

Gambato, Jacopo; Sandrini, Luca

Working Paper

Not as good as it used to be: Do streaming platforms penalize quality?

ZEW Discussion Papers, No. 24-045

Provided in Cooperation with:

ZEW - Leibniz Centre for European Economic Research

Suggested Citation: Gambato, Jacopo; Sandrini, Luca (2024) : Not as good as it used to be: Do streaming platforms penalize quality?, ZEW Discussion Papers, No. 24-045, ZEW - Leibniz-Zentrum für Europäische Wirtschaftsforschung, Mannheim

This Version is available at:

<https://hdl.handle.net/10419/300272>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DISCUSSION

// NO.24-045 | 06/2024

DISCUSSION PAPER

// JACOPO GAMBATO AND LUCA SANDRINI

Not as Good as It Used to Be: Do Streaming Platforms Penalize Quality?

Not as Good as it Used to Be: Do Streaming Platforms Penalize Quality?*

Jacopo Gambato[†]

Luca Sandrini[‡]

First version: 03.08.2023

This version: 19.06.2024

Abstract

In this study, we analyze the incentives of a streaming platform to bias consumption when products are vertically differentiated. The platform offers mixed bundles of content to monetize consumer interest in variety and pays royalties to sellers based on the effective consumption of the generated content. When products are not vertically differentiated, the platform has no incentive to bias consumption in equilibrium. With vertical differentiation, royalties can differ, and the platform biases recommendations in favor of the cheapest content, hurting consumers and high-quality sellers. Biased recommendations, if unconstrained, eliminate sellers' incentives to increase the quality of their content, but if constrained, may lead to the inefficient allocation of R&D efforts.

Keywords: platform economics, media economics, content aggregator, recommendation bias, innovation

JEL Codes: D4, L1, L5

*We thank the IOEK division at ZEW Mannheim, Marc Bourreau, Francesco Clavorà Braulin, Leonardo Madio, Andrea Mantovani, Goonj Mohan, Bernhard Ganglmair, Axel Gautier, Martin Peitz, Bianca Sanesi, Nicolas Schutz, and Christian Soegaard, as well as the participants to the 11th CRC Young Researchers Workshop, 11th Oligo Workshop, XII Ibeo Workshop, 50th EARIE Annual Conference, 12th MaCCI Annual Conference, 15th Paris Conference on Digital Economics, and 2024 EAYE Annual Meeting for their insightful comments.

[†]University of Mannheim, ZEW Mannheim, MaCCI - ja.gambato@gmail.com

[‡]ZEW Mannheim - Luca.Sandrini@zew.de

1 Introduction

Recommendation systems are ubiquitous in digital markets; they are a critical feature of digital platforms that host numerous sellers and consumers who often do not know each other. The adequate functioning of a recommendation system (and the consequent quality of the proposed matching) may determine whether a platform will thrive in the market.¹

Streaming platforms have found great success in the content market in the digital era thanks in part to their personalized features. In the music industry, for example, consumers rely on algorithmically generated playlists, such as the well-known “Discover Weekly” on Spotify, which is automatically generated for each user every week.² Such features are extremely popular and affect consumption patterns: [Aguiar et al. \(2021\)](#) show in their empirical investigation that inclusion in automatically generated playlists such as Spotify’s “New Music Friday” boosts future popularity compared to similar songs that are not included. The two observations strongly point to the ability of such platforms to affect an individual’s effective consumption bundle.³ Similarly, the news market heavily depends on algorithmic recommendations for connecting readers and news outlets. Using data from consumers’ browsing history, news aggregators can suggest the most appropriate news to users searching for a specific topic. This service is common, for example, on news aggregators such as Google News, Facebook, and X (formerly Twitter).⁴

The importance of recommendation systems for the digital economy at large, combined with the lack of transparency that characterizes them, has raised antitrust concerns about potential misconducts. Practices such as self-preferencing, price discrimination, and differential seller treatment generate significant distortions and, therefore, have made regulators worldwide wary. The recently enforced EU-based Digital Markets Act, which prohibits gatekeepers from engaging in behaviors deemed anti-competitive, is a clear example of the importance of this topic, especially given recent evidence of how common such practices seem to be in digital marketplaces ([Waldfoegel, 2024](#)).

This paper explores the incentives of digital platforms to bias consumption through algorithmic recommendation, as well as the effects of such bias on the innovation incentives of the platforms’ complementors (i.e., content creators). We show that the platform optimally designs a recommendation bias that shifts demand from high-quality products to low-quality ones. By artificially lowering the demand of high-quality content, the platform lowers the price premium of that content and minimizes the marginal costs of streaming it. The equilibrium bias is a function of the quality differential between the available goods and, if unconstrained, completely suppresses the incentive of sellers to vertically differentiate. Under some conditions, however,

¹Allegedly, Google gained its dominant position in search engines through its superior technology in recommending solutions for user queries compared to established rivals such as Yahoo.

²[Popper \(2016\)](#) reports that 40 million of Spotify’s (at the time) 100 million users accessed it in 2016. More recently, Spotify reported that Discover Weekly streamed 2.3 billion hours of music in the five years since its launch. See <https://newsroom.spotify.com/2020-07-09/spotify-users-have-spent-over-2-3-billion-hours-streaming-discover-weekly-playlists-since-2015/>.

³The most overt example is Spotify’s Discovery Mode, a service provided by Spotify to artists: “When an artist or label turns on Discovery Mode for a song, Spotify charges a commission on streams of that song in areas of the platform where Discovery Mode is active. All other streams of the same song in other areas of the platform are commission-free.” See <https://artists.spotify.com/en/discovery-mode>.

⁴See [Jeon and Nasr \(2016\)](#) and [Freimane \(2022\)](#) on the effects of competition between news aggregators and news outlets on their incentives to invest and market performance.

the platform may need to limit the recommendation bias to retain users and/or sellers.

We propose a framework in which two content providers (henceforth labeled “ a ” and “ b ”) sell their horizontally and vertically differentiated bundle goods to a unit mass of consumers uniformly distributed on the $[0, 1]$ line for a price p_j , $j \in \{a, b\}$. Each seller offers a bundled good that consists of only the content they produce (i.e., Bundle Good a (respectively b) is entirely made of Content a (b)). One can think about the bundle goods as representing music albums or newspapers. Each bundle incorporates a quality attribute that results from a costly investment in R&D. The sellers are assumed to be located at the extreme of the Hotelling line. Along the entire line, a platform (labeled p) offers a subscription-based service: upon paying a uniform fee f , a consumer can access a mix of content from a and b .⁵

The platform remunerates the sellers and pays them a royalty rate per share of the content shown to each consumer. Thus, consumers can choose among three bundles: Sellers a and b offer pure bundles, whereas the platform offers mixed bundles. If a consumer buys either of the pure bundles, only the content owned by one seller is consumed. Instead, by subscribing to the platform service, the consumers are offered a mix of content based on their preferences (Anderson and Neven, 1989; Hoernig and Valletti, 2007, 2011) and the platform’s recommendation system (Bourreau and Gaudin, 2022). The platform can be understood as an intermediary that smooths consumption for those who value a balanced mix of content.

Because consumers value high quality, we show that, without the platform’s generation of bias in its recommendation system, the equilibrium outcome has a higher price for (and more consumption of) the higher-quality product. While the platform can raise fees to monetize the higher average quality of its bundles, its ability to do so is limited because the rival is forced to offer a lower price than under no vertical differentiation. When consumers are offered their optimal consumption bundle, the quality differential hurts the platform. Thus, the platform always has the incentive to bias consumption away from the better, more expensive product if one is present. Because biasing consumption affects the sellers’ equilibrium prices, the platform trades off consumption bias and the ability to monetize efficiently on the consumer side.

When the quality difference is substantial, the platform is limited by consumer participation constraints and offers a personalized biased bundle that makes each consumer indifferent between joining and leaving the platform.⁶ At the same time, when the platform biases consumption away from the better and more expensive product, the seller offering such a product is penalized. If this penalty is severe enough, the seller could choose not to make their product available on the platform and compete directly with the other seller. Whenever this happens, the platform cannot remain active.⁷

Streaming platforms are known to popularize less-known artists who, therefore, benefit from substantial demand expansion when they join the service. To capture this additional dimension,

⁵More broadly, the platform we study can be understood as a content aggregator similar to online news aggregators. Subscription fees are common in music streaming platforms. Besides Spotify, notable examples include Deezer and Pandora.

⁶This assumption is strong but realistic. Platforms such as Spotify offer personalized content as playlists based on past consumption. The assumption is simply a reversal of what is already known: the platform being aware of a consumer’s taste instructs how much bias the consumer would be willing to tolerate.

⁷Consumers join the platform to mix content produced by both music labels: if the platform cannot attract both sellers, no consumer would be interested in joining.

we split the unit mass of consumers into two: some consumers are assumed to be *ex-ante* aware of the artists represented on the platform, while others are not. The latter group only learns about the artists and consumes their product if the platform manages to attract both music labels. Critically, when deciding whether to join the streaming service, the sellers’ options depend on the degree of demand expansion provided by the platform. Thus, the platform’s ability to bias consumption depends on the additional consumption it generates, a result that echoes the findings of [Jeon and Nasr \(2016\)](#).

Our findings have relevant implications both in the context of consumption steering digital markets and concerning the effect of subscription-based business models on sellers’ incentives to innovate. Steering emerges in equilibrium, not because of sellers competing for prominence but as a response of the platform to soften competition: the platform has the incentive to contain the price premium generated by the higher quality and the strong market presence of the better product. Biasing the recommendation system negatively impacts the seller with the high-quality product, whereas it benefits the runner-up by skewing consumption toward the latter, leading to substantial distortions in the incentives to innovate compared to a counterfactual scenario in which the platform is inactive.

The rest of the paper is structured as follows: After a review of the relevant literature, we introduce the model in [Section 2](#). [Section 3](#) presents our main results. In [Section 3.1](#), we solve the model for an exogenous level of vertical differentiation (modeled as an additional, fixed, stand-alone utility to all consumers). After solving and discussing the seller’s decision to participate as a function of the demand generated by the platform, we endogenize the seller’s decision to invest in quality ([Section 3.2](#)). In [Section 3.3](#), we discuss the robustness of our results. We complement the main analysis by considering several extensions in [Section 4](#). Finally, [Section 5](#) concludes the paper.

1.1 Related literature

The impact of recommendation systems on consumer choice has been the focus of many empirical investigations. Among these, the aforementioned [Aguiar et al. \(2021\)](#) and companion paper [Aguiar and Waldfogel \(2021\)](#) speak of the effect of inclusion in automatically generated playlists on the popularity of new songs on Spotify (see also [Aguiar et al., 2023](#)).⁸ Recently, research concerned with intermediaries’ incentives to strategically skew recommendations in a way that systematically harms consumers has been on the rise (see, for example, recent work by [Peitz and Sobolev, 2022](#) and [Motta, 2023](#)).

In this vein, [Bourreau et al. \(2021\)](#) analyze competition for prominence on digital platforms by comparing the bias generated when prominence is gained via monetary or data-based compensation. We contribute to this body of literature in various ways: We capture bias not by manipulating the search query but by manipulating the composition of the available bundles. Then, we assume the platform already has relevant information on the buyer’s side by building competition on the Hotelling line. Moreover, we tailor our framework to investigate the effect

⁸Generally, recommendation systems widen the range of products consumed, a phenomenon referred to as the “long-tail effect” ([Fleder and Hosanagar, 2009](#); [Brynjolfsson et al., 2011](#); [Oestreicher-Singer and Sundararajan, 2012](#); [Datta et al., 2018](#)).

of subscription-based business models on the sellers' incentive to innovate.

From this perspective, [Bourreau and Gaudin \(2022\)](#) is the paper that more closely relates to ours. While the two models share several aspects, they differ in ways that substantially affect the equilibrium analysis and results. In particular, we consider vertically differentiated content offered by competing sellers to highlight the role of quality in content aggregation. By contrast, [Bourreau and Gaudin \(2022\)](#) considers a baseline specification in which content is of similar quality. While they extend the analysis to incorporate different underlying marginal costs of production (which leads to bias in equilibrium, a result we also highlight in an extension), we endogenize the investment in quality, and with it, the vertical differentiation, where the authors take the wedge in marginal costs as a primitive of the model.

Furthermore, while the literature predicts a constant bias equally imposed on all consumers who join the platform, we unveil a more complex mechanism. The presence of a direct channel generates asymmetry in consumers' willingness to withstand the bias, with those in the middle of the distribution of tastes — who benefit more from mixing products — being more willing than those closer to the extremes to accept a more pronounced bias. By contrast, consumers closer in taste to the sellers' pure bundle are more likely to receive a recommendation that matches their preferences. Through this taste-driven outside option, the bias level (i.e., the amount of demand shifted from the cheapest product to the other) is non-monotonic in the quality differential. Quality is a demand shifter that incentivizes the producer of high-quality content to raise its price. Simultaneously, and for the same reason, a large quality differential between the existing goods implies that consumers are less willing to accept a biased recommendation that promotes the worse content. When the quality differential is small, the latter force is weak and dominated by the incentives to minimize costs. Eventually, when the quality differential reaches a certain threshold, either the consumer or the seller participation constraints start to bind.

This result is particularly interesting because it demonstrates the complexity of the problem faced by the platform, which must govern both sides of the market (users and sellers) and, simultaneously, weaken the negotiation position of its complementors to minimize costs. These results fit the music industry, where vertical quality is generally less evident than horizontal differentiation. In such a context, the platform has strong incentives to bias its recommendation system and can safely do so when the demand expansion it provides is substantial.

While we assume that sellers offer goods that are differentiated both horizontally and vertically, we explicitly consider consumers only sensible to horizontal variations of the product but who do not differ in their willingness to pay for quality.⁹ We do so to better relate to the literature on innovation, which often includes both vertical and horizontal dimensions of differentiation (see [Chen and Schwartz, 2013](#)).

Our findings suggest that the platform has the means and the incentive to bias consumption in favor of cheaper content and echo those in [Freimane \(2022\)](#), who examines the impact of a regulatory change affecting the bargaining process behind content provision on the Google

⁹In other words, we depart from [Mussa and Rosen \(1978\)](#), where consumers' income is considered (see also [Cremer and Thisse \(1991\)](#) and [Sutton \(1986\)](#)), and instead we follow the textbook definition of vertical differentiation in [Pepall et al. \(2014\)](#): Two goods are vertically differentiated if, when offered at the same price, all consumers prefer buying more of one than of the other.

News platform. The paper shows that the change, aimed at granting higher bargaining power to publishers vis-à-vis Google News, led the platform to change the composition of articles shown to readers, substituting content provided by larger publishers with cheaper alternatives. Even though the channel through which the asymmetry arises differs (we keep the bargaining power fixed and focus on vertical differentiation instead), the outcome aligns with our equilibrium predictions.

Widening the scope of the discussion, this paper relates to the evolving literature on the economics of media markets. While most previous contributions have focused on the mix of content and advertising in media (Anderson and Coate, 2005; Anderson and Gabszewicz, 2006; Peitz and Valletti, 2008; Thomes, 2013; Halmenschlager and Mantovani, 2017), we ignore this dimension altogether. We do so for two reasons: First, while many streaming platforms offer free subscriptions with ads as an alternative to the ad-free “premium” subscriptions, the latter represent an enormous and still growing market.¹⁰ Second, while the literature on media advertisement contraposes content and ads, bringing positive and negative utility to consumers, respectively, we focus on content bias because of the inherent alignment of interests it breaks. Consumers value good content and are willing to pay more for it. While the trade-off between content and ads is intuitive, the fact that the platform has an incentive to penalize high-quality products it is not competing against is not. The paper also relates to the literature on vertical relations and, in particular, to the coexistence of retailers and direct sale channels available to manufacturers, as well as the strategic interaction of a platform with its suppliers when competing against them. Recent work by Aguiar et al. (2023) is closely related to our paper. They analyze the platform’s incentives to include certain types of artists and songs in its playlist to leverage its market power and obtain better licensing deals with major music labels. We diverge from this work in two ways: we propose a theoretical investigation of the platform’s incentives to bias its recommendation system and focus on personalized recommendations.

Most papers consider consumers more or less sensible to prices and distribution channels.¹¹ We distance ourselves from this approach and, rather, follow recent work by Ronayne and Taylor (2022). Their paper studies the role of a competitive channel, like an online e-commerce platform, as an alternative distribution channel available to sellers. The authors focus their attention on the optimal governance structure of the competitive channel, assuming both this channel and the sellers have some captive consumers to extract fees from. By contrast, our market shares emerge endogenously in equilibrium. Captive consumers in our setting would allow the platform to bias more aggressively in equilibrium, a result proxied by the demand expansion we feature in our model: the more consumers stay inactive if not for the platform intermediation, the less constrained the platform is in designing the recommendation system.

¹⁰According to Spotify’s earning report to investors, the platform had 195 million premium subscribers in Q3 of 2022. Available at: https://s29.q4cdn.com/175625835/files/doc_financials/2022/q3/Q3-2022-Shareholder-Deck-FINAL-LOCKED.pdf

¹¹See, for example, Rhee and Park (2000), Chiang et al. (2003), and Kumar and Ruan (2006).

2 Model setup

There are two sellers (music labels), indexed by $j = a, b$, located at the left and right extremes of the $[0, 1]$ Hotelling line. In addition, there is a platform (p) that knows the consumers' location and offers them a personalized bundle of content from the two sellers. By doing so, the platform can better match consumer preferences (Adams and Yellen, 1976; Anderson and Neven, 1989; Bourreau and Gaudin, 2022).¹²

We consider two groups of consumers, informed and uninformed, each uniformly distributed on the Hotelling line, in the market for streamed products. The group of informed consumers has mass $\alpha \in [0, 1]$, whereas the group of uninformed consumers has mass $1 - \alpha$. The information they possess (or don't) refers to the existence and location of the firms operating in the market. Moreover, informed consumers know *ex-ante* their location on the line and the exact location of the sellers. By contrast, uninformed consumers only know about the streaming service and discover the two music labels if they are available on the platform. The two sellers produce one good each, a and b , respectively. We refer to them as the *pure bundles*, entirely made of contents produced *in-house*. These can be thought of as the albums produced by the two music labels. On the other hand, we define *mixed bundles* as the personalized goods that consumers can access via the streaming service, like a playlist that includes content produced by both sellers.

We indicate the location of consumers on the unit line with x . Then, we use $\lambda(x) + \varepsilon(x) \in [0, 1]$ to identify the share of Content a consumed by the user located at x if they join the platform service. In particular, $\lambda(x)$ is the preferred share consumers would choose to maximize utility, whereas $\varepsilon(x)$ is the personalized bias on the recommendation system imposed by the platform. In other words, $\varepsilon(x)$ is the extra share of Content a offered to each consumer by the platform's algorithm. By contrast, $1 - \lambda(x) - \varepsilon(x)$ represents the share of Content b offered to the same consumer.

Consumers purchase exactly one unit of the final good — either a pure bundle or the recommended mixed bundles. We use p_a and p_b to define the price of the pure bundles, paid directly to the music labels, and f to identify the subscription fee paid by consumers to access the platform's service instead.

The platform pays royalties (r_j) to the music labels per share of their content offered to consumers. We assume that the music labels charge a royalty rate equal to the market price: $r_j = p_j$. This assumption allows us to ignore any direct bargaining between sellers and the platform and any effect of the eventual differences in bargaining power. We assume that sellers have full bargaining power in the royalty-setting stage and, therefore, always select the highest rate possible given their price in the external market.¹³

¹²This framework has been used to study advertisements when consumers mix their consumption (Gal-Or and Dukes, 2003) and, more recently, to study welfare implications of different pricing structures (Hoernig and Valletti, 2007, 2011; Döpfer and Rasch, 2022).

¹³At the end of Section 3, we discuss the implications of relaxing this assumption. We show that while both sellers would have the incentive to unilaterally deviate by reducing their royalty rate and keep their price high, both doing so generates a cycle in which both sellers reduce their royalty rate by the same amount, leaving the bias unaffected but allowing the platform to lower its subscription fee and capture more consumers. By keeping the assumption that $r_j = p_j$, we select a feasible equilibrium that also corresponds to the best outcome for the sellers.

The utility functions of a consumer i , located in x_i , who purchases from a , b , or joins the platform, respectively, can be written as:

$$\begin{aligned} U_{i,a} &= V_a - p_a - tx_i^2, \\ U_{i,b} &= V_b - p_b - t(1 - x_i)^2, \\ U_{i,p} &= (\lambda(x_i) + \varepsilon(x_i))V_a + (1 - \lambda(x_i) - \varepsilon(x_i))V_b - f - t(x_i - (1 - \lambda(x_i) - \varepsilon(x_i)))^2, \end{aligned}$$

where $V_j = v + v_j$ is the intrinsic quality of the pure bundles — common to all consumers — and is composed of a common parameter $v > 0$, that we assume to be high enough to guarantee full coverage in the market (i.e., $v > 3t/2$) and a music label-specific parameter $v_j \geq 0$. Without loss of generality, we assume b to produce the weakly higher-quality content: $v_b \geq v_a = 0$. Finally, the parameter $t > 0$ represents the transportation costs that multiply the utility loss from taste mismatch. For tractability, we assume $v_b < 2t/3$ always holds.

The timing of the game is as follows:¹⁴

1. The platform chooses the level of bias of the recommendation system $(\varepsilon(x_i), \forall x_i)$ and commits to implementing it.¹⁵
2. Sellers observe the recommendation policy and decide whether to join the platform and serve both informed and uninformed consumers or stay out and compete for informed consumers only in a standard Hotelling setting.
3. Upon observing the entry decision and the quality attributes of the two contents, at Stage 3, the two sellers and the platform set the prices for the pure bundles and the streaming service (p_a , p_b , and f).
4. Given the prices and the recommendation system, consumers make their consumption decisions and profits are realized. The share α of informed consumers know their location on the Hotelling line and the sellers' locations. By contrast, the $1 - \alpha$ uninformed consumers only know that a platform exists. We assume that all consumers can sample the platform for free before subscribing.¹⁶

During the free sample period, uninformed consumers learn about the firms' locations and their own preferences. If the sellers decide not to join the platform at Stage 1, uninformed consumers do not learn anything and make no purchase. Our solution concept is Subgame Perfect Nash Equilibrium (SPNE). We solve the game by backward induction.¹⁷

¹⁴In Subsection 3.2, we augment the model with an additional Stage 0 in which either one or both sellers highly invest in innovation to generate $v_j \geq 0$.

¹⁵Once again, the assumption that the platform can commit to a certain level of bias is motivated by the evidence that platforms promise higher exposure to music labels in exchange for lower royalty rates and by fear of possible legal repercussions if the platform conditions bias on royalty rates.

¹⁶Many real-world streaming platforms, including Spotify, offer free trials to consumers. This assumption, therefore, matches the kind of platform we aim to model.

¹⁷For timing our model, we follow Fletcher et al. (2023): The platform commits to a recommendation system before prices are set and all agents are aware of the implied potential bias in equilibrium. See the testimony of Stephen McBride, Docket No. 14-CRB-0001-WR, regarding royalties-for-exposure deals between the online radio company Pandora and indie-label coalition Merlin. The testimony suggests that the platform might want to commit to its recommendation system regardless of fear of legal repercussions; if they were to condition the recommendation system on royalties, they would be more vulnerable to legal action being taken against them.

3 Content quality and recommendation bias

The music industry is both vertically and horizontally differentiated. Music labels experiment and research new ways of expressing their art. In other words, they innovate. Music labels backed by major recording companies generally have more resources than comparable independent music labels to produce better products — e.g., in terms of sound quality or music production — even when producing Content belonging to the same “genre”. Thus, we can assume that music labels compete with products with varying quality levels. Such a quality differential is assumed to be a primitive of the model. Afterward, we endogenize the choice of costly investment in quality to study which distortions, if any, the platform’s intervention leads to.

3.1 Exogenous quality differential

Since we assume $V_b = v + v_b \geq v = V_a$ with v_b exogeneously given, the utility functions of consumers purchasing from a , b , or joining the platform become:

$$\begin{aligned} U_{i,a} &= v - p_a - tx_i^2, \\ U_{i,b} &= v + v_b - p_b - t(1 - x_i)^2, \\ U_{i,p} &= v + (1 - \lambda(x_i) - \varepsilon(x_i))v_b - f - t(x_i - (1 - \lambda(x_i) - \varepsilon(x_i)))^2. \end{aligned}$$

Recall that $\lambda(x_i)$ indicates the optimal share of Content a in a consumer i ’s individual mix, whereas $\varepsilon(x_i)$ is the personalized bias introduced by the platform. As standard in these models, we derive the locations of indifferent consumers by equating the utility functions they obtain by choosing among the three options. These locations do not depend on the bias, as the personalized bias optimally selected for the indifferent consumers is $\varepsilon(x_{ap}) = \varepsilon(x_{pb}) = 0$. We write:

$$\begin{aligned} x_{ap} &= \frac{f - p_a + (t(1 - \lambda(x_{ap}) - v_b)(1 - \lambda(x_{ap})))}{2t(1 - \lambda(x_{ap}))} &\implies U_{i,a} &= U_{i,p}, \\ x_{pb} &= \frac{p_b - f + (t(2 - \lambda(x_{pb}) - v_b))\lambda(x_{pb})}{2t\lambda(x_{pb})} &\implies U_{i,p} &= U_{i,b}, \\ x_{ab} &= \frac{p_b - p_a + t - v_b}{2t} &\implies U_{i,a} &= U_{i,b}. \end{aligned}$$

We adopt the following notation: x_{jk} , which indicates a consumer who is indifferent between buying from firm j and firm k , with $j, k = a, b, p$ and $k \neq j$. Notice that the location of the consumer who is indifferent between the two pure bundles (a and b) must lie between the consumers indifferent between purchasing either the pure Bundle Goods a and b and joining the platform. In the analysis, we use x_{ab} mainly as a reference point.¹⁸

Notice that $\varepsilon(x_i)$ does not enter the location of the indifferent consumers. This is because the platform knows consumers’ locations and can offer them a personalized recommendation system. Indifferent consumers would change their consumption choices if subject to a bias that lowers their utility. Hence, the platform designs its algorithm to increase the bias in the distance

¹⁸The assumption $v > 3t/2$ is sufficient to ensure $U(x_{ab}) > 0$.

between the consumers and their preferred music labels. This bias is personalized and bound by consumer participation constraints. We define:

$$\bar{\varepsilon}(x_i) = \{\varepsilon(x_i) \in [0, 1 - \lambda(x_i)) \text{ s.t. } U_{i,p} |_{\lambda=\lambda(x_i)+\varepsilon(x_i)} = \max\{U_{i,a}, U_{i,b}\} \forall x_i \in (x_{ap}, x_{pb})\},$$

as the maximum bias level a consumer i located in x_i is willing to accept before leaving the platform and purchasing the pure bundle from their preferred music label. Through $\bar{\varepsilon}(x_i)$, we can see that the consumers in x_{ap} and x_{pb} would not accept any bias because for them $U_a = U_p(\lambda(x_{ap}))$ and $U_b = U_p(\lambda(x_{pb}))$, respectively. Formally, $\varepsilon(x_{ap}) = \varepsilon(x_{pb}) = 0$.

We can now derive the personalized recommendation system set by the platform. Intuitively, the platform aims to maximize the consumption of the streaming service. To do so, it offers the *efficient bundle* to indifferent consumers. We define the *efficient bundle* as the composite good that would be chosen by a consumer so that, for any prices p_a , p_b , and f , they would get the highest possible utility. By definition, the efficient bundle is not biased by the platform's recommendation system ($\varepsilon(x_i) = 0, \forall x_i$):

$$\lambda(x_i) = \arg \max_{\lambda \in (0,1)} (U_{i,s}) = 1 - x_i - \frac{v_b}{2t}.$$

Using this consumption choice, we update the indifferent consumers' location as:

$$x_{ap} = \sqrt{\frac{f - p_a}{t}} - \frac{v_b}{2t}, \quad x_{pb} = 1 - \sqrt{\frac{f - p_b}{t}} - \frac{v_b}{2t}. \quad (1)$$

The indifferent consumers' location allows us to compute the demand faced by each agent of the model. The platform's demand is given by $D_p = x_{pb} - x_{ap}$. Moreover, the demands of the two sellers change because of the different content proportions in the new biased bundles. This can be expressed as:

$$D_a = x_{ap} + \int_{x_{ap}}^{x_{pb}} (\lambda(x_i) + \varepsilon(x_i)) dx, \quad (2)$$

$$D_b = 1 - x_{pb} + \int_{x_{ap}}^{x_{pb}} (1 - \lambda(x_i) - \varepsilon(x_i)) dx. \quad (3)$$

Music Label b faces a demand that is decreasing in intensity of the bias. By contrast, Music Label a faces an increased demand because of the favorable bias. The variations in the music labels' demands are known before prices are set; thus, they affect the equilibrium prices of both the platform and the sellers. Music Label b is expected to lower its price in response to the decreased demand, whereas Music Label a would likely do the opposite because of the increased demand. The two prices would converge toward a common value if the bias is sufficiently intense. Because the consumers demand more content from the high-quality seller, lowering its price is in the platform's interest, absent any constraint on the seller's participation.

Because consumers have different tastes, the level of bias is personalized. Hence, the platform cares that the participation constraint of each consumer is satisfied. The platform ensures that it does not exceed the sum of the participation constraints of all consumers when deciding to shift the total mass of demand from one seller to the other (i.e., the total bias, see Figure 1).

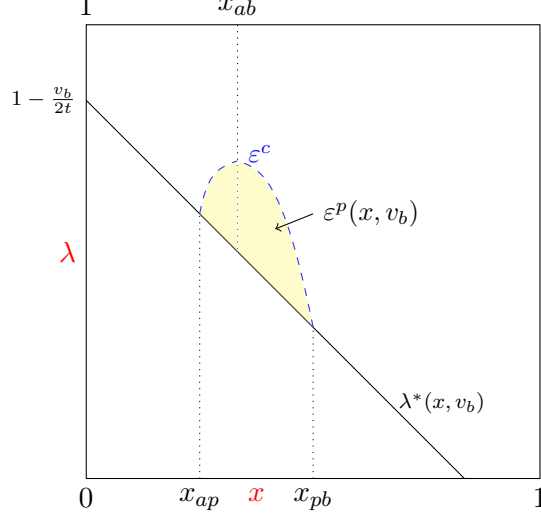


Figure 1: *Personalized biased bundle. The area in yellow is the total demand the platform can shift from Music Label b to Music Label a without losing users.*

Formally:

Condition 1. *Given $v_b > 0$, the aggregate personalized bias imposed by the platform to users of the streaming service can be identified by a general mass of bias ε^p , such that*

$$\varepsilon^p \equiv \int_{x_{ap}}^{x_{pb}} \varepsilon^p(x_i) dx \leq \int_{x_{ap}}^{x_{pb}} \bar{\varepsilon}(x_i) dx \equiv \varepsilon^c,$$

where the superscripts p, c indicate the total bias selected by the platform and the maximum bias that satisfies the consumer participation constraint, respectively; $\varepsilon^p(x_i)$ represents the individual level of bias the platform designs.

To solve the problem, and because what matters to the platform and sellers is the total mass of consumption shifted from one seller to the other, we forego solving the optimal individual level of bias for each consumer who joins the platform. Instead, we consider the total mass ε^p , noting that it must always be compatible with the total participation constraints of all buyers combined. This ensures a consistent solution while maintaining the problem tractable.

The new recommendation system can be written as $\int_{x_{ap}}^{x_{pb}} \lambda(x_i) dx + \varepsilon^p$. We adjust the profit functions accordingly:

$$\pi_p = f(x_{pb} - x_{ap}) - p_a \left(\int_{x_{ap}}^{x_{pb}} \lambda(x_i) dx + \varepsilon^p \right) - p_b \left(\int_{x_{ap}}^{x_{pb}} (1 - \lambda(x_i)) dx - \varepsilon^p \right), \quad (4)$$

$$\pi_a = p_a \left(x_{ap} + \int_{x_{ap}}^{x_{pb}} \lambda(x_i) dx + \varepsilon^p \right), \quad (5)$$

$$\pi_b = p_b \left(1 - x_{pb} + \int_{x_{ap}}^{x_{pb}} (1 - \lambda(x_i)) dx - \varepsilon^p \right). \quad (6)$$

Then, from the system of first-order conditions, we derive the profit-maximizing prices:

$$p_a(\varepsilon^p) = t - \frac{v_b}{3} + \frac{2\varepsilon^p t}{3}, \quad p_b(\varepsilon^p) = t + \frac{v_b}{3} - \frac{2\varepsilon^p t}{3}, \quad f(\varepsilon^p) = \frac{10t}{9} + \frac{v_b^2}{4t} - \varepsilon^p(v_b - \varepsilon^p t). \quad (7)$$

Therefore:

Lemma 1. *Consider the case in which the platform offers a biased mix $\lambda(x, v_b) + \varepsilon^p$ to the consumers. Then, the Stage 2 equilibrium prices are as derived in (7) and the profits of the music labels and the platform, respectively, are given by*

$$\pi_a(\varepsilon^p) = \frac{(t(3 + 2\varepsilon^p) - v_b)^2}{18t}, \quad \pi_b(\varepsilon^p) = \frac{(t(3 - 2\varepsilon^p) + v_b)^2}{18t},$$

$$\pi_p(\varepsilon^p) = \frac{t + 3\varepsilon^p(v_b - t\varepsilon^p)}{27} - \frac{v_b^2}{36t},$$

and the indifferent consumers are located in:

$$x_{ap} = \frac{1}{3} - \varepsilon^p, \quad x_{pb} = \frac{2}{3} - \varepsilon^p, \quad x_{ab} = \frac{1}{2} - \frac{v_b}{6t} - \frac{2\varepsilon^p}{3}.$$

Proof. See the Appendix. ■

The bias affects the two sellers in opposite ways. Seller *a* benefits from the platform selecting a biased mix, as it allows it to sell more of its content to the platform's subscribers, mitigating the quality gap. The increase in the demand for the content of Seller *a* exerts positive pressure on the price p_a . The indifferent consumer shifts to the left, but the price effect and the larger share of Content *a* in the biased mix more than compensate for the reduced demand on the direct channel.

On the other hand, Seller *b* suffers from the recommendation bias. Consumers are exposed to a lower-than-optimal level of Content *b* on the platform. To compensate for this loss, Seller *b* lowers the price p_b , inducing more consumers to purchase its pure bundle. However, the negative price effect and the reduced exposure of Content *b* in the mixed bundle dominate the demand expansion on the direct channel.

The platform does not lose demand but reshuffles its cost function. A positive bias $\varepsilon^p > 0$ makes sense provided that $p_b > p_a$, which in this case requires $v_b > 2\varepsilon^p t$. If that were not the case, the recommendation bias would backfire: if $\varepsilon^p > 0$ such that $p_a > p_b$, then the platform would find itself shifting demand toward the most expensive content, a non-optimal outcome.

The bias is set before the game starts and the platform commits to that level. Hence, once decided, it cannot be modified to adjust for the new price ordering. Moreover, the bias cannot exceed the maximum value described in Condition 1; the platform anticipates the bias's effects on the entry decision of consumers and on pricing and sets it consistently with their participation constraints.

3.1.1 Sellers' participation decision

At Stage 2, the sellers must decide whether to join the platform. If at least one of the sellers decides not to enter, they only compete with each other and with only $\alpha \in [0, 1]$ active consumers. By contrast, if they both decide to list their products on the platform, the remaining $(1 - \alpha)$ uninformed consumers join the market as well. Uninformed consumers learn of the existence of the sellers or their relative position only if the platform is active, which can happen only if the platform manages to attract both sellers.¹⁹

In all sub-games where at least one of the sellers decides not to join the platform, only the informed consumers are active. With no streaming service available, consumers cannot mix their consumption and are therefore limited to purchasing a pure bundle from either a or b . In these sub-games, sellers compete in a standard Hotelling setting. Given $v_b \geq 0$ and $\alpha \in [0, 1]$, equilibrium prices and profits when the platform is inactive are:

$$\begin{aligned} p_a^{out} &= t - \frac{v_b}{3}, & p_b^{out} &= t + \frac{v_b}{3}, \\ \pi_a^{out} &= \alpha \frac{(3t - v_b)^2}{18t}, & \pi_b^{out} &= \alpha \frac{(3t + v_b)^2}{18t}, \end{aligned} \quad (8)$$

where the superscript out indicates the case where only consumption outside the platform is possible.

When Seller $j = a, b$ considers whether to join the platform, it compares profit π_j^{out} and π_j anticipating equilibrium pricing and any consumption bias the platform might introduce. Notice that, compared to Seller a , Seller b has the better outside option if the platform is inactive. Moreover, Seller b would be penalized if the platform biased consumption. It is therefore sufficient to consider Seller b 's participation decision to determine whether the platform can be active in equilibrium, which depends on the share of informed consumers, α , and the quality difference, v_b . The platform's ability to bias consumption is also limited in that it must induce both sellers and buyers to join. In other words, the equilibrium bias the platform can design is bound by two constraints: the consumer participation constraints, addressed above, and the seller participation constraints.

Intuitively, the latter becomes stricter as α increases. If there are many informed consumers, the high-quality seller has stronger leverage at the entry stage. Suppose there are no uninformed consumers (i.e., $\alpha = 1$). In that case, the platform's optimal bias policy would hurt Music Label b . Moreover, b anticipates that joining the platform does not expose its product to more consumers. Then, b would rationally choose not to join a platform that commits to any positive level of bias. Therefore, the platform would reduce its optimal bias to zero to induce both sellers to join. On the other hand, suppose that $\alpha = 0$. If b does not join the platform, sales in the direct market become impossible because there are no consumers who are aware of b 's existence. Regardless of how biased the recommendation system is in favor of its rival, b would always choose to join the platform. This implies that, when establishing its bias policy, the platform is only constrained by the consumer's participation decision.

¹⁹Intuitively, because the platform's core service is to allow consumers to mix their streams/purchases. The condition needed for the platform to be active is that at least two goods (or, in this case, both goods) are listed.

Formally, we can distinguish a threshold for α as a function of v_b and the chosen bias ε^p that discerns when it is profitable to join the platform and when it is profitable to operate only on the direct market:

$$\alpha^* = \{\alpha \in [0, 1] \mid \text{s.t. } \pi_b(\varepsilon^p) = \alpha \pi_b^{\text{out}}\} \iff \alpha^* = \frac{\pi_b(\varepsilon^p)}{\pi_b^{\text{out}}} = \frac{((3 - 2\varepsilon)t + v_b)^2}{(3t + v_b)^2}.$$

Because the platform is inactive unless both sellers join the streaming service, it must choose a bias such that:

$$\varepsilon \leq \varepsilon^s = \frac{(3t + v_b)(1 - \sqrt{\alpha})}{2t}, \quad (9)$$

where the superscript s indicates the *seller's constraint*.

3.1.2 Equilibrium bias

We can finally proceed backward to Stage 1 and determine the equilibrium level of bias that the platform includes in its recommendation system. According to the above analysis, the platform's problem can be written as:

$$\begin{aligned} \max_{\varepsilon} \quad & \pi_p(\varepsilon) = \frac{t + 3\varepsilon(v_b - t\varepsilon)}{27} - \frac{v_b^2}{36t}, \\ \text{subject to} \quad & \varepsilon < \min\{\varepsilon^c, \varepsilon^s\}. \end{aligned}$$

The unconstrained maximization leads to $\varepsilon^p = \frac{v_b}{2t}$. Using the prices in (7) to determine the maximum bias consumers are willing to accept before leaving the platform, we obtain:

$$\begin{aligned} \varepsilon^c = \int_{x_{ap}}^{x_{ab}} & \frac{2t + 3v_b - \sqrt{72t^2x^2 - 4t^2 - 18tv_b^2 + 72tv_b x - 12tv_b + 27v_b^2}}{12t} dx + \\ & \int_{x_{ab}}^{x_{bp}} \frac{-2t + 3v_b - \sqrt{72t^2x^2 + t^2(68 - 144x) + 72tv_b x - 6tv_b(3v_b + 10) + 27v_b^2}}{12t} dx, \end{aligned}$$

where the superscript c indicates the *consumer participation constraint*.

Based on the location of indifferent consumers in Proposition 1, tedious calculations reveal that ε^c decreases in v_b . This result is intuitive: when the quality gap increases, the number of consumers on the platform who prefer Content a over Content b decreases (i.e., $x_{ab} - x_{ap}$ decreases in v_b). By contrast, the number of consumers who prefer Content b increases. As the quality gap v_b rises, a decreasing number of people are willing to accept a recommendation that promotes the low-quality product and hides the high-quality one.

The negative relation between v_b and ε^c implies that the equilibrium unconstrained bias and the consumers' constraint react differently to an increase in the quality gap: the former increases in v_b . A larger quality differential results in a stronger effect on costs than on prices, and greater incentive of the platform to bias consumption in equilibrium. On the other hand, consumers benefit from the high-quality product offered by b and become less willing to accept a biased bundle that contains less of it as the quality gap increases.

Altogether, the platform must design its recommendation system to attract sellers and retain

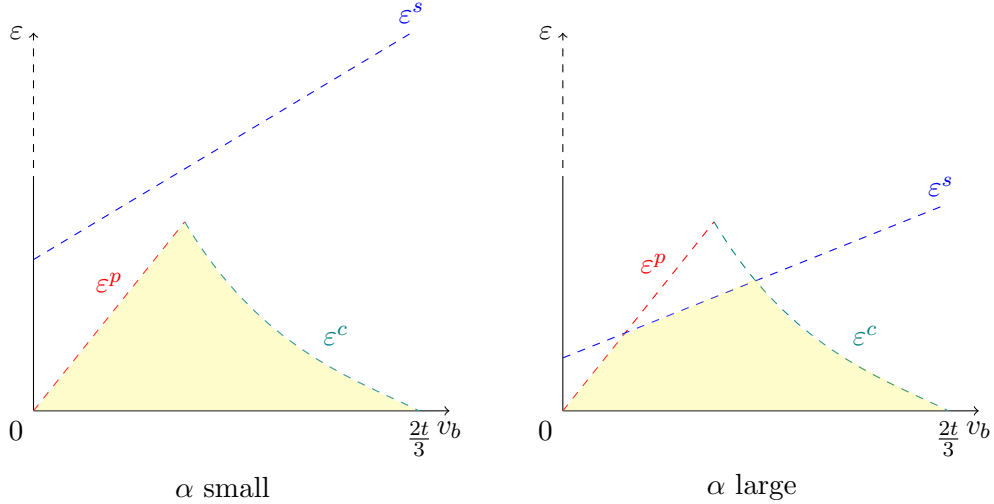


Figure 2: The constrained equilibrium level of the recommendation bias designed by the platform. When the share of informed consumers is sufficiently small (left panel), the optimal bias (ε^p) is bound by the consumer participation constraint (ε^c). When it begins to take effect, the platform needs to reduce the bias as quality increases because fewer users are willing to substitute Content a with the high-quality product (b). Otherwise (right panel), the platform must consider an additional constraint (ε^s) and needs to secure sellers' participation in the streaming service.

consumers. In other words, the total bias can never exceed ε^c nor ε^s . The resulting equilibrium is illustrated in Figure 2 and can be summarized as:

Proposition 1. *In equilibrium, the platform sets a bias such that all sellers join the platform and the market is fully covered. Moreover, the equilibrium bias is:*

$$\varepsilon^* = \min\{\varepsilon^p, \varepsilon^c, \varepsilon^s\}.$$

Proof. See the Appendix. ■

In choosing the bias intensity, the platform trades off cost minimization and competition intensity. Biasing consumption away from b leads the platform to optimally reduce p_b ; thus, the platform must update its optimal price downwards because the better product has a stronger competitive effect than the worse one. Hence, the platform aims to reduce the bias to set higher fees for consumers and raise the bias to minimize costs. This trade-off is optimally resolved at ε^p . This dynamic echoes the findings of Arya et al. (2007): because sellers compete with the platform on the margin, they set lower prices than they would if their only revenue stream was through the platform. These prices translate to lower operational costs for the streaming platform even without the adjustment that follows the strategic use of the recommendation system discussed here.

The unconstrained level of bias ε^p is such that, in equilibrium, the two sellers behave as Hotelling rivals with homogeneous goods:

Corollary 1. *In the unconstrained equilibrium, the sellers and the platform set prices*

$$p_a^* = p_b^* = t, \quad f^* = \frac{10t}{9}.$$

Moreover, their payoffs are:

$$\pi_a^* = \pi_b^* = \frac{t}{2}, \quad \pi_p^* = \frac{t}{27}.$$

Under normal conditions, the quality differential would increase the demand for the high-quality product and decrease, by the same amount, the demand for the low-quality one. Such asymmetry incentivizes the seller of the high-quality product to increase its price. If price increase materializes, the platform would have to pay more for the content consumers stream the most, with negative effects on its costs.

Anticipating this, the platform shifts demand from high-quality to low-quality content, compensating for the demand asymmetry caused by the vertical differentiation. Absent any demand asymmetry, the two sellers compete as if they were symmetric, and standard Hotelling payoffs ensue.

The results have important implications for determining the firms' incentives to invest in quality, which we present in the next section. Here, we anticipate that any constraint on the maximum amount of bias chosen by the platform ensures the producer of high-quality content a demand premium because of the quality differential, which translates into a higher price, a better payoff, and, a stronger incentive to invest in quality. By contrast, if no constraint limits the platform, the ability to bias recommendations disincentivizes sellers to invest in quality.

3.2 Endogenous investment in quality

While there are good reasons to believe that quality in music derives from the innate talents of the artists and not from strategic considerations, investments in better instrumentation, team training, and equipment help develop a good product. In this section, we endogenize the sellers' choice to invest in quality. To maintain the direction of the asymmetry studied above, we assume Seller b is always more efficient than Seller a . For tractability, we ignore the effect of consumer participation constraints on equilibrium quality levels when solving for the equilibrium investment in quality.²⁰

The timing of the interaction between sellers and the platform remains unchanged but augmented by an earlier stage (we refer to it as Stage 0) in which sellers independently and simultaneously select v_a and v_b . To model this investment, we consider the standard convex cost function $I(v_j) = \phi_j v_j^2$. In our simplified setting, a firm chooses how much to invest in R&D and these investments uniquely determine the quality of the product sold in the market, $v + v_j$. Therefore, a firm maximizes profits by either choosing its investment amount, $I(v_j)$, or its degree of innovativeness, v_j . For this reason, with a slight abuse of terminology, we also refer to v_j as the level of investment in R&D chosen by seller j .

²⁰This exercise aims to qualitatively evaluate the effect of recommendation bias on innovation investments. As discussed above, both consumer and seller participation constraints limit the bias's role for a high enough quality differential. In what follows, we compare the results of the unconstrained scenario and the scenario where the seller constraint binds the platform's choice. The consumer participation constraint leads to the same qualitative outcome.

The updated seller's profits are given by:

$$\pi_a^{eq} = D_a p_a - \phi_a v_a^2, \quad \pi_b^{eq} = D_b p_b - \phi_b v_b^2, \quad (10)$$

where superscript eq stands for “endogenous quality” and $\phi_b < 1 = \phi_a$ captures the higher efficiency of Seller b compared to a . We update the formulas detailed in Lemma 1 by including the cost functions for the two firms and the cost differential $v_b - v_a$:

$$\pi_a^{eq}(\varepsilon^p) = \frac{(t(3 + 2\varepsilon^p) - (v_b - v_a))^2}{18t} - v_a^2, \quad \pi_b^{eq}(\varepsilon^p) = \frac{(t(3 - 2\varepsilon^p) + (v_b - v_a))^2}{18t} - \phi v_b^2,$$

$$\pi_p^{eq}(\varepsilon^p) = \frac{t + 3\varepsilon^p((v_b - v_a) - t\varepsilon^p)}{27} - \frac{(v_b - v_a)^2}{36t},$$

We consistently obtain the unconstrained bias as a function of the quality differential from the first-order conditions of the platform's maximization problem: $\varepsilon^p = \frac{(v_b - v_a)}{2t}$.

We use the platform's equilibrium bias to obtain the equilibrium investment levels by plugging the bias in the equations for $\pi_a^{eq}(\varepsilon^p)$ and $\pi_b^{eq}(\varepsilon^p)$. Then, we proceed backward to Stage 0, when innovation investments are decided. From Corollary 1, we know that the sellers' payoffs at Stage 1 are equivalent to the standard Hotelling profits with homogeneous goods. Thus, in the unconstrained equilibrium, the payoffs of sellers and platforms at Stage 1 are updated to:

$$\pi_a^{eq} = \frac{t}{2} - v_a^2, \quad \pi_b^{eq} = \frac{t}{2} - \phi_b v_b^2, \quad \pi_p^{eq} = \frac{t}{27}.$$

Proposition 2 illustrates the sellers' incentives to invest in content quality in the unconstrained equilibrium and compares them to the incentives in the scenario where the platform is inactive, only a fraction α of consumers is active, and sellers compete in standard Hotelling fashion.

Proposition 2. *Suppose both sellers invest in quality, and Seller b is more efficient than Seller a by a factor ϕ_b . The unconstrained equilibrium levels of investment in quality are:*

$$v_a^{eq} = 0, \quad v_b^{eq} = 0.$$

By contrast, the equilibrium levels of investment outside the platform are:

$$v_a^{eq,out} = \frac{\alpha(9t\phi_b - \alpha)}{3((18t - \alpha)\phi_b - \alpha)}, \quad v_b^{eq,out} = \frac{\alpha(9t - \alpha)}{3((18t - \alpha)\phi_b - \alpha)}.$$

Proof. See the Appendix. ■

Proposition 2 presents the second main result of the paper: the presence of a recommendation bias mitigates the sellers' incentives to invest in content quality. This is true for the high-quality content producer that suffers from the traditional hold-up problem. Being unable to fully appropriate the benefits of selling a better product, its incentives to vertically differentiate drop, and so do the equilibrium level of investments. This is also true for the low-quality seller. Given $v_b^{eq} = 0$, any positive investment in quality would reverse the bias's sign with the

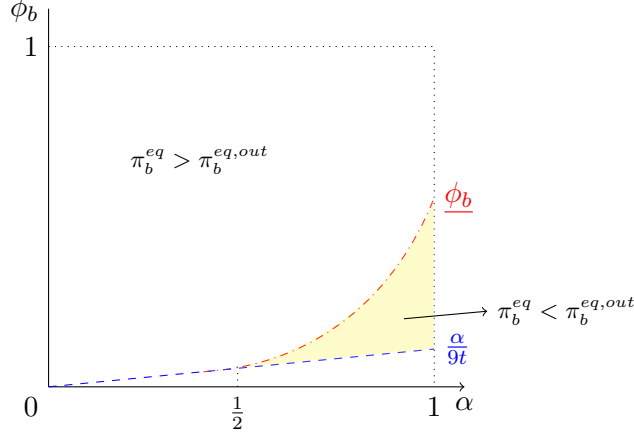


Figure 3: Corollary 2. The highlighted area where $\phi_b < \underline{\phi}_b$ indicates the parameter region where the unconstrained bias is not optimal from the platform's perspective, as it discourages sellers' participation.

consequent loss of demand. In other words, the low-quality seller benefits from a *low-quality premium* that activates whenever the high-quality rival invests in innovation and raises its price accordingly.

Overall, recommendation biases distort and suppress incentives to compete based on quality by rewarding the low-quality producer and punishing the high-quality one, leading to no R&D investment in equilibrium. A natural question arises at this point: do sellers always accept the unconstrained bias when innovation costs are considered?

Corollary 2. *The high-quality content producer prefers joining the streaming platform regardless of the bias if and only if $\phi_b > \underline{\phi}_b$.*

Proof. See the Appendix. ■

Corollary 2 complements Proposition 2 and expands the result derived in Expression 9, which shows how the seller's participation decision is updated once the investment costs are considered. If the two firms are sufficiently similar regarding innovation efficiency (ϕ_b is close to 1), then both sellers decide to withhold investments and join the streaming platform with an initial homogeneous quality $v > 0$. By contrast, if the efficiency gap is large (i.e., $\phi_b < \underline{\phi}_b$), the (potential) high-quality seller earns higher profits by staying off the platform and investing a positive amount $v_b^{eq,out}$. Condition $\phi_b < \underline{\phi}_b$ is feasible only if the demand expansion provided by the streaming platform is small ($\alpha > \frac{1}{2}$, see Figure 3).

If these conditions are satisfied, the platform is constrained in its ability to introduce bias, and investment in the platform takes place. Both conditions seem particularly strict when translated into real-world examples of the music industry or news media market.

Corollary 3. *Assume $\phi_b < \underline{\phi}_b$. The platform sets a constrained bias:*

$$\varepsilon^s = \frac{3(1 - \sqrt{\alpha})\phi_b(9t - \alpha)}{18t\phi_b - \alpha(\phi_b + 1)},$$

and the sellers invest $v_a = v_a^{eq,out}$ and $v_b = v_b^{eq,out}$ as they would do had they decided not to join the platform.

Proof. See the Appendix. ■

Corollary 3 complements our primary result. It shows that when the seller participation constraint is binding, the platform needs to provide the high-quality seller a sufficient payoff to be indifferent between joining the platform or going with the exit option — which implies the platform would not be active in equilibrium. Hence, the platform sets a positive but not too strong bias that restores the incentives to invest in quality the sellers would have outside its ecosystem. However, those incentives are lower than the ones the sellers would have had on a platform without bias. Thus, joining the streaming service does not allow the sellers to monetize the demand expansion produced by the platform ($1 - \alpha$ uninformed users).

Things are less clear when the consumer participation constraint is the one binding. In such cases, the platform needs to guarantee that each user gets at least as much utility as they would by directly buying the product from either seller. By imposing an upper bound on the equilibrium level of bias, consumers ensure that enough demand is driven toward the high-quality product, restoring some of the incentives to invest in quality. To analytically derive the conditions that apply to this scenario, as well as the equilibrium results, proves cumbersome. Thus, we limit our analysis to what can be obtained logically.

If the platform sets a bias lower than the unconstrained solution, the consequent shift in demand from the high-quality to the low-quality product does not fully compensate for the quality premium, as in the case with unconstrained bias. Thus, when the consumer participation constraint is binding, sellers are incentivized to invest in improving quality and do so. The equilibrium level of the investment will be lower than if any bias is absent, but still positive. The analysis in Section 3.1 revealed that a higher quality differential results in a tougher consumer participation constraint. Therefore, the constraint is more likely to bind when the efficiency gap between the two firms is pronounced (i.e., ϕ_b is low). Otherwise, the two goods would be perceived as qualitatively similar and the platform would be freer to set an unconstrained bias, discouraging innovation altogether.

3.3 Robustness: royalties and prices

In the analysis above we assumed that sellers could not separate prices and royalties, limiting their ability to best react to the bias. In reality, royalties and market prices do not coincide. Here, we briefly show how relaxing this assumption exacerbates the negative effects of biased recommendation on the sellers' incentives to invest in innovation. To do so, we modify the game's timing and allow the seller to strategically set the royalty rates just before the price competition stage.

At the royalty-setting stage, sellers observe the quality gap between the two pure bundle goods (a and b) and the bias ε . Each seller knows that, by asking a lower royalty rate r than its rival, it could benefit from the expansion in demand induced by the favorable bias. However, the seller also anticipates that a lower royalty rate implies a reduction in marginal costs for

its direct rival: the platform (if the cost of content decreases, the platform can price more aggressively and expand its direct demand).

Despite the bias expanding the demand for the cheapest content, it does so *on the platform*. In other words, the seller who lowers the royalty rate the most faces both an expansion of the total demand and a reduction of the direct demand (i.e., the number of consumers who buy the product directly at the market price). On the other hand, the share of consumers who purchase the content on the platform increases. Because the seller cannot efficiently monetize the on-platform demand (the royalty rate is forced down to *win the bias*), the overall effect is ambiguous. Thus, the standard trade-off between the *extensive margin* (total demand) and the *intensive margin* (net profit per sale) emerges.

Our model suggests that, for small negative variations in the royalty rate, the seller gains a discrete positive amount of extra demand, which makes this trade-off worthwhile. However, this is true for both firms. The competition for the bias that ensues is likely to drive the royalty rate down without altering the size of the bias. Eventually, this unraveling mechanism generates a prisoner dilemma where undercutting the rival constitutes a dominant strategy for sellers, but the ensuing equilibrium is largely inefficient. Furthermore, the competition for bias can potentially break the equilibrium and lead sellers to decide against joining the platform ecosystem because of their binding participation constraints. Such an outcome would be detrimental to consumers, who benefit from the possibility of mixing the content from the two sellers.

To prevent this scenario, the platform could — as observed in the aforementioned Spotify’s “Discovery Mode” — set a royalty rate it is willing to accept to promote certain content. Such a rate, provided that it satisfies sellers’ individual rationality, lowers the average cost for content. Our analysis captures the effects of such strategies on incentives to innovate: because the competition for bias would mitigate the positive relationship between investment and profit, the incentives to engage in costly R&D would also decrease. This is largely driven by the sellers’ inability to efficiently price the indirect demand generated online, which constitutes a large share of their total demand. In other words, our model represents a lower-bound scenario where firms do not lose their ability to set prices as a consequence of the bias. The emerging bias negatively affecting investments in this generous scenario suggests that the same effect would be stronger in a context where bias induces a price distortion.

4 Extensions

We extend the analysis in several directions. We consider the costly implications of letting consumers produce their favorite mix of content on the platform. We show that this ability tightens the constraint governing the maximum bias consumers are willing to tolerate, but otherwise leaves the analysis unchanged. This implies that if the platform could manipulate the search costs through design choices, it would increase the said cost to force consumers to use its recommendation system.

Furthermore, we consider a different timing of the interaction. Following [Bourreau and Gaudin \(2022\)](#), we assume that the recommendation system is not set up at the beginning, but

instead, after the agents have selected prices. We take a reduced form approach and show that the platform has the incentive to bias more in this case than in the baseline model because agents cannot condition prices on the equilibrium recommendation system.

Finally, we consider a different source of vertical differentiation in the form of asymmetric production costs by the sellers. We demonstrate that because the most efficient seller selects lower prices in equilibrium, the platform biases consumption toward this seller. This suggests that streaming platforms reward efficiency. We argue, however, that this also indicates that streaming platforms penalize experimentation and, instead, incentivize the production of modified content to avoid the added penalty linked to inefficiency.

4.1 On-platform search

Our baseline model relies on the implicit assumption that consumers are passive agents concerning the design of the mix they consume. In other words, we assume that consumers join the streaming service and take the bundle offered by the platform, with no opportunity to modify it.

In the real world this is hardly the case: consumers generally have some freedom in choosing the digital content they consume. Here, we address this dimension of the problem by allowing consumers to incur a fixed cost k to produce their optimal bundle on the platform. Spending this “search cost” allows consumers to prevent the bias from affecting them: if all consumers make this choice, no consumption bias can emerge.

The consumers’ utility when they join the platform becomes:

$$U_{i,p} = \begin{cases} v + (1 - \lambda(x_i) - \varepsilon(x_i))v_b - f - t(x_i - (1 - \lambda(x_i) - \varepsilon(x_i)))^2 & \text{if she does not pay,} \\ v + (1 - \lambda(x_i))v_b - f - k - t(x_i - (1 - \lambda(x_i)))^2 & \text{if she pays } k > 0, \end{cases}$$

where k is the search cost that consumers must incur to avoid the recommendation bias. Considering that the efficient bundle is $\lambda(x_i) = 1 - x_i - \frac{v_b}{2t}$, we can derive the cost level that makes any consumers on the platform indifferent between building the efficient mix and accepting the biased bundle:

$$\bar{k} = t(\varepsilon(x_i))^2.$$

The platform always has an incentive to bias recommendations when a difference in content quality emerges. Thus, the platform would take k as a third implicit constraint when setting up the bias: Suppose the platform did not account for the ability of consumers to generate their bundles. Consumers who anticipate that the individual bias they will face reduces their utility beyond k will join the streaming platform, as they knowing that the search cost is smaller than the opportunity cost of the search. Moreover, by Proposition 1, consuming the recommended Bundle must be slightly preferred to leaving the platform.

In this case, several consumers subjected to biased recommendations by the platform would instead choose their optimal bundle, skewing consumption in favor of the higher-quality content and exacerbating the royalty differential the bias was meant to reduce. Thus, the platform faces a profitable deviation and can construct a bundle such that $t(\varepsilon(x_i))^2 \leq k, \forall x_i$.

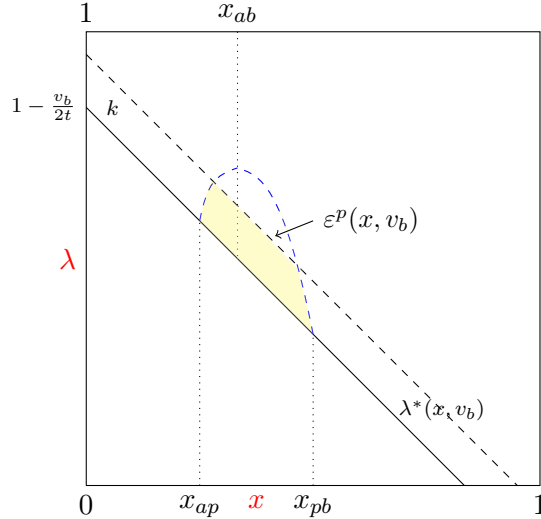


Figure 4: *Personalized biased bundle with search costs k (the area in yellow is the further constrained total demand that the platform can shift from Music Label b to Music Label a without losing users).*

Then, k contributes to the consumer constraint ε^c ; in particular, the highest individual bias any consumer would be willing to withstand becomes:

$$\begin{aligned} \bar{\varepsilon}_s(x_i) &= \{\varepsilon(x_i) \in [0, 1 - \lambda(x_i)] \text{ s.t. } U_{i,p}|_{\lambda=\lambda(x_i)+\varepsilon(x_i)} = \\ &= \max\{U_{i,a}, U_{i,b}, U_{i,p}|_{\varepsilon(x_i)=0} - k\} \forall x_i \in (x_{ap}, x_{pb}), \} \end{aligned}$$

and because consumers near x_{ap} and x_{pb} are willing to tolerate close to no bias, if the constraint implied by search cost k binds, the total bias must be strictly below a uniform bias such that $\varepsilon(x_i) = k, \forall x_i$. Therefore, $\varepsilon^p < (x_{pb} - x_{ap})k$ must apply as Figure 4 illustrates.

However, the ease with which consumers can find their preferred content is part of the platform's design (i.e., its architecture). Thus, if unconstrained, the platform could simply design the search process in a way such that $k = \bar{k}$ for a consumer i , who is willing to accept the largest bias (who is located in x_{ab}). Under our assumption of a monopoly platform, there are no incentives to design the search process in a way that limits its ability to bias the recommendation system and save on costs. However, the search process' efficiency is relevant when competition between platforms is accounted for. This additional dimension of the problem is left to future research efforts.

4.2 Different timing

In our baseline model, we adopt a timing that implies sellers are *ex-ante* aware of the *ex-post* level of bias because the platform announces it in the first stage and commits to it. The question of what would occur were the sellers unaware of the recommendation bias naturally emerges.

In this section, we address this question, assuming that the platform does not commit to a bias level, but firms are aware of the platform's incentives to steer demand toward the cheapest product. Hence, the new timing is as follows: at Stage 1, the sellers anticipate the bias level and, simultaneously with the platform, set the prices. Then, at Stage 2, the platform observes

prices and quality and adjusts its recommendation bias. Finally, consumers make their choices, and payoffs are realized.

According to Proposition 1, equilibrium prices would be asymmetric if there was a difference in the quality of the goods offered and, therefore, it would be profitable for the platform to bias its recommendation system. Moreover, as sellers cannot react after the recommendation bias is chosen, the platform would always steer as much demand as possible, subject to the consumer participation constraint (see the definition of $\bar{\varepsilon}(x_i)$ in Section 3).

The sellers anticipate this incentive and adjust their prices accordingly. The seller of the superior good reacts to the anticipated bias by lowering its price. By contrast, the seller of the inferior good, anticipating the demand expansion following the recommendation bias, is incentivized to increase its price. This adjustment occurs until it stops being profitable.

An equilibrium exists only when the prices do not “cross” (i.e., provided that the most valuable good is also the most expensive one): Assume that sellers anticipate the bias and adjust the prices to such an extent that the inferior content is now as expensive as the superior one. The platform observes the prices and reacts by recommending the cheapest high-quality content to more consumers. However, this reaction goes in the opposite direction of what was anticipated by the sellers, who would like to change their strategies ex-post.

Define the prices chosen by Sellers a and b in anticipation of the total recommendation bias $p_a(\varepsilon)$ and $p_b(\varepsilon)$, respectively. We can state the following:

Lemma 2. *Assume that $p_a|_{\varepsilon=0} < p_b|_{\varepsilon=0}$. Then, if the maximum bias the platform can impose (ε^c) is such that $p_a(\varepsilon^c) < p_b(\varepsilon^c)$, an equilibrium exists in which the platform sets the maximum bias, and the price difference between the two sellers shrinks. Otherwise, if the bias is such that $p_a(\varepsilon^c) \geq p_b(\varepsilon^c)$, a pure-strategy equilibrium does not exist.*

Proof. See the Appendix. ■

Because ε^c decreases in v_b , a pure-strategy equilibrium is more likely to emerge when v_b is sufficiently large.

4.3 Asymmetric costs

We now consider a different source of vertical differentiation, namely, an asymmetry in the cost functions of Sellers a and b . The quality of the content produced by the sellers is now constant and equal to v , assumed to be high enough to guarantee full coverage. Sellers maximize:

$$\pi_j = D_j(p_j - C_j), \quad j \in \{a, b\},$$

where C_j is a measure of the marginal cost of producing the content sold by a and b . To preserve the direction of the asymmetry in the primary model, we assume that $C_a = c_a > 0$, $C_b = 0$.

The framework differs from the primary model in two substantial ways. First, because the asymmetry lies in the costs rather than the value consumers attach to the content, the optimal mix of consumers is not affected by it and $\lambda^*(x_i) = 1 - x_i$. Second, when the better seller is more efficient, rather than offering higher-quality content, it would select a lower price in the absence of bias than the one offered by its competitor. Therefore, the platform would have an

incentive to penalize the worst of the two sellers with biased recommendations rather than the best seller as was in the primary model.

The analytical steps to solve the model mirror the ones for the main model. As before, the difference in equilibrium prices is tempered by the platform's intervention: ε^* is selected to reduce the distance between p_a and p_b . Unlike in our main model, however, because $C_a > C_b$, $\varepsilon^* \leq 0$, the platform introduces bias to boost the consumption of Seller b 's content. In doing so, the platform induces higher p_b and lower p_a to emerge in equilibrium compared to the case in which bias was not introduced. The platform balances the incentive to increase the bias to lower operational costs and reduce the bias to increase its subscription fee f .

Because both consumers and sellers must choose to join the platform, constraints ε^c and ε^s still have to be taken into account. Unlike in the main model, however, consumers do not want to purchase the content of one of the two sellers disproportionately. Because the bias is introduced to penalize the content of the least efficient seller and because this seller charges a higher equilibrium price due to inefficiency, consumers are less sensitive to the bias than before. Thus, the constraint represented by ε^c still decreases in the cost differential but is less tight than in the main model.

The constraint that induces sellers to join is also looser than the one considered in the primary model. Unlike before, because the bias penalizes the inefficient seller, the highest possible bias the platform can introduce must make the worse seller — not the better one — indifferent between joining or not. According to the standard Hotelling logic, the seller with higher marginal costs would see lower profits in the cost differential. Because the seller penalized by the bias is the one with the worst outside option, the platform can ignore the constraint represented by ε^s for a wider range of values α . We, therefore, state the following:

Proposition 3. *When Sellers a and b have different marginal costs of production, the platform introduces a positive bias in favor of the most efficient of the two and increases this seller's equilibrium profits. In particular:*

$$\varepsilon^* = \min\{\varepsilon^p, \varepsilon^c, \varepsilon^s, \}$$

where

$$\frac{\partial|\varepsilon^p|}{\partial\Delta_c} > 0, \quad \frac{\partial|\varepsilon^c|}{\partial\Delta_c} < 0, \quad \frac{\partial|\varepsilon^s|}{\partial\Delta_c} < 0,$$

and $\Delta_c = |c_a - c_b|$.

Proof. See the Appendix. ■

This last exercise serves two purposes. It highlights the difference between vertical differentiation driven by consumer taste and efficiency. Because the two approaches generate bias of opposite signs (favoring the least liked content and the most efficient seller, respectively), considering only efficiency as a driver of asymmetry may lead to the partial conclusion that intervention by streaming platforms is socially desirable because it creates the incentive to reduce marginal costs or production and, with them, equilibrium prices.

On the other hand, our cost modeling serves as a proxy for the choice of music labels to experiment with their content, which can be expected to increase the production costs instead

of optimizing the content creation process. From this perspective, the exercise’s outcome agrees with the one presented in the main model: the platform discourages risk but rewards the “assembly line” production of content. The overall takeaway becomes straightforward: if we assume that novelty and experimentation are costlier than producing commodified content, the platform always has the incentive to penalize creative content through the strategic manipulation of what consumers are exposed to.

5 Discussion and conclusion

In this paper, we investigate the incentives a streaming platform has to bias its content bundles to achieve optimal profitability. The platform has the potential to generate utility for consumers who value a balanced mix of content. When the content is of equal quality, sellers choose uniform prices and the platform has no incentive to bias consumption. By contrast, when sellers offer vertically differentiated products, the platform has the incentive to set exploit different royalties to minimize costs. The seller with the higher-quality product wants royalties to be raised because consumers value its product more. When this happens, the platform has an incentive to bias consumption toward the “cheaper”, lower quality product to minimize costs. This comes at the detriment of consumers, who lose the additional utility generated through efficient content mixing, and the higher-quality seller, whose demand is artificially shrunk. In equilibrium, the latter would set a lower price than in the case without intervention. The platform hampers the incentive to introduce higher-quality products by punishing them with reduced exposure. Furthermore, platform intervention can significantly distort equilibrium R&D efforts, potentially suppressing all investments.

Based on several real-life examples, we assume that the platform cannot price-discriminate consumers. Otherwise, the platform would have the incentive to offer different bundles at different prices to extract the rent it helps generate. The ability to price-discriminate does not eliminate the incentive to bias. However, because consumers must be convinced to join the platform, personalized pricing removes the ability to bias consumption. Price discrimination and consumption bias are substitute strategies. If personalized pricing were possible, the higher-quality seller would be better off in equilibrium. On the other hand, consumers would necessarily be worse off. The reason for this is straightforward: even when the platform has no incentive to bias consumption, it always has the incentive to price-discriminate to make consumers indifferent between joining or not the platform if possible.

Furthermore, the result is carried forward by assuming that sellers bargain their royalty rates individually. The incentive to bias consumption is derived from the cost difference for the platform to stream the sellers’ content. Suppose, however, that the sellers were both represented by an intermediary, such as a copyright-collecting agency, bargaining royalty rates for both. Such an agent would have the incentive to set equal royalties to reduce the incentive to bias consumption toward the cheaper product. It’s uncertain that this would not be detrimental to the seller of the higher-quality product.

It should be stressed that the mechanism studied in this paper requires the platform’s algorithmic component to be relevant. In a world where consumers have no access to automatically

generated content susceptible to manipulation but are always in perfect control of what they consume, the distortions predicted by the model would not materialize. On the other hand, if the algorithmic recommendation of streaming platforms is biased in a way that damages consumers, regulatory intervention would prevent such distortions from arising. Thus, in the broadest sense, the paper follows in the footprints of many others, calling for the inspection and direct regulation of the algorithms used by digital platforms to provide their service. While we acknowledge that this may disincentivize R&D expenditure and innovation to improve these algorithms, the loss could be compensated by stronger incentives to innovate on the content that algorithms would no longer be able to penalize.

In our model, the platform has to provide the sellers with at least the same payoffs they would gain if they were not joining the streaming service. Hence, profit-wise, sellers' participation in the platform implies higher incentives to invest in content quality. Therefore, our analysis examines the distortion of investments from the *potential* level the seller would select in the absence of the recommendation bias. Thus, we take a benevolent view toward the platform which, by design, never directly harms sellers. The shares of informed and uninformed consumers are exogenously determined and, in our model, act as a proxy of the popularity of music labels. Debuting music labels are unknown to the public and would likely be unable to succeed outside the platform service. By contrast, established music labels benefit from a large network of consumers interested in their content and willing to purchase it regardless of whether they join the streaming service.²¹

To conclude, the implications of the platform's ability to make content accessible to more consumers should be considered. Streaming platforms represent a substantial portion of the market they host and, therefore, many music labels (especially new ones) have little hope of reaching the public without being hosted on one of these platforms. Joining, however, requires coming to terms with the platform's ability to act as a gatekeeper. If music labels and content creators need the platform to reach interested consumers and the platform is designed to punish good content if it comes at a higher price, incentives to vertically differentiate are weakened. Hence, the model highlights a potential risk embodied in the platform ecosystem: if a significant share of users is held "captive" by the platform, the content available for streaming may become more commodified, representing a loss for society that is difficult to quantify.

²¹Notably, Joni Mitchell and Neil Young's collections were unavailable on Spotify between 2022 and 2024. Their motivation to temporarily delist from Spotify was not driven by commercial disputes but by the debate concerning Covid-19 misinformation in Joe Rogan's controversial podcast, which is also streamed on the platform. Yet, their decision to abandon the streaming service reflects the greater freedom celebrities enjoy vis-à-vis beginners.

References

- Adams, W. J. and Yellen, J. L. (1976). Commodity Bundling and the Burden of Monopoly. *The quarterly journal of economics*, pages 475–498.
- Aguiar, L. and Waldfogel, J. (2021). Platforms, Power, and Promotion: Evidence from Spotify Playlists. *The Journal of Industrial Economics*, 69(3):653–691.
- Aguiar, L., Waldfogel, J., and Waldfogel, S. (2021). Playlisting Favorites: Measuring Platform Bias in the Music Industry. *International Journal of Industrial Organization*, 78:102765.
- Aguiar, L., Waldfogel, J., and Zeijen, A. (2023). Platform Power Struggle: Spotify and the Major Record Labels, 2017-2020. *Mimeo*.
- Anderson, S. P. and Coate, S. (2005). Market Provision of Broadcasting: A Welfare Analysis. *The review of Economic studies*, 72(4):947–972.
- Anderson, S. P. and Gabszewicz, J. J. (2006). The Media and Advertising: a Tale of Two-sided Markets. *Handbook of the Economics of Art and Culture*, 1:567–614.
- Anderson, S. P. and Neven, D. J. (1989). Market Efficiency with Combinable Products. *European Economic Review*, 33(4):707–719.
- Arya, A., Mittendorf, B., and Sappington, D. E. (2007). The Bright Side of Supplier Encroachment. *Marketing Science*, 26(5):651–659.
- Bourreau, M. and Gaudin, G. (2022). Streaming Platform and Strategic Recommendation Bias. *Journal of Economics & Management Strategy*, 31(1):25–47.
- Bourreau, M., Kraemer, J., and Hofmann, J. (2021). Prominence-for-data Schemes in Digital Platform Ecosystems: Implications for Platform Bias and Consumer Data Collection. *Available at SSRN 3867580*.
- Brynjolfsson, E., Hu, Y., and Simester, D. (2011). Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales. *Management Science*, 57(8):1373–1386.
- Chen, Y. and Schwartz, M. (2013). Product Innovation Incentives: Monopoly vs. Competition. *Journal of Economics & Management Strategy*, 22(3):513–528.
- Chiang, W.-y. K., Chhajed, D., and Hess, J. D. (2003). Direct Marketing, Indirect Profits: A Strategic Analysis of Dual-channel Supply-chain Design. *Management science*, 49(1):1–20.
- Cremer, H. and Thisse, J.-F. (1991). Location Models of Horizontal Differentiation: A Special Case of Vertical Differentiation Models. *The Journal of Industrial Economics*, 39(4):383–390.
- Datta, H., Knox, G., and Bronnenberg, B. J. (2018). Changing Their Tune: How Consumers’ Adoption of Online Streaming Affects Music Consumption and Discovery. *Marketing Science*, 37(1):5–21.

- Döpfer, H. and Rasch, A. (2022). Combinable Products, Price Discrimination, and Collusion. Technical report, DICE Discussion Paper.
- Fleder, D. and Hosanagar, K. (2009). Blockbuster Culture’s Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity. *Management science*, 55(5):697–712.
- Fletcher, A., Ormosi, P. L., Savani, R., and Castellini, J. (2023). Biased Recommender Systems And Supplier Competition. *Available at SSRN 4319311*.
- Freimane, M. (2022). Substituting Away? The Effect of Platform Bargaining Regulation on Content Display. Technical report, mimeo.
- Gal-Or, E. and Dukes, A. (2003). Minimum Differentiation in Commercial Media Markets. *Journal of Economics & Management Strategy*, 12(3):291–325.
- Halmenschlager, C. and Mantovani, A. (2017). On the Private and Social Desirability of Mixed Bundling in Complementary Markets with Cost Savings. *Information economics and policy*, 39:45–59.
- Hoernig, S. and Valletti, T. M. (2011). When Two-part Tariffs are Not Enough: Mixing with Nonlinear Pricing. *The BE Journal of Theoretical Economics*, 11(1).
- Hoernig, S. H. and Valletti, T. M. (2007). Mixing Goods with Two-part Tariffs. *European Economic Review*, 51(7):1733–1750.
- Jeon, D.-S. and Nasr, N. (2016). News Aggregators and Competition among Newspapers on the Internet. *American Economic Journal: Microeconomics*, 8(4):91–114.
- Kumar, N. and Ruan, R. (2006). On Manufacturers Complementing the Traditional Retail Channel with a Direct Online Channel. *Quantitative Marketing and Economics*, 4:289–323.
- Motta, M. (2023). Self-preferencing and Foreclosure in Digital Markets: Theories of Harm for Abuse Cases. *International Journal of Industrial Organization*, 90:102974.
- Mussa, M. and Rosen, S. (1978). Monopoly and Product Quality. *Journal of Economic Theory*, 18(2):301–317.
- Oestreicher-Singer, G. and Sundararajan, A. (2012). Recommendation Networks and the Long tTail of Electronic Commerce. *Mis quarterly*, pages 65–83.
- Peitz, M. and Sobolev, A. (2022). Inflated Recommendations. *CRC TR 224 Discussion Paper Series*.
- Peitz, M. and Valletti, T. M. (2008). Content and Advertising in the Media: Pay-TV Versus Free-to-air. *international Journal of industrial organization*, 26(4):949–965.
- Pepall, L., Richards, D., and Norman, G. (2014). *Industrial Organization: Contemporary Theory and Empirical Applications*. John Wiley & Sons.

- Popper, B. (2016). Spotify's Discover Weekly Reaches 40 Million Users and 5 Billion Tracks Streamed. *The Verge*, May, 25.
- Rhee, B.-D. and Park, S.-Y. (2000). Online Store as a New Direct Channel and Emerging Hybrid Channel System. *Unpublished manuscript*.
- Ronayne, D. and Taylor, G. (2022). Competing Sales Channels with Captive Consumers. *The Economic Journal*, 132(642):741–766.
- Sutton, J. (1986). Vertical Product Differentiation: Some Basic Themes. *The American Economic Review*, 76(2):393–398.
- Thomes, T. P. (2013). An Economic Analysis of Online Streaming Music Services. *Information Economics and Policy*, 25(2):81–91.
- Waldfogel, J. (2024). Amazon Self-preferencing in the Shadow of the Digital Markets Act. Technical report, National Bureau of Economic Research.

A Appendix

Proof of Lemma 1, Proposition 1, and Corollary 1.

Proof. Starting from the indifferent consumers' locations, we calculate the demands of the three firms. The platform's demand comprises the consumers who join the platform (i.e., $D_p = x_{pb} - x_{ap}$). On the other hand, the expressions of the two music labels' demands are more complex because they include not only the direct consumption by users who buy directly from them but also the share of their content streamed to the platform's users. This can be represented as:

$$D_a = x_{ap} + \int_{x_{ap}}^{x_{pb}} (\lambda(x_i) + \varepsilon(x_i)) dx_i$$

$$D_b = 1 - x_{pb} + \int_{x_{ap}}^{x_{pb}} (1 - \lambda(x) - \varepsilon(x_i)) dx_i$$

What matters from the platform's perspective is not the individual level of bias that each consumer accepts but the total mass of demand that, via the biased recommendation, can shift from one seller to another. In other words, provided that the total mass of demand does not exceed the aggregate consumer participation constraint, we can treat it as a uniform value ε . Accordingly, the demand functions of the two music labels change as follows:

$$D_a = x_{ap} + \int_{x_{ap}}^{x_{pb}} \lambda(x_i) dx_i + \varepsilon$$

$$D_b = 1 - x_{pb} + \int_{x_{ap}}^{x_{pb}} (1 - \lambda(x)) dx_i + \varepsilon$$

We use these demands and the indifferent consumers' locations (see the main text) to obtain the profit functions of the three firms:

$$\pi_a = \left(\frac{pb - pa + t - v_b}{2t} + \varepsilon \right) p_a, \quad \pi_b = \left(\frac{pa - pb + t + v_b}{2t} + \varepsilon \right) p_b,$$

$$\pi_p = \frac{t - \sqrt{f - p_a} - \sqrt{f - p_b}}{t} f - \frac{t + p_b - p_a - 2\sqrt{(f - p_a)t}}{2t} p_a - \frac{t + p_a - p_b - 2\sqrt{(f - p_b)t}}{2t} p_b$$

Simple maximization concerning prices yields the following:

$$p_a(\varepsilon^p) = t - \frac{v_b}{3} + \frac{2\varepsilon^p t}{3}; \quad p_b(\varepsilon^p) = t + \frac{v_b}{3} - \frac{2\varepsilon^p t}{3}; \quad f(\varepsilon^p) = \frac{10t}{9} + \frac{v_b^2}{4t} - \varepsilon^p(v_b - \varepsilon^p t)$$

Using these prices in the functions of firms' profits and consumers' locations, we obtain Lemma 1.

From Lemma 1, we obtain the equilibrium bias selected by the platform and prove Proposition 1. No consumer is willing to accept a mix that contains too much Content a (i.e., $\varepsilon(x_i) \leq \bar{\varepsilon}(x_i)$). We assume that the platform can re-distribute the bias toward consumers according to their participation constraints, which is implicitly known by the platform because it is assumed that the platform knows consumers' individual locations. In other words, the con-

dition we impose is that the total demand shifted by the platform cannot exceed the aggregate participation constraint of the consumers, as stated in Condition 1. Moreover, we know the bias must also satisfy consumer participation as derived in expression (9).

Taking all these conditions into consideration, the platform's problem is as follows:

$$\begin{aligned} \max_{\varepsilon^p} \quad & \pi_p(\varepsilon) = \frac{t + 3\varepsilon^p(v_b - \varepsilon^p)}{27} - \frac{v_b^2}{36t} \\ \text{s.t.} \quad & \varepsilon^p < \min\{\varepsilon^c, \varepsilon^s\} \end{aligned}$$

Standard maximization yields the unconstrained profit-maximizing level of bias:

$$\frac{\partial \pi_p(\varepsilon)}{\partial \varepsilon} = 3v_b - 6t\varepsilon \implies \varepsilon^p = \frac{v_b}{2t}$$

which, combined with the aforementioned constraints, leads to Proposition 1.

Corollary 1 is derived by inputting $\varepsilon^p = \frac{v_b}{2t}$ into the prices in (7) and into the profits in Lemma 1 ■

Proof of Proposition 2 and Corollary 2.

Proof. The values in Proposition 2 are obtained from the solutions of the seller optimization problems when they operate within the platform (Corollary 1) and when they decide not to join (expression 8). This can be expressed as:

$$\max_{v_a} \pi_a^{eq} |_{\varepsilon=\varepsilon^p} = \frac{t}{2} - v_a^2 \quad \max_{v_b} \pi_b^{eq} |_{\varepsilon=\varepsilon^p} = \frac{t}{2} - \phi_b v_b^2 \quad \text{A.1}$$

$$\max_{v_a} \pi_a^{eq,out} = \alpha \frac{(3t - v_b + v_a)^2}{18t} - v_a^2 \quad \max_{v_b} \pi_b^{eq,out} = \alpha \frac{(3t + v_b - v_a)^2}{18t} - \phi_b v_b^2 \quad \text{A.2}$$

The maximization yields $v_a^{eq} = v_b^{eq} = 0$ are taken straightforward from (A.1). This is because the bias fully compensates for the demand premium the high-quality product would get and suppresses all the incentives to vertically differentiate.

However, outside the platform, the maximization process in (A.2) yields:

$$\begin{aligned} \frac{\partial \pi_a}{\partial v_a} = \frac{\alpha(3t + v_a - v_b)}{9t} - 2v_a & \implies v_a^{eq,out} = \frac{\alpha(9t\phi_b - \alpha)}{3((18t - \alpha)\phi_b - \alpha)} \\ \frac{\partial \pi_b}{\partial v_b} = \frac{\alpha(3t - v_a + v_b)}{9t} - 2v_b\phi_b & \implies v_b^{eq,out} = \frac{\alpha(9t - \alpha)}{3((18t - \alpha)\phi_b - \alpha)} \end{aligned}$$

The second-order conditions are satisfied when $\phi_b > \frac{\alpha}{18t}$. Moreover, both values are positive if and only if $\phi_b > \frac{\alpha}{9t}$. In the rest of the analysis, we assume this last condition to hold.

This proves Proposition 2. Corollary 2 is derived using the values obtained to compare profits:

$$\pi_b^{eq} |_{v_a=v_b=0} \equiv \frac{t}{2} > \frac{\alpha\phi_b(9t - \alpha)^2(18t\phi_b - \alpha)}{9(18t\phi_b - \alpha(\phi_b + 1))^2} \equiv \pi_b^{eq,out} |_{v_a=v_a^{eq,out}, v_b=v_b^{eq,out}}$$

This condition is verified if and only if:

$$\phi_b > \underline{\phi}_b \equiv \frac{9\alpha t}{-\alpha^3 + \alpha(\alpha - 9t)\sqrt{\alpha^2 + 81t^2} - 18(\alpha - 1)t - 81(\alpha - 2)t^2 + 9(2\alpha - 1)\alpha t}$$

Tedious calculations reveal that $\underline{\phi}_b < 1$ for all $\alpha \in [0, 1]$. Moreover:

$$\underline{\phi}_b > \frac{\alpha}{9t} \iff \alpha > \frac{1}{2}.$$

Thus, the efficient seller is better off not joining the streaming platform in the parameter region where both $\alpha \in (\frac{1}{2}, 1)$ and $\phi_b < \underline{\phi}_b$. This region is shown in Figure 3. This concludes the proof. ■

Proof of Corollary 3.

Proof. When the seller participation constraint binds, the equilibrium bias at Stage 1 is given by:

$$\varepsilon^s = \frac{(3t + v_b - v_a)(1 - \sqrt{\alpha})}{2t}.$$

With this level of bias, the sellers are certain to get a payoff as large as the one they would receive if they did not join the streaming service (i.e., $\pi_a^{eq,out}$ and $\pi_b^{eq,out}$). Intuitively, this level of bias implies firms have the same incentives to invest in quality as they would have if the platform did not exist. Therefore, the maximization problem is exactly as in A.2, and so are the solutions: $v_a = v_a^{eq,out}$ and $v_b = v_b^{eq,out}$. Using these solutions in ε^s yields the equilibrium bias at Stage 0 as defined in Corollary 3. ■

Proof of Lemma 2.

Proof. Define the demand of the sellers and the platform as D_a , D_b , and D_p , respectively. Consider a situation in which $v_b > v_a = 0$, so that, in equilibrium and absent any bias, $D_a < D_b$ and $p_a < p_b$. Because the platform pays p_a and p_b to the sellers in royalties, it has an incentive to increase the share of Content a (the cheapest) in the mix offered to consumers. Also, define ε^c as the total demand on the platform that can be steered toward the cheapest, inferior product (a) without altering D_p .

The bias enters the profit functions of the sellers by altering their demand function. Hence, $D'_{a,\varepsilon} > 0$ and $D'_{b,\varepsilon} < 0$. The two sellers anticipate the bias and modify their prices accordingly. Seller a , who benefits from the demand shock, increases the price to $p_a(\varepsilon^c) > p_a$, whereas Seller b lowers the price to $p_b(\varepsilon^c) < p_b$. This is because the bias enters the demand function inelastically (i.e., as long as $p_a(\varepsilon^c) < p_b(\varepsilon^c)$, the entire mass ε^c shifts toward Product a).

Thus, the two scenarios described in the Lemma emerge. First, the demand shift is insufficient to change the price ranking. In this case, the resulting equilibrium is such that

$$p_a^* \equiv p_a(\varepsilon^c) < p_b(\varepsilon^c) \equiv p_b^*$$

so that Seller a and the platform are better off. By contrast, Seller b and consumers are worse off (increasing the recommendation bias lowers consumer surplus).

Second, the demand shift is so great that the price ranking changes. In such a case, a pure-strategy equilibrium no longer exists. Thus, in anticipation of ε^c , sellers change their prices to such an extent that

$$p_a^* \equiv p_a(\varepsilon^c) \geq p_b(\varepsilon^c) \equiv p_b^*$$

Observing these prices, the platform implements a recommendation bias that goes in the opposite direction of the one anticipated by the seller, promoting Content b , which is now the cheapest. Clearly, this doesn't represent equilibrium. ■

Proof of Proposition 3

Proof. We derive the relevant functions as we did for the primary model. We assume that $V_a = V_b = v$ holds in this scenario; thus, we update the indifferent consumers' locations using expression (1) accordingly:

$$x_{ap}^{dc} = \sqrt{\frac{f - p_a}{t}}; \quad x_{pb}^{dc} = 1 - \sqrt{\frac{f - p_b}{t}}.$$

The sellers' demands are likewise derived in the baseline with homogeneous goods. The platform's demand is given by $D_p^{dc} = x_{pb}^{dc} - x_{ap}^{dc}$, where the apex dc stands for "different costs".

We assume $c_a > c_b = 0$; thus, the profit function of a must account for the marginal cost and in particular:

$$\pi_a^{dc} = (p_a - c_a) \left(x_{ap}^{dc} + \int_{x_{ap}^{dc}}^{x_{pb}^{dc}} \lambda^*(x) dx - \varepsilon \right)$$

The profit functions of b and p remain unchanged:

$$\pi_b^{dc} = p_b \left(1 - x_{pb}^{dc} + \int_{x_{ap}^{dc}}^{x_{pb}^{dc}} (1 - \lambda^*(x)) dx + \varepsilon \right)$$

$$\pi_p^{dc} = p_p \left(x_{pb}^{dc} - x_{ap}^{dc} \right) - p_a \left(\int_{x_{ap}^{dc}}^{x_{pb}^{dc}} \lambda^*(x) dx - \varepsilon \right) - p_b \left(\int_{x_{ap}^{dc}}^{x_{pb}^{dc}} (1 - \lambda^*(x)) dx + \varepsilon \right)$$

Seller a will see the platform bias consumption moving away from it because, by standard Hotelling logic, $p_a > p_b$ whenever $c_a > c_b$.

After substituting $\lambda^*(x) = (1 - x)$, standard F.O.C. arguments lead to equilibrium prices:

$$p_a^{dc} = \frac{2c_a}{3} - \frac{2t\varepsilon}{3} + t, \quad p_b^{dc} = \frac{c_a}{3} + \frac{2t\varepsilon}{3} + t,$$

$$p_p^{dc} = \frac{c_a^2}{16t} - \frac{1}{2}c_a(1 - \varepsilon) + t \left(\varepsilon^2 + \frac{10}{9} \right),$$

and profits:

$$\pi_a = \frac{(t(3 - 2\varepsilon) - c_a)^2}{18t}, \quad \pi_b = \frac{(t(3 + 2\varepsilon) + c_a)^2}{18t}, \quad \pi_p = \frac{1}{54} (3c_a\varepsilon + t(2 - 6\varepsilon^2)) - \frac{c_a^2}{144t};$$

The latter equation immediately leads to the platform's profit-maximizing bias by the stan-

standard F.O.C. argument:

$$\varepsilon^p = \frac{c_a}{4t}$$

To make consumers join, the following must hold true:

$$\varepsilon^p < \varepsilon^c = \int_{x_{a^c}^{d^c}}^{x_{b^c}^{d^c}} \bar{\varepsilon}(x) dx$$

where $\bar{\varepsilon}(x)$ is defined as the larger (absolute) bias a consumer x is willing to accept before choosing to leave the platform.

Finally, equilibrium profits in the subgame in which sellers choose not to join the platform are:

$$\pi_a^{dc,out} = \frac{\alpha(3t - c_a)^2}{18t}, \quad \pi_b^{dc,out} = \frac{\alpha(3t + c_a)^2}{18t},$$

so ε^p is constrained by ε^s satisfying:

$$\varepsilon^s = \frac{(3t - c_a)(1 - \sqrt{\alpha})}{2t}$$

The result, as stated in Proposition 3, follows immediately. ■



Download ZEW Discussion Papers:

<https://www.zew.de/en/publications/zew-discussion-papers>

or see:

<https://www.ssrn.com/link/ZEW-Ctr-Euro-Econ-Research.html>

<https://ideas.repec.org/s/zbw/zewdip.html>



IMPRINT

ZEW – Leibniz-Zentrum für Europäische Wirtschaftsforschung GmbH Mannheim

ZEW – Leibniz Centre for European
Economic Research

L 7,1 · 68161 Mannheim · Germany

Phone +49 621 1235-01

info@zew.de · zew.de

Discussion Papers are intended to make results of ZEW research promptly available to other economists in order to encourage discussion and suggestions for revisions. The authors are solely responsible for the contents which do not necessarily represent the opinion of the ZEW.