinely prefer buying directly from doing so. The fee f governs the depth of the discount under the baseline algorithm and therefore the severity of the first distortion; $\underline{\delta}$ governs consumers' intrinsic preference for the direct channel and therefore the cost of the second. When f is large relative to $-\underline{\delta}$, the first distortion dominates and OPP raises welfare; when f is small, the foreclosure cost dominates.

3.3 PPC algorithm

A PPC algorithm cannot impose an arbitrary pricing condition for exposure, but it can exclude any seller who offers a strictly positive off-platform discount. It therefore either has no effect or binds at price parity.

Proposition 3. *Under a PPC algorithm, the seller's period-2 off-platform discount is*

$$\Delta_2^{PPC} \;=\; \min\{0,\, \Delta_2^*\}.$$

If the baseline algorithm already has $\Delta_2^* \leq 0$, the PPC is slack and outcomes are identical. If instead $\Delta_2^* > 0$, the PPC binds and $\Delta_2^{PPC} = 0$. Consumer surplus under PPC is weakly lower than under the baseline algorithm but weakly higher than under OPP. Welfare under PPC is weakly higher than under both alternatives.

Proof. When $\Delta_2^* \leq 0$, the constraint $\Delta_2 \leq 0$ is slack and all outcomes coincide with the baseline algorithm. When $\Delta_2^* > 0$, the constraint binds and the seller sets $\Delta_2 = 0$. Consumer surplus is weakly lower than under the baseline algorithm because consumers either face the same discount or face a zero discount rather than a positive one; it is weakly higher than under OPP because OPP yields zero consumer surplus. Welfare is higher under the PPC algorithm than under the OPP or baseline algorithm. This is because welfare under the PPC algorithm is the same as under the baseline algorithm for $f < -\underline{\delta}$, in which case the latter is in turn higher than under OPP, or is at first-best if $f > -\underline{\delta}$ and thus weakly better than under OPP or the baseline algorithm.

□

When it binds, a PPC equalizes prices across channels without shutting down the direct channel. Since the market is covered, welfare then changes only through the sales channel:

consumers sort into the channel they prefer, rather than being steered by a price wedge. In the static benchmark, PPC is therefore weakly welfare-superior to both the baseline and OPP alternatives. That static comparison will serve as a benchmark for the dynamic welfare analysis in Section 5. This static welfare ranking relies on the covered-market environment: with partial market coverage, parity would also affect whether some consumers buy at all.

The objects Δ_2^*, q_2^*, and π_2^* defined in this section recur throughout the analysis. In period 1, however, pricing affects not only current channel choice but also future exposure through its impact on observed on-platform sales—the career-concerns distortion studied next.

4 Dynamic Equilibrium

We now turn to the full dynamic problem. Neither the platform nor the seller knows the realized type θ before period 1. As a result, first-period on-platform sales play a dual role: they are the platform's signal about seller appeal and, at the same time, the seller's instrument for influencing future exposure. This is what generates the career-concerns force at the heart of the paper. By raising her off-platform price, the seller shifts demand on-platform and thereby improves her chances of being featured in period 2. We begin with the baseline algorithm. There we first characterize the seller's period-1 problem and then derive the platform's optimal algorithm. We then compare the OPP and PPC alternatives.

4.1 Baseline algorithm

Broadly, a tougher retention rule induces the seller to keep more demand on-platform in period 1, but it also risks inefficiently replacing the seller with the reserve seller in period 2. This is the trade-off that the platform balances in its design of the algorithm, incorporating its expectation of the seller's behavior.

Before characterizing first the seller's problem and then the platform's problem, we prepare the ground with a lemma that simplifies the set of period-2 allocation rules that need to be considered.

Under the baseline algorithm, period-2 exposure can depend only on the period-1 on-platform sales signal $S_1 = \theta q(\Delta_1)$. We can therefore summarize the platform's period-2 policy by the seller's induced allocation as a function of S_1. The next lemma shows that, without loss, this allocation rule can be taken to be a *threshold rule*: there exists a threshold T such that the seller is featured in period 2 if and only if $S_1 \geq T$.

The intuition for this result (formalized in Appendix A.1) is that if the platform does

not use a threshold algorithm, it can redesign the algorithm to shift probability mass from featuring low-sales seller types to featuring high-sales seller types, while preserving the seller's off-platform pricing incentives. Such a reallocation necessarily improves promotion efficiency and raises the platform's period-2 profits, since higher-quality sellers are then more likely to be featured.

Lemma 1. *Among baseline algorithms, there exists an optimal threshold rule.*

Proof. See Appendix A.1. The appendix shows that any optimal non-threshold rule can differ from a threshold rule only off the equilibrium support of sales or on a measure-zero subset of that support. □

Lemma 1 reduces the platform's continuation rule to a one-dimensional object, the sales threshold T.

4.1.1 The seller's problem

Given Lemma 1, let T denote the period-2 sales threshold. The seller's expected profit over both periods, anticipating the period-2 benchmark discount Δ_2^* from Section 3, is

$$\Pi_{\text{seller}}(\Delta_1; T) \;=\; \tfrac{1}{2}\pi(\Delta_1) \;+\; \beta\pi_2^* \int_{\min\{1,\, T/q(\Delta_1)\}}^{1} \theta \, d\theta. \tag{4}$$

The coefficient $\frac{1}{2}$ in the first term is the prior expectation $\mathbb{E}[\theta]$, and the second term is the expected continuation payoff from being featured in period 2.

It is convenient to re-parametrize the platform's choice by the *induced cutoff* type of seller

$$\tau(T) \;\equiv\; \frac{T}{q(\Delta_1^*(T))},$$

where $\Delta_1^*(T)$ is the seller's equilibrium discount under threshold T. In equilibrium, the seller is featured in period 2 if and only if $\theta \geq \tau$. Appendix A.2 shows that on the relevant region the map $T \mapsto \tau(T)$ is one-to-one, so we index equilibria by τ rather than by T.

Not every cutoff is implementable: if the platform sets the threshold too high, the seller may prefer to price myopically at Δ_2^* and give up on continuation. Let $\bar{\tau}$ denote the highest implementable cutoff, characterized in Appendix A.2. The appendix also shows that the platform never chooses a non-implementable cutoff, so only $\tau \leq \bar{\tau}$ are relevant. For ease of exposition, we focus in the main text on the parameter region in which implementability does not bind at the platform's optimum and provide a characterization of the case where it does bind in the Appendix, showing qualitatively similar effects.

Assumption 1. $\beta > \underline{\beta} \equiv (f - \underline{\delta})(2 + f - 3\underline{\delta})/(2 - \underline{\delta} - f)^2$.

Economically, $\underline{\beta}$ is the level of continuation importance above which the seller's career-concerns incentives are strong enough to push her to the lower bound $\underline{\delta}$ before the implementability constraint binds. What is ruled out is the low-β region in which the platform would like a tougher rule, but the seller would instead give up on continuation and revert to myopic pricing. Appendix A.2 treats that case (where Assuption 1 is violated) separately.

We can now characterize the seller's period-1 discount for any implementable cutoff.

Proposition 4. *Suppose Assumption 1 holds, and let $\tau \in [0, \bar{\tau}]$. Define*

$$\bar{\beta}(\tau) \;\equiv\; \frac{\Delta_2^* - \underline{\delta}}{\pi_2^* \tau^2}.$$

If $\beta \geq \bar{\beta}(\tau)$, then $\Delta_1^*(\tau) = \underline{\delta}$. Otherwise, the optimal discount is interior and given by

$$\Delta_1^*(\tau) \;=\; \Delta_2^* \;-\; \tfrac{1}{4}(2 - f - \underline{\delta}) \left(\sqrt{1 + 4\beta\tau^2} - 1 \right), \tag{5}$$

which is strictly decreasing in β and τ, and strictly increasing in f and $\underline{\delta}$.

Proof. For any implementable cutoff, the seller's choice is either the corner $\Delta_1 = \underline{\delta}$ or an interior optimum. In the interior case, differentiating (4) with respect to Δ_1 and evaluating the derivative at an equilibrium with induced cutoff τ yields

$$\Delta_1 \;=\; \Delta_2^* \;-\; \frac{\beta \pi_2^* \tau^2}{q(\Delta_1)}. \tag{6}$$

Solving (6) gives the interior expression (5). The interior solution is feasible if and only if it exceeds the lower bound $\underline{\delta}$. Rearranging $\Delta_1^*(\tau) \leq \underline{\delta}$ is equivalent to

$$\beta \geq \frac{2(1 - \underline{\delta})(f - \underline{\delta})}{(2 - \underline{\delta} - f)^2 \tau^2} = \frac{\Delta_2^* - \underline{\delta}}{\pi_2^* \tau^2} = \bar{\beta}(\tau),$$

and in that case the seller chooses the corner solution $\Delta_1^*(\tau) = \underline{\delta}$. The comparative statics follow immediately from (5). □

The threshold $\bar{\beta}(\tau)$ has a simple interpretation. The numerator $\Delta_2^* - \underline{\delta}$ is the additional distortion, relative to the myopic benchmark, needed to push the seller all the way to the lower bound $\underline{\delta}$, while the denominator $\pi_2^* \tau^2$ measures the continuation-value gain from increasing period-1 sales around the cutoff type. Thus higher continuation stakes—either because period 2 matters more or because the cutoff is more demanding—lead the seller to cut the

period-1 off-platform discount further. In particular, for every $\tau > 0$ we have $\Delta_1^*(\tau) < \Delta_2^*$: this is the formal career-concerns distortion in the model.

The comparative statics are intuitive. Higher β or higher τ both strengthen the seller's incentive to shade her off-platform discount downward: in the first case future exposure is more valuable, in the second current on-platform sales are more pivotal for securing it. A higher fee f or higher $\underline{\delta}$, by contrast, raises the myopic benchmark $\Delta_2^* = (f + \underline{\delta})/2$ and lowers the continuation payoff π_2^* from being featured in period 2. The seller therefore both starts from a higher static optimum and has less to gain from distorting period-1 pricing to secure future exposure, so Δ_1^* rises.

4.1.2 Optimal algorithm

We now turn to the platform's profit-maximizing algorithm. The platform chooses the induced cutoff τ to maximize expected two-period profits, trading off period-2 selection against period-1 fee revenue through the cutoff's effect on the seller's off-platform pricing. Recall from Section 3 that q_2^* is the period-2 on-platform sales share under the baseline benchmark. Featuring the reserve seller under the baseline benchmark yields period-2 platform revenue $f\alpha q_2^*$. As discussed above, period 1 starts from exposing the seller because ex ante $E[\theta] = 1/2 > \alpha$ and the seller does not know her type before period 1. Absent any effect on period-1 pricing, the platform would in period 2 feature the seller if and only if $\theta \geq \alpha$. A platform without commitment power would therefore set $\tau = \alpha$. With commitment, however, the platform may distort the cutoff upward in order to influence the seller's period-1 pricing.

To compare the baseline, OPP, and PPC algorithms on a common scale, from now on we hold fixed the reserve seller's type α. Under OPP and PPC, the reserve seller is subject to the same period-2 pricing rule as the seller under the algorithm being analyzed.

As discussed above, under Assumption 1, every cutoff in a neighborhood of the platform's optimum is implementable, so the platform's problem reduces to choosing $\tau \in [0, 1]$ to maximize

$$\Pi_{\text{platform}}(\tau) \;=\; f\left[\tfrac{1}{2}q(\Delta_1^*(\tau)) \;+\; \beta\left(\alpha\tau + \tfrac{1-\tau^2}{2}\right) q_2^*\right]. \tag{7}$$

The first term in the square brackets is expected period-1 on-platform revenue. The second is expected period-2 revenue, combining reserve-seller revenue when the seller's type is below τ and the platform's revenue from the seller when $\theta \geq \tau$.

Differentiating (7) with respect to τ gives

$$\Pi'_{\text{platform}}(\tau) \;=\; f\left[(\alpha - \tau)\beta q_2^* \;-\; \tfrac{1}{2(1-\underline{\delta})}\tfrac{\partial \Delta_1^*}{\partial \tau}\right].$$

Absent the effect of τ on the seller's discount, the optimum would be $\tau = \alpha$. This is intuitive:

if the induced cutoff had no impact on the seller's off-platform discount, then the platform would feature the seller if and only if it were better than the reserve seller. However, since the threshold affects the first period discount, specifically $\partial\Delta_1^*/\partial\tau < 0$, the platform has a local incentive to raise the cutoff above α: a higher cutoff induces the seller to reduce her period-1 off-platform discount, shifting more sales on-platform and increasing fee revenue.

On the interior region where $\Delta_1^*(\tau) > \underline{\delta}$, the first-order condition simplifies to

$$\alpha \;=\; \tau\left(1 - \frac{1}{\sqrt{1+4\beta\tau^2}}\right). \tag{8}$$

Proposition 5. *Let Assumption 1 hold. Define*

$$\bar{\alpha} \;\equiv\; \sqrt{\frac{\Delta_2^* - \underline{\delta}}{\beta\pi_2^*}}, \qquad \underline{\alpha} \;\equiv\; \frac{2(1-q_2^*)}{2-q_2^*}\bar{\alpha}.$$

Then:

(i) If $0 \le \alpha < \underline{\alpha}$, the optimal cutoff $\tau^* \in [0, \bar{\alpha})$ is the unique solution to (8), with $\tau^* = 0$ if and only if $\alpha = 0$, and otherwise $\Delta_1^*(\tau^*) > \underline{\delta}$.

(ii) If $\underline{\alpha} \le \alpha < \bar{\alpha}$, then $\tau^* = \bar{\alpha}$ and $\Delta_1^*(\tau^*) = \underline{\delta}$.

(iii) If $\alpha \ge \bar{\alpha}$, then $\tau^* = \alpha$ and $\Delta_1^*(\tau^*) = \underline{\delta}$.

The second or third case may be empty under the maintained restriction $\alpha < 1/2$.

Proof. See Appendix A.3.

□

Figure 1 illustrates the three regions in Proposition 5. For low values of the reserve seller's type α, the platform distorts the cutoff upward above the efficient benchmark, accepting some promotion inefficiency in exchange for lower off-platform discounts and higher period-1 fee revenue. For intermediate values of α, the seller's discount has bottomed out at $\underline{\delta}$, so further upward distortions raise promotion inefficiency without any offsetting gain from period-1 pricing; the platform therefore chooses the kink cutoff $\bar{\alpha}$. For sufficiently high values of α, the efficient cutoff $\tau^* = \alpha$ already induces the seller to foreclose the off-platform channel, so no further distortion is useful.

A notable feature of Proposition 5 is that, on the interior region, the optimal induced cutoff depends only on the efficient benchmark α and the learning parameter β: the fee f and the channel-preference parameter $\underline{\delta}$ drop out of the platform's first-order condition in

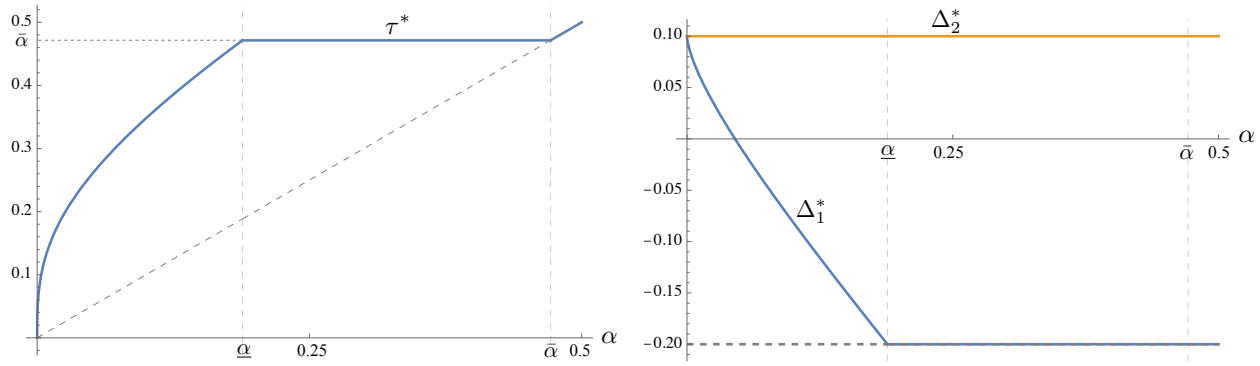

(a) Optimal induced cutoff

(b) First-period discount

Figure 1: Optimal baseline-algorithm outcomes as functions of the reserve seller's type α when $\underline{\delta} = -0.2$, $f = 0.4$ and $\beta = 2$. The left panel plots the induced cutoff τ^*; the right panel plots the optimal first-period discount (blue) and the myopic/second-period discount (orange).

Equation 8. Intuitively, q_2^* enters both sides of the platform's first-order condition symmetrically and cancels, so the pricing primitives affect outcomes only through the endogenous boundaries $\underline{\alpha}$ and $\overline{\alpha}$ that determine which region applies.[11]

4.2 OPP algorithm

When the platform can condition its algorithm on off-platform prices, as in the static analysis, it can require $\Delta_t \leq \underline{\delta}$ as a prerequisite for exposure. The same requirement applies to the reserve seller, so whenever it is featured it behaves exactly as in the static OPP benchmark in Section 3. The algorithm therefore effectively forecloses the off-platform channel in both periods.

Proposition 6. *Suppose $\alpha < 1/2$. Under the optimal OPP algorithm,*

$$\Delta_1^{OPP} = \Delta_2^{OPP} = \underline{\delta}, \qquad \tau^{OPP} = \alpha.$$

Proof. Whenever a seller is featured, the platform wants all consumers who find the product appealing to buy on-platform, since it earns fees only on on-platform sales. Under an OPP algorithm, it can achieve this directly by conditioning exposure on $\Delta_t \leq \underline{\delta}$. Any seller that violates this condition is not featured and therefore makes no sales in that period, so any featured seller optimally complies. Thus, under commitment both sellers are disciplined to set the minimum admissible discount, and competition for the slot is not through softer

[11] The interior-region boundary $\underline{\alpha}$ is strictly increasing in f and strictly decreasing in $\underline{\delta}$ and β; the same is true of the kink threshold $\overline{\alpha}$. Thus a higher fee or lower $\underline{\delta}$ enlarges the interior region (making an interior cutoff more likely), while a higher β shrinks it.

pricing but through which seller delivers higher platform value once both comply. Thus, in each featured period the seller sets $\Delta_t^{OPP} = \underline{\delta}$.

In period 1, this implies that all consumers who find the product appealing buy on-platform, so observed sales satisfy $S_1 = \theta$. The platform therefore learns the seller's type perfectly after period 1. In period 2, conditional on again enforcing $\Delta_2 \leq \underline{\delta}$, featuring a seller of type θ yields higher profits than featuring the reserve seller if and only if $\theta \geq \alpha$. Hence the platform's optimal second-period rule is to feature if and only if $S_1 = \theta \geq \alpha$, which implies the induced cutoff $\tau^{OPP} = \alpha$. □

The result extends the static OPP benchmark in two ways. First, by enforcing $\Delta_1 \leq \underline{\delta}$ in period 1, the platform shifts all interested consumers on-platform, so observed sales satisfy $S_1 = \theta$. Second, because current pricing is now disciplined directly, the platform no longer needs to manipulate the promotion cutoff to affect period-1 incentives, and can therefore implement the efficient period-2 cutoff $\tau^{OPP} = \alpha$ with no promotion inefficiency. OPP therefore eliminates promotion inefficiency, but it does so by foreclosing the direct channel, which may worsen the sales-channel misallocation when some consumers intrinsically prefer buying off-platform.

4.3 PPC algorithm

The PPC algorithm is more subtle than OPP because the constraint $\Delta_t \leq 0$ introduces an additional corner into the seller's period-1 problem: rather than choosing between an interior discount and the lower bound $\underline{\delta}$, the seller may also be pinned at price parity, $\Delta_1 = 0$. The same PPC restriction applies to the reserve seller, so whenever it is featured it behaves exactly as in the static PPC benchmark in Section 3. In the main text we focus on two polar cases, which deliver the sharpest insights. Appendix A.4 characterizes the intermediate cases in which PPC binds in one period but not the other.

First consider the case in which the PPC is slack in period 2. The next proposition shows that this implies that the PPC is also slack in period 1, so the optimal PPC algorithm coincides with the optimal baseline algorithm.

Proposition 7. (PPC slack in both periods.) *If $f < -\underline{\delta}$, then the optimal PPC algorithm coincides with the optimal baseline algorithm, and so:*

$$\Delta_1^{PPC} = \Delta_1^*, \qquad \Delta_2^{PPC} = \Delta_2^*, \qquad \tau^{PPC} = \tau^*.$$

Proof. If $f < -\underline{\delta}$, then $\Delta_2^* = (f + \underline{\delta})/2 < 0$, so the PPC has no bite in period 2: both the seller and the reserve seller choose Δ_2^* whenever featured. Proposition 4 then implies that,

for any implementable cutoff, $\Delta_1^*(\tau) \leq \Delta_2^* < 0$; the career-concerns force only pushes the period-1 discount further below zero. Hence the PPC constraint $\Delta_t \leq 0$ is slack in both periods, and for every implementable cutoff the platform's profit coincides with that under the baseline algorithm. □

More interesting is the case where the PPC binds in both periods.

Proposition 8. (PPC binding in both periods.) *If the PPC binds in both periods, then at the optimal PPC algorithm*

$$\Delta_1^{PPC} = \Delta_2^{PPC} = 0, \qquad \tau^{PPC} = \alpha.$$

In particular, when $\underline{\delta} = 0$, the PPC binds in both periods if and only if $f > 0$. More generally, whenever the PPC binds in the static problem, it binds in both periods for α sufficiently close to 0.

Proof. In the case that the PPC binds in both periods, then whenever either seller is featured it sets $\Delta_t = 0$. In particular, $\Delta_1^{PPC} = 0$ for the seller, and Proposition 3 gives $\Delta_2^{PPC} = 0$ in period 2. In a neighborhood of the optimal cutoff, Δ_1^{PPC} is constant at 0, so a marginal change in τ affects only period-2 expected revenue. If $\tau^{PPC} < \alpha$, raising τ increases period-2 expected revenue while leaving Δ_1^{PPC} unchanged, contradicting optimality; if $\tau^{PPC} > \alpha$, lowering τ has the same effect. Hence $\tau^{PPC} = \alpha$. When $\underline{\delta} = 0$, the seller never prefers a negative discount, so the PPC binds in both periods exactly when $f > 0$. The claim for the case in which the PPC binds in the static problem and α is close to 0 follows from the analysis of intermediate PPC cases in Appendix A.4. □

When the PPC binds in both periods, the seller's period-1 discount is fixed at parity, so the platform does not manipulate pricing through the cutoff. The only remaining role for τ is period-2 selection, and the platform therefore chooses the efficient cutoff $\tau^{PPC} = \alpha$. Thus, in this regime, the PPC removes promotion inefficiency while also eliminating the price wedge across channels. This is the key positive result for the PPC and it drives the welfare comparison in the next section.

5 Welfare

Two distortions drive the welfare comparisons in this paper. *Sales-channel misallocation* arises whenever on- and off-platform prices differ, causing some consumers to buy through the channel they value less. *Promotion inefficiency* arises when the platform sets the induced

cutoff above the efficient benchmark α in order to reduce the seller's first-period discount. Section 3 isolated the first of these forces in a static benchmark; Section 4 showed how the baseline algorithm can generate the second by setting $\tau^* > \alpha$. This section compares the three algorithm classes along exactly those two margins.

Welfare accounting. Because the market is covered conditional on discovery, and because we hold the on-platform price and fee fixed, welfare differences in our setting come from channel allocation and promotion rather than from changes in total trade or fee setting. Transfers also cancel across the platform and seller, so welfare per interested consumer is expected match value minus expected off-platform channel disutility.[12] Consumer surplus, by contrast, is the utility enjoyed by consumers who buy off-platform: on-platform buyers pay their reservation value and obtain zero surplus. Define the expected channel-disutility cost and the expected consumer surplus per interested consumer under discount Δ as

$$cd(\Delta) \;\equiv\; \int_{\underline{\delta}}^{\Delta} \frac{\delta}{1-\underline{\delta}}\, d\delta \;=\; \frac{\Delta^2 - \underline{\delta}^2}{2(1-\underline{\delta})}, \qquad cs(\Delta) \;\equiv\; \int_{\underline{\delta}}^{\Delta} \frac{\Delta-\delta}{1-\underline{\delta}}\, d\delta \;=\; \frac{(\Delta-\underline{\delta})^2}{2(1-\underline{\delta})}.$$

Note that $cd(\Delta)$ may be negative when some consumers strictly prefer the off-platform channel.

As above, the reserve seller is subject to the same period-2 pricing regime as the seller under the algorithm being evaluated.

Let

$$m(\tau) \;\equiv\; \frac{1-\tau^2}{2} \;+\; \alpha\tau$$

denote the expected type of the seller featured in period 2 given the threshold τ. The seller is featured with probability $1-\tau$ and has expected type $(1+\tau)/2$, while the reserve seller of type α is featured with probability τ. It is immediate that $m(\tau)$ is maximized at $\tau = \alpha$, so that deviations from the efficient cutoff reduce expected period-2 match value by

$$m(\alpha) - m(\tau) \;=\; \frac{(\tau-\alpha)^2}{2}. \tag{9}$$

This will be the formal measure of promotion inefficiency.

For any $(\Delta_1, \Delta_2, \tau)$, total welfare and total consumer surplus are

$$W(\Delta_1, \Delta_2, \tau) \;\equiv\; \tfrac{1}{2}\big(1 - cd(\Delta_1)\big) \;+\; \beta\big(1 - cd(\Delta_2)\big)m(\tau),$$

[12] Allowing for partial market coverage, e.g. as a result of downward-sloping demand, would raise a further welfare effect that our model does not capture, as we discuss in Section 6.5.

$$CS(\Delta_1, \Delta_2, \tau) \;\equiv\; \tfrac{1}{2} cs(\Delta_1) \;+\; \beta cs(\Delta_2) m(\tau).$$

The first-best benchmark sets both discounts to zero (eliminating channel misallocation) and uses the efficient cutoff (eliminating promotion inefficiency):

$$W^{FB} \;\equiv\; W(0, 0, \alpha).$$

We write the welfare associated with the baseline, OPP, and PPC algorithms as

$$W^* \;\equiv\; W(\Delta_1^*, \Delta_2^*, \tau^*), \qquad W^{OPP} \;\equiv\; W(\underline{\delta}, \underline{\delta}, \alpha), \qquad W^{PPC} \;\equiv\; W(\Delta_1^{PPC}, \Delta_2^{PPC}, \tau^{PPC}),$$

and similarly for consumer surplus. Throughout, $m^* \equiv m(\tau^*)$ denotes the expected featured type under the baseline algorithm.

5.1 OPP: promotion inefficiency versus channel foreclosure

The OPP algorithm eliminates promotion inefficiency. By forcing $\Delta_1 = \underline{\delta}$, it disciplines period-1 pricing directly and so removes the platform's incentive to distort the continuation cutoff upward. In the present noiseless setting, this also implies that first-period on-platform sales equal the seller's type exactly, so the platform implements the efficient second-period cutoff $\tau^{OPP} = \alpha$. The cost is that foreclosing the off-platform channel prevents consumers, if there are any, who intrinsically prefer buying directly from doing so. Whether OPP raises welfare therefore depends on whether the promotion gain outweighs the sales-channel misallocation it creates.

Proposition 9. *Consumer surplus is weakly lower under the OPP algorithm than under either the algorithm or the PPC algorithm. Relative to the algorithm, the welfare difference decomposes as*

$$W^{OPP} - W^* \;=\; \underbrace{\tfrac{1}{2} cd(\Delta_1^*) \;+\; \beta m^* cd(\Delta_2^*)}_{\textit{sales-channel effect}} \;+\; \underbrace{\beta \tfrac{(\tau^* - \alpha)^2}{2}}_{\textit{promotion-efficiency gain}} \qquad . \tag{10}$$

The promotion-inefficiency gain is non-negative; the sales-channel effect may have either sign.

Proof. Under OPP, every interested consumer buys on-platform at the reservation value, so $CS^{OPP} = 0$, whereas the baseline and PPC algorithms generate weakly positive consumer surplus. For the welfare decomposition, Proposition 6 gives $\Delta_1^{OPP} = \Delta_2^{OPP} = \underline{\delta}$ and $\tau^{OPP} = \alpha$, and $cd(\underline{\delta}) = 0$. Substituting into $W(\cdot)$ and using $m(\alpha) - m(\tau^*) = (\tau^* - \alpha)^2/2$ from (9) yields (10). □

The two terms in (10) formalize the trade-off. The sales-channel effect is positive when the algorithm sets a sufficiently large positive discount, so that OPP's foreclosure corrects a misallocation toward the off-platform channel; it is negative when consumers intrinsically prefer the direct channel and the baseline benchmark sets a negative discount, so that OPP prevents consumers from buying where they would rather. The promotion-inefficiency gain is the β-weighted cost of any upward cutoff distortion under the algorithm. The OPP–PPC comparison has the same structure: OPP removes promotion inefficiency when PPC does not, but forecloses the direct channel that PPC preserves.

A useful benchmark is the limit $\alpha \to 0$. The first-order condition (8) gives $\tau^* = 0$ and Proposition 4 implies $\Delta_1^* = \Delta_2^*$, so the promotion term vanishes and

$$W^{OPP} - W^*\big|_{\alpha=0} \;=\; \tfrac{1+\beta}{2}\, cd(\Delta_2^*) \;=\; \frac{(1+\beta)(f-\underline{\delta})(f+3\underline{\delta})}{16(1-\underline{\delta})}.$$

Hence OPP raises welfare near $\alpha = 0$ if and only if $f > -3\underline{\delta}$, recovering the static sufficient statistic from Proposition 2. With the dynamic promotion force switched off, the sign of the welfare comparison is pinned down entirely by the sales-channel term.

5.2 PPC: first-best welfare and consumer surplus

The PPC algorithm fixes the sales-channel misallocation whenever it binds in both periods. In this case, on- and off-platform prices coincide and every consumer sorts into her preferred channel. Crucially, binding PPC also eliminates promotion inefficiency, because once the first-period discount is pinned at zero, the platform has no reason to distort the cutoff upward and optimally sets $\tau^{PPC} = \alpha$. Both distortions vanish simultaneously.

Proposition 10. *If the PPC binds in both periods, then*

$$W^{PPC} \;=\; W^{FB}.$$

In particular, $W^{PPC} \geq W^{OPP}$ with strict inequality whenever $\underline{\delta} < 0$, and $W^{PPC} \geq W^*$ with strict inequality whenever the algorithm exhibits either sales-channel misallocation ($\Delta_1^* \neq 0$ or $\Delta_2^* \neq 0$) or promotion inefficiency ($\tau^* \neq \alpha$).

Proof. Lemma 8 gives $\Delta_1^{PPC} = \Delta_2^{PPC} = 0$ and $\tau^{PPC} = \alpha$. Substituting into $W(\cdot)$ and using $cd(0) = -\underline{\delta}^2/[2(1-\underline{\delta})]$, equal prices on both channels deliver the first-best sales-channel outcome and the cutoff $\tau^{PPC} = \alpha$ maximizes $m(\cdot)$. Hence $W^{PPC} = W(0,0,\alpha) = W^{FB}$. The comparisons with OPP and the algorithm follow because (i) OPP has $\Delta_t^{OPP} = \underline{\delta} < 0$ under $\underline{\delta} < 0$, inducing strict channel misallocation, and (ii) any deviation of $(\Delta_1^*, \Delta_2^*, \tau^*)$ from

$(0, 0, \alpha)$ strictly reduces $W(\cdot)$ by strict concavity of $-cd$ around zero and strict concavity of m around α. □

The result reframes what a PPC does. The PPC literature typically evaluates parity clauses as restrictions on static pricing and often highlights an anticompetitive effect. Here, a binding PPC shuts down the platform's dynamic incentive to manipulate the ranking rule: once the period-1 discount is pinned at zero, the cutoff can no longer be used as a pricing instrument, so the platform optimally sets it at the efficient level. A single contractual constraint neutralizes both distortions.

Consumer surplus, however, does not follow welfare. When PPC binds in both periods, it lowers the off-platform discount in the second period (reducing consumer surplus from off-platform buyers), and may lower or raise it in the first period. At the same time, it eliminates any promotion inefficiency, and so, on average, features higher seller types in period 2. These forces may go in opposite directions. To see them cleanly, decompose:

$$CS^{PPC} - CS^{*} = \underbrace{\tfrac{1}{2}\big(cs(0) - cs(\Delta_1^*)\big)}_{\text{(I) period-1 price effect}} + \underbrace{\beta m^*\big(cs(0) - cs(\Delta_2^*)\big)}_{\text{(II) period-2 price effect}} + \underbrace{\beta\,\tfrac{(\tau^*-\alpha)^2}{2} cs(0)}_{\text{(III) promotion-efficiency gain}} . \tag{11}$$

Term (III) is always non-negative and captures the consumer-surplus value of featuring higher seller types. Terms (I) and (II) capture the change in per-consumer surplus from the PPC-induced changes in off-platform discounts. (II) is necessarily negative, the sign of term (I) depends on whether the baseline benchmark has positive or negative discount.

When $\alpha = 0$, Proposition 5 that characterizes the algorithm gives $\tau^* = 0$ and $\Delta_1^* = \Delta_2^*$, so the promotion term (III) vanishes. Because in addition $f > -\underline{\delta}$, then $\Delta_1^* = \Delta_2^* > 0$ and both per-consumer terms are negative (PPC collapses a positive discount to zero, reducing surplus from off-platform buyers). Hence $CS^{PPC} < CS^{*}$. Conversely, when α is large enough that the algorithm has driven the first-period discount below zero, term (I) becomes strictly positive, and combined with the non-negative promotion term this can outweigh the negative period-2 term. A parametric case of this sort is illustrated in Figure 3.

The lesson is that consumer surplus and welfare need not rank algorithms the same way, even when the welfare ranking is sharp. A regulator focused on either measure in isolation may reach different conclusions about which algorithmic regime to favor.

5.3 Illustration

Figure 2 illustrates these welfare comparisons for PPC and OPP under a representative parameterization.

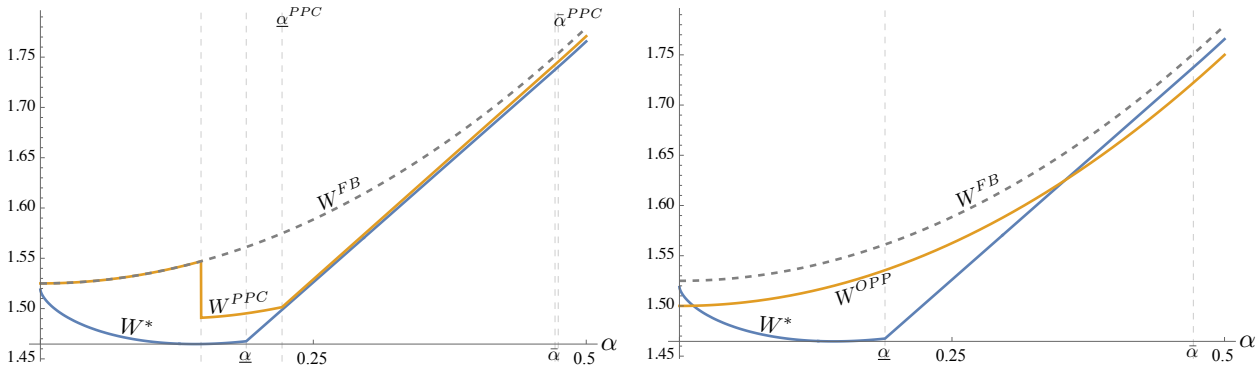

(a) PPC welfare

(b) OPP welfare

Figure 2: Welfare comparisons as functions of the reserve seller's type α. The left panel compares the first-best (dashed gray), algorithm (blue), and PPC algorithm (orange). The right panel compares the first-best (dashed gray), algorithm (blue), and OPP algorithm (orange). Parameters are $\underline{\delta} = -0.2$, $f = 0.4$, and $\beta = 2$.

The left panel of Figure 2 shows that PPC attains first-best welfare for low α — the region where the binding case of Proposition 10 applies — and then tracks the baseline algorithm closely for higher α. The discrete drop in the PPC curve marks the point at which the platform finds it profitable to abandon the binding regime and adopt an intermediate PPC policy in which the PPC binds in period 2 but is slack in period 1, so the platform partially benefits from price discipline without fully eliminating the career-concerns distortion; the details are worked out in Appendix A.4. The right panel shows OPP's ambiguous welfare comparison with the baseline benchmark: OPP lies above the baseline curve for intermediate values of the reserve seller's type α (where the sales-channel cost is small relative to the promotion inefficiency of the algorithm).

Figure 3 illustrates the divergence between welfare and consumer surplus. On the part of the interval where binding PPC attains first-best welfare, consumer surplus under PPC can lie *below* as well as *above* the consumer surplus under the baseline algorithm. In the latter case, PPC raises both welfare and consumer surplus relative to the algorithm.

5.4 Policy implications

The welfare analysis clarifies what is at stake in the regulatory debates surrounding ranking algorithms and price-parity clauses. A natural reading of recent policy discussion is that algorithms should not condition on off-platform prices, and that contractual price restrictions should be viewed with suspicion. Our results suggest a more nuanced picture.

First, restricting a platform from conditioning on off-platform prices does not neutralize the platform's ability to shape cross-channel pricing. A purely sales-based algorithm can

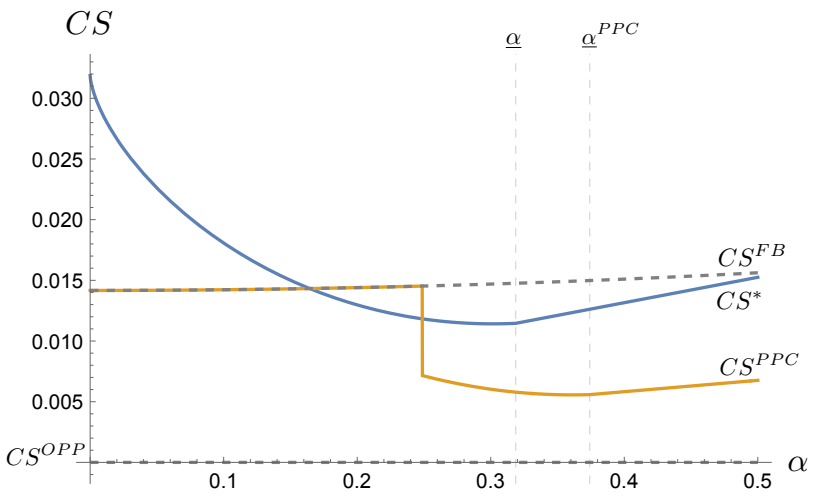

Figure 3: Consumer surplus as a function of the reserve seller's type α for the welfare first-best allocation (dashed gray), algorithm (blue), and PPC algorithm (orange), for $\underline{\delta} = -0.2$, $f = 0.4$, and $\beta = 0.7$.

replicate parity-like incentives by threatening to withhold exposure from sellers whose on-platform sales are low, and in doing so it creates promotion inefficiency that an explicit PPC would avoid when it binds. Second, a PPC embedded in a learning algorithm can deliver first-best welfare by simultaneously eliminating both distortions. Third, even when it does, PPC can reduce consumer surplus relative to the baseline algorithm, so welfare and consumer surplus give different guidance about policy.

Together, these observations caution against evaluating parity restrictions or data-input constraints in isolation. Auditing whether an algorithm uses off-platform information tells a regulator little about whether the algorithm distorts promotion or misallocates sales across channels. Assessing a ranking system's welfare effects requires market-wide context, because off-platform pricing shapes on-platform outcomes even when it is not a direct input to the ranking rule. When the algorithm is also learning from sales, this becomes harder to detect: a rule that steers demand toward the platform may look observationally similar to a rule that merely improves inference about seller quality.

6 Model Variations

In this section we consider a number of model variations. Subsections 6.1, 6.2, 6.3, and 6.4 vary the baseline while preserving the paper's core mechanism. Formal statements and derivations for these are collected in Appendix B. Subsection 6.5 then records the main restrictions we keep for tractability and briefly discusses what would change if they were

relaxed.

6.1 Who knows θ?

The baseline assumes that neither the platform nor the seller observes θ before period 1. Two variants clarify which parts of the analysis rely on that information structure.

The platform observes θ. Suppose the platform learns the realized type before period 1. The role played by the baseline threshold then splits in two. First, there is a *launch cutoff* τ_k: if $\theta < \tau_k$, the platform prefers the period-specific reserve seller in each period. Second, conditional on launch, there is a period-1 *sales target* $T^*(\theta)$, used purely to discipline pricing rather than to learn the seller's type. Appendix B.1 shows that with moderate continuation incentives the platform sets $T^*(\theta) = \theta(1 + \sqrt{\beta})q_2^* < \theta$ and the seller meets that target exactly with an interior discount, whereas with sufficiently strong continuation incentives it sets $T^*(\theta) = \theta$, which induces $\Delta_1 = \underline{\delta}$. In both cases the sales target exceeds the static benchmark θq_2^*.

The comparison with the baseline is

$$\tau_k \;\leq\; \alpha \;\leq\; \tau^*.$$

This inequality is informative. In the baseline, the threshold τ^* must do two jobs at once: it screens on expected continuation value and it provides incentives because period-1 sales both reveal θ and can be manipulated through pricing. That bundling is what drives the baseline upward distortion relative to α. Once θ is observed, these two jobs separate. The platform still uses the continuation rule to discipline period-1 pricing, but it no longer needs to distort the continuation cutoff upward in order to do so. The distortion instead shifts to the launch margin: once the seller's type is observed, featuring the known seller yields higher current fee revenue than the reserve benchmark because the platform can discipline period-1 pricing through the sales target. That extra current payoff makes the platform willing to feature some types below α.

The seller observes θ. Now suppose the seller privately knows her type before period 1. This raises a signaling concern: one might worry that off-platform pricing becomes a way of revealing information. Appendix B.2 studies this case and, for clarity, focuses on the parameter region $\beta \geq 1$, in which any type that can profitably hit the target does so. In that case, types in the interval $(\tau, \tau/q_2^*]$ bunch at the sales target: they choose a discount that meets the threshold exactly. Types above this interval choose the myopic discount

because they would clear the threshold anyway, while types below it also choose the myopic discount because they cannot profitably reach it. The platform again distorts the induced cutoff upward relative to α, now because raising τ extracts rents from the bunching types. So the basic career-concerns force survives even when the seller, rather than the platform, holds the private information.

Neither learning nor private information is essential to the paper's central mechanism. What is essential is that the platform commits to a rule that rewards high on-platform sales and that the seller controls a cross-channel pricing instrument through which those sales can be shifted. Changing the information structure changes the form of the distortion, but not the force that generates it.

6.2 Beyond two periods with known type

The two-period horizon is a useful benchmark, but the pricing-discipline force is not inherently two-period. Appendix B.3 studies an infinite-horizon version in which a seller of known type θ interacts repeatedly with the platform, both sides discount the future at rate $R \in (0, 1)$, and the platform commits to a grim-trigger retention rule: it keeps featuring the seller as long as observed sales imply an on-platform share at least q, and after any shortfall it switches permanently to the reserve seller. In that stationary environment, the highest target the seller is willing to satisfy is

$$q^*(R) \;=\; \min\Big\{1,\, (1+\sqrt{R})q_2^*\Big\},$$

so higher patience strengthens the same force that a larger continuation weight β does in the two-period model.

This extension is straightforward precisely because θ is known. If θ were initially unknown but the signal structure remained as in our baseline model, little would change: because sales are noiseless conditional on pricing, θ would effectively be revealed after one period, so from period 2 onward the problem would collapse to the known-type repeated game. To obtain gradual learning of the sort studied in the classic career-concerns literature, one would need to add an idiosyncratic noise term to sales. The posterior over seller quality would then become a payoff-relevant state, and a non-degenerate steady state would in turn require θ itself to be periodically redrawn or hit by exogenous shocks. In that respect, the richer many-period extension resembles the classic career-concerns problem of an agent facing an evolving reputation state (Holmström, 1999). We therefore view the known-type repeated game as the clean way to see that the paper's mechanism extends beyond two periods, while leaving the noisy unknown-type many-period problem for future work.

6.3 Off-platform fees

Stronger cross-channel competition amplifies the career-concerns force rather than blunting it. Appendix B.4 provides the formal analysis. Suppose the alternative channel itself charges a fee $f_0 < f$, as when a seller uses an enabler platform such as Shopify to facilitate direct sales. The seller's problem is then equivalent to the baseline evaluated at the effective fee

$$\tilde{f} \equiv \frac{f - f_0}{1 - f_0},$$

which is strictly decreasing in f_0. A lower off-platform fee therefore acts like a higher effective platform fee from the seller's perspective: it raises the myopic gain from diverting demand off-platform and strengthens the incentive to manipulate the platform's signal.

The platform's response then follows the same career-concerns logic as in the baseline. Since f is fixed, the platform can only react to a change in f_0 by adjusting the algorithm. Lowering f_0 raises $\tilde{f}$, strengthening the seller's incentive to divert demand off-platform. The platform-side response is therefore exactly the one implied by the baseline model evaluated at $\tilde{f}$: in the interior region the optimal cutoff is unchanged, while through the endogenous boundaries a lower f_0 can induce a more stringent algorithm. In substantive terms, a stronger reserve seller does not necessarily discipline the platform; it can instead strengthen the platform's incentive to use ranking as an indirect pricing instrument.

6.4 Repeat purchases

A natural concern is that consumers who discover a seller on the platform may return directly in later periods. Under a simple formalization, this does not change the baseline characterization; Appendix B.5 provides the formal analysis. Suppose only period-1 consumers can become repeat buyers in period 2, that they can return directly without relying on the platform's period-2 ranking, and that the seller posts common period-2 terms rather than conditioning on the channel through which a consumer first purchased. Repeat purchases then adds an additive constant to both the seller's and the platform's expected continuation values in period 1. As a result, the second-period discount, the first-period discount, and the platform's optimal cutoff are unchanged.

The invariance rests on two assumptions. First, repeat buyers do not need the platform to rediscover the seller in period 2. Second, the seller cannot condition period-2 offers on the buyer's period-1 purchase channel. Relaxing either assumption would add a new dynamic margin: ranking would matter directly for repeat demand in the first case, while period-1 channel choice would become informative about period-2 demand composition in the second.

Those are interesting extensions, but they are orthogonal to the mechanism we isolate here.

6.5 Other variations

Distributional assumptions. We assume uniform distributions for seller type and off-platform disutility to keep the pricing and cutoff formulas transparent and the welfare decompositions in closed form. We view this as a tractability device rather than a substantive restriction. In an earlier version of the paper, we analyzed non-uniform distributions of θ and δ and found the same qualitative insights, albeit without closed-form solutions. Richer distributions therefore change the exact formulas and can shift the quantitative region boundaries, but they do not remove the paper's core force: pricing matters when it can move on-platform sales, and those sales shape future exposure.

Exclusive discovery through the platform. In the baseline, consumers discover the seller only through the platform. If some consumers could find the seller directly, ranking would do less work and the career-concerns force would be attenuated in proportion to the importance of off-platform discovery. The effects would be similar, but weaker, because the platform would control a smaller share of total awareness.

Market coverage conditional on discovery. This is an economically substantive restriction. In the baseline, every consumer who finds the product appealing buys somewhere. That shuts down a deadweight-loss margin from higher prices and lets us isolate two distortions: sales-channel misallocation and promotion inefficiency. Relaxing full-market coverage — for instance, by allowing heterogeneous reservation values on-platform — would add a third welfare margin. OPP and PPC would then affect not only where trade occurs and which sellers receive exposure, but also whether trade occurs at all. The welfare ranking would therefore depend on the elasticity of on-platform demand as well as on the two margins studied in the main model. In a sense, the consumer-surplus results already point in this direction: higher prices can reduce consumer surplus even when the welfare comparison is unchanged or ambiguous in the baseline environment.

Exogenous platform fee f. We treat the on-platform fee f as exogenous here. This is a natural simplification for two reasons. First, platform fees are often fairly uniform across broad ranges of products rather than tailored to the information and demand structure of a particular product and, in practice, have been stable over long periods.[13] Second, the

[13]This assumption is common in the literature including in Zhou and Zou (2023) and Hervas-Drane and Shelegia (2025). A stream of literature has examined endogenous fees including Wang and Wright (2020) ,

covered-market environment of the model is not especially well suited to study fee design itself, because it shuts down the demand-extensive-margin effects that are often central to optimal fee setting.

None of these are essential to the paper's main message. The key ingredients are simply that ranking shapes future exposure and that current cross-channel pricing affects the signal on which ranking is based. When those conditions hold, algorithms exploit career-concerns incentives that create both sales-channel misallocation and promotion inefficiency.

7 Conclusion

We study how ranking design shapes seller pricing when a platform learns from sales that the seller can influence. In our environment, first-period on-platform sales are both the platform's signal about seller appeal and the seller's instrument for affecting future exposure. That link creates a career-concerns force: by restraining off-platform discounts, the seller can shift demand on-platform and improve her chances of being featured later. Anticipating this response, the platform may design its ranking rule not only to allocate exposure, but also to shape cross-channel pricing.

This perspective yields three main conclusions. First, restricting a platform to use only on-platform data does not make the ranking rule neutral with respect to off-platform competition. A purely sales-based algorithm can induce parity-like pricing incentives even without explicitly conditioning on off-platform prices, and it can do so in a way that may distort promotion upward relative to the efficient benchmark. Second, allowing the platform to condition directly on off-platform prices can improve welfare by separating pricing discipline from promotion, although it may do so by foreclosing the direct channel. Third, explicit price-parity restrictions can, in natural parameter regions, outperform the baseline benchmark precisely because they shut down the platform's incentive to use the continuation rule as an indirect pricing instrument.

The welfare analysis highlights two distinct distortions: the sales-channel misallocation and promotion inefficiency. Sales-channel misallocation reflects the interaction of three forces that shape cross-channel pricing: the platform fee pushes the seller to shift sales off-platform, stronger direct-channel demand (captured by lower $\underline{\delta}$) works in the opposite direction by allowing the seller to sustain a smaller off-platform discount or even an off-platform premium, and career concerns further tilt period-1 pricing toward the platform. Promotion inefficiency arises when the platform responds to those pricing incentives by distorting the promotion

Wang and Wright (2025), and Bar-Isaac and Shelegia (2025).

cutoff upward. OPP removes the promotion distortion but may worsen the sales-channel misallocation by forcing all trade on-platform. PPC, when it binds in both periods, eliminates both distortions. At the same time, welfare and consumer surplus need not move together: OPP can lower consumer surplus by forcing consumers onto the platform, while PPC can lower consumer surplus even as it raises welfare, because eliminating the off-platform discount also removes the surplus enjoyed by consumers who would otherwise have bought directly.

Broadly, the paper suggests caution in evaluating ranking systems solely on the variables they use as direct inputs. A regulator who prohibits the use of off-platform prices in ranking may leave untouched — or even strengthen — the incentive to achieve similar outcomes through sales-based learning rules. Conversely, a contractual restriction such as PPC may look harmful when viewed as a static pricing restraint, yet improve outcomes once one accounts for the dynamic interaction between pricing, learning, and future exposure. The economic object that matters is not only what the algorithm observes, but how sellers respond to the incentives it creates.

References

Bar-Isaac, Heski, and Sandro Shelegia. 2025. "Monetizing Steering." *Journal of Political Economy Microeconomics*.

Bar-Isaac, Heski, and Steven Tadelis. 2008. "Seller reputation." *Foundations and Trends® in Microeconomics*, 4(4): 273–351.

Bergemann, Dirk, Alessandro Bonatti, and Nicholas Wu. 2025. "Bidding with Budgets: Data-Driven Bid Algorithms in Digital Advertising." *International Journal of Industrial Organization*, 102: 103172.

Bergemann, Dirk, and Alessandro Bonatti. 2024. "Data, competition, and digital platforms." *American Economic Review*, 114(8): 2553–2595.

Boik, Andre, and Kenneth S. Corts. 2016. "The Effects of Platform Most-Favored-Nation Clauses on Competition and Entry." *The Journal of Law and Economics*, 59(1): 105–134.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello. 2025. "Artificial Intelligence, Algorithmic Recommendations and Competition." *SSRN Electronic Journal*.

Castellini, Jacopo, Amelia Fletcher, Peter L Ormosi, and Rahul Savani. 2023. "Recommender Systems and Competition on Subscription-Based Platforms." *SSRN Electronic Journal*.

Eisfeld, Luise. 2023. "How Do Online Product Rankings Influence Sellers' Pricing Behavior?" Working paper.

Enache, Andreea, and Andrew Rhodes. 2025*a*. "Platform Disintermediation with Repeated Transactions." *Management Science*.

Enache, Andreea, and Andrew Rhodes. 2025*b*. "Price Parity Clauses and Platform Data Acquisition." *Economics Letters*, 256: 112635.

Ennis, Sean, Marc Ivaldi, and Vicente Lagos. 2023. "Price-parity clauses for hotel room booking: Empirical evidence from regulatory change." *The Journal of Law and Economics*, 66(2): 309–331.

Franck, Jens-Uwe, and Martin Peitz. 2024. "The Digital Markets Act and the Whack-a-Mole Challenge." *Common Market Law Review*, 61(2): 299–344.

Hagiu, Andrei, and Julian Wright. 2024. "Marketplace leakage." *Management Science*, 70(3): 1529–1553.

Hervas-Drane, Andres, and Sandro Shelegia. 2025. "Retailer-led marketplaces." *Management Science*, 71(11): 9650–9669.

Holmström, Bengt. 1999. "Managerial incentive problems: A dynamic perspective." *The review of Economic studies*, 66(1): 169–182.

Hunold, Matthias, Reinhold Kesler, and Ulrich Laitenberger. 2020. "Rankings of online travel agents, channel pricing, and consumer protection." *Marketing Science*, 39(1): 92–116.

Hunold, Matthias, Reinhold Kesler, Ulrich Laitenberger, and Frank Schlütter. 2018. "Evaluation of best price clauses in online hotel bookings." *International Journal of Industrial Organization*, 61: 542–571.

Hunold, Matthias, Ulrich Laitenberger, and Guillaume Thébaudin. 2022. "Bye-box: An analysis of non-promotion on the Amazon marketplace." Centre de Recherche en Economie et Droit (CRED) 4.

Jiang, Baojun, Kinshuk Jerath, and Kannan Srinivasan. 2011. "Firm Strategies in the "Mid Tail" of Platform-based Retailing." *Marketing Science*, 30(5): 757–775.

Johnson, Justin P. 2017. "The Agency Model and MFN Clauses." *Review of Economic Studies*, 84(3): 1151–1185.

Johnson, Justin P, Andrew Rhodes, and Matthijs Wildenbeest. 2023. "Platform design when sellers use pricing algorithms." *Econometrica*, 91(5): 1841–1879.

Karle, Heiko, Marcel Preuss, and Markus Reisinger. 2026. "Selling on Recommender Platforms: Demand Boost Versus Customer Migration." CESifo.

Madsen, Erik, and Nikhil Vellodi. 2025. "Insider imitation." *Journal of Political Economy*, 133(2): 652–709.

Mantovani, Andrea, Claudio A Piga, and Carlo Reggiani. 2021. "Online platform price parity clauses: Evidence from the EU Booking. com case." *European Economic Review*, 131: 103625.

Ma, P, Andrea Mantovani, Carlo Reggiani, Annette Broocks, and Néstor Duch-Brown. 2024. "The price effects of prohibiting price parity clauses: evidence from international hotel groups." University of Oxford 1043.

Peitz, Martin. 2022. "The prohibition of price parity clauses and the Digital Markets Act." *Competition Policy International*.

Peitz, Martin, and Anton Sobolev. 2025. "Product recommendations and price parity clauses." University of Bonn and University of Mannheim, Germany.

Scott Morton, Fiona M. 2012. "Contracts that reference rivals." *Antitrust*, 27: 72.

Ursu, Raluca M. 2018. "The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions." *Marketing Science*, 37(4): 530–552.

Wang, Chengsi, and Julian Wright. 2020. "Search Platforms: Showrooming and Price Parity Clauses." *the RAND Journal of Economics*, 51(1): 32–58.

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

Zhou, Bo, and Tianxin Zou. 2023. "Competing for recommendations: The strategic impact of personalized product recommendations in online marketplaces." *Marketing Science*, 42(2): 360–376.

A Proofs and additional results

A.1 Proof of Lemma 1

Let $\phi : [0,1] \to [0,1]$ be any measurable baseline algorithm, and let Δ_1^* be an equilibrium discount under ϕ. Write $q^* \equiv q(\Delta_1^*)$. Since $S_1 = \theta q(\Delta_1)$, the seller's continuation payoff from choosing discount Δ_1 can be written as

$$\beta\pi_2^* \int_0^1 \theta\phi\left(\theta q(\Delta_1)\right) d\theta = \frac{\beta\pi_2^*}{q(\Delta_1)^2} \int_0^{q(\Delta_1)} S\phi(S) dS.$$

Hence the seller's problem depends on ϕ only through the cumulative weighted integral

$$M_\phi(q) \equiv \int_0^q S\phi(S) dS.$$

At the equilibrium discount Δ_1^*, the platform's expected profit is

$$\Pi^{\text{platform}}(\phi) = \frac{f}{2} q^* + f\beta q_2^* \left[\alpha + \frac{M_\phi(q^*)}{(q^*)^2} - \frac{\alpha}{q^*} \int_0^{q^*} \phi(S) dS\right].$$

The term $M_\phi(q^*)/(q^*)^2$ pins down the expected type of the seller when the seller is featured in period 2, while $\frac{1}{q^*}\int_0^{q^*} \phi(S) dS$ is the probability that the seller, rather than the reserve seller, is featured.

Suppose ϕ is not a threshold algorithm on $[0, q^*]$. Then there exist measurable sets $\Sigma_1, \Sigma_2 \subset [0, q^*]$ of equal positive measure m such that every element of Σ_1 is below every element of Σ_2, $\phi(S) > 0$ on Σ_1, and $\phi(S) < 1$ on Σ_2. Define

$$l \equiv \frac{\int_{\Sigma_2} S\, dS}{\int_{\Sigma_1} S\, dS} > 1,$$

and choose $\varepsilon > 0$ small enough that $\phi(S) - \varepsilon l \in [0,1]$ on Σ_1 and $\phi(S) + \varepsilon \in [0,1]$ on Σ_2. Now define a new algorithm $\tilde{\phi}$ by

$$\tilde{\phi}(S) = \begin{cases} \phi(S) - \varepsilon l, & S \in \Sigma_1, \\ \phi(S) + \varepsilon, & S \in \Sigma_2, \\ \phi(S), & \text{otherwise.} \end{cases}$$

This moves featuring probability from lower-sales realizations to higher-sales realizations.

Let

$$D(q) \equiv \int_0^q S\big(\tilde{\phi}(S) - \phi(S)\big)dS.$$

Because all of the mass removed lies below all of the mass added, we have $D(q) \leq 0$ for every $q \leq q^*$. Moreover,

$$D(q^*) = \varepsilon \int_{\Sigma_2} S\, dS - \varepsilon l \int_{\Sigma_1} S\, dS = 0,$$

and since both sets are contained in $[0, q^*]$, it also follows that $D(q) = 0$ for every $q \geq q^*$. Therefore, for every feasible discount Δ_1, the continuation payoff under $\tilde{\phi}$ is weakly below the continuation payoff under ϕ, with equality at $\Delta_1 = \Delta_1^*$. The current-profit term $\frac{1}{2}\pi(\Delta_1)$ is unchanged, so Δ_1^* remains an equilibrium discount under $\tilde{\phi}$.

Evaluating the platform's payoff at this same equilibrium discount, the weighted term $M_\phi(q^*)$ is unchanged by construction, while the overall probability of featuring the seller falls:

$$\int_0^{q^*} \tilde{\phi}(S)dS = \int_0^{q^*} \phi(S)dS - \varepsilon(l-1)m.$$

Hence

$$\Pi^{\text{platform}}(\tilde{\phi}) - \Pi^{\text{platform}}(\phi) = f\beta q_2^* \frac{\alpha}{q^*}\varepsilon(l-1)m \geq 0,$$

with strict inequality whenever $\alpha > 0$. Thus any non-threshold rule can be weakly improved by shifting featuring probability upward in this way.

Applying this exchange argument repeatedly eliminates every pair of positive-measure regions in which lower sales are featured with positive probability while higher sales are left unfeatured with positive probability. Therefore, when $\alpha > 0$, any optimal rule must coincide almost everywhere on the equilibrium support $[0, q^*]$ with a threshold rule: if it differed on a positive-measure subset of that support, the exchange above would raise platform profits strictly.

It remains to consider $\alpha = 0$. In that case the exchange argument still produces a threshold rule with the same equilibrium discount and the same platform payoff. By Proposition 5, the unique optimal threshold at $\alpha = 0$ has induced cutoff $\tau^* = 0$, so the seller's equilibrium discount is $\Delta_1^* = \Delta_2^*$ and hence $q^* = q_2^*$. For any rule ϕ, the platform's continuation payoff at this equilibrium quantity is proportional to $M_\phi(q_2^*)/(q_2^*)^2$, and since $0 \leq \phi \leq 1$ we have

$$M_\phi(q_2^*) = \int_0^{q_2^*} S\phi(S)dS \leq \int_0^{q_2^*} S\, dS = \frac{(q_2^*)^2}{2},$$

with equality if and only if $\phi(S) = 1$ for almost every $S \in [0, q_2^*]$. Therefore any optimal rule that ties the optimal threshold payoff at $\alpha = 0$ must agree with the threshold rule almost

everywhere on the equilibrium support. Any residual non-threshold behavior can survive only off the equilibrium support $S > q^*$ or on measure-zero sets. This proves Lemma 1.

A.2 Implementability of the induced cutoff and omitted cases in Proposition 5

A sales threshold does not necessarily implement every induced cutoff, because the seller may prefer to act myopically and forgo second-period promotion. It is therefore useful to begin with the threshold itself and then define the corresponding induced cutoff.

For a fixed sales threshold T, the seller's objective in (4) is piecewise concave in Δ_1. When $q(\Delta_1) < T$, the continuation term is zero and the seller solves the strictly concave static problem; when $q(\Delta_1) \geq T$, the seller's objective is again strictly concave, with second derivative

$$\frac{1}{2}\pi''(\Delta_1) - \frac{3\beta\pi_2^* T^2}{(1-\underline{\delta})^2 q(\Delta_1)^4} \;=\; -\frac{1}{1-\underline{\delta}} - \frac{3\beta\pi_2^* T^2}{(1-\underline{\delta})^2 q(\Delta_1)^4} < 0.$$

The issue is that the objective changes regime at $q(\Delta_1) = T$: once T exceeds $q(\Delta_1)$, the retention probability drops to zero. Accordingly, the implementability analysis below compares the best promotion payoff with the myopic payoff directly, rather than appealing to global quasi-concavity.

Let $V(T)$ denote the seller's equilibrium payoff under sales threshold T. A higher threshold weakly lowers the continuation probability for every discount choice, so $V(T)$ is weakly decreasing in T. At $T = 0$ the seller is featured for sure in period 2 and therefore chooses the myopic discount Δ_2^*, yielding

$$V(0) = \frac{1}{2}\pi_2^* + \frac{\beta\pi_2^*}{2} \geq \frac{1}{2}\pi_2^*.$$

At $T = 1$, continuation is impossible except on a zero-probability event, so the seller again chooses the myopic discount and obtains

$$V(1) = \frac{1}{2}\pi_2^*.$$

By continuity of the seller's problem, there is therefore a well-defined cutoff threshold at which the seller is just indifferent between pursuing second-period promotion and acting myopically.

Definition A.1. *Let $\bar{T}$ denote the highest sales threshold at which the seller is indifferent between pursuing second-period promotion and acting myopically.*

Thresholds above $\bar{T}$ are not relevant for the platform's problem, because Lemma A.2 below shows that the platform never chooses them.

It is convenient to characterize the corresponding highest implementable cutoff type

$$\bar{\tau} \equiv \tau(\bar{T}).$$

To derive $\bar{\tau}$ explicitly, note that at any equilibrium cutoff τ the seller's profit from choosing discount Δ_1 can be written as $\pi(\Delta_1) = \pi_2^* - \frac{(\Delta_1 - \Delta_2^*)^2}{1-\underline{\delta}}$, since Δ_2^* maximizes π. The seller's total expected payoff when pursuing promotion at cutoff τ is

$$\Pi_{seller} = \frac{1}{2}\pi(\Delta_1^*(\tau)) + \frac{\beta\pi_2^*(1-\tau^2)}{2}.$$

At the highest implementable cutoff, the seller must be indifferent between this payoff and the myopic payoff $\frac{1}{2}\pi_2^*$ (choosing $\Delta_1 = \Delta_2^*$ and forgoing promotion). Value matching gives

$$\frac{(\Delta_1^*(\bar{\tau}) - \Delta_2^*)^2}{1-\underline{\delta}} = \beta\pi_2^*(1-\bar{\tau}^2). \tag{12}$$

Interior case ($\Delta_1^*(\bar{\tau}) > \underline{\delta}$). Substituting the interior solution (5), obtained from the seller's first-order condition (6), into (12), we have

$$(\Delta_1^*(\bar{\tau}) - \Delta_2^*)^2 = \frac{(2-f-\underline{\delta})^2}{16}\left(\sqrt{1+4\beta\bar{\tau}^2} - 1\right)^2$$

and $\pi_2^* = \frac{(2-\underline{\delta}-f)^2}{4(1-\underline{\delta})}$, so (12) simplifies to

$$\left(\sqrt{1+4\beta\bar{\tau}^2} - 1\right)^2 = 4\beta(1-\bar{\tau}^2).$$

Expanding and collecting terms, the economically relevant root of this condition is

$$\bar{\tau} = \sqrt{\frac{3+\sqrt{1+8\beta}}{2\left(1+\sqrt{1+8\beta}\right)}}. \tag{13}$$

This case is valid when $\Delta_1^*(\bar{\tau}) > \underline{\delta}$. From (5), $\Delta_1^*(\bar{\tau}) = \underline{\delta}$ when

$$\beta = \underline{\beta} \equiv \frac{(f-\underline{\delta})(2+f-3\underline{\delta})}{(2-\underline{\delta}-f)^2}.$$

Thus the interior case applies when $\beta < \underline{\beta}$.

Corner case ($\Delta_1^*(\bar{\tau}) = \underline{\delta}$). If $\beta \geq \underline{\beta}$, then $\Delta_1^*(\bar{\tau}) = \underline{\delta}$, so $(\Delta_1^*(\bar{\tau}) - \Delta_2^*)^2 = (\Delta_2^* - \underline{\delta})^2 = \frac{(f-\underline{\delta})^2}{4}$. Substituting into (12) gives

$$\bar{\tau}^2 = 1 - \frac{(f-\underline{\delta})^2}{\beta(2-\underline{\delta}-f)^2} = 1 - \frac{(1-q_2^*)^2}{\beta(q_2^*)^2}.$$

Combining the two cases,

$$\bar{\tau} = \begin{cases} \sqrt{\dfrac{3+\sqrt{1+8\beta}}{2\left(1+\sqrt{1+8\beta}\right)}} & \text{if } \beta < \underline{\beta}, \\ \sqrt{1 - \dfrac{(1-q_2^*)^2}{\beta(q_2^*)^2}} & \text{if } \beta \geq \underline{\beta}. \end{cases} \tag{14}$$

This is the highest induced cutoff that can be implemented in equilibrium: for any higher cutoff, the seller prefers to act myopically and forgo second-period profits.

Recovering the primitive sales threshold from the induced cutoff gives

$$\bar{T} = \begin{cases} \dfrac{q_2^*}{4}\left(3+\sqrt{1+8\beta}\right)\sqrt{\dfrac{3+\sqrt{1+8\beta}}{2\left(1+\sqrt{1+8\beta}\right)}} & \text{if } \beta < \underline{\beta}, \\ \sqrt{1 - \dfrac{(1-q_2^*)^2}{\beta(q_2^*)^2}} & \text{if } \beta \geq \underline{\beta}. \end{cases}$$

The second line follows because in the corner region $\Delta_1^*(\bar{\tau}) = \underline{\delta}$, so $q(\Delta_1^*(\bar{\tau})) = 1$ and the sales threshold coincides with the induced cutoff: $\bar{T} = \bar{\tau}$. In the interior region, substituting the expression for $\bar{\tau}$ above into the corresponding threshold-cutoff mapping yields the first line.

The next lemma justifies the reduced-form use of the induced cutoff τ in the main text, even though the primitive algorithmic object is the sales threshold T.

Lemma A.1. *On the relevant region $\tau \in [0, \bar{\tau}]$, the map from a sales threshold T to the induced equilibrium cutoff $\tau(T)$ is one-to-one.*

Proof. Given the seller's equilibrium response, the sales threshold that implements a cutoff τ is

$$T(\tau) = \tau q(\Delta_1^*(\tau)).$$

If $\Delta_1^*(\tau) = \underline{\delta}$, then $T(\tau) = \tau$, which is clearly strictly increasing.

If the solution is interior, then by Proposition 4,

$$q(\Delta_1^*(\tau)) = \frac{q_2^*}{2}\left(1+\sqrt{1+4\beta\tau^2}\right),$$

and therefore

$$T(\tau) = \frac{q_2^*}{2}\tau\left(1 + \sqrt{1 + 4\beta\tau^2}\right).$$

Differentiating yields

$$T'(\tau) = \frac{q_2^*}{2}\left(1 + \sqrt{1 + 4\beta\tau^2} + \frac{4\beta\tau^2}{\sqrt{1 + 4\beta\tau^2}}\right) > 0.$$

Thus $T(\tau)$ is strictly increasing in the interior region as well. Since the interior and corner expressions coincide at the cutoff where $\Delta_1^*(\tau) = \underline{\delta}$, it follows that $T(\tau)$ is strictly increasing on the entire relevant region. Hence the mapping between T and τ is one-to-one. □

Lemma A.2. *The platform chooses a sales threshold T^* satisfying $T^* \leq \bar{T}$.*

Proof. Suppose instead that the platform chooses a threshold $T^* > \bar{T}$. By definition of $\bar{T}$, the seller then prefers the myopic discount $\Delta_1 = \Delta_2^*$ and forgoes second-period promotion, so it is not featured in period 2. The platform's profits are therefore

$$\frac{1}{2}(1 + 2\alpha\beta)fq_2^*.$$

By setting $T = 0$, the platform instead induces the seller to set the same myopic discount Δ_2^* in period 1 and features the seller for sure in period 2, earning

$$\frac{1}{2}(1 + \beta)fq_2^*,$$

which is strictly higher because $\alpha < 1/2$. Thus the platform strictly prefers $T = 0$, and therefore any optimal threshold satisfies $T^* \leq \bar{T}$. □

By Lemma A.1, this threshold bound translates directly into the cutoff bound $\tau^* \leq \bar{\tau}$ on the implementable region.

The main text characterizes the optimal baseline algorithm under Assumption 1, that is, for $\beta > \underline{\beta}$, when implementability does not bind. It remains to consider the complementary region $\beta \leq \underline{\beta}$. At the knife-edge $\beta = \underline{\beta}$, the implementability boundary coincides with the lower-bound kink, so the main-text characterization extends directly with $\bar{\alpha} = \bar{\tau}$. We therefore focus here on the genuinely omitted case $\beta < \underline{\beta}$, in which implementability binds before the seller reaches the lower bound $\underline{\delta}$.

Proposition A.1. *Suppose $\beta < \underline{\beta}$, and define*

$$\underline{\alpha} \equiv \bar{\tau}\left(1 - \frac{1}{\sqrt{1 + 4\beta\bar{\tau}^2}}\right).$$

Then there are two possible cases for the optimal baseline algorithm:

1. *If* $0 \leq \alpha < \underline{\alpha}$*, then* $\tau^* \in [0, \bar{\tau})$ *is the unique solution to* (8)*. Moreover,* $\tau^* = 0$ *if and only if* $\alpha = 0$*.*

2. *If* $\alpha \geq \underline{\alpha}$*, then* $\tau^* = \bar{\tau}$*.*

Proof. Since $\beta < \underline{\beta}$, implementability binds before the seller reaches the lower bound $\underline{\delta}$, so the entire feasible region $\tau \in [0, \bar{\tau}]$ is interior. Define

$$g(\tau) \equiv \tau \left(1 - \frac{1}{\sqrt{1 + 4\beta\tau^2}}\right).$$

Then

$$\Pi'_{platform}(\tau) = f\beta q_2^* \left(\alpha - g(\tau)\right),$$

and g is strictly increasing because

$$g'(\tau) = 1 - \frac{1}{(1 + 4\beta\tau^2)^{3/2}} > 0$$

for every $\tau > 0$. Thus $\Pi'_{platform}(\tau)$ is strictly decreasing on $[0, \bar{\tau}]$, so there is at most one interior solution. Moreover, when $\alpha > 0$, we have $\Pi'_{platform}(0) = f\beta q_2^* \alpha > 0$, so the optimum is at $\tau = 0$ only if $\alpha = 0$.

Evaluating the derivative at the upper bound gives

$$\Pi'_{platform}(\bar{\tau}) \geq 0 \quad \Longleftrightarrow \quad \alpha \geq g(\bar{\tau}) = \underline{\alpha}.$$

Thus, if $\alpha \geq \underline{\alpha}$, profits are increasing throughout the feasible region and the platform chooses the highest implementable cutoff, $\tau^* = \bar{\tau}$. If instead $0 \leq \alpha < \underline{\alpha}$, then $\Pi'_{platform}(0) \geq 0$ and $\Pi'_{platform}(\bar{\tau}) < 0$, so strict monotonicity implies a unique interior maximizer solving (8). This proves both cases. □

The condition $\beta > \underline{\beta}$ implies

$$\bar{\tau}^2 - \bar{\alpha}^2 = \frac{\beta(q_2^*)^2 - (1 - q_2^*)(2 - q_2^*)}{\beta(q_2^*)^2} > 0,$$

so $\bar{\alpha} < \bar{\tau}$. Moreover, on the region $\beta \geq \underline{\beta}$, the expression for $\bar{\tau}^2$ is increasing in β, so its minimum is attained at $\beta = \underline{\beta}$, where

$$\bar{\tau}^2 = \frac{1}{2 - q_2^*} > \frac{1}{4}$$

because $\underline{\delta} \leq 0$ and $f < 1$. Hence $\bar{\tau} > 1/2$. Combined with the maintained restriction $\alpha < 1/2$, this shows that when $\beta > \underline{\beta}$ the implementability constraint never binds in the platform's optimization problem, yielding the simplified characterization in Proposition 5.

A.3 Proof of Proposition 5

On the interior region where $\Delta_1^*(\tau) > \underline{\delta}$, define

$$g(\tau) \equiv \tau \left(1 - \frac{1}{\sqrt{1 + 4\beta\tau^2}}\right),$$

so that

$$\Pi'_{platform}(\tau) = f\beta q_2^* \left(\alpha - g(\tau)\right).$$

The first-order condition is therefore $\alpha = g(\tau)$, and g is strictly increasing because

$$g'(\tau) = 1 - \frac{1}{(1 + 4\beta\tau^2)^{3/2}} > 0$$

for every $\tau > 0$. Thus $\Pi'_{platform}(\tau)$ is strictly decreasing on the interior region, and there is at most one interior solution. Moreover,

$$\Pi'_{platform}(0) = f\beta q_2^* \alpha,$$

so $\tau^* = 0$ if and only if $\alpha = 0$.

By Proposition 4, the seller first reaches the maximal discount $\underline{\delta}$ at

$$\tau = \bar{\alpha} = \sqrt{\frac{2(1-\underline{\delta})(f-\underline{\delta})}{\beta(2-\underline{\delta}-f)^2}} = \sqrt{\frac{\Delta_2^* - \underline{\delta}}{\beta\pi_2^*}}.$$

By Assumption 1, we have $\bar{\alpha} < \bar{\tau}$; see Appendix A.2. Thus the kink at $\bar{\alpha}$ is reached before the implementability constraint. For every $\tau \geq \bar{\alpha}$ that is relevant to the platform's problem, the seller sets $\Delta_1^*(\tau) = \underline{\delta}$, and therefore

$$\Pi'_{platform}(\tau) = f\beta q_2^*(\alpha - \tau).$$

Appendix A.2 also shows that under Assumption 1 and the maintained restriction $\alpha < 1/2$ we have $\bar{\tau} > 1/2 > \alpha$, so $\tau = \alpha$ is implementable. Hence, if $\alpha \geq \bar{\alpha}$, profits increase up to $\tau = \alpha$ and decrease thereafter, so $\tau^* = \alpha$, proving the third case.

Now suppose $\alpha < \bar{\alpha}$. Since $\Pi'_{platform}(\tau)$ is strictly decreasing on $(0, \bar{\alpha})$, whether the

platform stops at the kink or at an interior solution is determined by the left derivative at $\bar{\alpha}$. Using Proposition 4, this derivative is nonnegative if and only if

$$\alpha \geq \frac{2(f-\underline{\delta})}{2+f-3\underline{\delta}}\bar{\alpha} = \frac{2(1-q_2^*)}{(2-q_2^*)}\bar{\alpha} \equiv \underline{\alpha}.$$

If $\underline{\alpha} \leq \alpha < \bar{\alpha}$, then profits are increasing on $(0, \bar{\alpha})$ and decreasing immediately after the kink because $\alpha - \tau < 0$ for every $\tau \geq \bar{\alpha}$. Hence $\tau^* = \bar{\alpha}$. If instead $0 \leq \alpha < \underline{\alpha}$, then $\Pi'_{platform}(0) \geq 0$ and $\Pi'_{platform}(\bar{\alpha}^-) < 0$, so the strict monotonicity of $\Pi'_{platform}(\tau)$ on $(0, \bar{\alpha})$ implies a unique interior maximizer in $[0, \bar{\alpha})$ solving (8). This proves the first case and completes the proof.

A.4 Further characterization of PPC cases

Suppose $f > -\underline{\delta}$, so the PPC binds in period 2 and therefore $\Delta_2^{PPC} = 0$. If the period-1 PPC constraint is slack, the seller's problem is the same as in Proposition 4, except that continuation payoffs are evaluated at $\Delta_2^{PPC} = 0$ rather than at Δ_2^*. Let $q(0)$ and $\pi(0)$ denote the corresponding second-period PPC demand and seller profit objects. Define

$$\bar{\beta}^{PPC}(\tau) \equiv \frac{\Delta_2^* - \underline{\delta}}{\pi(0)\tau^2}.$$

Then if $\beta \geq \bar{\beta}^{PPC}(\tau)$ the seller chooses $\Delta_1^{PPC}(\tau) = \underline{\delta}$, while for $\beta < \bar{\beta}^{PPC}(\tau)$ the first-period discount is interior and given by

$$\Delta_1^{PPC}(\tau) = \Delta_2^* - \frac{1}{4}(2-f-\underline{\delta})\left(\sqrt{1+4\frac{\pi(0)}{\pi_2^*}\beta\tau^2} - 1\right). \tag{15}$$

This mirrors the interior seller solution (5), with π_2^* replaced by $\pi(0)$. For a given τ, it is strictly above $\Delta_1^*(\tau)$ because $\pi(0) < \pi_2^*$ whenever $f > -\underline{\delta}$.

On the interior region, the platform's first-order condition is therefore the PPC analogue of (8):

$$\alpha = \tau\left(1 - \frac{\frac{q_2^*}{q(0)}\frac{\pi(0)}{\pi_2^*}}{\sqrt{1+4\frac{\pi(0)}{\pi_2^*}\beta\tau^2}}\right). \tag{16}$$

Accordingly, the PPC problem has an implementability constraint exactly as in the baseline benchmark. Define

$$\bar{\alpha}^{PPC} \equiv \sqrt{\frac{\Delta_2^* - \underline{\delta}}{\beta\pi(0)}}, \qquad \underline{\alpha}^{PPC} \equiv \frac{3f-\underline{\delta}}{2+f-3\underline{\delta}}\bar{\alpha}^{PPC}, \qquad \underline{\beta}^{PPC} \equiv \frac{(1-q(0))(2-q(0))}{q(0)^2}.$$

where $\underline{\alpha}^{PPC}$ is obtained from the left derivative of the PPC profit function at the kink $\tau = \bar{\alpha}^{PPC}$ using (16). When $\beta > \underline{\beta}^{PPC}$, the slack-period-1 PPC problem has the same three-region structure as Proposition 5: for $0 \leq \alpha < \underline{\alpha}^{PPC}$ the optimal cutoff is interior; for $\underline{\alpha}^{PPC} \leq \alpha < \bar{\alpha}^{PPC}$, the platform chooses $\tau^{PPC} = \bar{\alpha}^{PPC}$ and induces $\Delta_1^{PPC} = \underline{\delta}$; and for $\alpha \geq \bar{\alpha}^{PPC}$, it chooses the efficient cutoff $\tau^{PPC} = \alpha$ with $\Delta_1^{PPC} = \underline{\delta}$.

To determine when the binding candidate $(\tau, \Delta_1) = (\alpha, 0)$ can be optimal, define $\tilde{\alpha}^{PPC}$ by setting (15) equal to zero at $\tau = \alpha$:

$$\tilde{\alpha}^{PPC} \equiv \sqrt{\frac{\Delta_2^*}{\beta\pi(0)(1-\underline{\delta})}}.$$

Since $\Delta_1^{PPC}(\tau)$ is decreasing in τ and $0 > \underline{\delta}$, the cutoff at which the interior solution reaches 0 must lie strictly below the cutoff at which it reaches $\underline{\delta}$, so

$$\tilde{\alpha}^{PPC} < \bar{\alpha}^{PPC}.$$

Because the binding regime requires $\Delta_1^{PPC}(\alpha) \geq 0$, it can only be feasible for $\alpha \leq \tilde{\alpha}^{PPC}$. Since $\tilde{\alpha}^{PPC} < \bar{\alpha}^{PPC}$, the platform must switch away from binding weakly before the efficient-cutoff slack regime $\tau^{PPC} = \alpha$, $\Delta_1^{PPC} = \underline{\delta}$ becomes relevant. Therefore, when $\beta > \underline{\beta}^{PPC}$, the only remaining ambiguity is whether the profit comparison between the binding candidate $(\tau, \Delta_1) = (\alpha, 0)$ and the optimal slack-period-1 PPC solution changes sign before or after $\underline{\alpha}^{PPC}$. If it changes sign for $\alpha < \underline{\alpha}^{PPC}$, the switch is from binding to the interior slack regime. If instead it changes sign for $\underline{\alpha}^{PPC} \leq \alpha \leq \tilde{\alpha}^{PPC}$, the switch is from binding directly to the corner regime $\tau^{PPC} = \bar{\alpha}^{PPC}$ with $\Delta_1^{PPC} = \underline{\delta}$.

Finally, when $\alpha = 0$, no slack-period-1 PPC solution can satisfy the interior first-order condition: the PPC analogue of (8) has no positive solution at $\alpha = 0$, so the best slack candidate is the boundary cutoff $\tau = \tilde{\alpha}^{PPC}$. At that boundary, $\Delta_1^{PPC} = 0$, so the slack candidate gives the same first-period on-platform share as the binding candidate $(\tau, \Delta_1) = (0, 0)$, but strictly lower second-period profits because it features the seller on a smaller set of types. Hence the binding candidate is optimal at $\alpha = 0$. By continuity, if $f > -\underline{\delta}$ the PPC binds in both periods for α sufficiently close to 0.

B Model Variations: Details

This appendix collects the variations discussed in Section 6. The goal is twofold. First, the model variations verify that the paper's core mechanism does not rely on the baseline information structure. Second, they isolate which assumptions are doing work in the comparative statics and welfare analysis.

B.1 The platform observes θ

Suppose $\theta \sim U[0,1]$ as in the baseline, but the platform observes the realized draw before period 1. It is useful to separate the launch decision from pricing discipline. Conditional on launching the seller in period 1, the platform can require a period-1 on-platform share $x \in [q_2^*, 1]$ by committing to feature the seller in period 2 if and only if

$$S_1 \;\geq\; \theta x.$$

Because θ is known, this target has no learning role.

Lemma B.1. *Fix a realized type θ and suppose the platform has launched the seller in period 1. Given a target on-platform share $x \in [q_2^*, 1]$, the seller meets the target by choosing $q(\Delta_1) = x$ if and only if*

$$\beta \;\geq\; 1 - \frac{\pi(1 - x(1-\underline{\delta}))}{\pi_2^*}.$$

Otherwise, the seller chooses the myopic discount $\Delta_1 = \Delta_2^*$.

Proof. If the seller meets the target, it chooses the least distortion consistent with doing so, namely the unique Δ_1 satisfying $q(\Delta_1) = x$, or equivalently $\Delta_1 = 1 - x(1-\underline{\delta})$. Since θ multiplies both current and continuation profits, the seller prefers meeting the target to choosing the myopic benchmark if and only if

$$\theta\pi(\Delta_1) + \beta\theta\pi_2^* \;\geq\; \theta\pi_2^*,$$

which simplifies to the stated condition. □

Define

$$\beta_k \;\equiv\; 1 - \frac{1-f}{\pi_2^*} \;=\; \left(\frac{1-q_2^*}{q_2^*}\right)^2 \;<\; 1.$$

If $\beta \geq \beta_k$, the platform can induce full foreclosure of the off-platform channel in period 1 by setting $x = 1$. If $\beta < \beta_k$, the seller's incentive constraint binds at an interior target share.

Proposition B.1. *Suppose the platform observes θ before period 1.*

1. *If $\beta < \beta_k$, then conditional on launch the platform chooses*

$$x^* \;=\; (1+\sqrt{\beta})q_2^*,$$

so the period-1 sales target is

$$T^*(\theta) \;=\; \theta(1+\sqrt{\beta})q_2^*,$$

and the seller sets

$$\Delta_1^* \;=\; \Delta_2^* - \sqrt{\beta}(1-\Delta_2^*).$$

2. *If $\beta \geq \beta_k$, then conditional on launch the platform chooses $x^* = 1$, so*

$$T^*(\theta) \;=\; \theta,$$

and the seller sets $\Delta_1^* = \underline{\delta}$.

3. *The platform launches the seller in period 1 if and only if $\theta \geq \tau_k$, where*

$$\tau_k \;\equiv\; \begin{cases} \alpha\dfrac{1+\beta}{1+\sqrt{\beta}+\beta}, & \text{if } \beta < \beta_k, \\ \alpha q_2^*\dfrac{1+\beta}{1+\beta q_2^*}, & \text{if } \beta \geq \beta_k. \end{cases}$$

Proof. Consider first the case $\beta < \beta_k$. The platform wants the highest target share that the seller is willing to meet. By Lemma B.1, the optimal target therefore satisfies the seller's incentive constraint at equality:

$$\pi(\Delta_1^*) = (1-\beta)\pi_2^*.$$

Using

$$\pi(\Delta) = \pi_2^* - \frac{(\Delta - \Delta_2^*)^2}{1-\underline{\delta}},$$

the root below Δ_2^* is

$$\Delta_1^* = \Delta_2^* - \sqrt{\beta}(1-\Delta_2^*).$$

Substituting into $q(\Delta_1)$ gives

$$x^* = q(\Delta_1^*) = (1+\sqrt{\beta})q_2^*,$$

and therefore $T^*(\theta) = \theta(1+\sqrt{\beta})q_2^*$.

Conditional on launch, platform profit is

$$f\theta x^* + \beta f\theta q_2^*,$$

whereas using the period-specific reserve seller in each period yields

$$(1+\beta) f\alpha q_2^*.$$

Comparing these expressions implies launch if and only if

$$\theta \;\geq\; \alpha \frac{(1+\beta)q_2^*}{x^* + \beta q_2^*} \;=\; \alpha \frac{1+\beta}{1+\sqrt{\beta}+\beta}.$$

Now consider $\beta \geq \beta_k$. In this region the seller is willing to accept the strongest possible target, namely $x^* = 1$, which implies $T^*(\theta) = \theta$ and $\Delta_1^* = \underline{\delta}$. Conditional on launch, platform profit is then

$$f\theta + \beta f\theta q_2^*,$$

so launch is optimal if and only if

$$\theta \;\geq\; \alpha q_2^* \frac{1+\beta}{1+\beta q_2^*}.$$

This proves the stated expressions for τ_k. □

Corollary B.1. *Relative to the baseline learning environment,*

$$\tau_k \;\leq\; \alpha \;\leq\; \tau^*.$$

Moreover, conditional on launch, the period-1 off-platform discount when the platform observes θ is weakly lower than in the baseline learning environment.

Proof. The first inequality follows directly from the formulas for τ_k: in both regions the coefficient multiplying α is at most one. The second inequality is the baseline result from Proposition 5.

For the discount comparison, in the known-type case the seller either chooses $\underline{\delta}$ or the solution to

$$\pi(\Delta_1) = (1-\beta)\pi_2^*.$$

In the baseline learning case, the continuation value is diluted because the same threshold must simultaneously screen and provide incentives. At induced cutoff τ, the seller's equilib-

rium discount must weakly dominate the myopic benchmark, so

$$\pi(\Delta_1) \geq \left(1 - \beta(1-\tau^2)\right) \pi_2^*.$$

and since $\beta(1-\tau^2) \leq \beta$, the right-hand side is weakly larger than in the known-type case. Because $\pi(\cdot)$ is strictly increasing on $[\underline{\delta}, \Delta_2^*]$, this implies that the known-type seller's optimal discount is weakly lower. □

Remark B.1 (Observing θ weakly expands the feasible set)**.** *In this variant, observing θ weakly expands the platform's feasible set: after learning θ, the platform can always choose not to condition on it and thereby replicate any algorithm available under ignorance. What changes is not the value of information, but the form of the distortion: the baseline upward distortion of the pooled continuation cutoff disappears, while the launch cutoff falls below α.*

Implication for OPP and PPC. The ranking across algorithm classes is qualitatively unchanged. OPP remains the strongest instrument because it can directly impose $\Delta_1 \leq \underline{\delta}$ conditional on launch. PPC remains weaker but can still pin down period-1 pricing and thereby eliminate the platform's incentive to distort the continuation rule.

B.2 Only the seller knows θ

Now suppose $\theta \sim U[0,1]$ and is observed by the seller but not by the platform before period 1. Because θ enters profits multiplicatively, no pricing deviation separates types: the seller's ranking over feasible discounts is independent of type except through the feasibility of hitting the platform's target. Passive off-equilibrium beliefs — the platform does not update its belief about θ in response to unexpected pricing — are therefore the natural refinement here, and we adopt them throughout this subsection.

Lemma B.2. *Suppose the platform commits to a threshold rule T and maintains passive off-equilibrium beliefs. Then there is an equilibrium with no additional signaling through off-platform pricing, in which the seller sets $\Delta_2 = \Delta_2^*$ in period 2, and conditional on being featured in period 1 a seller of type θ behaves as follows:*

1. *If $T \leq \theta q_2^*$, the myopic discount already secures period-2 featuring, so $\Delta_1 = \Delta_2^*$.*
2. *If $T > \theta$, no feasible discount can secure period-2 featuring, so again $\Delta_1 = \Delta_2^*$.*
3. *If $\theta q_2^* < T \leq \theta$, the seller meets the threshold by choosing*

$$\Delta_1 = 1 - \frac{T(1-\underline{\delta})}{\theta}$$

if and only if

$$\beta \;\geq\; 1 - \frac{\pi\left(1 - \frac{T(1-\delta)}{\theta}\right)}{\pi_2^*};$$

otherwise it sets $\Delta_1 = \Delta_2^*$.

Proof. Because θ multiplies both current and continuation profits, it affects the seller's choice only through whether the target is feasible. The three cases therefore follow exactly as in the known-type seller problem. □

To simplify the characterization, suppose $\beta \geq 1$, so that any type for which meeting the threshold is feasible prefers to do so. Then the primitive threshold and the induced cutoff coincide, and every $\tau \in [0,1]$ is implementable by setting $T = \tau$.

Conditional on being featured in period 1, types $\theta \in (\tau, \tau/q_2^*]$ bunch at the sales target τ, while types outside this interval set the myopic discount. Platform profits can therefore be written as

$$\Pi^{\text{platform}} \;=\; f\left[q_2^*\frac{\tau^2}{2} + \tau\left(\min\left\{\frac{\tau}{q_2^*}, 1\right\} - \tau\right) + q_2^*\frac{1 - \min\{\tau/q_2^*, 1\}^2}{2} + \beta q_2^*\left(\alpha\tau + \frac{(1-\tau^2)}{2}\right)\right].$$

Assuming an interior cutoff with $\tau^* < q_2^*$, the first-order condition is

$$\beta q_2^*(\alpha - \tau) + \left(\frac{\tau}{q_2^*} - \tau\right) - \tau(1 - q_2^*) = 0, \tag{17}$$

or equivalently

$$\beta q_2^*(\alpha - \tau) + \frac{(1-q_2^*)^2}{q_2^*}\tau = 0,$$

giving

$$\tau^I \;=\; \frac{\beta}{\beta - \beta_k}\alpha.$$

Because $\beta > \beta_k$ under Assumption 1, this interior candidate is strictly above α. Thus the platform again distorts the induced cutoff upward relative to the efficient benchmark. The mechanism differs from the baseline in form but not in substance: here the distortion extracts rents from the bunching types rather than operating through the seller's marginal pricing choice.

OPP and PPC. The comparison across algorithms mirrors the baseline. OPP yields $\Delta_1^{OPP} = \Delta_2^{OPP} = \underline{\delta}$ and $\tau^{OPP} = \alpha$. PPC has no effect if $\Delta_2^* \leq 0$; if instead $\Delta_2^* > 0$, then $\Delta_2^{PPC} = 0$, and in period 1 PPC either leaves the seller at parity or weakly reduces the discount for all types relative to the baseline benchmark.

B.3 Beyond two periods with known type

To see that the paper's mechanism is not specific to a two-period horizon, suppose now that the seller's type θ is common knowledge and that a unit mass of consumers arrives in each period $t = 1, 2, \ldots$. The seller and the platform discount future payoffs at rate $R \in (0, 1)$. The reserve seller remains short-lived, so if the platform turns away from the seller it uses the period-specific reserve seller in every future period.

Consider stationary grim-trigger algorithms of the following form. If the seller is launched in period 1, then after any history in which it was featured and delivered on-platform share at least $q \in [q_2^*, 1]$, it continues to be featured next period. After any shortfall, the platform switches permanently to the reserve seller. Because θ is known, requiring on-platform share at least q is equivalent to requiring period sales at least θq.

Let

$$\Delta(q) \;\equiv\; 1 - q(1 - \underline{\delta})$$

denote the least distorted off-platform discount that delivers on-platform share q.

Lemma B.3. *Fix a target $q \in [q_2^*, 1]$. Under the grim-trigger algorithm, the seller is willing to meet the target in every period if and only if*

$$\frac{\pi(\Delta(q))}{1 - R} \;\geq\; \pi_2^*.$$

Equivalently, the target is feasible if and only if

$$q \;\leq\; \bar{q}(R) \;\equiv\; \min\left\{1,\, (1 + \sqrt{R})q_2^*\right\}.$$

Proof. If the seller meets the target, it chooses the least distorted discount consistent with doing so, namely $\Delta(q)$. Its value from complying in every period is therefore

$$\theta\pi(\Delta(q)) + \sum_{t=1}^{\infty} R^t \theta\pi(\Delta(q)) \;=\; \frac{\theta\pi(\Delta(q))}{1 - R}.$$

If instead the seller deviates in the current period, it chooses the myopic discount Δ_2^*, earns $\theta\pi_2^*$ today, and is then replaced forever by the reserve seller. Hence compliance is optimal if and only if

$$\frac{\theta\pi(\Delta(q))}{1 - R} \;\geq\; \theta\pi_2^*,$$

which is the stated incentive constraint.

Using

$$\pi(\Delta) \;=\; \pi_2^* - \frac{(\Delta - \Delta_2^*)^2}{1-\underline{\delta}},$$

the constraint becomes

$$(\Delta(q) - \Delta_2^*)^2 \;\le\; R(1-\underline{\delta})\pi_2^* \;=\; R(1-\Delta_2^*)^2.$$

Since $\Delta(q) \le \Delta_2^*$ for $q \ge q_2^*$, this is equivalent to

$$\Delta(q) \;\ge\; \Delta_2^* - \sqrt{R}(1-\Delta_2^*).$$

Substituting $\Delta(q) = 1 - q(1-\underline{\delta})$ and rearranging gives

$$q \;\le\; (1+\sqrt{R})q_2^*.$$

Imposing $q \le 1$ yields the stated upper bound $\bar{q}(R)$. □

Proposition B.2. *In the repeated known-type extension, conditional on launching the seller the platform chooses*

$$q^*(R) \;=\; \bar{q}(R) \;=\; \min\Big\{1,\, (1+\sqrt{R})q_2^*\Big\}.$$

The platform launches the seller in period 1 if and only if

$$\theta \;\ge\; \tau_R \;\equiv\; \frac{\alpha q_2^*}{q^*(R)}.$$

Proof. Conditional on featuring the seller, the platform's per-period revenue is $f\theta q$, which is strictly increasing in q. By Lemma B.3, the largest feasible target is therefore optimal, which proves the expression for $q^*(R)$.

If the seller is launched and complies forever, the platform obtains discounted profit

$$\frac{f\theta q^*(R)}{1-R}.$$

If instead the platform uses the reserve seller forever, it obtains

$$\frac{f\alpha q_2^*}{1-R}.$$

Comparing these values implies launch if and only if

$$\theta q^*(R) \;\ge\; \alpha q_2^*,$$

which is equivalent to the stated threshold τ_R. □

Remark B.2 (Relation to the two-period continuation weight)**.** *The seller's incentive constraint can be written as*

$$\pi(\Delta(q)) + \frac{R}{1-R}\pi(\Delta(q)) \;\geq\; \pi_2^*.$$

In the repeated-game incentive constraint, the continuation payoff enters with the multiplier

$$\beta_\infty(R) \;\equiv\; \frac{R}{1-R}.$$

The closed-form target $q^*(R) = \min\{1, (1+\sqrt{R})q_2^*\}$ *is the stationary analogue of the known-type two-period target derived in Appendix B.1.*

B.4 Off-platform fees

Suppose the off-platform channel itself charges a fee $f_0 < f$, for instance because it is operated by an enabler platform such as Shopify rather than by the seller directly. Then seller profit per interested consumer becomes

$$q(\Delta)(1-f) + (1-q(\Delta))(1-\Delta)(1-f_0).$$

Factoring out $(1-f_0)$ gives

$$(1-f_0)\left[q(\Delta)\left(1-\frac{f-f_0}{1-f_0}\right) + (1-q(\Delta))(1-\Delta)\right].$$

Hence the seller's problem is equivalent to the baseline problem with effective fee

$$\tilde{f} \;\equiv\; \frac{f-f_0}{1-f_0}.$$

Proposition B.3. *With off-platform fee* $f_0 < f$*, the seller's pricing problem is identical to the baseline model evaluated at the effective fee* $\tilde{f} = (f-f_0)/(1-f_0)$*. In particular, the comparative statics for seller behavior with respect to* f *carry over after replacing* f *by* $\tilde{f}$*.*

Proof. The factor $(1-f_0)$ is multiplicative and therefore does not affect the seller's choice of Δ. The remaining expression is exactly the baseline seller objective with fee $\tilde{f}$. □

Implication for the algorithm comparison. Because Proposition B.3 is a relabeling, the welfare ranking across the baseline, OPP, and PPC algorithms carries over unchanged, with $\tilde{f}$ in place of f. Lowering f_0 raises $\tilde{f}$, pushes the static benchmark discount Δ_2^* upward,

and therefore enlarges the region in which PPC binds and delivers the first-best welfare result of Proposition 10. Thus the welfare case for PPC becomes stronger, not weaker, as competition from alternative channels intensifies.

B.5 Repeat purchases

Suppose only period-1 consumers can become repeat buyers in period 2, and that any consumer who purchases in period 1 can return directly to the seller in period 2 without relying on the platform's ranking. Assume also that repeat buyers retain both whether they find the product appealing and their channel preference, and that the seller posts common period-2 terms rather than conditioning on period-1 purchase histories.

For a given induced cutoff τ, seller profits become

$$\Pi_{\text{seller}}(\Delta_1, \Delta_2; \tau) = \frac{1}{2}\pi(\Delta_1) + \beta\pi(\Delta_2)\frac{1-\tau^2}{2} + \frac{1}{2}\pi(\Delta_2).$$

Relative to the baseline, the only new term is the final one, reflecting repeat demand from the period-1 cohort. Since this term is a positive multiple of $\pi(\Delta_2)$, it does not change the seller's optimal period-2 discount. Platform profits become

$$\Pi^{\text{platform}}(\tau) = f\left[\frac{1}{2}q(\Delta_1^*(\tau)) + \beta\left(\alpha\tau + \frac{1-\tau^2}{2}\right)q_2^* + \frac{1}{2}q_2^*\right],$$

and the new term is additive and independent of τ, so the optimal cutoff is unchanged.

Proposition B.4. *Under the repeat-purchase assumptions above, repeat demand adds an additive constant to continuation values but does not affect the seller's optimal period-2 discount, the seller's optimal period-1 discount, or the platform's optimal cutoff. The baseline characterization of Δ_2^*, $\Delta_1^*(\tau)$, and τ^* continues to apply.*

Proof. The second-period seller problem maximizes a positive multiple of $\pi(\Delta_2)$ and is therefore unchanged. Once Δ_2^* is unchanged, the repeat-demand term in platform profits is additive and does not affect the choice of τ. Since the platform problem and the continuation value are otherwise identical to the baseline, the seller's first-period problem is unchanged as well. □

Remark B.3. *The invariance in Proposition B.4 depends on two substantive assumptions:*

1. *Repeat buyers can return directly to the seller in period 2 without relying on the platform's period-2 ranking.*

2. *Sellers cannot condition period-2 offers on the period-1 channel through which a consumer purchased.*

Relaxing (i) would make period-2 ranking relevant for repeat demand, reintroducing the career-concerns force on a new margin. Relaxing (ii) would make period-1 channel choice informative about consumers' preference for direct purchase, which the seller could exploit through type-dependent period-2 pricing. In either case, period-1 pricing would re-enter the platform's dynamic problem through a channel that the baseline rules out.