

Gambato, Jacopo; Sandrini, Luca

Working Paper

Dance to my tune! Discovery mode and built-in recommendation bias

ZEW Discussion Papers, No. 25-037

Provided in Cooperation with:

ZEW - Leibniz Centre for European Economic Research

Suggested Citation: Gambato, Jacopo; Sandrini, Luca (2025) : Dance to my tune! Discovery mode and built-in recommendation bias, ZEW Discussion Papers, No. 25-037, ZEW - Leibniz-Zentrum für Europäische Wirtschaftsforschung, Mannheim

This Version is available at:

<https://hdl.handle.net/10419/323931>

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.25-037 | 07/2025

DISCUSSION PAPER

// JACOPO GAMBATO AND LUCA SANDRINI

Dance to My Tune! Discovery Mode and Built-in Recommendation Bias

“Dance to my Tune!”

Discovery Mode and Built-In Recommendation Bias*

Jacopo Gambato[†]

Luca Sandrini[‡]

July 15, 2025

Abstract

We examine the strategic considerations and effects of product placement services in digital content aggregators, focusing in particular on Spotify’s “Discovery Mode”. Discovery Mode introduces bias in users’ consumption bundle that content providers pay through discounted royalties, and triggers users to actively adjust to it in response. The platform’s ability to manipulate consumption bundles (and adjust users’ participation fee consistently) leads to promotion of the cheapest content available and degradation of users’ effective consumption bundles. In equilibrium, Discovery Mode can either benefit or harm the provider of the cheapest content to stream, with the harm arising whenever Discovery Mode threatens to revert preexisting bias against said provider. Importantly, Discovery Mode always forces users to costly adjust their consumption more, unequivocally generating loss of efficiency in the market. We further highlight an indirect increase in the risk of market concentration stemming from the platform’s ability to bias consumption.

Keywords: platform economics, Discovery Mode, content aggregator, recommendation bias, streaming platforms

JEL Codes: D4, L1, L5

*We thank Luis Aguiar, Bernhard Ganglmair, Germain Gaudin, Marteen Janssen, Leonardo Madio, Andrea Mantovani, Simon Martin, Evgenia Motchenkova, Martin Peitz, Lorian Sabatino, Moritz Schwarz, Julian Wright, and all participants to the MaCCI/Epos 2025 Workshop on Digital Markets and CRESSE Annual Conference 2025 for useful discussions and comments.

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

[‡]ZEW Mannheim - luca.sandrini@zew.de

1 Introduction

Recommendation systems are ubiquitous in digital markets and are a core feature of many streaming platforms. Recommendation systems have been documented to expand the range of products consumed in a variety of settings, a phenomenon referred to as the “long-tail effect” (Fleder and Hosanagar, 2009; Brynjolfsson et al., 2011; Oestreicher-Singer and Sundararajan, 2012a; Datta et al., 2018). They therefore represent a powerful instrument to boost consumption of specific products (Aguiar and Waldfogel, 2021; Aguiar et al., 2021).

Recommendation systems do not just improve matching, though. Due to the sheer number of products available on digital marketplaces and streaming platforms, these tools contribute to the success or failure of products by boosting their visibility or systematically hiding them from users. Content moderation is a prime example of the ability of platforms to systematically affect visibility of hosted content: posts and other kinds of content that are (or are suspected to be) harmful for the community can have their diffusion drastically reduced (Madio and Quinn, 2024). More recently, however, the literature has explored more harmful ways in which biased recommendations can be used by intermediaries. Self-preferencing or differential treatment based on economic conditions and contractual terms, for example, seem to be largely diffused as well (Reimers and Waldfogel, 2023; Waldfogel, 2024; Chen and Tsai, 2024).

This paper analyzes the incentives of a digital platform to bias its recommendation system for profit-maximizing purposes through contractually stipulated bias, focusing specifically on the particular case of streaming platforms in the music industry. The music industry has traditionally been characterized by powerful intermediaries largely in control of the diffusion of new and old content (the radio, then television through popular programs such as MTV, and, eventually, digital platforms).¹ Contractually stipulated bias is not a recent phenomenon: artists and labels paying to place specific content in high-rotation on intermediation channels, a practice referred to as “Payola”, were already common in the radio era.²

Streaming platforms, however, are special in their ability to shape consumption: in their investigation of Spotify’s recommendation system, Aguiar and Waldfogel (2021) show that early inclusion in popular playlists boosts future performance of songs, indicating an uncanny ability of modern intermediation to affect consumption choices. The finding raises concerns because all of the most popular streaming services are offered by digital behemoths such as Apple, Google, Amazon, and Spotify, all of which are (openly or not) suspected to be engaging in active distortion of their recommendation systems (Bourreau et al., 2021; Bourreau and Gaudin, 2022; Freimane, 2022; Aguiar et al., 2024). In a recent investigation published on The Harper’s Magazine, for example, journalist Liz Pelly suggests that Spotify might be outsourcing the production of cheap music content to external companies with the purpose of promoting these artificial “ghost artists” in Spotify-curated playlists.³ Such practices make obvious economic

¹Ironically (although purposely) the clip of the song “Video killed the radio star” by The Buggles was the first video ever broadcasted by MTV during its official launch on August 1, 2001.

²See, for instance, The Independent, “Payola: One of music’s oldest arrangements back with a bang on streaming playlists” by Adam Sherwin, August 20, 2015. Available at: <https://www.independent.co.uk/arts-entertainment/music/news/payola-one-of-music-s-oldest-arrangements-back-with-a-bang-on-streaming-playlists-10464513.html>. Last access on July 15, 2025.

³See The Harper’s Magazine, “The Ghosts in the Machine” by Liz Pelly, January 2025. Available at: https://harpers.org/archive/2025/01/the-ghosts-in-the-machine-liz-pelly-spotify-musicians/?utm_source

sense from the platform’s perspective: many Spotify-curated playlists are themed for “background” consumption for whenever users are not actively listening to music.⁴ By including financially convenient content instead of more relevant but expensive alternatives, Spotify can lower the amount it pays in royalty during these sessions and, therefore, in general.

In this paper we study the strategic considerations behind, and implications of, a specific contractually stipulated distortion of recommendations: Spotify’s “Discovery Mode”. Designed as a product placement service, the Discovery Mode (DM henceforth) allows artists (or music labels owning artists’ content’s intellectual property rights) to select songs to be promoted to consumers. Whenever promoted songs are consumed as a consequence of the promotion, the platform obtains a discount on the royalty rate owed to the IP holder.⁵ We build a theoretical model with this mechanism in mind: a monopoly platform informed about consumers’ taste offers two competing labels to bias consumption in their favor compared to a baseline without DM in exchange for a discount in her operational costs (that is, the royalties owed to the labels). Consumers can both costly engage with content by actively selecting their consumption bundle and costlessly consume the bundle selected for them by the platform.

We show that in this environment bias always favors the platform: bias acts as a demand shifter, boosting consumption of whatever content the platform has an incentive to push on consumers, namely the cheaper option available. Consumers, however, can actively adjust their bundle if recommendations stray too much from their optimal one. Because such adjustments are costly, the bias constrains the fee that the platform can extract from consumers. The DM, then, acts as a way to relax this constraint: if some content becomes cheaper to stream through labels purchasing the DM, the platform can afford to increase the bias and reduce the fee compared to the baseline outcome without DM. While this generates gains from the platform, it ultimately hurts efficiency as the consumers are induced to curate their consumption more than they would otherwise.

The introduction of DM can have different effects on the profits of the competing labels. As already mentioned, the platform has the incentive to inflate consumption of the content that is cheapest to provide. Her ability to do so depends on consumers’ reaction to the bias and consequent tightening of their participation constraint. When the DM allows the platform to introduce bias that would not have been profitable to introduce without the discount, both the platform and the recipient of said bias benefit from it. Because the platform has an incentive to bias consumption even in the absence of the DM, however, when bias can be introduced without it the same seller might not have the incentive to purchase the DM in a vacuum. In this case, however, his competitor always does. In equilibrium, the threat of a reversal in the direction of the bias through aggressive discounting makes the provider already benefiting from the bias accept the DM despite it leading to a loss of margin on the biased portion of his demand. While the former case can be thought of as a gracious promoting of smaller labels, penalized by their

=pocket_saves. Last access on July 15, 2025.

⁴Playlists designed to be a background for workouts, strolls, chores, homework, and more, and playlists themed around different moods or artists are extremely popular.

⁵According to the dedicated website, “*with Discovery Mode, artists, and labels identify songs that are a priority, and our system adds that signal to the algorithms that power personalized playlists.*” In other words, music labels and artists pay for visibility in the form of “*a percentage of [the] revenue generated on those selected streams [...]*.” See <https://artists.spotify.com/discovery-mode>. Last access July 15, 2025.

smaller portfolio of artists, the latter leads to gains at the exclusive benefit of the platform at the expense of these very same smaller labels. While this outcome hinges on systematic bias already be commonplace on digital platforms without contracts like the DM, we argue that recent empirical findings (most prominently, [Reimers and Waldfogel, 2023](#) and [Aguilar et al., 2024](#)) strongly points at bias pre-dating the introduction of the DM, raising concerns about the implications of such promotional contracts.⁶

Finally, we broaden the scope of the discussion by investigating the incentives of major labels to expand their catalog by poaching independent artists from smaller competitors when the platform is able to bias consumption and when the platform is not. We show that, besides the clear incentive of expanding to increase consumption of content owned by them, major labels have an additional incentive to expand as this makes consumers more sensitive to the bias all else equal. Because the platform sets the bias to balance operational costs paid to the labels and utility extracted from consumers through the fee, poaching independent artists tightens the participation constraint of consumers for all levels of bias introduced, ultimately reducing the equilibrium bias.

The rest of the paper is structured as followed: after a review of the relevant literature, we introduce our main framework in Section 2. Section 3 contains our main analysis, and Section 3.2 in particular presents our results through a simplified illustrative model. We extend the analysis by discussing the role of rationality for our results and the interaction between bias and incentives to expand in Section 4. Section 5 concludes.

Related literature. The paper contributes to several intertwined strands of literature. Media markets have always been an important topic in economic research. Until recently, the literature has been focused primarily on the interaction between content and advertisement ([Anderson and Gabszewicz, 2006](#); [Peitz and Valletti, 2008](#); [Thomes, 2013](#)). In recent years, however, there has been a shift in focus when studying streaming platforms, with more emphasis being given to choices related to mixing content from different sources. The literature has largely settled at [Anderson and Neven \(1989\)](#) as the workforce model: the “mixing Hotelling” framework allows to seamlessly capture consumers’ preferences for consuming different mixes of content.⁷

A limitation of this approach, however, comes in the stark difference that mechanically arises whenever any kind of asymmetry between the competing sellers is incorporated. In contrast, we develop a model of heterogeneous consumption modes. We build on the model in [De Corniere and Sarvary \(2023\)](#) and allow consumers to either actively listen to music (a costly mode of consumption) or to reproduce it in background. This specification allows us to have a more nuanced perspective on the effect of asymmetries between (in our case) content providers: unlike in papers that build on [Anderson and Neven \(1989\)](#), we show that, unless consumers are largely unresponsive to bias, differentiation must be either substantial or contractually induced by the platform for biased recommendations to emerge in equilibrium.

⁶Our findings align with the concerns previously expressed by professionals within the music industry. See The Guardians, “Pay to get playlist? The accusations against Spotify’s Discovery Mode” by Liz Pelly, February 2025. Available at: <https://www.theguardian.com/music/2025/feb/19/spotify-discovery-mode-payola-playlist>, last access on July 15, 2025.

⁷A non-exhaustive list includes [Hoernig and Valletti \(2007, 2011\)](#); [Döpper and Rasch \(2022\)](#); [Bourreau and Gaudin \(2022\)](#) and [Gambato and Sandrini \(2023\)](#).

Recommendation systems, their design, and their role in shaping consumption have been extensively studied. Much of the early empirical literature focused on quantifying the effect of recommendation systems on sales (Tucker and Zhang, 2010; Oestreicher-Singer and Sundararajan, 2012a,b). In more recent years, instead, scholars have attempted to quantify bias in recommendations in various settings. Besides the aforementioned Aguiar and Waldfogel (2021) and Aguiar et al. (2024), notable contributions outside of the streaming platform environment can be found in Chen and Tsai (2024) and Waldfogel (2024), which highlight bias in Amazon’s recommendations, and Freimane (2022), which instead investigates bias in news aggregators. The theoretical literature on bias in recommendations largely models the incentives to introduce bias by an intermediary through a consequent increase in revenue (Hagi and Jullien, 2011, 2014; De Corniere and Taylor, 2019; Peitz and Sobolev, 2022). This approach is, however, not well suited to model streaming platforms, where monetization is generally applied to the consumer side. We follow Bourreau and Gaudin (2022) instead, and consider a reduction in operational costs, namely the royalties owed to content providers, as the channel creating the incentive to introduce bias in consumption.

We augment the framework and contribute to the existing literature by introducing two different modes of consumption, one of which (labelled “background”) can be steered through recommendations. Users can also consume content actively; “active” consumption is costly, and can be understood as the active choice of what music to listen to, for example by building a playlist or picking songs sequentially. Background consumption, instead, is free and relies on a query generated by the platform. Consumers still have preferences over their consumption bundle, and these preferences are assumed to be known by the platform. If recommendations do not match consumers’ preferences, they are induced to increase their costly active consumption, which limits both the ability of the platform to introduce bias and her ability to monetize the streaming service through a fee.

We argue that such a modeling choice is both realistic and supported by the existing literature. As shown in Aguiar et al. (2021), platform-generated playlists are extremely popular and the evidence of bias in their composition is substantial. Bourreau and Gaudin (2022), additionally, presents an extension in which consumers can either follow the recommendation of the platform or costly search for content directly. The authors show that introducing the option to avoid the recommendation reduces the equilibrium bias, a prediction that our model generates as well. Differently from Bourreau and Gaudin (2022), however, we consider marginal increments of efforts in response to the introduction of the bias. We do so for two reasons: first, this approach allows us to avoid the extreme prediction of user either exclusively searching or exclusively following the recommendations.

Second, and more importantly, we believe this approach to better match the Discovery Mode as it is implemented in reality. As mentioned, the DM leads to lower royalties for songs consumed exclusively through specific channels, in particular personalized “thematic” playlists produced and published directly by Spotify. The DM, therefore, is only relevant when consumers engage with these bundles of generic content, a choice very different from that of consuming specific content selected by the user directly. Crucially, this allows us to consider users that are fundamentally “naive” in the context of background content consumption, that is, not

anticipating the bias introduced in this kind of bundles, arguing that background consumption is fundamentally inattentive by definition. We therefore relate to recent literature on inattention in digital markets, of which [Heidhues et al. \(2023\)](#), which explicitly considers consumers who make mistakes when evaluating offers, [Groh and von Wangenheim \(2023\)](#), which demonstrate how consumers not anticipating price discrimination changes competing firms' pricing strategy, and [Johnen and Somogyi \(2024\)](#), which considers consumers who do not anticipate shrouded features in an e-commerce setting, are notable examples.

2 A model of Discovery Mode

Consider a unit mass of homogeneous consumers who decide whether to subscribe to a streaming platform S paying a uniform subscription fee F . If they do not, they receive their exit option, which is normalized to zero. The subscription allows users to listen to music produced by two music labels $i = a, b$. Without loss of generality, we assume that content produced by label b is of higher quality.⁸ On the streaming platform, the two labels list their catalogs, which include many artists and genres. We assume that their payoffs (per consumption time) consist of royalty payments for the right to stream their content and are: $\pi_b \geq \pi_a > 0$, which are exogenously given. The first (weak) inequality reflects the larger catalog of the mainstream label and, hence, a relatively stronger bargaining power than the rival *vis-à-vis* the platform.

Consumers. Platform users have a preference for variety and prefer to listen to a mixed content bundle that includes music from both labels. Because the catalog of the mainstream label b is assumed to be larger, we impose that consumers optimally design a bundle that includes more of content b . Formally, we impose that consumers listen to one time-unit worth of music. We define $\lambda \in (0, 1/2)$ as the share of time that consumers choose to spend listening to content a , whereas the remaining $1 - \lambda$ of the time they spend listening to their preferred content b .

Music consumption occurs in two modes: active and background. In the former mode, consumers personally select and control the music they want to listen to. In other words, they build their own playlist with no interference from the platform recommendation system. In the latter mode, instead, consumers let the algorithm pick the queuing list for them, starting from a prompt (a song, an artist, or a genre they want to listen to). In this segment of demand, the platform has the possibility to (slightly) adjust the music bundle offered to consumers, possibly distorting their preferred mix away from the bundle implicitly defined by λ .

We define $\xi \in [-\lambda, 1 - \lambda]$ the recommendation bias that the platform may impose on the background demand segment. Intuitively, ξ represents the distortion away from λ . Moreover, it represents the additional demand that the platform transfers to content a in the users' consumption bundle.⁹ We use $v(\lambda; \xi)$ to label the utility of consumers from listening to a bundle

⁸We can interpret the quality of the two music labels as the number of artists they have signed exclusive contracts with, making the quality of the two labels equivalent to the breadth of their catalogs. Following with interpretation, one can think of a as either a minor label signing a smaller number of artists, or as a ghost label. We thank Andrea Mantovani and Lorien Sabatino for suggesting this interpretation.

⁹Notice that, when the value of ξ is negative, there is a subtraction of content a , whereas when it is positive, the platform suggests more of that content.

made of $\lambda + \xi$ of content a and $1 - \lambda - \xi$ of content b . The platform cannot bias the recommendation when the consumers are actively choosing their queuing list.¹⁰ Instead, in the background segment of the demand, bias is possible. We assume $v(\lambda; 0) > v(\lambda; \xi) \forall \xi \neq 0$, where the former is the utility in the active demand segment with no bias, and the latter is the utility when the recommendation is biased.¹¹ In other words, users unambiguously suffer from the bias. Crucially, consumers choose their preferred bundle only in the active segment $\lambda = \arg \max \{v(\lambda; 0)\}$. This behavioral assumption implies that consumers are myopic in that they do not account for the bias in the background segment of the market.

The choice of how to allocate their time between active and background consumption is endogenous. Formally, consumers face the following problem:

$$\max_{\alpha} U(\alpha, \xi, F) = \underbrace{\alpha v(\lambda; 0)}_{\text{active consumption}} + \underbrace{(1 - \alpha)v(\lambda; \xi)}_{\text{background consumption}} - \underbrace{F}_{\text{subscription fee}} - \underbrace{\frac{k\alpha^2}{2}}_{\text{cost of effort}}, \quad (1)$$

where $\alpha \in [0, 1]$ is the share of active consumption chosen by the users, and $1 - \alpha$ is the complementary share of background consumption. Active consumption is costly and $k > 0$ represents the cost of the effort it implies. Finally, F is the uniform subscription fee consumers have to pay to access the streaming platform.

The solution to the maximization problem is:

$$\alpha(\xi) = \frac{v(\lambda; 0) - v(\lambda; \xi)}{k}.$$

Users react to deviations from the optimal bundle by increasingly taking control of the queuing list. Intuitively, should the platform choose a no-bias policy ($\xi = 0$), then consumers would be able to costlessly consume their optimal mix as provided by the platform (which the platform can derive from users browsing and consumption history). If, instead, bias is introduced in their consumption bundle (in whatever direction this bias is introduced), users can turn to their device and adjust the queuing list costly.

Music Labels. The platform compensates the sellers (music labels) by paying a royalty π_i per streamed time. Moreover, the platform offers a Discovery Mode service (DM hereafter) that works as follows: if a label chooses to opt in the service, her content is promoted in the background demand segment. In exchange, the adopting label agrees to discount the royalty rate for any incremental demand in that segment by a factor $\rho_i \in \{1, \rho\}$. If seller i refuses to opt in for the DM, then $\rho = 1$, otherwise, the platform applies a discount rate $\rho \in [0, 1)$. It is important to remark that only the content in excess of the optimal mixes λ and $1 - \lambda$ can be discounted. The discount does not apply to the standard recommendation that reproduces the

¹⁰If a consumer picks a song, the platform cannot reproduce something different without drastically harming the users' utility.

¹¹This assumption can be easily microfounded as the rational choice of consumers maximizing a utility function à la Singh and Vives (1984). Define the problem of the consumers as $\max_{q_a, q_b} v(\lambda) = q_a v_a + q_b v_b - \frac{1}{2}(q_a^2 + q_b^2 + 2\phi q_a q_b)$, where q_i and v_i represent the quantity demanded of content i and its value, respectively, whereas $\phi \in (0, 1)$ is a differentiation parameter. From the system of first-order conditions, we obtain the optimal demands for the two contents $q_a = \frac{v_a - \phi v_b}{(1 - \phi^2)}$ and $q_b = \frac{v_b - \phi v_a}{(1 - \phi^2)}$. Then, we simply define the share $\lambda = \frac{q_a}{q_a + q_b} = \frac{v_a - \phi v_b}{(v_a + v_b)(1 - \phi)}$. Any deviation from this consumption bundle lowers the utility of consumers.

optimal mixes.¹² We can write the payoff of the labels as:

$$\begin{aligned}\Pi_a(\xi, \rho_a) &= \underbrace{\alpha(\xi) \lambda \pi_a}_{\text{active demand}} + \underbrace{(1 - \alpha(\xi)) \left((\lambda + \xi) \pi_a - \underbrace{\max\{\xi, 0\} ((1 - \rho_a)) \pi_a}_{\text{Discovery Mode discount}} \right)}_{\text{background demand}}, \\ \Pi_b(\xi, \rho_b) &= \underbrace{\alpha(\xi) (1 - \lambda) \pi_b}_{\text{active demand}} + \underbrace{(1 - \alpha(\xi)) \left((1 - \lambda - \xi) \pi_b + \underbrace{\min\{\xi, 0\} ((1 - \rho_b)) \pi_b}_{\text{Discovery Mode discount}} \right)}_{\text{background demand}}.\end{aligned}$$

To interpret the payoffs above, assume $\xi > 0$ and the platform biases consumption in the background demand toward the cheaper content a . Then, in that demand segment, label a obtains $\lambda + \xi$ of the demand. In exchange for this boost in the demand, the platform requires label a to discount the excess $(1 - \alpha(\xi))\xi$ portion of the demand generated by the promotion. The resulting royalty rate in that subsegment is $\rho_a \pi_a \leq \pi_a$. If $\xi < 0$, the same (symmetric) reasoning applies to label b . Under the DM, the platform biases its recommendation system and drives demand in the background segment of the demand toward content b . In exchange, the platform pays a discounted royalty rate for the biased demand $\rho_b \pi_b \leq \pi_b$.

Platform. The platform is an intermediary between the users and the labels. It receives a subscription payment from each user and compensates the labels by paying a royalty rate for each stream. Because the operational costs of the platform are the payoffs of the labels, there is a clear misalignment of interests between the suppliers of content and the aggregator.

We write the platform's maximization problem as:

$$\begin{aligned}\max_{F, \xi} \quad & \Pi_S(\xi, F, \rho) = F - \Pi_a(\xi, \rho_a) - \Pi_b(\xi, \rho_b) \\ \text{s.t.} \quad & F \leq \bar{F}(\xi) = \{F > 0 \mid U(\xi, F)|_{F=\bar{F}} = 0\} \\ \Rightarrow \quad & \bar{F}(\xi) - \lambda \pi_a - (1 - \lambda) \pi_b + (1 - \alpha(\xi)) \underbrace{(\max\{\xi, 0\}(\pi_b - \rho_a \pi_a) - \min\{\xi, 0\}(\pi_a - \rho_b \pi_b))}_{\text{cost minimization benefit}},\end{aligned}\tag{2}$$

where the constraint on F simply imposes that the subscription fee is compatible with the consumers' participation constraint.

Timing. The game is divided in two main stages unfolding within a single period. In the B2B stage, the platform and the labels agree on the contractual terms of the service. More in detail, at time $t = 0$ the DM is announced (for given ρ). Then at time $t = 1$, upon observing ρ , labels decide whether to opt in the DM or not. If they do, they agree that a discount ρ is applied to the compensation of each stream in excess of the shares λ and $1 - \lambda$ dictated by an unbiased algorithm. Formally, at this stage, labels decide $\rho_i = \{\rho, 1\}$.

¹²This is consistent with the way Spotify advertises this type of service on its website, for instance: “When Discovery Mode is turned on for a song, Spotify charges a commission on streams of that song in Discovery Mode contexts. All streams in other contexts are commission-free.” This feature is what differentiates this class of more traditional, legal promotion services from illegal practices such as *payola* schemes. See <https://artists.spotify.com/discovery-mode>. Last access July 15, 2025.

The B2C stage ensues: the platform interacts with its consumer base. At time $t = 2$, given the structure of payments π_i and ρ_i , the platform selects the subscription fee (F) and the recommendation bias (ξ). Finally, at time $t = 3$, users decide to join the platform and how to allocate their consumption time (α). The equilibrium concept is SPNE; we solve the model backwards.

3 Results

3.1 Main model: equilibrium analysis

From section 2 above, we know that $\alpha(\xi)$ is increasing in the size of the recommendation bias ξ and decreasing in the cost of effort k . Proceeding backwards, we can use the value of $\alpha(\xi)$ in the participation constraint of the consumer to derive the maximum fee that the platform can charge. Formally,

$$U(\xi, F) = \frac{(v(\lambda; 0) - v(\lambda; \xi))^2}{2k} + v(\lambda; \xi) - F \implies \bar{F}(\xi) = \frac{(v(\lambda; 0) - v(\lambda; \xi))^2}{2k} + v(\lambda; \xi).$$

Notice that the maximum subscription fee without bias is $\bar{F}(0) = v(\lambda; 0)$, and that the subscription fee is decreasing in $|\xi|$ for any $\xi \neq 0$. Because the profit of the platform in (2) is linearly increasing in F , the platform sets the highest possible fee that is consistent with the participation constraint. In other words, $F = \bar{F}(\xi)$. In what follows, we first derive the results in a benchmark case where the discovery mode is inactive (i.e., $\rho_i = 1$). Then, we will compare the results with the case where the platform offers a DM service.

Benchmark: no DM. Assume the platform does not offer a DM service. Then, the platform could only consider biasing the recommendation system towards cheaper content a , as doing otherwise would increase operational costs. Thus, in our benchmark, $\xi \geq 0$. Evaluated at $t = 2$, the platform's payoff can be rewritten as:

$$\Pi_S(\xi, \rho_i)|_{\rho_i=1} = \underbrace{\frac{(v(\lambda; 0) - v(\lambda; \xi))^2}{2k} + v(\lambda; \xi)}_{\bar{F}(\xi)} - \lambda \pi_a - (1 - \lambda) \pi_b + \underbrace{\left(1 - \frac{v(\lambda; 0) - v(\lambda; \xi)}{k}\right)}_{1 - \alpha(\xi)} \xi (\pi_b - \pi_a).$$

Biasing the recommendation system from the most expensive content b to the least expensive a yields a benefit of $\pi_b - \pi_a$ for each unit of demand transferred, which constitutes the marginal revenue from the bias. However, the bias negatively affects $\bar{F}(\xi)$ — both by reducing the utility of listening to music and increasing the effort to engage with the queue directly. Also, by increasing $\alpha(\xi)$, biasing the recommendation system results in a smaller portion of total demand being available for such a strategy, thus reducing the scope for biasing the demand further.

This trade-off regulates the decision of the platform. Formally:

$$\frac{\partial \Pi_S}{\partial \xi} = \underbrace{\left(1 - \frac{v(\lambda; 0) - v(\lambda; \xi)}{k} + \frac{\xi(\pi_b - \pi_a)}{k}\right)}_{\text{positive}} \underbrace{\frac{\partial v(\lambda; \xi)}{\partial \xi}}_{\text{negative}} + \underbrace{\left(1 - \frac{v(\lambda; 0) - v(\lambda; \xi)}{k}\right)}_{\text{positive}} (\pi_b - \pi_a).$$

It is possible to state the following:

Proposition 1. *Define ξ^{no} as the solution of the maximization problem of the platform absent the DM. We have that:*

1. *A strictly positive bias exists provided that $|\partial v(\lambda; \xi)/\partial \xi| < (\pi_b - \pi_a)$;*
2. *If the two labels are symmetric and $\pi_a = \pi_b$, then $\xi^{no} = 0$;*
3. *As the gap $\pi_b - \pi_a$ increases, the equilibrium bias ξ^{no} also increases, ceteris paribus.*

Proof. The proof of this proposition is obtained as follows: first, a strictly positive interior solution exists only if a solution of the first-order condition does not exist when evaluated at $\xi = 0$. In such case, one can see that the derivative above collapses to a simple comparison of $|\partial v(\lambda; \xi)/\partial \xi|$ and $(\pi_b - \pi_a)$. When the latter term is larger, the marginal revenues from increasing the bias at the margin are larger than the marginal costs. Relatedly, if $\pi_b = \pi_a$, there exists no marginal benefit from biasing the demand, and the equilibrium bias is $\xi^{no} = 0$. Similarly, when $\pi_b - \pi_a$ gets larger, the marginal revenues from biasing the recommendation system towards the cheapest content increases, and so does the incentives to set a larger bias. ■

One should notice that point 2 of the proposition above does not imply that the bias is strictly different from zero if $\pi_b \neq \pi_a$. Indeed, a necessary and sufficient condition for a strictly positive bias is that the marginal (financial) gains from biasing the demand are higher than the marginal loss of utility of consumers exposed to the bias (and the consequent reduction in the subscription fee).

Discovery Mode. The analysis of the B2C stage is substantially the same if we allow the platform to offer a promotion service such as the DM. Indeed, consumers' choice at stage $t = 3$ does not depend on the contractual relation between the platform and the labels. Hence, the maximum fee users are willing to pay is independent of the royalty rates π_a and π_b , as well as of any discount rate ρ .

Define $i = a, b$ the content favored by the bias, and $j = a, b$ (with $j \neq i$) the content that is harmed. Starting from (2), the platform's trade-off can be rewritten as:

$$\frac{\partial \Pi_S}{\partial \xi} = \underbrace{\left(1 - \frac{(v(\lambda; 0) - v(\lambda; \xi))}{k} + \frac{\xi(\pi_j - \rho\pi_i)}{k}\right)}_{\text{positive}} \underbrace{\frac{\partial v(\lambda; \xi)}{\partial \xi}}_{\text{negative}} + \underbrace{\left(1 - \frac{v(\lambda; 0) - v(\lambda; \xi)}{k}\right)(\pi_j - \rho\pi_i)}_{\text{positive}}.$$

From which we derive the following proposition:

Proposition 2. *Define ξ^* as the solution of the maximization problem of the platform. We have that:*

1. *A strictly positive bias (in absolute value) exists provided that $|\partial v(\lambda; \xi)/\partial \xi| < (\pi_j - \rho\pi_i)$;*
2. *For any given ρ , and assuming both labels opt in the DM, the platform always promote the originally cheaper content;*

3. *Raising the discount rate ($\rho \downarrow$) generates more asymmetry and increases the gap $\pi_j - \rho\pi_i$. In turns, this creates more incentives to bias the demand in favor of content i in equilibrium.*

Proof. The proof of this proposition follows the same logic as the ones for Proposition 1. Point 2 of this proposition is however different and is obtained by simply comparing the marginal benefits from biasing the demand in both directions. The platform chooses the direction that is more profitable. Formally, it biases the recommendation system towards the cheaper content if and only if $\pi_a - \rho\pi_b < \pi_b - \rho\pi_a$. This inequality implies $(1 - \rho)(\pi_b - \pi_a) > 0$, which is always satisfied for $\pi_b > \pi_a$. ■

Before proceeding to the B2B stage, it is important to analyze the similarities, as well as the differences, between this proposition and Proposition 1. Both propositions state the same message: a bias is possible if the marginal financial benefits from shifting the demand toward the cheapest content dominates the negative effect that the shift exerts on users' utility. However, this comparison of marginal benefits and costs is now governed by the discount rate ρ and this generates some interesting new results. First, depending on the size of ρ , the bias could potentially go in both direction. Indeed, when $\rho < \pi_a/\pi_b$, biasing the demand toward the originally more expensive content generates a financial benefit to the platform (although, if possible, the platform would always choose to promote content a). We will see later that this possibility allows the platform to induce the (reluctant) label a to accept the DM even when the ρ is very low.

Second, for low enough ρ , the platform can generate enough marginal gains to make introducing bias in consumption profitable even if, without the DM, the platform would choose not to introduce any bias. In other words, the incentives of the platform to bias the recommendation do not depend only on parameters, but also on a variable that the platform can be assumed to be able to adjust to some extent. Third, and relatedly, because the DM generates asymmetry endogenously, a strictly positive bias (in absolute value) can emerge in equilibrium even when the labels are ex-ante symmetric. This is, currently, the results that mostly diverge from the existing literature without explicitly contradicting it. We formalize this last result in the following:

Corollary 1. *With the DM, a recommendation bias can emerge also when contents are ex-ante symmetric from a cost perspective ($\pi_a = \pi_b$).*

We can now proceed backward to the B2B stage and analyze the incentives of the labels to commit to a DM service at $t = 1$. To do so, we first need to distinguish between two main scenarios: whether a positive bias (in absolute value) would emerge absent the DM or not. From Proposition 1, we know that this is indeed the case when $\pi_b - \pi_a > |\partial v(\lambda; \xi)/\partial \xi|$ — i.e., when the asymmetry of contents' costs is high enough that it is convenient for the platform to shift part of the demand in the background segment toward the cheapest content.

Let us start from the scenario in which the condition is not satisfied, and $\xi^{no} = 0$. In this case we can state the following:

Proposition 3. Assume that $\pi_b - \pi_a \leq |\partial v(\lambda; \xi)/\partial \xi|$. Then, for both labels, opting in the DM is a weakly dominant strategy. For any $\rho \in (0, \bar{\rho}_a)$, the platform biases the recommendation towards the originally cheaper content a . Otherwise, for $\rho \in [\bar{\rho}_a, 1)$, the platform does not bias the recommendation in any direction.

Proof. To prove this proposition, we need first to define $\bar{\rho}_i$ as the discount rate that makes it convenient (from the platform's standpoint) to bias recommendation towards content i . However, one should recall that we are in the scenario where $\xi^{no} = 0$, that is to say, $\pi_j - \rho\pi_i|_{\rho=1} \leq |\partial v(\lambda; \xi)/\partial \xi|$. From Proposition 2, we know that lowering ρ increases the left-hand side of the inequality leaving unchanged the right-hand side. Hence, we need to find:

$$\bar{\rho}_i = \left\{ \rho \in (0, 1) \mid \pi_j - \rho\pi_i|_{\rho=\bar{\rho}_i} = \left| \frac{\partial v(\lambda; \xi)}{\partial \xi} \right| \right\} \equiv \frac{\pi_j - \left| \frac{\partial v(\lambda; \xi)}{\partial \xi} \right|}{\pi_i}.$$

Because $\pi_j - \rho\pi_i \leq |\partial v(\lambda; \xi)/\partial \xi|$, it follows that $\bar{\rho}_i < 1$. Also, provided that $\pi_j > |\partial v(\lambda; \xi)/\partial \xi|$, we have that $\bar{\rho}_i > 0$. Finally, because $\pi_b \geq \pi_a$, it follows that $\bar{\rho}_a \geq \bar{\rho}_b$.

Then, we can turn to the incentives of the labels. Notice that the DM is designed (by assumption) to alter the price of the demand in excess of the optimal share λ — excess demand that materializes exclusively in the background segment of the demand. It means that neither labels are asked to lower their royalty rates on the demand generated by an unbiased recommendation system, nor labels are compensated+ for unrealized demand that would match the consumer optimal mix. With this in mind, one can see that firms are asked to accept an option: if the realized mix is more favorable than the consumers' optimum, firms accept to discount the demand in excess. This excess demand would not materialize absent the DM. Notice also that opting out would not protect the labels from a bias favoring their rival. Thus, labels are comparing 0 with some positive revenue if the bias is in their favor, whereas their payoffs are independent of their choices if the bias is in favor of their rival. If $\rho \geq \bar{\rho}_a$, labels know that the platform would not bias the recommendation system and are therefore indifferent. If $\rho \in (\bar{\rho}_b, \bar{\rho}_a)$, then label a anticipates that the platform would promote its content a if it accepts the commitment. Otherwise, the platform would not bias the recommendation system. For label a , thus, it is strictly dominant to accept the commitment, while firm b is still indifferent. Finally, for $\rho < \bar{\rho}_b$, the platform always biases the recommendation system. If both firms accept the option, the bias is in favor of the originally cheaper content a , whereas the bias is in favor of the originally more expensive content b if label a opts out. Accepting the DM is therefore strictly dominant for both labels. ■

Proposition 3 expands on the results in Proposition 2. Specifically, it states that, if the DM is necessary to generate the bias, it can do it provided that the discount rate chosen by the platform is sufficiently large ($\rho < \max\{\bar{\rho}_a, \bar{\rho}_b\}$). Because of the design of the DM that monetize on demand in excess of the consumer optimum, both labels are always willing to opt in the program, as it can be profitable but never causes harm. Existing asymmetry between royalties dictates which label will be systematically rewarded by the platform (i.e., the originally cheaper one).

Corollary 2. *Assume that $\pi_b - \pi_a \leq |\partial v(\lambda; \xi)/\partial \xi|$. For any ρ , label a (b) is weakly better (worse) off with the DM. The platform is always weakly better off.*

Let us now move to the other case, where $\xi^{no} > 0$. This scenario is fundamentally different from the previous one, for two main reasons. First, label a is not always better off by opting in the DM. Indeed, absent the DM, label a benefits from an existing bias. Opting in the DM requires label a to lower the margin on an excess demand that was (at least partly) already there. For a sufficiently large value of ρ , it is still worth opting in for label a , as the discount comes with some additional excess demand (Proposition 2). For low enough values of ρ , however, label a might, ceteris paribus, prefer not to opt in. Label b is in a very different position. The DM, in this case, becomes an opportunity to recover the part of demand that was shifted away from b to a by the platform prior to introduction of the DM, as well as to gain some excess demand at a discounted rate. Thus, from label b 's perspective, the DM is always convenient. Therefore, even if a has no direct incentive to pay for the DM, the threat of b paying for it and re-appropriating the bias forces a to take up the offer as well.

Second, but relatedly, because of the incentives of the two labels, the size of the bias might be non-monotonic in the size of the discount rate (ρ). We can therefore write the following:

Proposition 4. *Assume that $\pi_b - \pi_a > |\partial v(\lambda; \xi)/\partial \xi|$. Label a opts in the DM if and only if $\rho \in (0, \bar{\rho}_b)$ or $\rho \in (\bar{\rho}_a, 1)$. Label b is always willing to opt in the DM. In equilibrium, the platform sets bias $\xi = \min\{\xi^*, 1 - \lambda\}$ towards label a , where:*

$$\xi^* = \begin{cases} \xi(\rho) & \text{if } \rho \leq \bar{\rho}_b, \\ \xi(1) & \text{if } \rho \in (\bar{\rho}_b, \bar{\rho}_a), \\ \xi(\rho) & \text{if } \rho \geq \bar{\rho}_a, \end{cases}$$

with $\xi(\rho) > \xi(1) \forall \rho < 1$ and $\bar{\rho}_a = \{\rho \in (0, 1) \mid \xi(\rho) \rho = \xi(1)\}$.

Corollary 3. *Assume that $\pi_b - \pi_a > |\partial v(\lambda; \xi)/\partial \xi|$. The DM generally favors label a in terms of total demand. However, if $\rho < \bar{\rho}_b$, the platform uses the treat of biasing the demand towards label b to induce label a to accept a deal that is, ultimately, detrimental in terms of payoffs.*

3.2 Illustrative model

To illustrate the results more clearly, let us assume that the utility from listening to music is linearly decreasing in the intensity of bias:

$$v(\lambda; \xi) = v(\lambda) - \gamma|\xi|, \quad \text{with } \gamma \in (0, k),$$

where γ measures how sensitive the value of the bundle is to the size of the bias. Notice that, in this simplified model, $|\partial v(\lambda; \xi)/\partial \xi| = \gamma$. Starting from the maximization of utility in stage $t = 3$, we can derive the equilibrium of the game without and with the discovery mode. First, consumers allocate time to consumption modes regardless of the contractual relations specified in the B2B stage. Formally, $\alpha(\xi) = \frac{\gamma|\xi|}{k}$. Using this, it is possible to calculate the maximum fee

that users are willing to pay for the streaming service:

$$\bar{F} = v(\lambda) - \gamma|\xi| + \frac{\gamma^2 \xi^2}{2k}.$$

Consider first the benchmark with no discovery mode. The problem of the platform becomes:

$$\Pi_S(\xi, \rho_i)|_{\rho_i=1} = \underbrace{v(\lambda) - \gamma|\xi| + \frac{\gamma^2 \xi^2}{2k}}_{\bar{F}(\xi)} - \lambda \pi_a + (1 - \lambda) \pi_b + \underbrace{\left(1 - \frac{\gamma|\xi|}{k}\right)}_{1 - \alpha(\xi)} \xi (\pi_b - \pi_a),$$

where the first-order condition is satisfied if:¹³

$$\frac{\partial \Pi_S}{\partial \xi} = (\pi_b - \pi_a - \gamma) - \frac{\gamma(2(\pi_b - \pi_a) - \gamma)}{k} \xi = 0 \iff \xi = \frac{k(\pi_b - \pi_a - \gamma)}{\gamma(2(\pi_b - \pi_a) - \gamma)}.$$

Consistently with Proposition 1:

Remark 1. *Absent the DM, the platform sets a strictly positive bias $\xi > 0$ if $\pi_b - \pi_a > \gamma$ and no bias ($\xi = 0$) otherwise.*

If instead the discovery mode is allowed, the bias in favor of a label i can be written as

$$\frac{\partial \Pi_S}{\partial \xi} = (\pi_j - \rho \pi_i - \gamma) - \frac{\gamma(2(\pi_j - \rho \pi_i) - \gamma)}{k} \xi = 0 \iff |\xi| = \frac{k(\pi_j - \rho \pi_i - \gamma)}{\gamma(2(\pi_j - \rho \pi_i) - \gamma)}.$$

As for the general case, the only thing that changes is that the gap between costs is adjusted with the discount rate ρ .

Remark 2. *With the DM, the platform sets a strictly positive bias $|\xi| > 0$ if $\pi_j - \rho \pi_i > \gamma$ and no bias ($\xi = 0$) otherwise. All results in Proposition 2 holds.*

To proceed backward to the B2B stage, we first need to define the thresholds below which the discount rate generates an incentive to bias the recommendation:

$$\bar{\rho}_i = \frac{\pi_j - \gamma}{\pi_i}.$$

Notice that when $\pi_j - \pi_i < \gamma$ it holds that $\bar{\rho}_i < 1$, while $\bar{\rho}_i$ is otherwise bigger than 1 (meaning that there is bias in the no DM benchmark).

Starting from the scenario in which there is no bias in the benchmark, that is $\pi_j - \pi_i < \gamma$, we have:

Remark 3. *If $\pi_j - \pi_i < \gamma$, both labels opt in the DM. The platform sets the recommendation system such that $\xi = \min\{\xi^*, 1 - \lambda\}$, where:*

$$\xi^* = \begin{cases} \frac{k(\pi_b - \rho \pi_a - \gamma)}{\gamma(2(\pi_b - \rho \pi_a) - \gamma)} & \text{if } \rho \in [0, \bar{\rho}_a), \\ 0 & \text{if } \rho \in [\bar{\rho}_a, 1]. \end{cases}$$

Both the platform and label a are weakly better off in equilibrium for any $\rho \in [0, 1]$. Instead, label b and consumers are weakly worse off.

¹³The second-order condition for a maximum is satisfied for $2(\pi_b - \pi_a) > \gamma$.




Figure 1: Equilibrium bias when $\pi_b - \pi_a > \gamma$. Linearity is imposed for the sake of exposition, ξ^* is concave in ρ . Notice that $\tilde{\rho}_a$ could potentially be either smaller than $\bar{\rho}_b$ or larger than 1; in the former case, the bias shows no discontinuity and label a opts in the DM $\forall \rho$. In the latter case, the adoption of the DM only occurs for low $\rho < \bar{\rho}_b$ and is, therefore, always detrimental to label a .

Finally, turning our attention to the case in which a bias would have existed even absent the DM, i.e., when $\pi_b - \pi_a > \gamma$, we have the following:

Remark 4. *If $\pi_j - \pi_i > \gamma$, Label b always opts in the discovery mode, whereas label a opts in if $\rho \in [0, \bar{\rho}_b)$ or $\rho \in (\tilde{\rho}_a, 1)$. The platform sets $\xi = \min\{\xi^*, 1 - \lambda\}$, where:*

$$\xi^* = \begin{cases} \frac{k(\pi_b - \rho\pi_a - \gamma)}{\gamma(2(\pi_b - \rho\pi_a) - \gamma)} & \text{if } \rho \in [0, \bar{\rho}_b), \\ \frac{k(\pi_b - \pi_a - \gamma)}{\gamma(2(\pi_b - \pi_a) - \gamma)} & \text{if } \rho \in [\bar{\rho}_b, \tilde{\rho}_a), \\ \frac{k(\pi_b - \rho\pi_a - \gamma)}{\gamma(2(\pi_b - \rho\pi_a) - \gamma)} & \text{if } \rho \in [\tilde{\rho}_a, 1], \end{cases}$$

with $\tilde{\rho}_a = \frac{(2\pi_b - \gamma)(\pi_b - \pi_a - \gamma)}{2(\pi_b - \pi_a) - \gamma}$. See Figure 1 for a graphical representation. All results in Corollary 3 hold.

4 Extensions

4.1 Fully rational consumers

We first extend the model by assuming that users do not incur any cost for active consumption ($k = 0$). This scenario coincides with the fully rational specification in which consumers are always able to maximize their utility by selecting the most appropriate consumption bundle.

One can see that, if $k = 0$, the optimal level of α chosen by the consumers, i.e., the optimal level of active consumption, is $\alpha(\xi)|_{k=0} = 1$. In other words, if consumers face no costs for active consumption (which is the mode of consumption that provides the highest payoff), they would actively select music all the time.

This result, which is not surprising, carries important implications for the equilibrium of the game. In fact, $\alpha(\xi) = 1$ implies that no bias is possible at all. The platform cannot shift any part of the demand towards the cheapest product and, hence, cannot operate any cost-minimization strategy by altering the demand schedule of consumers. Consequently, firms payoff in this scenario can be written as:

$$\Pi_a^{ben} = \lambda\pi_a; \quad \Pi_b^{ben} = (1 - \lambda)\pi_b; \quad \Pi_S^{ben} = \bar{F}(0) - \lambda\pi_a - (1 - \lambda)\pi_b.$$

Because no bias is technically possible — i.e., the platform cannot reproduce something different than the song the consumers actively select — the scope for any product placement service such as the discovery mode falls drastically. Indeed, in the context of our model, the scope for such a built-in recommendation bias disappears.

Remark 5. *In a fully rational model with no behavioral costs ($k = 0$), the platform cannot bias its recommendation system. Indeed, a recommendation system is not even necessary as consumers know what they want to listen to and actively choose the most appropriate content.*

This first extension qualifies the results of the model by removing the cost of active consumption: when active consumption is the only effective consumption, bias cannot emerge. The outcome is tautological: If consumers can costlessly control what content they consume, any deviation from their optimal mix λ becomes immaterial as consumers would always optimally select their optimal mix themselves. However, this would also imply that, if consumers were acting as such, nobody would ever consume platform-generated mixes of content, which is demonstrably false. We take the result as a sign that consumption of platform-generated mixes of content is necessary for bias to emerge.

4.2 Bias and market structure

An interesting question that arises from this framework is whether the platform’s ability to bias consumption creates incentives for labels to expand by signing more artists. Several high-profile acquisitions by major labels like Sony and Warner Music have taken place in the last few years.¹⁴ One of the most controversial cases can be traced to Sony’s acquisition of the British recording company *AWAL* after poaching “some of the most commercially successful artists” affiliated with *AWAL* in the past.¹⁵ It is therefore important to understand whether the intermediation of platforms such as Spotify (or Amazon Music or Apple Music) can have an impact on the incentives of major labels, the ones generally biased against on these platforms (as shown by Aguiar et al., 2024 and confirmed by Reimers and Waldfogel, 2023), to increase their catalog through acquisitions of independent artists and labels.

It is not obvious how to translate a poaching strategy in the language of the model above. Therefore, to extrapolate the incentives of the labels to expand by poaching artists, we rely on our microfoundation: Recall that, following Singh and Vives (1984), we can define the problem of the consumers as $\max_{q_a, q_b} v(\lambda) = q_a v_a + q_b v_b - \frac{1}{2} (q_a^2 + q_b^2 + 2\phi q_a q_b)$, where q_i and v_i represent the quantity demanded of content i and its value, respectively, and $\phi \in (0, 1)$ is a differentiation parameter. Furthermore, recall that v_i can be interpreted as a measure of the “size” of the catalog of a label. Then, an artist migrating from label a to label b is consistent with a decrease (increase) of v_a (v_b).

First, there is a clear incentive to expand if consumers are fully rational. Imagine that label b decided to acquire one of the artists active under label a . In doing so, the major label can

¹⁴For example: Sony recently acquired the Czech music label *Supraphon*, and Warner Music recently acquired the catalog of Italian indie label *DWA Record*. See <https://www.sonymusic.com/sonymusic/supraphon-acquisition-czech-republic/> and <https://www.musicbusinessworldwide.com/warner-music-group-acquires-dwa-records-expanding-italian/>. Last access on July 15, 2025.

¹⁵See <https://www.musicbusinessworldwide.com/if-you-wanted-to-know-why-sony-music-bought-awal-this-is-why/>. Last access on July 15, 2025.

increase the share of their content actively selected by the users (or, q_b increases, q_a decreases, and $\lambda^* = \frac{q_a}{q_a + q_b}$ decreases). In the conservative case where this does not alter the royalty rate asked by the major, the incentive is simply determined by the additional stream of royalty revenue due to the presence of the indie artists in the major label's catalog. It follows that the incentive to expand does not derive solely from the presence of bias.

That said, we are interested in understanding if the presence of the bias alters the incentives highlighted above. A bias of magnitude ξ can be incorporated into the microfoundation through an increase in q_a and a decrease in q_b of size ξ . As this shift in consumption generates a new effective $\lambda = \frac{q_a + \xi}{q_a + q_b} > \lambda^*$, it also generates a different (lower) utility for consumers. Formally, by including ξ in the utility function, we derive:

$$v\left(\frac{q_a + \xi}{q_a + q_b}\right) = \frac{v_b^2 + v_a^2 - 2v_a v_b \phi}{2(1 - \phi^2)} - (v_b - v_a)\xi = v(\lambda^*) - (v_b - v_a)\xi.$$

The exercise instructs the following modeling choice: the act of poaching artists from a label by a rival makes the utility of consumers more reactive to the bias. Indeed, the microfoundation highlights that the effect of introducing bias ξ on consumers' utility is proportional to $v_b - v_a$, which grows as label b acquires artists and label a loses them. Henceforth, we refer to this through the lenses of the illustrative model. Formally, we assume that poaching is captured by increases to γ in the consumers' utility, $v(\lambda; \xi) = v(\lambda) - \gamma|\xi|$. For the sake of clarity, it is useful to think of γ as a positive function of the difference in quality between the content of the two labels ($v_b - v_a$), and, by extension, of the difference in the sizes of their catalogs. We can now discuss how the bias affects the incentives of b to poach artists from a .¹⁶

From the illustrative model, one can see that γ affects the ability to bias consumption in two ways. First, bias is only feasible if the difference in effective royalties, $\pi_b - \rho\pi_a$, $\rho \leq 1$, is larger than γ . Second, the equilibrium bias ξ decreases in γ as per Remarks 3 and 4. Intuitively, because the equilibrium bias resolves the trade-off between operational costs and revenue generated by the fee (which, in equilibrium, equals the utility of consumers accounting for the equilibrium bias), a larger γ constrains the bias due to an enhanced effect on the fee. This is in addition to the additional stream of royalty revenue already present in the fully rational case mentioned above. Effectively, this provides the labels with a tool to mitigate the bias. Indeed, the firm that is more vulnerable to the cost-minimization strategy of the platform — in our model, the major — can poach artists who signed with the indie label and, by doing so, increase its own catalog. Provided that γ is indeed a positive function of $v_b - v_a$, this practice translates into a lower bias in equilibrium.

The model, then, suggests that the ability of the platform to introduce bias in the consumers' bundles to minimize operational costs might strengthen the incentives of major labels to expand. Ignoring for the moment the possibility of entry by new artists, this may lead to higher concentration in the music industry compared to an hypothetical counterfactual in which bias was unfeasible or otherwise absent. The results echo concerns expressed by some of the

¹⁶We focus our attention on the major label's incentives for two reasons. First, major labels are the ones more likely to expand by acquiring smaller competitors. Second, and relatedly, because they are the ones most penalized by the bias, it is natural to ask whether they would have a way to counteract it. While we acknowledge that indie label a would also have the incentive to expand, we argue that if both a and b have the same incentive to expand, b would be the one to be able to follow through.

most influential names in independent music in the last years that strongly suggest that major labels’ incentives to expand their indie catalogs is strengthened by their interaction with digital platforms such as Spotify. In 2017, Sony bought two independent rival distributors, the German *Finetunes* and the Norwegian *Phonofiles*. When asked to comment on the expansion of major labels through acquisition of indie artists and recording companies, Charles Caldas, CEO of independent music rights organization *Merlin*, commented that this “is not positive news for the indie labels and artists [...]. Merlin has long been vocal about our concerns that the majors, via their faux-indie imprints, are land-grabbing independents rights in order to bolster their market shares and use the value of those indie artists to extract disproportionate value from the market in their negotiations with digital services.”¹⁷

5 Conclusion

In this paper, we examined the incentives behind, mechanisms justifying, and implications of contractually stipulated bias in digital platforms. We focused on Spotify’s Discovery Mode, a recent addition to Spotify’s suite of services targeted to the content provider hosted on the platform. Discovery Mode allows creators to select specific content to be prioritized by the platform when building personalized consumption bundles in exchange for a reduction of the royalties owed by the platform to the creators. We showed that this service increases bias in the recommendations at best, and hurts independent labels and artists by leveraging the incentive of larger players to pay for visibility at worse. Furthermore, we highlight a positive relation between the ability of a platform to bias consumption and the incentives of labels to try to expand through acquisition of rivals’ catalogs of content. The analysis could be further enhanced with the introduction of a bargaining stage between the platform and the labels over ρ , a dimension currently outside of the scope of this paper.

Our model captures a largely overlooked feature of music streaming platforms: their cost structure, characterized by per-stream royalty payments, creates strong incentives to favor content that minimizes payout obligations, a dynamic absent in traditional radio’s fixed licensing model. Unlike the radio, streaming platforms also benefit from low user exit costs: listeners dissatisfied with a song can simply skip or change playlists without leaving the platform. This reduces the risk associated with biased curation, making biased algorithmic recommendations easier to implement.¹⁸ Somehow paradoxically, while digital platforms promised democratized access to music, streaming platforms’ cost structure enhances both the incentive and capacity for biased promotion, potentially surpassing that of traditional media.

This observation makes studying bias in the context of digital platforms more important than ever. “Pay for play” deals are not novel in the music industry: payola deals have been documented since the age in which consumption of music was dominated by commercial radio, and is classified as commercial bribery (and, therefore, illegal) in the United States. Because the unprecedented reach and commercial presence of intermediaries has grown significantly, making gatekeepers object of countless scholarly investigations and regulatory interventions,

¹⁷See <https://www.musicbusinessworldwide.com/the-orchard-bolsters-distribution-might-europe-finetunes-phonofile-acquisition/>. Last access on July 15, 2025.

¹⁸This argument is also proposed in Aguiar et al. (2024).

with the European DMA and DSA being the most obvious and recent examples. The ability of intermediaries such as Spotify to maintain a critical mass of users due to this reduced risk associated with biased curation, the immense amount of data at their disposal, and their de facto role of gatekeepers for digital music consumption makes the prospect of a laissez-faire approach to their design choice dubious.

While proposing specific interventions is, at this stage, premature, we stress that a blanket ban of promotional services is hardly an optimal solution. Independent labels and the artists they represent can benefit greatly from such services, especially in an industry as concentrated and saturated as the music industry. Biasing recommendations towards lesser known artists and groups could be beneficial to consumers who are unlikely to be exposed to less popular artists organically. At the same time, these “pay for play” deals can be detrimental to both competition and user experience when the ones to benefit from it are already popular creators. Anecdotally, the latter case seem to be more common: in 2018 heavy allegations against Spotify were raised after the release of Drake’s album “Scorpion” on the platform. In the weeks that followed, songs by the Canadian rapper were included in a variety of Spotify-generated playlists of dubious fit, and even in playlists explicitly labeled “Best of British [artists].” Given the outrage that followed, it is at least plausible that any attempt at introducing payola afterwards would be carried out in ways designed to be harder to spot for users and regulators, making the matter of investigating these deals more relevant than ever.

References

- 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. (2024). Platform Power Struggle: Spotify and the Major Record Labels. *NBER Working paper 33048*.
- 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.
- 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, N. and Tsai, H.-T. (2024). Steering via algorithmic recommendations. *The RAND Journal of Economics*, 55(4):501–518.
- 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 Sarvary, M. (2023). Social media and news: Content bundling and news quality. *Management Science*, 69(1):162–178.
- 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.
- Freimane, M. (2022). Substituting Away? The Effect of Platform Bargaining Regulation on Content Display. Technical report, mimeo.

- Gambato, J. and Sandrini, L. (2023). Not as good as it used to be: Do streaming platforms penalize quality? *ZEW-Centre for European Economic Research Discussion Paper*, (23-032).
- Groh, C.-C. and von Wangenheim, J. (2023). Pigou meets wolinsky: Search, price discrimination, and consumer sophistication. *Price Discrimination, and Consumer Sophistication (August 24, 2023)*.
- Hagiü, A. and Jullien, B. (2011). Why do intermediaries divert search? *The RAND Journal of Economics*, 42(2):337–362.
- Hagiü, A. and Jullien, B. (2014). Search diversion and platform competition. *International Journal of Industrial Organization*, 33:48–60.
- Heidhues, P., Köster, M., and Köszegi, B. (2023). Steering fallible consumers. *The Economic Journal*, 133(652):1430–1465.
- 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.
- Johnen, J. and Somogyi, R. (2024). Deceptive features on platforms. *The Economic Journal*, 134(662):2470–2493.
- Madio, L. and Quinn, M. (2024). Content moderation and advertising in social media platforms. *Journal of Economics & Management Strategy*.
- Oestreicher-Singer, G. and Sundararajan, A. (2012a). Recommendation Networks and the Long tTail of Electronic Commerce. *Mis quarterly*, pages 65–83.
- Oestreicher-Singer, G. and Sundararajan, A. (2012b). The visible hand? demand effects of recommendation networks in electronic markets. *Management science*, 58(11):1963–1981.
- 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.
- Reimers, I. and Waldfogel, J. (2023). A framework for detection, measurement, and welfare analysis of platform bias. Technical report, National Bureau of Economic Research.
- Singh, N. and Vives, X. (1984). Price and quantity competition in a differentiated duopoly. *The Rand journal of economics*, pages 546–554.
- Thomes, T. P. (2013). An Economic Analysis of Online Streaming Music Services. *Information Economics and Policy*, 25(2):81–91.

- Tucker, C. and Zhang, J. (2010). Growing two-sided networks by advertising the user base: A field experiment. *Marketing Science*, 29(5):805–814.
- Waldfogel, J. (2024). Amazon Self-preferencing in the Shadow of the Digital Markets Act. Technical report, National Bureau of Economic Research.



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.