Ranking algorithms, learning, and pricing*

Heski Bar-Isaac†and Sandro Shelegia‡

September 10, 2025

Abstract

We study ranking algorithms that determine product exposure when sellers can set different on- and off-platform prices, raising the prospect of misallocation across sales channels. Initial uncertainty about product desirability is resolved through sales experience, leading to sellers having incentives akin to career concerns. Ranking design that relies on historical sales can replicate contractual price parity clauses (PPCs) but, with uncertainty, neglecting off-platform prices in ranking decisions can lead to distorted promotion of sellers who aim to boost on-platform sales and exposure. Explicitly conditioning on off-platform prices (including through PPCs) can raise both sales-channel allocation efficiency and promotion efficiency.

1 Introduction

Ranking algorithms—which determine the order in which consumers see sellers, or whether they see them at all—play a critical role in e-commerce. Because visibility is a necessary condition for making sales and earning profits, sellers devote considerable effort to understanding and influencing these algorithms.[1] Platforms, in turn, strategically design these algorithms, and their objectives and consequences have attracted growing scrutiny from regulators and policymakers. The European Union's Digital Markets Act and Digital Services Act—supported by the European Centre for Algorithmic Transparency—have established regulatory frameworks for the oversight and auditing

*Thanks to Audiences at MACCI, U Toronto, ESSET 2025, Mad-Bar 2025, IO Winter Symposium - Bergamo 2024, Pricing, Algorithms and Competition Conference, 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 (rant 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]Evidence of seller interest in ranking optimization is abundant. For example, 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 the search rankings. As one illustration, Webfx notes 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!" https://www.webfx.com/blog/marketing/what-is-amazon-seo/ accessed March 5. 2025.

of algorithms used by digital platforms and gatekeepers to ensure "fairness and contestability," and there has been broad interest in auditing algorithms (CMA, 2022).[2]

A particular concern involves the use of prices on other sales channels to influence within-platform rankings. This issue has received heightened attention amid broader regulatory interest in contractual price restrictions such as price parity clauses that govern prices at other sales venues.[3] In the academic literature, such practices have been examined in the context of online travel agents (Hunold et al. 2018) and the Amazon marketplace (Hunold, Laitenberger and Thébaudin 2022) and have led some to question whether provisions banning price parity clauses can be circumvented through the design of ranking algorithms (Peitz 2022 and Franck and Peitz 2024).

While such practices are clearly of interest to regulators, the literature on ranking algorithms has typically focused on within-platform competition. Instead, we present a model where consumers' interest in a seller is uncertain and learned through experience. In that environment, we contrast ranking algorithms that rely on off-platform prices to those that do not.

An initial observation is that off-platform prices can affect on-platform sales—even when the algorithm does not directly condition on them. In the extreme case where off-platform prices are prohibitively high, the seller effectively disables that channel. All sales must take place on the platform; it is clear that, as a result, on-platform sales may be higher than when off-platform prices are lower.[4] Thus, algorithms that do not condition on off-platform prices directly may nonetheless respond to them indirectly via their influence on platform sales. In turn, this affects a seller's incentives. By raising her off-platform price today, a seller can increase the likelihood of being promoted in the future. This mechanism is reminiscent of the classic career concerns model (Holmström (1999)), with the off-platform price playing the role of hidden effort—one that is ultimately futile in equilibrium, since the platform fully anticipates it. However, we argue that in a natural context, banning algorithms from using off-platform prices is not neutral. It can actually lead to worse outcomes.

Platforms take time to learn the quality of products and decide whether to promote them through the ranking algorithm based on historical sales. The force described in the paragraph above suggests that a platform may design its algorithm—mapping first-period sales to second-

[2]The Digital Markets Act is the common designation for "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 and amending Directives (EU) 2019/1937 and (EU) 2020/1828"; the Digital Services Act is "Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and amending Directive 2000/31/EC (Digital Services Act)."

[3]A notable example is Booking.com's use of a "most favored nation" clause, which was deemed illegal by the German Federal Court of Justicein May 2021. The clause had attracted regulatory scrutiny as early as 2015. see Mantovani, Piga and Reggiani (2017) and Hunold et al. (2018) for analysis and further discussionof this case.

[4]This is not lost on the search engine optimization (SEO) industry and sellers. https://www.bigcommerce.com/blog/amazon-seo-strategy/#amazon-optimizations-that-get-a-320-increase-in-sales-in-less-than-10-minutes (accessed on March 5, 2025) lists as the first optimization strategy "1 Sales are King ... Your first option for generating sales is by driving both internal and external traffic to your Amazon listing". And https://www.bigcommerce.com/articles/omnichannel-retail/how-to-sell-on-amazon/#3-steps-to-selling-successfully-on-amazon(accessed on March 5, 2025) includes a section titled "Point all product links to Amazon" noting that "You may ask why you should send customers to Amazon if you already have a proprietary website. The reason? 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."

period exposure—to encourage higher off-platform prices and greater on-platform sales. To do so most profitably, however, can involve "punishing" products that do not generate enough sales. As a result, the platform may not promote a product that should be promoted by the ranking algorithm (a force we term "promotion inefficiency"). This would appear to be a practice that is "unfair" to sellers, contravening the EU's Digital Markets Act, and can reduce consumer surplus and overall welfare.

We demonstrate how different constraints on algorithms—such as banning the use of off-platform prices, or using them only in a way that ensures off-platform are not lower than on-platform (that is a price parity clause)—affect both promotion efficiency and (through differential impact on prices) the efficient allocation of transactions across the different available sales channels. The extent of "sales channel efficiency" naturally depends on the distribution of consumer preferences. Our central analysis simplifies the pricing problem by supposing that any interested consumer has the same unit demand for an on-platform purchase. As a result, if consumers derive more utility from on-platform purchases, then welfare is higher if the off-platform is foreclosed and only promotion inefficiency can arise. In these circumstances, algorithms that rely directly on off-platform prices can maximize welfare. More generally, there can be a trade-off between sales channel efficiency and promotion efficiency. Sales channel efficiency is always maximized when there are identical prices on and off the platform, providing intuition for our result that an algorithm that imposes a price parity clause can raise both types of efficiency.

Thus, our results highlight that prohibiting platforms from using information on off-platform prices can reduce welfare and suggest when such concerns might be most pronounced. Further, auditing algorithms to ensure promotion efficiency and sales channel efficiency may require off-platform information, suggesting that simply auditing algorithms without wider market context may be of limited value.

Our paper makes three primary contributions to the literature. First, we provide a dynamic model in which a platform commits to a ranking algorithm that influences seller behavior across channels, even without explicitly conditioning on off-platform prices. This contrasts with existing work that either assumes one-shot mechanisms or abstracts away from seller-type uncertainty. Second, we show how ranking algorithms serve as commitment devices that shift seller incentives in the presence of asymmetric information, drawing a novel link between platform design and the career concerns framework. Third, we analyze the welfare and policy implications of algorithmic design. In particular, we demonstrate conditions under which a platform-optimal algorithm leads to inefficient exclusion or effectively implements a most-favored-nation (MFN) clause—raising potential concerns for regulators but in a more distortive way than an explicit contractual clause and highlighting the possibilities of promotion and sales-channel inefficiencies. In this way, our results contribute to ongoing debates about algorithmic transparency and fairness, particularly in light of recent regulatory initiatives such as the EU Digital Markets Act.

Roadmap After reviewing some related literature in the next section, we introduce a model featuring a monopoly platform that is the only channel through which consumers can discover a product. However, consumers may purchase the product either on-platform or through an alternative sales channel. The seller sets prices for both channels. An algorithm, which may or may not be allowed to use off-platform prices, determines whether to feature the product, and in a second period, historical sales on the platform can also be used to inform the algorithm.

Our analysis builds intuition by beginning with the case in which the seller's type is known. By committing not to feature the seller in the second period unless certain sales targets are met, the platform can indirectly influence off-platform pricing in the first period. If the second period is sufficiently important, the platform can fully shift first-period demand to itself by inducing the seller to raise the off-platform price to a level that, in effect, forecloses alternative sales channels. Next, in Section 5, we turn to our main analysis: the case of an unknown seller type, and analyze the seller's behavior before tracing through the implications for the platform's profit-maximizing design of its algorithm. This allows us to contrast algorithms that do or do not use off-platform prices and conduct comparative statics and welfare exercises. Specifically, in addition to algorithms restricted from using off-platform information, we consider both fully unrestricted algorithms that can use off-platform prices in an unrestricted way ("OPP algorithms") and algorithms that can only use off-platform prices to impose price parity clauses ("PPC algorithms").

The baseline model aims to highlight forces as clearly as possible, and as a result, is naturally a stylized one. We address a number of natural considerations, in Section 7, including alternative fee structures, sellers who know their types, fees in the off-platform channel (that can be understood as addressing the intensity of competition from alternative platforms), repeat purchases, and allowing for downward-sloping demand. Finally, we conclude.

1.1 Literature Review

Our work contrasts with Castellini et al. (2023) and Calvano et al. (Calvano et al. (2025)), who also examine the role of platform ranking algorithms. Their focus is on how such algorithms affect competition among products within a platform, as does Johnson, Rhodes and Wildenbeest (2023) who focuses on how algorithms promote competition within the platform. Bar-Isaac and Shelegia (2024) contrasts algorithm-based versus auction-based product ranking in terms of their impact on consumers. Instead, we study how algorithmic design shapes competition between the platform and off-platform sales channels.

There are a few recent papers that consider this interaction (albeit in a static setting in contrast to our setup where historical sales play a key role). Hagiu and Wright (2024) suppose that the off-platform channel is always less efficient and, while less focused on the ranking algorithm, consider platform strategies to avoid consumers making off-platform purchases. Karle, Preuss and Reisinger (2025) and Peitz and Sobolev (2025) also consider this interaction but highlight different trade-offs and effects. Karle, Preuss and Reisinger (2025) consider a seller with existing consumers, who may migrate to the platform, and focus on commission rates. Peitz and Sobolev (2025) consider an environment where prohibiting PPCs is always welfare-decreasing in a model where a platform has information about consumers and can condition the recommendation strategy based on whether consumers are picky and how well-matched they are for the product. Bergemann and Bonatti (2024) and Bergemann, Bonatti and Wu (2025) consider the interaction between on-platform and off-platform prices as shaped by platform mechanisms. These papers analyze profit-maximizing designs in one-shot settings that allow for price discrimination and ad-based allocations but abstract from the dynamic incentive effects created by sales history. Their focus is on multi-product sellers and the role of consumer data in efficient matching, rather than on commitment and the contrast between algorithms that rely on off-platform prices and those that rely on historical sales.

A small literature examines platforms interacting with sellers of unknown type. Jiang, Jerath and Srinivasan (2011), and Madsen and Vellodi (2025) study environments in which the platform may become a competitor, leading sellers to obscure their high types to deter platform entry. We study the opposite incentive: sellers seek to convey high type to secure continued platform exposure. In this respect, our model is closely related to the career concerns literature, in which early actions influence later outcomes by affecting beliefs about agent quality.[5]

Finally, our paper is related to another literature that studies the use and impact of explicit most-favored-nation (MFN) and price-parity clauses (PPC), which have been a focus of regulatory scrutiny in online markets. Notable papers in this literature include Scott Morton (2012), Boik and Corts (2016), Johnson (2017), Hunold et al. (2018), and Calvano and Polo (2020). Our contribution is to show how and when algorithmic design can replicate the *incentive effects* of MFNs—altering seller behavior across channels—without any explicit contractual restriction. This suggests that algorithmic design may warrant similar regulatory attention, especially under emerging digital platform laws such as the EU's Digital Markets Act.

2 Model

We study a setting with a monopoly platform that gives a seller access to consumers in exchange for an ad valorem fee $f \in [0, 1)$. The seller may also operate a direct, off-platform channel, accessible to consumers who become aware of the product through the platform.

A mass of consumers arrives in each of two periods. In period 1, the mass is normalized to 1; in period 2, the mass is $\beta \geq 0$.[6] The platform decides in each period whether or not to feature the seller. If featured, consumers learn of the seller and can observe both the on-platform price p_t, and the off-platform price $\tilde{p}_t$ for the product. If not featured, the seller makes no sales in that period through either channel.

Consumers are characterized by two types of heterogeneity:

- **Taste for the product**: Upon seeing the seller, a consumer finds the product appealing with probability $\theta \in [0, 1]$, where $\theta \sim G(\cdot)$ with finite density $g(\cdot)$ and mean μ. If the product is appealing, the consumer's willingness to pay is 1; otherwise, it is 0. Thus, θ can be interpreted

[5] See the seminal workHolmström 1999, or Bar-Isaac and Tadelis (2008)for a survey of some of the literature.

[6] This can also be understood to account for any time discounting. Broadly β can be interpreted as sales occurring after learning concludes and may reflect that learning occurs among a fraction of consumers before potential wider deployment.

as the fraction of consumers who derive value from the product. All consumers can obtain zero utility by abstaining.

- **Channel preference**: Conditional on being aware of the seller and liking the product, a consumer may purchase either on or off the platform. Purchasing off the platform incurs a disutility $\delta \sim H(\cdot)$ distributed over the interval $[\underline{\delta}, 1]$ where $\underline{\delta} \leq 0$.[7] We assume H is log-concave and that δ is independent of θ. Notably, we allow for $\underline{\delta} < 0$ and that the disutility δ may be negative. That is, we allow some consumers to prefer off-platform channel (for example, a traveler may prefer to book directly with a hotel, or a consumer may anticipate better service by purchasing directly with the manufacturer).

We suppose that the seller incurs no costs and assume that $f < 1$, so that profitable trade is possible on the platform.[8] The platform incurs an opportunity cost $A \geq 0$ when featuring the seller, reflecting the revenue (before applying the fee f) that it could have earned by allocating the same exposure to an alternative seller or use.[9] This outside option allows for the possibility of promotion inefficiency, and therefore plays a key role in welfare analysis, as we discuss in Section 6.

Let $r(\theta) \equiv \theta^2 g(\theta)$. This is the integrand in the raw (un-centered) second moment of distribution G and will be useful in characterizing the informativeness of sales realizations. We assume that $r(\theta)$ is strictly quasi-concave on $[0, 1]$ and increasing on $[0, \mu]$. This condition is satisfied by symmetric single-peaked distributions and many common cases such as the Beta family of distributions.[10]

2.1 Timing and Information

The game unfolds over two periods, $t = 1, 2$.

Before any pricing or consumption decisions, the platform sets an algorithm, as described in Section 2.2. The seller observes the algorithm, and then in each period, the following events occur:

1. **Pricing**: The seller sets the on-platform price p_t and the off-platform price $\tilde{p}_t$.
2. **Seller featured**: As per the algorithm, the seller is either featured or not.
3. **Consumer decisions** (if seller is featured): Consumers observe p_t and $\tilde{p}_t$, decide whether to buy, and choose a channel.
4. **Sales realization**: The seller and platform observe platform sales s_t. Off-platform sales are not observed by the platform.

[7] Capping disutilities by 1 is without much loss because 1 is the highest willingness to pay of which disutility is subtracted.

[8] Allowing for a positive marginal cost introduces additional notation and messier expressions but does not affect any of our qualitative results.

[9] That is, the platform earns Af from the outside option.

[10] The derivative of $x^2 g(x)$ is $x^2 g'(x) + 2xg(x)$, which for a symmetric single-peaked distribution is positive for $x < \mu = 1/2$.

In the incomplete information version of the model, the platform does not know the seller's type θ, but uses observed first-period sales s_1 together with observed or expected prices to update its belief. In the complete information benchmark, θ is common knowledge.

Neither the seller nor the platform discounts the future. That is, both value profits across periods equally.[11]

2.2 Different Forms of Algorithms

A ranking algorithm is a rule that maps available information into a ranking. In our model, the ranking takes a particularly simple form: a binary decision about whether a seller is featured or not. We consider three classes of algorithms, distinguished by the type of information on which they rely:

- A **(restricted) algorithm** uses only information observable on the platform. In the initial period, this corresponds to the platform price p_1. In the second period, this corresponds to current and historical prices on the platform, p_2 and p_1, as well as the period 1 on-platform sales, s_1.
- An **OPP algorithm** (off-platform price algorithm) conditions on both on- and off-platform prices. In period 1, it may use $\tilde{p}_1$ in addition to p_1; in period 2, it can condition on all current and past prices, $p_1, p_2, \tilde{p}_1, \tilde{p}_2$ and historical on-platform sales, s_1.[12]
- A **PPC algorithm** (price parity clause algorithm) is similar to a restricted algorithm but imposes the following constraint: any seller whose off-platform price is lower than their on-platform price is excluded from being featured. Otherwise, it relies on no further off-platform information.

The first two classes of algorithms are natural benchmarks. In practice, a platform may not have access to all off-platform prices; for example, a potential guest can obtain an idiosyncratic quote by calling a hotel directly—an action that may be costly for a platform to replicate, especially in settings such as hotels or airlines where the number of prices is very large. However, regulators and commentators have expressed concerns over the use of off-platform information broadly. In particular, price parity clauses have been a focal point of concern and are often more feasible for platforms to monitor and enforce than fully general OPP algorithms. These considerations motivate our interest in PPC algorithms.

The analysis proceeds as follows. We first characterize the seller's second-period pricing decision. We then examine first-period pricing under the assumption that the seller's type is known. Finally, we consider the case in which the seller's type is unknown. We begin by analyzing the profit-maximizing restricted algorithm as our baseline case, before turning to optimal OPP and PPC algorithms.

[11] Coefficient β incorporates any discounting that may take place.

[12] As we will soon emerge, access to off-platform sales, in addition to prices, would not improve platform outcomes.

3 Period 2 equilibrium

We solve the game by backward induction.

We begin by noting that, in many cases, in any period, it is optimal for the platform to induce the seller to set the platform price at $p_t^* = 1$, the highest price that interested consumers (who have identical unit demand for a good on the platform) are willing to pay. This is intuitive: the platform seeks to maximize the on-platform revenue, on which it earns a proportionate fee.

We continue our analysis under the assumption that this is the case and show that this holds when $H(\cdot)$ is the uniform distribution in Appendix A.1.[13]

It is convenient to define $\Delta_t \equiv p_t - \tilde{p}_t = 1 - \tilde{p}_t$. This is the off-platform discount. Note that we do permit the discount to be negative, in which case the off-platform price is higher than the platform price. Indeed, this possibility is important to our analysis.

A consumer who finds the product appealing (which occurs with probability θ) and has disutility δ from buying off-platform chooses to buy on the platform if:

$$p_2 \leq \tilde{p}_2 + \delta \Leftrightarrow \Delta_2 \leq \delta.$$

Thus, the fraction $1 - H(\Delta_2)$ of consumers who find the product appealing buys on the platform; the remaining fraction $H(\Delta_2)$ buys off-platform.

The seller's profit per consumer for whom the product has appeal will be denoted by

$$\pi(\Delta) \equiv (1 - H(\Delta))(1 - f) + H(\Delta)(1 - \Delta).$$

We can rewrite this as

$$\pi(\Delta) = 1 - f + H(\Delta)(f - \Delta). \tag{1}$$

This expression can be interpreted as follows. The seller earns $1 - f$ on every consumer who buys, and then obtains a further $f - \Delta_2$ on the fraction $H(\Delta_2)$ of consumers that buys off-platform. In particular, if the seller gives no discount $\Delta_2 = 0$, she stands to gain f per consumer she attracts to the off-platform channel, and if she sets the 'full' discount, $\Delta_2 = f$, she makes no extra money on it.

It is convenient to assume that $\lim_{x \to \underline{\delta}} H(x)/h(x) = 0$, which holds for the uniform distribution (and all others with $h(\underline{\delta}) > 0$) but also for the triangular and beta distributions.[14] Given that H is strictly log-concave, we know that $\pi'(\Delta) = h(\Delta)\,(f - \Delta) - H(\Delta)$ changes sign at most once. Recall that H has support $[\underline{\delta}, 1]$, and $\pi'(\Delta)$ is non-negative at $\Delta = \underline{\delta}$ (since $f - \underline{\delta} \geq 0$) and is strictly

[13] The force that we must consider is that a lower on-platform price may induce a smaller off-platform discount and, therefore, higher on-platform sales. This force is also relevant for downward-sloping which we consider in Section 7.5.

[14] This assumption ensures that $\Delta_2^* > \underline{\delta}$ while in its absence $\Delta_2^* = \underline{\delta}$ may hold.

negative at $\Delta = 1$ (since $f \leq 1$). Thus, the unique maximizer Δ_2^* solves

$$\Delta = f - \frac{H(\Delta)}{h(\Delta)}. \tag{2}$$

We summarize this discussion as follows:

Proposition 1. *In period 2, seller sets* $p_2^* = 1$ *and* Δ_2^* *as the unique solution to (2).*

Proof. This follows from the discussion above. □

We can see that $\Delta_2^* < f$, which is intuitive: an off-platform discount can only be profitable if it is lower than the fee that the seller pays to the platform. However, the off-platform "discount" may actually be negative (so that there is a premium associated with buying directly from the seller). For example, if $f = 0$ and $\underline{\delta} < 0$, then $\Delta_2^* < 0$: the seller sets a higher off-platform price to extract surplus from consumers who prefer to buy off-platform. It will be convenient to use $\pi_2^* \equiv \pi(\Delta_2^*)$ for period two profits per consumer.

For welfare considerations, it is useful to find the circumstances where the discount is zero. Given that the market is covered, a zero price differential between channels ensures that consumers buy from whichever sales channel generates greater utility and, so, maximizes welfare. This requires that $f = H(0)/h(0)$. For this condition to hold, it must be that $\underline{\delta} < 0$ if $f > 0$ or $\underline{\delta} = 0$ if $f = 0$. Notably, if some consumers prefer buying off the platform (that is, $\underline{\delta} < 0$), then a strictly positive platform fee is welfare-maximizing. In this case, starting from $f = 0$, an infinitesimal fee increase leads to higher platform profits, welfare, and consumer surplus. Intuitively, as the fee increases, the seller becomes less interested in selling through the platform and reduces the off-platform price accordingly. This brings the off-platform discount closer to the welfare-maximizing one since, in the absence of platform fees, the off-platform price is higher than the on-platform one as the seller uses its market power over consumers.

3.1 OPP Algorithm

We now discuss the possibility that the platform's algorithm is also based on the off-platform price $\tilde{p}_t$. Since the platform has commitment power, it can deny the seller access to consumers based on this price. In particular, the platform would choose to feature the seller only if $\Delta_2 < \underline{\delta}$ so that no consumer buys off-platform and all sales take place on the platform, where it earns fees. That is, the platform, in effect, would require exclusivity. The seller would have to comply, or else it would not be featured on the platform, no consumers would be aware of the product, and, so, the seller would make no sales on- or off-platform. Thus, it is clear that observing realized sales without any uncertainty about the seller's type, while equivalent to observing the seller's off-platform historical price, is nevertheless fundamentally different from observing the off-platform price directly because the latter allows the algorithm to condition on the price at the time of determining whether or not to feature the product. Instead, a restricted algorithm must wait for sales to be realized (i.e. until period 2 in our model) to react to off-platform prices.

Proposition 2. *Under the optimal OPP algorithm, all consumers buy from the platform at the price $p_2 = 1$. Consumer surplus is lower in this case than under the (restricted) algorithm that does not condition on the off-platform price. Welfare may be higher or lower.*

Proof. The first statement follows directly from the discussion in the text. Clearly, with all purchases on the platform and at a price of $p_2 = 1$, no consumer surplus is generated at all. This is not the case for the restricted algorithm, where some consumers choose to purchase off-platform. Since they vary in δ but purchase at a common off-platform price, there must be strictly positive consumer surplus. With regard to welfare, note that given price discounts off-platform, some consumers with $\delta > 0$ purchase under the restricted algorithm, reducing welfare. Instead, if $\underline{\delta} < 0$ then some consumers who strictly prefer to purchase off-platform may do so under the restricted algorithm but not if the platform can condition on the off-platform price. Either force may dominate, leading to an ambiguous welfare result. □

3.2 PPC Algorithm

A PPC algorithm does not have the full flexibility of an OPP algorithm to affect off-platform prices, but can ensure that the off-platform price $\tilde{p}_2$ is no lower than the on-platform price $p_2 = 1$. Consequently, if under the restricted algorithm, a seller would set $\tilde{p}_2 > 1$, the PPC would have no impact at all; otherwise, that is, if the PPC was binding, then the seller would set $\tilde{p}_2 = p_2$.[15] The following result arises immediately.

Proposition 3. *When the platform uses a PPC algorithm, then the off-platform discount is determined by 2 when its solution, Δ_2^*, is negative; and, otherwise, the off-platform price and on-platform price are identical: $p_2 = \tilde{p}_2 = 1$. Consumer surplus under the PPC algorithm is lower than for the restricted algorithm, but higher than the OPP algorithm; welfare is higher than under both alternatives.*

Proof. Total welfare is weakly lower under the PPC algorithm than under the restricted algorithm since either consumers face the same discount in both cases or a positive discount under the restricted algorithm and no discount under the PPC. Consumers are worse off—they face weakly higher off-platform prices. The same force that can lead to the OPP platform raising welfare relative to the restricted algorithm—namely, ensuring that consumers who do not suffer a positive disutility from buying off-platform nevertheless do so because of a lower off-platform price—also applies for the PPC algorithm, but whereas the OPP algorithm precludes all off-platform purchases, the PPC algorithm can allow at least some of those consumers with $\delta < 0$ to buy off-platform leading to higher welfare and greater consumer surplus than the OPP algorithm. □

Intuitively, a welfare distortion that arises in our setting is that consumers may end up buying through one sales channel, even though the alternative channel brings them greater utility. This

[15] Note that since the PPC is an explicit contractual clause, even if breach is only discovered later, a fine on discovering breach ex-post can be sufficient to ensure that the seller is compliant throughout. This is what we assume throughout our discussion on PPCs.

allocative distortion arises when the prices at these different sales channels differ, even though there are no differences in costs. As Proposition 3 highlights, either the PPC has no impact on prices (if it does not bind) or it brings the off-platform price in line with the on-platform price, and, thereby, eliminates such sales channel misallocation and raises welfare.

4 Period 1 equilibrium for known seller-type

We now analyze period 1 pricing under the assumption that the seller's type θ is known. We will assume that $\theta \geq A/(1-H(\Delta_2^*))$ to ensure that the seller is more profitable than the outside option, because we wish to focus on the platform's extractive ability vis à vis the seller.[16] In this case, the platform observes on-platform sales given by:

$$s_1 = \theta(1 - H(\Delta_1)),$$

and the platform may commit to feature the seller in period 2 if these sales exceed a threshold T.[17]

The seller cannot meet the threshold if $T > \theta$ and thus would set the first-period discount equal to the myopic optimum Δ_2^*. Instead, if $T \leq \theta(1 - H(\Delta_2^*))$, then setting the myopically optimal discount ensures being featured in period 2. This is what the seller would do; that is, the seller would optimally choose its first-period off-platform discount to be Δ_2^*.

In other cases, to meet the threshold T, the seller θ must set an off-platform discount $\Delta_1 < \Delta_2^*$ such that $T = \theta(1 - H(\Delta_1))$, or equivalently

$$\Delta_1 = H^{-1}(1 - T/\theta).$$

The seller will optimally do so if

$$\theta \left(\pi(H^{-1}(1 - T/\theta)) + \beta\pi_2^*\right) \geq \theta\pi_2^*.$$

The left-hand side of the inequality corresponds to the profit from setting Δ_1 and, so, being featured in the second period (and earning per consumer profits of π_2^* in that period). Instead, the right-hand side involves setting the discount at the optimal myopic level and foregoing the opportunity to be featured and to make sales in the second period.

Note that the seller's type drops out of this equation since it is a common multiplicative factor in both cases. Simplifying the inequality, we obtain the following condition.

$$\pi(H^{-1}(1 - T/\theta)) \geq (1 - \beta)\pi_2^*. \tag{3}$$

[16]In the absence of this assumption, the seller may well not be featured. We will make an analogous assumption for the case of the unkown seller-type.

[17]A threshold rule is equivalent to simply choosing the off-platform price in this setting. Trivially, it is optimal. As we show later, a threshold rule is optimal when the seller's type is unknown. We, therefore, maintain a threshold rule in this section for ease of comparison with the analysis in Section 5.

This inequality trivially holds for $\beta \geq 1$. For $\beta < 1$, whether or not the inequality is satisfied depends on other parameters, as summarized in the following result, which is immediate following the discussion above.

Lemma 1. *A seller of a known type θ when faced with a sales threshold T sets the first period discount $\Delta_1 = H^{-1}(1 - T/\theta)$ if $\frac{T}{\theta} \in (1 - H(\Delta_2^*), 1]$ and $\beta < 1 - \frac{\pi(H^{-1}(1-T/\theta))}{\pi_2^*}$, otherwise $\Delta_1 = \Delta_2^*$.*

The seller meets the threshold either because meeting it entails no further discount beyond that which it would otherwise choose, i.e. $T \leq \theta(1 - H(\Delta_2^*))$, or when it is feasible to meet the threshold ($\theta \geq T$) and the second period is important enough (that is β is high enough) to warrant doing so.

4.1 Optimal algorithm for known seller type

The platform prefers that the seller set $\Delta_1 = \underline{\delta}$ and divert all sales to the platform. In this case, the seller would earn $1 - f$ per consumer. To characterize when this is incentive compatible, define the threshold:

$$\underline{\beta} \equiv 1 - \frac{1-f}{\pi_2^*} < 1.$$

If $\beta \geq \underline{\beta}$, the seller would prefer to set $\Delta_1 = \underline{\delta}$ and forego (optimal) first-period off-platform sales for the opportunity of on-platform second-period sales (and the opportunity to earn π_2^* in the second period) rather than forego second-period sales and earn the maximal possible, π_2^*, in the first period . The platform can effectively induce the seller to do this by choosing a sales threshold rule with the threshold set at $T = \theta$, which the seller can attain by foregoing off-platform sales by setting $\Delta_1 = \underline{\delta}$. If, instead, $\beta < \underline{\beta}$, then (3) does not hold at $\Delta_1 = \underline{\delta}$, and the platform cannot force this level of discount. In this case, it implements a higher discount $\Delta_1^* \in (\underline{\delta}, \Delta_2^*]$ through a corresponding sales threshold $T^* = \theta(1 - H(\Delta_1^*))$ that ensures that the seller is indifferent between choosing this discount and foregoing second-period sales, since this is the lowest discount that the platform can force. Thus, the optimal discount (from the platform's perspective) in this case is implicitly defined by the following equation.

$$\pi(\Delta_1^*) = (1 - \beta)\pi_2^*. \tag{4}$$

The following proposition summarizes this discussion.

Proposition 4. *With known seller type θ:*

- *If $\beta < \underline{\beta}$, the platform commits to feature the seller in period 2 if $s_1 \geq T^* = \theta(1 - H(\Delta_1^*))$, where Δ_1^* solves (4) and Δ_2^*, as above is the solution to (2).*
- *If $\beta \geq \underline{\beta}$, the platform commits to feature the seller if $s_1 = \theta$, and the seller sets $\Delta_1^* = \underline{\delta}$.*

Proof. Immediate following the discussion above. □

In the high-β case, the platform captures all demand through a sales-based threshold rule—even without observing the off-platform price. Although the platform does not observe the seller's off-platform price, it can use a sales threshold to reward sellers who drive demand towards the platform. Anticipating this, the seller raises the off-platform price, thereby steering consumers toward the platform. The algorithm effectively forecloses the off-platform channel in the first period without explicit pricing constraints. When β is low, the platform cannot induce such a high off-platform price (or equivalently low price discount) in the first period. Instead, it allows the seller to retain some off-platform demand—resulting in lower on-platform sales—in order to ensure continued participation in the second period.

Platform profits from a seller of known type θ are:

$$\Pi = f\theta \left[(1 - H(\Delta_1^*)) + \beta(1 - H(\Delta_2^*))\right].$$

First-best welfare requires that there is no sales-channel misallocation and, so, $\Delta_1^* = \Delta_2^* = 0$. As previously established, $\Delta_2^* = 0$ when $f = H(0)/h(0)$, which in turn requires $\underline{\delta} < 0$ if $f > 0$. In this case, to have $\Delta_1^* = 0$ as well, Equation (3) implies $\beta = 1$. Generically, therefore, such an outcome does not arise, and so there is a distortion from the socially optimal outcome.

4.2 OPP algorithm

Suppose that the platform can condition its algorithm on the seller's off-platform price. In this case, just as for the second period, it can then require $\Delta_1 < \underline{\delta}$ as a prerequisite for exposure, effectively foreclosing the off-platform channel. However, this conditioning is inconsequential for the first period when $\beta \geq \bar{\beta}$, as the same outcome arises through a commitment to a sales threshold rule.

4.3 PPC algorithm

When the platform can impose a PPC, this will affect the second-period profit that a seller can anticipate if featured. With a PPC, the seller anticipates setting $\Delta_2^{PPC} = max\{0, \Delta_2^*\}$ and earning $\pi(max\{0, \Delta_2^*\})$. Analogously to the analysis above, define

$$\underline{\beta}_{PPC} \equiv 1 - \frac{1 - f}{\pi(max\{0, \Delta_2^*\})} \leq \underline{\beta}.$$

Then clearly, if $\beta > \underline{\beta}$, the platform can foreclose the off-platform in the first period through any of the restricted algorithm, OPP algorithm, or PPC algorithm. However, there are circumstances (specifically, where $\underline{\beta}_{PPC} > \beta > \underline{\beta}$) in which a PPC algorithm forecloses the off-platform channel in the first period, but the restricted algorithm does not; and there are circumstances in which the OPP algorithm forecloses the off-platform channel, but the PPC algorithm does not. Note that when both the PPC and the restricted algorithm allow off-platform prices that do not foreclose sales, the prices resulting from these algorithms will be at different levels.[18]

5 Period 1 equilibrium for unknown seller-type

Now we turn to the central case for our analysis, where the platform is uncertain about the seller's type. In this section, we argue that an algorithm based on a sales threshold is optimal, derive the seller's optimal pricing given a threshold algorithm, and then characterize the optimal algorithm. We suppose that it is profitable for the platform to feature the seller in the first period; else, the algorithmic design is irrelevant. The following assumption is sufficient to ensure that this is the case.[19]

Assumption 1. *The platform's alternative option satisfies* $A < \mu(1 - H(\Delta_2^*))$

It will be useful to define $\alpha \equiv A/(1 - H(\Delta_2^*))$ as the type of the seller that would give the platform the same revenue as the outside option in the second period, thus the assumption above is equivalent to $\alpha < \mu$.

With the seller's type unknown, the platform observes $s_1 = \theta(1 - H(\Delta_1))$ and, depending on the level of sales, can choose to feature the seller in the second period with some probability. As we now show, under some (sufficient) conditions on g it is optimal to use a threshold strategy so that the platform features the seller in period 2 if and only if her period 1 sales s_1 exceed some T.

Lemma 2. *A threshold algorithm is optimal when the density g has decreasing elasticity.*

Proof. We prove the result by contradiction in Appendix A.2 .

We can think of the algorithm as assigning a probability $\phi(s)$ of being featured in the second period associated with any sales level s. Our approach is to suppose that an optimal but non-threshold rule ϕ exists. The idea of the proof is to show that this entails two sets of sales levels Σ_1 (low s) and Σ_1 (high s) such that types with $s \in \Sigma_1$ are featured with positive probability while those with $s \in \Sigma_2$ are not always featured. Shift a small probability mass ε from Σ_1 to Σ_2, scaling it by a factor $l > 1$ so the seller's expected continuation probability—and hence her chosen discount—remains constant. Because g has decreasing elasticity, the ratio of weighted type densities attached to Σ_2 versus Σ_1 increases after the shift, which leaves the seller's payoff unchanged but strictly raises the platform's expected profit, yielding a contradiction. □

Many common distributions, including the uniform distribution, satisfy the decreasing-elasticity property. We proceed under the assumption that the distribution function induces a threshold algorithm, with threshold, T. In this case, if the seller sets the discount Δ_1, she will be featured in

[18] In particular, we can define $\hat{\Delta}_{PPC}$ as the solution to $\pi(\hat{\Delta}_{PPC}) = (1 - \beta)\pi(max\{0, \Delta_2^*\})$ and then the first-period off-platform discount under the PPC algorithm is given by $max\{0, \hat{\Delta}_{PPC}\}$.

[19] This follows since the seller never gives a discount lower than Δ_2^* in the first period, because this is the discount it would give if it were sure that it would never be featured again, and it can never raise its likelihood of being featured by offering a lower discount (and reducing the on-platform sales).

period 2 if

$$\theta \geq \frac{T}{(1 - H(\Delta_1))}$$

It follows that we can write down the seller's expected profit over two periods (anticipating the optimal period 2 discount) as follows:

$$\Pi_{seller}(\Delta_1, T) = \mu\pi(\Delta_1) + \beta\pi_2^* \int_{\min\left\{1, \frac{T}{(1-H(\Delta_1))}\right\}}^{1} \theta g(\theta)\ d\theta. \tag{5}$$

The first term is the expected seller type, μ, multiplied by the profit per interested consumer, $\pi(\Delta_1)$, and the integral in this expression is corresponds to the expected type of the seller, θ, when clearing the sales threshold multiplied by the probability of being featured in the second period.

In principle, the discount Δ_1 that a seller chooses can result in $T/(1-H(\Delta_1)) \geq 1$. In such a case, a seller forgoes second-period sales altogether. It is useful to introduce the notation $\bar{\Delta}_T \equiv H^{-1}(1-T)$ as the first-period discount above which the seller will not be featured in the second period for a given T.

The profit Π_{seller} may not be quasi-concave in Δ_1 if G is not sufficiently dispersed, even for $\Delta_1 < \bar{\Delta}_T$. To proceed, we will assume that it is for any $T \in (0, 1)$ and $\Delta_1 \in [\underline{\delta}, \bar{\Delta}_T]$. This is easily shown to hold for $G \sim U(0, 1)$ and $H \sim U(\underline{\delta}, 1)$ (our running example) and can be verified for other distributions as necessary.[20] However, since there is a discontinuous derivative at $\bar{\Delta}_T$ even if profits are quasi-concave to the left and right, the overall function may not be quasi-concave and the seller's profit may be maximized to the left or to the right of $\bar{\Delta}_T$, where in the latter case the seller gives up on period 2 sales.

We define $\bar{T}$ as the threshold at which the seller is indifferent between foregoing second-period sales and not; that is, it is the unique solution to the following equation:[21]

$$\mu\pi_2^* = \max_{\Delta_1 \in [\underline{\delta}, \bar{\Delta}_T]} \left\{ \mu\pi(\Delta_1) + \beta\pi_2^* \int_{\frac{T}{(1-H(\Delta_1))}}^{1} \theta g(\theta)\ d\theta \right\}.$$

Moreover, this solution must satisfy $\bar{\Delta}_{\bar{T}} < \Delta_2^*$. For $T > \bar{T}$ the seller will set $\Delta_1 = \Delta_2^*$, and for $T \leq \bar{T}$ the seller will set $\Delta_1 \leq \bar{\Delta}_T$. As we show next, only the former case is relevant for the remainder of the paper.

Lemma 3. *The platform chooses a threshold T^* that satisfies $T^* \leq \bar{T}$.*

[20] A sufficient condition is that $\frac{h(\Delta) g\left(\frac{T}{1-H(\Delta)}\right)}{(1-H(\Delta))^3}$ is non-decreasing in Δ_1 for any $T \in (0, 1)$ and any $\Delta_1 \in [\underline{\delta}, H^{-1}(1-T)]$.

[21] For T sufficiently close to zero, for example, for $T \leq 1 - H(\Delta_2^*)$, the optimal discount satisfies $\Delta_1 < \bar{\Delta}_T$, that is, the seller surpasses the sales threshold for a strictly positive measure of types. In contrast, for T sufficiently close to 1, the optimal discount is $\Delta_1 = \Delta_2^* > \bar{\Delta}_T$, that is, the seller is never featured in the second period and so sets the discount at the myopic level. Increasing T reduces $\Pi_{seller}(\Delta_1, T)$ for $\Delta_1 < \bar{\Delta}_T$, and reduces $\bar{\Delta}_T$ without altering seller profit for $\Delta_1 > \bar{\Delta}_T$, in which case $\Pi_{seller}(\Delta_1, T) = \mu\pi_2^*$, thus there must exist a solution $\bar{T}$ and our assumption on quasi-concavity ensures it is unique.

Proof. Suppose that $T > \bar{T}$ is optimal. For any such $T > \bar{T}$ the seller sets $\Delta_1 = \Delta_2^*$ and is not featured in the second period. The platform's profits are then $\mu f(1-H(\Delta_2^*))+\beta f\alpha(1-H(\Delta_2^*))$. By setting $T = 0$, the platform induces the seller to set a discount Δ_2^* in both periods and features the seller for sure, earning $\mu(1+\beta)f(1-H(\Delta_2^*))$, which by Assumption 1 (which ensures that the seller is ever worth being featured at all) is strictly higher than the profit from $T > \bar{T}$, a contradiction. □

Given this lemma, we only need to characterize the first-period discount Δ_1^* for $T \leq \bar{T}$. We can write the first-order condition for interior $\Delta_1^* \in (\underline{\delta}, \bar{\Delta}_T)$ as

$$\Delta_1 = f - \frac{H(\Delta_1)}{h(\Delta_1)} - \frac{\beta\pi_2^* r\left(\frac{T}{(1-H(\Delta_1))}\right)}{\mu(1-H(\Delta_1))}. \tag{6}$$

Recall that $r(\theta) \equiv \theta^2 g(\theta)$. Since the right-hand side of (6) is decreasing in Δ_1 given the log-concavity of H and the maintained assumption that $r'(\cdot) > 0$, while the left-hand side is increasing in Δ_1, this equation has at most one solution.

As to whether (6) has a solution or not, there are two cases to consider. If β is too large, then there will be no solution. Here, the optimal discount is $\underline{\delta}$ (in this case, the seller forgoes off-platform sales in period 1, which is as much as it can do to increase the chances of being featured in the second period). This is the case when

$$\beta \geq \hat{\beta} \equiv \frac{(f-\underline{\delta})\mu}{\pi_2^* r(T)}.$$

It is intuitive that the seller will be willing to forego off-platform sales in the first period to increase the chance of being promoted in the second period when the second period is sufficiently important, the fee is low (so that it is not too costly to divert sales to the platform), and second period profits are high. Otherwise, if $\beta < \hat{\beta}$, then $\Delta_1^* \in (\underline{\delta}, \bar{\Delta}_T)$ and it solves (6).

Proposition 5. *Suppose the seller's type is unknown and $T \leq \bar{T}$. If $\beta \geq \hat{\beta}$, then $\Delta_1^* = \underline{\delta}$; otherwise, Δ_1^* solves (6). Δ_1^* is lower than Δ_2^*, and if $\beta < \hat{\beta}$ then it is decreasing in β and T, and is increasing in f.*

Proof. The characterization of the solution is immediate following the discussion above. The comparative statics described follow on inspecting (6). □

The comparative statics described at the end of Proposition 5 are intuitive. The seller sets a lower off-platform price (a higher discount) when the on-platform fee is higher in order to divert more sales from the platform to a venue where it does not have to incur fees. The seller sets a higher off-platform price when it cares more for the future and is willing to sacrifice more to be featured in the next period, or when facing a "tougher" threshold to be featured.

It is worth discussing the optimal (interior) discount in detail and comparing it to the optimal period 2 discount that is characterized by Equation (2). The new (third) term in (6) is always negative, which means that the first-period discount, which is induced by algorithmic manipulation, is smaller than the second-period discount (where there is no future and no prospect of manipulating the algorithm). That is, $\Delta_1^* < \Delta_2^*$.

Notably, as the next result shows, Δ_1^* is (weakly) larger than the discount we obtained for the known type. So the seller with a known type is forced to give a smaller discount than the seller with an unknown type, or, equivalently, charges a higher off-platform price.[22] Intuitively, the platform cannot distinguish an off-platform price deviation from its uncertainty about the seller's type and thus is not able to extract the seller's rents as efficiently as when the type is known.

Corollary 1. *The first-period off-platform discount when the seller is of known type is weakly lower than that when the seller's type is unknown.*

Proof. Since the seller optimally chooses Δ_1^* rather than choosing Δ_2^* and foregoing the opportunity of second period profits, Δ_1^* for a seller of unknown type must satisfy

$$\mu\pi(\Delta_1^*) + \beta\pi_2^* \int_{\min\left\{1, \frac{T}{(1-H(\Delta_1))}\right\}}^1 \theta g(\theta)\, d\theta \geq \mu\pi_2^*$$

or, alternatively,

$$\pi(\Delta_1^*) \geq \left(1 - \beta\frac{1}{\mu}\int_{\min\left\{1, \frac{T}{(1-H(\Delta_1))}\right\}}^1 \theta g(\theta)\, d\theta\right)\pi_2^*$$

Since the integral divided by $\mu = \int_0^1 \theta g(\theta)\, d\theta$ is smaller than 1, it follows that $\pi(\Delta_1^*) > (1-\beta)\pi_2^*$. For the known seller type, the discount is either $\underline{\delta}$, the lowest possible, or satisfies $\pi(\Delta_1) = (1-\beta)\pi_2^*$, thus it must in all cases be lower than Δ_1^*. □

Given the equilibrium discount Δ_1^*, it will be useful to define $\tau \equiv T/(1 - H(\Delta_1^*(T))$ as the marginal type that just clears the sales threshold for the equilibrium first-period discount. It can be shown that τ is increasing in $T < \bar{T}$, even when taking into account the dependence of Δ_1^* on T.[23] It will be convenient to work with τ as the choice variable for the platform. We also define $\bar{\tau} \equiv \bar{T}/(1 - H(\Delta_1^*(\bar{T}))$ as the equilibrium type corresponding to the highest threshold that the platform may ever set, $\bar{T}$.

5.1 Optimal algorithm for unknown seller type

The above analysis was conducted for a fixed threshold T, or equivalently, the equilibrium threshold type τ. We now turn to finding the platform's profit-maximizing (restricted) algorithm.

[22] Note that the platform may, therefore, benefit from such uncertainty.

[23] This can be shown using log-concavity of H. Namely, from (6) we can find $\frac{\partial \Delta_1}{\partial T} = -\frac{\beta\pi_2^* h(\Delta_1)^2 r'(\tau)}{-\mu(1-H(\Delta_1))^2 H(\Delta_1)h'(\Delta_1)+2\mu h(\Delta_1)^2(1-H(\Delta_1))^2+\beta\pi_2^* h(\Delta_1)^3(\tau r'(\tau)+r(\tau))}$, which is easily shown to be negative by log-concavity of H. Thus $\frac{\partial\tau}{\partial T} = \frac{\partial \frac{T}{1-H(\Delta_1(T))}}{\partial T} = \frac{1+\tau h(\Delta_1)\frac{\partial\Delta_1}{\partial T}}{(1-H(\Delta_1))}$, which can be shown to be positive with simple manipulations and using log-concavity of H again.

If the platform cannot commit to the algorithm, then it will set $\tau = \alpha$. That is, a platform that cannot commit retains sellers if and only if they would earn the platform more profit in the second period than the outside option. However, given the inherent commitment associated with an algorithm, the platform can set a different threshold (that is, $\tau \neq \alpha$) because the threshold, τ, affects Δ_1^*. Sellers react to the threshold understanding that the discount they offer off-platform affects sales on the platform and the probability that they sell more than the threshold level, and earn profits in the second period. In this way, a different threshold can lead to a different off-platform discount.

From Lemma 3 we know that the optimal threshold type is below $\bar{\tau}$ and, therefore, induces the seller to set an optimal discount, Δ_1^*, that ensures that the product is featured with strictly positive probability in period 2.

If the platform features the seller in period 1 and keeps featuring it in period 2 if s_1 exceeds T, the platform's profits (given that $\tau \leq \bar{\tau}$) are

$$\Pi_{platform}(\tau) = f\left[\mu(1 - H(\Delta_1^*(\tau)) + \beta\left(\alpha G(\tau) + \int_{\tau}^{1} \theta g(\theta) d\theta\right)(1 - H(\Delta_2^*))\right]$$

In this expression, the platform earns a proportionate fee f on all sales on the platform, as captured in the square brackets. The first term inside brackets corresponds to the first period of sales on the platform, given that the expected type of seller is μ and the discount off the platform is $\Delta_1^*(\tau)$. The second term can be understood as, first, reflecting that the outside option (equivalent) to a type α is featured when sales are below T or, equivalently, the seller's type is lower than $\tau = \frac{T}{1-H(\Delta_1^*(T))}$; and, second, the seller's expected type conditional on being featured. Finally, conditional on the consumer liking the product, a fraction $1 - H(\Delta_2^*)$ buys on the platform.

The platform can gain from featuring the seller in the second period: it can assure itself the seller's upside compared to the outside option by setting $\tau = \alpha$ (but, as we show, it may earn even more with $\alpha \neq \tau$). In the first period, the expected profit that the platform earns when it features the seller is $f\mu(1 - H(\Delta_1^*))$, which is at least as high as $f\mu(1 - H(\Delta_2^*))$, and so by Assumption 1 is even higher than α.

In order to proceed, we make two assumptions. First, we will maintain that platform profit is quasi-concave in the threshold τ in the relevant range where $\tau \leq \bar{\tau}$.[24] The second assumption is related to the possibility that the platform optimally sets $\tau = \bar{\tau}$ because $\Pi'_{platform}(\bar{\tau}) > 0$. That is, the seller is pushed to the limit where she is indifferent between trying to be retained in the second period and giving up on the second period altogether. We will proceed by focusing on the parameters when this is not the case to simplify exposition and focus on central economic mechanisms. That is, in what follows, we assume that $\tau^* < \bar{\tau}$. In the examples that we provide, we check that this is indeed the case.[25]

[24] This is an indirect assumption on primitives but can be verified for our running example that features $G \sim U(0,1)$ and $H \sim U(\underline{\delta}, 1)$. See Appendix X for the proof.

[25] This circumstance may arise even for the example where both $G(.)$ and $H(.)$ are uniform when β is sufficiently low.

A key force in our analysis is that the threshold τ may be distorted in the direction that reduces Δ_1^* in order to raise the level of on-platform sales (and fee revenue for the platform). Since $Sign\left(\frac{\partial \Delta_1^*}{\partial \tau}\right) = -Sign\left(r'(\tau)\right) < 0$, and by our assumption that $r(\theta)$ is increasing on $[0, \mu]$, it follows that increasing τ induces the seller to reduce the first-period discount.

To derive the optimal threshold, which characterizes the optimal algorithm, take the derivative of platform profits with respect to the threshold-type:

$$\Pi'_{platform}(\tau) = f\left[(\alpha - \tau)\beta g(\tau)(1 - H(\Delta_2^*)) - \mu h(\Delta_1^*)\frac{\partial \Delta_1^*}{\partial \tau}\right]$$

Absent the effect that τ has on Δ_1^*, that is, when $\frac{\partial \Delta_1^*}{\partial \tau} = 0$, it is clear and intuitive that $\tau^* = \alpha$—if the choice of threshold has no impact on the seller's off-platform discount, then the platform would feature the seller if and only if it was better than the outside option. Since $\frac{\partial \Delta_1^*}{\partial \tau} < 0$, then $\tau^* > \alpha$ for any $\alpha > 0$. The platform distorts the threshold upwards so that the seller raises its off-platform first-period price (or equivalently, reduces its off-platform discount).

For $\Delta_1^* > \underline{\delta}$ (that is, in the case that we are examining where an interior solution is optimal for the seller), the platform's first order condition, 5.1, can be written as:

$$\tau^* = \alpha - \frac{\partial \Delta_1^*}{\partial \tau}(\tau^*) \cdot \frac{\mu h(\Delta_1^*)}{\beta g(\tau^*)(1 - H(\Delta_2^*))}. \tag{7}$$

When the outside option is worthless, $\alpha = 0$, the platform will always feature the product in the second period; that is, $\tau^* = 0$.[26] If α is so high that, even at $\tau = \alpha$, the seller would choose to forego off-platform sales by setting $\Delta_1^* = \underline{\delta}$, then the platform will optimally set $\tau^* = \alpha$. In this case, the platform gets the best of both worlds: it features the seller if and only if it is better than the outside option thereby maximizing second-period profits, while simultaneously inducing the seller to sell only on the platform in the first period. This happens when

$$\beta \geq \frac{\mu(f - \underline{\delta})}{\pi_2^* r(\alpha)}.$$

If there is no value of τ that induces the seller to set $\underline{\delta}$, that is

$$\max_{\tau \in [0,1]} r(\tau) < \frac{\mu(f - \underline{\delta})}{\beta \pi_2^*} \tag{8}$$

then the platform's optimal threshold must be interior. In this case, set the thresholds that we introduce below, $\underline{\alpha}$ and $\bar{\alpha}$, as both equal to 1.

Assume, instead, that $\max_{\tau \in [0,1]} r(\tau) \geq \frac{\mu(f - \underline{\delta})}{\beta \pi_2^*}$ holds. Define $\bar{\alpha}$ as the smallest root of

$$r(\alpha) = \frac{\mu(f - \underline{\delta})}{\beta \pi_2^*}.$$

[26] $\Pi'_{platform}(0) < 0$ for $\alpha = 0$ because $r'(0) = 0$.

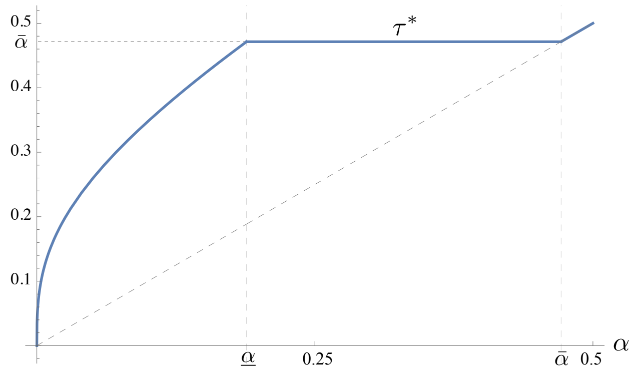

Figure 1: Optimal threshold type as a function of the outside option when $G \sim U(0,1)$, $H \sim U(-0.2, 1)$, $f = 0.4$, and $\beta = 2$.

If the platform's outside option exceeds $\bar{\alpha}$, then setting the threshold equal to the outside option $\tau^* = \alpha$ leads the seller to set $\Delta_1^* = \underline{\delta}$, clearly maximizing platform profits. If $\tau = \bar{\alpha}$, the seller sets $\Delta_1^* = \underline{\delta}$ and thus all period 1 sales are transacted on the platform.

There is a range of values of $\tau > \alpha$ at which the seller sets $\Delta_1^* = \underline{\delta}$, creating a kink in the platform's profit function. The platform may, therefore, choose $\tau^* = \bar{\alpha}$ when $\bar{\alpha} > \alpha$. For this case to arise, it must be the case that $\Pi'_{platform}(\bar{\alpha}) \geq 0$. This derivative is decreasing in α. The smallest value of the outside option that induces this case is then defined as

$$\underline{\alpha} \equiv \max \left\{ 0, \bar{\alpha} + \frac{\partial \Delta_1^*}{\partial \tau}(\bar{\alpha}) \cdot \frac{\mu h(\underline{\delta})}{\beta g(\bar{\alpha})(1 - H(\Delta_2^*))} \right\}.$$

Since $\frac{\partial \Delta_1^*(\tau)}{\partial \tau} < 0$, it is immediate that $\bar{\alpha} > \underline{\alpha}$. It follows that for $\alpha < \underline{\alpha}$, if any such cases exist, the seller sets its off-platform discount at $\Delta_1^* > \underline{\delta}$ and makes off-platform sales.

Proposition 6. *There are three possible cases for the optimal algorithm and associated seller behavior: (i) if $\alpha \geq \bar{\alpha}$ then $\tau^* = \alpha$ and $\Delta_1^* = \underline{\delta}$; (ii) if $\alpha \in [\underline{\alpha}, \bar{\alpha})$ then $\tau^* = \bar{\alpha}$ and $\Delta_1^* = \underline{\delta}$; (iii) if $\alpha < \underline{\alpha}$ then $\tau^* \in (\alpha, \underline{\alpha})$ and $\Delta_1^* > \underline{\delta}$.*

Proof. For $\alpha \geq \bar{\alpha}$, following Proposition 5 then for $\tau = \alpha$ the seller sets $\Delta_1^* = \underline{\delta}$, thus $\tau = \alpha$ is optimal. For $\alpha \in [\underline{\alpha}, \bar{\alpha})$ given the quasi-concavity of platform profits and the definition of $\underline{\alpha}$, $\Pi'_{platform}(\tau) \geq 0$ for all $\tau < \bar{\alpha}$, and platform profits are decreasing for $\tau > \bar{\alpha}$, thus they are maximized at $\bar{\alpha}$. For $\alpha < \underline{\alpha}$, given quasi-concavity of platform profits, there must be an interior maximizer given by Eq (7) on noting is never a corner solution at $\tau = 0$ because profit is increasing there. □

Figure 1 illustrates how the optimal threshold varies as a function of the outside option and illustrates the general qualitative features characterized in Proposition 6.

For low outside options (corresponding to low α or A) the platform excludes some profitable seller types in the second period to induce higher sales on the platform in the first period. For an intermediate outside option, the induced discount is $\underline{\delta}$ (foreclosing off-platform sales) and, thus, the

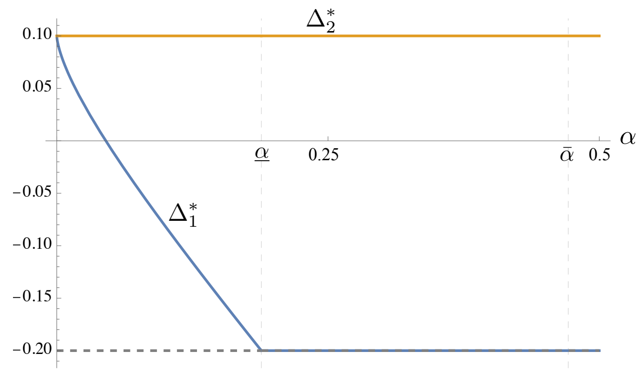

Figure 2: Optimal first period discount (blue) and the myopic/second-period discount (orange) when $G \sim U(0,1)$, $H \sim U(-0.2,1)$, $f = 0.4$, and $\beta = 2$.

platform chooses the threshold that just induces this discount. This threshold is independent of the outside option and, thus, the plot of τ^* is flat in this region. For high outside options, even efficient second-period promotion, $\tau = \alpha$, induces $\Delta_1^* = \underline{\delta}$ and so the platform sets the threshold at the myopically (and socially) optimal level.

Figure 2 illustrates qualitative features of the first-period discount set by the seller as a function of the outside option, as characterized by Proposition 6.

As is intuitive, the discount is decreasing in the outside option of the platform. The higher the outside option, the higher the sales threshold chosen, leading to a lower off-platform discount. For sufficiently large outside options, the discount becomes maximal, that is set at the level of $\underline{\delta}$ which in this example is equal to -0.2.

5.2 OPP algorithm

Suppose the platform can condition its algorithm on the seller's off-platform price. Then, as in the second period or in the case of a known-type seller, it can require $\Delta_1 < \underline{\delta}$ as a prerequisite for exposure, effectively foreclosing the off-platform channel. Moreover, the platform can perfectly infer the seller's type and choose to promote in the second period if and only if doing so is more profitable; that is, if $\tau = \alpha$.

5.3 PPC algorithm

As described in Proposition 3, there are circumstances in which the seller would choose a higher price off-platform than on-platform under the restricted algorithm, and in these cases, the PPC has no bite in the second period. As a result, in these cases, the seller's value of being featured in the second period—along with the platform's profit from featuring a seller of a given type—is identical with or without a PPC.

Turning back to the first period in such cases: if, under the restricted algorithm, a seller would again choose a negative discount (i.e., a higher off-platform price), the PPC would have no bite at the same threshold for promotion. However, the platform may benefit from reducing τ, with less concern that this would induce a larger off-platform discount, since the PPC limits this discount. We can write the resulting discount as $\Delta_1^{PPC}(T) = \min\{0, \Delta_1^*(T)\}$. This suggests that the platform can lower τ, bringing it closer to α, thereby improving promotion efficiency while mitigating concerns about sales channel inefficiency.

Of course, the same logic applies if the seller would choose a positive discount under the optimal threshold for the restricted algorithm: the PPC would then ensure that prices are identical on- and off-platform, again eliminating sales channel inefficiency.

The analysis in which the PPC also binds in the second period is more involved, since the seller (and platform) now expects second-period profits that differ from those under the restricted algorithm. Consequently, the maximization problems faced by both the seller and the platform are altered, and their responses to changes in the promotion threshold differ. In principle, this leads to a more complex problem.

Analogous to Equation (6), for example, we can write:

$$\mu\pi(\min\{0, \Delta_2^*\}) = \max_{\Delta_1 \in [0, \bar{\Delta}_T]} \left\{ \mu\pi(\Delta_1) + \beta\pi(\min\{0, \Delta_2^*\}) \int_{\frac{T}{1-H(\Delta_1)}}^{1} \theta g(\theta)\, d\theta \right\}. \tag{9}$$

This expression reflects both of the aforementioned changes: the adjusted second-period profit term on the right-hand side, and the restricted domain over which the maximization is taken.

Accordingly, we introduce the notation $\overline{T}^{PPC}$ and $\hat{\beta}^{PPC}$ as natural analogues to those defining corner solutions in the restricted algorithm case.

We return to this characterization and notation in our welfare analysis below, but note that central forces discussed above remain. In particular, a PPC—when binding—brings on- and off-platform prices together, eliminating sales-channel misallocation, and can reduce discounts without requiring threshold manipulation, thereby improving promotion efficiency.

6 Welfare analysis

The discussion above highlights two key forces that influence welfare and serve as potential sources of inefficiency in our setting. First, there is misallocation across sales channels, whereby consumers may purchase through a channel that offers lower value when on- and off-platform prices differ. Second, the platform may choose not to feature a seller even when doing so would generate greater social value than featuring the alternative.

To evaluate this second source of inefficiency—and to conduct a broader welfare analysis—we must take a stance on the welfare implications of the platform's outside option.

Before we do so, it is convenient to introduce notation for the off-platform consumer disutility (per interested consumer), suffered through purchasing off-platform rather than directly from the platform (which may be negative), given a discount Δ:

$$CD(\Delta) \equiv \int_{\underline{\delta}}^{\Delta} \delta h(\delta)\ d\delta.$$

To the extent that the platform's alternative to featuring a given seller is to feature some other seller, both the platform's profit and the associated welfare impact of this alternative may vary depending on the type of algorithm the platform is allowed to use. That is, when comparing the restricted algorithm with an OPP or PPC algorithm, the value of the alternative, A, may differ—leading us to define corresponding terms A^{OFF} and A^{PPC}, each of which may generate different welfare and consumer surplus outcomes.

In this context, a natural benchmark is to assume that the alternative is a known seller of type α. Then, in the relevant second period,

$$A^{OFF} = \alpha(1 - H(\Delta_2^*)),$$

with no consumer surplus or disutility under the OPP algorithm, since there are no off-platform sales (that is, $CD(\underline{\delta}) = 0$). Under the PPC algorithm,

$$A^{PPC} = \alpha(1 - H(\max\{0, \Delta_2^*\})),$$

with associated consumer disutility given by $CD(\max\{0, \Delta_2^*\})$.[27]

Total welfare across both periods under the restricted algorithm can be written as:

$$W^* = \mu(1 - CD(\Delta_1^*)) + \beta(1 - CD(\Delta_2^*)) \left(\int_{\tau^*}^{1} \theta g(\theta)\ d\theta + G(\tau^*)\alpha \right)$$

This follows from the observation that, in the first period, there are expected to be μ interested consumers, and the surplus associated with each is given by the value of the platform (1), minus any consumer disutility from off-platform purchases ($CD(\Delta_1^*)$). The corresponding expression for the second period includes a final term in brackets, which can be interpreted as the likelihood that a consumer is interested in the featured product. This term captures both the probability and conditional expectation associated with a seller being featured in the second period, as well as the corresponding terms for the alternative option being featured instead.

For the OPP algorithm, which forecloses the off-platform channel in both periods and features a seller in the second period if and only if it is better than the alternative, welfare is given by:

$$W^{OPP} = \mu + \beta \left(\int_{\alpha}^{1} \theta g(\theta)\ d\theta + \alpha G(\alpha) \right).$$

Similarly, for the PPC algorithm, we can write:

$$W^{PPC} = \mu(1 - CD(\Delta_1^{PPC}) + \beta(1 - CD(\Delta_2^{PPC}) \left(\int_{\tau_{PPC}^*}^{1} \theta g(\theta)\ d\theta + \alpha G(\tau_{PPC}^*) \right),$$

where the notation τ_{PPC}^* and Δ_t^{PPC} reflect that the threshold for promotion and the seller's first-

[27] We characterize these expressions for the second period, since our focus is on the distortion associated with featuring the seller in that period and its relation to the extent of sales-channel misallocation.

period discount under the optimal PPC algorithm may differ from those of the restricted algorithm.

6.1 Comparing the restricted and OPP algorithm

We begin by writing the difference in welfare in these two cases.

$$W^{OPP} - W^* = \mu CD(\Delta_1^*) + \beta \left[CD(\Delta_2^*) \left(\alpha G(\tau^*) + \int_{\tau^*}^{1} \theta g(\theta)\ d\theta \right) + \int_{\alpha}^{\tau^*} \theta g(\theta)\ d\theta \right]$$

There are two key differences between these scenarios. First, the OPP algorithm eliminates the consumer disutility from the alternative (off-platform) channel, which may or may not be welfare-improving depending on the value of $\underline{\delta}$. Second, it always leads to more efficient promotion of sellers: since $\tau^* \geq \alpha$, the final term in the welfare expression is positive. Indeed, the OPP algorithm always maximizes promotion efficiency.

If $\underline{\delta} = 0$, then both $CD(\Delta_1^*)$ and $CD(\Delta_2^*)$ are positive, so eliminating them is unambiguously beneficial. In this case, any off-platform purchase is inefficient, and the OPP algorithm—by foreclosing the off-platform channel—eliminates sales-channel misallocation. Thus, when $\underline{\delta} = 0$, the OPP algorithm, by both eliminating sales-channel inefficiency and maximizing promotion efficiency, ensures maximal welfare.

However, when $\underline{\delta} < 0$, these results no longer necessarily hold—and it is straightforward to construct counterexamples. In general, there may be a trade-off between promotion efficiency (which is maximized by the OPP algorithm) and sales-channel efficiency (where the comparison between the two algorithms depends on parameters).

We summarize this discussion as follows.

Proposition 7. *For $\underline{\delta} = 0$, the OPP algorithm maximizes welfare and therefore increases social welfare relative to the restricted algorithm. For $\underline{\delta} < 0$, the welfare associated with the OPP algorithm may be higher or lower.*

Proof. See discussion above. □

Figure 3 compares welfare outcomes, showing first-best welfare (which, of course, dominates the alternatives) alongside welfare under the OPP algorithm and welfare under the restricted algorithm, as functions of the outside option α. As expected, the higher the quality of the alternative seller that might be featured, the higher the first-best welfare. Similarly, welfare under the OPP algorithm also increases in α.

The OPP algorithm always eliminates all off-platform consumer disutility and achieves efficient promotion (and thus will feature the outside option in the second period more often as the quality of this outside option improves). By contrast, under the restricted algorithm, welfare can fall in α over a range where promotion inefficiency worsens. In this example, the OPP algorithm dominates the restricted algorithm when α is intermediate and the promotion distortion associated with the restricted algorithm is significant. However, for high α the restricted algorithm generates higher

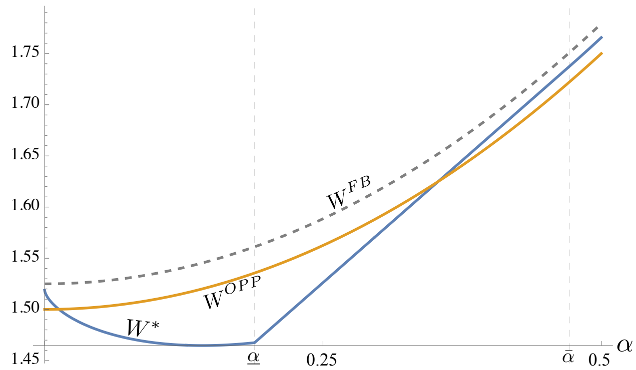

Figure 3: Welfare in the first-best (dashed gray), restricted algorithm (blue) and OPP algorithm (orange) as functions of the outside option. $G \sim U(0,1)$, $H \sim U(-0.2,1)$, $f = 0.4$, and $\beta = 2$.

welfare: at sufficiently high α, both algorithms foreclose the off-platform channel in the first period and promote efficiently, but in the second period the OPP algorithm induces greater sales-channel misallocation. The dominance of OPP in the intermediate range illustrates that simply prohibiting algorithms from relying on off-platform prices is not an unambiguously beneficial policy.

6.2 Comparing the restricted and PPC algorithm

Trivially, in cases where the PPC has no bite in any period, outcomes are identical to those under the restricted algorithm, and both consumer surplus and welfare are unchanged.

Almost as immediate is the case where $\underline{\delta} = 0$. Here, the seller would never choose a negative discount, as this would imply setting a higher off-platform price, which is strictly suboptimal. Thus, the PPC would bind in both periods. Moreover, in this case, no consumer would buy off-platform, since doing so would entail paying the same price as on the platform while incurring a strict disutility $\delta > 0$. Therefore, the PPC would fully foreclose off-platform sales and function identically to the OPP algorithm. As in the first part of Proposition 7, and for the same reasons, the PPC algorithm maximizes welfare in this case. Proposition 8 below further suggests that the PPC can maximize welfare in a broader set of environments.

As discussed in Section 5.3, the optimal PPC algorithm may constrain prices, even if it would not do so under the threshold chosen by the restricted algorithm. If $\Delta_2^* \leq 0$ without PPC (which implies that also $\Delta_1^* < 0$ following Proposition 5), and in response to PPC the platform adjusts the threshold, then the PPC must bind at the new threshold, $\Delta_1^{PPC}(\tau^{PPC}) = 0$. Moreover, in such a case, the platform necessarily sets an efficient threshold $\tau^{PPC} = \alpha$, as can be seen in the proof of Proposition 8 , below. This means that PPC induces efficient channel allocation in the first period and efficient promotion in the second. If, instead, the platform chooses not to adjust the threshold, which happens if restricted algorithm profits are higher than the profits associated with $\Delta_1^{PPC}(\tau^{PPC}) = 0$ and $\tau^{PPC} = \alpha$, then PPC has no effect.

More interesting cases involve $\underline{\delta} < 0$ (so that some consumers prefer buying off-platform) and where $\Delta_2^* > 0$ (so that PPC binds in the second period). In this case, since the PPC reduces the seller's second-period profit (that is, $\pi(0) < \pi_2^*$), career concerns are lessened, and the seller would set a higher discount in the first period. This can be seen from the new FOC for the interior solution, which can be compared to 6 and uses $\pi(0)$ rather than π_2^* in the numerator of the final term.

$$\Delta_1 = f - \frac{H(\Delta_1)}{h(\Delta_1)} - \frac{\beta\pi(0)r\left(\frac{T}{(1-H(\Delta_1))}\right)}{\mu(1-H(\Delta_1))}. \tag{10}$$

The following is immediate.

Lemma 4. *Suppose the seller's type is unknown, $\underline{\delta} < 0$, $T \leq \bar{T}^{PPC}$, and under the restricted algorithm $\Delta_2^* > 0$. If $\beta \geq \hat{\beta}^{PPC}$, then $\Delta_1^{PPC} = \underline{\delta}$; otherwise, Δ_1^{PPC} solves (10) or is set to 0 if the solution is positive.*

For a given sales threshold, T we can see that efficiency is increased by PPC everywhere. If under the restricted algorithm $\Delta_1^* \geq 0$, then the PPC algorithm either sets $\Delta_1^{PPC} = 0$, or $0 \geq \Delta_1^{PPC} > \Delta_1^*$. In either case, misallocation across channels is necessarily reduced.

Note, however, the PPC also affects the platform's optimal threshold and, therefore, promotion efficiency. For a given Δ_1, the platform cares more about the second period because it retains higher demand in the second period; $1 - H(0) > 1 - H(\Delta_2^*)$, this suggests that it would adjust the threshold in the direction of α raising promotion efficiency.

In general, the platform has to take into account not only the corner solution of the seller's problem at $\underline{\delta}$ but also a new corner solution at 0 induced by PPC. We write the platform's problem as

$$\Pi_{\text{Platform}}(T) = f\Bigg[\mu\left(1 - H\left(\Delta_1^{PPC}(T)\right)\right) + \\ \beta\left(\alpha G\left(\frac{T}{1 - H\left(\Delta_1^{PPC}(T)\right)}\right) + \int_{\frac{T}{1-H\left(\Delta_1^{PPC}(T)\right)}}^{1} \theta g(\theta)\, d\theta\right)(1 - H(0))\Bigg].$$

We focus on two interesting special cases. First, we will study what happens if PPC fully achieves its goal in that it binds in both periods (after the platform has adjusted the threshold optimally). Second, we study the situation where equilibrium pricing is interior for the seller, $\Delta_1^{PPC}(\tau^{PPC}) \in (\underline{\delta}, 0)$, and the platform's threshold is also interior, $T^{PPC} < \bar{T}^{PPC}$.

Proposition 8. *If PPC binds in both periods, then $\tau^{PPC} = \alpha$ and the PPC algorithm maximizes welfare.*

Proof. If the platform optimally sets τ such that PPC binds, $\Delta_1^{PPC} = 0$, then $\tau^{PPC} = \alpha$ must hold. If it does not, then consider $\tau^{PPC} < \alpha$. This cannot be optimal because increasing τ improves second-period profits and also (weakly) decreases Δ_1^{PPC} by Lemma 4. Now consider $\tau^{PPC} > \alpha$. Reducing τ improves second-period profits and does not alter Δ_1 which remains at 0 (again by Lemma 4). Profits would rise overall, and, so, $\tau^{PPC} = \alpha$ must hold. Since in addition to

$\tau^{PPC} = \alpha$ we have $\Delta_1^{PPC} = \Delta_2^{PPC} = 0$ by the supposition that PPC binds in both periods, welfare is maximized because all distortions (both sales channel misallocation and promotion inefficiency) are eliminated. □

If under the PPC algorithm, the threshold, τ, is such that the seller sets $\Delta_1^{PPC} < 0$ so that PPC binds in period two but does not in period one, this does not necessarily mean that $\Delta_1^{PPC} = \Delta_1^*$ because both the seller and platform problems are altered by PPC binding in period 2.

The platform's FOC for the interior case is

$$\tau^{PPC} = \alpha - \frac{\partial \Delta_1^{PPC}}{\partial \tau}(\tau^{PPC}) \cdot \frac{\mu h(\Delta_1^{PPC})}{\beta g(\tau^{PPC})(1 - H(0))}.$$

This is analogous to the platform's FOC in the baseline model with two differences: first, the platform retains more of the second-period demand if it features the seller, corresponding to the $1 - H(0)$ term; and, second, the seller's response to a change of τ is governed by (10). The former force suggests that the platform cares more about the future and, so, has a diminished incentive to distort the threshold as discussed above. The latter force is in general ambiguous, but for uniform H we can show that it goes in the same direction: reducing the seller's second-period profits (through the PPC) makes the seller's optimal first-period discount less responsive to the threshold, τ, in absolute magnitude, thus reducing τ further.

Proposition 9. *Assume that H is uniform and that the first-period discount is interior with or without a PPC; that is, $\Delta_1^*, \Delta_1^{PPC} \in (\underline{\delta}, 0)$. Then $\tau^{PPC} < \tau^*$ and $\Delta_1^{PPC} > \Delta_1^*$, so imposing a PPC increases welfare compared to the restricted algorithm. OPP algorithm may be better or worse than either PPC or the restricted algorithm.*

Proof. See Appendix A.3 □

Taken together, the results above highlight many cases in which a PPC can either maximize welfare or improve upon the restricted algorithm. Intuitively, a binding PPC brings on- and off-platform channel prices together, eliminating sales-channel inefficiency. As argued above, it often does so in a way that also eliminates or reduces promotion inefficiency.

Figure 4 illustrates welfare comparisons for the uniform case. For these parameters, the PPC binds in the second period and leads to equal on- and off-platform prices, whereas without the PPC the discount would be positive.[28]

As in Proposition 8, for low α the PPC binds in both periods, so that $\tau^{PPC} = \alpha$ and first-best welfare obtains.[29] The discrete drop in the W^{PPC} curve occurs when α is high enough that the platform abandons the efficient promotion threshold and the PPC-induced zero discount. Instead, the platform profitably distorts τ to the level it would have chosen if the PPC applied only in

[28]Specifically, it would be 0.1, as can be seen from Figures 2 and 1, which plot discounts and thresholds for the same parameters in the baseline model.

[29]For α close to zero, Δ_1^* is close to Δ_2^*, which in turn is positive by assumption.

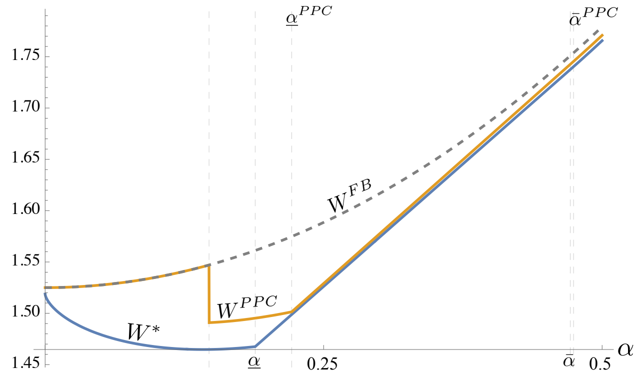

Figure 4: Welfare in the first-best (dashed gray), restricted algorithm (blue) and PPC algorithm (orange) as functions of the outside option. $G \sim U(0, 1)$, $H \sim U(-0.2, 1)$, $f = 0.4$, and $\beta = 2$.

period 2, while gaining from inducing a negative period 1 discount—that is, a higher off-platform price than the on-platform price.[30]

Because even in this case there is less promotion distortion with PPC, the seller responds with a larger (i.e., closer to zero) discount. Consequently, just after the drop, the orange curve lies above the blue one. This is consistent with Proposition 9, which compares welfare when discounts (and thus thresholds) are interior in both cases.

As the figure shows, in this example the PPC algorithm dominates the restricted algorithm for all α, including values outside the scope of Propositions 8 and 9, such as where both curves slope upward for $\alpha \in [\underline{\alpha}^{PPC}, \bar{\alpha}]$.[31] In that range, under both regimes, the seller diverts all sales to the platform, but the PPC algorithm achieves this with a more distortive threshold. Nevertheless, the PPC algorithm still delivers higher welfare because it generates lower sales-channel misallocation in the second period—in fact, it eliminates it entirely. This latter effect outweighs the greater promotion inefficiency of the PPC algorithm, keeping the orange curve above the blue curve throughout.

7 Extensions

The core insights of the model—concerning how ranking algorithms affect off-platform pricing and seller selection—remain robust across various extensions. While additional considerations may arise in richer settings, the key forces we identify are grounded in the model's primitives and continue to operate in natural generalizations.

In particular, our framework highlights how algorithmic design can induce sales-channel misallocation through its impact on off-platform prices, including via career-concerns-like incentives. It also shows how algorithms influence promotion efficiency through their determination of rankings. These mechanisms are not specific to our baseline assumptions and carry over to extended

[30] As α increases, setting $\tau = \alpha$ leads to a lower discount than the seller would choose in the absence of PPC. For sufficiently high α, the discount is close enough to zero so that the PPC's impact is minimal, making it profitable to set a much higher threshold—one that would have been optimal without PPC in period 1.

[31] This ordering is clearly example-specific and may change with different parameters or functional forms, which is why we cannot state that the PPC algorithm always improves welfare.

environments.

Furthermore, some welfare properties of the PPC and OPP algorithms follow almost directly from their definitions. A PPC, by equalizing on- and off-platform prices, reduces sales-channel misallocation by construction. An OPP algorithm, which removes the off-platform option altogether, avoids the need to distort the promotion threshold and thus maximizes promotion efficiency. These intuitions—highlighted and clarified in our baseline analysis—provide a foundation for understanding the role of algorithmic design in a broader class of settings.

7.1 Alternative fee structures

Our model is not well-suited for a rich analysis of endogenous fee-setting, given the simplifying assumption that all interested consumers have unit demand and share a common valuation for the product. In this setting, it is immediate that an OPP algorithm would choose a fee of $f = 1$, extracting the full surplus. Similarly, the PPC algorithm would also extract full surplus in many cases. In particular, in this context, fee-setting does not affect our welfare conclusions: both the OPP and PPC algorithms maximize welfare (as shown in Propositions 7 and 8) at any fee level up to this threshold, provided that $\underline{\delta} = 0$.

Of course, in practice, fee-setting is shaped by many additional considerations. These include competitive pressures, regulatory concerns, and the use of uniform fees across a broad range of products that may differ substantially in their cost-to-valuation ratios. Such forces, which are not addressed in our model, may constrain fee levels in important ways.

With a per-unit fee, because of the unit demand, nothing of substance changes in our analysis. The platform still looks to maximize on-platform sales, and so the per-unit fee would enter our expressions, for example 6, in a way identical to the proportionate fee.

When fees are fixed—so that the platform earns the same amount regardless of whether a sale occurs on- or off-platform—the platform does not care about the type of the seller beyond their willingness to pay the fee. This makes the notion of rankings and the algorithms that allocate them somewhat moot. In this context, perhaps fee determination through auctions is a more interesting question, but one beyond the scope of our analysis.[32]

7.2 Off-platform fees

In this section, we explore the possibility that the off-platform channel is itself another platform that charges a fee $f_0 < f$. We do not allow this alternative platform to attract new consumers; though, of course, a more symmetric treatment in which multiple platforms serve as venues both for product discovery and as sales channels would be an interesting extension. Instead, as in the baseline model, we assume that consumers learn about the product only through the primary platform. One can think of Shopify.com as an example of this kind of "enabler" platform—one that facilitates transactions but does not itself generate consumer traffic.

[32] In Bar-Isaac and Shelegia (2024) we addressed the question of whether to use an ad auction or an algorithm with a proportionate fee in a model without learning.

Given fees on the primary and enable platforms, the seller's profit per consumer is given by $(1-H(\Delta))(1-f)+H(\Delta)((1-\Delta)(1-f_0))$. This can be written as

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

Thus the analysis of seller behavior above (in Sections 3 4, 4) can be applied directly but under an adjusted fee $\tilde{f}\equiv\frac{f-f_0}{1-f_0}$.[33] Since $\tilde{f}$ is decreasing in f_0, its increase has the same effect on the seller's behavior for a given threshold τ as a reduction in f does. As we know from the platform's optimization over T, it depends on f only through Δ_1^*, thus the same change of variable can be applied throughout. Consequently, not only can the characterization results be readily modified but our welfare results from Section 6 apply directly.

Following the discussion above, comparative statics on f_0 are the inverse of those on f. In the second period, increasing f_0 induces the seller to decrease the discount, which harms consumers but may or may not increase welfare.[34] The effect on the first-period interior discount of an unknown seller type is more nuanced, even when the platform threshold is held constant.[35] On the one hand, increasing f_0 (reducing $\tilde{f}$) reduces π_2^*, making being retained in period 2 less attractive, and thereby putting upward pressure on Δ_1. On the other hand, for the same reason as in period 2, reducing the discount becomes more profitable in period 1. In addition to these two effects, f_0 affects the platform's threshold-setting incentives because it changes how Δ_1^* responds to τ on the margin. Thus, the overall effect on Δ_1^* is, in general, ambiguous.[36]

7.3 Repeat purchases

In many marketplaces, consumers who discover sellers in one period return to their direct channel for repeat purchases, bypassing the platform where they discovered these sellers. Platforms exert efforts to avoid such disintermediation by not allowing sellers and buyers to interact to the extent possible.[37] In this section, we will explore the effect of repeat purchases in our environment.

In order to maintain the two-period assumption, and in line with the fact that it's the first-period consumers that play a key role for learning, we assume that only first-period consumers are interested in repurchasing the good in period two. Otherwise, the model is unaltered. We assume that repeat consumers retain both of their characteristics across periods(that is, whether they find the good of the seller appealing, and the disutility they suffer from purchasing off-platform).

There are two plausible modeling choices appropriate for different settings: either all consumers who purchase in period one can go directly to the seller's off-platform offering (natural in service

[33]The way to understand this adjusted fee is to think that the relative share of the revenue that the seller keeps on the platform compared to off-platform is $\frac{1-f}{1-f_0}$ which then leads to the effective fee $\tilde{f}=1-\frac{1-f}{1-f_0}$.

[34]Welfare in the second period goes up with f_0 if $\tilde{f}>\frac{H(0)}{h(0)}$ and goes down otherwise.

[35]If $\Delta_1^*=\underline{\delta}$, then there is no effect.

[36]For H uniform, as can be seen from (14), increasing f_0 reduces period 1 discount for a constant τ, and τ is independent of f_0, thus the discount falls in both periods.

[37]For a wide-ranging discussion and analysis of the use of fees but not algorithms in such a context, see Enache and Rhodes (2025)

sectors where the purchase results in a visit to the seller) or only those who buy off-platform in the first period can bypass the platform in period 2.

We focus on the former assumption, which leads to a neutrality result. To simplify, we assume that repeat consumers do not engage with any of the platform algorithms in period 2. If they do not have a match with the seller in period 1, either they do not return to the platform, or, if they do, they get the same price as "fresh" visitors in period 2, and buy wherever the net utility is higher.

For a given threshold T, the seller's profits over the two periods are

$$\Pi_{seller}(\Delta_1, \Delta_2, T) = \mu\pi(\Delta_1) + \beta\pi(\Delta_2) \int_{\min\left\{1, \frac{T}{(1-H(\Delta_1))}\right\}}^{1} \theta g(\theta)\; d\theta + \pi(\Delta_2)\mu.$$

Note that repeat consumers return to the seller whether or not it is featured.[38] Their presence does not alter second-period pricing, where Δ_2^* remains optimal.

Similarly, the platform's profits can be written as follows

$$\begin{aligned} \Pi_{platform}(T) &= f\Bigg[\mu(1 - H(\Delta_1^*(T)) + \beta\Bigg(\alpha G\left(\frac{T}{(1 - H(\Delta_1^*(T))}\right) \\ &+ \int_{\frac{T}{(1-H(\Delta_1^*(T))}}^{1} \theta g(\theta) d\theta\Bigg)(1 - H(\Delta_2^*)) + \mu(1 - H(\Delta_2^*))\Bigg]. \end{aligned}$$

Clearly, the presence of repeat consumers does not alter the seller's or the platform's optimization problem: for both, repeat consumers simply generate an additional fixed profit. Thus, our baseline analysis applies without modification.

7.4 Seller knows its type

In this section we allow the seller to know its type θ. First, it is relatively straightforward to show that with appropriate beliefs no signaling takes place, so the seller of any type will set $p_t = 1$ in both periods on the platform.[39] The intuition is that the seller's pricing incentives are independent of its type, which merely scales demand for any given price.

We proceed by examining off-platform pricing, assuming that the platform uses a threshold T. While this threshold need not be optimal in this setting, it facilitates comparison with the baseline model. To further simplify, we assume $\beta \geq 1$, which ensures that any seller with $\theta > T$ has the incentive to meet the threshold.

Following our analysis in Section 4, it is clear that the seller will use the following strategy: if $\theta \leq T$ or $\theta > \frac{T}{1-H(\Delta_2^*)}$ then $\Delta_1 = \Delta_2^*$; that is, if the seller cannot meet the threshold even if diverting all sales, or the myopic discount allows it to meet the threshold, it simply chooses the

[38] In writing the expression, we assume that the product is available on the platform, but not seen by fresh consumers, if not featured and we interpret β as representing the masses of different consumer populations rather than involving any discounting.

[39] Assume the platform holds passive beliefs regarding Δ_t and θ for any p_t, and promotes the seller if $p_t = 1$. Any deviation from $p_t = 1$ is then unprofitable.

myopic discount. Otherwise, it will set the first period discount at a level that allows it to just meet the threshold, $\Delta_1 = H^{-1}(1 - T/\theta)$, so all these intermediate types bunch at the sales target T. Given such pricing by the various seller types, the platform's profits are given by

$$\begin{aligned}\Pi_{platform} &= f\Bigg[(1-H(\Delta_2^*))\int_0^T \theta g(\theta)\,d\theta + T\int_T^{\min\left\{\frac{T}{1-H(\Delta_2^*)},1\right\}} g(\theta)\,d\theta \\ &+ (1-H(\Delta_2^*))\int_{\min\left\{\frac{T}{1-H(\Delta_2^*)},1\right\}}^1 \theta g(\theta)\,d\theta + \beta(1-H(\Delta_2^*))\left(\alpha G(T) + \int_T^1 \theta g(\theta)\,d\theta\right)\Bigg]\end{aligned}$$

The first-order condition at an interior threshold (assuming $T^* < 1 - H(\Delta_2^*)$) takes a new form:

$$\beta g(T)(1-H(\Delta_2^*))(\alpha - T) + \left(G\left(\frac{T}{1-H(\Delta_2^*)}\right) - G(T)\right) - Tg(T)H(\Delta_2^*) = 0. \tag{11}$$

The first term in (11) reflects a familiar incentive in the second period where, in the absence of the effect T has on Δ_1^*, the platform would set T at α.[40] The second term in brackets reflects the mass of sellers' types who bunch at T, thus giving the platform extra sales when T is higher. The last term reflects the losses from the types who, instead of bunching at the threshold (and the marginal such type does so by diverting all sales to the platform), give up on trying to attain this sales threshold and take a fraction $H(\Delta_2^*)$ of consumers off the platform. The combined effect of the last two terms may be positive or negative, thus T may be distorted upwards or downwards depending on $g(\cdot)$, just as in the base model where we assumed $\theta^2 g(\theta)$ is increasing to obtain that the distortion is upwards. If $g(\cdot)$ is non-decreasing (e.g. the uniform distribution) then the two terms are strictly positive; thus, the threshold is distorted upwards if $g(\cdot)$ is not too decreasing.

In this version of the model, familiar trade-offs between sales channel misallocation and promotion efficiency arise as in the baseline model.

The effect of OPP and PPC algorithms is also similar. OPP leads to $\Delta_1^{OPP} = \Delta_2^{OPP} = 0$ and $\tau^{OPP} = \alpha$.

PPC has no effect if $\Delta_2^* \leq 0$. If, instead, $\Delta_2^* > 0$ without PPC, then $\Delta_2^{PPC} = 0$. In the first period, for a given threshold T the seller sets $\Delta_1 = 0$ if $\theta < T$ or $\theta > \frac{T}{1-H(0)}$, and otherwise sets $\Delta_1 = H^{-1}(1 - T/\theta)$ as without PPC. Thus, PPC reduces the discount for all cases, strictly so for $T \notin (T, T/(1 - H(0)))$.

Welfare effects are also familiar. For instance, if $\underline{\delta} = 0$ then both OPP and PPC achieve first-best welfare by inducing $\Delta_1 = \Delta_2 = 0$ and $\tau = \alpha$. If $\underline{\delta} < 0$ then OPP may or may not improve welfare depending on whether it sufficiently worsens channel misallocation or not. As before, if PPC binds in both periods (once the threshold has been optimally adjusted), it will necessarily achieve first-best because $\tau^{PPC} = \alpha$ must hold.[41]

[40] In this setting T ends up being the threshold type as well.

[41] The proof is similar to the case where the seller does not observe its type. As there, if PPC binds in period 1, then $\tau \neq \alpha$ cannot be optimal.

7.5 Downward-sloping demand

As in the main model, suppose each consumer has unit demand and gets utility $v-p$ on the platform and $v-\tilde{p}-\delta$ off the platform and that all consumers have an outside option of 0. But in contrast with the main model, suppose that consumers are heterogeneous in their (on-platform) valuations for the good, v, leading to a downward-sloping demand on the platform. Suppose that v and δ are independent with well-behaved CDFs $H(\delta)$ and $F(p)$, where we denote $q(p) \equiv 1-F(p)$.[42]

The seller's second-period profits are

$$(p_2-\Delta_2)\int_{\underline{\delta}}^{\Delta_2} q(p_2+\delta-\Delta_2)h(\delta)\,d\delta + (1-f)(1-H(\Delta_2))q(p_2)p_2,$$

As we explain below, the algorithm will be able to dictate p_2, thus the seller's only choice variable is Δ_2. Its first-order condition with respect to Δ_2 is

$$\Delta_2 \quad = \quad fp_2 - \frac{\int_{\underline{\delta}}^{\Delta_2}[q(p_2-\Delta_2+\delta)+(p_2-\Delta_2)q'(p_2-\Delta_2+\delta)]h(\delta)\,d\delta}{q(p_2)h(\Delta_2)}.$$

Note that the optimal off-platform price for a given on-platform price is below the one that maximizes off-platform profits because a higher discount diverts sales from the platform (where the seller incurs the fee).[43]

The platform's second period profits are proportional to on-platform sales, $q(p_2)(1-H(\Delta_2))$. The algorithm can effectively dictate p_2 by threatening not to feature the seller if it chooses a different price on the platform. In effect, the platform is a price leader who sets p_2 and the seller is a price follower who reacts with Δ_2. Thus, the platform maximizes its profits with respect to p_2, taking into account how Δ_2 is set by the seller, yielding the first order condition:

$$p_2 \quad = \quad \frac{q(p_2)}{-q'(p_2)+q(p_2)\frac{h(\Delta_2)}{1-H(\Delta_2)}\frac{\partial \Delta_2}{\partial p_2}}$$

where $\frac{\partial \Delta_2}{\partial p_2}$ is implicitly defined by the seller's first-order condition, as above.

If $\frac{\partial \Delta_2}{\partial p_2} > 0$, as is the case for H uniform and $q(p) = 1-p$, then the on-platform price that the algorithm induces is below the revenue-maximizing price because lowering p_2 induces the seller to set a lower off-platform discount and therefore shifts demand to the platform.[44]

If $\Delta_2 < 0$, then increasing f improves the allocation between channels but also creates a deadweight loss. In our baseline setting, this latter effect is absent, so such an increase in f is always welfare-improving, whereas in the more general case considered here, it may not be.

[42]This model is equivalent to the one where all consumers share the same downward-sloping demand $q(p)$ and δ is a disutility of purchasing off platform per unit purchased.

[43]This condition generalizes (2) which obtains by setting $q(p) = 1$.

[44]In the baseline model with identical consumers, the revenue-maximizing price is 1. In general, for the same reason as here, in the baseline model the platform may choose to force $p_2 < 1$ in order to induce a lower off-platform discount. Appendix A.1 outlines conditions where this is not the case, e.g. H is uniform. Since demand is downward-sloping here, no such corner solution is available.

Under the OPP algorithm, the outcome is simpler. First, the platform forecloses off-platform sales by requiring $\Delta_2 < \underline{\delta}$. Second, to maximize on-platform sales, the algorithm induces the seller to set p_2 at the revenue-maximizing level, which—as argued above—is higher than the price under the restricted algorithm. Consequently, the on-platform price is higher, raising deadweight loss, while off-platform sales are foreclosed. The latter could either increase or decrease sales-channel misallocation relative to the restricted algorithm.

The PPC case is more involved here than in the earlier analysis. Even if $\Delta_2^* \leq 0$ under the restricted algorithm, the platform may wish to induce a different p_2 so that $\Delta_2 \geq 0$ binds under PPC.[45] If the PPC binds at the platform-optimal p_2 and $\frac{\partial \Delta_2}{\partial p_2} > 0$, then p_2^{PPC} is optimally set at the revenue-maximizing level. Increasing p_2 from a lower level would not change the discount (which remains at 0) but would raise platform revenues. If PPC does not bind, then necessarily p_2^{PPC} and Δ_2^{PPC} are set at the same level as under the restricted algorithm.

If $\underline{\delta} = 0$, then while the OPP and PPC algorithms are equivalent, it is no longer clear that they deliver higher welfare than the restricted algorithm. The reason is that with downward-sloping demand, welfare depends not only on the price difference across channels but also on the absolute price level. With the OPP or PPC algorithm, the platform would induce the seller to set p_2 at the revenue-maximizing level, which is higher than the price under the restricted algorithm if $\frac{\partial \Delta_2}{\partial p_2} > 0$, and shut down off-platform sales. The latter effect, just as in our baseline model, eliminates sales channel misallocation. However, the price effect introduces deadweight loss. Consequently, overall welfare effects are ambiguous.

Given that even the myopic second-period pricing problem is involved, incorporating downward-sloping demand into the full analysis—with first-period pricing and learning—remains analytically cumbersome, even though it is numerically tractable.

8 Conclusion

We have presented a novel model of algorithmic learning that embeds a career concerns framework within a model of platform disintermediation. We have shown that learning algorithms can be used to influence sellers' pricing strategies on channels outside the platform. Perhaps surprisingly, allowing algorithms to directly incorporate off-platform prices into their decisions can, in some cases, increase welfare. Our analysis highlights the roles of sales-channel misallocation and promotion efficiency in determining these outcomes. In doing so, it contributes to the broader understanding of how algorithm design shapes seller behavior and market efficiency in the context of online platforms, revealing both potential gains and trade-offs. In particular, our results show that seemingly simple and intuitive policies—such as banning algorithms from using price parity clauses—can have detrimental consequences for welfare.

[45] This is immediate if $\Delta_2^* = 0$ without PPC: increasing p_2 beyond the pre-PPC level p_2^* towards the revenue-maximizing price must be optimal, since Δ_2 remains equal to zero for all such prices while platform profits increase.

References

Bar-Isaac, Heski, and Sandro Shelegia. 2024. "Monetizing Steering."

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: Algorithmic and Data-Driven Bids in Digital Advertising."

Bergemann, Dirk, and Alessandro Bonatti. 2024. "Data, competition, and digital platforms."

Boik, Andre, and Kenneth S. Corts. 2016. "The Effects of Platform MFNs on Competition and Entry." *Journal of Law and Economics*, 59(1).

Calvano, Emilio, and Michele Polo. 2020. "Market Foreclosure and MFN Clauses in Platform Competition." *Journal of Industrial Economics*, 68(2): 304–341.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello. 2025. "Artificial intelligence, algorithmic recommendations and competition." *Algorithmic Recommendations and Competition (February 19, 2025).*

Castellini, Jacopo, Amelia Fletcher, Peter L Ormosi, and Rahul Savani. 2023. *Recommender systems and competition on subscription-based platforms.* SSRN.

Enache, Andreea, and Andrew Rhodes. 2025. "Platform Disintermediation with Repeated Transactions."

Franck, Jens-Uwe, and Martin Peitz. 2024. "The digital markets act and the whack-a-mole challenge." *Common Market Law Review*, 61(2).

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

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

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."

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. 2025. “Selling on Platforms: Demand Boost versus Customer Migration.” Working paper.

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

Mantovani, Andrea, Claudio A Piga, and Carlo Reggiani. 2017. “The dynamics of online hotel prices and the EU Booking. com case.” *Com Case (September 30, 2017) .NET Institute Working Paper*, , (17-04).

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.

A Omitted proofs and further results

A.1 The optimality of setting the on-platform price at 1

In our analysis, we assume that the $p_1 = p_2 = 1$, that is the on-platform price is set optimally at a level that extracts consumers' full valuation on the platform. In this appendix, we provide conditions under which it is true. The reason why this may not be the case for the restricted algorithm is that the on-platform price can affect the seller's off-platform discount.

First, we show that for the uniform case the platform induces $p_2 = 1$ via the algorithm. In the second period, for a given p_2 the seller's discount is $\Delta_2^* = \frac{fp_2+\delta}{2}$. The platform's period two profits are $fp_2\left(1-\frac{\Delta_2^*-\delta}{1-\underline{\delta}}\right) = fp_2\left(1-\frac{fp_2-\delta}{2(1-\underline{\delta})}\right)$. This is clearly increasing for all $p_2 \leq 1$, thus a corner solution obtains. From the perspective of period 1, higher p_2 increases the seller's period 2 profits, thus can only increase the platform's period 1 profits by inducing a higher Δ_1^*.

Turning to the first period (assuming $p_2 = 1$), the seller sets

$$\Delta_1^* \quad = \quad \frac{1}{4}\left(2+fp_1+\underline{\delta}-\sqrt{\frac{2\beta}{\mu}r(\tau)\left(2-\underline{\delta}-f\right)^2+\left(2-\underline{\delta}-fp_1\right)^2}\right)$$

The platform's period 1 profits holding τ fixed (second-period profits are independent of p_1) are

$$f\mu p_1\left(1-\frac{\Delta_1^*-\underline{\delta}}{1-\underline{\delta}}\right).$$

Substituting for Δ_1^* and taking the derivative of this expression with respect to p_1, we obtain

$$\mu f\frac{\left(L-fp_1\right)\left(2+L-\underline{\delta}-fp_1\right)}{4L\left(1-\underline{\delta}\right)},$$

where $L \equiv \sqrt{\frac{2\beta}{\mu}r(\tau)\left(2-\underline{\delta}-f\right)^2+\left(2-\underline{\delta}-fp_1\right)^2}$. This derivative is positive since $fp_1 \leq 1$, $\underline{\delta} \leq 0$, and $L > fp_1$. Thus, the platform will optimally induce $p_1 = 1$.

In general, for distributions other than the uniform, we can obtain $\frac{\partial \Delta_2^*}{\partial p_2}$ implicitly from $\Delta_2 = fp_2 - \frac{H(\Delta_2)}{h(\Delta_2)}$ (generalization of (2) for $p_2 \leq 1$) and then find the derivative of period 2 platform profits $\frac{\partial(fp_2(1-H(\Delta_2^*(p_2)))}{\partial p_2}$. Suppose that it is positive for all $p_2 \leq 1$ this would ensure that the algorithm would ensure that $p_2 = 1$. After some algebraic manipulations, a sufficient condition is that $\frac{h(x)^2(xh(x)+H(x))}{2h(x)^2-H(x)h'(x)} + H(x) < 1$ holds for the x that is the unique solution of $\frac{H(x)}{h(x)} + x = f$. For example, for $H(x) = x^l$ with $l > 0$, this condition holds for all $l \leq 1$, but for $l > 1$ we require that $f < \frac{1}{l}\left(\frac{1}{l+1}\right)^{\frac{1-l}{l}}$.

The general analysis for the first period becomes unwieldy. However, for any particular parameterization, it is simple to verify, at least numerically, that setting $p_1^* = 1$ is optimal.

A.2 Proof of Lemma 2

The algorithm $\phi(s)$ induces an equilibrium discount Δ_1^* associated with the seller's problem. Specifically, Δ_1^* is the solution to

$$\max_{\Delta_1 \in [\underline{\delta},1]} \mu\pi(\Delta_1) + \beta\pi_2^* \int_0^1 \theta\phi(\theta(1 - H(\Delta_1))g(\theta)d\theta$$

In this expression, the first term corresponds to the expected first-period profits (recall that $\mu = E[\theta]$). The second term corresponds to the second period. π_2^* is the profit per consumer who is exposed to and likes the good. The likelihood that a consumer likes the good is θ, and the likelihood that the good is featured in the second period, so that a consumer is exposed, is given by $\phi(\theta(1 - H(\Delta_1))$ when the seller sets first-period discount Δ_1. Of course, these likelihoods must be integrated, accounting for all possible values of θ and the probabilities with which they arise.

In general, the platform's problem is to choose the function ϕ in order to maximize its expected profits, which are given by the following expression.

$$f\mu(1 - H(\Delta_1^*)) + f\beta(1 - H(\Delta_2^*))\left(\alpha + \int_0^1 (\theta - \alpha)\phi(\theta(1 - H(\Delta_1^*))g(\theta)d\theta\right).$$

Again, the first expression corresponds to the first period in which the platform earns a fee on the expected on-platform sales. The second expression incorporates that the platform may feature the outside option, in which case, on-platform sales would be $(1 - H(\Delta_2^*))\alpha$, or it may not, in which case, there are additional sales of $(\theta - \alpha)(1 - H(\Delta_2^*))$. The platform's expected probability of featuring the seller of type θ is given by $\phi(\theta(1 - H(\Delta_1^*))$ where the anticipated discount Δ_1^* is itself a function of ϕ.

Suppose that ϕ is not a threshold strategy and maximizes the platform's profits. There must exist positive measure sets Σ_1 and Σ_2 in $[0,1]$ such that for any $s_1 \in \Sigma_1$ and any $s_2 \in \Sigma_2$, the following properties hold: $s_1 < s_2 < (1 - H(\Delta_1^*))$,[46] $\phi(s_1) > 0$, $\phi(s_2) < 1$, and

$$\int_{\Sigma_1} g\left(\frac{s_1}{1 - H(\Delta_1^*)}\right) ds_1 = \int_{\Sigma_2} g\left(\frac{s_2}{1 - H(\Delta_1^*)}\right) ds_2 > 0.$$

That is, Σ_2 corresponds to higher types/sales levels, all elements in Σ_1 have some chance of being featured, and all in Σ_2 have some chance of not being featured; and for sets Σ_1 and Σ_2 the corresponding types (associated with discount Δ_1^*) have equal measure.[47] Let

$$l = \frac{\int_{\Sigma_2} \frac{s_2}{1-H(\Delta_1^*)} g\left(\frac{s_2}{1-H(\Delta_1^*)}\right) ds_2}{\int_{\Sigma_1} \frac{s_1}{1-H(\Delta_1^*)} g\left(\frac{s_1}{1-H(\Delta_1^*)}\right) ds_1}$$

[46] Taking Δ_1^* as given, sales above $(1 - H(\Delta_1^*))$ never arise, and so assigning $\phi(s) = 1$ for $s \geq (1 - H(\Delta_1^*))$ does not alter platform payoffs. Moreover, it can only induce lower Δ_1^* and so would clearly be preferable for the platform.

[47] The existence of such sets is guaranteed by the supposition that ϕ is not threshold type, and then the measures of Σ_1 and Σ_2 can be adjusted to make corresponding type measures equal.

be a scaling factor. Using Eq (12) and that $s_1 < s_2$ for any $s_1 \in \Sigma_1$ and $s_2 \in \Sigma_2$, it follows that $l > 1$. Consider a new algorithm, with associated $\tilde{\phi}(.)$, which is identical to $\phi(.)$ outside of Σ_1 and Σ_2; that is $\tilde{\phi}(s) = \phi(s)$ for $s \in [0,1] \backslash (\Sigma_1 \cup \Sigma_2)$. Otherwise, it is less likely to feature a seller who generates a sales level in Σ_1 and more likely to feature a seller who generates a sales level in Σ_2 but does so in a way that ensures that the solution to the seller's problem is a discount level at, or below, the level that it would choose under $\phi(.)$, Δ_1^* . In particular, $\tilde{\phi}(s) = \phi(s) + \varepsilon$ for $s \in \Sigma_2$ and $\tilde{\phi}(s) = \phi(s) - \varepsilon l$ for $s \in \Sigma_1$ where ε is sufficiently small such that $\phi(s_1) - \varepsilon l \geq 1$ and $\phi(s_2) + \varepsilon \leq 1$ hold on respective sets.

By construction, profit for the seller from choosing Δ_1^* is the same at $\phi(s)$ and $\tilde{\phi}(s)$. To see this, write the difference between profits at ϕ and $\tilde{\phi}$ for discount Δ_1 as

$$M(\Delta_1) \equiv \beta \pi_2^* \varepsilon \left[\int_{\Sigma_2} \frac{s_2}{1 - H(\Delta_1)} g\left(\frac{s_2}{1 - H(\Delta_1)}\right) ds_2 - l \int_{\Sigma_1} \frac{s_1}{1 - H(\Delta_1)} g\left(\frac{s_1}{1 - H(\Delta_1)}\right) ds_1 \right].$$

From the definition of l, $M(\Delta_1^*) = 0$. Thus, the transformation does not affect profit at Δ_1^*.

For any $s_2 > s_1$, $\frac{g(xs_2)}{g(xs_1)}$ is weakly decreasing in x if the elasticity of the density g, defined as $\frac{\theta g'(\theta)}{g(\theta)}$, is weakly decreasing. Take s between Σ_1 and Σ_2 and consider $\Delta_1 > \Delta_1^*$. Write

$$\frac{\int_{\Sigma_2} \frac{s_2}{1-H(\Delta_1)} g\left(\frac{s_2}{1-H(\Delta_1)}\right) ds_2}{\int_{\Sigma_1} \frac{s_1}{1-H(\Delta_1)} g\left(\frac{s_1}{1-H(\Delta_1)}\right) ds_1} = \frac{\int_{\Sigma_2} \frac{s_2}{1-H(\Delta_1)} \frac{g\left(\frac{s_2}{1-H(\Delta_1)}\right)}{g\left(\frac{s}{1-H(\Delta_1)}\right)} ds_2}{\int_{\Sigma_1} \frac{s_1}{1-H(\Delta_1)} \frac{g\left(\frac{s_1}{1-H(\Delta_1)}\right)}{g\left(\frac{s}{1-H(\Delta_1)}\right)} ds_1} < \frac{\int_{\Sigma_2} \frac{s_2}{1-H(\Delta_1^*)} \frac{g\left(\frac{s_2}{1-H(\Delta_1^*)}\right)}{g\left(\frac{s}{1-H(\Delta_1^*)}\right)} ds_2}{\int_{\Sigma_1} \frac{s_1}{1-H(\Delta_1^*)} \frac{g\left(\frac{s_1}{1-H(\Delta_1^*)}\right)}{g\left(\frac{s}{1-H(\Delta_1^*)}\right)} ds_1} = l$$

where the inequality follows from the decreasing elasticity of $g(.)$. Thus the profit difference $M(\Delta_1)$ for any $\Delta_1 > \Delta_1^*$ is negative. An identical argument demonstrates that when $\Delta_1 < \Delta_1^*$ it is positive. This means that optimal Δ_1 (which maximizes the seller's profit) cannot increase.

Holding Δ_1^* fixed (which may, thus, only fall further increasing the platform's profits), the change in the platform's profits is

$$f\alpha(1 - H(\Delta_2^*))\varepsilon(l - 1) \int_{\Sigma_1} g\left(\frac{s_1}{1 - H(\Delta_1^*)}\right) ds_1 > 0. \tag{12}$$

The platform is better off because a switch from $\phi(s)$ to $\tilde{\phi}(s)$ increases the probability of the outside option being selected without changing the average type of the seller that is retained. We have arrived at a contradiction.

A.3 Proof of Proposition 9

Since period 1 discount is interior in both regimes, it is given by the respective FOC. Similarly, the optimal platform threshold is characterized by the corresponding first-order condition (FOC).

Without PPC, FOCs were derived in (14) and (15). With PPC, instead, we get

$$\tau^{PPC} = \alpha + \frac{r'(\tau)}{2g(\tau)\sqrt{\left(\frac{2-\underline{\delta}-f}{2(1-\underline{\delta}-f)}\right)^2 + \frac{2\beta r(\tau)}{\mu(1-\underline{\delta}-f)}}}$$

$$\Delta_1^{PPC} = \frac{1}{4}\left(2+\underline{\delta}+f-(2-\underline{\delta}-f)\sqrt{1+\frac{4(1-\underline{\delta}-f)}{(2-\underline{\delta}-f)^2}\cdot\frac{2\beta r(\tau)}{\mu}}\right).$$

For the same τ, since $\frac{4(1-\underline{\delta}-f)}{(2-\underline{\delta}-f)^2} < 1$, it follows that $\Delta_1^{PPC} > \Delta_1^*$. Moreover, since the left-hand side of the platform's first-order condition crosses the right-hand side from above (by the assumed quasi-concavity), and given that $f > \underline{\delta}$, we have

$$1+\frac{2\beta r(\tau)}{\mu} < \left(\frac{2-\underline{\delta}-f}{2(1-\underline{\delta}-f)}\right)^2 + \frac{2\beta r(\tau)}{\mu(1-\underline{\delta}-f)}.$$

It follows that $\tau^{PPC} < \tau^*$. Since Δ_1^{PPC} is decreasing in τ^{PPC}, it follows that $\Delta_1^{PPC} > \Delta_1^*$, even though each is evaluated at its respective threshold.

To summarize, we have:

$$\alpha < \tau^{PPC} < \tau^* \quad \text{and} \quad \Delta_2^* > \Delta_2^{PPC} = 0 > \Delta_1^{PPC} > \Delta_1^*.$$

Thus, distortions are lower under PPC in both periods, and therefore welfare must be higher with PPC.

A.4 Proof of concavity of the platform profits for uniform $H(\cdot)$ and $G(\cdot)$

Assume $H(\cdot) \sim U(\underline{\delta}, 1)$ and $G(\cdot) \sim U(0,1)$. Then

$$\Pi_{platform}(\tau) = \frac{2-\underline{\delta}-f}{8(1-\underline{\delta})}\left(1+\beta\big(2+4\alpha\tau-2\tau^2\big)+\sqrt{1+4(1-\underline{\delta})^2\beta\tau^2}\right),$$

Taking the second derivative gives

$$\Pi''_{platform}(\tau) = \frac{(2-\underline{\delta}-f)\beta}{2(1-\underline{\delta})}\left(-1+\frac{(1-\underline{\delta})^2}{(1+4(1-\underline{\delta})^2\beta\tau^2)^{3/2}}\right).$$

Since $(1-\underline{\delta})^2 \in [0,1]$ we have

$$-1+\frac{(1-\underline{\delta})^2}{(1+4(1-\underline{\delta})^2\beta\tau^2)^{3/2}} < 0.$$

It follows that

$$\Pi''_{platform}(\tau) < 0 \quad \text{for all} \quad \tau \in (0,1].$$

B Uniform $H(\cdot)$

In this appendix, we restrict attention to the case in which $\delta \sim H(\cdot) = U(\underline{\delta}, 1)$ with $\underline{\delta} \in [-1, 0]$. This allows for some simpler characterizations used to verify assumptions and underlining our numerical illustrations.

B.1 Analogs of Section 3 results

For the second period, it can be verified that $\Delta_2^* = (f + \underline{\delta})/2$ and $\pi_2^* = \frac{(2-\underline{\delta}-f)^2}{4(1-\underline{\delta})}$. Note that even if $f = 1$ (the highest possible fee by assumption), whereby the seller makes zero profits on the platform, it nevertheless chooses not to give such a large discount that all consumers buy off-platform, preferring instead to charge a higher price and sell to fewer consumers. Welfare is maximized at $f = -\underline{\delta}$ and consumer surplus is maximized at $f = 1$.[48] If the platform can set the fee optimally, then it would choose $f^* = 1 - \frac{\underline{\delta}}{2}$. Even though the market is covered, the platform chooses an interior fee which balances higher fee vs leakage off platform.

B.2 Analogs of Section 4 results

For the known type, the first period discount is given by $\Delta_1 = 1 - (1 - \underline{\delta})\frac{T}{\theta}$ if $\frac{T}{\theta} \in \left(\frac{2-\underline{\delta}-f}{2(1-\underline{\delta})}, 1\right]$ and $\beta < \left(1 - \frac{2(1-\underline{\delta})T}{\theta(2-f-\underline{\delta})}\right)^2$, otherwise $\Delta_1 = (f + \underline{\delta})/2$.

Turning to the platform's behavior,

$$\underline{\beta} = \frac{(f - \underline{\delta})^2}{(2 - \underline{\delta} - f)^2}$$

is the critical threshold such that if $\beta \geq \underline{\beta}$, the platform sets $T^* = \theta$ and subsequently the seller sets $\Delta_1^* = \underline{\delta}$ in the first period. If, instead, $\beta < \underline{\beta}$, then the optimal sales threshold is

$$T^* = \frac{\theta \left(4 - \underline{\delta} - f + \sqrt{\beta}\,(2 - \underline{\delta} - f)\right)}{4\,(1 - \underline{\delta})}$$

and the corresponding discount is

$$\Delta_1^* = \frac{1}{2}\left(\underline{\delta} + f - \sqrt{\beta}\,(2 - \underline{\delta} - f)\right). \quad (13)$$

B.3 Analogs of Section 5 results

For the unknown seller, the interior discount is

$$\Delta_1^* \quad = \quad \frac{1}{4}\left(2 + \underline{\delta} + f - (2 - \underline{\delta} - f)\sqrt{1 + \frac{2\beta r(\tau)}{\mu}}\right). \quad (14)$$

[48] Although a higher fee is ruled out by assumption, the seller may still be willing to incur a loss on platform sales, as this is the only way to reach consumers and generate revenue from off-platform sales.

This is readily shown to be higher than the discount for the known type in (13).

The interior platform threshold solves

$$\tau^* \quad = \quad \alpha + \frac{r'(\tau)}{2g(\tau)\sqrt{1+\frac{2\beta r(\tau)}{\mu}}}. \tag{15}$$

The critical threshold, $\bar{\alpha}$, is the smaller root of [49]

$$r(\alpha) = \frac{4\mu(f-\underline{\delta})\,(1-\underline{\delta})}{\beta\,(2-\underline{\delta}-f)^2},$$

and $\underline{\alpha}$ is equal to

$$\underline{\alpha} = \max\left\{0, \bar{\alpha} - \frac{2\mu\,(1-\underline{\delta})\,(2-\underline{\delta}-f)\,r'\,(\bar{\alpha})}{g\,(\bar{\alpha})\left(8\mu\,(1-\underline{\delta})^2 + \beta\,(2-\underline{\delta}-f)^2\,r\,(\bar{\alpha})\right)}\right\}$$

[49] Recall that if no root exists then $\underline{\alpha} = \bar{\alpha} = 1$.