

Gambato, Jacopo; Sandrini, Luca

Working Paper

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

ZEW Discussion Papers, No. 23-032

Provided in Cooperation with:

ZEW - Leibniz Centre for European Economic Research

Suggested Citation: Gambato, Jacopo; Sandrini, Luca (2023) : Not as good as it used to be: Do streaming platforms penalize quality?, ZEW Discussion Papers, No. 23-032, ZEW - Leibniz-Zentrum für Europäische Wirtschaftsforschung, Mannheim

This Version is available at:

<https://hdl.handle.net/10419/275671>

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.23-032 | 08/2023

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[‡]

August 3, 2023

Abstract

We study the incentives of a streaming platform to bias consumption when products are vertically differentiated. The platform offers mixed bundles of content to monetize consumers' interest in variety and pays royalties to sellers based on the effective consumption of the content they produce. When products are not vertically differentiated, the platform has no incentive to bias consumption in equilibrium: the platform being active represents a Pareto-improvement compared to the case in which she is not. With vertical differentiation, royalties can differ; the platform always biases recommendations in favor of the cheapest content, which hurts consumers and the high-quality seller. Biased recommendation always diminishes the incentives of a seller to increase the quality of her content for a given demand. If a significant share of the users is ex-ante unaware of the existence of the sellers the platform can bias recommendations more freely, but joining the platform encourages investment in quality. The bias, however, can lead to inefficient allocation of R&D efforts. From a policy perspective, we propose this as a novel rationale for regulating algorithmic recommendations in streaming platforms.

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

JEL Codes: D4, L1, L5

*We thank the IOEK division at ZEW Mannheim, Marc Bourreau, Leonardo Madio, Martin Peitz, Francesco Clavorà Braulin, and Christian Soegaard, as well as the participants to the 11th Oligo Workshop and XII Ibeo Workshop for their insightful comments.

[†]University of Mannheim, ZEW Mannheim, MaCCI - Jacopo.Gambato@zew.de

[‡]Budapest University of Technology and Economics - sandrini.luca@gtk.bme.hu

1 Introduction

Streaming platforms found great success in the digital era thanks in part to personalized features such as the well-known Spotify “Discover Weekly” playlist, automatically generated for each user every week¹. Such features are extremely popular and have a real impact on consumption patterns. For example, [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 not included. The two observations strongly point to the ability of such platforms to affect individual consumers’ effective consumption bundle.

In this study, we focus our attention on the ability such a platform has to bias consumption through algorithmic recommendation: we study the role of a streaming platform in enabling consumers to “mix” different products in individually optimal proportions and its ability to profitably introduce bias. We consider a framework in which two horizontally differentiated sellers compete not with each other but with a monopoly platform offering a streaming service to all the consumers in the market. The platform attracts consumers that want to enjoy the content produced by both sellers, cashes in a subscription fee, and pays royalties to the sellers based on the effective consumption of their product. The set-up well represents economic agents such as Spotify and different music labels who sell a product that is available both on the platform and outside of it, for example in the form of CDs or audio files up to purchase.

While most of the current literature on the topic addresses concerns regarding anti-competitive steering practices², we adapt the framework to investigate the effect of subscription-based business models on the incentive of content providers to innovate. As shown in recent work by [Bourreau and Gaudin \(2022\)](#), streaming platforms have strong incentives to bias recommendations to reduce consumption of products that carry higher royalties.³ The result is intuitive: since the platform charges a fixed fee to all consumers, she has the incentive to minimize costs once their participation is ensured. The finding speaks in favor of royalties being strategically used to gain prominence on this kind of platform ([Bourreau et al., 2021](#)). It is however unclear how this dynamic is affected by vertical differentiation in the products offered. On the one hand, a higher quality product is more desirable and, therefore, allows the platform to charge more to access it. On the other hand, as the seller with the superior product would demand to be paid a higher royalty rate, the incentive to bias away from his product would be higher as well. Which effect dominates and the overall impact on the platform, the sellers, and the buyers are at the core of our analysis.

More precisely: we propose a framework in which two horizontally differentiated content providers (here labeled “*a*” and “*b*”) sell their bundle good 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

¹[Popper \(2016\)](#) reports that 40 million out of Spotify’s, at the time, 100 million users used it in 2016. More recently, Spotify reported that in the five years since its launch, Discover Weekly streamed 2.3 billion hours of music. See <https://newsroom.spotify.com/2020-07-09/spotify-users-have-spent-over-2-3-billion-hours-streaming-discover-weekly-playlists-since-2015/>.

²See, for example, [Peitz and Sobolev \(2022\)](#).

³In recent work, [Aguiar et al. \(2023\)](#) provide empirical support to the theoretical results. Using data from playlists on Spotify, the authors show that the share of songs owned by major music labels in prominent playlists offered by Spotify decreases over time. The interpretation of the finding is that Spotify is leveraging its power in the playlist market to obtain better deals in royalty payments.

consists of only the content they produce (i.e., bundle good a (respectively b) is entirely made of content a (b)). The sellers are assumed to be located at the extreme of the Hotelling line. Along the entire line, a platform for streamed content (labeled p) offers a subscription-based service: upon paying a uniform fee p_p , a consumer can access a mix of content from a and b .⁴

The platform remunerates the sellers and pays them a royalty rate per share of the content shown to each consumer. Consumers, therefore, can choose between three bundle goods: sellers a and b offer pure bundles, whereas the platform offers mixed bundle goods. If a consumer buys either of the pure bundle goods, she consumes only the content owned by one seller. 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.

We show that when products are vertically homogeneous, the existence of such an intermediary represents a Pareto-improvement compared to the alternative competitive outcome. The platform attracts consumers located in the middle of the Hotelling line: these are the consumers with the highest willingness to pay for the possibility of mixing products. Hence, the platform can charge a price higher than the sellers’ and still make positive profits after paying royalties. This result holds under the extreme assumption that sellers have full bargaining power when setting the royalty rate.

In the baseline specification with uniform products, the platform has no incentives to introduce a bias for its users. When products are not vertically differentiated, the two sellers optimally select the same price in equilibrium anticipating the consumption taking place both in and out of the platform. Biasing consumption, in this case, would skew the demand in favor of one of the two sellers, inducing him to raise his price and monetize from it. The rival would, instead, choose to reduce the price of his product to induce consumers closer to him to leave the platform. Such a strategy cannot be optimal for the platform, as it would effectively bias consumption in favor of the most expensive option rather than the cheaper one.

The model’s predictions change drastically when products are vertically differentiated. Since consumers value high quality, without platform intervention, the equilibrium outcome features a higher price and larger share of consumption for the high-quality product. While the platform can raise its price to monetize the higher average quality of her bundles, her ability to do so is limited since the rival is forced to offer a lower price than under no vertical differentiation. When consumers are offered their optimal consumption bundle, the platform is hurt by the quality differential. In this specification, we show that the platform always has the incentive to bias consumption away from the better, more expensive product. Since 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 consumers’ participation constraints and offers a biased bundle that makes consumers indifferent between joining and leaving the platform. The result follows from the assumption that the platform can discriminate consumers and offers targeted recommendation

⁴Such fees are the most common when it comes to music streaming platforms. Besides Spotify, other notable examples are Deezer and Pandora.

bias⁵.

When the platform biases consumption away from the better and more expensive product, the respective seller is penalized. If this penalty is severe enough, the seller could choose not to make his product available on the platform and compete with the other seller directly instead. Whenever this happens, it is clear that the platform cannot be active. Consumers join the platform to mix content produced by both music labels: if the platform cannot attract both sellers, no consumer is interested in joining. Streaming platforms, however, have been known to popularize less-known artists and, therefore, generate demand. To capture this additional dimension, we split the unit mass of consumers in 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. When deciding to join, the sellers' outside option is worse if the group of consumers that only consumes if the platform is active is larger. It follows that the ability of the platform to bias consumption depends on the additional consumption she generates.

The findings have relevant implications both in the context of consumption steering in digital markets and in regard to the effect of subscription-based business models on the incentives of sellers to innovate. First, steering emerges in equilibrium, not because of sellers competing for prominence but rather as a response of the platform to soften competition: the platform has the incentive to contain the price effect generated by the difference in quality and the stronger market presence of the better product. The overall effect hurts consumers and the seller with the high-quality product and benefits the runner-up by skewing consumption towards him.

Incentives to innovate and produce higher-quality goods are weakened when a platform that can bias consumption is present in the market for given demand. In particular, the platform always selects a positive level of bias when products are vertically differentiated. The bias is constrained by sellers' and consumers' participation decisions. In equilibrium, consumers who join the platform are exposed to more of the cheaper, low-quality content that they would optimally select. Furthermore, we show that the bias introduced by the platform distorts incentives to innovate and can lead to inefficient allocation of R&D efforts, with the more (resp., less) efficient seller investing less (resp., more) than it would have if the platform had been inactive.

The rest of the paper is structured as follows: after a review of the relevant literature, we introduce the model and solve the baseline specification with homogeneous quality products in Sections 2 and 3, respectively. In Section 4 we introduce vertical differentiation by allowing one of the products to provide additional fixed, stand-alone utility to all consumers. After solving and discussing the seller's participation decision as a function of the demand generated by the platform, we endogenize the decision to invest in quality by the sellers (Subsection 4.2). We complement the main analysis by considering several extensions in Section 5. Section 6 concludes.

⁵The assumption is strong but realistic. It is well known that platforms such as Spotify offer personalized content in the form of playlists based on past consumption. The assumption, then, is simply a reversal of what is already known: the platform being aware of a consumer's taste instructs how much bias she would be willing to tolerate.

1.1 Related Literature

Recommendation systems represent a core feature of digital platforms, and streaming platforms like the ones we study in this paper are no exception. The impact of these 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 impact that inclusion in automatically generated playlists has on the popularity of new songs on Spotify. Generally, recommendation systems have been shown to greatly widen the range of consumed products, a phenomenon generally referred to as the “long-tail effect” ([Fleder and Hosanagar, 2009](#); [Brynjolfsson et al., 2011](#); [Oestreicher-Singer and Sundararajan, 2012](#); [Datta et al., 2018](#)). More recently, literature concerned with the incentives of intermediaries to strategically skew recommendations in a way that systematically harms consumers (see, for example, recent work by [Peitz and Sobolev, 2022](#)) has been on the rise.

It seems clear that the impact that these systems have on consumption makes them an obvious candidate for strategic manipulation. In this spirit, [Bourreau et al. \(2021\)](#) studies competition for prominence on digital platforms, comparing bias generated when prominence is gained via monetary or data-based compensation. We distance ourselves from this setting in various ways: first, we capture bias not through manipulation of the search query but through manipulation of the composition of available bundles. Second, we assume the platform already has relevant information on the buyers’ side by building competition on the Hotelling line. From this perspective, the paper more closely resembles [Bourreau and Gaudin \(2022\)](#), who consider a market where the platform does not directly compete against the sellers. Instead, we explicitly model competition through the ability for consumers to purchase directly from the sellers active in the market. Furthermore, we allow the platform to condition the bias imposed on consumers based on their location while [Bourreau and Gaudin \(2022\)](#) focuses on uniform biases.

Novel to the literature is also the fact that we introduce vertical differentiation between the sellers: in our model, we assume sellers offer goods that are differentiated both horizontally and vertically. Moreover, we assume consumers are only sensible to horizontal variations of the good, whereas they do not differ in their willingness to pay for quality. In that sense, we depart from [Mussa and Rosen \(1978\)](#), where consumers’ income is taken into consideration (see also [Cremer and Thisse \(1991\)](#) and [Sutton \(1986\)](#)). Here, we adopt a model of spatial competition à la Hotelling, in which we assume that one seller provides a good with higher intrinsic value than his rival⁶. We do so to better relate to the existing literature on innovation, which often includes both vertical and horizontal dimensions of differentiation (see [Chen and Schwartz, 2013](#)).

Our findings that the platform has the means and the incentive to bias consumption in favor of cheaper content echo those detailed in [Freimane \(2022\)](#). The author examines the impact of a regulatory change affecting the bargaining process behind content provision on the Google News platform. In the paper, it is shown 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 towards cheaper

⁶In other words, 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 strictly prefer buying one over the other.

alternatives. Even though the channel through which the asymmetry arises is different (we hold bargaining power fixed, and focus on vertical differentiation instead), the outcome is well-aligned with our equilibrium predictions.

For the timing of 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. The choice is motivated both by technical sensibility and by the inner working of the music streaming industry. On one hand, it is sensible to assume recommendation systems to be implemented not in response to a specific interaction with a specific seller, but through an algorithm the platform commits to. Further, there is anecdotal evidence that streaming platforms and sellers therein trade-off royalties for exposure. In 2014, the online radio company Pandora admitted to agreeing with indie-label coalition Merlin to exchange lower royalty rates for an increase in exposure⁷. The testimony highlights that these agreements are common practice, and justify a timing such that sellers set royalties aware of the bias the platform might generate in response. In passing, the testimony suggests that the platform might want to commit to its recommendation system regardless for fear of legal repercussions: if they were to condition the recommendation system on royalties, they would realistically be more vulnerable to legal action being taken against them.

Widening the scope of the discussion, the paper relates to the evolving literature on the economics of media markets. While most past contributions focused on the mix of content and advertising in media ([Anderson and Coate, 2005](#); [Anderson and Gabszewicz, 2006](#); [Peitz and Valletti, 2008](#); [Thomes, 2013](#)), we ignore this dimension altogether. We do so for two reasons: first, while it is true that many streaming platforms offer free subscriptions with ads in alternative to the ad-free “Premium” ones, the latter in itself represents an enormous and still growing market⁸.

Second, while the literature on advertisement in media contraposes content and ads, bringing positive and negative utility to consumers respectively, we focus on the 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 would have the incentive to penalize high-quality products she is not competing with is not. To our knowledge, this is the first paper to explore this dimension of the problem. The emerging result is in conflict with recent work by [De Corniere and Taylor \(2019\)](#): despite the alignment of interest of sellers and buyers to produce and consume better quality products, the “congruence” case studied by the authors, bias in our case can never be CS improving as they suggest. The reason follows from the discussion in the introduction: the platform uses bias to strategically reduce cost rather than inflate revenue, which reverts the incentives and the direction of the bias that congruence would suggest.

More in general, the paper relates to the growing literature on platforms initiated by [Armstrong \(2006\)](#). Streaming platforms find their footing and generate network effects by facilitating mixing in addition to facilitating contact between different sides – be these sides buyers and

⁷See the testimony of Stephen McBride, Docket No. 14-CRB-0001-WR for a full recounting

⁸According to Spotify’s earning report to investors, the number of the premium subscriber in Q3 of 2022 was 195 million. Available at: https://s29.q4cdn.com/175625835/files/doc_financials/2022/q3/Q3-2022-Sharholder-Deck-FINAL-LOCKED.pdf

sellers, or users and advertisers. This distinction separates our work from other models studying platform steering: [Teh and Wright \(2022\)](#) show that steering can benefit consumers when searching for a product that represents a good enough match is very costly. In our context, instead, the platform profits by offering a service, that is, by allowing consumers to reach mixed bundles and represents a net welfare gain when she cannot, or chooses not to, bias consumption. Whenever she intervenes, however, she does so to the detriment of consumers.

To model our environment, we build on early work by [Adams and Yellen \(1976\)](#) and, more closely, by [Anderson and Neven \(1989\)](#): we consider a Hotelling framework in which location on the unit line uniquely determine the optimal mix of consumption. Intuitively, consumers closer to a (respectively b) want to purchase a higher share of the product produced by a (b). Consumers equidistant from the two find it optimal to consume the two in equal proportions. The framework has been used in the past 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](#)).

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 when competing against its suppliers. The recent work by [Aguiar et al. \(2023\)](#) relates very closely to our paper. They analyze the incentives of the platform 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 differentiate from this work in two ways. First, we propose a theoretical investigation of the platform’s incentives to bias its recommendation system to minimize costs. Second, we focus on individual recommendations.

Most of the literature considers consumers more or less sensible to prices and distribution channels⁹. In contrast, [Tsay and Agrawal \(2004\)](#) studies the manufacturer’s optimal choice of distribution channels between direct, retail-based, and hybrid, with a focus on market penetration. While retailers extract part of the rent generated by the sale, the additional market penetration they lead to can be worth pursuing. We consider a similar dynamic: music labels, the manufacturers in our setting, can choose to sell through the platform on top of directly because the platform generates additional demand that would remain inactive otherwise.

More closely related to this paper is recent work by [Ronayne and Taylor \(2022\)](#). The 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 rent from. In contrast, our market shares emerge endogenously in equilibrium. The presence of captive consumers in our setting would allow the platform to bias more aggressively in equilibrium, a result that is proxied by the demand expansion we feature in our model: the more consumers would stay inactive if not for the platform intermediation, the less constrained the platform is in designing the recommendation system.

Finally, this paper encompasses the literature on the effect of rent-sharing mechanisms on innovation incentives (see [Berton et al., 2021](#), for a review of the literature). Most of the

⁹See, for example, [Rhee and Park \(2000\)](#), [Chiang et al. \(2003\)](#), and [Kumar and Ruan \(2006\)](#)

research on this topic focused on the effects of institutions such as unions on the incentives of firms to invest in R&D. [Grout \(1984\)](#) first analyzed this topic and concluded that unions might act as rent-seekers, thus lowering firms' incentives to innovate. By appropriating part of the innovation-generated revenues, the argument goes, unions exert negative pressure on firms' incentives and introduce the well-known hold-up problem. Furthermore, [Haucap and Wey \(2004\)](#) focus on unionization structure (i.e., the degree of wage centralization) and its effects on innovation incentives. The authors show that centralized wage-setting institutions are the most efficient in generating innovation incentives. Indeed, under some conditions, a centralized union could also outperform a market where wages are determined competitively. On the contrary, [Mukherjee and Pennings \(2011\)](#) find that unions centralization increases the incentive for technology licensing, which, under some conditions, may boost the investments in innovation by firms. In the same spirit, [Kline et al. \(2019\)](#) find that firms obtaining patent protection observe a rise in workers' compensation and productivity¹⁰.

Our paper contributes to the literature by analyzing the issue in a B2B setting. More specifically, we consider the strategic interaction between innovators and a platform that can steer consumers' demand toward the most convenient good. We show that the platform may severely hinder incentives to invest in innovation even if the innovators have full bargaining power in determining their royalty rates. Moreover, we argue that the platform can appropriate part of the innovation value by biasing its recommendation system and artificially raising competitive pressure on the innovator. Finally, we show that the intervention of the platform can lead to severe inefficiency of the equilibrium allocation of R&D effort.

2 Model setup

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

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 has mass $1 - \alpha$. The information they possess (or don't possess) refers to the existence and location of the firms operating in the market. Moreover, informed consumers know ex-ante their own location on the line, as well as the exact location of the sellers. On the contrary, 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*, which are entirely made of contents produced *in-house*. These can be thought of as the albums produced by the two music labels. Instead, we define *mixed bundles* as the personalized good that the consumers can access via the streaming service, like a playlist that contains content produced by both sellers.

¹⁰Furthermore, they estimate that workers capture roughly 30% of patent-induced surplus.

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 consumer located at x if she joins the platform service. In particular, $\lambda(x)$ is the preferred share that the consumers would choose to maximize utility, whereas $\varepsilon(x)$ is the personalized bias on the recommendation system imposed by the platform. Put differently, $\varepsilon(x)$ is the extra share of content a offered to each consumer by the platform's algorithm. Conversely, $1 - \lambda(x) - \varepsilon(x)$ 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. We use p_p to identify the subscription fee paid by consumers to access the platform's service instead.

Finally, 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: $r_j = p_j$. The assumption allows us to ignore any direct bargaining between sellers and the platform and any effect of eventual differences in bargaining power. In a way, we assume that sellers have full bargaining power in the royalty setting stage and, therefore, always select the highest rate possible given their own price in the external market.

The utility function of consumer i located in x_i can be written as:

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

where $V_j = v + v_j$ is the intrinsic quality of the pure bundles (which is 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, and a music label-specific parameter $v_j \geq 0$. In what follows, we analyze the benchmark case of $v_a = v_b = 0$ and the asymmetric scenario where $v_b > v_a = 0$. 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. Figure 1 shows the diagram of the model.

Importantly, we allow the platform to bias the bundles offered to consumers as a mean to influence music labels' price decisions. By doing so, the platform alters the shares of content in the personalized mixed bundles: we analyze the incentives of the platform to steer consumers away from high-quality, and expensive, content and offer them a mixed bundle that is disproportionately rich in low-quality, and cheap, content.

Timing. The timing of the game is as follows: at stage 1, the platform chooses the level of bias of the recommendation system (ε) and commits to implementing it.¹¹ At stage 2, sellers observe the recommendation policy and decide whether to join the platform and serve both informed and uninformed consumers or to stay out and compete for informed consumers only

¹¹Once 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 order to achieve better economic conditions, and by fear of possible legal repercussions if the platform conditioned bias on royalty rates.

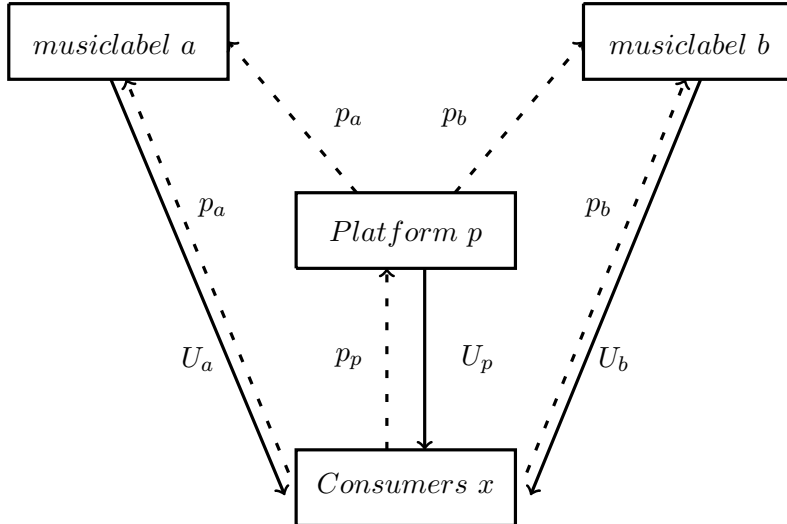


Figure 1: *The diagram of the model with payments and services when the platform is active, and both informed and uninformed consumers participate.*

in a standard Hotelling setting. 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 p_p). In Subsection 4.2, we augment the model with an additional stage 0 in which either one or both sellers costly invest in innovation to generate $v_j \geq 0$.

At stage 4, given the prices and the recommendation system, consumers make their consumption decision and profits realize. One should remember that the share α of informed consumers know both their location on the Hotelling line as well as the locations of the sellers. Instead, the $1 - \alpha$ uninformed consumers only know that a platform exists. We assume that all consumers can sample the platform for free before subscribing¹². During the free sample period, uninformed consumers learn the location of the firms and their preferences. If the sellers decide not to join the platform at stage 1, it is clear then that uninformed consumers do not learn anything and make no purchase.

Our solution concept is Sub-game Perfect Nash-Equilibrium. We solve the game by backward induction.

3 Baseline: Homogeneous quality

We begin the analysis by focusing on the baseline case where the two music labels produce content of identical quality – i.e., $V_j = v \forall j$. We start from the demand faced by the three sellers (the two music labels and the platform) in the last stage of the game. The assumption

¹²Many real-world streaming platforms, including Spotify, offer free trials to consumers. The assumption, therefore, well matches the kind of platform we aim to model.

$v_a = v_b = 0$ simplifies the utility functions to:

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

where the apex bln indicates we are in the baseline model specification.

As standard in these models, we derive the locations of indifferent consumers by equating the utility functions they obtain by choosing between the three options:

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

We adopt the following notation: x_{jk} indicates the indifferent consumer between buying from firm j and firm k , with $j, k = a, b, s$ 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 other two. In the proceeding of the analysis, we use x_{ab} mainly as a reference point.

Notice that $\varepsilon(x)$ does not enter the location of the indifferent consumers. This is simply because the platform knows consumers' locations and can offer them a personalized recommendation system. Indifferent consumers would change their consumption choice if subject to a bias that lowers their utility. Hence, the platform designs its algorithm to increase the bias in the distance between the consumers and their preferred music labels. This bias is personalized and it is bounded by the participation constraint of the consumers. More in detail, the personalized bias for each platform user is $\varepsilon(x_i) < \bar{\varepsilon}(x_i)$, where

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

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

With this in mind, we can now derive the personalized recommendation system set by the platform. Intuitively, the platform aims at maximizing consumption of the streaming service. To do so, it offers the *efficient bundle* to indifferent consumers. We define *efficient bundle* as the composite good that would be chosen by a consumer so that, for any prices p_a , p_b , and p_p , she would get the highest possible utility. By definition, the efficient bundle is not biased by the platform recommendation system ($\varepsilon = 0$). Formally:

$$\lambda^*(x_i) = \arg \max_{\lambda \in (0,1)} (U_{i,p}^{bln}) = 1 - x_i$$

Using this consumption choice, it is possible to update the location of the indifferent consumers as:

$$x_{ap}^{bln} |_{\lambda(x_{ap})=\lambda^*(x_{ap})} = \sqrt{\frac{p_p - p_a}{t}}; \quad x_{pb}^{bln} |_{\lambda(x_{pb})=\lambda^*(x_{pb})} = 1 - \sqrt{\frac{p_p - p_b}{t}}. \quad (1)$$

This information allows us to compute the demand and expected profits of all agents that we use to derive in the Appendix the following Lemma:

Lemma 1. *The profits of the music labels and the platform for any given level of $\varepsilon(x_i)$ are*

$$\pi_a^{bln} = \frac{1}{18}t(3 + 2\varepsilon)^2; \quad \pi_b^{bln} = \frac{1}{18}t(3 - 2\varepsilon)^2; \quad \pi_p^{bln} = \frac{t(1 - 39\varepsilon^2)}{27};$$

and the indifferent consumers are located in:

$$x_{ap}^{bln} = \frac{1}{3} - \varepsilon; \quad x_{pb}^{bln} = \frac{2}{3} - \varepsilon; \quad x_{ab}^{bln} = \frac{1}{2} - \varepsilon;$$

Proof. See the appendix. ■

Finally, we proceed backward to the first stage of the game, when the platform announces and commits to a level of bias $\varepsilon(x_i)$. It follows from the Lemma 1 above that:

Proposition 1. *The equilibrium recommendation system with homogeneous quality is the one that recommends the efficient bundle to all consumers.*

Proof. See the Appendix. ■

To understand Proposition 1, one should look at the prices set by the music labels given the level of bias — i.e., $p_a^{bln} = t + \frac{2\varepsilon t}{3}$ and $p_b^{bln} = t - \frac{2\varepsilon t}{3}$. Interestingly, the bias exerts a positive effect on the price of the seller that it favors. Hence, if the two music labels offer bundles of the same quality, the demand differential implied by a biased recommendation system would result in higher prices for the content that the platform offers more to consumers. It goes without saying that the platform would never engage in such behavior, as it would be detrimental to her profitability. It follows that the only optimal recommendation schedule possible is the unbiased one. Put differently: when goods are homogeneous the platform always offers the *efficient bundle* to all consumers.

Corollary 1. *In equilibrium, all firms make positive profits. In particular, music labels obtain the standard Hotelling outcomes. Formally:*

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

Proof. The proof follows from plugging $\varepsilon = 0$ in the results derived in Lemma 1. ■

4 Heterogeneous quality

In this section, we relax the assumption that goods are homogeneous. Indeed, the music industry is both vertically and horizontally differentiated.¹³ Music labels experiment and research new ways of expressing their art. In other words, they innovate. It is, therefore, credible to assume that music labels compete with products that embed different quality levels. Therefore, we repeat the analysis assuming that music label b produces content of better quality that ensures higher utility to all consumers irrespective of their preferences for the available varieties. We first assume such quality differential to be a primitive of the model; afterward, we endogenize the choice of costly investment in quality to study which distortions, if any, the intervention of the platform leads to.

4.1 Exogenous quality differential

First, we assume $V_b = v + v_b > v = V_a$ with v_b exogeneously given. The utility functions become:

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

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

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

absent the price effect the quality gap v_b moves the indifferent consumers towards the location of the music label a , thus shrinking her demand: quality shifts demand¹⁴.

We define the *efficient bundle* in this scenario as:

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

where the apex ^{hq} stands for ‘‘heterogeneous quality’’.

¹³Music labels backed by major labels generally have more resources than comparable independent music labels to produce better products – e.g., in terms of sound quality or international collaborations.

¹⁴As we are interested in gradual innovations, we focus on values of v_b for which there is always at least a consumer that prefers the pure bundle b but would derive positive utility from mixing. Formally, $v_b < 2t/3$.

We use this information to update the location of the indifferent consumers:

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

The location of the indifferent consumers helps us compute the demand faced by each agent of the model. In particular, the platform's demand is simply given by $D_p = x_{pb} - x_{ap}$. Instead, the demands of the two sellers change because of the different proportions of content in the new biased bundles. Formally:

$$D_a^B = x_{ap} + \int_{x_{ap}}^{x_{pb}} (\lambda^{hq}(x_i) + \varepsilon(x_i)) dx \quad (3)$$

$$D_b^B = 1 - x_{pb} + \int_{x_{ap}}^{x_{pb}} (1 - \lambda^{hq}(x_i) - \varepsilon(x_i)) dx \quad (4)$$

The apex B indicates the scenario where the platform offers a biased mix to consumers.

From the above demand functions, it is possible to anticipate that music label b faces a demand which is decreasing in the intensity of the bias ε ; consequently, music label a faces an increased demand because of the favorable bias. This is analogous to what we observed in the baseline model, and it explains why in that case bias would not emerge in equilibrium. However, in this model specification, the variation in goods' quality makes the two demands asymmetric, to begin with, and the platform may have incentives to increase the demand for the music label that offers the cheapest good.

Because the variation in music labels' demands is known before prices are set, 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 as a consequence of the increased demand. Possibly, the two prices would converge towards a common value if the bias is sufficiently intense. Because the consumers demand more content from the high-quality seller, reducing its price is indeed in the interest of the platform, absent any constraint on the sellers' 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. In deciding the total mass of demand to shift from one seller to the other — i.e., the total bias — the platform pays attention that it does not exceed the sum of the PC of all the consumers (see Figure 2). 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 apexes p, c indicate the total bias selected by the platform and the maximum bias that satisfies consumers' participation constraint, respectively; $\varepsilon^p(x_i)$ represents the individual level of bias the platform designs.

In order to solve the problem, and since 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

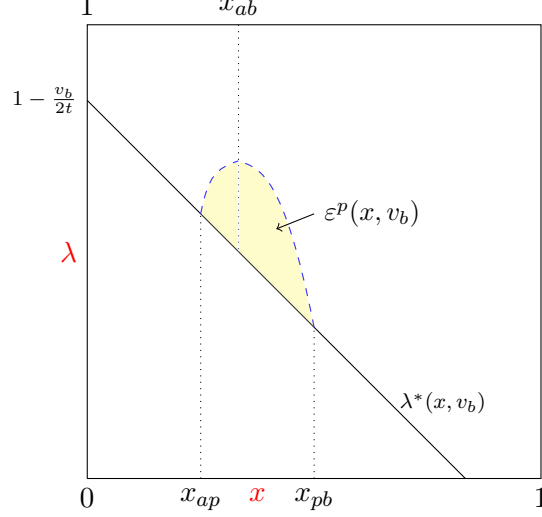


Figure 2: *Personalized biased bundle.* The area in yellow is the total demand that the platform can shift from music label b to music label a .

individual level of bias for each consumer that 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 consistency of the solution while maintaining the problem tractable.

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

$$\pi_p^B = p_p (x_{pb}^{hq} - x_{ap}^{hq}) - p_a \left(\int_{x_{ap}^{hq}}^{x_{pb}^{hq}} \lambda^{hq}(x_i) dx + \varepsilon^p \right) - p_b \left(\int_{x_{ap}^{hq}}^{x_{pb}^{hq}} (1 - \lambda^{hq}(x_i)) dx - \varepsilon^p \right) \quad (5)$$

$$\pi_a^B = p_a \left(x_{ap}^{hq} + \int_{x_{ap}^{hq}}^{x_{pb}^{hq}} \lambda^{hq}(x_i) dx + \varepsilon^p \right) \quad (6)$$

$$\pi_b^B = p_b \left(1 - x_{pb}^{hq} + \int_{x_{ap}^{hq}}^{x_{pb}^{hq}} (1 - \lambda^{hq}(x_i)) dx - \varepsilon^p \right) \quad (7)$$

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

$$p_a(\varepsilon^B) = t - \frac{v_b}{3} + \frac{2\varepsilon^p t}{3}; \quad p_b(\varepsilon^B) = t + \frac{v_b}{3} - \frac{2\varepsilon^p t}{3}; \quad p_p(\varepsilon^B) = \frac{10t}{9} + \frac{v_b^2}{4} - \varepsilon^p (v_b - \varepsilon^p t) \quad (8)$$

The next Lemma follows directly:

Lemma 2. *Consider the case in which the platform offers a biased mix $\lambda^{hq}(x, v_b) + \varepsilon^p$ to the consumers. Then the stage 2 equilibrium prices are as derived in (8), the profits of the music labels and the platform are*

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

$$\pi_p(\varepsilon^p) = \frac{t + 3\varepsilon^p (7v_b - 13t\varepsilon^p)}{27} - \frac{v_b^2}{36t}$$

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. ■

Before proceeding backward to derive the sellers' entry decision and the equilibrium bias, let us discuss the results in Lemma 2. The bias affects the two sellers in opposite ways. Seller a benefits from employing a biased mix, as it allows it to sell more of its content to the platform subscribers, mitigating the quality gap. As seen in the benchmark case, the increase in the demand for the seller a 's content generates a 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 reduction of demand on the direct channel.

Conversely, seller b suffers from 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 the pure bundle good b . However, the negative price effect and the reduced exposure of content b in the mixed bundle good dominate the demand expansion on the direct channel.

The platform does not lose demand but reshuffles its cost function more conveniently. It is worth mentioning that a positive bias $\varepsilon > 0$ 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 herself shifting demand toward the most expensive content, which is obviously non-optimal.

Recall that the bias is set before the game starts and that the platform commits to that level. Hence, once decided, it cannot be modified to adjust for the new ordering of the prices. Recall also that the bias cannot exceed the maximum one (Condition 1): the platform anticipates the effects of the bias on the entry decision of consumers and on pricing and sets it consistently with their participation constraints.

4.1.1 Sellers' participation decision

Let us now proceed backward and consider the sellers' participation decision given the bias ε^p . At stage 1, the sellers must decide whether to compete with each other in a market where $\alpha \in [0, 1]$ consumers are aware of their products and locations, or to join the platform and also reach the other $1 - \alpha$ consumers who are ex-ante unaware of the two sellers. As mentioned in Section 2, we refer to the former group as the "informed consumers". We refer to the latter group as the "uninformed consumers" instead. Uninformed consumers learn of the existence of the sellers or their relative position only if the platform is active, which can only happen if the platform manages to attract both sellers¹⁵.

¹⁵Consider the case in which only one seller $j = a, b$ joins the platform. Uninformed consumers learn about her and her position during the free trial of the streaming service. After the trial, they decide what to purchase (the subscription to the streaming service or the pure bundle). However, the platform operates as a retailer here (it only offers the pure bundle of the music label, as there are no other goods to include in the mix). Because the royalty rate is $r_j = p_j$, and royalties enter the cost structure of the platform, it must be that the subscription fee $p_p \geq p_j$, which means all consumers weakly prefer purchasing the pure bundle j directly by the seller.

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:

$$p_a^{out} = t - \frac{v_b}{3}; \quad p_b^{out} = t + \frac{v_b}{3}$$

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

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

When seller $j = a, b$ decides whether to join the platform, he 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, b is the seller that would be penalized if the platform biased consumption. It follows that it is sufficient to consider the participation decision of b to determine whether the platform can be active or not in equilibrium. This decision depends on the share of informed consumers, α , and the quality difference, v_b .

In the baseline specification with homogeneous goods ($v_b = 0$), it is clear that the platform is always active: because firms make the Hotelling profits in equilibrium (Corollary 1) they are strictly better off if they are exposed to the uninformed consumers. In the limit case in which $\alpha = 1$ (that is, there are no uninformed consumers), moreover, sellers are indifferent between joining or not; in this case, we assume that the indifference is split in favor of the platform, which can then become active.

The prediction changes drastically if the products are vertically differentiated. In particular, the ability of the platform to bias consumption is 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 consumers' participation constraints, addressed above, and the sellers' participation constraints.

Intuitively, the latter becomes stricter the higher α is. If there are many informed consumers, the high-quality seller has a stronger bargaining chip at the entry stage. Suppose that there are, in fact, no uninformed consumers — i.e., $\alpha = 1$. The platform's optimal bias policy would negatively affect music label b . Clearly, b anticipates that joining the platform does not expose his product to more consumers. Then, b would rationally choose not to join the platform if she commits to any positive level of bias. Therefore, the platform would reduce her optimal bias to zero to induce both sellers to join.

At the opposite limit, suppose that $\alpha = 0$: if the music label b does not join the platform, he cannot make any sale in the direct market as there are no consumers who are aware of her existence. Regardless of how biased the recommendation system is in favor of his rival, he would always optimally choose to join the platform. In turn, this implies that the platform is only constrained by the consumer participation decision when setting up her bias policy.

Formally, we can distinguish a threshold for α as a function of v_b and the chosen bias ε^p that sorts between 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 not active unless both sellers join the streaming service, she must then choose a bias such that:

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

where the apex s indicates it is the *sellers' constraint*.

4.1.2 Equilibrium Bias

We can finally proceed backward to stage 0 and determine the equilibrium level of bias that the platform includes in her recommendation system. From the analysis above, the problem of the platform can be written as

$$\begin{aligned} \max_{\varepsilon} \quad & \pi_p(\varepsilon) = \frac{t + 3\varepsilon(7v_b - 13t\varepsilon)}{27} - \frac{v_b^2}{36t} \\ \text{subject to} \quad & \varepsilon < \min\{\varepsilon^c, \varepsilon^s\} \end{aligned}$$

It is easy to observe that the unconstrained maximization leads to $\varepsilon^p = \frac{7v_b}{26t}$. Using the prices in (8) to evaluate the maximum bias consumers are willing to accept before leaving the platform, we obtain:

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

where the apex c indicates that it is the *consumers' participation constraint*. Using the location of indifferent consumers in Proposition 2, tedious calculations reveal that ε^c is decreasing in v_b .

Proposition 2. *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. Moreover, see Figure 3 for an illustration of the result. ■

In choosing the bias intensity, the platform trades off cost minimization and the intensity of competition. Since biasing consumption away from b leads it to optimally reduce p_b , the platform must update its own optimal price downwards because the better product has a stronger competitive effect than the worse one. The platform, therefore, wants to reduce the bias to set higher fees for consumers and raise the bias to minimize costs. The trade-off is optimally

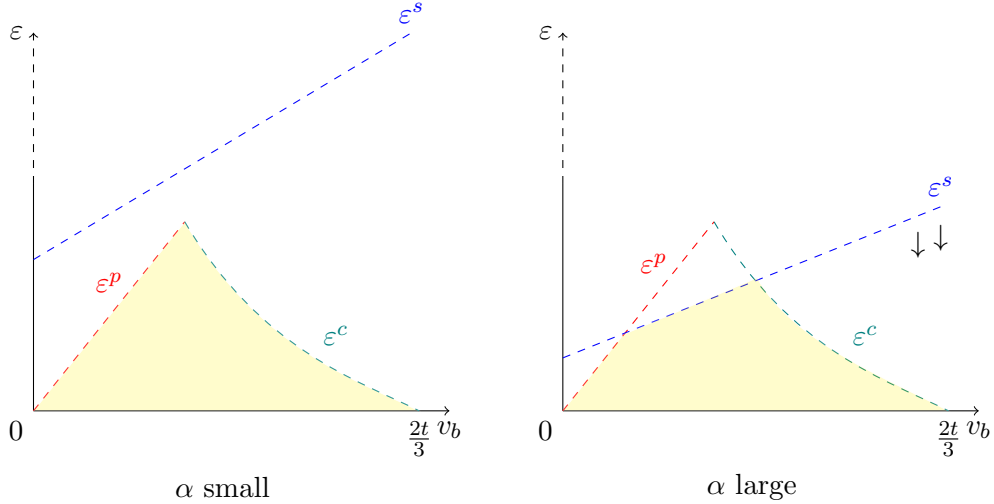


Figure 3: *The constrained equilibrium level of recommendation bias designed by the platform. When the share of informed consumers is sufficiently small (left panel), the optimal bias (ε^p) is bounded by the consumers' participation constraint (ε^c). When it kicks in, the platform needs to reduce the bias as quality increases because fewer users are willing to substitute content a with the high-quality good b . Otherwise (right panel), the platform must consider an additional constraint (ε^s) and needs to secure sellers' participation in the streaming service.*

resolved at ε^p . The dynamic echoes the result found in [Arya et al. \(2007\)](#): since 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 from the strategic use of the recommendation system discussed here.

Crucially, the platform is aware of several constraints. In particular, consumers and sellers must be offered conditions that make them want to join. Consumers' participation constraint (ε^c) arises because each biased bundle offered to each consumer must provide weakly more utility than what they could obtain outside of the platform.

Sellers face a different trade-off when choosing to join the platform or staying out: the more consumers are already aware of their existence and product, the less they benefit from exposure. Vice versa, if there are many uninformed consumers in the market, joining the platform boosts the visibility of the sellers and, consequently, their sales. This trade-off is incorporated in the constraint represented by ε^s . The higher α is, the harder it is to convince sellers to withstand the bias on the platform. This, in practice, limits the total bias that the platform can realistically design and commit to. Notably, this implies that the minimum level α that makes seller b strictly better off on the platform than outside of it is non-monotonic.

One can notice that the unconstrained bias is increasing in v_b : the larger the quality differential is, the stronger the effect it has on costs relative to the effect it has on prices, and the more the platform is interested in biasing consumption in equilibrium. Instead, the threshold ε^c decreases in v_b : the higher v_b is, the more consumers benefit from consuming the product offered by b , and the less flexible they are in accepting a biased bundle. To attract them, the platform must design its recommendation system such that the total bias never exceeds ε^c .

Finally, for the sellers to be better off on the platform, they need to be exposed to more new

consumers to compensate for the lower individual exposure. When the consumer participation constraints start binding — that is, for v_b high enough — ε^* decreases. As a consequence, the high-quality seller’s penalty generated by the platform gets smaller the better his product is since more consumers want to purchase more of its content.

Overall, the model highlights a subtle interaction between the business model of most streaming platforms and the incentive of creators to produce high-quality content. It is clear that absent the bias, sellers would have a strong incentive to innovate and compete with high-quality products: consumers value better content and are willing to pay more for it. The intermediation of the platform dilutes these incentives by introducing a bias that penalizes better content when it is more expensive to offer it. We assumed so far that V_j is exogenously given for both sellers. However, it is easy to extend the analysis to a hypothetical stage 0 in which the sellers had to costly invest in the creation of higher-quality content. In what follows, we show that unless the platform generates significant new demand (that is unless α is small), the incentive to costly invest in higher quality is lessened when the platform can bias consumption. Further, we show that the equilibrium allocation of R&D effort can be made inefficient by the platform’s intervention.

4.2 Endogenous investment in quality

We now endogenize the choice of sellers to costly invest in quality. To maintain the direction of the asymmetry studied above, we assume seller b always to be more efficient than seller a . We consider two scenarios: in the first, we assume seller a to be incapable of investing in quality, while b is free to invest in quality at the beginning of the game. Afterward, we assume both sellers simultaneously select their quality investment from the same baseline value $V_j = V$. For tractability reasons, we ignore agents’ participation constraints’ effect on equilibrium quality levels when solving for the equilibrium investment in quality¹⁶.

Formally, the timing of the interaction between sellers and the platform is 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 standard convex cost function $I(v_j) = \phi_j v_j^2$. Notice that in our simplified setting, a firm chooses how much to invest in R&D and these investments uniquely determine the quality of the good sold in the market, $v + v_j$. As a consequence, a firm maximizes profits by either choosing the amount of its investments, $I(v_j)$, or its degree of innovativeness, v_j ; for this reason, with a slight abuse of terminology, in what follows, we often refer to v_j as the level of investment in R&D chosen by seller j .

The sellers’ profits change to:

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

where apex eq stands for “endogenous quality”, $\phi_a = \infty$ or $\phi_a = 1$ (depending on whether a

¹⁶The aim of this exercise is to qualitatively evaluate the effect recommendation bias has on innovation investments. As discussed above, both consumer and seller participation constraint limits the role of the bias for high enough quality differential. We acknowledge this effect, but choose to ignore it for the sake of clarity of the exposition of the results.

is assumed to be able to invest in quality or not), and $\phi_b < 1$ captures the higher efficiency of seller b compared to a . We make the following technical assumption to ensure the existence of an interior solution:

Assumption 1. *The investment cost functions are sufficiently convex: $\phi_b > \tilde{\phi} \approx \frac{0.1183}{t}$.*

We update the formulas detailed in Lemma 2 by including the cost functions for the two firms and the cost differential $v_b - v_a$:

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

$$\pi_p^{eq}(\varepsilon^p) = \frac{t + 3\varepsilon^p(7(v_b - v_a) - 13t\varepsilon^p)}{27} - \frac{(v_b - v_a)^2}{36t}$$

As before, we obtain $\varepsilon^p = \frac{7(v_b - v_a)}{26t}$ by simple F.O.C. argument; we use the equilibrium platform bias to obtain equilibrium investment levels by plugging it in the equations for $\pi_a^{eq}(\varepsilon^p)$ and $\pi_b^{eq}(\varepsilon^p)$. Then, we proceed backward to stage 0, when investments in innovation are decided.

4.2.1 Only one seller invests

First, we directly extend the analysis brought forth, assuming v_b to be exogenously selected by endogenizing the choice of the more efficient seller with respect to the optimal quality differential on and off the platform. In this scenario, we impose $\phi_a = \infty$ to guarantee $v_a = 0$ in equilibrium. After plugging in $\varepsilon^p = \frac{7v_b}{26t}$ in the profit function of b , by F.O.C. we obtain the optimal investment in quality:

$$\left. \frac{\partial \pi_b^{eq}}{\partial v_b} \right|_{\varepsilon^p = \frac{7v_b}{26t}} = 0 \quad \iff \quad v_b^{eq} = \frac{13t}{169t\phi_b - 2}.$$

Unsurprisingly, v_b^{eq} is decreasing in ϕ_b : the more efficient seller b is, the higher the optimal quality differential, all else being equal. To give meaning to this value, however, we must compare it with the relevant counterfactual. First, one should consider what the optimal investment level would be in the sub-game in which sellers do not join the platform. In this case, only a fraction $\alpha \in (0, 1)$ of consumers is active. Thus, the rent that the sellers can generate depends on the size of this group. It is easy to show that the higher the number of informed consumers α is, the stronger the incentives to invest in quality outside the platform ecosystem. If we update the payoff of the innovative seller b when it decides not to join the platform ($\pi_b^{eq,out} = \pi_b^{out} - I(v_b)$), the standard maximization procedure yields the following value of investment:

$$\frac{\partial \pi_b^{eq,out}}{\partial v_b} = 0 \quad \iff \quad v_b^{eq,out} = \frac{3\alpha t}{18t\phi_b - \alpha}.$$

As anticipated, for all (valid) values of ϕ_b , it holds:

$$\frac{\partial v_b^{*,out}}{\partial \alpha} > 0, \quad \text{and} \quad v_b|_{\alpha=0} = 0.$$

Second, we compare our equilibrium result with the profit-maximizing level of investments assuming that the platform is not able to bias the recommendation system ($\varepsilon^p = 0$). In this

hypothetical scenario, both sellers have a clear incentive to join as argued in the previous subsection. Furthermore, as shown above, when the platform does not introduce a bias, the competitive stage is equivalent to a standard Hotelling framework (Corollary 1).

It follows that in this scenario firm b optimally invests:

$$v_b^{eq,nobias} = v_b^{eq,out}|_{\alpha=1} > v_b^{eq} \quad \forall \phi_b,$$

where the last relationship is established by direct comparison.

In summary, firm b anticipates the effect of the bias and lowers its investment in quality v_b . The result above highlights that for some α joining the platform may reduce the incentives to invest in innovation. This happens if the share of informed consumers is sufficiently large. In this case, the positive shock in demand from the additional uninformed users who only listen to music via the platform is insufficient to compensate for the demand shift towards the low-quality rival entailed by the biased recommendation system. However, for low enough values of α , the opposite holds. The increased demand creates strong enough incentives to invest despite the distortion introduced by the platform.

4.2.2 Both sellers invest

Suppose now that both sellers are able to costly invest in quality, that is, $\phi_a = 1 > \phi_b$. After plugging in ε^p in the equations above, it is once again straightforward to obtain the equilibrium levels of investment by simple F.O.C. argument. It is also immediate to obtain the equilibrium level of investment in the sub-game in which the platform is inactive. Remember that, in this case, only a fraction α of consumers is active and sellers compete in standard Hotelling. Then:

Lemma 3. *Consider the case in which both sellers can invest in quality, and seller b is more efficient than seller a by a factor ϕ_b . Equilibrium levels of investment on the platform are:*

$$v_a^{eq} = \frac{169t\phi_b - 4}{(2197t - 26)\phi_b - 26}; \quad v_b^{eq} = \frac{169t - 4}{(2197t - 26)\phi_b - 26};$$

equilibrium levels of investment outside the platform is:

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

and equilibrium levels of investment, if the platform is assumed not to bias, is:

$$v_a^{eq,nobias} = v_a^{eq,out}|_{\alpha=1}; \quad v_b^{eq,nobias} = v_b^{eq,out}|_{\alpha=1}.$$

Proof. See the Appendix. ■

The bias imposed by the platform distorts the incentives to innovate even if we allow both sellers to invest in innovation. In particular, for given demand, the intervention of the platform dilutes the incentives of seller b to innovate, which in turn leads to an equilibrium investment v_a lower than the one a would have selected in the absence of bias. This can be seen by direct comparison of v_a^{eq} and $v_a^{eq,nobias}$, and of v_b^{eq} and $v_b^{eq,nobias}$. The result is once tempered when

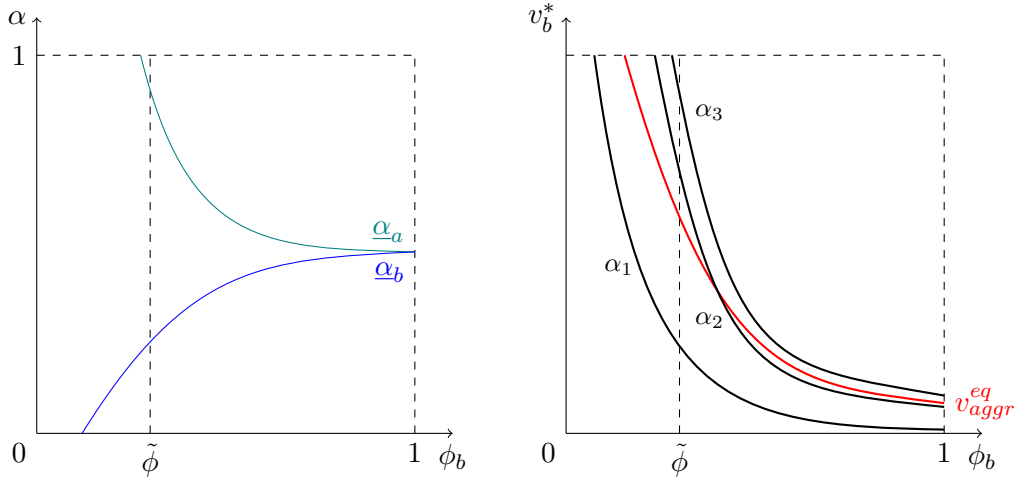


Figure 4: *On the left: the minimum value $\underline{\alpha}_a$ and $\underline{\alpha}_b$ such that, for $\alpha > \underline{\alpha}_j$, seller j would have invested more in quality outside of the platform; On the right: aggregate investment level for increasing values of α ($\alpha_1 = 0.1$ $\alpha_2 = 0.45$ $\alpha_3 = 1$). $t = 1$ in both graphs.*

we account for the additional demand generated by the platform: if α is low enough, joining the platform exposes the labels to enough new demand that, despite the bias, equilibrium investment in quality is higher on the platform than off of it.

The equations in Lemma 3 allow us to make some statements regarding the overall effect on aggregate investment of quality, which can be shown to crucially depend on the relative values of α and ϕ_b . In particular, the lowest value α such that equilibrium aggregate investments are higher off the platform than on it is increasing in ϕ_b : the more marked the difference in efficiency is, the stronger the distortion on equilibrium investment levels introduced by the platform is. Vice versa, for α high enough, aggregate investment would be higher outside of the platform than it is on the platform. Furthermore, For all values of ϕ_b , there are values α such that b would have invested more and a would have invested less outside of the platform than on it. The last remark implies that platform intervention can also lead to inefficient allocation of R&D effort by innovators, with efficient (resp. inefficient) sellers investing less (resp. more) than they would have without the bias in the recommendation system. The results are illustrated in Figure 4 and summarized as follows:

Proposition 3. *If the platform does not generate enough new demand, recommendation bias leads to lower aggregate investment in innovation; the distortion is stronger when the difference in efficiency between labels is not too high. Furthermore, for all values $\phi_b \in (\tilde{\phi}, 1)$, there exist values α such that equilibrium allocation of R&D effort is inefficient.*

Proof. See the Appendix. ■

5 Extensions

We extend the analysis in several directions. First, we consider the implications of letting consumers costly produce their own favorite mix of content on the platform. We assume that

the associated cost is a design choice of the platform and, therefore, is selected strategically at the beginning of the game. We show that the platform has the incentive to increase the said price to force consumers to use her recommendation system, which follows seamlessly from the analysis above.

Second, 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 of the game but, rather, after the agents have selected prices. We take a reduced form approach and show that the platform has the incentive to bias more than in the baseline model since agents cannot condition prices on the equilibrium recommendation system.

Finally, we consider a different source of vertical differentiation in the form of asymmetric costs of production by the sellers. We show that since the most efficient seller selects lower prices in equilibrium, the platform optimally biases consumption towards him. On one hand, this suggests that streaming platforms reward efficiency. We argue however that this, too, indicates that streaming platforms penalize experimentation and, instead, create the incentive to produce commodified content to avoid the added penalty linked to inefficiency.

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

In real-world examples, this is hardly the case. Consumers, once they join a streaming platform, have some freedom in choosing the music or digital content they want to consume.

It is however apparent that the extent to which it is easy to search for the desired digital content is part of the design of a platform, of its architecture. Hence, a natural question arises: what is the optimal “degree of freedom” that the platform should give to consumers to build their own consumption mix?

In this extension, we address this question by assuming that the consumer may (or may not) decide to pay a cost in order to get the efficient bundle and get rid of any recommendation bias imposed by the platform.

Formally, the utility of the consumers when they join the platform becomes:

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

where k is the search cost that consumers must incur in order to avoid recommendation bias. Keeping in mind that the efficient bundle is $\lambda^{hq}(x_i) = 1 - x_i - \frac{v_b}{2t}$, it is possible to derive the cost level that makes any consumers on the platform indifferent between building consuming the efficient mix and get the biased bundle:

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

If unconstrained, the platform can simply design the search process in a way such that $k = \bar{k}_i$

for the consumer i who is willing to accept the largest bias (who is located in x_{ab}). Thus, the consumers are induced to accept the bias set by the platform, and the analysis in the previous section goes through. Otherwise, if constrained in some ways, the platform has to set a smaller k . In such a case, the consumers located closer to x_{ab} will be willing to pay the search cost in order to get the efficient mix. Thus, the ability of the platform to bias the recommendation is limited to consumers at the margins of the platform's demand.

Under our assumption of a monopoly platform, we observe no incentives to design its search process in a way that limits its ability to bias the recommendation system and save on costs. However, how efficient the search process looks to consumers may be a relevant characteristic when competition between platforms kicks in. We leave this question open to further research.

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 issue assuming that the platform does not commit to a bias level, but firms are aware of its incentives to steer demand toward the cheapest product. Hence, the new timing is the following: 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 realize.

In the baseline specification with homogeneous goods, the platform offers an efficient mix to all consumers, and the game unfolds as in Proposition 1. However, Proposition 2 reveals that prices are not symmetric if there is a variation in the qualities of the goods offered. Hence, the platform finds biasing its recommendation system profitable, as there is a demand mass that can be steered toward the cheapest (inferior) good without impacting the number of consumers joining the streaming service. Moreover, as sellers are not able to react after the choice of the recommendation bias, the platform would always steer as much demand as possible, subject to the consumers' participation constraint (see the definition of $\bar{\varepsilon}(x_i)$ in Section 4).

The sellers anticipate this incentive and adjust their prices accordingly. On the one hand, the seller of the superior good reacts to the anticipated bias by lowering her price. On the other hand, the seller of the inferior good, anticipating the demand expansion following the recommendation bias, has the incentive to increase its price. Eventually, this adjustment occurs until profitable.

One should notice that an equilibrium exists only provided that the prices do not "cross" - i.e., provided that the most valuable good is also the most expensive one.

Otherwise, assume the sellers anticipate the bias and adjust the prices to such an extent that the inferior good is now as expensive as the superior one. The platform observes the prices and optimally 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.

In other words, consider the prices $p_a(\varepsilon)$ and $p_b(\varepsilon)$ chosen by the sellers in anticipation of

the total recommendation bias. We can state the following:

Lemma 4. *Assume that $p_a|_{\varepsilon=0} < p_b|_{\varepsilon=0}$. Then if the maximum bias that 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 $p_a(\varepsilon^c) \geq p_b(\varepsilon^c)$, an equilibrium in pure strategy does not exist.*

One should notice that because ε^c is decreasing in v_b , we can say that an equilibrium in pure is more likely to emerge when v_b is sufficiently large.

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:

$$\pi_i = D_i(p_i - C_i), \quad i \in \{a, b\}$$

where C_i 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 main model, we assume that $C_a = c_a > 0$, $C_b = 0$.

The framework differs from the main model in two substantial ways. First, since the asymmetry lies in the costs rather than the value consumers attach to the content, the optimal mix of consumers is not affected by the asymmetry, and $\lambda^*(x_i) = 1 - x_i$. Second, when the better seller is more efficient rather than offering higher quality content, the price he would optimally select in the absence of bias would be lower than the one offered by his competitor. The platform, therefore, would have an incentive to penalize the worst of the two sellers with her biased recommendations rather than the best one as it was in the main model.

The analytical steps to solve the model mirror the ones made explicit for the main model. As before, the difference in equilibrium prices is tempered by the intervention of the platform: ε^* is selected to reduce the distance between p_a and p_b . Unlike in the main model, however, since $C_a > C_b$, $\varepsilon^* \leq 0$: the platform introduces bias to boost consumption of seller b 's content. In doing so, she 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 reduce operational costs and reduce to bias to increase her own subscription fee p_p .

Finally, since 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, consumers do not want to consume the content of one of the two sellers disproportionately. Since the bias is introduced to penalize the content of the least efficient seller, and since this seller charges a higher price in equilibrium because of this inefficiency, consumers are less sensitive to the bias than before. It follows that the constraint represented by ε^c is still decreasing in the cost differential, but is less tight than in the main model.

The constraint relevant to induce sellers to join is also looser than the one considered in the main model. Unlike before, since the bias penalizes the inefficient seller, the highest possible bias the platform can introduce must make the worse seller indifferent between joining or not, rather than the better one as it was before. From standard Hotelling logic, the seller with higher marginal costs would see his profits decrease in the cost differential. Since the seller penalized

by the bias is the one with the worse outside option, then, it is clear that the platform can ignore the constraint represented by ε^s for a wider range of values α .

We can state the following result:

Proposition 4. *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 his 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. On one hand, it highlights the difference between vertical differentiation driven by consumer taste and efficiency: in particular, 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 might lead to the partial conclusion that streaming platforms' intervention might be socially desirable since it creates the incentive to reduce marginal costs or production and, with them, equilibrium prices.

On the other hand, however, we believe that modeling costs as we have proxies the choice of music labels to experiment with their content, which can be expected to increase costs of production, instead of optimizing the process of creating said content. From this perspective, the outcome of the exercise is well in line with the one presented in the main model: the platform discourages risk and, instead, rewards “assembly line” production of content. The overall takeaway, then, becomes straightforward: if we assume as true that novelty and experimentation are costlier than producing commodified content, the platform always has the incentive to penalize it through strategic manipulation of what consumers are exposed to.

6 Discussion and conclusion

In this paper, we study the incentives of a streaming platform to bias bundling in an effort to achieve optimal economic conditions. The platform has the potential to generate utility for consumers that value a balanced mix of content. When content is of equal quality, sellers select uniform prices, and the platform has no incentive to bias consumption. When sellers offer vertically differentiated products, instead, they have the incentive to set different royalties. In particular, the seller with the higher quality product wants to raise royalties since consumers value his product more. When this happens, the platform has an incentive to bias consumption towards the “cheaper”, lower quality product to minimize costs. This comes at the detriment of consumers, that lose the additional utility generated through efficient content mixing, and the higher quality seller, who sees his demand artificially shrunk. In equilibrium, the latter would set a lower price than without intervention: the platform dampens the incentive to introduce higher

quality products by punishing them with reduced exposure. Further, platform intervention can significantly distort equilibrium R&D efforts.

Based on several real-life examples, we assumed that the platform cannot price discriminate consumers. If she could, it is clear that she would have the incentive to offer different bundles at different prices in an effort to extract the rent she helps generate. The ability to price discriminate does not eliminate the incentive to bias. However, since consumers must be convinced to join the platform, personalized pricing would remove 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 be as well off if products were vertically differentiated; they would also be strictly worse off if products were of the same quality. The reason is straightforward: the platform has no incentive to bias consumption under the baseline specification, but she would still have the incentive to price discriminate if it was possible.

More subtly, the result is carried forward by the assumption of sellers bargaining their royalty rate individually. The incentive to bias consumption follows directly from the difference in cost for the platform to stream the content of the sellers. Suppose, however, that the sellers were both represented by an intermediary, such as a copyright collecting agency, bargaining royalty rates for both. It is clear that such an agent would have the incentive to set equal royalties to reduce the incentive to bias consumption towards the cheaper product. It is less clear that this would not be to the detriment of the higher quality product's seller.

The model, as it stands, is limited by the assumption that sellers set prices and royalties equally. Separating the two, letting sellers set prices outside the platform and royalty rates inside of it, might lead to new insights.

On the one hand, it is natural to think that the platform would exploit this opportunity by inducing sellers to compete for its internal demand. The sellers would, accordingly, lower their royalties, as they would be trapped in a prisoner-dilemma-like situation. Possibly, the platform's subscription fee may approach the prices of the seller, forcing them to lower their prices also on the captive segments.

On the other hand, one can also envision that sellers may have the incentives to increase their prices in their *captive* segments to increase the surplus extraction from their loyal consumers. In order for this strategy to be profitable, the sellers have to increase their royalties, which enters the platform objective function as marginal costs. By doing so, the sellers ensure their price cannot be matched by the platform and operate as monopolists in their captive segments. The platform might be limited in its possibility to bias demand in this case, as sellers would be using royalties to increase their rival's costs rather than to extract surplus directly. A formal analysis of these scenarios is left as an open question for further research.

It is important to stress that the mechanism studied in this paper requires the platform's algorithmic component to be relevant. In a world in which consumers had no access to automatically generated content susceptible to manipulation but were always in perfect control of what they consume, the distortions predicted by the model would not bite. From the opposite perspective, if the algorithmic recommendation of streaming platforms were provably biased in a way that damages consumers, regulatory intervention would prevent such distortions from

arising. The broadest takeaway from the paper follows in the footsteps of many others, calling for inspection and direct regulation of the algorithms used by digital platforms to provide their service. While we acknowledge that this might disincentivize R&D expenditure and innovation to ameliorate these algorithms, perhaps the loss could be more than compensated by the stronger incentives to innovate on the content that algorithms would no longer be able to penalize.

It must be noticed that, 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 the investments from the *potential* level the seller would select absent the recommendation bias.

In doing so, we take an overly benevolent view towards the platform, which, by construction, never directly harms sellers. Consistently, the shares of informed and uninformed consumers are exogenously determined and, in our model, proxy the popularity of music labels. Debuting music labels are unknown by the large public and would likely be unable to thrive outside the platform service. On the contrary, 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 or not.¹⁷

Evaluating the net effect on the music industry imposed by the advent of the platform is well beyond the scope of our model. Here, our goal is to understand whether the platform can distort sellers' incentives to invest in quality by changing the algorithmic design of the recommendation system.

Abstracting from the model, it is important to consider the implications of the ability of the platform to make content accessible to more consumers. It is difficult to argue that streaming platforms do not represent a significant portion of the market they host. It follows that 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 ability of the platform 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 weaken. The model, then, highlights a potential risk embodied in the platform ecosystem: if a significant share of the users is held "captive" by the platform, the content available for streaming may become more commodified, which occurs at a loss for society that is difficult to quantify.

¹⁷Notably, Joni Mitchell and Neil Young's collections are not available on Spotify since 2022. Their motivation to delist from Spotify was not driven by commercial disputes but rather by the debate of Covid-19 misinformation in Joe Rogan's controversial podcast streamed on the platform. Yet, their decision to abandon the streaming service reflects the larger freedom celebrities enjoy vis-a-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.
- Armstrong, M. (2006). Competition in two-sided markets. *The RAND journal of economics*, 37(3):668–691.
- Arya, A., Mittendorf, B., and Sappington, D. E. (2007). The bright side of supplier encroachment. *Marketing Science*, 26(5):651–659.
- Berton, F., Dughera, S., and Ricci, A. (2021). Are Unions Detrimental to Innovation: Theory and Evidence. Technical report, IZA Discussion Papers 14102, Institute of Labor Economics (IZA).
- 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.
- De Corniere, A. and Taylor, G. (2019). A model of biased intermediation. *The RAND Journal of Economics*, 50(4):854–882.
- 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.
- Grout, P. A. (1984). Investment and Wages in the Absence of Binding Contracts: A Nash Bargaining Approach. *Econometrica*, 52(2):449–460.
- Haucap, J. and Wey, C. (2004). Unionisation structures and innovation incentives. *The Economic Journal*, 114(494):C149–C165.
- 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.
- Kline, P., Petkova, N., Williams, H., and Zidar, O. (2019). Who Profits from Patents? Rent-Sharing at Innovative Firms. *The Quarterly Journal of Economics*, 134(3):1343–1404.
- 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.
- Mukherjee, A. and Pennings, E. (2011). Unionization structure, licensing and innovation. *International Journal of Industrial Organization*, 29(2):232–241.
- 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 tail 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.
- Teh, T.-H. and Wright, J. (2022). Intermediation and steering: Competition in prices and commissions. *American Economic Journal: Microeconomics*, 14(2):281–321.
- Thomes, T. P. (2013). An economic analysis of online streaming music services. *Information Economics and Policy*, 25(2):81–91.
- Tsay, A. A. and Agrawal, N. (2004). Channel conflict and coordination in the e-commerce age. *Production and operations management*, 13(1):93–110.

A Appendix

Proof of Lemma 1 and Proposition 1.

Proof. From the indifferent consumers' locations as derived in expressions (1), we build the demands of the three firms. First, the demand platform is simply composed by the consumers who join the platform, i.e., $D_p^{bln} = x_{pb}^{bln} - x_{ap}^{bln}$. Instead, the expressions of the two music labels' demands are more complex, as 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. Formally:

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

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

What matters from the platform perspective is not the individual level of bias that each consumer support, but the total mass of demand that, via the biased recommendation, it is able to shift from one seller to the other. In other words, provided that the total mass of demand does not exceed the aggregate participation constraint of the consumers, we can treat it as a uniform value ε .¹⁸ Accordingly, the demand functions of the two music labels change as follows:

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

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

We use these demands and the consumers' locations in expression (1) to obtain the profit functions of the three firms:

$$\pi_a^{bln} = \frac{pb - pa + t^2}{2t^2} p_a, \quad \pi_b^{bln} = \frac{pa - pb + t^2}{2t^2} p_b,$$

$$\pi_p^{bln} = \frac{t - \sqrt{p_p - p_a} - \sqrt{p_p - p_b}}{t} p_p - \frac{p_b - p_a - t(2\sqrt{p_p - p_a} - t)}{2t^2} p_a - \frac{p_a - p_b - t(2\sqrt{p_p - p_b} - t)}{2t^2} p_b$$

Simple maximization with respect to the prices yields the following:

$$p_a^{bln} = t + \frac{2\varepsilon t}{3}, \quad p_b^{bln} = t - \frac{2\varepsilon t}{3}, \quad p_p^{bln} = \frac{10t}{9} + \varepsilon^2 t$$

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

We can now go backward to the first stage of the game, where the platform sets and commits

¹⁸For a thorough discussion of the conditions on the bias, refer to Section 4.

to a specific intensity of bias. From Lemma 1, the problem of the platform is the following:

$$\max_{\varepsilon} \pi_p^{bln} = \frac{t(1 - 39\varepsilon^2)}{27}$$

It is straightforward to observe that the function has a unique maximum in $\varepsilon = 0$. This proves Proposition 1. \blacksquare

Proof of Lemma 2 and Proposition 2

Proof. Proof of Lemma 2 follows a similar logic applied to prove Corollary 1. Assume the platform biases her recommendation system by transferring additional demand ε to the firm with low-quality content (firm a). The new demand functions are

$$\begin{aligned} D_a^{\varepsilon,hq} &= x_{ap}^{\varepsilon,hq} + \int_{x_{ap}^{hq}}^{x_{pb}^{\varepsilon,hq}} \lambda^*(x) dx + \varepsilon \\ D_b^{\varepsilon,hq} &= 1 - x_{pb}^{\varepsilon,hq} + \int_{x_{ap}^{\varepsilon,hq}}^{x_{pb}^{\varepsilon,hq}} (1 - \lambda^*(x)) dx - \varepsilon \\ D_p^{\varepsilon,hq} &= x_{pb}^{\varepsilon,hq} - x_{ap}^{\varepsilon,hq} \end{aligned}$$

We use the location of indifferent consumers derived in (2) and apply the same backward induction logic used in the case with no bias. From the system of the first-order condition, the profit-maximizing prices are as derived in equation (8). We use those prices in the payoff functions in equations (5)-(6)-(7) and in (2) to derive Lemma 2.

As mentioned in the main text, we treat ε as a general mass of demand that the platform can shift from one seller to the other. One should notice that, individually, each consumer is not willing to accept a mix that contains too much content a — i.e., $\varepsilon(x_i) \leq \bar{\varepsilon}(x_i)$. We account for that assuming that the platform can re-distribute the bias towards consumers according to their participation constraints. In other words, the condition 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 satisfied the participation of the consumers as derived in expression (9).

Taking all these conditions into consideration, we can now proceed backward to the first stage of the game. As before, from Lemma 2, the problem of the platform is the following:

$$\begin{aligned} \max_{\varepsilon} \pi_p^{hq} &= \frac{t + 3\varepsilon^p(7v_b - 13\varepsilon^p)}{27} - \frac{v_b^2}{36t} \\ s.t. \quad \varepsilon^p &< \min\{\varepsilon^c, \varepsilon^s\} \end{aligned}$$

Standard maximization yields the unconstrained profit-maximizing level of bias: $\varepsilon^p = \frac{7v_b}{26t}$. We use this value and prove Proposition 2. \blacksquare

Proof of Lemma 3 and Proposition 3 The proof of Lemma 3 is identical to the ones provided for Lemma 2 up until stage 0, with two adjustments: quality differential, $v_b - v_a = v_b$, is now equal to unconstrained $v_b - v_a$, and sellers' profit is reduced by $I(v_j) = \phi_j v_j^2$.

From the proof of Lemma 2, we know that adjusted unconstrained bias can be written as:

$$\varepsilon^p = \frac{7(v_b - v_a)}{26t},$$

which is positive (resp. negative) if $v_b > v_a$ (resp. if $v_b < v_a$): the bias favors the lower quality product. Plugging in this expression in equations 10, one obtains:

$$\begin{aligned}\pi_a^{eq} &= \frac{2(v_a - v_b)^2}{169t} + \frac{t}{2} + \frac{1}{13}((2 - 13v_a)v_a - 2v_b) \\ \pi_b^{eq} &= \frac{(13t + 2(v_b - v_a))^2}{338t} - \phi_b v_b^2\end{aligned}$$

The system of equations that solves these two expressions together leads to the first set of equations in Lemma 3. The second set of equations comes from F.O.C. of the system of the following equations:

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

Proposition 3 presents two results: aggregate investment in quality is higher off the platform than on the platform for α high enough, and for all values of ϕ_b there exist values of α such that b invests more off the platform than on it, and a invests less off the platform than on it. We prove the two results one at a time.

For the former: aggregate investments on and off the platform are simply:

$$\begin{aligned}v_{aggr}^{eq} &= v_a^{eq} + v_b^{eq} = \frac{169t(1 + \phi_b) - 8}{13((169t - 2)\phi_b - 2)} \\ v_{aggr}^{eq,out} &= v_a^{eq,out} + v_b^{eq,out} = \frac{\alpha(9t(1 + \phi_b) - 2\alpha)}{3((18t - \alpha)\phi_b - \alpha)}\end{aligned}$$

The latter equation is increasing in α : the numerator is increasing by virtue of the fact that $\alpha \in (0, 1)$; the denominator shrinks in α . We can then find the value $\underline{\alpha}$ that equates v_{aggr}^{eq} and $v_{aggr}^{eq,out}$:

$$\begin{aligned}\underline{\alpha} &= -\frac{(1 + \phi_b)(-19773t^2\phi_b - 273t(1 + \phi_b) + 24)}{52((169t - 2)\phi_b - 2)} + \\ &\quad + \frac{\sqrt{9(1 + \phi_b)^2(8 - 13t((507t + 7)\phi_b + 7))^2 - 5616t\phi_b((169t - 2)\phi_b - 2)(169t(1 + \phi_b) - 8)}}{52((169t - 2)\phi_b - 2)}\end{aligned}$$

Finally, it can be shown that $\underline{\alpha} \in (0, 1)$ for ϕ_b is high enough, and strictly increasing in ϕ_b , which proves the first result.

To prove the second statement we follow the same intuition: first, we find $\underline{\alpha}_a$ and $\underline{\alpha}_b$ at

which seller a and b respectively would have invested the same on and off the platform:

$$\alpha_a = \frac{39t(507t+7)\phi_b^2 + 3(91t-4)\phi_b - 3}{26((169t-2)\phi_b - 2)} + \frac{\sqrt{(\phi_b(13t((507t+7)\phi_b+7)-4)-4)^2 - 312t\phi_b(169t\phi_b-4)((169t-2)\phi_b-2)} - 12}{26((169t-2)\phi_b - 2)}$$

$$\alpha_b = \frac{3(6591t^2\phi_b + 91t(1 + \phi_b))}{26((169t-2)\phi_b - 2)} + \frac{3\left(\sqrt{4(1 + \phi_b) - 13t((507t+7)\phi_b+7)^2} - 312t(169t-4)\phi_b((169t-2)\phi_b-2) + 4(1 + \phi_b)\right)}{26((169t-2)\phi_b - 2)}$$

It can be showed that $\lim_{\phi_b \rightarrow 1} \alpha_a = \lim_{\phi_b \rightarrow 1} \alpha_b$, and in particular:

$$\lim_{\phi_b \rightarrow 1} \alpha_a = \lim_{\phi_b \rightarrow 1} \alpha_b = \frac{3}{26} \left(\frac{6591t^2 - 494t + 8}{4 - 169t} + 39t + 2 \right) = \frac{6}{13}$$

Finally, we show that α_a (resp. α_b) is decreasing (resp. increasing) in ϕ_b , which proves the result. To simplify the proof, assume $t = 1$ without loss of generality. One can see that in this case, α_a exists provided that $\phi_b \geq \tilde{\phi} \approx 0.1287$. Moreover, the two thresholds become:

$$\alpha_a|_{t=1} = \frac{3\left(\phi_b(6682\phi_b + 87) - 4 - \sqrt{\phi_b(\phi_b(52\phi_b(858637\phi_b - 146979) + 267985) - 3192) + 16}\right)}{26(2 - 167\phi_b)}$$

$$\alpha_b|_{t=1} = \frac{3\left(\sqrt{4(1 + \phi_b) - 13(514\phi_b + 7)^2} - 51480\phi_b(167\phi_b - 2) - 6591\phi_b - 87(1 + \phi_b)\right)}{26(2 - 167\phi_b)}$$

The expressions of the first derivatives w.r.t. ϕ are cumbersome, but with some algebra, it is possible to show that they write:

$$\frac{\partial \alpha_a|_{t=1}}{\partial \phi_b} \equiv \frac{495(27602\phi_b + 393)}{2(2 - 167\phi_b)^2 \sqrt{3999836\phi_b^2 + 140548\phi_b + 841}} - \frac{6435}{2(2 - 167\phi_b)^2} > 0$$

$$\frac{\partial \alpha_b|_{t=1}}{\partial \phi_b} \equiv \frac{771\phi_b(167\phi_b - 4) + 57}{(2 - 167\phi_b)^2} + \frac{3\phi_b(\phi_b((31414589 - 286784758\phi_b)\phi_b - 881874) + 10363) - 60}{(2 - 167\phi_b)^2 \sqrt{\phi_b(\phi_b(52\phi_b(858637\phi_b - 146979) + 267985) - 3192) + 16}} < 0$$

the inequalities hold $\forall \phi_b \geq \tilde{\phi}$.

Extension: On-platform search

Proof. The level of search costs that makes the platform's users indifferent between accepting the bias or paying in order to get the efficient mix is obtained by solving the following:

$$v + (1 - \lambda(x_i) - \varepsilon(x_i))v_b - p_p - t(x_i - (1 - \lambda(x_i) - \varepsilon(x_i)))^2 = v + (1 - \lambda(x_i))v_b - p_p - k - t(x_i - (1 - \lambda(x_i)))^2$$

$$k < \varepsilon(x_i)(v_b - t(2(1 - \lambda(x_i) - x) - \varepsilon(x_i)))$$

Using the efficient mix bundle $\lambda^{hq}(x_i) = 1 - x_i - \frac{v_b}{2t}$ into the above threshold, we obtain:

$$k < \bar{k} \equiv t(\varepsilon^2(x_i))$$

■

Extension: Different timing - Lemma 4

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, 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 good a without altering D_p .

The bias enters the profit functions of the sellers by altering their demand function. In fact, $D'_{a,\varepsilon} > 0$ and $D'_{b,\varepsilon} < 0$. The two sellers anticipate the bias and modify their prices accordingly. Seller a , who is benefiting from the demand shock, increases the price to $p_a(\varepsilon^c) > p_a$, whereas seller b lower her price to $p_b(\varepsilon^c) < p_b$. This is so 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 good a .

Intuitively, the two scenarios described in the Lemma emerge. First, the demand shift is not enough to change the ranking of the prices. In this case, the resulting equilibrium is such that

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

such that seller a and the platform are better off. In contrast, seller b and consumers are worse off (increasing the recommendation bias lowers consumer surplus).

Second, the demand shift is so significant that the ranking of prices changes. In such a case, an equilibrium in pure strategy no longer exists. In fact, in anticipation of ε^c , sellers change their prices but 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 cannot be an equilibrium. ■

Extension: Asymmetric costs - Proposition 4

Proof. We derive the relevant functions as we did for the main model. Since by assumption it holds $V_a = V_b = v$, we use the indifferent consumer as derived in expression (1). The demands of sellers are likewise the ones derived in the baseline with homogeneous goods. Demand of the platform is given by $D_p^{dc} = x_{pb}^{dc} - x_{ap}^{dc}$, where the apex dc stands for “different costs”.

Since by assumption it holds $c_a > c_b = 0$, profit function of a must account for the marginal

cost and in particular:

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

Profit functions of b and p are unchanged:

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

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

We notice that a will see the platform bias consumption away from him since, 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:

$$p_a^{dc} = \frac{2c_a}{3} - \frac{2t\varepsilon}{3} + t, \quad p_b^{dc} = \frac{1}{3}(c_a + t(2\varepsilon + 3)),$$

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

and profits:

$$\pi_a = \frac{(c_a + t(2\varepsilon - 3))^2}{18t}, \quad \pi_b = \frac{(c_a + t(2\varepsilon + 3))^2}{18t}, \quad \pi_p = \frac{t}{27} - \frac{(c_a - 52t\varepsilon)(c_a - 4t\varepsilon)}{144t};$$

The latter equation immediately leads to the platform profit maximizing bias by standard F.O.C. argument:

$$\varepsilon^p = \frac{7c_a}{52t}$$

To make consumers join, it must hold the following:

$$\varepsilon^p < \varepsilon^c = \int_{x_{ap}^{dc}}^{x_{pb}^{dc}} \bar{\varepsilon}(x) dx$$

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

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

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

ε^p is constrained by ε^s satisfying:

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

The result as stated in Proposition 4 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.