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

Jacopo Gambato* Luca Sandrini†

March 14, 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 has the incentive to bias recommendations in favor of the cheapest content. The distortion hurts consumers and the seller with the best content. The platform always introduces a bias in equilibrium when products are vertically differentiated but can only be active if it generates enough additional consumption to induce sellers to join its streaming service despite the bias.

Keywords: Platform Economics, Media Economics, Recommendation Bias, Innovation

JEL Codes: D4, L1, L5

*University of Mannheim and ZEW Mannheim, Jacopo.Gambato@zew.de
†Budapest University of Technology and Economics, sandrini.luca@gtk.bme.hu

1 Introduction

We study the role of a streaming platform in enabling consumers to "mix" different products in individually optimal proportions and its ability to profitably bias consumption. We consider a framework in which 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 mix the product of both sellers, cashes in a subscription fee and pays royalties to the sellers based on effective consumption of their product. The set-up well represents economic agents such as Spotify and different artists 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.

We focus our attention on the ability such a platform has to bias consumption through algorithmic recommendation. Streaming platforms found great success in the digital era thanks in part to personalized features such as the well-known "Discover Weekly" playlist, automatically generated for each user every week.[1] Such features are not only very popular. They also have a real impact on consumption patterns: Aguiar et al. (2021) show in their empirical investigation that inclusion in automatically generated playlists such as "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.

While most of the current literature on the topic addresses concerns regarding anti-competitive steering practices, 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. 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 differences 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, 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 labelled "a" and "b") sell their bundle good to a unit mass of consumers uniformly distributed on the $[0, 1]$ line for a price p_i, $i \in \{a, b\}$. Each seller offers a bundled good that consists of only the content they produce (i.e., bundle good a (respectively b) is entirely made of content a (b)). The sellers are assumed to be located at the extreme of the Hotelling line. Along the entire line, a platform for streamed content (labelled s) offers a subscription-

[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/.

based service: upon paying a uniform fee p_s, a consumer can access a mix of content from a and b.[2]

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 buyers who value a balanced mix of content.

We show that when consumers are free to mix consumption based solely on their preference, without intervention by the platform, the existence of such an intermediary represents a Pareto-improvement when products have the same quality. 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 also holds under the extreme assumption that sellers have full bargaining power.

Because the platform cannot price discriminate through a fixed subscription fee and must compete for attention with the sellers, only the marginal consumers are indifferent between joining the platform and purchasing the closest pure bundle. The others consumers who are more centrally located strictly prefer the former option.

In this specification, 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 can mix their consumption autonomously, 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. Moreover, we show that whenever this is the case, the optimal level of bias is set to make consumers indifferent between staying on the platform and only purchasing on the external market. The result follows from the assumption that the platform can discriminate consumers and bias content differently depending

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

on their taste[3].

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 consumption of both artists: 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. In equilibrium, the latter group only learns about the artists and consumes their product if the platform manages to attract both artists. 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. In particular, the platform always selects a positive level of bias when products are vertically differentiated. The bias is constrained by the consumers' participation constraint. In equilibrium, consumers who join the platform are exposed to more of the cheaper, low-quality content that they would optimally select. An interesting conflict emerges: even though an individual seller and the consumers are aligned in their interest for a high-quality product to be made available in a vacuum, the intermediation by the platform can make consumers worse off because of the consequent difference in quality compared to the symmetric case.

The rest of the paper is structured as follows: after a review of the relevant literature, we introduce the model and solve the benchmark specification with homogeneous quality of the products in Section 2. In Section 3 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 (Section 4), Section 5 concludes.

[3]The 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 instruct how much bias he would be willing to tolerate

1.1 Related Literature

Recommendation systems represent a core feature of digital platforms, and streaming platforms like the one 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 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).

It seems clear that the impact 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 by 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 bias consumers based on their location while Bourreau and Gaudin (2022) focuses on uniform biases. Finally, we introduce vertical differentiation between the sellers.

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 market[4].

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 platform would have the incentive to penalize 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

[4]According to Spotify's earning report to investors, the number of the premium subscriber in Q3 of 2022 was 195 millions. Available at: https://s29.q4cdn.com/175625835/files/doc_financials/2022/q3/Q3-2022-Shareholder-Deck-FINAL-LOCKED.pdf

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 they buyers and 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; Döpper and Rasch, 2022).

When ad intensity is a strategic choice in this kind of models, the location game leads to minimum differentiation as the unique equilibrium; otherwise, maximum differentiation is the general equilibrium outcome. We assume the latter and take the extremes of the Hotelling line as the exogenous locations for the sellers. As the platform's role is to smooth the consumption of differentiated products, the choice is a natural one. Rather than focusing on different pricing, however, we make two simplifying assumptions to study the strategic use of bias in isolation: first, we consider only linear pricing. Second, we forego direct modeling of the bargaining process between sellers and the platform and assume royalties to be equal to the price on the external market. The effects we find are implicitly limited, in magnitude, by the assumption of sellers holding full market power, suggesting that the distortions the platform can achieve in reality might be even more severe if the platform can be assumed to hold significant bargaining power compared to the artists she hosts.

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 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 analysed 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 productivity.[5]

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 towards 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.

2 Model setup

We consider two groups of consumers (informed and uninformed), each uniformly distributed on the $[0, 1]$ Hotelling line in a 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 location of the firms operating in the market. There are two firms (sellers or artists), $j = a, b$, who are located at the left and right extremes of the unit line, respectively.

In addition to them, we assume that there is a digital platform (s) that knows the consumers' location and offers them a personalized bundle of contents from the two artists. By doing so, the platform can better match consumers' preferences (Anderson and Neven, 1989; Bourreau and Gaudin, 2022). We define the bundles a and b sold by the two sellers as *pure bundles*, entirely made of contents produced *in-house*. Instead, we define *mixed bundles* the personalized good that the platform offers to each consumer on the line. We define $\lambda(x) \in (0, 1)$ as the share of content a offered to the consumer in x. Conversely, $1 - \lambda(x)$ represents the share of content b offered to the same consumer.

Informed consumers know ex-ante that artist a is located in zero and artist b is located in 1. Moreover, they know their own positions on the line. Instead, uninformed consumers only know about the platform and discover the two artists after they try the platform's streaming service. Consumers purchase exactly one unit of the final good — either the pure or the recommended mixed bundles. We use p_a and p_b to define the price of the pure bundles paid directly to the artists (one can think about pure bundles as CDs, Vinyl, or digital albums). We use p_s to identify the subscription fee paid by consumers to access the platform's service.

Finally, the platform pays royalties (r_j) to the artists per share of their content offered to consumers. We assume that the artists 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.

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

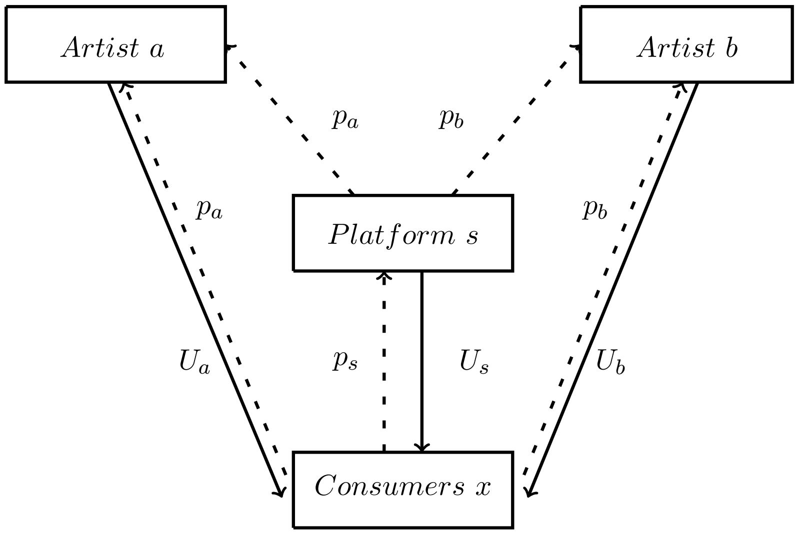

Figure 1: The diagram of the model with payments and services when the platform is active, and both informed and uninformed consumers participate.

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

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

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 $v > 0$ and an artist-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 $v_b < t$ always holds. Figure 1 shows the diagram of the model.

Importantly, we allow the platform to bias the bundles offered to consumers in reaction to artists' 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.

The timing of the game is as follows: at stage 1, artists decide whether to join the platform and serve both informed and uninformed consumers or to stay out and compete for informed consumers only. Upon observing the entry decision and the quality attributes of the two contents, at stage 2, the two artists and the platform set the prices for the pure bundles and the streaming service (p_a, p_b, and p_s). Then, at stage 3, the platform observes the prices (hence, the royalties) and chooses the recommendation system of the streaming service ($\lambda(x)$). Crucially, in this stage, the platform chooses its *bias policy* – i.e., to include more of one content than optimal from the consumers' perspective. Finally, given the prices and the recommendation system, consumers make their consumption decision and profits realize. One should remember that informed consumers know both their location on the Hotelling line as well as the locations of the firms. Instead, the $1-\alpha$ uninformed consumers only know that a platform exists. We assume that all consumers can sample the platform for free before subscribing[6]. During the free sample, uninformed consumers learn the location of the firms and their preferences. Intuitively, if the firms decide not to join the platform at stage 1, 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.

2.1 Homogeneous quality

We begin the analysis by focusing on the benchmark case where the two artists produce content of identical quality – i.e., $V_j = v \; \forall j$. We start from the demand faced by the three sellers (the two artists and the platform) in the last stage of the game. The assumption $v_a = v_b = 0$ simplifies the utility functions to:

$$\begin{aligned} U_{i,a}^{ben} &= v - p_a - t x_i^2 \\ U_{i,b}^{ben} &= v - p_b - t(1-x_i)^2 \\ U_{i,s}^{ben} &= v - p_s - t(x_i - (1-\lambda(x_i)))^2 \end{aligned}$$

where the apex ^{ben} indicates we are in the benchmark scenario. 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:

$$\begin{aligned} x_{as}^{ben} &= \frac{p_s - p_a + t(1-\lambda(x_{as}))^2}{2t(1-\lambda(x_{as}))} && \Longrightarrow \; U_{i,a}^{ben} = U_{i,s}^{ben} \\ x_{sb}^{ben} &= \frac{p_b - p_s + t(2-\lambda(x_{sb}))\lambda(x_{sb})}{2t\,\lambda(x_{sb})} && \Longrightarrow \; U_{i,s}^{ben} = U_{i,b}^{ben} \\ x_{ab}^{ben} &= \frac{p_b - p_a + t}{2t} && \Longrightarrow \; U_{i,a}^{ben} = U_{i,b}^{ben} \end{aligned}$$

We adopt the following notation: x_{jk} indicates the consumer that is indifferent 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.

Efficiently mixed bundle. Intuitively, the best recommendation system that the platform can employ from the consumer's perspective is the one that offers to consumers their preferred mixed bundles so to maximize utility by minimizing transportation costs. We define *efficient bundle* the composite good offered to a consumer so that, for any prices p_a, p_b, and p_s, she gets the highest possible utility. Formally:

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

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

Using this recommendation system, it is possible to update the location of the indifferent consumers as:

$$x_{as}^{ben}|_{\lambda(x_{as})=\lambda^*(x_{as})} = \sqrt{\frac{p_s - p_a}{t}}; \qquad x_{sb}^{ben}|_{\lambda(x_{sb})=\lambda^*(x_{sb})} = 1 - \sqrt{\frac{p_s - p_b}{t}}.$$

We are now ready to derive the demand functions faced by the two artists and the platform. Notice first that in this sub-game the platform is active: sellers have agreed to join, which in turn allowed uninformed consumers to discover them and their taste. Total demand at this stage, then, is the full unit mass. All consumers located between $x_i \in [x_{as}^{ben}$ and $x_{sb}^{ben}]$ obtain a higher utility joining the platform at the subscription fee p_s than buying the pure bundles directly from either artist. This result derives from the negative effect of preferences mismatch on utility ($\partial U_{i,a}^{ben}/\partial x_i < 0$ and $\partial U_{i,b}^{ben}/\partial x_i > 0$). In words, this relation means that no consumer farther away from a than x_{as}^{ben} (respectively, from b than x_{sb}^{ben}) prefers the pure bundle a (respectively, b) to the efficient bundle made available by the platform streaming service. Hence, the demand of the platform is given by $D_s^{ben} = x_{sb}^{ben} - x_{as}^{ben}$.

The demand functions of the two artists, instead, are not made by a single segment, as in the platform's case. The two artists derive revenues from the direct sale of their pure bundles and the royalties paid by the platform, which are paid proportionally to the total consumption of their content by consumers on the platform. Formally, the demands of artists a and b are, respectively:

$$D_a^{ben} = x_{as}^{ben} + \int_{x_{as}^{ben}}^{x_{sb}^{ben}} \lambda^*(x)\, dx \tag{1}$$

$$D_b^{ben} = 1 - x_{sb}^{ben} + \int_{x_{as}^{ben}}^{x_{sb}^{ben}} 1 - \lambda^*(x)\, dx \tag{2}$$

Equations (1) and (2) help us understand more about the potential effect of a bias recommendation system employed by the platform. Consider a consumer i who pays the subscription fee to use the streaming service. Her efficient bundle comprises $(1 - x_i)$ share of content a and x_i share of content b. Assume the platform biases her bundle and increases the proportion of content a from $(1 - x_i)$ to $(1 - x_i + \varepsilon_i)$. Not only would the consumer suffer a utility loss from this bias, the artists would also incur a demand distortion. As prices are given at this stage, the artist cannot react and suffers from a profit loss. On the contrary, artist a would gain from the bias, as now more content would be consumed at the same price/royalty rate.

Suppose for now that the platform selects the efficient mix for all consumers. Using the demands of the artists and the platform (D_s^{ben}, D_a^{ben}, and D_b^{ben}), it is possible to write the objective function of the two artists and the platforms as:

$$
\begin{aligned}
\pi_s^{ben} &= p_s\left(x_{sb}^{ben} - x_{as}^{ben}\right) - p_a\left(\int_{x_{as}^{ben}}^{x_{sb}^{ben}} \lambda^*(x)\,dx\right) - p_b\left(\int_{x_{as}^{ben}}^{x_{sb}^{ben}} 1-\lambda^*(x)\,dx\right)\\
\pi_a^{ben} &= p_a\left(x_{as}^{ben} + \int_{x_{as}^{ben}}^{x_{sb}^{ben}} \lambda^*(x)\,dx\right)\\
\pi_b^{ben} &= p_b\left(1-x_{sb}^{ben} + \int_{x_{as}^{ben}}^{x_{sb}^{ben}} 1-\lambda^*(x)\,dx\right)
\end{aligned}
$$

From the system of first-order conditions, we derive the equilibrium prices:

$$
p_a^{ben} = p_b^{ben} = t; \qquad p_s^{ben} = \frac{10t}{9} \tag{3}
$$

The next proposition follows directly:

Proposition 1. *Consider the case in which the platform offers the efficient mix $\lambda^*(x)$ to the consumers, then the prices are as derived in (3), the profits of the artists and the platform are*

$$
\pi_a^{ben} = \pi_b^{ben} = \frac{t}{2}; \qquad \pi_s^{ben} = \frac{t}{27}
$$

and the indifferent consumers are located in:

$$
x_{as}^{ben} = \frac{1}{3}; \quad x_{sb}^{ben} = \frac{2}{3}; \quad x_{ab}^{ben} = \frac{1}{2};
$$

Proposition 1 illustrates the market outcome of the game when the platform offers the efficient bundle to each consumer that chooses to join. The two artists set the standard Hotelling prices and obtain the standard Hotelling profits. This is partially due to the assumption of full bargaining power. As royalty rates are assumed to match the respective prices on the external market, they do not suffer from platform competition.

Intuitively, the platform plays an important role from the Social Welfare standpoint. Indeed, because a and b are located at the extremes, consumers located "in the middle" in a traditional Hotelling model with firms located at the extremes of the line suffer strong utility losses due to large transportation costs. The platform offers all consumers the possibility to consume precisely the variety they want (the efficient bundle), which increases consumers' utility. However, because the platform cannot price discriminate the consumers, it cannot extract all consumers' informational rent. Consequently, consumers are better off in equilibrium than when the platform is absent. Interestingly, the platform charges a larger subscription price to all consumers who want to use its service and still make positive profits. Intuitively, the platform achieves this by monetizing the transportation costs of consumers.

Finally, because the platform pays the sellers their share of revenues, there is no profit loss by the sellers. This follows from the fact that the total consumption of each seller's product is the same with or without the platform. If the platform is not available, standard Hotelling logic implies that all consumers closer to a (respectively, b) would consume his pure bundle. If the platform offers an efficient mix to consumers, the total quantity sold made by each seller does not change. Therefore, the platform's presence may represent a Pareto improvement over the standard Hotelling competition model.

Biased bundle. The efficient bundle is an equilibrium in all sub-games if there is not a recommendation system that allows the platform to earn higher profits by employing it. In other words, we need to ensure that the platform has no incentive to bias the recommendation system to maximize profits. Because biasing the recommendation system negatively affects consumers' surplus, this question is also important from a policy perspective. In what follows, we define *biased bundles* all the composite goods offered to each consumer that contain a different proportion of the two contents than the efficient bundle.

Because the platform pays royalties to the artists proportional to the total consumption of their content by consumers via the streaming service, it has the incentives to recommend the cheapest content to a broader audience and to hide the most expensive one to reduce its consumption.

Intuitively, such a strategy makes only sense if the prices of the two content are not equal. Otherwise, the platform has no reason to bias its recommendation system. From Proposition 1, we know that in the case of homogeneous quality, the prices set by the two artists are symmetric. Hence:

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

To prove this, suppose the platform decides to alter the proportions of a and b by recommending a bundle that contains a larger share $1 - x_i + \varepsilon$ of content a and a lower one $x_i - \varepsilon$ of content b than the efficient bundle to a specific consumer i. Assume also that $\varepsilon > 0$ does not violate the participation constraint of such consumer, i.e., $U_{i,s}|_{\lambda=\lambda^*+\varepsilon} \geq \max\{U_{i,a}; U_{i,b}\}$. In this case, the platform's cost structure changes. In particular, the platform cost function modifies by $\Delta Cost = +\varepsilon p_a - \varepsilon p_b$, as the platform substitutes a share ε of content a with the same share of content b in the bundle offered to that consumer. From Proposition 1, we know that $p_a = p_b = t$, which implies a cost differential $\Delta Cost = 0$. The platform has no incentives to bias its recommendation system.

3 Innovation and Bias

The music industry is both vertically and horizontally differentiated. Artists backed by major publishers generally have more resources than independent artists to produce better products – e.g., in terms of sound quality or international collaborations. Moreover, artists experiment and research new ways of expressing their art. In other words, they innovate.

It is, therefore, credible to assume that artists compete with products that display different quality levels. In this section, we modify the analysis by assuming that consuming the content by artist b guarantees a larger utility to all consumers irrespective of their preferences for varieties.

Formally, we assume $V_b = v + v_b > v = V_a$ and rewrite the utility functions as:

$$U_{i,a} = v - p_a - tx_i^2$$
$$U_{i,b} = v + v_b - p_b - t(1 - x_i)^2$$
$$U_{i,s} = v + (1 - \lambda(x_i))v_b - p_s - t(x_i - (1 - \lambda(x_i)))^2$$

Recall that $\lambda(x_i)$ indicates the share of content a in each consumer's individual mix. As before, we derive the locations of indifferent consumers by equating the utility functions they obtain by choosing between the three options:

$$x_{as} = \frac{p_s - p_a + (t(1 - \lambda(x_{as}) - v_b)(1 - \lambda(x_{as}))}{2t(1 - \lambda(x_{as}))} \implies U_{i,a} = U_{i,s}$$
$$x_{sb} = \frac{p_b - p_s + (t(2 - \lambda(x_{sb}) - v_b))\lambda(x_{sb})}{2t\,\lambda(x_{sb})} \implies U_{i,s} = U_{i,b}$$
$$x_{ab} = \frac{p_b - p_a + t - v_b}{2t} \implies U_{i,a} = U_{i,b}$$

As expected, absent the price effect the quality gap v_b moves the indifferent consumers towards the location of the artist a, thus shrinking her demand. Quality is a demand shifter.

Efficiently mixed bundle. The analysis of the market outcome when the platform recommends efficient bundles $(\lambda^*(x_i, v_b))$ to all consumers can be readily adapted from the previous section. Formally:

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

The platform offers more content b to all consumers than in the benchmark case. This is because all consumers, irrespective of their location, perceive the quality of content b as higher than the quality of content a.

Using the same backward induction logic, we use the efficient recommendation system to update the location of the indifferent consumers:

$$x_{as} = \sqrt{\frac{p_s - p_a}{t}} - \frac{v_b}{2t}; \qquad x_{sb} = 1 - \sqrt{\frac{p_s - p_b}{t}} - \frac{v_b}{2t}.$$

The demand functions of sellers and platforms follow immediately. As before, the demand of the platform is given by all consumers located in $x_i \in [x_{as}, x_{sb}]$. Formally, $D_s = x_{sb} - x_{as}$ indicates the total demand of the platform.

As for the artists, their demand functions are made of two segments, i.e., the consumers that purchase the pure bundles directly and the total consumption by consumers who use the streaming service. Formally:

$$D_a = x_{as} + \int_{x_{as}}^{x_{sb}} \lambda^*(x_{as})\, dx \tag{4}$$
$$D_b = 1 - x_{sb} + \int_{x_{as}}^{x_{sb}} 1 - \lambda^*(x_{sb})\, dx \tag{5}$$

From a simple comparison of equations (1) - (2) and (4) - (5), one can easily see that, ceteris paribus, the quality gap v_b implies a reduction of the total demand of artist a and a general expansion of the demand of artist b. This is due to two main effects. First, because the quality of content b is higher than that of content a, more people want to substitute the former with the latter. Hence, the direct sales of the artist b increase, and the direct sales of the artist a fall as more consumers join the platform. Second, within the streaming service of the platform, more content of artist b is offered to consumers, increasing her demand and lowering the one of her rival.

Using the demands of the artists and the platform (D_s^{ben}, D_a^{ben}, and D_b^{ben}), it is possible to write the objective function of the two artists and the platforms as:

$$\begin{aligned}
\pi_s &= p_s\left(x_{sb} - x_{as}\right) - p_a\left(\int_{x_{as}}^{x_{sb}} \lambda^*(x, v_b)\, dx\right) - p_b\left(\int_{x_{as}}^{x_{sb}} 1 - \lambda^*(x, v_b)\, dx\right) \\
\pi_a &= p_a\left(x_{as} + \int_{x_{as}}^{x_{sb}} \lambda^*(x, v_b)\, dx\right) \\
\pi_b &= p_b\left(1 - x_{sb} + \int_{x_{as}}^{x_{sb}} 1 - \lambda^*(x, v_b)\, dx\right)
\end{aligned}$$

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

$$p_a = t - \frac{v_b}{3}; \quad p_b = t + \frac{v_b}{3}; \quad p_s = \frac{10t}{9} + \frac{v_b^2}{4t} \tag{6}$$

The next proposition follows directly:

Proposition 2. *Consider the case in which the platform offers the efficient mix $\lambda^*(x, v_b)$ to the consumers, then the prices are as derived in (6), the profits of the artists and the platform are*

$$\pi_a = \frac{(3t - v_b)^2}{18t} \quad \pi_b = \frac{(3t + v_b)^2}{18t}; \quad \pi_s = \frac{t}{27} - \frac{v_b^2}{36t}$$

and the indifferent consumers are located in:

$$x_{as} = \frac{1}{3}; \quad x_{sb} = \frac{2}{3}; \quad x_{ab} = \frac{1}{2} - \frac{v_b}{6t};$$

Proposition 2 highlights two important results. First, the number of consumers purchasing the two pure bundles does not change. Due to the difference in the quality of the two available products, seller b has the incentive to raise p_b to extract a higher revenue from consumers who prefer to consume his pure bundle to mix on the platform. Seller a, instead, lowers the price of its product to induce more consumers to leave the platform and purchase the pure bundle. The difference in prices compensates for the difference in quality, and consumption on the direct channel is unchanged. The share of demand of each seller on the platform, instead, is affected by the change in quality: consumers who mix the products optimally consume more of b's product than before.

Second, the platform raises its price to monetize the higher quality of content b. This is not enough to compensate for the larger costs that the platform faces: because the efficient recommendation system increases the proportion of the most expensive content in the bundles of all consumers, the platform's operational costs rise more than the subscription fee. As a consequence, the platform is hurt by the difference in quality because of the composition of the efficient bundle, now richer in b's product, and the uniform participation fee.

Biased bundle. Intuitively, the efficient bundle when quality is not symmetric cannot be an equilibrium. It is easy to prove that, given the prices p_a, p_b, and p_s, at stage 2, the platform has strong incentives to offer each consumer a biased bundle that contains less of the most expensive content(without violating their participation constraint). To prove it, suppose the platform decides to alter the proportions of a and b by recommending a bundle that contains a larger share $1 - x_i - \frac{v_b}{2t} + \varepsilon$ of content a and a lower one $x_i + \frac{v_b}{2t} - \varepsilon$ of content b than the efficient bundle to a specific consumer i. Assume also that $\varepsilon > 0$ does not violate the participation constraint of such consumer, i.e., $U_{i,s}|_{\lambda=\lambda^*(x_i,v_b)+\varepsilon} \geq \max\{U_{i,a}, U_{i,b}\}$. In this case, the platform's cost structure changes. In particular, the platform cost function modifies by $\Delta Cost = +\varepsilon p_a - \varepsilon p_b$, as the platform substitutes a share ε of content a with the same share of content b in the bundle offered to that consumer. From Proposition 2, we know that $p_a < p_b$, which implies a cost differential $\Delta Cost = \varepsilon(p_a - p_b) < 0$. Hence,

Corollary 2. *The efficient bundle with asymmetric quality is not an equilibrium. The platform can lower costs by manipulating the recommendation system without altering its subscription revenues.*

When prices are different, the optimal strategy from the platform's perspective is to favor the cheapest content. It does so by biasing the recommendation system to include the largest possible share of it in the mixed bundles offered to consumers. Notice that consumers are not willing to accept any recommendation system passively. Any mixed bundle offered to a consumer must be such that she weakly prefers joining the platform to buying the pure bundle from the closest artist. Define the mix $\bar{\lambda}(x_i, v_b)$ that favors content a and makes a consumer i indifferent between the streaming service and one of the pure bundles as:

$$\bar{\lambda}(x_i, v_b) = \{\lambda \in [\lambda^*(x_i, v_b), 1) \text{ s.t. } U_{i,s}|_{\lambda=\bar{\lambda}(x_i,v_b)} = \max\{U_{i,a}, U_{i,b}\} \; \forall x_i \in (x_{as}, x_{sb})\}$$

Intuitively, consumers located at the center of the market segment of the platform are willing to sustain a larger bias. In contrast, indifferent consumers located at x_{as} and x_{sb} would leave the platform if anything other than the efficient bundle is offered to them. Using the utility functions, we can rewrite the maximum individual bias as:

$$\bar{\lambda}(x_i, v_b) = 1 - x_i - \frac{v_b}{2t} + \min\left\{\frac{\sqrt{4t(p_a - p_s) + (v_b + 2tx_i)^2}}{2t}, \frac{\sqrt{4t(p_b - p_s) + (v_b + 2t(1 - x_i))^2}}{2t}\right\}$$

It is easy to prove that $\bar{\lambda}(x_i, v_b) = \lambda^*(x_i, v_b)$ when $x_i = x_{as}$ or $x_i = x_{sb}$. Moreover, the two terms in the brackets are equal when $x = x_{ab}$.

We can now update the demand functions of the three sellers. First, notice that the locations of the indifferent consumers do not change as they consume the efficient bundles (the platform cannot bias their bundles without losing them). Hence, $D_s = x_{sb} - x_{as}$ is unaltered.

Instead, the demand of the artists changes because of the different proportions of content in the new biased bundles. Formally:

$$D_a^B = x_{as} + \int_{x_{as}}^{x_{sb}} \bar{\lambda}(x_i, v_b)\, dx \tag{7}$$

$$D_b^B = 1 - x_{sb} + \int_{x_{as}}^{x_{sb}} 1 - \bar{\lambda}(x_i, v_b)\, dx \tag{8}$$

The apex ^{B} indicates the scenario where the platform offers a biased mix to consumers. The demands can be used to derive the profit of the two artists and the platform:

$$\pi_s^B = p_s\,(x_{sb} - x_{as}) - p_a \left(\int_{x_{as}}^{x_{sb}} \bar{\lambda}(x_i, v_b)\, dx\right) - p_b \left(\int_{x_{as}}^{x_{sb}} 1 - \bar{\lambda}(x_i, v_b)\, dx\right)$$

$$\pi_a^B = p_a \left(x_{as} + \int_{x_{as}}^{x_{sb}} \bar{\lambda}(x_i, v_b)\, dx\right)$$

$$\pi_b^B = p_b \left(1 - x_{sb} + \int_{x_{as}}^{x_{sb}} 1 - \bar{\lambda}(x_i, v_b)\, dx\right)$$

The maximization of the profits with respect to the prices proves to be too complex to be solved analytically. Therefore, we must adopt an alternative method to approximate the game's equilibrium. Still, from the above profit functions, it is possible to observe two important features of the solution. First, artist b faces a lower demand due to the biased recommendation system; second, artist a faces an increased demand because of the favorable bias. These two variations in artists' demands exert a pressure of opposite signs on the equilibrium prices.

On the one hand, artist b lower her price in response to the lower demand. On the other hand, artist a raises her price in response to the higher demand. Possibly, the two prices converge towards a common value. It can be proved that if this convergence occurs, the biased mix is not an equilibrium. In fact, upon observing equal prices, the platform is better off offering efficient bundles to all consumers. Therefore, for an equilibrium to occur, the price of the high-quality content must remain above the price of the low-quality one regardless of the bias adopted by the platform. This situation can only occur if v_b is sufficiently large. In order to solve the game, we need to adopt an alternative method to model the bias. In particular, instead of assuming that the platform sets a personal bias to each subscriber, we consider a uniform level of bias $\varepsilon > 0$ that applies to all subscribers. Of course, this violates the participation constraint of some of the subscribers, and it is, therefore, not a viable option for the platform. However, we bypass this problem assuming the bias is redistributed to consumers according to their participation constraints.

Lemma 1. *Given $v_b > 0$, for any p_a, p_b, and p_s, there exists a uniform bias $\varepsilon^*(v_b)$ equal to the sum of the maximal individual biases that can be sustained by consumers on the platform:*

$$\int_{x_{as}}^{x_{sb}} \varepsilon(x, v_b) dx = \varepsilon^*(v_b)(x_{sb} - x_{as})$$

where $\varepsilon(x, v_b) = \min\left\{\frac{\sqrt{4t(p_a - p_s) + (v_b + 2tx)^2}}{2t}, \frac{\sqrt{4t(p_b - p_s) + (v_b + 2t(1-x))^2}}{2t}\right\} - \lambda^*(x_i, v_b)$

Intuitively, the individual level of bias only matters to the consumer it applies to. From

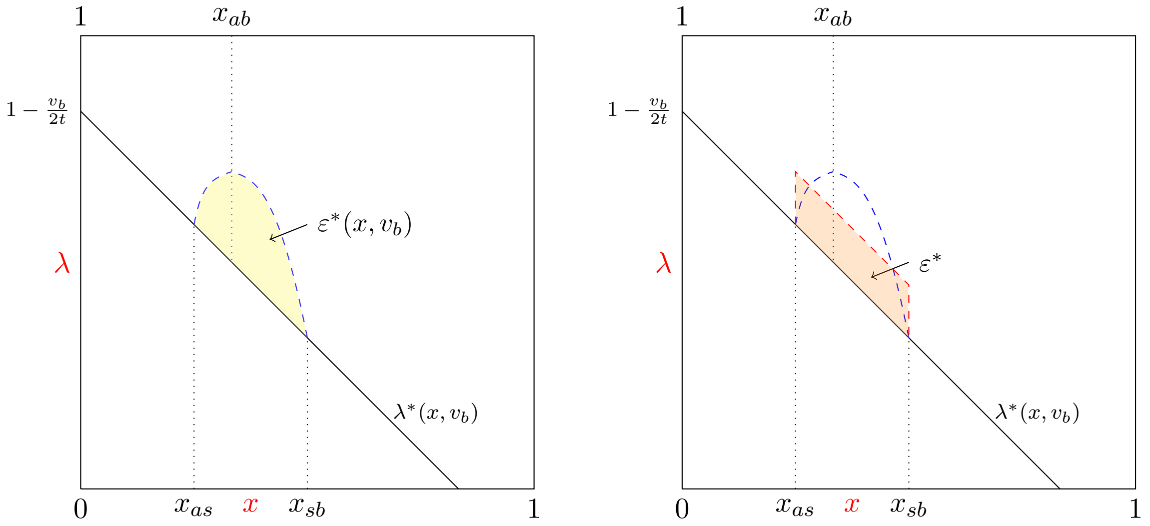

Figure 2: On the left: Personalized biased bundle. The area in yellow is the total demand that the platform can shift from artist b to artist a. On the right: Uniform biased bundle. Each consumer receives a bias ε. The two colored areas must be exactly equal for the uniform bias to be payoff equivalent to the personalized one.

the perspective of sellers and the platform, the total bias is what matters. Solving the problem assuming a uniform level of bias to be redistributed, then, does not break consistency as long as the total bias is incompatible with the constraints of all buyers combined. Lemma 1 shows that for all price levels p_a, p_b, there exists a unique mass $\varepsilon^*(v_b)$ that minimizes platform costs and does not break the consumer participation constraint. This level $\varepsilon^*(v_b)$, selected observing prices, is equivalent to the equilibrium bias imposed by the platform. As such, it dictates the final demand of the sellers and must be incorporated into their maximization problem to obtain equilibrium prices. Figure 2 provides a graphical representation of the assumption.

The new recommendation system can be written as $\lambda^*(x_i, v_b) + \varepsilon^*(v_b)$. For brevity, we denote the uniform bias $\varepsilon^*(v_b)$ as ε^*. Hence, we adjust the profit functions as:

$$
\begin{aligned}
\pi_s^B &= p_s\,(x_{sb} - x_{as}) - p_a \left(\int_{x_{as}}^{x_{sb}} \lambda^*(x_i, v_b) + \varepsilon^*\, dx \right) - p_b \left(\int_{x_{as}}^{x_{sb}} 1 - \lambda^*(x_i, v_b) - \varepsilon^*\, dx \right) \\
\pi_a^B &= p_a \left(x_{as} + \int_{x_{as}}^{x_{sb}} \lambda^*(x_i, v_b) + \varepsilon^*\, dx \right) \\
\pi_b^B &= p_b \left(1 - x_{sb} + \int_{x_{as}}^{x_{sb}} 1 - \lambda^*(x_i, v_b) - \varepsilon^*\, dx \right)
\end{aligned}
$$

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

$$
p_a = t - \frac{v_b}{3} + \frac{2\,\varepsilon^*\,t}{3}; \quad p_b = t + \frac{v_b}{3} - \frac{2\,\varepsilon^*\,t}{3}; \quad p_s = \frac{10t}{9} + \frac{v_b^2}{4} - \varepsilon^*(v_b - \varepsilon^*\,t) \tag{9}
$$

The next proposition follows directly:

Proposition 3. *Consider the case in which the platform offers a biased mix $\lambda^+(x, v_b) + \varepsilon^*$ to the consumers. Then the prices are as derived in (9), the profits of the artists and the platform* *are*

$$\pi_a = \frac{(t(3+2\varepsilon^*)-v_b)^2}{18t}; \qquad \pi_b = \frac{(t(3-2\varepsilon^*)+v_b)^2}{18t}; \qquad \pi_s = \frac{t}{27} - \frac{v_b^2}{36t} + \frac{\varepsilon^*}{9}(7v_b - 13t\varepsilon^*)$$

and the indifferent consumers are located in:

$$x_{as} = \frac{1}{3} - \varepsilon^*; \qquad x_{sb} = \frac{2}{3} - \varepsilon^*; \qquad x_{ab} = \frac{1}{2} - \frac{v_b}{6t} - \frac{2\varepsilon^*}{3}$$

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, this increase in the demand for the seller a's content generates positive pressure on the price of the pure bundle a in the direct channel. 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.

Instead, seller b suffers from the bias in the recommendation. Consumers are exposed to a lower-than-optimal level of content b on the platform. To compensate for this loss, seller b lowers her price, 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. Finally, 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 $v_b > 2\varepsilon t$. Recall that the bias is chosen upon observing the prices. Also, recall that the bias is either the maximum one ($\lambda = \min\{\lambda^a_{as}; \lambda^a_{sb}\}$) when prices are different, or zero ($\lambda = \lambda^+(x, v_b)$) if prices are not different. Hence, if the condition above is not satisfied, the platform has no incentives to steer consumption toward the content a, as it would not be cheaper.

Finally, notice that the equilibrium value of the bias depends on the consumer participation constraint as per Lemma 1. In particular, sellers set prices anticipating the bias, which affects consumer's participation decision and, therefore, the value of the constraint. The overall effect is governed by v_b: it can be shown that the higher v_b is, the tighter the constraint gets: consumers want to be exposed to b's content more the better it is comparatively, which makes them more sensitive to the bias itself. Therefore, as v_b grows and the consumer participation constraint tightens, the penalty imposed on seller b by the platform becomes milder and milder.

Finally, in the Appendix, we show that

Corollary 3. *Consider the case in which the platform offers a biased mix $\lambda^+(x, v_b) + \varepsilon^*$ to consumers. The high-quality seller b does not have a profitable deviation in the pricing strategy. Moreover, imitating the rival's price to zero the bias out is not profitable.*

4 Sellers' participation decision

Let us now proceed backward and consider the sellers' participation decision. At the beginning of the game, consumers are divided in two distinct groups, both distributed uniformly on the $[0, 1]$ line. A first group of mass $\alpha \in [0, 1]$ is aware of the existence of the sellers and their relative location – they are the informed consumers. The second group, with mass $1 - \alpha$, is uninformed. Uninformed consumers learn of the existence of the sellers or their relative position only if they join the platform. This can only happen if the platform manages to attract both sellers.[7]

In the sub-game where either or both sellers decide not to join the platform, only a proportion α of consumers are active. With no platform active, consumers cannot mix their consumption and are therefore limited to purchasing a pure bundle from either a or b. In this sub-game, 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 then:

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

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

Where the apex ^{out} indicates the scenario in which 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 benchmark case in which $v_b = 0$, it is clear that the platform is always active: since there is not bias in equilibrium, sellers 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. Suppose again that there are no uninformed consumers, that is, that $\alpha = 1$. Artist b anticipates the bias policy of the platform and the losses he would suffer as a consequence of it. He also knows that joining the platform does not expose his product to more consumers. Clearly, then, b would rationally choose not to join the platform. At the opposite limit, suppose that $\alpha = 0$: if artist b does not join the platform, he cannot make any sale. Regardless of how biased the recommendation system is in favor of his rival, he would always optimally choose to join the platform. For intermediate levels of α, direct comparison of π_b^{out} and π_b leads to the following equilibrium result:

[7]Consider 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 artist, as there is 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 $p_s \geq p_j$, which means all consumers weakly prefer purchasing the pure bundle j directly by the seller.

Proposition 4. *Assume the two contents are vertically differentiated. Then if the number of informed consumers is sufficiently low*

$$\alpha \leq \alpha^* \equiv \frac{((3-2\varepsilon^*)t+v_b)^2}{(3t+v_b)^2}$$

Then, in equilibrium, both sellers join the platform, and the recommendation system is biased in favor of the low quality content.

The platform suffers from a lack of commitment: since the bias policy is chosen after pricing decisions have already been made, sellers anticipate it and decide whether to join the streaming service accordingly. The higher α is, the lower the sellers' benefit from exposure on the platform, and the less b is willing to accept the bias. α^* represents the threshold value that makes seller b indifferent between joining the platform or not.

The comparative statics of the threshold α^* requires some discussion. To derive the equilibrium outcomes of the game, we treated the level of bias as a parameter. However, we know from the previous discussion (see Lemma 1) that whenever prices differ, the platform is willing to favor the cheapest seller and implement a biased recommendation system that makes all its users indifferent between joining the platform and leaving. In other words, the platform adopts a maximum-bias policy. Thus, the actual level of bias moves together with the participation constraint of the consumers.

Intuitively, as v_b increases and the quality difference becomes more pronounced, the high-quality seller b charges a larger price to monetize its quality. Thus, from logic, as raising v_b exerts a positive pressure on the price gap $p_b - p_a$, the platform is willing to increase the level of bias as its best cost-minimizing strategy. However, when the quality gap increases, consumers are less willing to substitute content a with content b, as they derive less utility from it. Consequently, the "potential" bias, i.e., the set of available bias policies, shrinks. This leads to the counterintuitive result that, despite the fact that platform would like to bias more its recommendation system when v_b increases, it is forced to reduce it, as consumers would otherwise be discouraged from joining the platform.

Formally, the discussion above translates in ε^* being decreasing in v_b, and α^* being unequivocally increasing in v_b. In the limit, $\lim_{v_b \to \infty} \alpha^* = 1$. Indeed, in such a theoretical case, the seller would derive infinite payoffs in both scenarios and be, therefore, indifferent between the two. In general, the seller b's penalty generated by the platform gets smaller the better his product is (because more consumers want to purchase its content). As v_b grows, the return from being exposed to more consumers becomes larger since the seller can increase prices due to v_b. Since the price reduction caused by the seller anticipating the bias in equilibrium becomes less impactful, the former effect dominates. The seller is willing to withstand the platform bias for more values of α as v_b grows as a result.

Overall, the platform's incentive to bias consumption away from the better (and more expensive) product makes any equilibrium in which she is active only feasible if she generates enough additional consumption to induce sellers to join despite the bias. Abstracting from the model, one can think of α as an indirect measure of popularity: the more an artist is known, the less he needs a streaming platform to reach consumers. Less-known artists, on the other hand, greatly benefit from the exposure a streaming platform provides. They are likely to join even if they suspect that they would be penalized by the recommendation system. Because of the threat of biasing, moreover, any incentive they might have to provide above-average products is dampened: if they do, and raise price (and royalties) accordingly, they know that they would be penalized by the recommendation system. These effects, combined, raise concerns over the impact streaming services can have on the quality of the medium they monetize.

5 Conclusion

In this paper, we study the incentives of a streaming platform to bias bundling in an effort to minimize operational costs. 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 shrink. 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.

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 was 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.

Recall that the high-quality seller never wants to condition his optimal price to match that of his rival. The result follows from the assumption that sellers set prices and royalties equally. Separating the two, letting sellers set prices outside the platform and the agency set royalties inside might lead to new insights. The observation extends the discussion to the relative bargaining power of such agencies, and what their royalty-setting incentives truly are. A thorough discussion of this additional dimension is beyond the scope of this paper. Nonetheless, the observation suggests that representation in bargaining could improve consumer welfare by reducing or eliminating the incentives the platform has to bias consumption towards low-quality content.

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.

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.

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.

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.

Döpper, 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.

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 Bargining 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.

Mukherjee, A. and Pennings, E. (2011). Unionization structure, licensing and innovation. *International Journal of Industrial Organization*, 29(2):232–241.

Oestreicher-Singer, G. and Sundararajan, A. (2012). Recommendation networks and the long tail of electronic commerce. *Mis quarterly*, pages 65–83.

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.

Popper, B. (2016). Spotify's discover weekly reaches 40 million users and 5 billion tracks streamed. *The Verge, May*, 25.

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.

A Proofs – under construction

Proof of Proposition 1 and Corollary 1.

Proof. ■

Proof of Proposition 2 and Corollary 2.

Proof. ■

Proof of Lemma 1

Proof. ■

Proof of Proposition 3

Proof. ■

Proof of Corollary 3

Proof. Artist b is worse off by the possibility of recommendation bias. The intensity of the distortion is monotonically increasing in the size of the demand loss due to the bias ε. Anticipating the choice of the platform in the bias setting stage, artist b may decide to imitate the price of artist a in stage 1 to induce the platform to offer the efficient bundle to consumers.

In this case, the profit functions are adjusted as follows:

$$\begin{aligned}
\pi_a &= p_a \left(x_{as} + \int_{x_{as}}^{x_{sb}} \lambda^+(x, v_b)\, dx \right); \\
\pi_b &= p_a \left(1 - x_{sb} + \int_{x_{as}}^{x_{sb}} (1 - \lambda^+(x, v_b))\, dx \right); \\
\pi_s &= p_s(x_{sb} - x_{as}) - p_a \left(\int_{x_{as}}^{x_{sb}} \lambda^+(x, v_b)\, dx + \int_{x_{as}}^{x_{sb}} (1 - \lambda^+(x, v_b))\, dx \right);
\end{aligned}$$

The standard profit maximization procedure yields the following prices:

$$p_a = p_b = t - v_b; \quad p_s = \frac{10t}{9} - v_b; \tag{10}$$

In turn, prices in (10) imply:

$$\pi_a = \frac{(t - v_b)^2}{2t}; \quad \pi_b = \frac{(t - v_b)(t + v_b)}{2t}; \quad \pi_s = \frac{t}{27}; \tag{11}$$

Standard comparison of the profits of artist b in the two cases shows that imitating the rival's price p_a is a profitable strategy if and only if $v_b < \frac{t}{10}\left(3\sqrt{1 + 12\epsilon - 4\epsilon^2} + 2\epsilon - 3\right) < 2\varepsilon\, t$.

It follows that artist b never imitates artist a's price, as it is not profitable for the admitted value of parameters.

The analysis above suggests that by biasing the recommendation system, the platform may lower the incentives for an artist to invest in innovation with respect to the benchmark case with the efficient bundle. ■

Proof of Proposition 4

Proof. ■