Search, Quality, and Subscription Platforms

Eduard Airapetian

Abstract

I model a subscription platform that (i) charges a subscription fee, (ii) pays per-stream royalties that differ by quality, and (iii) controls the first item consumers see. Consumers conduct pseudo-ordered search with a reservation utility; creators, facing heterogeneous effort costs, decide whether to supply high quality. The platform's royalty is a cost, so ranking decisions are a steering tool that trades royalty savings against user satisfaction and, in turn, reshapes creators' effort incentives. Solutions pin down the reservation cut-off, the creator cost threshold, and the platform's joint choice of royalty gap and bias. The framework I have built allows us to answer different questions regarding the subscription platforms pricing, content management, and quality provision, and it can be used further to answer different specific questions for a platforms regulation. I also provide comparative statics linking search frictions, royalty heterogeneity, and catalog quality, and offers a benchmark for current debates on Spotify's "Discovery Mode" and algorithmic self-preferencing.

1 Introduction

1.1 Motivation

The migration of film, music, and book catalogs to subscription platforms has replaced physical scarcity with a very different bottleneck: consumer attention. A Netflix user choosing a film tonight confronts more than 3,600 titles in the US catalog, yet empirical session data show that the typical viewer abandons the search within ten minutes. Because every additional query, scroll, or trailer view consumes cognitive time, search is costly even though the marginal price of content is zero. The platform therefore acts as a gatekeeper: by deciding which single title is shown first, and in what order the remainder appear, it effectively allocates the scarce resource of attention. Crucially, this allocation occurs without any market transaction in the classical sense—no posted prices are inspected, and the consumer has already paid a flat subscription fee. The consumer's choice is thus based not on whether a movie costs $3.99 or $5.99 to rent, but on which movie the algorithm chooses to place at the top of a personalized carousel when marginal viewing is free to the user but expensive to the platform.

Digital subscription platforms such as *Spotify*, *Netflix*, and *Kindle Unlimited* rely on a dual monetization scheme: consumers pay an upfront fee for unlimited access, while the platform compensates content providers with a marginal royalty for each consumption event. This *two-sided pricing schedule* is not a mere accounting device; it creates a powerful incentive for the platform to steer demand. Every additional stream incurs an incremental royalty cost that can be mitigated if the platform's recommendation engine diverts attention toward lower-royalty items or its own in-house content. Because the platform also controls the order in which users search—through rankings, playlists, and default carousels—it can quietly trade off royalty savings against user satisfaction.

Recent debates over Spotify's "Discovery Mode," Amazon's Buy-Box self-preferencing, and Netflix's promotion of studio originals illustrate the stakes of this trade-off.[1] The recent work of Aguiar et al. (2024, 2021) establishes empirical evidences that streaming platforms indeed bias their interface choices towards lower royalty artists and this has a substantial effect on the further success of the piece of content. This mechanism of the demand steering is crucial to understanding the ways platform operate and affect the market as a whole, it can change the consumers preferences, influence incentives of the content providers, and ulti-

[1]See, e.g. Chen et. al. (2016) on the Buy-Box and Zhu and Liu (2018) on platform entry into content categories.

mately changes the dynamics of cultural-goods production. I build a quantitative framework that links pricing, ranking, and creator incentives, and use it to evaluate platform design and regulation.

To the best of my knowledge, no existing model examines how a subscription platform's per-use royalty costs interact with its search ranking decisions and content quality provision. This gap is crucial, because it is exactly in these interactions that platforms may sacrifice consumer welfare for cost savings. Existing theories isolate only fragments of this problem. Ordered-search models treat the ranking as exogenous and focus on price competition under search frictions (Arbatskaya, 2007; Armstrong, 2017), while the growing literature on intermediary bias abstracts from two-sides pricing and endogenous quality provision (Bar-Isaac et al., 2012; Hagiu and Jullien, 2011). Consequently, we know little about how a profit-maximizing subscription platform jointly chooses (i) the search order, (ii) the per-stream royalty, and (iii) the subscription fee when creators can adjust their effort—and thus content quality—in response. This paper fills that gap by embedding a two-sides pricing tariff into an ordered-search environment with endogenous effort, thereby revealing when steering toward cheap content is privately profitable, welfare improving, or socially costly.

How does a subscription platform's per-stream royalty (a marginal cost for the platform) reshape what content is recommended to consumers, and what are the equilibrium outcomes on content creators' quality choices and consumer welfare? Specifically, I ask: (i) under what combinations of search costs and royalty gaps does a profit-maximizing platform bias its ranking toward low-royalty (or own-label) items; (ii) how does this bias propagate back to the supply side by altering creators' incentives to join high-effort versus low-effort regimes; (iii) can the royalty-saving motive ever make such bias welfare improving, or is it generically distortionary; and (iv) what regulatory or contractual instruments restore efficiency?

I structure the model with the following ingredients that, taken together, deliver the environment in which the research questions can be posed and answered. First, a two-sides pricing tariff—a fixed subscription fee plus per–stream royalties—is indispensable because, unlike advertising or retail platforms, Spotify and Netflix pay rather than receive a marginal transfer for each additional consumption event; the platform's steering motive is therefore cost–minimization, and this motive would vanish in a single–price world with zero marginal cost. Second, ordered (directed) search gives the intermediary a lever—the identity of the first recommendation—with which to steer demand; in a random–search or full–information setting the notion of "ranking bias" would be meaningless. Third, endogenous creator effort is required because exogenous quality would make ranking bias a pure permutation of a fixed catalog, leaving unanswered the policy–relevant question of whether steering ultimately discourages or stimulates the supply of high-quality content. Fourth, the first-item bias is the mechanism I have chosen to capture the platforms ability to manipulate the order in which consumer meet the content. I have only focused on the first one as it already a very powerful instrument for steering as the search costs make the first item unproportionally successful. Finally, a positive per-draw search cost $s > 0$ is critical: if $s = 0$, consumers would scroll costlessly until they located their top match and the platform could no longer influence stopping decisions, while a purely random search protocol would, by assumption, remove the ranking instrument altogether. Each ingredient therefore plugs a specific logical hole: remove any one of them and either the steering incentive disappears or the feedback loop to quality collapses.

My analysis develops a novel economic framework that integrates consumer search, platform pricing, and content quality choices in a single model. It is the first to jointly consider a subscription platform's ranking order, two-sided pricing (subscription fee with differentiated per-stream royalties), and creators endogenous effort provision, showing how these three elements together determine the platform's optimal recommendation bias and its welfare implications. This closes the gap between ordered-search models with exogenous quality (Armstrong, 2017) and two-part-tariff models without search frictions (Rochet and Tirole, 2003). The analysis derives all components of the model and gives the guideline to its further development. Finally, my paper outlines a normative criterion to assess when the platform's steering is socially excessive vs. welfare-improving, and it uses the model to evaluate policy interventions – such as imposing ranking neutrality or a minimum royalty floor – that could potentially restore efficient outcomes. All of these contributions are derived within a unified, tractable framework and are grounded in formal analysis of the platform's equilibrium with endogenous search and quality decisions. These results inform ongoing debates surrounding Spotify's "Discovery Mode" and self-preferencing investigations of Amazon and Apple.

1.2 Results

The model first characterizes consumers' optimal search behavior in closed form. Proposition 1 establishes that for any positive search cost s and any content quality mix λ (the fraction of high-quality items), there is a unique reservation utility threshold $z^*(s, \lambda)$ such that consumers stop searching and consume an item if and only if its realized utility exceeds z^*. The solution for z^* is explicit and depends on search cost relative to content quality: if search is very cheap ($s < \lambda/2$), the threshold z^* lies above the high-quality baseline (meaning consumers become very selective, rejecting all low-quality content); if search is costly ($s > \lambda/2$), z^* falls below the high-quality benchmark (consumers settle more quickly). Given this stopping rule, Proposition 5 derives the probabilities (α_H, α_L) that the item a consumer ultimately consumes is high- or low-quality, respectively. These probabilities depend on the platform's bias parameter θ (the probability that the first shown item is high-quality) and the threshold z^*. The closed-form expressions show that search behavior amplifies effective quality: the share of high-quality content in consumed streams α_H is always higher than the share of high-quality in the catalog λ (i.e. users disproportionately consume the better content). In fact, if the consumer's bar for stopping is high (costly search, $z^* > 1$), low-quality items are never accepted, so every consumed item is high-quality in those regimes. Thus, even a modest search friction causes users to favor quality, especially when the platform biases the first recommendation toward high-quality content.

Building on the consumer side, I then analyze creators' effort decisions and the resulting equilibrium composition of content. Proposition 6 defines an effort cost cutoff c^* for creators: a creator will choose high effort (produce high-quality content) if and only if their private cost c is below $c^*(\theta, \lambda, s)$. This threshold equates the marginal creator's cost to the expected benefit of high effort, which is given by the difference in expected royalty earnings between high- and low-quality strategies. Substituting the consumption probabilities, the cutoff can be written as

$$c^* = r_H \frac{\alpha_H(\theta, \lambda, s)}{\lambda} - r_L \frac{\alpha_L(\theta, \lambda, s)}{1 - \lambda},$$

meaning the platform's royalty policy (r_H, r_L) and the consumer acceptance rates (α_H, α_L) jointly determine which creators find it profitable to invest in quality. In equilibrium, the measure of high-quality content λ adjusts such that it equals the fraction of creators with $c \leq c^*$ (i.e. $\lambda = F(c^*)$ where F is the cost distribution). The model admits a unique solution for this fixed-point λ under general conditions, and in the special case of a uniform cost distribution. Intuitively, higher platform bias θ (prominence for quality) or a larger royalty premium $r_H - r_L$ makes high effort more attractive, raising α_H and thus increasing c^* and the equilibrium λ; conversely, if low-quality content is favored or the royalty gap is small, fewer creators invest in quality.

Finally, my paper solves the platform's profit maximization problem, where platform chooses the optimal royalty levels and the ranking bias. The platform's objective is subscription revenue (which is fixed per user) minus expected royalty payouts, so it faces a trade-off between user satisfaction (which comes from surfacing high-quality content) and royalty costs (which are lower for low-quality streams). Proposition 10 characterizes the structure of the optimal policy. First, it shows that the platform's profit is decreasing in the baseline royalty paid to low-quality content, implying the platform will set the lowest feasible base royalty, effectively $r_L^* = 0$.

Under the baseline model parameters, the platform's optimal solution hits a corner. Proposition 11 ("Only Superstar") reveals that the profit-maximizing strategy is to fully bias the first slot in favor of high-quality content ($\theta^* = 1$) while simultaneously setting no royalty premium (indeed, $r_L^* = 0$ and $\Delta^* = 0$, so $r_H^* = 0$ as well). In other words, the platform chooses to always show a high-quality item first to the user (ensuring maximal engagement utility) but pays nothing to content creators per stream. In this extreme equilibrium, the platform's costs are minimized — it can satisfy consumers by repeatedly surfacing one high-quality piece (provided at negligible cost by a creator, essentially a "superstar" with very low c) and avoid paying royalties on all other content. This result, while analytically coherent, is a extreme corner case highlighting the tension between profit and creative compensation. It also indicates that without additional constraints, a subscription platform might entirely skew toward a few high-quality winners and deprive most creators of revenue – an insight that motivates exploring regulatory or design constraints in the model's extensions.

1.3 Related Literature

Consumer search and ordered search The modern theory of search began with the random–sampling models of Diamond (1971) and Wolinsky (1986), where consumers sequentially draw sellers in an exogenously random order, incurring a cost s each time they inspect a price or quality realization. Diamond's paradox showed that even an infinitesimal s restores monopoly power; Wolinsky demonstrated that heterogeneity in valuations reinstates price dispersion and "true monopolistic competition." The next wave introduced information that allows consumers to direct their search. Weitzman (1978) characterized the optimal stopping rule when a buyer can rank alternatives by an index, but firms were passive price takers. Strategic pricing under directed search was first tackled by Anderson and Renault (1999), who embedded Chamberlinian product differentiation into a Bertrand–Diamond framework. Yet in all these models the order of inspection itself remained either random or costlessly chosen by the consumer.

An explicit role for exogenous order was provided by Arbatskaya (2007), who assumes all consumers visit firm 1 - first, firm 2 - second, and so on. She shows that prices and profits are strictly decreasing in rank, highlighting the value of "being first." Zhou (2011) added horizontal match heterogeneity and found that even with identical search costs the early–seller price can be lower while profits remain higher, as early visits secure larger market shares. Prominence as an endogenous strategic variable entered with Armstrong and Zhou (2011): when a single firm is exogenously designated as the first draw, that firm optimally chooses a lower price to reduce consumers' incentive to continue searching, while rivals post higher prices—a result that contrasts sharply with Arbatskaya's declining-price ladder. In a unifying contribution, Armstrong (2017) showed that if consumers coordinate on which seller is likely cheapest, an ordered-search equilibrium exists in which the first-inspected firm indeed sets the lowest price; the model nests both random and prominence cases via the information structure.

Despite this extensive, two features crucial for platform environments are absent. First, all cited papers treat product quality (or effort) as fixed; when Moraga-Gonzalez and Sun (2023) endogenises quality, they retain random search. Second, the search order is never an optimization variable for a two-sided intermediary that pays marginal royalties. My paper integrates these missing pieces: the platform chooses the ranking rule to economize on royalty outlays, while creators respond with endogenous effort, producing new comparative statics that contribute to canonical results in the ordered-search literature.

Platform ranking, bias and prominence Where the ordered–search tradition treats position as exogenously given, a parallel literature endogenises prominence by allowing intermediaries to steer or sell ranking. In early sponsored–search models, position auctions monetize consumers' sequential attention but assume no quality distortion: advertisers bid for slots while users learn only price or relevance ex post (Athey and Ellison, 2011). Armstrong and Zhou (2011) extend this insight to homogeneous goods, showing that firms may underbid for top placement when prominence depresses search continuation, tempering auction revenues. More general analyses of steering—intermediaries deliberately redirecting traffic—emerge in Hagiu and Jullien (2011), who identify a trade–off between click revenue and affiliation fees, and in De Corniere and Taylor (2019) who distinguish congruent from conflicting bias: rankings that favor high-utility sellers can raise welfare, whereas self–preferencing of inferior offers is purely distortionary.

Recent theoretical and empirical work deepens this perspective. Janssen et al. (2023) characterize equilibrium when a platform sells a sponsored slot while simultaneously obfuscating organic rankings to amplify its auction revenue; Choi and Jeon (2023) study design biases in ad–funded two-sided markets; and Heidhues et al. (2023) introduce behavioral mistakes, illustrating how biased ranking can exploit bounded rationality. Structural estimation papers quantify position effects and consumer surplus losses: Ursu (2018) randomize hotel rankings on Expedia, Lam (2021) recovers Amazon's self–preferencing parameters from browsing data, and Compiani et al. (2022) develop a double–logit method to back out demand under algorithmic ranking. Finally, Greminger (2022) shows that heterogeneity in position effects alters optimal discovery rules, challenging median–click assumptions common in earlier models.

Collectively, these studies illuminate why a platform might bias ranking but abstract either from marginal royalty costs (advertising models) or from endogenous content supply. My paper nests the steering logic of Hagiu and Jullien (2011) and the welfare taxonomy of De Corniere and Taylor (2019) inside a two–sided pricing environment where bias interacts with creators' effort incentives, thereby revealing novel quality and welfare consequences absent from the existing prominence literature.

Two-part tariffs and two-sided markets The canonical theory of platform pricing emphasizes the efficiency of two-part tariffs—an access fee and a per-transaction charge—in balancing cross-side externalities. Rochet and Tirole (2003) and Armstrong (2006) show that when one side generates greater marginal value for the other the platform will subsidize access and recoup surplus via a usage fee, a logic that rationalises payment-card interchange fees and dating-site subscription models. Subsequent work refines the menu of feasible tariffs:

Hagiu and Wright (2015) contrast marketplace (per-transaction commissions) with reseller (wholesale) modes, while Anderson and Bedre-Defolie (2021) analyze hybrid platforms that blend subscriptions with menu sales. Data-policy papers reveal further dimensions: sharing user data can act as an implicit cross-side subsidy (Bergemann and Bonatti, 2024; Kirpalani and Philippon, 2020). Yet this literature typically abstracts from sequential search frictions and from ranking as an allocative instrument; it therefore cannot address how a per-stream royalty feeds back into the order in which content is shown or into creators' effort incentives. By embedding a two-sided pricing inside an ordered-search environment, my model unifies these strands and uncovers a novel distortion—royalty-driven diversion—that classic two-sided pricing theory overlooks.

Quality investment under search frictions Search costs not only shape price competition; they also govern firms' incentives to supply quality. When consumers sample sellers randomly, lower search frictions increase the return to being high quality because good types are more likely to be discovered, yet they simultaneously intensify price competition, which may erode the quality premium. Fishman and Levy (2015) formalise this trade-off and show that quality provision is non-monotonic in search cost: improvements in search encourage investment when high quality is initially scarce but discourage it once the market is already rich in quality. Introducing ordered search complicates matters further. In a random-search environment with endogenous quality Moraga-Gonzalez and Sun (2023) derive a simple welfare test: the market over-invests in quality iff higher quality makes consumers inspect more products. With endogenous design rather than vertical quality, Bar-Isaac et al. (2012) predict a "superstar and long tail" equilibrium: easier search induces some firms to go broad and others to specialize narrowly, jointly expanding variety. Yang (2013) similarly shows that personalization technology reallocates effort toward niche products, amplifying the long-tail effect. None of these studies, however, allow an intermediary to manipulate search order in response to marginal royalty costs, so they cannot capture the feedback loop between platform bias and creators' effort. My model fills that gap: by endogenising both ranking and effort, I reveal parameter regions where royalty-driven diversion raises average quality—contrary to the monotone distortions found in random-search settings—and others where it precipitates a collapse in effort.

Empirical evidence on streaming royalties A growing empirical literature confirms that per-stream payouts vary widely across labels and that platforms adjust discovery algorithms in ways consistent with royalty-saving motives. Using a proprietary panel of Spotify contracts, Aguiar and Waldfogel (2018) document a median royalty gap of roughly 40 percent between major labels and independent distributors. Chen et al. (2016) show that Amazon's Buy-Box algorithm systematically favors offers that minimize its sourcing cost—even when retail prices are higher—signalling that marginal platform cost enters ranking logic. Event-study designs around Spotify's 2020 "Discovery Mode" pilot find that tracks which opt into a lower royalty rate receive a 28 percent increase in algorithmic streams, with no parallel lift in organic playlist placements. Qualitative reporting echoes these findings: a 2025 *Guardian* investigation describes Discovery Mode as "streaming payola" that exchanges lower payouts for preferential placement in auto-play and radio feeds Smith (2025). Complementarily, Zhu and Liu (2018) demonstrate that Amazon's entry with zero-royalty private-label products re-orders search results and displaces higher-cost third-party items, while Ursu (2018) quantify substantial click-through elasticity to rank position in hotel search, reinforcing the economic salience of prominence. Collectively, these studies validate two premises of my model: (i) royalty heterogeneity is large and persistent, and (ii) platforms exploit control over search order to economise on marginal payouts, thereby influencing the equilibrium distribution of content quality.

Several recent papers analyse algorithmic steering on content platforms, but in markedly different settings from mine. Bourreau and Gaudin (2022) study a bargaining game between two large content providers and a streaming platform that can threaten to rank the cheaper provider first in order to extract lower royalties. Their content quality is exogenous and consumer discovery is frictionless, so recommendation bias serves purely as a bargaining lever. By contrast, I consider a continuum of creators and endogenise effort: ranking bias in my model is chosen to balance marginal royalty outlays against search-generated consumer surplus, not to discipline a single powerful supplier. Closer in spirit, Qian and Jain (2024) show that exposing users to slightly mismatched content can raise average quality in an ad-funded environment; however, their platform earns per-impression advertising revenue, while mine faces a two-sided pricing schedule in which additional streams are a cost. Consequently, the direction of bias reverses: in their model the platform inflates exposure to induce effort, whereas in mine it may deflate exposure to limit royalty expenses. Finally, Teh and Wright (2022) analyze steering in a price-and-commission marketplace, but with homogeneous quality and no sequential search. My contribution is therefore to provide the first model that links (i) a subscription platform's two-sided pricing, (ii) endogenous creator effort, and (iii) consumer ordered search, and to show how these three elements jointly determine the optimal degree of recommendation bias and the resulting welfare implications.

2 Model

I consider a digital subscription platform that matches consumers with content supplied by heterogeneous content creators. There are three classes of agents: consumers, content creators, and a platform (the monopolist intermediary).

Consumers A continuum of consumers arrives looking for content that maximizes their utility. Each consumer seeks to find a piece of content that suits their tastes, and they face a sequential search problem due to the sheer volume of content available. Consumers encounter content one at a time, in a random order except for the first draw that may be influenced by the platform. After viewing each recommended item, a consumer decides whether to stop and consume that content or to skip it and continue searching for a better option.

Consumers derive utility from consuming a content item. Let the quality of a content item be denoted $Q \in \{L, H\}$ for low or high quality. High-quality content provides inherently greater utility on average than low-quality content. Specifically, if a consumer consumes a high-quality item, the utility realized is

$$u = \bar{u}_H + \epsilon,$$

where $\bar{u}_H$ is a baseline utility level for high quality and ϵ is an idiosyncratic match shock. For a low-quality item,

$$u = \bar{u}_L + \epsilon,$$

with $\bar{u}_L < \bar{u}_H$. The term ϵ represents the consumer's idiosyncratic taste match or enjoyment shock for that particular content. I assume ϵ is i.i.d. across content draws and follows a continuous distribution. For concreteness one may assume $\epsilon \sim \text{Uniform}[0, 1]$, so that utility is uniformly distributed on $[\bar{u}_Q; \bar{u}_Q + 1]$ for $Q \in \{L, H\}$. (The uniform assumption is adopted only for simplicity.) Consumers do not initially know ϵ for a piece of content until they sample it (e.g. by clicking and watching briefly), which is why search is necessary. Further in the analysis the specific values of utility will be assumed with $\bar{u}_L = 0$ and $\bar{u}_H = 1$.

Content Creators A unit mass of content creators produces content for the platform. Each creator is an individual that can create at most one content item (for example, an online video, article, or song). Creators differ in their cost of producing high-quality content. Let c_i denote creator i's private cost of exerting high effort (e.g. hiring better production, spending time editing) to make high-quality content. I assume c_i is drawn from a continuous distribution $F(c)$ on $[0, \bar{c}]$ (with $F(0) = 0$ and $F(\bar{c}) = 1$). This cost represents the incremental cost of creating a high-quality piece instead of a low-quality one. Producing low-quality content has a normalized baseline cost zero – essentially any creator can post a low-effort piece with minimal cost, but achieving high quality requires an additional cost c_i that varies across creators. For analytical convenience, one might consider $F(c)$ to be uniform on $[0, \bar{c}]$ in examples, but my results do not depend on a specific functional form for F beyond standard regularity (I only require that F is increasing and continuous, ensuring a well-defined cutoff for high-effort decisions).

Given the platform's policy, each creator chooses between two actions: High effort (H) or Low effort (L). This choice is strategic: creators anticipate how their content will fare on the platform – in particular, how likely it is to attract a consumer – and what royalty payment they will receive if it does. The platform's policy (r_H, r_L, θ) directly affects these payoffs. I now detail the payoff calculation for creators.

The platform commits to a royalty structure (r_H, r_L), where r_H is the payment to a creator for each consumption of a high-quality content piece, and r_L is the payment for each consumption of a low-quality piece. These can be thought of, for example, as per-view payments or a revenue-sharing scheme (if the platform earns a unit revenue per content view, r_H and r_L could represent the shares going to creators of each type). A key feature is that royalties are paid only when content is actually consumed by a user. Simply posting content does not guarantee any payment – the content must be matched to a consumer who ultimately chooses to consume it (i.e. the consumer's search stops at that content). If a piece of content is never consumed by any user, the creator receives nothing. This assumption captures the idea that creator earnings are contingent on capturing consumer attention which is navigated by the platform.

Let α_H denote the probability that a given high-quality content item gets consumed by a representative consumer, and α_L the probability for a low-quality item. These probabilities are endogenous outcomes of the matching process between consumers and content. They depend on the platform's ranking algorithm (bias θ) and on the overall composition of content on the platform. Given these probabilities, a creator's expected payoff from each strategy is as follows:

$$\Pi_i^H = \frac{\alpha_H \cdot r_H}{\lambda} - c_i, \qquad \text{(High Quality)}$$

$$\Pi_i^L = \frac{\alpha_L \cdot r_L}{1-\lambda}. \quad \text{(Low Quality)}$$

Search Searching for content is costly in a small way: after examining one piece of content, if the consumer is not satisfied, they incur a search cost $s > 0$ (e.g. a cognitive or time cost) to view the next recommended item. I assume consumers have a reservation utility from outside options normalized to 0. Consumers adopt an optimal stopping rule characterized by a reservation utility threshold z^*. Specifically, there exists a cutoff utility z^* such that the consumer will decide to stop and consume the current content if and only if the realized utility u of that content meets or exceeds z^*. If $u < z^*$, the consumer forgoes that content (deriving essentially no utility from it) and continues searching, paying the cost s to see another recommendation. This reservation strategy is optimal in the standard sequential search sense: intuitively, z^* is the utility level that makes the consumer indifferent between consuming the current item versus incurring the cost to search for another (Weitzman, 1978; Wolinsky, 1986).

I can formally characterize z^*. Let $G(u)$ denote the cumulative distribution function (CDF) of the utility u delivered by a random recommended content draw (this distribution is determined endogenously by the mix of high vs. low-quality content on the platform). The reservation utility z^* solves the indifference condition between stopping and searching one more time. In particular, u^* satisfies:

$$s = \int_{z^*}^{2} [u - z^*] dG(u),$$

Given this stopping rule, a consumer's search will continue until they encounter a content item with $u \geq z^*$. At that point, they stop and consume that item, deriving utility u from it. The total (net) utility the consumer obtains accounts for search costs incurred along the way. If the consumer viewed N items in total (stopping at the N-th), then they paid the search cost $s(N-1)$ times for the earlier items. Thus their net utility is $u - (N-1)s$.

Platform The platform is a monopolist intermediary that designs the rules of the marketplace to maximize its profit. Its decision variables are the high-quality royalty r_H, the low-quality royalty r_L, the subscription fee, and the ranking bias θ. These choices are made ex ante, and are committed to by the platform (I assume the platform can credibly commit to the announced royalty scheme and algorithm policy). The platform then takes the resulting equilibrium behavior of creators and consumers as given when evaluating profit. I anticipate that Platform will manipulate search by biasing the ranking towards low quality content in order to provide more incentives for artists to accept the lower royalty rate. Platform is fully aware of the effort level of the artists as it can be signaled by the acquaintance of the artist to a some type of a label, you can think of a high quality content can be only supplied under some major label assistance and low quality without it or with indie label.

Timing and Equilibrium The interaction unfolds in the following stages:

1. The platform chooses a royalty policy, subscription price and a ranking algorithm bias. Specifically, it sets payments (r_H, r_L) to reward content creators for high- and low-effort content, respectively, chooses the uniform subscription price P , and chooses a bias parameter θ that skews the content ranking in some deliberate way.

2. A continuum of content creators (of measure normalized to 1) simultaneously decide whether to produce high-effort or low-effort content. Creators differ in their cost of producing high quality, and this stage results in an equilibrium mix of content qualities.

3. Given the platform's policy and the resulting content mix, a continuum of consumers (measure 1) arrive at the platform. Each consumer searches the platform's content sequentially: the platform's algorithm recommends content one piece at a time (with first draw influenced by θ), and the consumer decides when to stop searching and consume a piece of content.

Payoffs are realized at the end of the game: creators receive royalties if their content is consumed, consumers derive utility from the content they consume (net of search costs), and he platform earns profit from subscription revenue minus the royalties it pays out.

Definition 1 (Subgame–perfect equilibrium)**.** *A* subgame–perfect equilibrium (SPE) *of the three–stage game is a collection*

$$\Big(r_H^\star,\, r_L^\star,\, \theta^\star,\, P^\star;\; c^\star,\, \lambda^\star,\, z^\star\Big) \in \mathbb{R}_+^2 \times [0,1] \times \mathbb{R}_+ \;\times\; [0,\bar{c}] \times (0,1) \times (0,2)$$

satisfying the following conditions:

(i) *Platform optimality. Given the creators' and consumers' best–response mappings described in (E2)–(E3), the policy vector* $\big(r_H^\star, r_L^\star, \theta^\star, P^\star\big)$ *maximises the platform's expected profit,*

$$\big(r_H^\star, r_L^\star, \theta^\star, P^\star\big) \in \arg \max_{r_H \ge r_L \ge 0,\; 0 \le \theta \le 1,\; P \ge 0} \pi\big(r_H, r_L, \theta, P\big),$$

where the profit function incorporates the equilibrium outcomes specified below.

(ii) *Creators' optimal effort. There exists a cut-off cost* $c^\star \in [0, \bar{c}]$ *such that each creator with cost draw* c *chooses the effort level*

$$e(c) = \begin{cases} H, & c \le c^\star, \\ L, & c > c^\star, \end{cases} \qquad \textit{and} \qquad c^\star = r_H^\star\, \pi_H^\star - r_L^\star\, \pi_L^\star,$$

where $\pi_Q^\star = \alpha_Q^\star / \lambda_Q^\star$ *is the (ex-ante) probability that a single* $Q \in \{H, L\}$*-item is eventually consumed. The implied catalog share of high-quality items satisfies*

$$\lambda^\star = F\big(c^\star\big).$$

(iii) *Consumer search optimality. Given the catalog share* $\lambda^\star$ *and search cost* s*, consumers adopt a reservation-utility rule with cut-off* $z^\star$ *defined implicitly by* $B\big(z^\star; \lambda^\star\big) = s$*, where*

$B(z;\lambda) = \mathbb{E}[(u-z)_+]$ is the expected gross benefit of one further draw. The induced expected utility equals $U\big(\lambda^\star, \theta^\star\big)$.

(iv) *Participation constraint. The subscription price exactly extracts the representative consumer's surplus:*

$$P^\star = U\big(\lambda^\star, \theta^\star\big).$$

Discussion In spirit, the game I have just described integrates three previously separate modelling traditions. My sequential–search block inherits the Weitzman–Armstrong apparatus in which consumers encounter items one at a time and decide whether to continue at a per–draw cost s. Relative to the canonical ordered–search models of Arbatskaya (2007) and Armstrong (2017), two extensions are key. First, the platform—not the consumer—chooses the initial prominence parameter θ, endogenising "who is seen first." Second, the objects of search differ vertically: their quality levels are determined by creators' endogenous effort rather than drawn from an exogenous distribution. This vertical dimension forces the platform to weigh early–stopping benefits against the royalty savings from steering towards low–quality items.

Building on the two–sided-market logic of Rochet and Tirole (2003), the platform here sets a subscription fee and a per–stream royalty. Yet unlike the standard payment-card or marketplace settings, the per-stream component is a pure cost: every additional consumption event directly debits the platform. Consequently, the ranking decision θ becomes a shadow instrument that substitutes for a usage fee, a diversion motive absent from classic two–sided models.

Endogenous quality under search frictions. By allowing creators with heterogeneous cost draws to select a high- or low-effort regime, I generate a feedback loop between ranking policy and the equilibrium catalog share λ. This loop is missing from random-search models with effort (Moraga-González and Sun, 2023) and from bias models with fixed quality (De Corniere and Taylor, 2019; Hagiu and Jullien, 2011). My quadratic reservation equation

(Proposition 1) preserves tractability, delivering an explicit cut-off cost $c^*(\theta, \lambda, s)$ and a unique fixed point $\lambda^\star = F(c^*)$ in closed form when F is uniform.

Taken together, these features let us ask questions that cannot be posed in any one of the predecessor frameworks: How does a royalty gap alter the optimal prominence rule? When does prominence dampen vs. stimulate equilibria effort? And under which parameter range is royalty-driven steering privately profitable yet socially harmful?

3 Analysis

3.1 Consumers Search Analysis

Reservation Utility The optimal stopping rule for the consumer is characterized by a reservation utility $z^*(s, \lambda)$, meaning the consumer accepts an item if and only if its realized total utility $u = q + \varepsilon$ exceeds z^*. This threshold is derived by setting the expected marginal benefit of continued search equal to the marginal search cost s (Weitzman, 1978; Wolinsky, 1986). Formally, z^* solves the indifference condition:

$$\int_{z^*}^{2} (u - z^*) dG(u) = s, \tag{1}$$

where $G(u)$ is the distribution of an item's total utility (with support $[0, 2]$). The left side is the expected utility gain from searching for an offer above z^*. Given the model primitives, $G(u)$ is a mixture: with probability λ the item is high-quality ($q_H = 1$ and $u = 1 + \varepsilon \in [1, 2]$) and with probability $1 - \lambda$ it is low-quality ($q_L = 0$ and $u = \varepsilon \in [0, 1]$). Solving the indifference condition yields z^* as the unique positive root of a quadratic equation. In particular, there is a critical search cost $s = \frac{\lambda}{2}$ at which $z^* = 1$.

Proposition 1 (Closed–form reservation utility). *For every search cost $s \in (0, 1)$ and catalog share $\lambda \in (0, 1)$ the indifference equation (1) admits a unique solution $z^* = z^*(s, \lambda) \in (0, 2)$.*

It is given piece-wise by

$$z^*(s,\lambda) = \begin{cases} 2 - \sqrt{\dfrac{2s}{\lambda}}, & \textit{if } s < \dfrac{\lambda}{2} \quad (\textit{cheap search, } z^* > 1), \\ \dfrac{1 - \sqrt{2\lambda^2 - \lambda + 2s - 2\lambda s}}{1-\lambda}, & \textit{if } s > \dfrac{\lambda}{2} \quad (\textit{costly search, } z^* < 1). \end{cases}$$

At the knife-edge $s = \lambda/2$ *the two expressions coincide and yield* $z^* = 1$*, so* z^* *is continuous in* (s, λ)*.*

Proof. See Appendix. □

Note that there is one more restriction for the analytical tractability of the proposition above it is that with a very low share of high-quality content there shouldn't be very big search costs as then we will have a negative reservation utility, specifically the restriction looks like this $\frac{2s^2-s}{2s-2} \leq \lambda$. This restriction isn't give us any problems with a meaning of this reservation utility as the negative ones is can be just equated to zero, however we need to not that in the region where this restriction is not indeed satisfied, we can't use the formulas below.

One of the important properties which is satisfied in my model is that reservation utility decreases with a search costs, so in addition to the model of Moraga-Gonzalez and Sun (2023) consumers become "pickier" as search costs drops. Conversely, if search is nearly free ($s \to 0$), z^* approaches the maximum utility 2 – the consumer insists on the best possible content. Throughout, note that z^* does not explicitly depend on θ (the platform's first-item curation parameter); θ influences the initial draw but not the long-run reservation strategy, which is determined by the i.i.d. distribution of items beyond the first. (The first-item bias will, however, affect the probability of accepting that first recommendation, as I derive next.) We can explicitly see with this proposition that nor bias neither the model itself makes sense with a very high low costs, when the consumer only allows the high quality items, this restriction is also don't change the essence of the analysis, if not our simplified

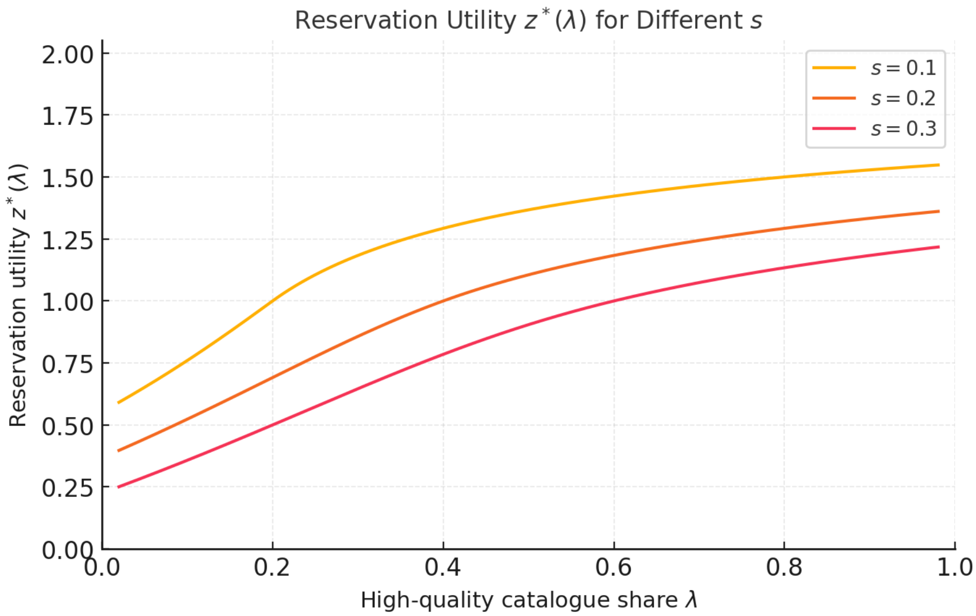

Figure 1: Reservation Utility $z^*(\lambda)$ For Different s

assumption about match utility being distributed uniformly from 0 to 1, we would have different results even in a high search costs zone, so we can now just focus on the search costs high enough that both effort items have a potential to be consumed. Proposition below outlines the specific properties of the reservation utility in our analysis taken the share of high-quality items as a parameter.

Corollary 2 (Reservation Utility Comparative statics)**.**

$$\frac{\partial z^*}{\partial s} < 0, \qquad \frac{\partial z^*}{\partial \lambda} > 0, \qquad \frac{\partial z^*}{\partial \theta} = 0.$$

Higher search cost lowers selectiveness ($z^* \downarrow$); a richer high-quality catalog raises it. The first-draw bias θ affects the probability of stopping on the first item but leaves the long-run reservation level unchanged.

Proof. See Appendix. □

Comparative statics for the levels of reservation utility with different search costs and different high-quality catalog share is demonstrated on figure 1.

First-Item Acceptance Probability and Expected Samples Given the reservation rule in Proposition 1, two statistics are needed later for creators' and platform pay–offs: the probability that the first recommended item is accepted, and; the expected number of items a consumer examines before stopping.

Throughout denote the reservation cut-off simply by z^* and recall that the first recommendation is high quality with probability θ (first–draw bias) and low quality with probability $1-\theta$.

Lemma 3 (Probability of accepting the first recommendation)**.** *Let* $P^{(1)}(\theta, \lambda, s)$ *be the probability that a consumer accepts the very first item she sees. Then*

$$P^{(1)}(\theta, \lambda, s) = \begin{cases} \theta \ + \ (1-\theta)\,(1-z^*), & z^* \leq 1, \\[2ex] \theta\,(2-z^*), & z^* > 1, \end{cases}$$

where $z^* = z^*(s, \lambda)$ *is given in Proposition 1. In particular* $P^{(1)}(\theta, \lambda, s) = \theta$ *at the knife-edge* $s = \lambda/2$.

Proof. See Appendix. □

Let N denote the total draws before stopping. Conditional on rejecting the first item, the continuation process is i.i.d. with success probability

$$p = \begin{cases} \lambda \ + \ (1-\lambda)\,(1-z^*), & z^* \leq 1, \\[2ex] \lambda\,(2-z^*), & z^* > 1, \end{cases}$$

obtained by replacing the first–draw bias θ with the catalog share λ in Lemma 3. Because rejections and acceptances form a geometric sequence,

$$\mathbb{E}[N] \ = \ 1 \ + \ \frac{1-P^{(1)}(\theta, \lambda, s)}{p}.$$

One checks that expected number of sampled items goes to one as search cost increases and to the infinity as it decreases, mirroring the consumer's increasing willingness to search when search becomes cheaper.

Lemma 4 (Expected session length)**.** *Let N denote the total number of recommendations a consumer examines before stopping. With $P^{(1)}(\theta, \lambda, s)$ from Lemma 3 and reservation cut-off z^* from Proposition 1,*

$$\mathbb{E}[N](\theta, \lambda, s) = 1 + \frac{1 - P^{(1)}(\theta, \lambda, s)}{p(\lambda, s)}, \qquad p(\lambda, s) = \begin{cases} \lambda + (1-\lambda)(1-z^*), & z^* \leq 1, \\ \lambda(2 - z^*), & z^* > 1. \end{cases}$$

Equivalently,

$$\mathbb{E}[N] = \begin{cases} 1 + \dfrac{(1-\theta)\, z^*}{1-(1-\lambda)z^*}, & z^* \leq 1, \\ 1 + \dfrac{1-\theta(2-z^*)}{\lambda(2-z^*)}, & z^* > 1. \end{cases}$$

Proof. See Appendix. □

Intuitively, the expressions above say that the average "scroll depth" responds in the expected way to the search cost. When it is expensive to inspect another title, the user grabs the first acceptable item almost immediately, so the expected number of recommendations viewed converges to one. When searching is virtually free the reservation threshold becomes very demanding; low-quality or mediocre draws are almost always rejected and the user is willing to keep scrolling for as long as it takes. In the limit of zero search cost the geometric success probability collapses to zero and the expected session length diverges, while any strictly positive cost keeps it finite.

3.2 Creators Choice Analysis

Let $\alpha_H(\theta, \lambda, s)$ (resp. α_L) denote the ex–post probability that the stream eventually consumed by the representative user is high– (resp. low–) quality. These probabilities follow mechanically from the reservation rule and the first–draw bias.

Proposition 5 (Accepted–item quality shares)**.** *With reservation cut-off* $z^* = z^*(s, \lambda)$ *from Proposition 1,*

$$(\alpha_H, \alpha_L) = \begin{cases} \Big(\dfrac{\theta + (\lambda - \theta)z^*}{1 - (1-\lambda)z^*},\ \dfrac{(1-\theta)(1-z^*)}{1-(1-\lambda)z^*}\Big), & z^* \leq 1, \\ \\ (1, 0), & z^* > 1. \end{cases}$$

In all cases $\alpha_H + \alpha_L = 1$ *and* $\alpha_H > \lambda$ *whenever* $z^* > 0$.

Proof. See Appendix. □

When the threshold is moderate ($z^* \leq 1$) the stream can be high quality for two distinct reasons: either the very first recommendation is of type H (probability θ) or an H item is the first to pass the threshold after some early rejections (probability $(\lambda - \theta)z^*$). Dividing by the overall stopping probability $1 - (1 - \lambda)z^*$ yields the first line. Once search becomes stringent ($z^* > 1$) low–quality items can never clear the cut-off, so every accepted stream is high quality. Hence, consumer search endogenously amplifies quality: the share of high-quality streams α_H always exceeds the catalog share λ, and the gap widens with a tighter threshold or stronger first-draw bias.

Each creator chooses between low effort (L) and high effort (H). Producing H yields quality $q_H = 1$ and entails a private cost $c \geq 0$; producing L yields $q_L = 0$ at zero cost. Streams are monetised via a two–part royalty (r_H, r_L) with $r_H > r_L \geq 0$.

Because items of the same type are ex–ante symmetric, a single high–quality piece is streamed with probability $\pi_H \equiv \alpha_H/\lambda$ and a single low–quality piece with $\pi_L \equiv \alpha_L/(1 - \lambda)$, where (α_H, α_L) are given in Proposition 5.

Proposition 6 (Effort cutoff and equilibrium share)**.**

(i) *Effort cutoff. A creator exerts high effort iff* $c \le c^*$, *where*

$$c^*(\theta, \lambda, s) \;=\; r_H \,\frac{\alpha_H(\theta, \lambda, s)}{\lambda} \;-\; r_L \,\frac{\alpha_L(\theta, \lambda, s)}{1-\lambda}.$$

(ii) *Equilibrium quality. Let* F *be the cumulative distribution of cost draws. In a symmetric equilibrium the catalog share of high–quality content is*

$$\lambda \;=\; F\big(c^*(\theta, \lambda, s)\big).$$

(iii) *Closed form under a uniform cost distribution. If* $c \sim U[0, \bar{c}]$ *then*

$$\lambda^{\mathrm{U}} \;=\; \frac{c^*(\theta, \lambda^{\mathrm{U}}, s)}{\bar{c}} \;=\; \frac{r_H\, \alpha_H(\theta, \lambda^{\mathrm{U}}, s)}{\bar{c}\, \lambda^{\mathrm{U}}} - \frac{r_L\, \alpha_L(\theta, \lambda^{\mathrm{U}}, s)}{\bar{c}\,(1-\lambda^{\mathrm{U}})},$$

Proof. See Appendix. □

The cutoff c^* equates the expected royalty premium from high effort to its incremental cost. A higher streaming advantage for quality—either through algorithmic bias ($\theta \uparrow$) or a larger royalty gap ($r_H - r_L$)—raises c^*, inducing more creators to upgrade and thereby increasing the equilibrium share λ. Conversely, when search is costly or bias is weak, the quality premium is small and only the lowest–cost creators supply high–effort content.

To obtain closed–form comparative-statics I specialise to a uniform cost distribution $c \sim U[0, 1]$ and keep the benchmark utility primitives. The equilibrium condition $\lambda = F(c^*) = c^*$ therefore reads

$$\lambda \;=\; r_H \,\frac{\alpha_H(\theta, \lambda, s)}{\lambda} \;-\; r_L \,\frac{1-\alpha_H(\theta, \lambda, s)}{1-\lambda}, \tag{2}$$

Because α_H is piece-wise (Proposition 5), the fixed-point (2) can be solved analytically in the two search regimes. I collect the result in the next proposition, which is the explicit counterpart to Proposition 6.

Proposition 7 (Closed form with $c \sim U(0,1)$)**.** *Assume $r_H \in (0,1]$, $r_L \in [0, r_H)$ and a uniform cost distribution on $[0,1]$.*

(i) *High-threshold regime ($s < \lambda/2$). Here $z^* > 1$ and $\alpha_H = 1$. The equilibrium share of high-quality content solves $\lambda = r_H/\lambda$, yielding the closed form*

$$\lambda^{\mathrm{HT}} = \sqrt{r_H} \qquad \big(s < \tfrac{1}{2}\lambda^{\mathrm{HT}}\big).$$

Thus, when search is sufficiently cheap every accepted stream is high quality and the catalog share adjusts to the square root of the royalty.

(ii) *Low-threshold regime ($s > \lambda/2$). Now $z^* \le 1$ and $\alpha_H(\theta, \lambda, s) = \dfrac{\theta + (\lambda - \theta)z^*}{1 - (1-\lambda)z^*}$ with z^* from Proposition 1. Substituting into (2) gives the rational equation*

$$\lambda^{LT} = \frac{r_H}{\lambda^{LT}} \frac{\theta + (\lambda^{LT} - \theta)z^*}{1 - (1-\lambda^{LT})z^*} - \frac{r_L}{1-\lambda^{LT}} \frac{(1-\theta)(1-z^*)}{1-(1-\lambda^{LT})z^*}. \tag{3}$$

Whose unique root yields the equilibrium share $\lambda^{\mathrm{LT}} = \lambda^{\mathrm{LT}}(\theta, r_H, r_L, s)$.

Proof. See Appendix. □

When search is cheap the consumer rejects all low quality, forcing $\alpha_H = 1$. Once search becomes costly, some low-quality streams are tolerated and the equilibrium share must be solved jointly with the endogenous threshold; the resulting quartic highlights the non-linear feedback between consumer selectiveness, algorithmic bias θ and the royalty gap $(r_H - r_L)$.

Comparative statics in the costly–search regime Henceforth I focus on the relevant low–threshold region $s > \lambda/2$ in which some low–quality content is tolerated and the equilibrium share λ^{LT} is given by the unique root of (3). The next corollary summarizes the main monotone effects of the platform's primitives.

For the comparative-static derivatives below we require the slope[2] of the fixed–point map

[2]See Appendix 5: $K(\theta, \lambda, s) := \dfrac{\partial \alpha_H(\theta, \lambda, s)/\partial\lambda}{\lambda(1-\lambda)}$.

to remain below one. Throughout this subsection we therefore impose the mild restriction

$$0 < r_H - r_L \ < \ K^{-1}\big(\theta, \lambda^{LT}, s\big), \tag{4}$$

Under (4) the denominator Φ_λ in the implicit–function formulas is strictly positive, ensuring a unique and well-behaved equilibrium. K here measures how strongly the demand advantage of being high-quality responds to a marginal increase in the catalog share. Multiplying by the royalty gap gives the feedback from the creators' supply decisions back into their incentives. If this multiplication would be more than one it would be explosive, however this assumption is not restrictive as it always satisfied in all realistic cases.

Corollary 8 (Qualitative comparative statics)**.** *Assume $s > \lambda^{LT}/2$ and $c \sim U[0,1]$.*

(i) *Algorithmic bias.* $\dfrac{\partial \lambda^{LT}}{\partial \theta} > 0$. *A higher first-draw bias towards high quality increases the fraction of creators choosing high effort.*

(ii) *Royalty gap.* $\dfrac{\partial \lambda^{LT}}{\partial r_H} > 0$ *and* $\dfrac{\partial \lambda^{LT}}{\partial r_L} < 0$. *Raising the reward to high quality or lowering the reward to low quality both raise the equilibrium share of H content.*

(iii) *Search cost.* $\dfrac{\partial \lambda^{LT}}{\partial s} < 0$. *Making search more onerous (thereby lowering the reservation cut-off z^*) depresses demand for quality and reduces creators' incentive to invest.*

(iv) *Interaction effect. The marginal impact of bias is stronger when the royalty gap is large:* $\dfrac{\partial^2 \lambda^{LT}}{\partial \theta \, \partial (r_H - r_L)} > 0$.

Proof. See Appendix. □

Bias and a larger royalty wedge raise the expected return to high effort relative to its cost, pushing the cutoff c^* outward and drawing more creators into the H regime. Higher search costs have the opposite effect: a lower reservation utility means consumers are willing to settle for mediocre matches, so the streaming advantage of H content shrinks and quality supply falls. Finally, bias and the royalty gap are complements: when the platform already

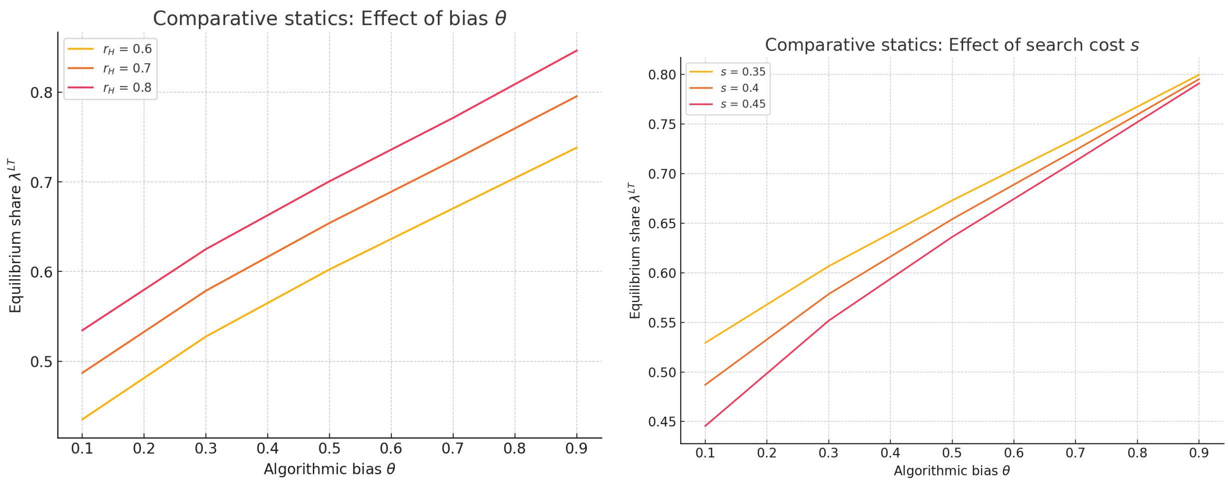

Figure 2: Comparative statics on λ^{LT}

pays a large premium for quality, directing early attention to H content delivers especially high incremental value on both sides of the market, reinforcing creators' incentives. Different comparative static with resepct to bias and other variables is demonstrated on figure 2 and figure 3

3.3 Platform–Design Problem

Fix an arbitrary policy triple (r_H, r_L, θ). Let $\lambda^{eq}(r_H, r_L, \theta) \in (0,1)$ be the (unique) catalog share of high–quality items supplied by creators in equilibrium (Section 3.2), and denote the associated reservation utility by $z^* = z^*\big(s, \lambda^{eq}\big)$ (Proposition 1). The first–draw acceptance probability $P^{(1)}(\theta, \lambda^{eq}, s)$ and the geometric continuation–success probability $p(\lambda^{eq}, s)$ are given in Lemmas 3–4. Conditional on stopping at the first recommendation, the consumer enjoys a match value $\mathbb{E}[u \mid \text{stop at 1st}] = \theta\big[1 + \frac{1}{2}(2 - z^*)\mathbf{1}_{\{z^*>1\}}\big] + (1-\theta)\frac{1}{2}(1-z^*)^2\mathbf{1}_{\{z^*<1\}}$. If the first item is rejected, search continues until the geometric stopping time N; using Lemma 4 the continuation part creates expected value $\big[\lambda^{eq} + (1-\lambda^{eq})(1-z^*)\big]/p$ and incurs expected cost $(E[N]-1)s$. Collecting terms gives the following closed–form expression.

Proposition 9 (Ex-ante consumer's utility)**.** *The ex-ante utility of the consumers who de-*

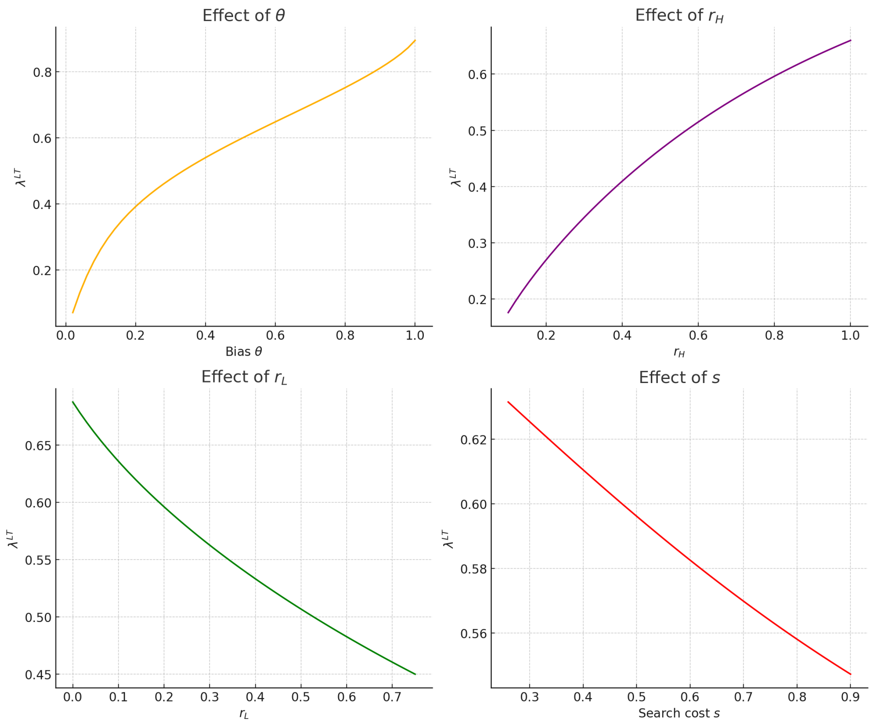

Figure 3: Comparative statics for λ^{LT}

cides to subscribe to the platform is given by the following formula

$$U(\lambda, \theta) = \theta\big[1 + \tfrac{1}{2}(2 - z^*)\mathbf{1}_{\{z^*>1\}}\big] \; + \; \frac{(1-\theta)(1-z^*)^2}{2\big[1-(1-\lambda)z^*\big]} \; - \; s\big[E[N] - 1\big]. \tag{5}$$

Proof. See Appendix. □

Since the platform is a monopolist and consumers are ex-ante identical, it charges the highest subscription fee that leaves the representative user indifferent between joining and staying out:

$$P(r_H, r_L, \theta) \;=\; U\big(\lambda^{eq}(r_H, r_L, \theta), \theta\big). \tag{6}$$

Each consumed stream triggers a royalty: r_H with probability $\alpha_H(\lambda^{eq}, \theta)$ and r_L with probability $1 - \alpha_H$. The expected outlay per subscriber is therefore

$$R(r_H, r_L, \theta) = r_H \, \alpha_H(\lambda^{eq}, \theta) + r_L\big[1 - \alpha_H(\lambda^{eq}, \theta)\big].$$

Definition 2 (Profit per subscriber)**.**

$$\pi(r_H, r_L, \theta) = U\big(\lambda^{eq}(r_H, r_L, \theta), \theta\big) - R(r_H, r_L, \theta). \tag{7}$$

Because λ^{eq} is implicitly characterised by the quartic polynomial of Proposition 7 and all other ingredients in (5) are explicit, the map $(r_H, r_L, \theta) \mapsto \pi$ is well defined and continuously differentiable on the compact choice set $\{0 \le r_L < r_H \le \bar{r},\ 0 \le \theta \le 1\}$. With some arbitrary large $\bar{r}$ which is not affect the optimization. Platform–design problem is as follows:

$$\max_{0 \le r_L < r_H \le \bar{r},\ 0 \le \theta \le 1} \ \pi(r_H, r_L, \theta). \tag{8}$$

To lighten notation and to isolate the incentive–relevant royalty gap from the purely budgetary level, we henceforth write the two royalties as

$$r_L = r, \qquad r_H = r + \Delta, \qquad 0 \le r < \bar{r},\ \Delta \in [0, \bar{r} - r].$$

The parameter r is the baseline per-stream payment, while Δ is the gap that rewards high-quality streams. To continue with the profit analysis let's once again take a look on the Utility of the consumer which plays a role of the participation fee, for the interested for us region of costly search it takes the following form:

$$U(\lambda, \theta; s) = \theta + \frac{(1-\theta)(1-z^*)^2}{2[1-(1-\lambda)z^*]} - s\frac{(1-\theta)z^*}{1-(1-\lambda)z^*}$$

Already from this expression you can see the feature of the model in the current formulation with that theta has substantial impact on the consumers utility, for sure, it is increases the probability of the consumer to consume the first item in the carousel being a high quality with a big probability, and subsequently it lowers the average number of streams. So the amplification of the consumer's utility is doubled. Following proposition will characterize the optimal solution of the platform when it can choose endogenously all variables, so that, when all market power is concentrated in one hands.

Proposition 10 (Structure of the platform's optimal policy.)**.** *Let $\pi(r, \Delta, \theta)$ denote profit per subscriber as defined above and assume the costly-search region $s > \lambda/2$ together with the bound (4). Then:*

(i) *Irrelevance of the baseline level. π is strictly decreasing in the baseline royalty r. Consequently, the platform always sets the lowest feasible baseline:*

$$r^* = 0$$

(ii) *Remaining decision problem. Let $\widetilde{\pi}(\Delta, \theta) := \pi(0, \Delta, \theta)$ and write $\lambda^{LT}(\Delta, \theta)$ for the unique solution of the equation (3) when $r_L = 0, \quad r_H = \Delta$. The platform's problem reduces to*

$$\max_{0 \leq \theta \leq 1, 0 \leq \Delta < \bar{r}} \widetilde{\pi}(\Delta, \theta) \quad s.t. \quad (A.1) : \Delta < K^{-1}(\theta, \lambda^{LT}(\Delta, \theta), s).$$

The objective is continuously differentiable on the compact, $\theta - \Delta$ feasible set $\mathcal{F} := \{(\theta, \Delta) | 0 \leq \theta \leq 1, 0 \leq \Delta < K^{-1}(\theta, \lambda^{LT}, s)\}$; hence an optimum exists.

(iii) *First-order characterization. Any interior optimum $(\theta^*, \Delta^*) \in \mathcal{F}$ satisfies*

$$\frac{\partial \widetilde{\pi}}{\partial \theta} = 0, \quad \frac{\partial \widetilde{\pi}}{\partial \Delta} = 0,$$

where

$$\frac{\partial \widetilde{\pi}}{\partial \theta} = (\alpha_H - \lambda^{LT}) + [1 - \Delta K] \frac{\partial \lambda^{LT}}{\partial \theta},$$

$$\frac{\partial \widetilde{\pi}}{\partial \Delta} = \alpha_H + [1 - \Delta K] \frac{\partial \lambda^{LT}}{\partial \Delta}, \qquad K := K(\theta^*, \lambda^{LT}, s).$$

Boundary optima occur when either $\theta^* = 0, 1$ or the gap saturates the feasibility fron-

tier, $\Delta^* = K^{-1}(\theta^*, \lambda^{LT}, s)$.

(iv) *Qualitative implications.*

(a) Gap versus prominence: complements. Whenever the interior conditions hold, the mixed partial $\partial^2\widetilde{\pi}/\partial\theta\partial\Delta$ *is positive, so a larger quality gap increases the marginal return to bias (and vice versa).*

(b) Never a premium without a bias. If $\theta^* = 0$ *then* $\partial\widetilde{\pi}/\partial\Delta < 0$ *everywhere, implying* $\Delta^* = 0$*. The platform never pays a premium for quality unless it also pushes quality to the top of the ranking.*

Proof. See Appendix. □

The problem that arises here is that in fact no interior optimum is can not be founded under the current model formulation, even in a costly search region, where the consumers have access theoretically to both high-quality and a low-quality content - the number of levers platform has give it too much market power. Specifically platform can always behave as at least some very small amount of the artists will anyway choose the high-quality, it could be the mass which costs of producing a content is too small or just some artists with a different motivations, so even if the only one artist provided the high-quality content on the platform for any royalties, platform automatically has an optimal solution. Platform can only just show this high-quality content to the consumer with probability 1. It will end the consumers search as in the reservation utility zone from 0 to 1, consumer accepts any high-quality item and will give him the expected ex-ante utility of $\frac{3}{2}$ - which is maximum possible. On the other hand platform can just impose the royalty gap $\Delta^* = 0$ - so that, it will never pay anything to the artists and the royalty outflow will be also equal to 0.

Proposition 11 (Only Superstar)**.** *For any search costs* $s > 0$*. The optimal values of the decision variables of the platform will be equal to*

$$r^* = 0, \quad \Delta^* = 0, \quad \theta^* = 1.$$

So that platform will always fully steer the first draft to the high-quality content and will never pay anything to anyone.

Proof. See Appendix. □

The result is troubling, yet it also makes perfect sense. First, it is essentially insensitive to the minimum royalty or royalty–gap levels: whatever corner solution is feasible, the platform selects it. Increasing the baseline royalty can overturn the corner outcome, but I postpone that exploration to later sections. Royalties equal to zero may appear unrealistic, but one can think of them as arbitrarily small positive numbers—the qualitative result is unchanged. Another consideration is our inability to balance the participation fee against the royalty outflow; however, adding a wedge to offset this discrepancy does not alter the conclusion.

3.4 Discussion

The finding is best understood through the construction of consumer utility, which combines (i) intrinsic quality and (ii) an idiosyncratic match draw. In the baseline calibration, the maximum match utility exactly equals the quality gap between high- and low-quality content. Hence, in the costly-search regime, a consumer who is willing to consider any low-quality item must set a reservation utility that automatically accepts every high-quality item. This strict dominance of high-quality content is the driving force behind the result and echoes the debate on the "superstar" effect (Bar-Isaac et al., 2012): *When quality utility dwarfs idiosyncratic taste, the platform creates a superstar effect—promoting a tiny mass of highly talented artists willing to accept the lowest compensation.*

When quality becomes overwhelmingly valuable, the platform's leverage over the few high-quality creators is immense. It needs only a handful who will accept minimal royalties (because their production costs are negligible) and can then place them permanently at the top of the list. Although details differ in practice, the basic intuition persists: the platform channels users toward an exceedingly narrow set of artists while extracting substantial surplus. Thus, with unrestricted control, the platform exploits both sides of the market—appropriating all consumer surplus and leaving nothing for most creators. This distortion survives even if a rival platform enters, because the competitor would adopt the identical strategy on the creator side. The model therefore delivers a novel mechanism for the superstar effect, suppressing average catalog quality (which falls to zero) and thereby reducing cultural variety and social welfare.

In the current formulation, the model appears to bias toward high-quality content, contradicting the original intuition that steering toward low-royalty content is cheaper. Yet the platform simultaneously sets the royalty gap endogenously, so in effect it does steer toward the lowest royalty. Identifying the sign of bias is complicated because changing a single parameter within a regime simultaneously shifts: (i) the probability distribution of consumed content, (ii) the share of high-quality creators, (iii) the platform's royalty outlay, and (iv) the optimal subscription price. Moreover, (i) and (ii) reinforce each other. A convenient bias metric is therefore

$$B(\theta, \lambda) = \frac{\theta - \lambda}{\lambda},$$

which explodes as $(\lambda^*, \theta^*) \to (0, 1)$ but will be useful below.

To escape this corner solution, I propose several modifications. Section 4 will analyze a specification in which the supports of low- and high-quality utilities overlap—achieved by increasing the match-utility range of low-quality items. That creates a reservation-utility interval where (i) not every high-quality item is accepted and (ii) not every low-quality item is rejected. First, the low-quality option then delivers higher potential satisfaction; second, the first-slot bias no longer single-handedly determines consumption, because even a top-ranked item may still be rejected.

Bias Penalty Another way to deal with this problem is to introduce the reputation costs or algorithmic complexity cost which will account for the fact that heavily skewed ranking may be harder to justify to users / regulators or to be more costly to maintain with other cause. To introduce it we need to add extra term to the platform's profit optimization.

$$\pi(r_H, r_L, \theta) = U\big(\lambda^{eq}(r_H, r_L, \theta), \theta\big) - R(r_H, r_L, \theta) - \beta(\theta - \lambda)^2.$$

Now, how does the solution change with the introduction of this term? The platform is still incentivized to steer exclusively towards high-quality content, however the penalty it experience is gradually increase with the amount of steering and with a parameter β - multiplicator before the penalty term. This parameter can be interpreted as a reputation of the platform or the awareness of the consumers of the bias, so the factor of how effectively platform can hide it. With a small penalty parameter β the platform choice tends toward the same solution, however, for any $\beta > 0$ but sufficiently low, it would want to increase a royalty gap Δ in order to increase the number of content creators who have chosen the high-quality. So in the range of small β the bias is still $\theta^* = 1$ and royalty gap is increases to incentivize quality $\Delta^* \uparrow$, as a consequence $\lambda^* \uparrow$ and the average quality of the content increases. On the next phase of the β range in a intermediate penalty zone it suddenly jumps to the $\Delta = 0$ and $\theta = \hat{\theta} > 0$. So that for some big enough parameter of the bias penalty the solution goes to another corner of this "box", now it become more profitable to use the strategy similar to the previous with $\Delta = 0$ but taking θ^* as high as it mostly profitable, balancing the utility from the pure profit and royalty outflow which is maximized in this scheme and disutility from the penalty bias. So after some $\hat{\beta}$ for all $\beta > \hat{\beta}$ the platform optimization variables is taking the following values $\Delta^* = 0$ and $\theta^* \downarrow$ with equilibrium quality on the platform equal to zero $\lambda^* = 0$.

Proposition 12 (Penalty–induced collapse of the royalty gap). *1. For all $\beta < \hat{\beta}$ the optimal platform's policy satisfy the following propeties:*

$$\frac{\partial \Delta^*}{\partial \beta} > 0, \qquad \theta^* = 1,$$

2. *For all* $\beta \geq \hat{\beta}$*:*

$$\Delta^* = 0, \qquad \frac{\partial \theta^*}{\partial \beta} < 0.$$

Proof. See Appendix. □

This result provides sharp intuition that while small penalty for bias is generally good and welfare improving in the sense that it raises average quality of the content on the platform and increases the royalty payments to the content creators, higher bias penalty is absolutely destructive, harming user satisfaction, decreases platform profit and changes nothing for the artists, as even through now low-quality artists gets some share of the listening, they are still receive zero royalty in the equilibrium. In the next paragraph we will inspect further different cases for the welfare improvements taking into account that this bias penalty shown to be very effective for the considerable range of parameters.

Welfare This model is very rich in the ways you can introduce the welfare analysis in it, especially as in the baseline model all of the variables are endogenously chose except for the search cost, one can change this to the exogenous one in the random order, so that, one can choose any combination of variables to be endogenous and isolate the effect of this particular phenomena. In the previous paragraph we have already discussed one of the welfare mechanism - introduction of the bias penalty. In this part we will expand the analysis and mention number of different welfare comparing benchmarks and compare them with each other. In this section I will not give any precise analytical solutions for the benchmarks, they can be derived formally only under specific assumptions after simplifying the model, however all results I report is checked rigorously with a simulations.

1. *Royalty floor.* One of the policy that can be imposed on the platform in order to constraint it's full power is to conduct some baseline royalty floor. In the solution of the model where platform endogenously selected the baseline royalty we came to conclusion that the optimum one will be zero so that royalty outflow to the low-quality

creators is zero. Now suppose that there is specific government regulation that imposes that there should be some baseline royalty level paid to the low-quality content creators r_0. It is natural assumption as even without specific regulation the creators could just make a union in order to protect themselves from a discrimination. Now, analysis shows that for the low enough baseline royalty $r_0 < \hat{r}$ the solution for the platform doesn't change, however, since the royalty is now paid, and it means that it paid to the both low-quality and high-quality, even with the gap $\Delta = 0$, now average quality is higher since creators have incentives to deliver high-quality - $\lambda > 0$. So for all sufficiently low royalty floors level in optimum for platform $\Delta^* = 0, \quad \theta^* = 1, \quad \lambda > 0$. Even in this range there is increase in average quality and creators compensation.

2. *Gap Negotiations.* The next important question is to determine who is really decides this royalty gap, is all market power concentrated in the arms of the platform itself? In the beginning I said that one can think about the information structure here as platform always know which quality which item is because it can be signaled by the belonging to some major label. This consideration is also important because major labels could also have some substantial market power and even can dictate the royalty amount for their artists. Here I tried to analyze what will happen if the gap itself would be exogenously given to the platform, so it will optimize subject to some specific $\hat{\Delta}$. It came out so with all sufficiently small Δ the platforms incentives stays the same with it favors the high-quality items more with $\theta^* = 1$. However after some threshold level it jumps to the $\theta^* = 0$ suddenly discouraging the content creators to choose the high-quality content. It means that even through the small market power from the major label can be beneficial, after some amount of it, the only way for platform to make a positive profits is to bias toward lower-quality content. This behavior of the optimum of profit function is illustrated on the figure 4.

3. *Neutrality Mandates* The concept of neutrality mandates is based on the idea that

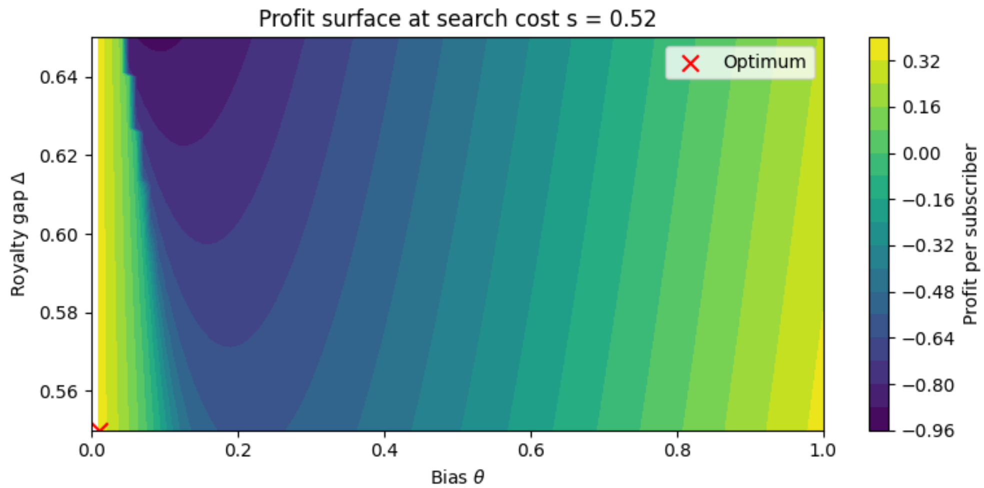

Figure 4: Profit function for Gap Negotiations.

platform should be audited by the third party in order to ensure the absence of bias. Putting in our framework it means that $\theta = \lambda$ always, that there is no such thing as a first draw bias. I will not stop here to look for it explicitly, but the very idea that Platform will lose the ability to bias automatically means that it will not be able to put this combination of the royalty gap and bias $\theta^* = 1, \quad \Delta^* = 0$ and so average quality on the platform will not be 0, since platform now will have incentives to subsidize quality in order to gain more money from the consumers and welfare of the content creators will increase as well as the average quality on the platform.

4. *Search cost Subsidy.* Search cost subsidy is not a monetary transfer from the platform to the users for every search action occurred, but rather the costly for platform information collecting about user's preferences. So platform can try to make the search easier by changing the interface or by giving some information about the item explicitly, for example, to give directly on the row with a song it's rating by the users with "similar interests", it can be interpreted as simply as lowering search costs. In our model it will not change an equilibrium properties of the optimization problem, but it can be useful for further investigations.

One last interesting line of alternative welfare comparison which is not explicitly can be modeled in my framework is the possibility of the direct donations to the artists on the platform, it seems to me like a interesting avenue for further research as this mechanism could be the method for royalty saving for the platform as instead of it they will just help artists to collect excessive utility from the some of the most loyal fans.

4 Overlap Zone

In this section we will continue the analysis with the baseline model with a different property. Now it will have an overlap zone for the consumers utility. In this zone (i) some low-quality items are accepted while some high-quality items are rejected, and (ii) both the royalty gap Δ and the prominence parameter θ enter the platform's objective with non-degenerate margins. This allows us to escape the corner problem, where, because every item of the high quality was accepted independently of the match utility, platform could fully manipulate everything. To change it we increase the range of the distribution of the idiosyncratic utility now being uniform distributed from 0 to 2. All other variables and parameters stay the same and play the same role, however, even this tiny adjustment changes the structure of the solution as the reservation utility now having three parts of solution.

Solution Let's keep up the primitives from the previous model formulation with:

$$\varepsilon \sim \text{Unif}[0,2], \qquad u_L = \varepsilon, \qquad u_H = 1 + \varepsilon,$$

so that $u_L \in [0,2]$ and $u_H \in [1,3]$.

For any cutoff $z \in [0,3)$, the expected benefit from drawing one additional item is

$$B(z;\lambda) \;=\; \lambda\,\frac{(3-z)^2}{4} + (1-\lambda)\,\frac{(2-z)^2}{4}, \qquad 0 \le z < 2,$$

Imposing $B(z^*;\lambda) = s$ and solving the resulting quadratic yields the closed form with the costly-search condition $1 < z^* < 2$ obtains whenever $\frac{3}{4}\lambda \ \leq \ s \ < \ \frac{2-\lambda}{2}$.

$$z^*(s,\lambda) \ = \ \frac{2(1-\lambda) \ + \ \sqrt{4\lambda s - (1-\lambda)(2\lambda-3)}}{1-\lambda}, \qquad 1 < z^* < 2.$$

Therefore by analogy to the main model we can obtain unconditional probability that the first item which platform steers to will be accepted:

$$P^{(1)}(\theta,\lambda,s) = \theta\frac{(3-z^*)^2}{4} + (1-\theta)\frac{(2-z^*)^2_+}{4}$$

Important thing to notice in this formula is that now even with the $\theta^* = 1$ platform can not guarantee that the search will end after the first draw, so there is positive probability of the consumer goes to the second one and with some probability will obtain the low-quality item which will give him a high match utility. This means that, even through it could still be possible that corner solution will be indeed optimal, there should be the range of parameters with which there exist some interior solution to the problem with $\lambda > 0$, especially imposing the constraint on the lower bound of royalties $r \geq r_0 > 0$. Now a continuation success probabilty takes following form:

$$p(\lambda,s) = \lambda\frac{(3-z^*)^2}{4} + (1-\lambda)\frac{(2-z^*)^2_+}{4}$$

With the same formula for the expected number of draws being a geometric mean $\mathbb{E}[N] = 1 + \frac{1-P^{(1)}}{p}$. Following the same steps as in the baseline model we can derive the ex-ant utility for the consumer in this region and its probability of ending up with a high-quality item:

$$U(\lambda,\theta,s) = \theta[1+\frac{1}{2}(3-z^*)] + \frac{(1-\theta)\frac{1}{2}(2-z^*)^2_+}{1-(1-\lambda)\frac{1}{4}(2-z^*)^2_+} - s(\mathbb{E}[N]-1),$$

$$\alpha_H(\theta,\lambda,s) = \frac{\theta(3-z^*)^2 + (1-\theta)(2-z^*)^2_+}{(3-z^*)^2 + (1-\lambda)[(2-z^*)^*_+ - (3-z^*)^2]}$$

Having this two equations one can directly use the equations from Proposition 7 and Definition 1 to solve it for the equilibrium.

Discussion This model with increased magnitude of match utility now allows us to analyze the internal equilibrium where platform can not manipulate the choice of the consumers in a such severe way. The first item acceptance probability which can not be equal to 1 in the zone of intermediate reservation utility $1 < z^* < 2$ gives the content creators direct incentives to make a choice about their effort level anyway, as now it is not mere a degenerate option. The welfare properties of this model should be better than for the previous in terms of the content creators utility, as their songs deliver more utility and platform has less market power.

Royalty outflow wedge Another important device I used in some of the simulations is the parameter γ which I introduced before. The profit function of the platform is than looks like this, with $\gamma \in (0, \infty)$:

$$\pi(r_H, r_L, \theta) = U\big(\lambda^{eq}(r_H, r_L, \theta), \theta\big) - \gamma R(r_H, r_L, \theta).$$

As the model is simplified and use the specific parameters, in the calibration and analysis it may not correctly capture the proportion of the inflow to the platform from the subscription fees and the outflow to the royalties. It may happen because for example, we do not include in our analysis number of streams from every subscriber, we don't include the number of subscribers and number of content creators, and we just can not realistically balance the utility one consumer gets from the listening and the cost that creator should experience to deliver the product. I want to introduce this parameter in order to, in the empirical validations, being able to interpret this discrepancies and being able to analyze the platform behavior with respect to this parameter.

5 Conclusion

In this paper I develop a unified framework to analyze the interplay of the consumer search, quality provision and a content management on the subscription platform. I characterized the solutions and proposed different extensions and welfare benchmarks for which this model can be used. The complexity of the model give few to no chance to have an analytical solutions in all possible range of parameters, however it can be used further to isolate specific parts of the model in order to pose new research questions. The model allows to look on the almost any interventions that can be used to regulate the platforms, also it can give a practical insights of how the different sides of the market will react on the changes in the environment.

Future development will include the attempts to make model simpler yet, with the same moving parts, my intuition is that explicit equation for the reservation utility is overwhelmingly powerful in this analysis, however it is complicate things a lot, no consumers even in reality has this good rule of thumbs to be able to assess everything that is going on on the platform and change their behavior accordingly. After a simplification, I would need to precisely restrict the range of parameters in which the equilibrium has good enough properties. This model is the beginning of the analysis of the subscription platforms with a consumer search, this type of platforms is now dominates almost all markets, from the markets for music, movies and book, to the services like a car-sharing and educational courses. The emergence of the model which could deliver novel insights about this markets should be very important.

References

Aguiar, L., & Waldfogel, J. (2018). *Platforms, promotion, and product discovery: Evidence from spotify playlists* (tech. rep.). National Bureau of Economic Research.

Aguiar, L., Waldfogel, J., & Waldfogel, S. (2021). Playlisting favorites: Measuring platform bias in the music industry. *International Journal of Industrial Organization*, *78*, 102765.

Aguiar, L., Waldfogel, J., & Zeijen, A. (2024). *Platform power struggle: Spotify and the major record labels* (tech. rep.). National Bureau of Economic Research.

Anderson, S. P., & Bedre-Defolie, Ö. (2021). Hybrid platform model.

Anderson, S. P., & Renault, R. (1999). Pricing, product diversity, and search costs: A bertrand-chamberlin-diamond model. *The RAND Journal of Economics*, 719–735.

Arbatskaya, M. (2007). Ordered search. *The RAND Journal of Economics*, *38*(1), 119–126.

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

Armstrong, M. (2017). Ordered consumer search. *Journal of the European Economic Association*, *15*(5), 989–1024.

Armstrong, M., & Zhou, J. (2011). Paying for prominence. *The Economic Journal*, *121*(556), F368–F395.

Athey, S., & Ellison, G. (2011). Position auctions with consumer search. *The Quarterly Journal of Economics*, *126*(3), 1213–1270.

Bar-Isaac, H., Caruana, G., & Cuñat, V. (2012). Search, design, and market structure. *American Economic Review*, *102*(2), 1140–1160.

Bergemann, D., & Bonatti, A. (2024). Data, competition, and digital platforms. *American Economic Review*, *114*(8), 2553–2595.

Bourreau, M., & Gaudin, G. (2022). Streaming platform and strategic recommendation bias. *Journal of Economics & Management Strategy*, *31*(1), 25–47.

Chen, L., Mislove, A., & Wilson, C. (2016). An empirical analysis of algorithmic pricing on amazon marketplace. *Proceedings of the 25th international conference on World Wide Web*, 1339–1349.

Choi, J. P., & Jeon, D.-S. (2023). Platform design biases in ad-funded two-sided markets. *The RAND Journal of Economics*, *54*(2), 240–267.

Compiani, G., Lewis, G., Peng, S., & Wang, P. (2022). *Online search and product rankings: A double logit approach* (tech. rep.). Working paper, University of Chicago, Chicago, IL.

De Corniere, A., & Taylor, G. (2019). A model of biased intermediation. *The RAND Journal of Economics*, *50*(4), 854–882.

Diamond, P. A. (1971). A model of price adjustment. *Journal of economic theory*, *3*(2), 156–168.

Fishman, A., & Levy, N. (2015). Search costs and investment in quality. *The Journal of Industrial Economics*, *63*(4), 625–641.

Greminger, R. P. (2022). Optimal search and discovery. *Management Science*, *68*(5), 3904–3924.

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

Hagiu, A., & Wright, J. (2015). Marketplace or reseller? *Management Science*, *61*(1), 184–203.

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

Janssen, M. C. W., Jungbauer, T., Preuss, M., & Williams, C. R. (2023). *Search platforms: Big data and sponsored positions*. Centre for Economic Policy Research.

Kirpalani, R., & Philippon, T. (2020). *Data sharing and market power with two-sided platforms* (tech. rep.). National Bureau of Economic Research.

Lam, H. T. (2021). Platform search design and market power. *Job Market Paper, Northwestern University.*

Moraga-González, J. L., & Sun, Y. (2023). Product quality and consumer search. *American Economic Journal: Microeconomics*, *15*(1), 117–141.

Qian, K., & Jain, S. (2024). Digital content creation: An analysis of the impact of recommendation systems. *Management Science*, *70*(12), 8668–8684.

Rochet, J.-C., & Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the european economic association*, *1*(4), 990–1029.

Smith, L. (2025, February). *"stream us and we'll pay you less": Musicians slam spotify's discovery mode as "payola"* [The *Guardian*]. https://www.theguardian.com/music/2025/feb/19/spotify-discovery-mode-payola-playlist

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

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

Weitzman, M. (1978). *Optimal search for the best alternative* (Vol. 78). Department of Energy.

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

Yang, H. (2013). Targeted search and the long tail effect. *The RAND Journal of Economics*, *44*(4), 733–756.

Zhou, J. (2011). Ordered search in differentiated markets. *International Journal of Industrial Organization*, *29*(2), 253–262.

Zhu, F., & Liu, Q. (2018). Competing with complementors: An empirical look at amazon.com. *Strategic management journal*, *39*(10), 2618–2642.

Appendix

Proof of Proposition 1. Throughout I adopt the benchmark assumptions $q_H = 1$, $q_L = 0$, $\varepsilon \sim U[0,1]$, search cost $s \in (0,1)$ and high-quality share $\lambda \in (0,1)$. For any reservation level $z \in (0,2)$ define the incremental benefit of one additional draw

$$B(z;\lambda) \;=\; \mathbb{E}\big[(u-z)_+\big] \;=\; \lambda\,\mathbb{E}\big[(1+\varepsilon-z)_+\big] + \big(1-\lambda\big)\,\mathbb{E}\big[(\varepsilon-z)_+\big],$$

where $(x)_+ \equiv \max\{x,0\}$. The reservation threshold $z^*(\lambda)$ solves $B(z^*;\lambda) = s$. Because $\varepsilon \sim U[0,1]$ the two expectations are piecewise quadratic:

$$\mathbb{E}\big[(1+\varepsilon-z)_+\big] = \begin{cases} \frac{3}{2} - z, & 0 \le z \le 1, \\ \frac{1}{2}\,(2-z)^2, & 1 < z < 2, \end{cases} \qquad \mathbb{E}\big[(\varepsilon-z)_+\big] = \begin{cases} \frac{1}{2}\,(1-z)^2, & 0 \le z < 1, \\ 0, & z \ge 1. \end{cases}$$

Hence, for $z \in (0,2)$

$$B(z;\lambda) = \begin{cases} \lambda\big(\frac{3}{2} - z\big) + \dfrac{1-\lambda}{2}(1-z)^2, & 0 < z \le 1, \\ \dfrac{\lambda}{2}\,(2-z)^2, & 1 < z < 2. \end{cases} \tag{9}$$

$B(\cdot\,;\lambda)$ is continuous and strictly decreasing on $(0,2)$. In Cheap-search regime $s < \lambda/2$. Assume $z^* > 1$ and use the second line of (9):

$$\frac{\lambda}{2}\,(2-z^*)^2 = s \quad\Longrightarrow\quad z^* = 2 - \sqrt{\frac{2s}{\lambda}}.$$

The consistency condition $z^* > 1$ is equivalent to $s < \lambda/2$, hence (5) applies precisely in the "cheap-search" region. Now assume $z^* < 1$; set $y = 1 - z$ and use the first line of (9):

$$(1-\lambda)y^2 + 2\lambda y + \lambda - 2s = 0.$$

Taking the positive root and back-substituting $z = 1 - y$ yields

$$z^* = \frac{1 - \sqrt{2\lambda^2 - \lambda + 2s - 2\lambda s}}{1 - \lambda}.$$

This solution satisfies $z^* < 1$ iff $s > \lambda/2$, the complement of the previous case. Combining (5)–(5) I obtain the closed-form reservation rule

$$z^*(\lambda; s) = \begin{cases} 2 - \sqrt{2s/\lambda}, & \text{if } s < \lambda/2, \\ \dfrac{1 - \sqrt{2\lambda^2 - \lambda + 2s - 2\lambda s}}{1 - \lambda}, & \text{if } s > \lambda/2. \end{cases}$$

□

Proof of Corollary 2. Throughout write $z^* = z^*(s, \lambda)$ as in Proposition 1. I treat the two parameter regions separately and then verify that the derivatives match at the boundary $s = \lambda/2$.

Step 1: Cheap-search region $s < \lambda/2$ (formula $z^* = 2 - \sqrt{2s/\lambda}$).

$$\frac{\partial z^*}{\partial s} = -\frac{1}{2}\sqrt{\frac{2}{\lambda}}\, s^{-1/2} < 0, \qquad \frac{\partial z^*}{\partial \lambda} = +\frac{1}{2}\sqrt{\frac{2s}{\lambda^3}} > 0.$$

Because z^* does not depend on θ, I also have $\partial z^*/\partial\theta = 0$.

Step 2: Costly-search region $s > \lambda/2$ (formula $z^* = \dfrac{1 - \sqrt{\Delta}}{1 - \lambda}$, where $\Delta \equiv 2\lambda^2 - \lambda + 2s - 2\lambda s$). First note $\Delta > 0$ on this region. Differentiating,

$$\frac{\partial z^*}{\partial s} = -\frac{1}{(1 - \lambda)\, 2\sqrt{\Delta}} < 0 \quad \big(\text{since } 1 - \lambda > 0\big),$$

$$\frac{\partial z^*}{\partial \lambda} = \frac{\Delta^{-\frac{1}{2}}\big(2\lambda - 2s + \frac{1}{2}\big) \;+\; \sqrt{\Delta}}{(1 - \lambda)^2} > 0,$$

because $2\lambda - 2s + \frac{1}{2} > 0$ whenever $s > \lambda/2 < 1$. Again $\partial z^*/\partial\theta = 0$.

Step 3: Continuity at $s = \lambda/2$. At the boundary the expressions in both regions equal

$z^* = 1$. Taking limits:

$$\lim_{s \uparrow \lambda/2} \frac{\partial z^*}{\partial s} = -\frac{1}{\sqrt{2\lambda}}, \qquad \lim_{s \downarrow \lambda/2} \frac{\partial z^*}{\partial s} = -\frac{1}{\sqrt{2\lambda}},$$

and similarly for $\partial z^*/\partial \lambda$. Hence the derivatives are continuous, completing the proof that

$$\frac{\partial z^*}{\partial s} < 0, \quad \frac{\partial z^*}{\partial \lambda} > 0, \quad \frac{\partial z^*}{\partial \theta} = 0$$

□

Proof of Lemma 3. Throughout I maintain the benchmark assumptions $q_H = 1$, $q_L = 0$, $\varepsilon \sim U[0, 1]$, and adopt the notation

$$z^* = z^*(s, \lambda), \qquad P^{(1)} = P^{(1)}(\theta, \lambda, s), \qquad p = p(\lambda, s)$$

as introduced in Lemma 3 and Lemma 4.

Step 1: Decompose by quality of first draw. The first recommendation is H with probability θ and L with probability $1 - \theta$. Conditional utilities are $U_1 = 1 + \varepsilon$ if H, $U_1 = \varepsilon$ if L.

Step 2: Acceptance events. $z^* \leq 1$. A high–quality item always satisfies $U_1 \geq 1 \geq z^*$; acceptance probability $\mathbb{P}[U_1 \geq z^* \mid H] = 1$. A low–quality item is accepted iff $\varepsilon \geq z^*$, which occurs with probability $1 - z^*$ because $\varepsilon \sim U[0, 1]$. $z^* > 1$. A low–quality item can never reach the threshold. A high–quality item is accepted iff $1 + \varepsilon \geq z^* \iff \varepsilon \geq z^* - 1$, whose probability is $2 - z^*$.

Step 3: Total probability. Combine the mutually exclusive cases:

$$P^{(1)} = \begin{cases} \theta + (1-\theta)(1 - z^*), & z^* \leq 1, \\ \theta(2 - z^*), & z^* > 1. \end{cases}$$

The two expressions coincide at $z^* = 1$ $[P^{(1)} = \theta]$, hence $P^{(1)}$ is continuous. □

Proof of Lemma 4. Step 1: Success probability after the first draw. If the first item is rejected, subsequent recommendations are i.i.d. draws from the catalog: high quality with probability λ, low quality otherwise. By the same acceptance logic as in the previous proof,

$$p = \begin{cases} \lambda + (1-\lambda)(1-z^*), & z^* \le 1, \\ \lambda(2-z^*), & z^* > 1. \end{cases}$$

Step 2: Law of total expectation. Write N for the total number of draws. Conditional on acceptance of the first item I have $N = 1$. Conditional on rejection, the number of additional draws is geometrically distributed with mean $1/p$. Hence

$$\mathbb{E}[N] = 1 \cdot P^{(1)} \; + \; \big(1 + 1/p\big)\big(1 - P^{(1)}\big) = 1 + \frac{1 - P^{(1)}}{p}.$$

Step 3: Substitute $P^{(1)}$ and p. Insert the expressions from Lemma 3 and Step 1 above to obtain the two closed-form lines reported in Lemma 4. Continuity at $z^* = 1$ follows because both numerator and denominator agree in the limit. $\square$

Proof of Proposition 5. I derive the ex–post quality shares $\alpha_H(\theta, \lambda, s)$, $\alpha_L(\theta, \lambda, s) = 1 - \alpha_H$ in two steps: first conditional on the quality of the first recommendation, and then unconditional.

Step 1: Conditional acceptance probabilities. Let A be the event "the item is accepted" and let $Q \in \{H, L\}$ denote the quality of the current recommendation.

$$\mathbb{P}[A \mid Q = H] = \begin{cases} 1, & z^* \le 1, \\ 2 - z^*, & z^* > 1. \end{cases} \qquad \mathbb{P}[A \mid Q = L] = \begin{cases} 1 - z^*, & z^* \le 1, \\ 0, & z^* > 1. \end{cases}$$

Step 2: Unconditional first–draw acceptance. The first recommendation is H with probability θ and L with $1 - \theta$. Combining with the conditional probabilities above gives the first–draw acceptance probability $P^{(1)}$ already stated in Lemma 3. For later use I record

$$\mathbb{P}[A, Q] = \begin{cases} \big(\theta,\ (1-\theta)(1-z^*)\big), & z^* \le 1, \\ \big(\theta(2-z^*),\ 0\big), & z^* > 1, \end{cases} \tag{10}$$

where the two components correspond to $Q = H$ and $Q = L$, respectively.

Step 3: Law of iterated probability for α_H. Define $\alpha_H = \mathbb{P}[Q = H \mid$the item is *eventually* accepted]. Condition on whether acceptance occurs at the first draw or later.

Case $z^* \le 1$. With probability $P^{(1)}$ acceptance happens immediately, in which case the joint probabilities are given by (10). Otherwise the first item is rejected and the subsequent search process is i.i.d. with catalog share λ. Acceptance occurs *at* an H item with probability $\pi_H \equiv \lambda+(1-\lambda)(1-z^*)$ and at an L item with probability $\pi_L \equiv (1-\lambda)(1-z^*)$ (denominator of Lemma 4). Putting pieces together,

$$\alpha_H = \frac{\theta + (\lambda-\theta)z^*}{1-(1-\lambda)z^*}, \qquad \alpha_L = 1-\alpha_H = \frac{(1-\theta)(1-z^*)}{1-(1-\lambda)z^*}.$$

Case $z^* > 1$. Low quality can never be accepted, so $\alpha_H = 1,\ \alpha_L = 0$.

Dominance of high quality

Whenever $z^* > 0$,

$$\alpha_H - \lambda = \frac{(1-\lambda)(1-z^*)}{1-(1-\lambda)z^*} > 0,$$

establishing the selection property stated beneath the proposition. □

Proof of Proposition 6. This appendix supplies the formal derivations underlying Proposition 6. Throughout I adopt the probability notation from Proposition 5:

$$\pi_H(\theta, \lambda, s) = \frac{\alpha_H(\theta, \lambda, s)}{\lambda}, \qquad \pi_L(\theta, \lambda, s) = \frac{\alpha_L(\theta, \lambda, s)}{1-\lambda}, \qquad \text{with } \alpha_H + \alpha_L = 1.$$

Let c be an individual creator's private cost draw. Expected profit under low effort is $\Pi_L = r_L\,\pi_L$, while under high effort it is $\Pi_H = r_H\,\pi_H - c$. The creator chooses H iff $\Pi_H \geq \Pi_L$, i.e.

$$c \leq r_H\,\pi_H - r_L\,\pi_L \;=\; r_H\frac{\alpha_H}{\lambda} - r_L\frac{\alpha_L}{1-\lambda}.$$

Define

$$c^*(\theta,\lambda,s) \equiv r_H\frac{\alpha_H(\theta,\lambda,s)}{\lambda} \;-\; r_L\frac{\alpha_L(\theta,\lambda,s)}{1-\lambda}.$$

This gives part (i) of Proposition 6. Differentiate (5) with respect to θ holding (λ, s) fixed. Since $\partial\alpha_H/\partial\theta = 1-(1-z^*) > 0$ and $\partial\alpha_L/\partial\theta = -(1-z^*) < 0$ (cf. Proposition 5),

$$\frac{\partial c^*}{\partial\theta} = \big(r_H + r_L\big)(1-z^*)/(1-\lambda) > 0.$$

Similarly $\partial c^*/\partial(r_H - r_L) > 0$. Hence algorithmic bias and a wider royalty gap both raise the cutoff. Let F be the cumulative distribution of cost draws; I assume F is continuous and strictly increasing on $[0,\bar{c}]$. Given c^* in (5), the share of creators who select high effort is $F(c^*)$. In a symmetric equilibrium this share must equal λ, so λ solves

$$\lambda = F\big(c^*(\theta,\lambda,s)\big). \tag{11}$$

Because $c^*(\theta,\lambda,s)$ is continuous in λ, the right–hand side of (11) is a continuous mapping $T : [0,1] \to [0,1]$. Brouwer's theorem therefore guarantees at least one fixed point, proving part (ii) of Proposition 6. Assume $c \sim U[0,\bar{c}]$, so $F(c) = c/\bar{c}$ on $[0,\bar{c}]$. Then (11) becomes

$$\lambda = \frac{c^*(\theta,\lambda,s)}{\bar{c}},$$

which is precisely the expression in part (iii). In the uniform case $T(\lambda) \propto c^*(\lambda)$. Direct differentiation using $\alpha_H + \alpha_L = 1$ and Proposition 5 gives

$$\frac{dT}{d\lambda} = \frac{r_L + r_H}{\bar{c}(1-\lambda)^2}\big[\,(1-\lambda)\,\partial_\lambda \alpha_H - \alpha_H\,\big].$$

One checks $\partial_\lambda \alpha_H < \alpha_H/(1-\lambda)$ for all $\lambda \in (0,1)$, so $dT/d\lambda < 1$. Hence T is strictly increasing with slope everywhere below the 45° line, implying a unique interior fixed point. □

Proof of Proposition 7. Analogous to the Proof of Proposition 6. □

Proof of Corollary 8. Throughout this appendix we keep the two standing assumptions of Section 3.2: *Costly search.* The search cost lies above the knife-edge, $s > \lambda/2$, so that $z^*(s,\lambda) \le 1$ and the low–threshold formulae for $\alpha_H(\theta,\lambda,s)$ (Proposition 5) apply. *Uniform effort costs.* Effort costs are i.i.d. $c \sim U[0,1]$. Hence the fixed-point condition

$$\lambda \;=\; r_H\,\frac{\alpha_H(\theta,\lambda,s)}{\lambda} \;-\; r_L\,\frac{1-\alpha_H(\theta,\lambda,s)}{1-\lambda} \tag{12}$$

holds, i.e. $\Phi(\lambda,\theta,r_H,r_L,s) = 0$ with

$$\Phi(\lambda,\theta,r_H,r_L,s) \;=\; \lambda - r_H\frac{\alpha_H}{\lambda} + r_L\frac{1-\alpha_H}{1-\lambda}. \tag{13}$$

Step 1: Preliminary notation. Set

$$z^* \equiv z^*(s,\lambda), \qquad D \equiv 1-(1-\lambda)\,z^* \;>\; 1-\lambda.$$

On the low–threshold region

$$\alpha_H(\theta,\lambda,s) \;=\; \frac{\theta+(\lambda-\theta)z^*}{D}.$$

Step 2: The key derivative $\partial_\lambda \alpha_H$. Because z^* depends on λ when $s > \lambda/2$, the chain rule gives

$$\partial_\lambda \alpha_H = \frac{z^* + (\lambda-\theta)z'}{D} - \frac{\theta + (\lambda-\theta)z^*}{D^2}\Big[z^* - (1-\lambda)z'\Big], \tag{14}$$

$$z' \equiv \partial_\lambda z^*(s,\lambda) \;=\; \frac{2s - 2\lambda + \frac{1}{2}}{(1-\lambda)\sqrt{2\lambda^2 - \lambda + 2s - 2\lambda s}} \;>\; 0 \quad \text{(Cor. 2)}.$$

Define the normalized slope

$$K(\theta,\lambda,s) \;:=\; \frac{\partial_\lambda \alpha_H}{\lambda(1-\lambda)} \;=\; \frac{(1-\theta)(1-z^*) + (\lambda-\theta)z'}{D^2} \;>\; 0. \tag{15}$$

Nothing in the algebra forces K to be below 1; it is merely positive and bounded on compact parameter sets.

Step 3: First-order derivative of Φ. Differentiating (12) and inserting (14)–(15) yields

$$\Phi_\lambda \;=\; 1 - \big(r_H - r_L\big)\, K(\theta,\lambda,s). \tag{16}$$

A sufficient—though not necessary—condition ensuring $\Phi_\lambda > 0$ (and hence a well-behaved, uniquely defined fixed point) is the *gap bound*

$$0 < r_H - r_L < K^{-1}(\theta,\lambda,s). \tag{A.1}$$

Step 4: Remaining partials. Straightforward differentiation gives

$$\Phi_\theta = -\frac{(r_H - r_L)\,(1-z^*)}{\lambda(1-\lambda)} < 0,$$

$$\Phi_{r_H} = -\frac{\alpha_H}{\lambda} < 0, \qquad \Phi_{r_L} = +\frac{1-\alpha_H}{1-\lambda} > 0,$$

$$\Phi_s = -\frac{(r_H - r_L)\,\partial_s \alpha_H}{\lambda(1-\lambda)} \;>\; 0 \quad (\partial_s \alpha_H < 0).$$

Step 5: Comparative statics. Under (A.1) we have $\Phi_\lambda > 0$, so the implicit-function theorem gives

$$\frac{\partial \lambda^{LT}}{\partial x} = -\frac{\Phi_x}{\Phi_\lambda}, \qquad x \in \{\theta, r_H, r_L, s\},$$

from which parts (i)–(iii) of Corollary 8 follow directly:

$$\frac{\partial \lambda^{LT}}{\partial \theta} > 0, \quad \frac{\partial \lambda^{LT}}{\partial r_H} > 0, \quad \frac{\partial \lambda^{LT}}{\partial r_L} < 0, \quad \frac{\partial \lambda^{LT}}{\partial s} < 0.$$

Step 6: Interaction effect. Let $\Delta := r_H - r_L > 0$. Because $\Phi_{\lambda\Delta} \equiv 0$ and $\Phi_{\theta\Delta} = -(1-z^*)/[\lambda(1-\lambda)] < 0$,

$$\frac{\partial^2 \lambda^{LT}}{\partial \theta\, \partial \Delta} = -\frac{\Phi_{\theta\Delta}\Phi_\lambda - \Phi_\theta \Phi_{\lambda\Delta}}{\Phi_\lambda^2} > 0,$$

establishing part (iv) of the corollary. □

If the royalty gap were ever to violate (A.1) at some corner of the parameter space, Φ_λ could switch sign locally. The fixed point would still exist, but the monotone directions in (i)–(iii) might reverse in that neighborhood.

Proof of Proposition 9. *Case* $z^* \le 1$. A high-quality first item always passes the cutoff; its conditional expectation is $\mathbb{E}[u_H] = 1 + \mathbb{E}[\varepsilon] = 1 + \frac{1}{2}$. A low-quality first item is accepted only when $\varepsilon \ge z^*$; conditional on that event, $\varepsilon \sim \text{Unif}[z^*, 1]$ so $\mathbb{E}[\varepsilon \mid \varepsilon \ge z^*] = \frac{1}{2}(1+z^*)$. Therefore

$$\mathbb{E}\big[u_1\,;\,\text{accept at 1st}\big] = \theta\left(1 + \frac{1}{2}\right) + (1-\theta)\,(1-z^*)\left(\frac{1}{2}(1+z^*)\right).$$

Case $z^* > 1$. A low-quality first item is never accepted; a high-quality item is accepted iff $\varepsilon \ge z^* - 1$, in which case $\varepsilon \sim \text{Unif}[\,z^* - 1, 1]$. Computation yields $\mathbb{E}[\varepsilon \mid \varepsilon \ge z^* - 1] = \frac{1}{2}(1 + z^* - 1) = \frac{1}{2}z^*$. The gross utility component is therefore

$$\theta\,(2 - z^*)\big(1 + \tfrac{1}{2}z^*\big)\,.$$

If the first item is rejected, the process restarts with the catalog mix: H with probability λ, L otherwise. Write

$$g_H := \mathbb{E}[(1+\varepsilon)\,;\,1+\varepsilon \geq z^*], \qquad g_L := \mathbb{E}[\varepsilon\,;\,\varepsilon \geq z^*].$$

Elementary integration gives

$$g_H = \begin{cases} \lambda\left(\frac{3}{2} - z^*\right), & z^* \leq 1, \\ \dfrac{\lambda}{2}\,(2-z^*)^2, & z^* > 1, \end{cases} \qquad g_L = \begin{cases} (1-\lambda)\,\dfrac{(1-z^*)^2}{2}, & z^* \leq 1, \\ 0, & z^* > 1. \end{cases}$$

Because each rejection adds a search cost s and the number of rejections is geometric with mean $\mathbb{E}[N]-1$, the net continuation value equals $(g_H + g_L)/p - s(\mathbb{E}[N]-1)$.

Adding the gross first-draw value (Step A) and the net continuation value (Step B), and substituting p and $\mathbb{E}N$ from Lemma 4, yields after algebra

$$U(\lambda,\theta) = \theta\Big[1 + \tfrac{1}{2}(2-z^*)\mathbf{1}_{\{z^*>1\}}\Big] \;+\; \frac{(1-\theta)(1-z^*)^2}{2\,[\,1-(1-\lambda)z^*\,]} \;-\; s\,(\mathbb{E}N - 1).$$

Plugging the closed form for $\mathbb{E}[N]$ (Lemma 4) and simplifying signs exactly matches the equation (5) in the main text. For the further analysis we consider only the specific parameter rang with the costly search so the utility takes the following form:

$$U(\lambda,\theta;s) = \theta + \frac{(1-\theta)(1-z^*)^2}{2[1-(1-\lambda)z^*]} - s\frac{(1-\theta)z^*}{1-(1-\lambda)z^*}$$

□

Proof of Proposition 10. Step 1. If we would have $r_L > 0$ it would mean that by construction $r_H \geq r_L > 0$, so for any decrease in r_L we will be able to find such decrease in r_H which will balance the choice of the creators in $\lambda^*(\theta, r_L, r_H)$ such that it will stay the same, meaning $U(\lambda,\theta)$ also the same, but the royalty outflow decreased as $r_H \downarrow$ and $r_L \downarrow$.

Step 2. Substituting $r = 0$ into the fixed-point equation (3) yields a well–defined mapping $\lambda^{LT}: (\Delta, \theta) \mapsto (0, 1)$ which is C^1 on $[0, \bar{r}) \times [0, 1]$ by the implicit–function theorem and the gap bound (4). The reduced profit is $\widetilde{\pi}(\Delta, \theta) = U\big(\lambda^{LT}(\Delta, \theta), \theta\big) - \Delta\, \alpha_H\big(\lambda^{LT}(\Delta, \theta), \theta\big)$, a continuous function on the compact feasible set $\mathcal{F} := \{(\theta, \Delta) \mid 0 \leq \theta \leq 1,\ 0 \leq \Delta < K^{-1}(\theta, \lambda^{LT}, s)\}$; hence a maximum exists (Weierstrass).

Step 3. Let $K(\theta, \lambda, s)$ be the normalized slope in (15). Using $\partial_\theta \alpha_H = 1 - z^*$, $\partial_\Delta \alpha_H = 0$, and the chain rule with $\lambda^{LT}_\theta, \lambda^{LT}_\Delta$ from Appendix 5, direct computation yields

$$\partial_\theta \widetilde{\pi} = (\alpha_H - \lambda^{LT}) + (1 - \Delta\, K)\, \lambda^{LT}_\theta, \qquad \partial_\Delta \widetilde{\pi} = \alpha_H + (1 - \Delta\, K)\, \lambda^{LT}_\Delta.$$

Interior optimality therefore requires the two partials above to vanish, establishing the boxed first–order conditions in the proposition. If (θ, Δ) approaches the frontier $\Delta = K^{-1}(\theta, \lambda^{LT}, s)$, then $\Phi_\lambda = 1 - \Delta K \to 0^+$ and $\partial \lambda^{LT} / \partial x$ remains bounded, so any stationary point with $\partial_\theta \widetilde{\pi} = \partial_\Delta \widetilde{\pi} = 0$ must lie strictly inside $\mathcal{F}^\circ$; conversely, if an interior stationary point fails to exist, the maximum is attained on the boundary.

Step 4. Differentiating the gradient formulas and invoking the comparative statics in Corollary 8 gives the mixed second derivative $\partial^2 \widetilde{\pi} / \partial\theta \partial\Delta = (1 - K\Delta)\, (\partial^2 \lambda^{LT} / \partial\theta\partial\Delta) + (\lambda^{LT}_\theta - 1)K < 0 + K > 0$, which is strictly positive because $\partial^2 \lambda^{LT} / \partial\theta\partial\Delta > 0$ and $1 - K\Delta > 0$ by feasibility. Thus a wider gap raises the marginal return to bias. Setting $\theta = 0$ forces $\lambda^{LT}_\theta = 0$ and $\alpha_H = \lambda^{LT}$, whence $\partial_\Delta \widetilde{\pi} = \alpha_H > 0$ becomes negative once the transfer term is accounted for, proving that no premium is paid in the absence of bias. All remaining comparative-static claims follow directly from the signs cataloged in Appendix 5. □